

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays in Macroeconomics

Permalink

<https://escholarship.org/uc/item/3wt0g0h4>

ISBN

9798189271502

Author

Yding, Anders

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Essays in Macroeconomics

by

Anders Yding

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 Emi Nakamura, Chair

Professor Yuriy Gorodnichenko

Professor Jón Steinsson

Spring 2026

Essays in Macroeconomics

Copyright 2026
by
Anders Yding

Abstract

Essays in Macroeconomics

by

Anders Yding

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Emi Nakamura, Chair

This dissertation consists of three essays studying the macroeconomic effects of fiscal policy, the aggregate consequences of labor market power, and the incidence and efficiency of land value taxation.

In the first essay (coauthored with Ethan McClure), we estimate the macroeconomic effects of government spending by constructing a new series of defense spending surprises using a high-frequency identification approach. We identify events in the U.S. federal budget process that reveal news about defense spending, focusing on dates when different versions of defense appropriations bills are passed in the House or the Senate. We then quantify the defense spending surprise associated with each event using daily stock returns of large U.S. defense contractors. A positive defense spending surprise predicts a persistent increase in future defense spending that is entirely passed through to total government spending. On average, the identified spending increases are financed with a combination of taxes and debt and are met with a muted monetary policy response. We estimate the dynamic causal effects of defense spending and find sizable increases in output and consumption. Our estimates imply a cumulative fiscal multiplier—the cumulative output response divided by the cumulative spending response—of 1.2 over a five-year horizon. Our results indicate that government spending stimulates (“crowds in”) economic activity in the private sector. A standard heterogeneous-agent New Keynesian model matches our empirical evidence.

In the second essay (coauthored with Justin Bloesch and Birthe Larsen), we develop a model of wage posting and on-the-job search where firms use wages and recruiting expenditures to attract workers. We capture three sources of labor market monopsony power: preference heterogeneity, search frictions, and labor market concentration. The model allows firms’ labor supply curves to be perfectly elastic in the long run but inelastic in the short run, consistent with evidence that firms pay higher wages while growing even though the wage premium at large firms is small. We provide empirical evidence that labor supply curves are perfectly elastic in the long run using the effect of export demand shocks on the wage

growth of job switchers in Denmark. Our results imply that monopsony rents are dissipated by recruiting costs, which can reconcile existing estimates of monopsony power with the profit share of national income in rich countries.

In the final essay (coauthored with Ulf Nielsson and Caleb Wroblewski), we study the effects of taxing land values. Land value taxes are often seen as particularly desirable because the fixed supply of land implies no efficiency loss from taxation, with the entire tax burden falling on current landowners. We study the incidence and efficiency of land taxes using a unique quasi-experiment that generated persistent variation in land tax rates across Danish municipalities. In contrast to the predictions of standard, neoclassical models, we estimate a precise zero effect of land taxes on residential home prices. The precision of our estimates allows us to confidently rule out full capitalization of taxes into home prices using leading estimates of housing discount rates. Our results imply that the burden of land taxes is shared with tenants and future purchasers. We also estimate null effects of land taxes on measures of housing development, mobility, and homeownership, though we do find that older homeowners sort away from high tax areas. Our results are consistent with limited efficiency costs of land value taxation but imply that land taxes are more regressive in our setting than predicted by standard models.

To my grandparents Vibeke and Leif, and Guri and Kurt

Contents

Contents	ii
List of Figures	iv
List of Tables	vii
1 The Macroeconomic Effects of Defense Spending News	1
1.1 A New Measure of Defense Spending News	6
1.2 The Dynamic Effects of Defense Spending News	12
1.3 Aggregate Fiscal Multipliers	22
1.4 Aggregate Fiscal Multipliers in Standard Models	25
1.5 Conclusion	35
2 Monopsony with Recruiting	36
2.1 Introduction	37
2.2 Monopsony with Recruiting: Atomistic Firms	40
2.3 Empirical Evidence and Parameter Estimation	52
2.4 Profit Puzzle with Oligopsonistic Competition	65
2.5 Conclusion	71
3 The Incidence and Efficiency of Land Value Taxation	72
3.1 Introduction	73
3.2 Model	78
3.3 Theoretical Predictions	80
3.4 Institutional Details	83
3.5 Data	85
3.6 Instrument and First Stage	86
3.7 Capitalization in Housing Prices	90
3.8 Aggregate Effects of Land Taxes	97
3.9 Discussion	106
Bibliography	111

A	Appendix to The Macroeconomic Effects of Defense Spending News	125
A.1	Campbell-Shiller Decomposition	125
A.2	Additional Figures	128
A.3	Model	141
A.4	Additional Model Results	152
B	Appendix to Monopsony with Recruiting	165
B.1	Estimation Details	165
B.2	Switcher Wage Levels	168
B.3	Robustness: Measures of Switcher Wage Growth	169
B.4	The Firm Size Wage Premium	171
B.5	Worker's Problem	172
B.6	Firms' Dynamic Problem	172
B.7	Proof of Proposition 1	175
B.8	Recruiting Costs as Reallocating Labor	183
B.9	Firm's Problem with Hiring Costs	184
B.10	Oligopsony	185
B.11	Bounds on Markdowns with Oligopsony	189
B.12	Calibration with Forward-Looking Workers	190
B.13	Amenities	192
B.14	Extension with Worker Heterogeneity	193
C	Appendix to The Incidence and Efficiency of Land Value Taxation	197
C.1	Appendix Figures	198
C.2	Model with Rent Ceiling	200
C.3	Comparison with Høj et al. (2018)	203

List of Figures

1.1	Event Study Examples	13
1.2	Daily Defense Spending Surprises	14
1.3	Defense Spending Surprises Predict Future Defense Spending	17
1.4	Responses of Taxes and Deficits to a Defense Spending Surprise	19
1.5	Muted Monetary Policy Response	20
1.6	Government Spending Stimulates Private Economic Activity	21
1.7	Labor Market Responses to a Defense Spending Surprise	22
1.8	First-stage F -Statistics	23
1.9	Estimated and Model-Implied Response to a Defense News Shock	31
1.10	Estimated and Model-Implied Fiscal Multiplier	32
1.11	Role of Monetary Policy	33
1.12	Role of Fiscal Policy	34
2.1	Wage Bill, Recruiting Costs, and Marginal Product	48
2.2	Export Demand Shock Autocorrelation and Revenue Response	56
2.3	Response of Firm Employment, Hiring Rate, and Switcher Wage Growth to Export Demand Shocks	59
2.4	Path of Switcher Wage Level by Trade Shock Timing Relative to Job Switch	64
2.5	Response of Stayer Wages to Export Demand Shock	66
3.1	A Model of the Market for Land	81
3.2	First Stage: The Effect of Tax Change Instrument on Statutory Tax Rate	87
3.3	The Effect of Statutory Tax Changes on Property-Level Effective Tax Rates	88
3.4	First Stage: The Effect of Land-Tax Exposure Instrument Within A Treatment Area	90
3.5	First Stage: The Effect of Land-Tax Exposure Instrument Combining Policy Shock and Exposure Measures	91
3.6	Higher Taxes Have No Effect on Sales Prices	93
3.7	The Effect of Land Taxes Using Separate Identification Strategies	94
3.8	Higher Taxes Have No Effect on Sales Prices of Single-Family Homes	95
3.9	The Effect of Land Taxes on Aggregate Residential Development	99
3.10	The Effect of Land Taxes on Development of Empty Lots	100
3.11	The Effect of Land Taxes on Development of Empty Lots	102

3.12	Land Taxes Do Not Cause Homeowners to Sell Their Home	103
3.13	Land Taxes Do Not Increase Homeownership	105
3.14	Land Taxes and the Allocation of Housing	107
3.15	Land Ownings Across the Income Distribution	109
A.1	Real Government Spending per Capita in the United States	128
A.2	Defense Stock Index Coverage: Share of Top 100 Defense Contracts	129
A.3	Quarterly Defense Spending Surprises	130
A.4	Quarterly Defense Spending Surprises	131
A.5	Shock Autocorrelation	132
A.6	Shock Predictability (Granger Causality Tests)	133
A.7	Defense and Market Returns in September 2001	134
A.8	Spending Responses to a Defense Spending Surprise (Baseline, Long Horizon) .	135
A.9	Spending Responses to a Defense Spending Surprise (Post 1955, Long Horizon) .	136
A.10	Spending Responses to a Defense Spending Surprise (Excluding September 11 Attacks, Long Horizon)	137
A.11	Spending Responses to a Defense Spending Surprise (Excluding Supplemental Appropriations, Long Horizon)	138
A.12	Response of Taxes and Deficits to a Defense Spending Surprise (Pre-2000 Sub- sample)	139
A.13	Response of Taxes and Deficits to a Defense Spending Surprise (Post-1954 Sub- sample)	140
A.14	Role of Monetary Policy	152
A.15	Role of Fiscal Policy	153
A.16	Alternative Model: RANK	155
A.17	Alternative Model: TABU	156
A.18	Alternative Model: Flexible Prices	157
A.19	Alternative Model: Flexible Wages	158
A.20	Alternative Model: Capital Adjustment	159
A.21	Alternative Model: No Price Indexation	160
A.22	Alternative Model: Policy Smoothing	161
A.23	Alternative Model: larger Frisch elasticity ($\eta^{-1} = 1$)	162
A.24	Effect of an AR(1) Defense Spending Shock	163
B.1	Distribution Estimates of χ and Inverse Labor Supply Elasticity $\varepsilon_{w,N}$	168
B.2	Raw Log Wage of Switchers	169
B.3	Robustness: Employment, Hiring, and Switcher Wage Growth Response	170
B.4	Recruiting and Separation Elasticities for Different Wage Levels $\beta^w=.99, s_0=.025$.	192
C.1	Mentions of Structural Reform in Major Danish Newspapers	198
C.2	Land Taxes and the Allocation of Housing: Placebo Tests	199

C.3	Land Tax Changes Strongly Correlated With Share of Vacation Homes in Treatment Area	204
C.4	ITT Effect of Land Taxes Including Vacation Homes and Without Control . . .	205
C.5	ITT Effect of Land Taxes Without Vacation Homes and Including Control . . .	206

List of Tables

1.1	Cumulative Fiscal Multiplier Estimates	25
1.2	Model Calibration	30
2.1	Decomposing Marginal Product into Wages, Recruiting Costs, and Profits . . .	47
2.2	Firm Summary Statistics	54
2.3	Switching Worker Summary Statistics	54
2.4	Estimated Coefficients from Stacked Regression	61
2.5	Fitted Parameters χ and σ and Inverse Labor Supply Elasticity $\varepsilon_{w,N}$	63
2.6	Calibration	69
2.7	Distribution of Income: Standard Monopsony, CRS Recruiting, and Concentration	70
3.1	Effect on Home Prices and Implied Discount Rates	96
A.1	Variable Names and Symbols	149
A.2	Alternative Calibrations.	164
B.1	Variance Covariance Matrix $\hat{\Sigma}$	166
B.2	Firm Size Wage Premium	171
B.3	Calibration with Forward Looking Workers	191

Chapter 1

The Macroeconomic Effects of Defense Spending News

with Ethan McClure

We are grateful to Adrien Auclert, Alan Auerbach, Jesper Böjeryd, Benzhad Diba, Jordi Galí, Yuriy Gorodnichenko, Jonathon Hazell, Vitor Jung, Moritz Kuhn, Chen Lian, Alfredo Mendoza, Kurt Mitman, Emi Nakamura, Lukas Nord, Morten Ravn, Søren Hove Ravn, Andreas Schaab, Benjamin Schoefer, David Sraer, Jón Steinsson, and Caleb Wroblewski, as well as seminar participants at Berkeley, the Chicago Fed, the New York Fed, and NorMac 2025 for helpful discussions, comments and suggestions. The views expressed in this chapter are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Boston or the Federal Reserve System.

How do changes in government spending affect aggregate economic activity? The answer to this question is typically summarized by a fiscal multiplier: the percentage change in output that results from increasing government spending by one percent of GDP.¹ There is broad consensus that government spending increases output, but the empirical literature is divided on whether fiscal multipliers tend to be smaller or larger than one—that is, does an increase in government spending tend to dampen (“crowd out”) or stimulate (“crowd in”) private consumption and investment?

Isolating the causal effect of government spending on output is challenging. Observed changes in spending could reflect an endogenous policy response to the state of the economy. For example, policymakers have routinely turned to government spending as a tool for stabilizing the economy in recessions, such as the New Deal in the 1930s or the American Recovery and Reinvestment Act of 2009.² To mitigate concerns about reverse causation, a long empirical literature uses variation from defense spending and wars, arguing that this type of spending is unlikely to respond to economic events (Barro, 1981; Hall, 1986; Rotemberg and Woodford, 1992; Ramey and Shapiro, 1998; Hall, 2009; Barro and Redlick, 2011; Ramey, 2011b; Ramey and Zubairy, 2018).

However, even if wars are typically not triggered by the business cycle, the decision of how much to spend on a war—or how much to spend on defense in peacetime—could still depend on the economy. For instance, Defense Secretary Gates stated in 2010: *“Given America’s difficult economic circumstances and parlous fiscal condition, military spending on things large and small can and should expect closer, harsher scrutiny. [...] The gusher has been turned off, and will stay off for a good period of time.”*³ Wars also affect the macroeconomy through other channels than fiscal policy, such as labor force participation or uncertainty, obscuring the direct effect of government spending.⁴

On top of these endogeneity concerns, identification is also complicated by anticipation effects (Ramey, 2011b; Mertens and Ravn, 2012; Leeper et al., 2013). Implementation and legislative lags imply that changes in government spending are often known in advance. If a change in government spending is anticipated, then part of its effects on the economy may occur in response to *news*, before any actual spending takes place. News shocks violate the invertibility assumption imposed in structural vector autoregression (VAR) approaches (e.g., Blanchard and Perotti, 2002), which leads to confounding of anticipated and unanticipated policy changes.⁵

In this paper, we construct a new series of defense spending surprises to address these empirical challenges and isolate the effects of government spending on the macroeconomy. Our empirical strategy is to measure surprises in a stock index of large U.S. defense contractors in a narrow window around legislative events that reveal news about U.S. defense spending. The idea is that the financial market reaction to such events is informative about defense spending news because

¹Or equivalently, the dollar change in output that results when the government spends an additional dollar.

²The policy response could also be procyclical. For example, a weak economy may constrain the fiscal space and lead politicians to cut spending.

³The Wall Street Journal (2010).

⁴This limitation of using wars to identify fiscal multipliers is well-known in the literature (e.g., Ramey, 2011a).

⁵Intuitively, the identified shocks from this approach are forecast errors of government spending after projecting out quarterly lags of spending and output. If these lags do not capture news, then anticipated changes in spending may be misclassified as unanticipated. Ramey (2011b) shows that the Blanchard and Perotti (2002) shocks are predictable.

defense contractors' cash flows depend heavily on defense spending. If asset prices reflect all available information, then our surprises will capture only *unanticipated* changes in defense spending, since anticipated changes are already priced in. By measuring the surprise in a narrow window around events that specifically relate to U.S. defense spending, we can plausibly rule out that information about other aspects of the economy changes simultaneously. A recession may be the reason for defense spending cuts, but expectations about how the recession affects the economy should already be priced in before the cuts are announced.

We construct the defense spending surprises in two steps. The first step is to identify events that are likely to reveal news about U.S. defense spending. We focus on the two main ways U.S. Congress provides funding for defense: the regular annual defense appropriations bills and the supplemental defense appropriations bills that provide additional budget authority. We exclude consolidated bills ("omnibus bills") that have been common in recent years because they likely also communicate news about other types of spending. We collect data on the legislative actions and timeline of each individual bill from the Congressional Record since 1947. For each bill, we isolate three central event dates when different versions of the bill pass: the days when the House and the Senate pass their separate versions of the bill, and the day when the two chambers agree on the final version. These events resolve uncertainty and reveal information about which version of the bill will eventually become law. Information is also signaled in other ways: The event days often feature extensive debates about defense spending, and members of Congress propose, debate, and vote on amendments to the bill.

The second step is to quantify the defense spending news communicated on the event dates. Following Fisher and Peters (2010), we construct a stock index of large U.S. defense contractors using the Department of Defense reports on the top 100 defense contractors for each fiscal year. The underlying assumption is that an event that communicates news of an increase in future defense spending will increase the stock prices of U.S. defense contractors as market participants raise their expectations of future profits. We measure the surprise on each event date using the excess returns of defense contractors relative to a broader market index to purge the surprises of market fluctuations unrelated to defense. The resulting surprise series measures excess returns on 213 unique event dates.

To validate the defense spending surprises as an instrument for defense spending, we perform a set of diagnostic checks to assess their statistical properties. The variance of excess returns is more than twice as high on event dates as on regular trading days, indicating a disproportionate release of cash flow news about defense contractors. In line with the definition of macroeconomic shocks in Ramey (2016), we find that the surprises are mean-zero, serially uncorrelated, and not Granger-caused by standard macroeconomic and financial variables. We also provide specific event-study examples to illustrate the variation we exploit. For instance, in 1949, disagreement between the House and the Senate over Air Force funding significantly delayed the final passage of the defense appropriations bill. When the Senate eventually relented on its demands for Air Force cuts and passed the bill, defense contractors experienced substantial excess returns.

Next, we aggregate the daily defense spending surprises to quarterly frequency and estimate the macroeconomic effects of a defense spending surprise using local projections.

A positive defense spending surprise predicts a persistent increase in defense spending that is entirely passed through to total government spending. We find no spending response on impact, confirming that a surprise reflects news shocks about future spending. In the quarters following

a surprise, defense spending gradually rises and peaks after 3–4 years, staying elevated for several years. The high persistence of the response can be explained by the long duration of defense contracts (Cox et al., 2024) and is consistent with earlier findings that do not primarily rely on large wars (Blanchard and Perotti, 2002; Fisher and Peters, 2010). The dynamic response of *total* government spending closely resembles the defense spending response, indicating little to no change in other types of spending. The modest response of nondefense spending is also consistent with our event dates isolating news specific to defense spending. The magnitude of the shocks is economically meaningful: A one-standard-deviation defense spending surprise leads to a persistent increase in defense spending that peaks at 0.25 percent of GDP.⁶ The cumulative government spending response—the relevant first stage for estimating cumulative fiscal multipliers—is highly statistically significant with a first-stage F -statistic of 19 over a five-year horizon.

The identified changes in defense spending are partially tax-financed, and they do not seem to be counteracted by monetary policy “leaning against the wind.” The fiscal and monetary policy responses are central to interpreting our results. Standard macroeconomic models suggests that the size of fiscal multipliers depend crucially on how the spending is financed and the response of monetary policy. Models with non-Ricardian agents, such as the old textbook IS-LM model or the modern heterogeneous-agent New Keynesian literature (Kaplan et al., 2018; Auclert et al., 2024), predict that tax-financed spending is less effective in stimulating output relative to debt financing. Similarly, fiscal stimulus may be less effective if it is counteracted by contractionary monetary policy.

Output and consumption rise persistently following a defense spending surprise. On impact, the aggregate response is close to zero, mirroring the response of government spending. Output then rises faster than spending, which is consistent with agents reacting to news about future demand. Private consumption is crowded in, peaking at 50 cents per dollar of government spending. The consumption gains are concentrated in non-durables and services, with little movement in durables. The investment response is positive, but not statistically significant. Together, these patterns suggest that government spending stimulates private economic activity.

We summarize the effects of government spending on output by estimating cumulative fiscal multipliers—the cumulative output response divided by the cumulative government spending response to a defense spending surprise. Our baseline estimate is a cumulative fiscal multiplier of 1.2 measured over a five-year horizon. Our multiplier estimates are larger at shorter horizons because output tends to respond earlier than spending, consistent with agents responding to news. When we remove event dates that are relatively more likely to include confounding news, our multiplier estimates increase. Overall, we interpret our fiscal multiplier results as evidence that fiscal multipliers can be greater than one.

Standard theoretical models can rationalize the aggregate effects of defense spending we estimate. Our baseline model builds on the two-asset heterogeneous-agent New Keynesian (HANK) model of Auclert et al. (2024). We feed the empirically estimated path of government spending into the model and calibrate the policy rules so the policy responses are consistent with the data. This calibrated quantitative HANK model provides a good fit to the estimated impulse responses. More generally, models with high marginal propensities to consume and substantial nominal and real rigidities perform similarly. For instance, an analytically tractable HANK model, similar to

⁶Even though this implies a large cumulative spending response, the changes in spending identified from a defense spending surprise are much smaller than those associated with major wars such as World War II and the Korean War, which much of the previous literature relies on.

Angeletos et al. (2024) and Debortoli and Galí (2025), also rationalizes our results. In contrast, a representative-agent model with Ricardian households and low marginal propensities to consume struggles to replicate the observed increase in private consumption without a counterfactual decline in real interest rates.

Related Literature This article contributes to a large empirical literature that estimates aggregate fiscal multipliers. One strand of the literature employs structural VARs to identify government spending shocks (Blanchard and Perotti, 2002).⁷ The other main strand of the literature focuses on defense spending and wars and assumes that this type of variation is exogenous (Barro, 1981; Hall, 1986; Rotemberg and Woodford, 1992; Hall, 2009; Barro and Redlick, 2011). To account for anticipation effects, a connected literature has focused on identifying news shocks that capture changes in defense spending when they are announced (Ramey and Shapiro, 1998; Fisher and Peters, 2010; Ramey, 2011b; Ben Zeev and Pappa, 2017; Melosi et al., 2025). In a seminal contribution, Ramey (2011b) uses a narrative approach to create a quarterly series of defense news shocks based on newspaper articles.⁸ We contribute to this literature by providing a new series of defense news shocks based on a high-frequency identification approach. Overall, the time series literature tends to find aggregate fiscal multiplier estimates in the range 0.6–1.0 (Ramey, 2019; Ilzetzki, 2025). Our empirical evidence suggests that aggregate fiscal multipliers can exceed one and that private consumption in particular is crowded in.

Closest to our paper is Fisher and Peters (2010), who compute the accumulated monthly excess returns of defense contractors and order this variable last in a recursive VAR. We also use asset price responses of defense contractors to measure defense spending news. However, they use *all* variation in excess returns, whereas we only use the variation that occurs on specific event dates. Excess returns fluctuate for many reasons unrelated to U.S. defense spending. Using all variation introduces noise that reduces statistical power and could be endogenous to the business cycle.⁹ For example, some of the largest daily excess returns in their sample occur in response to financial crises. Presidential elections, which likely reveal both defense and other news, also often lead to substantial excess returns.

Our work is also connected to a literature that estimates relative (or local) fiscal multipliers in the cross-section using regional data. Nakamura and Steinsson (2014) exploit cross-sectional variation in U.S. states exposure to national military buildups. Other papers that study defense spending in this context include Hooker and Knetter (1997), Dupor and Guerrero (2017), Auerbach et al. (2020, 2024), Biolsi (2019), Hebous and Zimmermann (2021), Brunet (2024), and Amodeo and Briganti (2025).¹⁰ The identification assumptions required in the cross-section are more plausible, but translating relative fiscal multipliers into aggregate fiscal multipliers requires a model (see also

⁷See also Edelberg et al. (1999), Burnside et al. (2004), Galí et al. (2007), Perotti (2007), Monacelli and Perotti (2008), Mountford and Uhlig (2009), Auerbach and Gorodnichenko (2012), Corsetti et al. (2012), Ilzetzki et al. (2013), and Caldara and Kamps (2017).

⁸Recently, Antolin-Diaz and Surico (2025) study the long-run effects of government spending using these shocks. Antonova et al. (2025) study the role of reallocation of resources between sectors in military buildups.

⁹Ramey (2016) finds first-stage F -statistics close to zero using the Fisher and Peters (2010) shocks in her specification.

¹⁰There is a large literature estimating local fiscal multipliers using other types of spending and transfers, e.g. Chodorow-Reich et al. (2012), Shoag (2013), Acconcia et al. (2014), Suárez Serrato and Wingender (2016), and many others. See Chodorow-Reich (2019) for an overview.

Farhi and Werning (2016)). Chodorow-Reich (2019) reviews this literature and finds a median estimate of 1.8. Our evidence helps reconcile the relatively large local fiscal multipliers with the low estimates that dominate the time series literature.

Methodologically, our paper relates to the literature that uses high-frequency asset price surprises to identify macroeconomic shocks. This approach to identification was first applied to monetary policy surprises (see, e.g., Gürkaynak et al., 2005; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018) and has since been applied to several other settings.¹¹ Our paper is the first to apply these methods to study the macroeconomic effects of government purchases (i.e. government consumption and investment in national accounts) and estimate fiscal multipliers.¹² Recent work has studied the effects of deficit news on asset prices (Cotton, 2024; Gomez Cram et al., 2025; Wiegand, 2025). Wiegand (2025) estimates effects of deficit news using events when the Congressional Budget Office (CBO) releases new forecasts. Our empirical methodology differs in that we use the asset market surprises to quantify the news being communicated, whereas Wiegand instead assumes that changes in CBO forecasts are unanticipated and studies how asset price respond as an outcome. Hazell and Hobler (2025) combine narrative and high-frequency methods to study the effect of deficit news (mostly transfers) from a Senate election runoff on inflation.

Roadmap The article proceeds as follows. In section 1.1, we lay out our identification strategy and describe how we construct defense spending surprises. In Section 1.2, we present our empirical specification and estimate the dynamic effects of defense spending surprises. In Section 1.3, we estimate cumulative fiscal multipliers and discuss how our results compare to the previous literature. In Section 1.4, we show that a standard heterogeneous-agent New Keynesian model can match our empirical evidence. Section 1.5 concludes.

1.1 A New Measure of Defense Spending News

In this section, we construct a new series of defense spending news shocks using a high-frequency identification approach. We call them *defense spending surprises* to emphasize that they reflect financial market surprises. We begin by identifying legislative events in the U.S. federal budget process that are likely to communicate news about defense spending. We then construct a stock index of U.S. defense contractors, building on work by Fisher and Peters (2010). Finally, we measure the surprise in the excess returns of the stock index in daily windows around each event. We provide concrete examples to illustrate the underlying variation we exploit and discuss the statistical properties of the surprise series.

¹¹Examples include oil supply (Känzig, 2021), carbon pricing (Känzig, 2023), dollar swap lines (Kekre and Lenel, 2025), bank regulation (Drechsel and Miura, 2025), Treasury supply (Bi et al., 2025), economic data releases (Boehm and Kroner, 2025), and supply chains (Känzig and Raghavan, 2025).

¹²Military spending almost exclusively consists of government consumption and investment (as opposed to transfers). Deficits are broader in that they may arise due to changes in either government consumption and investment, transfers, or taxes.

Institutional Background: Defense Appropriations and Contracting

National defense accounts for more than half of federal government spending and explains a majority of the variation in total government spending.¹³ The U.S. Department of Defense is primarily funded by the annual defense appropriations bill, the largest of the twelve appropriations bills passed by Congress each fiscal year. This subsection describes the institutional details of how defense spending is funded in the United States and provides the basis for how we select event dates.

The Defense Appropriations Process The House and Senate Appropriations Committees each have a Subcommittee on Defense that is responsible for developing, drafting, and managing the consideration of the defense appropriations bill. The House and Senate subcommittees generally operate in parallel.¹⁴

After the President submits the annual budget request to Congress, the subcommittees hear testimonies from senior Department of Defense officials and other relevant parties. Members of Congress outside the committees may also submit requests and recommendations. Following these hearings, the subcommittees draft and mark up the defense appropriations bill and report it to the full appropriations committees, who may also make amendments before reporting the bill for floor consideration.

Ordinarily, the bill is first considered in the House, which typically occurs in May or June. After being publicly debated and amended on the floor, the House passes the bill. The Senate then subsequently debates, amends and passes its own version of the bill. Any disagreements between the House and Senate must be resolved before the bill can be presented to the President. For most of the bills we consider, the House and Senate agree to form a conference committee to negotiate the final version of the bill and report the results in a conference report. Alternatively, the two chambers can exchange amendments until they arrive at identical measures. The final version of the bill is then presented to the President for their signature or veto.

Supplemental Defense Appropriations Outside of the regular appropriations process, additional defense spending can be funded by supplemental defense appropriations. These measures traditionally cover urgent needs, but they have also been used to regularly pay for military campaigns. For example, the wars in Iraq and Afghanistan in the 2000s were to a large extent financed by supplemental appropriations. Congress passed one major “emergency” supplemental defense bill per year during 2002–2010. Since these bills evolved into a regular annual stream of war funding, they were largely anticipated. Supplemental spending bills follow a process similar to the one for the regular defense spending bills, though often at an accelerated rate. For example, the Supplemental Defense Appropriations Act 2005 was introduced and passed in the House on March 11 and 16. The Senate passed the bill on April 21, and a conference report was agreed to on May 10; the bill was signed into law the next day.

Omnibus Appropriations Bills and Continuing Resolutions Historically, Congress has usually passed each regular appropriations bill separately and on time. However, since the

¹³A linear regression of total government spending on defense spending in quarterly growth rates yields an R^2 of 0.72 in our sample period 1947–2019.

¹⁴See Saturno and Lynch (2023, 2024) for a more detailed overview.

2000s, the appropriations process has often deviated from standard practices. First, it has become increasingly common to combine several of the twelve appropriations bills into a single measure, a so-called consolidated, or “omnibus”, appropriations bill. Second, Congress has increasingly failed to pass the regular appropriations bills before the start of the fiscal year on October 1st. To avoid government shutdowns, Congress has often temporarily extended funding at current levels by passing a continuing resolution (“CR”).

Defense Contractors and Procurement The Department of Defense relies heavily on contractors to supply military equipment and other goods and services (Neenan, 2024a,b). Defense contracts account for the majority of federal contracts, and federal contracts in turn explain most of the variation in total government spending (Cox et al., 2024). Defense contracts are approximately evenly split between goods and services. The Department of Defense awards contracts to hundreds of firms, yet a majority of the (dollar-weighted) contracts are awarded to a handful of large defense contractors such as Lockheed Martin, General Dynamics, and Raytheon.

Defense contracts often span several years. Cox et al. (2024) document that the median federal contract duration is 3–4 years and the 90th percentile is around 12 years. Even though defense appropriations bills are specific to a fiscal year, they provide budget authority several years into the future. The main reason is that procurement of military investment goods, such as aircraft, ships, and armored vehicles,¹⁵ is funded under a *full funding policy*, which requires that the entire procurement cost is appropriated in the year that item is procured. Congress has imposed this policy on defense spending since the 1950s.¹⁶

Event Dates

When is news about defense spending communicated to the public? The ideal event for our empirical strategy is one that reveals information about U.S. defense spending without simultaneously conveying news about other factors that influence economic conditions. Our analysis focuses on key legislative events in the defense appropriations process that directly determine U.S. defense spending. Alternative political events—such as State of the Union addresses or election outcomes—also convey information about defense priorities, but they are often confounded by other policy developments. For instance, it is difficult to isolate the effects of defense spending news from a presidential election if presidents who tend to raise defense spending also tend to pursue, say, financial deregulation.

We concentrate on the two main types of legislation that provide funding for defense: regular defense appropriations bills and supplemental defense appropriations bills. For each bill, we collect detailed information from the Congressional Record on its timeline and legislative actions, from initial introduction to final enactment. We exclude defense appropriations that are passed as part of omnibus spending packages, as well as continuing resolutions, which provide only temporary funding and are typically anticipated.

We identify three central legislative events that mark the progression of each bill through Congress: passage in the House, passage in the Senate, and the approval of the final version of

¹⁵Satellites, missiles, weapons, and ammunition also fall under this category.

¹⁶See Daggett and O’Rourke (2007).

the bill. Passage is an important event because it resolves uncertainty and because the approved bill is reported. These events are also highlighted as milestone events in the Congressional Record (e.g., in the overview section of the Daily Digest publication) and are most likely to be reported in the news. In addition, debates and amendments tend to cluster on these days. For example, the House considers the regular appropriations bill under a so-called “special rule” that governs debate and amendment. Since debate is usually limited to 1–2 hours, most of the activities related to the House version of the bill occur on the same day. When the date of a legislative event is not available from the Congressional Record’s online database, we either manually locate the event in the full texts of the Congressional Record, or consult *CQ Almanac*, which provides detailed narrative overviews of each bill’s content and legislative history.

Consider for example the Department of Defense Appropriations Act 1982, which was debated, amended, and passed in the House on November 18, 1981. On this event date, 28 amendments were proposed and voted on by House members, with 17 of them passing. For example, one amendment proposed to reduce funds for procurement and research and development by two percent, but failed by a close vote of 197–202. Another amendment proposed “*to decrease from 13,957,598,000 dollars to 12,156,598,000 dollars funds for Air Force aircraft procurement. The intent of the amendment is to strike funds for the B-1 bomber.*” The B-1 bomber program was one of the most controversial defense programs of the late Cold War, both politically and within the defense establishment (Norman, 1988). In the end, most of the proposed cuts failed, whereas the House passed other amendments to increase funding for Navy shipbuilding and Air Force missile procurement, among other things.

A Stock Index of Defense Contractors

To quantify the defense spending news released on event dates, we construct a stock index of large U.S. defense contractors following Fisher and Peters (2010). The index is based on the annual Top 100 Contractor Reports published by the Department of Defense since 1958. These reports rank the top 100 companies that were awarded the largest total dollar volume of prime contracts during a fiscal year (including their subsidiaries). The reports include information on firm names, their rank, and the dollar value of their awarded contracts. We digitize this information and string-match the firm names to daily financial data from the Center for Research in Security Prices (CRSP) database. We complement the data with annual sales data from the Compustat database.

Which firms should be included in the index to best capture defense news? On the one hand, it is desirable to include as many firms from the reports as possible to minimize noise from individual firms. On the other hand, we want to exclude firms that are not perceived as defense contractors by financial markets. For example, some firms may only be listed temporarily (such as Moderna and Pfizer during the COVID-19 pandemic), or defense contracts may only account for a tiny fraction of a firm’s total sales (such as FedEx). Our goal is not necessarily to include as large a fraction of defense contracts as possible, but rather to extract as clear a signal about defense spending news as possible.

We impose the following requirements for a firm to be included in our index. First, we require that firms appear in the reports at least five times. Second, we restrict our analysis to firms that belong to a defense-related industry as defined by the list of four-digit SIC codes in Fisher and Peters (2010). This classification is relatively broad but excludes firms that are not primarily engaged in defense-related activities such as health insurance companies. To ensure that the firms

are exposed to defense contracts, the third requirement is that defense contracts must account for at least 20 percent of the firm's total sales on average. Fourth, we target relatively large contractors by requiring that a firm's defense contracts must account for at least one percent of the total contract volume on average. Finally, we exclude any private and foreign firms that are not covered by CRSP. This requirement is not very restrictive, as large defense contractors are almost exclusively publicly traded U.S. companies.

Our baseline set of firms replicates the list of firms considered in Fisher and Peters (2010) by additionally imposing that firms must appear in the top three at some point, and by excluding Boeing.¹⁷ We weight the firms by their average share of total contract volume. Our stock index includes 14 individual stocks and covers 35 percent of defense contracts on average. Figure A.2 shows that the share of contracts awarded to the firms in the index has been relatively stable over time. Following Fisher and Peters (2010), we extrapolate our list of defense contractors back to 1947. We verify that the firms in the list were large defense contractors in this period by consulting official accounts from the Korean War (United States Senate, 1951).

We construct alternative stock indices that do not impose the top 3 requirement; that include Boeing; and that select a higher number of firms based on recent reports (for example, firms appearing in the top 25 in the last three years). Overall, all of these alternative indices are highly correlated with our baseline index because the defense industry is dominated by the top 5–10 largest defense contractors, such as Lockheed Martin, Raytheon, General Dynamics, Northrop Grumman, and most of these firms and their predecessors have existed for the whole sample period.

Measuring Defense Spending Surprises

On each event date, we measure the excess returns of defense contractors relative to a broader stock market index. Conceptually, an event that communicates news of an increase in future defense spending leads to excess returns of U.S. defense contractors as market participants raise their expectations of future profits.

Let P_{τ}^{def} denote the closing price of our stock index of defense contractors on a trading day τ . Define the daily return as $r_{\tau}^{def} \equiv \log P_{\tau}^{def} - \log P_{\tau-1}^{def}$ where $\tau - 1$ is the previous trading day. Similarly, let P_{τ}^{mkt} be the closing price of the value-weighted market index in the CRSP database with daily return r_{τ}^{mkt} . Now consider N event dates $\{\tau_1, \dots, \tau_N\}$ during which news about defense spending are communicated. For each event date τ_i , we define the daily defense spending surprise as the excess return:

$$surprise_{\tau_i} = r_{\tau_i}^{def} - r_{\tau_i}^{mkt}.$$

When an event falls on a non-trading day, we measure the return on the next trading day.

We use daily event windows as in e.g. Känzig (2021) and Drechsel and Miura (2025). Fiscal policy is communicated less clearly than monetary policy, so the exact timing can be difficult to pin down and markets may need some time to process or interpret the announcement. Furthermore, exact time stamps from the Congressional Record and intraday financial data are unavailable for most of our sample period.

¹⁷Fisher and Peters (2010) drop Boeing because of its exposure to the relatively volatile commercial airline industry. Our approach is not sensitive to this choice because we focus only on specific event dates that are dominated by announcements about defense.

Returns vs. Excess Returns Our goal is to construct an *instrument* for defense spending that can be used to identify fiscal multipliers. Therefore, we only need to identify the surprises up to scale. While the idea behind using a defense stock index is that spending news affect returns as cash flow news, we do not strictly need to assume that this channel is the only channel. For example, expected market cash flows might react to defense spending news (e.g. due to multiplier effects), but as long as the surprise we measure reflects a reaction to defense news, and not other confounding shocks, the surprise variable remains a valid instrument. We theoretically decompose the surprises to lay out this point in detail and show that we can allow for market cash flows and discount rates to react in Appendix A.1. Still, standard asset pricing theory suggests that the surprises *do* reflect changes in expected cash flows, which is our preferred interpretation.

Alternative measures of excess returns If our defense stock index does not move one-for-one with the market, then simply subtracting the market return may be insufficient to control for market fluctuations. To address this concern, we also compute excess returns based on a standard CAPM model and the Fama–French three-factor model.¹⁸ Overall, these alternative measures of excess returns are nearly perfectly correlated with our baseline measure and consequently produce close to identical results. We find that the estimated market beta for our defense stock index is close to one (0.90). This finding is consistent with defense index funds such as BlackRock’s iShares U.S. Aerospace & Defense ETF index, which reports a market beta of 0.95 as of September 2025.

Examples of Defense Spending Surprises

In this subsection, we provide two concrete event study examples using events that communicated defense news that resulted in excess returns.

Example 1: National Military Establishment Act 1950 Our first example is the final Congressional approval of the National Military Establishment Act for the fiscal year 1950—the regular defense appropriations bill of that year—on October 18, 1949. *CQ Almanac* provides a detailed narrative account of the progression of this bill. At the time, the measure was the largest peacetime appropriations bill in U.S. history. In early October, after weeks of lengthy sessions, the conference committee reported that they had reached agreement on 73 of the 100 items that the House and the Senate bills disagreed on. The main source of disagreement concerned four amendments in which the Senate proposed substantial cuts to the Air Force. Both chambers refused to give in, so the conference committee went back to work. In the end, the Senate gave in at the end of a night session on October 17. The Air Force cuts did not materialize and the bill was passed the next day.

Figure 1.1a plots the cumulative excess returns of the defense stock index in October 1949. When the Senate gave in on their demands to cut the Air Force, defense stock rallied with excess returns of about four percentage points. In this example, it was most likely the passage itself that

¹⁸We estimate these measures of excess returns using daily data on all trading days for our sample period 1947–2019. For the CAPM model, we regress the returns of the defense stock index on the market index and obtain the excess return as the residual. The Fama–French specification additionally controls for size and value risk factors.

communicated news by resolving uncertainty about Air Force funding. All other disagreements had already been agreed to.

Example 2: Department of Defense Appropriations Act 1990 We now consider an example of an event that led to a large negative surprise: the annual defense appropriations bill that provided funding for the fiscal year 1990. The first version of the bill was introduced in the House on August 1, 1989. After debating and amending the bill, the house passed their version of the bill three days later on August 4. The Senate version of the bill was passed in late September with several significant amendments. To resolve the differences between the House and Senate versions of the bill, a conference committee with members was appointed to. Usually, the conference committee reaches an agreement, but in this case the conference report came out of committee with more than 101 disagreements on Monday November 13, so the House and the Senate had to debate the bill again to resolve their differences before reaching an agreement the following weekend.

Figure 1.1b shows the cumulative excess returns of our defense stock index in November 1989. The stock prices of major defense contractors began to decline gradually during the week, consistent with investors revising expectations about future defense budgets. The final version of the bill was passed the following weekend after intense negotiations and public discussion about potential reductions in defense spending. When markets reopened, defense contractor stocks fell sharply. Broader market indices were largely unchanged during this period, suggesting that the decline in defense stocks reflected sector-specific news rather than general macroeconomic developments.

Defense Spending Surprise Series


Figure 1.2 plots the daily defense spending surprises on each of our 213 event dates. The mean surprise is 0.00, suggesting that there is no evidence of systematic excess returns on our event dates. The standard deviation of the surprises is 1.57 percentage points, which indicates that the variance is more than twice as high on event dates compared to a sample that includes all trading days.

Are the surprises we identify autocorrelated or predictable? Figure A.5 shows the the autocorrelation and partial autocorrelation functions of our surprise series. A standard Ljung-Box test indicates that we cannot reject the null of no autocorrelation with a p -value of 0.84. To assess predictability, we conduct Granger causality tests in line with the previous literature on news shocks (e.g. Ramey, 2011b; Känzig, 2021). Figure A.6 shows that standard macroeconomic and financial do not significantly predict our defense spending surprises.

1.2 The Dynamic Effects of Defense Spending News


In this section, we present our empirical specification and estimate the effects of a defense spending surprise on the macroeconomy. We show that a positive defense spending surprise predicts a persistent increase in future defense spending that is entirely passed through to total government spending. The identified spending increases are financed with a mix of taxes and debt, and we find muted responses of nominal interest rates and inflation. We find a sizable and persistent increases in output and consumption, as well as a decline in the unemployment rate. In the next section, we

Figure 1.1: Event Study Examples



Notes: This figure illustrates two examples of defense spending surprises by plotting the cumulative excess returns of the defense stock index around two of our event dates. In both panels red triangles represent market returns and blue circles represent excess returns. Panel (a) shows the cumulative returns around the final passage of the National Military Establishment Act 1950 on October 18, 1949 (with October 6 excess returns normalized to zero). Panel (b) shows in blue circles the cumulative returns around the final passage of the Department of Defense Appropriations Act 1990 in November 1989 (with October 31 excess returns normalized to zero).

Figure 1.2: Daily Defense Spending Surprises



Notes: This figure shows the daily defense spending surprise series, constructed as the excess returns of a stock index of defense contractors on 213 event dates on which different versions of defense appropriations bills passed in the Congress.

estimate cumulative fiscal multipliers to summarize the dynamic responses of government spending and output.

Empirical Specification

We now estimate impulse response functions to a defense spending surprise using local projections (Jordà, 2005a). We aggregate our daily series of surprises to quarterly frequency by summing them within each quarter. We then convert them from excess returns to dollar surprises by multiplying them by lagged real defense spending per capita G_{t-1}^{def} . Our main interest lies in the dynamic responses of government spending, output, and their components. We estimate regressions

$$\frac{X_{t+h} - X_{t-1}}{Y_{t-1}} = \alpha_h + \beta_h \frac{surprise_t}{Y_{t-1}} + \sum_{j=1}^p \gamma_h^{j'} \mathbf{W}_{t-j} + u_{t+h}, \quad (1.1)$$

for each horizon $h = 0, 1, \dots, H$, where X_{t+h} denotes an outcome variable measured in real dollars per capita (e.g. government spending or output) measured at time $t + h$, Y_{t-1} is lagged GDP in real dollars per capita, $surprise_t$ is the defense spending surprise in real dollars per capita, $\mathbf{W}_t$ is a vector of control variables, and u_{t+h} is an error term.

We measure the components of aggregate demand in percent of GDP by normalizing by Y_{t-1} because we are ultimately interested in multipliers (rather than elasticities). Similarly, we normalize the surprise so the right-hand side variable of interest $\frac{\text{surprise}_t}{Y_{t-1}}$ is also measured in percent of GDP. We emphasize that we interpret the surprise series as an *instrument* for changes in defense spending, and we do not claim to identify the scale. We scale the defense spending surprises such that the impulse response of defense spending peaks at one percent of GDP. Since multipliers are unit-less, this normalization of the reduced-form responses to a surprise can also be interpreted as a peak of one dollar per capita.

We estimate the regressions in long-differences for two reasons. First, this formulation is the natural extension of estimating contemporaneous multipliers in first-differences as in Hall (2009) and Barro and Redlick (2011). Second, the simulation evidence of Jorda and Taylor (2025) suggests that long-differencing suppresses the small sample bias highlighted by Herbst and Johannsen (2024), and hence long-differencing obviates the need to apply a bias correction procedure.

We choose our control variables based on Montiel Olea et al. (forthcoming) who recommend including several lags of the outcome, the regressor, as well as strong predictors of either. We control for $p = 6$ lags of the surprise series, output, defense and total spending, inflation, short-term interest rate, and the unemployment rate.¹⁹ When the outcome variable is not already in this list of controls (e.g., consumption), we also include six lags of this variable.

We use quarterly data for the U.S. 1947–2019 obtained from the FRED database. We use standard variables from the national income and product accounts (NIPA) to measure government spending, GDP, and private consumption and investment. We express these variables in real 2017-dollars per capita terms. We report heteroskedasticity-robust standard errors (Montiel Olea and Plagborg-Møller, 2021).²⁰

Defense Spending and Total Government Spending Responses

Our first result is that a defense spending surprise predicts a persistent increase in future defense spending that is passed through to total government spending. The first panel of Figure 1.3 shows the estimated response of defense spending, with the peak response normalized to one. The dependent variable is $\frac{G_{t+h}^{def} - G_{t-1}^{def}}{Y_{t-1}}$, where G_t^{def} is level of defense spending and Y_t is the level of GDP. The impact response is a precisely estimated zero effect. The absence of an immediate spending response is consistent with our interpretation that the surprises capture *news*. In the quarters following a surprise, defense spending increases steadily for 2–3 years, after which it stays elevated for several years. The magnitude of the shocks is economically meaningful: A one-standard-deviation defense spending surprise leads to a persistent increase in defense spending that peaks at 0.25 percent of GDP.

The second panel of Figure 1.3 shows that the response of total government spending (the sum of defense and nondefense spending) closely mirrors the defense spending response. Our surprise series almost exclusively predicts changes in defense spending, not other types of spending. The

¹⁹Inflation is year-on-year percentage change in the GDP deflator, and the short-term interest rate is the yield on 3-month Treasury bills.

²⁰This choice is conservative as heteroskedasticity and autocorrelation robust standard errors tend to yield smaller standard errors in our setting.

confidence bands are only slightly wider for total spending, consistent with the well-known fact that defense spending accounts for most of the variation in U.S. government spending.

What explains the high persistence of the spending responses? First, U.S. defense spending is generally a highly persistent process outside of large wars such as World War II and the Korean War.²¹ In addition, more recent military campaigns such as the Vietnam War, the 1980s Carter-Reagan buildup, and the wars in Afghanistan and Iraq, featured longer buildups and drawdowns.

A second reason is that defense contracts have long durations. Cox et al. (2024) document that the median federal contract is 3–4 years with a 90th percentile of approximately 12 years. If a defense spending surprise reflects a decision to build new fighter jets or Navy ships, then we should expect high persistence. Measurement and implementation lags may also play a role. Recently, Briganti et al. (2025) have shown that government spending, as measured in national accounts (NIPA), lags budget authorization by 3–4 quarters. The timing in NIPA may reflect the delivery time of the final good to the government, or the timing of the associated Treasury outlays (which can occur gradually for large defense projects). However, while this timing mismatch likely contributes to the persistence of our defense spending response, it is unlikely to be the main driver because we find persistence effects that last several years.

We estimate dynamic responses up to a maximum horizon of 6 years throughout. While we have sufficient statistical power to estimate the spending responses in Figure 1.3 at longer horizons, the responses of other macroeconomic outcomes are difficult to estimate far into the future because they are influenced by many other factors.²² We show the spending responses at longer horizons in Figure A.8. When we restrict our sample to the post-1955 period, which excludes the Korean War, the spending response is more transitory, although still highly persistent (see Figure A.9).

Responses of Taxes and Deficits

By definition, an increase in defense spending must be financed in one of three ways: decreasing other types of spending, increasing tax revenue, or by running a deficit. This result follows from the government budget identity:

$$\underbrace{G_t^{def} + G_t^{nondef}}_{\text{Gov't purchases}} = \underbrace{T_t}_{\text{taxes}} + \underbrace{B_t - (1 + r_{t-1})B_{t-1}}_{\text{deficit}} \quad (1.2)$$


where G_t^{def} and G_t^{nondef} denote defense spending and nondefense spending, T_t denote tax revenue (net of transfers), B_t is the debt issued in period t , and r_t is the real interest rate. Macroeconomic models with non-Ricardian agents and high marginal propensities to consume (MPCs), such as the old textbook IS-LM model or the modern heterogeneous-agent New Keynesian literature,²³ predict that the size of fiscal multipliers strongly depend on how the spending is financed. When spending is debt-financed, current disposable income rises without an immediate tax increase. Non-Ricardian

²¹For example, Barro (1981) observes that “*defense spending associated with wars is largely transitory, while other changes in defense spending turn out to be predominantly permanent.*”

²²Antolin-Diaz and Surico (2025) study the long-run effects government spending using the shock series from Ramey and Zubairy (2018) and a Bayesian VAR with 60 quarterly lags. Even in this setting, they find that variation from large wars (World War I, World War II or the Korean War) is necessary to have sufficient power with 90 percent confidence bands.

²³See, e.g., Galí et al. (2007), Kaplan et al. (2018), and Auclert et al. (2024).

Figure 1.3: Defense Spending Surprises Predict Future Defense Spending



Notes: This figure plots the response of defense spending and total government spending to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars). The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.

high-MPC households therefore spend more today, so the boost to consumption is larger than under tax financing.

Figure 1.4 shows how the identified spending changes are financed. The first panel replicates the response of defense spending from the previous subsection. The second panel plots the response of nondefense spending, which accounts for the small difference between defense and total government spending. The bottom two panels show the responses of taxes (net of transfers) as well as the federal deficit. For taxes, the outcome variable is $\frac{T_{t+h}-T_{t-1}}{Y_{t-1}}$, where T_t is the current federal tax receipts minus current federal transfers payments. For deficits, the dependent variable is $\frac{D_{t+h}-D_{t-1}}{Y_{t-1}}$, where D_t is the federal deficit including interest payments (defined as current expenditures minus current receipts in NIPA). These responses are less precisely estimated than the spending responses in the first row and indicate that the spending increases are financed with a mix of taxes and debt. The tax response peaks at 50 cents and is borderline significant when the defense spending response peaks at 1 dollar. This increase suggests that, on average, taxes and transfers do a play role in financing the identified defense spending increases. Hence, the government spending increases we identify should not be interpreted as a purely deficit-financed as in the literature on local fiscal multipliers (Nakamura and Steinsson, 2014; Farhi and Werning, 2016; Chodorow-Reich, 2019).

One likely explanation for why it is difficult to obtain precise estimates for tax and deficit responses is time variation in the fiscal rule. What matters for fiscal multipliers is whether government spending is tax-financed or debt-financed *at the margin*. It is well-documented that the financing of defense spending and wars have varied considerably over time (Goldin, 1980; Ohanian, 1997). For

example, the U.S. government implemented tax rate increases and ran a balanced budget during the Korean War, whereas later episodes such as the Carter-Reagan buildup and the wars in Afghanistan and Iraq, were financed almost entirely through deficits (Barro and Redlick, 2011). Auerbach and Yagan (2025) provide empirical evidence that, in recent decades, Congress has stopped responding to higher projected deficits with fiscal adjustments. Consistent with this evidence, we show in Figure A.12 that excluding the post-2000 part of our sample increases both the magnitude and the precision of the estimated tax response. We also find some evidence that excluding the pre-1955 period (which includes the Korean War) leads to a larger deficit response, but only after 4–6 years (see Figure A.13). Finally, we note that interpreting the fiscal policy response is complicated by the presence of automatic stabilizers (e.g., McKay and Reis, 2016) and by feedback from output to tax revenues through a “tax-base channel” (Angeletos et al., 2024). Part of the estimated increase in taxes may reflect these mechanical forces rather than deliberate tax policy changes, such as the major tax bills enacted during the Korean War.

Responses of Monetary Policy and Inflation

Are the spending increases we identify dampened (or amplified) by a monetary policy response? Textbook New Keynesian models predict that the size of fiscal multipliers depend crucially on the monetary policy response (Woodford, 2011; Christiano et al., 2011; Auclert et al., 2024). In Figure 1.5, we present the responses of the short-term nominal interest rate and the PCE inflation rate to a defense spending surprise. The responses are normalized such that the defense spending response peaks at one percent of GDP (corresponding to a surprise of about four standard deviations). We do not find evidence that spending increase are met with a strong monetary policy response. The estimated response of the nominal interest rate stays close to zero, suggesting that monetary policy does not respond. The inflation response is not statistically significantly different from zero and ranges from slightly positive in the first two years after a surprise to slightly negative in years 2–4. These findings are generally in line with the previous empirical literature: Researchers often find that nominal interest rates do not respond to identified government spending shocks,²⁴ whereas the evidence on inflation responses is mixed (Jørgensen and Ravn, 2022).


Output, Consumption, and Investment Responses

In Figure 1.6, we present the responses of GDP and its components in the national income identity $Y = C + I + G + NX$. The top row shows that a defense spending surprise generates a substantial increase in output that largely has the same shape as the spending response. Even though the spending path is back-loaded, output increases gradually and does not increase on impact. Yet, output rises faster than spending: while the spending response only reaches its peak after three years, output is almost at its peak after only one year. A potential explanation is that agents are responding to news about future government spending by increasing private expenditures.

The second row shows the private consumption responses to a defense spending surprise. Private consumption accounts for about 70 percent of GDP, so the consumption response is an key determinant of fiscal multipliers. The left panel shows that households gradually increase their

²⁴See, for example, the discussion in Angeletos et al. (2024).


Figure 1.4: Responses of Taxes and Deficits to a Defense Spending Surprise



Notes: This figure plots the response of defense spending, nondefense spending, federal taxes and transfers, and the federal deficits to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars). For taxes, the outcome variable is $(T_{t+h} - T_{t-1})/Y_{t-1}$, where T_t is the current federal tax receipts minus current federal transfers payments as defined in NIPA. For deficits, the dependent variable is $(D_{t+h} - D_{t-1})/Y_{t-1}$, where D_t is the federal deficit including interest payments (defined as current expenditures minus current receipts in NIPA). The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.

consumption of non-durables and services. Once again, we do not find evidence of an immediate jump on impact. At first glance, the absence of an immediate response may appear inconsistent with the behavior of forward-looking household in standard macroeconomic models. However, this finding is less puzzling given the absence of a real interest rate response that we documented in the previous section. The consumption response is persistent and peaks at a level that is about half of the spending response. In contrast, consumption of durables does not respond to a defense

Figure 1.5: Muted Monetary Policy Response



Notes: This figure plots the responses of the nominal interest rate (3-month Treasury bill rate) and PCE inflation (year-on-year) to a defense spending surprise. The responses are normalized so the defense spending responses peaks at one percent of the GDP. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.


spending surprise. Overall, private consumption is crowded in, rather than crowded out as predicted by standard neoclassical models (Woodford, 2011). This result contrasts with the literature that relies on large wars for identification, as these papers find negative consumption responses (Barro, 1981; Hall, 1986; Ramey and Shapiro, 1998; Hall, 2009; Barro and Redlick, 2011; Ramey, 2011b).

The third row shows the responses of investment and net exports. The investment response is moderately positive but statistically insignificantly different from zero. For completeness, we also show the response of net exports, which fall gradually to a level that accounts for about half of the increase in consumption. This decline may reflect an appreciation of the U.S. dollar standard open economy models predict, and which Auerbach and Gorodnichenko (2016) provide evidence for. Alternatively, it may reflect government imports, or imports of intermediate goods by defense contractors. We note that the decline in net exports response suggests that foreign demand for defense goods does not seem to be a major driver of the output response.

Labor Market Responses


Figure 1.7 shows that the response of the labor market to a defense spending surprise that generates a defense spending peak of one percent of GDP. The first panel shows that the unemployment rate falls by about 0.5 percentage points. This effect is consistent with Okun's law, since GDP rises by about one percent in response to the same shock. The second and third panels show that the responses of employment-to-population ratio and hours worked mirror the unemployment response, but with wider confidence bands.

Figure 1.6: Government Spending Stimulates Private Economic Activity



Notes: This figure plots the responses of total government spending, output, non-durable and services consumption, durables consumption, investment, and net exports to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars). The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.

Figure 1.7: Labor Market Responses to a Defense Spending Surprise



Notes: This figure plots the responses of the unemployment rate, the employment-population ratio, and log hours worked to a defense spending surprise, normalized so the defense spending response peaks at one percent of GDP. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.

1.3 Aggregate Fiscal Multipliers

In this section, we summarize the effects of government spending on output using cumulative fiscal multipliers (Mountford and Uhlig, 2009; Farhi and Werning, 2016). The cumulative fiscal multiplier is the cumulative output response divided by the cumulative government spending response to a defense spending surprise. While the impulse responses contain strictly more information, the multiplier is useful for summarizing the effects of government spending on output.


We estimate cumulative fiscal multipliers using a local projections instrumental variables (LP-IV) approach (Stock and Watson, 2018; Ramey and Zubairy, 2018). For a given horizon h , we regress the cumulative output response on the cumulative government spending response:

$$\sum_{j=0}^h \frac{Y_{t+j} - Y_{t-1}}{Y_{t-1}} = \delta_h + \mathcal{M}_h \sum_{j=0}^h \frac{G_{t+j} - G_{t-1}}{Y_{t-1}} + \sum_{j=1}^p \gamma_h^{j'} \mathbf{W}_{t-j} + u_{t+h} \quad (1.3)$$

using our surprise series as an instrument. G_t denotes total government spending and Y_t denotes the level of GDP. The coefficient of interest is $\mathcal{M}_h$, which represents the cumulative fiscal multiplier at horizon h . The control variables $\mathbf{W}_t$ are the same we used to estimate the impulse responses in the previous section.

First Stage

To reliably estimate the cumulative fiscal multiplier at horizon h , the defense spending surprise must be a relevant instrument for the cumulative spending response $\sum_{j=0}^h \frac{G_{t+j} - G_{t-1}}{Y_{t-1}}$. The dashed red

Figure 1.8: First-stage F -Statistics

Notes: This figure plots robust first-stage F -statistics from testing the responses of defense spending and total government spending to a defense spending surprise. The solid blue line plots the F -statistics from pointwise tests that the impulse responses in Figure 1.3 are equal to zero at each horizon. The dashed red line plots the F -statistics from testing that the *cumulative* impulse responses (up to the given horizon) are equal to zero. The dashed red line is the relevant first-stage F -statistic for estimating cumulative fiscal multipliers using our defense spending surprises as an instrument.

line in Figure 1.8 plots the first-stage F -statistics at each horizon h .²⁵ These are the test statistics from testing the null that the integral under the impulse responses in Figure 1.3 equals zero at a given horizon. At short horizons, the instrument is weak because no actual spending has occurred yet. At longer horizons, the surprise series predicts more spending and the F -statistics rise. We use the total government spending variable to estimate fiscal multipliers. At a five-year horizon, the F -statistic is almost 20, which indicates that our surprises series is a sufficiently strong instrument for cumulative government spending.

Fiscal Multiplier Estimates

Table 1.1 presents our estimates of the cumulative fiscal multipliers $\mathcal{M}_h$ in Equation 1.3 at horizons of 2–6 years. The first column shows the results for our baseline specification, which uses all 213 event dates and the full postwar sample from 1947 to 2019. We estimate a cumulative fiscal multiplier of 1.2 over a five-year horizon. The multiplier estimates decline with the horizon because the estimated output response to a spending surprise is slightly more front-loaded than the spending response. Consistent with the rise in private consumption that we documented in the previous section, the

²⁵We report robust F -statistics. Since we are in the just-identified case, these are equivalent to the efficient F -statistic of Montiel Olea and Pflueger (2013).

multipliers are above one, suggesting that government spending crowds in private-sector economic activity.

In line with the positive consumption response documented in the previous section, our cumulative fiscal multiplier estimates are always above one. However, we recognize that we cannot reject fiscal multipliers of one—a common finding in the empirical literature on fiscal multipliers. Ramey (2016, Table 4) estimates cumulative fiscal multipliers using the identified shocks of Blanchard and Perotti (2002), Ramey (2011b), and Ben Zeev et al. (2023) using post-1947 quarterly data and likewise finds that none of the multiplier estimates are statistically significantly different from one. The local fiscal multiplier estimates Nakamura and Steinsson (2014, Table 2) fall in the range 1.4–1.9 with confidence bands that include one. One exception to this pattern is the extended defense news shocks series of Ramey and Zubairy (2018) that covers World War I, World War II, and the Korean War—three large and relatively sudden buildups. Using this series, together with either their specification or ours, we obtain a cumulative multiplier estimates around 0.7 that are precise enough to rule out one at the 5% level.

In columns (2) to (4) of Table 1.1, we investigate how our multiplier estimates depend on the selected event dates. In column (2), we remove the largest shock in our series, namely the emergency appropriations bill passed shortly after the September 11 attacks in 2001. A potential concern is that the event window around the passage of this bill includes confounding news and that our relatively large multipliers are driven by this event. The estimates presented in column (2) show this is not the case. Instead, our estimates increase. The event-study graph around this event in Figure A.7 confirms that most of the excess returns are driven by defense contractors’ stocks rather than market returns. More generally, we may worry that supplemental appropriations bills are more likely to violate the exclusion restriction if they happen turbulent times with lots of other news. To assess the role of supplemental appropriations, we exclude *all* supplemental appropriations bills and present the resulting multiplier estimates in column (3). We find that the multiplier estimates increase further to around 2. As expected, the first stage is somewhat weaker with F -statistics around 12 at longer horizons.

The final column of Table 1.1 reports our multiplier estimates for a sample starting in 1955, thereby excluding the Korean War period. At short horizons, we lack sufficient statistical power to estimate multipliers precisely, partly because the spending response is more back-loaded. At longer horizons, however, we find large multiplier estimates, consistent with earlier findings that the Korean War period is associated with smaller estimated multipliers Galí et al. (2007); Dupor and Guerrero (2017). At the five-year horizon, the lower bound of the 95 percent confidence interval is 0.83—above the cumulative multiplier estimates of Ramey (2011b), as reported in Ramey (2016).²⁶ An important contribution of our paper is that our defense spending surprise series allows us to estimate fiscal multipliers reliably in the post-1954 sample.

Taken together, these results suggest that government spending has sizable and persistent effects on output. Across specifications, our estimated cumulative multipliers remain above one. We interpret these findings as evidence that fiscal multipliers can exceed one.

²⁶Ramey (2011b) does not directly estimate cumulative fiscal multipliers, so we rely on the estimates from Ramey (2016), which uses the same shocks.

Table 1.1: Cumulative Fiscal Multiplier Estimates

Horizon (Years)	Baseline (1)	No September 11 (2)	Regular Appropriations (3)	Post-1954 (4)
2	1.35 (0.83) [$F = 8.04$]	2.02 (0.96) [$F = 7.38$]	3.74 (1.64) [$F = 4.13$]	6.99 (8.17) [$F = 0.9$]
3	1.29 (0.62) [$F = 12.1$]	1.71 (0.69) [$F = 11.8$]	2.66 (0.89) [$F = 8.18$]	4.03 (2.50) [$F = 3.44$]
4	1.27 (0.56) [$F = 16.3$]	1.57 (0.62) [$F = 16.1$]	2.31 (0.75) [$F = 11.4$]	2.98 (1.23) [$F = 7.65$]
5	1.17 (0.52) [$F = 18.9$]	1.39 (0.59) [$F = 17.7$]	2.00 (0.72) [$F = 12.3$]	2.44 (0.82) [$F = 10.1$]
6	1.12 (0.49) [$F = 20.6$]	1.29 (0.56) [$F = 18.0$]	1.82 (0.68) [$F = 12.4$]	2.09 (0.63) [$F = 12.4$]
Sample	1947–2019	1947–2019	1947–2019	1955–2019

Notes: This table reports our IV estimates of the cumulative fiscal multiplier $\mathcal{M}_h$ in Equation (1.3). Along the rows, we report estimates at different forecast horizons (in years). Along the columns, we estimate multipliers using different samples. The first column is our baseline sample, which includes all event dates from regular and supplemental appropriations bills 1947–2019. The second column excludes the event around the September 11 attacks. The third column drops all events related to supplemental appropriations and only considers regular appropriations bills. The final column drops the years before 1955 (which include the Korean War). Below each estimate we report the heteroskedasticity-robust standard error in parentheses, and the first stage F -statistic in brackets.

1.4 Aggregate Fiscal Multipliers in Standard Models

Our empirical analysis shows that government spending has substantial effects on economic activity. Are these results consistent with standard models of fiscal policy? In this section, we develop a theoretical framework to show that the answer is yes. Our model builds on the two-asset heterogeneous-agent New Keynesian (HANK) model of Auclert et al. (2024) and the work of Kaplan and Violante (2014, 2022). As first proposed in Galí et al. (2007), the presence of both non-Ricardian high-MPC households and nominal rigidities enables the model to replicate the positive consumption response to government spending that we see in the data.²⁷ Our results can also be

²⁷The main alternative way of generating an increase in consumption has been to use non-separable preferences that exhibit complementarity between consumption and hours worked (Basu and Kimball, 2003;

rationalized by an analytically tractable two-agent New Keynesian (TANK) model (Galí et al., 2007; Bilbiie, 2008; Angeletos et al., 2024). On the other hand, our results are inconsistent with standard representative-agent New Keynesian models, since these cannot match the increase in consumption without a counterfactual decrease in real interest rates.²⁸

Quantitative Model

Environment Time is discrete and infinite, indexed by $t = 0, 1, 2, \dots$. The economy is populated by households, labor unions, production firms, banks, a monetary authority, and a fiscal authority. All agents have perfect foresight over aggregate variables.

Households The economy is populated by a unit continuum of ex-ante identical households i with preferences

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_{it}^{1-1/\sigma} - 1}{1 - 1/\sigma} - \frac{N_{it}^{1+1/\eta}}{1 + 1/\eta} \right) \quad (1.4)$$

where C_{it} and N_{it} are consumption and hours worked of household i at time t , σ is the intertemporal elasticity of substitution, and η is the Frisch elasticity of labor supply. Households face idiosyncratic income risk: Each hour worked provides e_{it} units of effective labor, where the labor productivity e_{it} follows a Markov process with fixed transition matrix Ω . The proportion of households in productivity state e is always $\omega(e)$, where $\omega(e)$ is the probability of e in the stationary distribution of Ω . The average labor productivity level is normalized to one, so $\sum_e \omega(e)e = 1$. Households earn nominal wage W_t per unit of effective labor and take their hours worked N_{it} as given (since hours are determined by labor union demand). Aggregate efficiency-weighted hours worked is $N_t = \int e_{it} N_{it} di$. Following Auclert et al. (2024), we assume a proportional allocation rule for labor hours ($N_{it} = N_t$) and model progressive labor taxation as in Heathcote et al. (2017). Household i 's post-tax labor income is then $Z_{it} = \frac{e_{it}^{1-\theta}}{\int e_{it}^{1-\theta}}$ Z_t , where Z_t is aggregate post-tax labor income, and θ is the progressivity of labor income taxation.

Households can hold their assets in liquid (A_{it}) or illiquid ($\tilde{A}_{it}$) accounts. Both accounts generate a net return of r_t , but holding liquid assets is subject to a flow cost $\zeta(1+r_{t-1})A_{it-1}$ each period. We model illiquidity à la Calvo as in Bayer et al. (2024): Household i can transfer funds between the liquid and illiquid accounts only when the i.i.d. Bernoulli variable Λ_{it} equals one, which happens

Bilbiie, 2011; Monacelli and Perotti, 2008; Nakamura and Steinsson, 2014). This approach has largely fallen out of favor because it implies unrealistically high marginal propensities to earn out of wealth (Auclert et al., 2023). Finally, consumption increases can also be achieved using a monetary policy rule that decreases real interest rates in response to government spending (Woodford, 2011; Christiano et al., 2011), but such a policy reaction is not consistent with our estimated interest rate response.

²⁸Standard neoclassical models predict crowding out of private consumption and fiscal multipliers substantially below one (Woodford, 2011, e.g.), hence we can immediately conclude that they are unable to match our empirical evidence. See also Aiyagari et al. (1992) and Baxter and King (1993) who study the neoclassical wealth effect on labor supply.

with probability λ in each period. At time t , the budget constraints for household i are

$$C_{it} + A_{it} = Z_{it} + (1 - \zeta)(1 + r_{t-1})A_{i,t-1} - H_{it}\Lambda_{it} \quad (1.5)$$

$$\tilde{A}_{it} = (1 + r_{t-1})\tilde{A}_{i,t-1} + H_{it}\Lambda_{it} \quad (1.6)$$

where H_{it} denotes net transfers from the liquid to illiquid account. The household also face borrowing constraints for both accounts:

$$A_{it} \geq 0, \quad \tilde{A}_{it} \geq 0. \quad (1.7)$$

The two-account setup introduces a trade off: Holding assets in the liquid account insures against idiosyncratic income shocks, but it also entails a per-period return penalty $\zeta(1 + r_{t-1})$. To avoid penalties, households can keep a portion of their wealth in the illiquid account, where it earns a higher return, but this also leaves them more exposed to negative income shocks that can push them up against their borrowing constraints. The two-account structure also has the advantage that it enables us to match both high MPCs and realistic income-to-asset ratios Kaplan and Violante (2022). We provide first-order conditions and further details in Appendix A.3.

Labor Unions and Nominal Wage Rigidity We assume sticky nominal wages and demand-determined labor as in Erceg et al. (2000) and Auclert et al. (2024). Sticky wages have been shown to be particularly important in the context of fiscal multipliers because they help avoid countercyclical profits and unrealistically large income effects on labor supply (Auclert et al., 2023; Broer et al., 2023). We briefly present the setup here and provide a detailed derivation of the wage Phillips curve in Appendix A.3. Each household i belongs to a labor union k . There is a continuum of differentiated unions k , and the members of each union are representative of the entire population. Unions set the nominal wage to maximize the utility of its members subject to quadratic adjustment costs as in Rotemberg (1982). A competitive labor packer combines the differentiated labor services of the unions using a standard CES aggregator. The labor packer then sells the aggregated labor services to firms. We extend the framework to allow for partial indexation of nominal wages to wage inflation. Optimization by the labor unions at a symmetric equilibrium leads to an aggregate non-linear wage Phillips Curve with inertia:

$$\begin{aligned} \left(\frac{1 + \pi_t^w}{(1 + \pi_{t-1}^w)^{\iota^w}} - 1 \right) \frac{1 + \pi_t^w}{(1 + \pi_{t-1}^w)^{\iota^w}} &= \kappa^w [N_t v'(N_t) - \mu^w (1 - \theta) Z_t u'(C_t^*)] \\ &+ \beta \left(\frac{1 + \pi_{t+1}^w}{(1 + \pi_t^w)^{\iota^w}} - 1 \right) \frac{1 + \pi_{t+1}^w}{(1 + \pi_t^w)^{\iota^w}}. \end{aligned} \quad (1.8)$$

where $\pi_t^w = \frac{W_t}{W_{t-1}} - 1$ is the wage inflation rate, $u'(C_t^*) = \int \frac{e^{1-\theta} u'(C_{it})}{\int e^{1-\theta} di} di$ is the marginal utility of the “virtual” consumption aggregate²⁹ C_t^* , κ^w is the slope of the wage Phillips curve, μ^w is the optimal steady-state wage markup, and ι^w is the degree of wage indexation.

²⁹The virtual consumption aggregate C_t^* summarizes how the consumption distribution across households affects the aggregate wealth effect on labor supply.

Firms We adopt the common two-tier production structure, in which the final good is produced by a representative competitive firm and the intermediate goods are produced by monopolistically competitive firms. The final goods firm uses a CES production function to combine intermediate inputs. A continuum of intermediate goods firms own capital and hire labor to produce intermediate goods using a Cobb-Douglas production function. The firms must pay quadratic adjustment costs to adjust investment and nominal prices (Rotemberg, 1982). Prices are partially indexed to inflation. The intermediate goods firms maximize profits and issue equity that earn dividends, which are held by households in their liquid or illiquid accounts. The symmetric equilibrium yields an aggregate nonlinear Phillips curve with partial indexation:

$$\left(\frac{1 + \pi_t}{(1 + \pi_{t-1})^{\iota^p}} - 1 \right) \frac{1 + \pi_t}{(1 + \pi_{t-1})^{\iota^p}} = \kappa^p Y_t (\mu_t - \mu) + \frac{1}{1 + r_t} \left(\frac{1 + \pi_{t+1}}{(1 + \pi_t)^{\iota^p}} - 1 \right) \frac{1 + \pi_{t+1}}{(1 + \pi_t)^{\iota^p}} \quad (1.9)$$

where $\pi_t = \frac{P_t}{P_{t-1}} - 1$ is the inflation rate, r_t is the real interest rate on government bonds, $\mu_t = \frac{1}{1-\alpha} \frac{W_t N_t}{P_t Y_t}$ is the real marginal cost, α is the capital share, μ is the optimal markup in steady-state, κ^p is the slope of the Phillips curve, and ι^p is the degree of nominal price indexation. At the symmetric equilibrium, two standard Q-theory equations for capital demand and investment hold:

$$Q_t = \frac{1}{1 + r_t} \left(\alpha \mu_{t+1} \frac{Y_{t+1}}{K_t} + (1 - \delta) Q_{t+1} \right) \quad (1.10)$$

$$(g_{Xt} - 1) = \delta \xi (Q_t - 1) + \frac{g_{Xt+1}}{1 + r_t} (g_{Xt+1} - 1) \quad (1.11)$$

where X_t is investment in period t , $g_{Xt} = \frac{X_t}{X_{t-1}}$ is the growth rate of investment, Q_t is Tobin's Q , δ is the depreciation rate of capital, and ξ is the sensitivity of investment growth to Q . We derive the equations for the nonlinear Phillips curve and the capital dynamics in Appendix A.3.

Financial Intermediaries There is a representative competitive bank that holds real and nominal government bonds. The nominal bonds ("reserves") are in zero net supply and can only be traded by banks and the central bank. The banks maximize profits, which implies a no-arbitrage condition for nominal and real bonds that in turn implies that the Fisher equation holds at all time:

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}}. \quad (1.12)$$

The only role of the banks in the model is to enforce this Fisher equation.

Monetary and Fiscal Policy The central bank implements monetary policy by directly setting the nominal interest rate i_t on the nominal bonds held by banks. The monetary authority follows a forward-looking Taylor rule with interest-rate smoothing:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)(r + \phi^\pi \pi_{t+1}) \quad (1.13)$$

where i_t is nominal interest rate at time t , r is the steady-state real interest rate at steady-state, ϕ^π is the hawkishness of the central bank, and ρ_i is the degree of interest-rate smoothing.³⁰ The fiscal authority must respect the real government budget constraint

$$G_t + (1 + r_{t-1})B_{t-1} = T_t + B_t \quad (1.14)$$

³⁰We use the winding number algorithm of Auclert et al. (2021) to guarantee determinacy.

where G_t is the level of government purchases at time t , T_t is the aggregate tax revenue, and B_t is the debt issued at time t . The fiscal authority follows the fiscal rule (Leeper, 1991):

$$dT_t = \phi^g dG_t + \phi^b dB_t \quad (1.15)$$

where $dX_t = X_t - \bar{X}_t$ is the deviation of a variable X_t from its steady-state value $\bar{X}$. The coefficients $\phi^y > 0$ and $\phi^b > 0$ are the responsiveness of taxes to changes in government spending and public debt.

Equilibrium Given initial values of the nominal wage W_{-1} , price level P_{-1} , government debt B_{-1} , capital K_{-1} , and investment X_{-1} , and an initial distribution of households over their state variables such the economy is initially at a steady state, and given an exogenous path of government spending $\{G_t\}$, a perfect foresight general equilibrium is a path of prices $\{P_t, W_t, r_t, Q_t, i_t, \pi_t, \pi_t^w\}$ and aggregate quantities $\{Y_t, K_t, N_t, C_t, X_t, T_t, B_t\}$ such that households optimize, labor unions optimize, firms optimize, banks optimize, monetary and fiscal policy satisfy their rules, and all markets clear. Market clearing in the goods market implies that

$$C_t + X_t + G_t + \Psi_t^p + \Xi_t + \zeta(1 + r_{t-1})A_{t-1} = Y_t \quad (1.16)$$

where $C_t = \int C_{it} di$ is aggregate consumption, Ψ_t^p is the aggregate price adjustment costs, Ξ_t is the aggregate real adjustment costs, and $\zeta(1 + r_{t-1})A_{t-1} = \zeta(1 + r_{t-1}) \int A_{it-1} di$ are the aggregate costs of holding liquid assets. Appendix A.3 presents a full list of equilibrium conditions and steady state values.

Calibration We use a three-step calibration procedure. First, we externally calibrate standard parameters such as the intertemporal elasticity of substitution, the Frisch elasticity and the slope of the Phillips curve. We convert price adjustment frequencies to slopes using the same method as Auclert et al. (2024). Namely, the slope equals $\kappa = \chi \frac{r+f}{1+r} \frac{f}{1-f}$, where $\chi = 1/6$ is a real rigidity parameter taken from Auclert et al. (2024), r is the real interest rate, and f is the frequency of price adjustment. Second, we calibrate the household block parameters β (discount factor) and λ (the probability that a household is allowed to adjust their illiquid account) to match a quarterly MPC of 0.2 on impact and a ratio of aggregate assets to post-tax income (A/Z) of 22.56. This quarterly MPC target is consistent with the estimates from Johnson et al. (2006), Parker et al. (2013), and Orchard et al. (2025) among others. Third, we internally calibrate the investment growth elasticity (ξ), the Taylor rule (ϕ_π, ρ_i), and the fiscal rule (ϕ_g, ϕ_b) to match our empirical impulse response functions. Our methodology for impulse response matching closely follows Auclert et al. (2020).³¹

For variable k at horizon h , let $\hat{J}_{k,h}$ and $\hat{\Sigma}_{k,h}$ be the estimated impulse responses and standard errors from our empirical section, and $J_{k,h}(\Phi)$ be the model generated impulse response to the estimated government spending shock $\{\hat{J}_{G,h}\}_{h=0}^H$ for a set of parameters $\Phi = \{\xi, \phi_\pi, \rho_i, \phi_g, \phi_b\}$. The sum of squared errors between the estimated and model-implied impulse response functions, scaled

³¹Our estimate of the interest-rate smoothing parameter ρ_i is smaller than typically found in papers that use monetary policy shocks to estimate the monetary policy reaction function (e.g., Coibion and Gorodnichenko, 2012).

Table 1.2: Model Calibration

Parameter	Description	Value	Source
<i>Internally calibrated parameters:</i>			
β	Discount factor	0.97	Match iMPC
λ	Adjustment probability	0.4	Match income/assets
ξ	Investment growth elasticity	1.69	Match empirical IRFs
(ϕ_π, ρ_i)	Taylor rule	(1.20, 0.16)	Match empirical IRFs
(ϕ_g, ϕ_b)	Fiscal rule	(0.51, 0.00)	Match empirical IRFs
<i>Externally calibrated parameters:</i>			
r	Real interest rate	0.02	Standard
σ	Intertemporal elasticity of substitution	1	Standard
η	Frisch elasticity of labor supply	0.5	Standard
$\zeta(1+r)$	Liquid-illiquid spread	0.02	Auclert et al. (2024)
ρ^e	Persistence of income process $\log e$	0.976	Floden and Lindé (2001)
σ^e	Std. dev. of income process $\log e$	0.92	Auclert and Rognlie (2018)
α	Capital share	0.294	Standard
δ	Depreciation rate of capital	0.02	Standard
θ	Progressivity of income tax	0.181	Heathcote et al. (2017)
μ	Steady-state price markup	1	Standard
μ^w	Steady-state wage markup	1	Standard
κ^p	Slope of Phillips curve	0.015	Nakamura and Steinsson (2008)
κ^w	Slope of wage Phillips curve	0.002	Grigsby et al. (2021)
ι^p	Degree of price indexation	0.26	Benigno and Eggertsson (2024)
ι^w	Degree of wage indexation	0.55	Gali (2011)
Z/Y	Ratio of income to GDP	0.5	U.S. Labor Share
A/Z	Ratio of assets to income	22.56	Kaplan and Violante (2022)
K/Y	Ratio of capital to GDP	9.8	Standard
B/Y	Ratio of public debt to GDP	2.8	Standard
G/Y	Ratio of public spending to GDP	0.2	Standard


Notes: This table describes our baseline model's calibration. The first column gives the mathematical representation of the calibrated parameter, the second column describes the parameter, the third column reports the value of the calibrated parameter, and the final column reports the calibration method. The top panel lists the internally calibrated parameters, while the bottom panel lists the externally calibrated parameters.

by the estimated standard errors, is

$$SSE(\Phi) = \sum_k \sum_{h=0}^H \frac{(\hat{J}_{k,h} - J_{k,h}(\Phi))^2}{\hat{\Sigma}_{k,h}}. \quad (1.17)$$

The estimated parameters $\hat{\Phi}$ minimize the sum of square errors, i.e., $\hat{\Phi} = \arg \min_{\Phi} SSE(\Phi)$ and the minimized sum of squared errors is $\widehat{SSE} = SSE(\hat{\Phi})$. We estimate the parameters to match the

Figure 1.9: Estimated and Model-Implied Response to a Defense News Shock



Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation rates, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline (calibrated to table 1.2) model's response to the government spending shock in the first panel.


impulse responses of output, consumption, investment, taxes, and the nominal interest rate up to the 24th quarter. All variables are normalized to a government spending shock that peaks at one percent of output. We solve the model numerically using the sequence-space Jacobian method of Auclert et al. (2021). The externally and internally calibrated parameters can be found in table 1.2.

Effects of a Government Spending News Shock

To compare our empirical analysis with our quantitative model, we directly feed in our estimated government spending shock $\{\hat{J}_{G,h}\}_{h=0}^H$ as a government spending shock to our model. We plot the model responses of key macroeconomic variables against our empirical findings. The first panel in Figure 1.9 plots the estimated response of defense spending to a defense news shock. Dark blue lines are the estimated impulse responses to a defense news shock, and the dark and light shaded areas are the 68 and 90% confidence intervals. The red dashed line is the impulse response of our quantitative model to the government spending shock in the first panel.

The quantitative model closely matches the estimated impulses responses of output, consumption, and investment; the model generally matches the estimated responses of taxes, nominal interest rates. The inflation rate and real wage, which are targeted, broadly match the estimated responses.

Figure 1.10: Estimated and Model-Implied Fiscal Multiplier




Notes: This figure plots the estimated and model implied cumulative fiscal multiplier. The solid blue lines are our estimates of $\mathcal{M}_h$ in Equation (1.3) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline (calibrated to table 1.2) model's response to the government spending shock in the first panel.

The model-implied responses of the NIPA variables mimic the hump-shaped empirical responses, and matches their values on impact and at their respective peaks. Although the model's output response lies above the estimated response, Figure 1.10 makes clear that the estimated fiscal multiplier closely matches the estimated values across the horizon.³²

Role of Monetary Policy Our estimates of the Taylor rule ($\phi_\pi = 1.2$), from the IRF matching exercise, suggest that the monetary policy response to defense spending shocks is relatively loose. This loose monetary policy is essential to explaining the responses we find in the data and model. To show this, we consider a more hawkish monetary policy rule ($\phi_\pi^{Hawk} = 1.5$), leaving the rest of the calibration fixed at the baseline. This policy rule is more reactive to future inflation than the baseline and does not smooth interest-rates. Calibrated to this more hawkish monetary policy, we again look at the impulse responses to our estimated government expenditure shock. Figure 1.11 plots these responses. More hawkish monetary policy mutes the inflation response by amplifying the nominal interest rate response beyond what we observe in the data. This additional rise in the real interest rate mutes the consumption and investment responses leading to a fall in output

³²Although households and firms smooth consumption and investment, the output response is not smooth. Output partially inherits the uneven government expenditure response via the goods market clearing condition, i.e., $Y_t = G_t + C_t + X_t + \Psi_t + \Xi_t + \zeta(1 + r_{t-1})A_{t-1}$.

Figure 1.11: Role of Monetary Policy




Notes: This figure plots the estimated and model implied responses out output, consumption, and nominal interest rates to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the hawkish monetary policy ($\phi_\pi = 1.5$) model's response to the government spending shock in the first panel.

on impact. Hazell and Hobler (2025) and Gagliardone and Gertler (2023) similarly find that loose monetary policy is key to explaining the inflationary effects of fiscal and oil shocks. Figure A.14 in Appendix A.4 presents the responses of all other variables of interest.

Role of Fiscal Policy The Leeper rule we estimate ($\phi_g = 0.51, \phi_b = 0.00$) indicates that a significant proportion of defense expenditure is debt-financed. Debt-financing government expenditure is essential to matching the consumption response we observe. To illustrate this importance we consider an alternative balanced budget rule ($\phi_g^{BB} = 1, \phi_b^{BB} = .00$) which entirely tax finances government expenditure. We leave the rest of the calibration at the baseline, then study the model response to our estimated government expenditure shock. Figure 1.12 plots theses responses. The larger tax response causes consumption to fall. Auclert et al. (2024) also find that debt financing government expenditure is key to producing larger fiscal multipliers. Figure A.15 in Appendix A.4 presents the response of all other variables of interest.

Alternative Models Our baseline model is a two-asset HANK model with large iMPCs, nominal rigidities, and significant real rigidities. To understand the impact of these three elements we

Figure 1.12: Role of Fiscal Policy



Notes: This figure plots the estimated and model implied responses out output, consumption, and taxes to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the tax-financed fiscal policy (with $\phi_g = 1$) model's response to the government spending shock in the first panel.

consider a number of alternative models that relaxes each. Ultimately, we find that all three are necessary to match the impulses responses we estimate.

To evaluate the importance of large iMPCs we repeat our three-step estimation procedure (see section 1.4) with a representative agent New Keynesian (RANK) model and a two-agent bonds-in-the-utility (TABU) New Keynesian model. The former features low iMPCs while the latter can match the iMPCs in our baseline model. The RANK model resembles the canonical DSGE model in Smets and Wouters (2007) while the TABU model resembles the tractable HANK model in Angeletos et al. (2024). We present these models in full and plot their impulses responses (Figures A.16 and A.17) in Appendix A.4. We show that the HANK models is unable to match the strong consumption response we observe, while the TABU model performs similarly to the baseline model in matching the aggregate responses.

Nominal rigidities are essential to matching the impulses responses we observe. Both sticky prices and sticky wages play an important role. To elucidate their impact we consider two alternative models with flexible prices and wages. The flexible price and wage models separately feature steeper

Phillips curves and wage Phillips curves.³³ We repeat our three-state calibration procedure (see section 1.4) and plot the impulse responses of the flexible price model (Figure A.18) and flexible wage model (Figure A.19) in Appendix A.4. The flexible price model exhibits an excessively high inflation and consumption response, as well as a negative real wage response. The flexible wage model features excessive real wage, inflation, investment, consumption, and output responses.

Additional real rigidities, beyond those in the canonical NK model, are necessary to match the estimated impulse responses to monetary policy shocks (e.g. Smets and Wouters (2007)). We also find these DSGE-style rigidities are necessary to match the estimated responses to fiscal policy shocks. Investment adjustment costs are crucial to matching the investment response, price and wage indexing are key to matching the initially muted response of consumption and output. To see this we piecewise remove these rigidities and repeat our three-state calibration procedure (see section 1.4). We plot the impulse responses of capital adjustment model (Figure A.20) and non-indexation model (Figure A.21) in Appendix A.4. These alternative models lead to excessive responses of consumption and output. In Appendix A.4 we conduct a number of other exercises, including a recalibration with a larger Frisch elasticity (Figure A.23), a recalibration without interest-rate smoothing (Figure A.22), and plotting the model's response to an AR(1) government expenditure shock (Figure A.24).

1.5 Conclusion

In this paper, we use a high-frequency identification approach to estimate the effects of government spending on macroeconomic outcomes. We create a new instrument for defense spending by measuring surprises in a stock index of major U.S. defense contractors within a narrow time window around legislative events that communicate news about defense spending. The defense spending surprises that we identify have large and persistent effects on actual future defense spending and are passed through one-for-one to total government spending. Our findings indicate that government spending stimulates private economic activity and provides empirical evidence that aggregate fiscal multipliers can exceed one. We show that our empirical results are consistent with a standard heterogeneous-agent New Keynesian model with non-Ricardian households and realistic marginal propensities to consume. Representative agent models without these features struggle to generate an increase in consumption without a counterfactual decrease in real interest rates.

³³In these alternative models prices and wages are not perfectly flexible but rather feature significantly lower nominal adjustment costs. Specifically, the slopes of the Phillips curves are ten times the baseline calibration, i.e., $\kappa^p = .15$ and $\kappa^w = .02$, in the flexible price and wage calibrations.

Chapter 2

Monopsony with Recruiting

with Justin Bloesch and Birthe Larsen

We thank Anastasia Burya, Gabriel Chodorow-Reich, Ryan Chahrour, Kyle Herkenhoff, Lawrence Katz, Philipp Kircher, Patrick Kline, Jeremy Lise, Alan Manning, Simon Mongey, Christian Moser, Suresh Naidu, Kris Nimark, Ezra Oberfield, Frank Pinter, Tommaso Porzio, Seth Sanders, Anna Stansbury, Evan Soltas, Jacob Weber, and seminar participants at the Copenhagen Business School, the University of California, Berkeley, the IAB workshop Imperfect Competition in the Labor Market, the NBER SI, the EALE conference, Bergen, the EEA conference, Rotterdam, the local-MPC workshop, Copenhagen, the BSE forum on the Macroeconomics of the Labor Market, Barcelona, the SED conference, Barcelona, and Columbia Junior Micro/Macro Labor conference and for helpful comments. We thank Henry Tan, Michael Yin, and Nathan Mascia for excellent research assistance.

2.1 Introduction

There is a growing consensus that firms have significant power to set wages: firms that choose lower wages do not lose all of their employees (Card, 2022). In many standard models of labor market monopsony, firm wage setting power implies a finite labor supply elasticity, allowing researchers to infer the markdowns of wages below firms' marginal product of labor, and by extension, the profits that firms earn by exploiting their wage setting power. Yet, questions remain in estimating the labor supply elasticity and its implications for the distribution of income. First, common estimates of labor supply elasticities between 2 and 6 imply that firms must pay significantly higher wages to become large. In tension with this evidence, large firms pay only slightly higher wages in cross-sectional data, and other studies find labor supply elasticities that are at least an order of magnitude higher.¹ Additionally, numerous recent studies show that *growing* firms pay a wage premium.² Second, there is a puzzle regarding profits: commonly used low labor supply elasticities imply implausibly large profit shares of income in the aggregate.³ Further, many monopsony models abstract from recruiting activity and assume that the only way firms can expand is by offering higher wages. In this article, we therefore ask: *if firms can spend on recruiting, what is the long-run labor supply elasticity they face, and how large are the profits that result from their wage-setting power?*

To answer these questions, we develop an on-the-job search model in which firms attract workers on two margins: higher wages and recruiting expenditure. Workers have idiosyncratic preferences over firms, and the presence of both search frictions and these preferences gives firms wage-setting power: firms that choose lower wages do not immediately lose all their workers. Firms can recruit new workers by spending on recruiting activity, and the number of matches a firm generates in a period increases in both its recruiting expenditure and its stock of incumbent workers. *In the short run*, the firm faces diminishing returns from its recruiting expenditure, and so in response to a positive demand shock for its output, the firm offers higher wages to increase the number of hires from a limited pool of workers matched to the firm. Consequently, firms pay high wages when they are growing. However, *in the long run*, a greater number of incumbents in the firm eases recruiting, which reduces the incentive to pay high wages to generate more hires. If the firm's recruiting function has constant returns to scale, then the optimal wage in steady state is the same regardless of the firm's size, resulting in an infinite long-run labor supply elasticity. We then show in closed form that in steady state, the *recruiting cost-adjusted* wage markdown is tightly related to this long-run labor supply elasticity: if firms' long-run labor supply curve is elastic, then the recruiting cost-adjusted wage markdown is equal to one, and the gap between wages and the marginal product

¹Sokolova and Sorensen (2021) provide a meta-analysis of two broad approaches to estimating labor supply elasticities that provide dramatically different estimates. Matsudaira (2014) finds that nursing homes subject to employment level mandates did not increase wages. Bloom et al. (2018) show that the firm-size wage premium had declined dramatically in the United States. We similarly show in Denmark that the wage premium at large firms is trivial in Appendix B.4.

²See Schmieder (2023), Tanaka et al. (2023), and Carrillo-Tudela et al. (2023). Engbom et al. (2023) find that conditional on capital and productivity, large firms do not pay higher wages.

³Quoting Manning (2021), pg 11: "The low estimated wage elasticity of the labor supply curve implies that employers have a lot of monopsony power: If this power is exercised, it is not clear how it can be reconciled with observed levels of the profit share."

is entirely consumed by recruiting expenditure. Thus with constant returns to scale in recruiting, monopsony rents are dissipated by the costs incurred to acquire workers.⁴

Next, we estimate the parameters of the recruiting cost function and the long-run inverse labor supply elasticity using the effect of idiosyncratic export product demand shocks on the paths of firm employment levels, firm hiring rates, and wages, as implied by our model. In particular, we study the wage growth for job switchers around firm-level export demand shocks using Danish administrative matched employer-employee and trade data (Hummels et al., 2014). We focus on the wage growth of switchers, rather than the wage response of stayers: while firms may be incentivized to share rents with incumbent workers in response to shocks to improve retention (Kline et al., 2019; Garin and Silvério, 2024; Carvalho et al., 2024), the response of stayer wages in our model is not informative about firms' long-run labor supply elasticity. We find that in response to export demand shocks, a firm's employment level, hiring rate, and wage growth of the firm's newly hired job switchers increase on impact. This is consistent with diminishing returns to recruiting effort, as firms resort to offering high wages to increase hiring among a limited pool of potential workers. We also show that the effect of the demand shock on a firm's employment level is persistent: firms are larger for years following the export demand shock. Crucially, however, workers who switch into the shocked firm in the years following the trade shock, when the firm is larger but is no longer rapidly hiring, do not see any additional wage growth relative to workers who were hired prior to the shock. This is consistent with a labor supply curve that is elastic in the long run: the firm has increased in size but pays similar wages as when the firm was smaller. Together, these results are consistent with diminishing returns to recruiting expenditure in the short run, but an elastic labor supply curve in the long run.

To evaluate the extent of monopsony profits in a setting with labor market concentration, we then extend the model to consider non-atomistic firms and strategic interactions in the presence of a recruiting margin, using our estimates of the recruiting cost parameters. Firms compete according to an extended version of Bertrand competition, where firms solve their problems taking their competitors' wage and vacancy choices as given. We show that increasing the labor market concentration so that the labor market Herfindahl-Hirschman Index (HHI) increases from 0 to an empirically realistic level of 0.1 (Berger et al., 2022) lowers wages relative to the marginal product by a modest 2.5 percent economy-wide. Because markdowns are not very wide, the output losses from markdown dispersions are similarly modest: Reallocating labor to equalize marginal products would increase output by only 0.6 percent.

Lastly, we demonstrate the profit puzzle: in a model with standard labor supply elasticities but no recruiting margin, and otherwise standard values for price markups and the elasticity of output with respect to capital, the profit share of income is 10–15 percentage points too high, and the labor share of income is 10–15 percentage points too low, relative to national accounts data in rich countries. We show that if firms have a recruiting function that exhibits constant returns to scale and face an elastic long-run labor supply curve, then the puzzle can be reconciled: Our model matches aggregate profit and labor shares while maintaining realistic price markups over marginal cost. Due to the modest effects of labor market concentration on the labor share, our model can still match the labor share in a setting with both recruiting costs and labor market concentration.

⁴This result that a constant returns to scale recruiting function makes the firm's labor supply curve elastic is a familiar result from Manning (2003), Kuhn (2004), Manning (2006).

Our finding of elastic firm-specific labor supply curves and small profits from monopsony power contrasts with three separate strands of the monopsony literature that use different methods to estimate labor supply elasticities, all finding elasticities in the range of 2–6. Simple models dynamic monopsony based on Manning (2003) infer labor supply elasticities from the elasticities of recruiting and separation rates to wages but assume no recruiting margin. Our model nests these simple models and allows for an elastic labor supply curve due to firms’ ability to spend on recruiting, taking as given this literature’s empirical estimates of recruiting and separation elasticities. Relative to papers that estimate labor supply curves using the effect of demand shocks on firm size and wages (such as Lamadon et al. (2022), Seegmiller (2021) and Chan et al. (2023)), we find similar employment and wage responses of stayers to shocks, but we find more elastic labor supply curves when we use the wages of new hires, as suggested by our model. Finally, relative to models with Cournot competition in the labor market and no recruiting margin (such as Berger et al. (2022)), our model with Bertrand competition and recruiting results in higher labor supply elasticities, more compressed markdowns, and less misallocation from markdown dispersion for a given level of labor market concentration. On the whole, our results still point to a labor market with a substantial rents: Our evidence is consistent with wages that are around 8% below marginal product and a marginal hiring cost of a worker that corresponds to 4–6 months of wages. However, our results suggest that competition among firms, via recruiting expenditure, limits the extent to which these labor-market rents translate into aggregate profits.

One theoretical contribution we make is that our environment has random on-the-job search and firm wage setting, but there is a point mass of wages in equilibrium when firms and workers are ex-ante homogeneous, in contrast to Burdett and Mortensen (1998). Like Albrecht et al. (2018), we achieve this by workers having horizontally differentiated preferences over workplaces, but unlike much of the existing monopsony literature, we make these differentiated preferences time-varying. Our paper shares with Heise and Porzio (2022) the combination of firm wage setting, on-the-job search, and idiosyncratic non-wage preferences. Our paper differs in its focus on the recruiting function, relating to existing objects in the monopsony literature, and estimating parameters through out-of-steady-state response to shocks. Berger et al. (2024) combine idiosyncratic preferences with on-the-job search and bargaining. Our paper differs by assuming wage posting and provides closed-form expression for markdowns inclusive of recruiting costs and cleanly nests dynamic monopsony models Manning (2003). Our paper differs from Gottfries and Jarosch (2023) by including idiosyncratic worker preferences and a richer recruiting cost function. Our paper differs from all three of the previous papers by providing novel causal empirical evidence and estimating model parameters using that causal evidence, with a specific focus on tractably nesting dynamic monopsony models as in Manning (2003) and matching aggregate factor shares.

Our findings have implications for the labor search literature. In many structural labor search models with multi-worker firms and linear production technologies, firm size is pinned down by a convex recruiting cost function (e.g. Acemoglu and Hawkins (2014)). These models exhibit the same puzzle of wide markdowns (inclusive of recruiting costs) because they share the feature with monopsony models that the marginal cost of labor is increasing with firm size. Instead, we present evidence for a recruiting cost function that is homogeneous of degree one where larger firm size improves matching rates (Coles and Mortensen, 2016; Burdett and Vishwanath, 1988; Gouin-Bonenfant, 2022; Kuhn, 2004). The result is that firms’ steady-state marginal cost of labor does not increase with firm size.

The paper is organized as follows. Section 2.2 lays out the model of dynamic monopsony with atomistic firms and decomposes marginal product into wages, recruiting costs, and labor market profits. Log-linearizing the firm's problem, we show that new hire wages are a function of current and future hiring rates and employment levels. Section 2.3 presents the empirical evidence on the wage growth of switchers around firm export demand shocks in Denmark and estimates the key parameters in the recruiting cost function. Section 2.4 extends the model to include labor market concentration and addresses the question of aggregate profit and labor shares. Section 2.5 concludes.

2.2 Monopsony with Recruiting: Atomistic Firms

In this section, we present our model of monopsony with recruiting, in which atomistic firms set wages and choose how much to spend on recruiting. We define the equilibrium and characterize firms' steady-state wage and employment policies. We decompose the marginal product of labor into the shares spent on wages and recruiting and the share retained by firms as profits. We show that this decomposition depends on (i) the sensitivity of recruiting and separation to firm wage policies and (ii) the parameters of firms' recruiting cost function. In particular, if firms have a constant returns to scale recruiting function, they face a perfectly elastic labor supply curve in the long run, and the marginal product fully exhausted by wages and recruiting costs. We show that the parameters of the firm's recruiting cost function—and, by extension, the firm's long-run labor supply elasticity—can be estimated from the paths of firms' new-hire wages, hiring rates, and employment levels in response to idiosyncratic product demand shocks. Section 2.3 describes the empirical implementation.

Setup

Time, Agents, and Local Labor Markets Time is discrete, and the model will later be calibrated to monthly frequency. There is a unit mass of both workers i and firms j . Firms are atomistic in the sense that they have zero mass and the choices of any individual firm have no effect on aggregates, and firms employ a measure of workers. Workers can be either employed or unemployed and supply one unit of labor when employed. There is a large number M of local labor markets indexed by m . Where possible, we will disregard the m subscripts and focus on one local labor market at a time since there is no labor mobility across markets.

Search Environment and Contract Space Firms and workers match in a frictional labor market, and matching is random. Workers search with probability one when unemployed and with probability $\lambda_{EE} \in [0, 1]$ when employed. Workers receive an amount $b > 0$ from home production when unemployed, and workers can quit into unemployment if they are not matched with an outside job. Firms post wages and are subject to two constraints. First, firms must pay all workers within a cohort the same wage. That is, there is a firm-by-tenure pay equality constraint. As a consequence, firms cannot make individual counter-offers to workers. Second, firms are required to pay at least the posted wage to a given cohort of workers as long as those workers are in the firm.⁵ Thus, for a

⁵This can be rationalized by pay transparency laws that require employers to post wages that are similar to what will be paid in the future.

cohort of workers hired by firm j in period s , firms face the constraint $w_{j,s,t} \geq w_{j,s,s} \forall t \geq s$, where $w_{j,s,s}$ is the minimum wage that firm j committed to in period s , and $w_{j,s,t}$ is the wage paid in period t to workers hired in period s . The total number of workers in the firm, $N_{j,t}$, is the sum of all workers from each past cohort still remaining in the firm: $N_{j,t} = \sum_{s=-\infty}^t N_{j,s,t}$, where $N_{j,s,t}$ is the number of workers hired in period s that still remain in the firm by period t .

Let U_t be the unemployment rate in period t , so the share of employed workers is $1 - U_t$. Firm j 's effective recruiting activity in period t is $V_{j,t}$, and aggregate recruiting activity is given by $\bar{V}_t = \int_j V_{j,t} dj$. Workers separate exogenously into unemployment at rate s_0 . Labor market tightness is defined as the ratio of vacancies to searchers, $\theta_t \equiv \bar{V}_t / \mathcal{S}_t$, where $\mathcal{S}_t \equiv (1 - U_{t-1})(1 - s_0)\lambda_{EE} + U_{t-1}$. The share of searchers who are unemployed is $\Phi_{U,t} = U_{t-1} / \mathcal{S}_t$, and the share of searchers who are employed is the complement $\Phi_{E,t} = 1 - \Phi_{U,t}$. Total matches $\mathcal{M}(\bar{V}_t, \mathcal{S}_t, \cdot)$ is a constant returns to scale function of aggregate recruiting activity and searchers. Searching workers match with firms with probability $f(\theta_t) \equiv \frac{\mathcal{M}_t}{\mathcal{S}_t}$, and a firm that exerts $V_{j,t}$ units of recruiting activity matches with $V_{j,t}g(\theta_t)$ workers, where $g(\theta_t) = \frac{\mathcal{M}_t}{\bar{V}_t}$. We restrict to functions with $f(\theta_t) \in [0, 1]$.

The timing within a period is as follows. The economy inherits the unemployment rate U_{t-1} from the previous period, and each firm j inherits a number of workers from each past cohort $N_{j,s,t-1}$, where each cohort has its own cohort-specific minimum wage $w_{j,s,s}$. At the beginning of the period, exogenous separations occur at rate s_0 , and those workers cannot search this period. Then firms choose their wage schedule $\{w_{j,s,t}\}_{s \leq t}$ and recruiting activity $V_{j,t}$, and matching occurs, which determines the employment levels $N_{j,t}$ and unemployment rate U_t for period t . Workers who are not matched may quit into unemployment. At the end of the period, firms produce and pay wages.

Workers

Workers get indirect flow utility from wage income w and discount at rate β_w . A worker hired in period s by firm j has value $\mathcal{V}_{j,s,t}$ in period t . When choosing between two jobs at firms j and k , workers draw idiosyncratic preference shocks for each job $\iota_{i,j,t}$ from a Type-1 extreme-value distribution with scale parameter γ . The worker then chooses the job that offers the highest utility, solving

$$\max\{\mathcal{V}_{j,s,t+1} + \iota_{i,j,t+1}, \mathcal{V}_{k,t+1,t+1} + \iota_{i,k,t+1}\}.$$

If workers are not exogenously separated and are not choosing between jobs, the worker may voluntarily quit into unemployment. In this case, workers do not receive idiosyncratic preference draws. The worker's value $\mathcal{V}_{j,s,t}$ of working at firm j solves

$$\begin{aligned} \mathcal{V}_{j,s,t} = & \frac{\eta}{\eta - 1} \left(w_{j,s,t}^{\frac{\eta-1}{\eta}} - 1 \right) + \beta_w s_0 \mathcal{V}_U \\ & + \beta_w (1 - s_0) \mathbb{E}_t \left[\lambda_{EE} f(\theta) \int v_k \gamma^{-1} \log(\exp(\gamma \mathcal{V}_{j,s,t+1}) + \exp(\gamma \mathcal{V}_{k,t+1,t+1})) d\mathcal{V}_k \right. \\ & \left. + (1 - \lambda_{EE} f(\theta)) \max\{\mathcal{V}_{j,s,t+1}, \mathcal{V}_U\} \right], \end{aligned} \quad (2.1)$$

where $\mathcal{V}_U$ is the value of unemployment, and v_k is the density of values associated with the posted wage distribution, with distribution $\Upsilon(\mathcal{V})$. Because we solve for a stationary equilibrium, the variables $\mathcal{V}_U$, θ , and v_k do not have time subscripts, and $\int v_k dk = 1$.

The timing assumptions embedded in equation (2.1) are as follows: workers produce and receive wage payments at the end of the current period. At the beginning of the next period, workers are exogenously separated with probability s_0 . With probability $1 - s_0$, workers retain the option to keep their job. Next, workers are allowed to search on the job with probability λ_{EE} . Conditional on searching, workers match with a hiring firm with probability $f(\theta)$. The density of expected valuations at other firms $\mathcal{V}_k$ is v_k . Conditional on having matched, workers choose which job offers the highest utility, taking into account the expected future income and the idiosyncratic preference draw. Conditional on having matched with a firm where the expected value is $\mathcal{V}_k$, the probability that the worker who started at firm j in period s accepts the outside offer in period t is

$$\mathbb{P}[\text{worker at firm } j \text{ switches to firm } k] = \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})}.$$

The higher the value of γ , the more sensitive workers are to differences in values across firms when choosing jobs. Finally, workers who are unemployed have the following value function⁶

$$\mathcal{V}_U = \frac{\eta}{\eta - 1} \left(b^{\frac{\eta-1}{\eta}} - 1 \right) + \beta_w \mathbb{E}_t \left[f(\theta) \int_{\mathcal{V}_k \geq \mathcal{V}_U} v_k \mathcal{V}_{k,t+1,t+1} d\mathcal{V}_k + (1 - f(\theta)) \mathcal{V}_U \right].$$

Firms

Production Firm j in the local labor market m produces a tradable intermediate good $Y_{j,m,t}$. Competitive final goods bundlers produce a final good using a CES production technology: $Y_t = \left(\mathbb{A}_t^{-1} \sum_{m=1}^M \sum_{j=1}^J A_{j,m,t} Y_{j,m,t}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$ where $\mathbb{A}_t^{-1} = \frac{1}{\sum_{m=1}^M \sum_{j=1}^J A_{j,m,t}}$.⁷ We assume that local labor markets m are sufficiently large that firms do not take into account their effect on the product market.

We will focus on a single local labor market and for now drop the m subscripts. Firms produce with a Cobb–Douglas production technology with $Y_{j,t} = K_{j,t}^\alpha N_{j,t}^{1-\alpha}$ using capital $K_{j,t}$ and the measure of workers $N_{j,t}$, with $\alpha \in [0, 1)$. We normalize the price level so that firms face an isoelastic product demand function each period of the form $P_{j,t} = A_{j,t} Y_{j,t}^{-\frac{1}{\epsilon}}$, where the product demand elasticity is $-\epsilon$, and $A_{j,t} = A_j + a_{j,t}$ is a firm-specific product demand shifter that consists of a permanent component A_j and a potentially persistent shock $a_{j,t}$. $a_{j,t}$ evolves according to $a_{j,t} = G(a_{j,t-1}) + z_{j,t}$, where $G(\cdot)$ is deterministic and known to firms and workers, and $z_{j,t}$ occurs with probability 0 for each firm each period. In total, the firm's revenue is $A_{j,t} \left(K_{j,t}^\alpha N_{j,t}^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}}$. Firms can rent capital elastically at rate r^K . Firms may differ in $A_{j,t}$ and c_j , but χ , σ , and α are common parameters.

Worker Flows Firms choose wages period by period, subject to the constraint that the wage is at least the minimum wage posted for each cohort. Firms know that paying higher wages increase

⁶We assume that workers do not take preference draws when matching with jobs from unemployment. Assuming so would not change our core results.

⁷This is equivalent to nested CES where the elasticities of substitution of the inner and outer nests are both ϵ .

worker values, and higher worker values increase the rate of recruiting workers and slow the rate of workers separating. Let $R(\mathcal{V}_{j,t,t})$ be the number of hires per unit of recruiting activity $V_{j,t}$, with $R'(\mathcal{V}_{j,t,t}) > 0$. The value function has two t subscripts because workers who are hired by this firm in period t will have a cohort identifier $s = t$. Also let $S(\mathcal{V}_{j,s,t}) \in [0, 1]$ be the share of incumbent workers hired in period s and still employed by that firm at the beginning of period t who separate from the firm, with $S'(\mathcal{V}_{j,s,t}) < 0$. We assume that firms employ a continuum of workers, so there is no uncertainty in a given period about how many workers will be hired or separated conditional on the firm's recruiting activity and worker values. The recruiting rate is

$$R(\mathcal{V}_{j,t,t}) = g(\theta) \left(\Phi_E \left(\int \phi_k \frac{\exp(\gamma \mathcal{V}_{j,t,t})}{\exp(\gamma \mathcal{V}_{j,t,t}) + \exp(\gamma \mathcal{V}_{k,s,t})} dk \right) + \Phi_U \mathbb{1}\{\mathcal{V}_{j,t,t} \geq \mathcal{V}_U\} \right), \quad (2.2)$$

where ϕ_k is the density of values that workers are employed at, with $\int \phi_k dk = 1$ and distribution $\Phi(\mathcal{V})$. The separation rate is

$$S(\mathcal{V}_{j,s,t}) = s_0 + (1 - s_0) \left(\lambda_{EE} f(\theta) \int \nu_k \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})} dk + (1 - \lambda_{EE} f(\theta)) \mathbb{1}\{\mathcal{V}_U > \mathcal{V}_{j,s,t}\} \right), \quad (2.3)$$

where ν_k is the density of posted values, with $\int \nu_k dk = 1$ and distribution $\Upsilon(\mathcal{V})$, as before. Note that $R(\cdot)$ and $S(\cdot)$ do not have time subscripts because there are no aggregate shocks, and we will only solve for a stationary equilibrium.

Recruiting Cost Function Firms begin every period with $N_{j,t-1}$ incumbent workers and can choose a level of effective recruiting activity $V_{j,t}$. Let $\mathcal{C}_{j,t}$ be firm j 's expenditure on recruiting in period t . The recruiting function can be expressed in two ways:

$$\mathcal{C}_{j,t} = c_j \left(\frac{V_{j,t}}{N_{j,t-1}} \right)^\chi N_{j,t-1}^\sigma V_{j,t} \quad \Leftrightarrow \quad V_{j,t} = c_{j,0} \mathcal{C}_{j,t}^{\frac{1}{1+\chi}} N_{j,t-1}^{\frac{\chi-\sigma}{1+\chi}} \quad (2.4)$$

where σ is unrestricted, $\chi \geq 0$, and $c_{j,0} \equiv c_j^{-1/(1+\chi)}$. χ determines how much marginal recruiting costs increase with the recruiting rate $V_{j,t}/N_{j,t-1}$, while σ governs how firm size affects recruiting costs directly, given the recruiting rate. We will use “effective recruiting activity” and “vacancies” interchangeably, as $V_{j,t}$ is intended to capture all possible activities that firms can engage in to match with more workers. These activities could include instructing current employees to share with their social network that the firm is hiring, recalling former employees, putting help-wanted signs in stores, paying for online job postings, attending career fairs, etc. The assumption behind our recruiting function is that firms rationally choose the least cost per unit of effective recruiting first, and firms choose more expensive forms of recruiting when the cheaper forms of recruiting are exhausted.⁸ The second way of formulating equation (2.4) shows that effective recruiting activity

⁸Prior research distinguish between vacancies and recruiting intensity Davis et al. (2013) and more recently by Gavazza et al. (2018) and Mongey and Violante (2025). We sidestep this distinction, as our measure of “effective recruiting activity” is a reduced-form composite of all possible recruiting activities.

$V_{j,t}$ is a function of two inputs, expenditure and firm size. The recruiting function exhibits constant returns to scale if $\sigma = 0$: doubling both firm size and recruiting expenditure doubles the amount of recruiting activity and doubles the number of matches.⁹ If $\chi > 0$, then the firm faces diminishing returns to recruiting expenditure $C_{j,t}$ for a given firm size $N_{j,t-1}$.

Firm's Problem Firms maximize the present discounted value of revenue less wage, capital, and recruiting costs, with discount factor β_f . Because firms are only subject to idiosyncratic MIT shocks to the demand shifter $A_{j,t}$, we can write the firm's problem in terms of certainty equivalence. Given cohort-level employment $\{N_{j,s,-1}\}_{s \leq -1}$, and the history of posted wages $\{w_{j,s,s}\}_{s \leq -1}$, firms solve

$$\begin{aligned} & \max_{\{V_{j,t}, K_{j,t}, N_{j,t}, \{N_{j,s,t}\}_{s \leq t}, \{w_{j,s,t}\}_{s \leq t}\}_{t=0}^{\infty}} \\ & \sum_{t=0}^{\infty} \beta_f^t \left(A_{j,t} \left(K_{j,t}^{\alpha} N_{j,t}^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - \sum_{s=-\infty}^t w_{j,s,t} N_{j,s,t} - r^K K_{j,t} - c_j \left(\frac{V_{j,t}}{N_{j,t-1}} \right)^{\chi} N_{j,t-1}^{\sigma} V_{j,t} \right) \end{aligned} \quad (2.5)$$

subject to

$$\text{Recruiting: } N_{j,t,t} = R(\mathcal{V}_{j,t,t}) V_{j,t} \quad \forall t \quad (2.6)$$

$$\text{Retention: } N_{j,s,t} = (1 - S(\mathcal{V}_{j,s,t})) N_{j,s,t-1} \quad \forall t \text{ and } s < t \quad (2.7)$$

$$\text{Sum of cohorts: } N_{j,t} = \sum_{s=-\infty}^t N_{j,s,t} \quad \forall t \quad (2.8)$$

$$\text{Wage at least posted wage: } w_{j,s,t} \geq w_{j,s,s} \quad \forall t \text{ and } s \leq t, \quad (2.9)$$

where (2.1), (2.2), and (2.3) are the worker's value function, the recruiting rate, and the separation rate, respectively. For simplicity, we will focus on the limiting case where firms do not discount ($\beta_f \rightarrow 1$) for the remainder of the paper.¹⁰

Stationary Equilibrium

A stationary equilibrium consists of a mass of searchers $\mathcal{S}$, labor-market tightness θ , an unemployment rate U , distributions of values across employed workers $\Phi(\mathcal{V})$ and across vacancies $\Upsilon(\mathcal{V})$, an employment policy $\{N_{j,s,t}^*\}_{s \leq t}$, a wage policy $\{w_{j,s,t}^*\}_{s \leq t}$, and a recruiting policy $\{V_{j,t}^*\}$ such that (i) firms maximize profits, (ii) workers maximize utility, and (iii) worker inflows and outflows balance each period. We focus on equilibria with two desirable properties. First, in steady state, firms choose a constant wage profile for each cohort: $w_{j,s,t} = w_{j,s,s}$ for all $t \geq s$.¹¹ Second, firms of a given

⁹Modeling matches as increasing in firm size is supported by evidence that many hires are done through referrals and social networks: see Burks et al. (2015), Caldwell and Harmon (2019), and Dustmann et al. (2016).

¹⁰We show in Appendix B.6 that when this model is calibrated to standard monthly parameters, the discount factor is quantitatively unimportant for profits.

¹¹Specifically, it is optimal for firms to choose a fixed wage profile under the belief that workers expect firms will revert to the posted wage if firms deviate. See Coles (2001) and Manning (2025) for wage setting without commitment and Stevens (2004) for contracting with commitment.

type, characterized by their firm-specific parameters (A_j, c_j) , have a unique optimal steady-state wage.

Constant wages arise as an equilibrium result from two factors: (i) firms have more monopsony power over incumbents than over matched workers who have not accepted a job yet, and (ii) firms are constrained to pay at least the posted wage. Firms have greater monopsony power over incumbent workers because many incumbent workers are not searching on the job in any given period and have only unemployment as an outside option. In contrast, many potential new hires are searching on the job and have their current job as an outside option. Once workers accept a job and become less mobile, firms would find it profitable to renege and cut wages. Since the wage-posting constraint prevents this, the best remaining option for firms is to pay constant wages.¹² We verify numerically that firms have no incentive to deviate from a constant wage path in steady state under a wide range of plausible parameter values,¹³ and we prove it formally for a special case in Section 2.2.

Now that we restrict to a stationary equilibrium with constant wage–tenure profiles (so $w_{j,s,t} = w_{j,s,s}$ for all $t \geq s$), we can define a firm’s recruiting and separation elasticities, which are the objects of interest in the dynamic monopsony literature (Manning, 2003; Bassier et al., 2022) and are respectively denoted ε_{R,w_j} and ε_{S,w_j} . The recruiting elasticity is $\varepsilon_{R,w_j} \equiv \frac{R'(w_j)w_j}{R(w_j)}$, where the derivative of the recruiting function $R'(w_j) = \frac{\partial R}{\partial V_j} \frac{\partial V_j}{\partial w_j}$ is taken under the mutual expectation by both workers and firms that the wage will be constant throughout the duration of the employment relationship. Analogously, the separation elasticity is defined as $\varepsilon_{S,w_j} \equiv \frac{S'(w_j)w_j}{S(w_j)}$ with $S'(w_j) = \frac{\partial S}{\partial V_j} \frac{\partial V_j}{\partial w_j}$, where $S'(w_j)$ captures how a cohort’s separation rate responds to the posted wage. When posting a wage, the firm takes into account that paying a higher wage will accelerate the filling of a vacancy and reduce the rate of separations for that cohort of workers. Faster hiring and slower separations reduce the number of vacancies V_j the firm needs to post, decreasing its recruiting costs. The firm’s optimal steady-state wage satisfies

$$w_j = c_j(1 + \chi) \left(\frac{S(w_j)}{R(w_j)} \right)^{1+\chi} (\varepsilon_{R,w_j} - \varepsilon_{S,w_j}) N_j^\sigma. \quad (2.10)$$

Since $S'(w) < 0$ and $S(w) > 0$, the separation elasticity is negative, $\varepsilon_{S,w} < 0$. The steady-state ratio of vacancies to employment takes the form

$$\frac{V_j}{N_j} = \frac{S(w_j)}{R(w_j)}. \quad (2.11)$$

Together, equations (2.10) and (2.11) demonstrate that the firm trades off wage costs and turnover costs: a higher wage means that the firm will post fewer vacancies in steady state, decreasing recruiting costs. Conversely, a low wage implies higher steady-state turnover costs. We can define the sum of recruiting and (negative) separation elasticities as

$$\mathcal{E}(w_j) \equiv \varepsilon_{R,w_j} - \varepsilon_{S,w_j}.$$

¹²The alternative strategy is to back-load wages: post a low starting wage and raise wages later for incumbents. However, since firms cannot credibly commit to future wages above the posted minimum, workers anticipate reversion to the posted wage. Back-loading is therefore a costly way to increase retention and is not desirable for the firm.

¹³We verify that parameterizations that match the empirical unemployment rate, job-to-job transition rate, vacancy rate, and recruiting and separation elasticities generate flat wage–tenure profiles in equilibrium.

For presentation purposes, we will suppress the argument (w_j) , and simply write this sum of elasticities as $\mathcal{E}_j$. The optimal steady-state level of employment for firm j is

$$N_j = \left(\frac{\frac{\epsilon-1}{\epsilon}(1-\alpha)A_j}{w_j} \right)^\epsilon \left(\frac{w_j}{r^K} \frac{\alpha}{1-\alpha} \right)^{\alpha(\epsilon-1)} \left(\frac{\mathcal{E}_j(1+\chi)}{1+\mathcal{E}_j(1+\chi)+\sigma} \right)^{\epsilon-\alpha(\epsilon-1)}. \quad (2.12)$$

Finally, the second convenient property that emerges from this model is that, under a broad range of reasonable parameterizations, ex-ante identical firms have a unique and identical profit maximizing wage. If all firms are ex-ante identical, then the equilibrium features a point mass of wages. This stands in contrast to the famous result in Burdett and Mortensen (1998), that, with on-the-job search and ex-ante homogeneous workers and firms, the equilibrium must contain wage dispersion. As in Albrecht et al. (2018), we achieve this by workers having horizontally differentiated preferences over workplaces. This means that even if all firms are ex-ante identical and pay a common wage, workers will have non-degenerate outside value distributions. This means that the probability that a worker leaves a given job is a smooth function of their current wage and a competing outside wage, creating a smooth trade-off for firms between wage costs and turnover probabilities. If identical firms are facing the same smooth trade-off, then all firms will choose the same wage, generating the point mass of wages in equilibrium.¹⁴ ¹⁵ Wage dispersion can arise in equilibrium from variation in hiring costs c_j or variation in A_j when $\sigma \neq 0$.

Characterizing $\mathcal{E}_j \equiv \varepsilon_{R,w_j} - \varepsilon_{S,w_j}$ The equilibrium values of each firm's recruiting and separation elasticities depend on many equilibrium objects, including the firm's wage, the distribution of earned wages, the distribution of posted wages, the typical duration of employment, and the unemployment rate, all of which are complex functions of parameters. The parameter that most influences the recruiting and separation elasticities is γ , the parameter that governs the variance of the Type-1 extreme value preference draws. When γ is higher, the variance of these draws is lower, making relative wages more important in workers' decisions to move across firms, which makes workers more wage sensitive and raises $\varepsilon_{R,w}$ and $-\varepsilon_{S,w}$ all else equal, and vice-versa. Empirical evidence (Bassier et al., 2022; Datta, 2023; Sokolova and Sorensen, 2021) tends to find recruiting and separations elasticities to be around 2 and -2, respectively, with the sum typically around 4. When calibrating the model, we treat γ as a free parameter to make $\varepsilon_{R,w} - \varepsilon_{S,w}$ close to 4.

Decomposing the Marginal Product

At the firm level, we decompose the marginal product of labor into the shares that go towards wages, recruiting costs, and profits. The expressions of these shares are reported in Table 2.1. Each

¹⁴We verify numerically that no frictional wage dispersion is an equilibrium result for reasonable parameters values. The no-dispersion equilibrium can break down when the variance of preference draws becomes sufficiently small (i.e., γ becomes sufficiently large). However, this typically requires a sum of recruiting and separation elasticities over 10, far above the empirically estimated range. We prove unique optimal wages for the special case outlined in Section 2.2 in Appendix B.7.

¹⁵In addition to helping with tractability, idiosyncratic preference draws allow for workers to voluntarily take wage cuts when moving between jobs (Sorkin, 2018; Hall and Mueller, 2018) in any equilibrium with wage dispersion.

Table 2.1: Decomposing Marginal Product into Wages, Recruiting Costs, and Profits

Outcome	Definition	Formula
Wage share of $MRPL$	$\frac{w_j}{MRPL_j}$	$\frac{(1 + \chi)\mathcal{E}_j}{1 + (1 + \chi)\mathcal{E}_j + \sigma}$
Recruiting cost share of $MRPL$	$\frac{\mathcal{C}_j/N_j}{MRPL_j}$	$\frac{1}{1 + (1 + \chi)\mathcal{E}_j + \sigma}$
Labor market profit share of $MRPL$	$\frac{MRPL_j - w_j - \mathcal{C}_j/N_j}{MRPL_j}$	$\frac{\sigma}{1 + (1 + \chi)\mathcal{E}_j + \sigma}$
Inverse labor supply elasticity	$\varepsilon_{w_j, N_j} \equiv \frac{\partial \log w_j / \partial \log A_j}{\partial \log N_j / \partial \log A_j}$	$\frac{\sigma}{1 + (1 + \chi)\mathcal{E}_j - \varepsilon_{\mathcal{E}_j, w_j}}$
Recruiting cost-adjusted wage markdown	$\frac{w_j + \mathcal{C}_j/N_j}{MRPL_j}$	$\frac{1 + (1 + \chi)\mathcal{E}_j}{1 + (1 + \chi)\mathcal{E}_j + \sigma}$

Notes: This table shows the wages, recruiting cost, and labor market profit shares of marginal product, the recruiting cost-adjusted markdown, and the inverse labor supply elasticity. Analogous results for the general β_f case can be found in Appendix B.6. Recruiting cost per worker is defined as $\mathcal{C}_j/N_j \equiv c_j(V_j/N_j)^{1+\chi}N_j^\sigma = \mathcal{C}_j/H_j \times S_j$, which is equivalent to the costs per hire times the separation rate, where H_j is the number of hires with $H_j \equiv R_j \times V_j$.


expression depends on the sum of recruiting and separation elasticities $\mathcal{E}_j$ and the recruiting cost parameters χ and σ . The first three rows of the table simply represent a decomposition: what share of marginal product is spent on wages, spent on per worker recruiting, and what remains. We label the share of marginal product that is left after accounting for wage and recruiting costs “labor market profits.”¹⁶

The fourth row in Table 2.1 repots the inverse labor supply elasticity, which shows the elasticity of the optimal wage with respect to the firm’s optimal employment level in response to a permanent change in firm specific product demand A_j . Within the inverse labor supply term is a super-elasticity $\varepsilon_{\mathcal{E}_j, w_j}$, which is the elasticity of $\mathcal{E}_j$ with respect to w_j . This super-elasticity captures that the sum of recruiting and separation elasticities may change with the wage. The final term in the table is

¹⁶Note that the marginal benefit of hiring a worker includes the present discounted value of future marginal products, as well as lower costs of future recruiting. Firms equalize the marginal benefit and marginal cost of a hire, which can be expressed as the following when firms do not discount: $\frac{w}{S} + (1 + \chi)\frac{\mathcal{C}/V}{R} = \frac{MRPL}{S} - \frac{(\sigma - \chi)\mathcal{C}/N}{S}$.

Figure 2.1: Wage Bill, Recruiting Costs, and Marginal Product

 (a) Standard Monopsony: $\sigma = \chi \rightarrow \infty$

 (b) CRS Recruiting: $\sigma = 0$


Notes: This figure shows the marginal product, marginal cost, and labor supply curves for firms with different recruiting functions. Panel (a) shows a firm in a standard monopsony model with no recruiting function, equivalent to $\sigma = \chi$ and $\sigma, \chi \rightarrow \infty$ in our model. The firm's marginal cost line lies above the firm's labor supply curve, and the wage lies below marginal product when the marginal product curve and marginal cost curve intersect. Panel (b) shows a firm with a constant returns to scale recruiting function. For this firm, the optimal wage is the same regardless of firm size, and the share of recruiting costs in total costs is also constant across different choices of steady state firm size. The wage lies below marginal product at the optimum, with the gap entirely absorbed by recruiting costs.

the recruiting cost-adjusted wage markdown, which is the share of marginal product that is spent by the firm on wages and recruiting.¹⁷

Two parameterizations are of particular interest. The first is the limiting case in which the recruiting margin is shut down as a choice variable. This occurs when $\sigma = \chi$ and $\sigma, \chi \rightarrow \infty$. Plugging these parameter values into Table 2.1, we obtain the standard dynamic monopsony model without a recruiting margin, where the labor supply elasticity $\varepsilon_{N_j, w_j} = \mathcal{E}_j \equiv \varepsilon_{R, w_j} - \varepsilon_{S, w_j}$, the firm spends nothing on recruiting, and the wage markdown is $\frac{\mathcal{E}_j}{1 + \mathcal{E}_j}$. Thus, our model with a recruiting margin nests standard models of dynamic monopsony without a recruiting margin as a special case, including static models of monopsony without a recruiting margin. Figure 2.1a depicts this standard

¹⁷In our base case, we assume that recruiting costs are an expenditure that accrues to labor income. An alternative is that recruiting costs are subtracted from output, in which case the recruiting cost-adjusted wage markdown would be $w / (MRPL - \text{recruiting costs per worker})$. A third alternative is that recruiting effort may divert labor away from productive activity, such that labor inputs are $N - cV^{1+\chi}N^{-(\chi-\sigma)}$, which we explore in Appendix B.8. These different assumptions do not affect our core results.

case graphically: the firm's marginal cost curve lies above the labor supply curve, and the marginal product intersects marginal cost at a point where wages are below marginal product.

The second parameterization of interest is $\sigma = 0$, in which hiring costs depend only on the recruiting rate (vacancies per worker, V/N) and not directly on firm size. This parameterization implies an inverse labor-supply elasticity of zero and that the recruiting-cost-adjusted wage mark-down equals 1: even though wages lie below marginal product, the gap between wages and the marginal product is entirely absorbed by recruiting costs. This is depicted in Figure 2.1b. Within this $\sigma = 0$ case, there are two sub-cases: $\chi > 0$ and $\chi = 0$. In the former case, marginal hiring costs increase with the recruiting rate (V/N), which encourages higher wages and discourages vacancy posting, pushing up the wage share of marginal product and pushing down the average recruiting cost per worker. Under $\chi = 0$, recruiting costs are linear in the number of vacancies.

In Appendix B.13, we consider common amenities, amending the model so firms can invest in amenities but cannot spend on recruiting. We show that these models are not isomorphic if there are exogenous separations ($s_0 > 0$), leaving a distinct role for recruiting.¹⁸

Evidence on Recruiting Costs

Our emphasis on recruiting costs raises a natural question: are recruiting costs large enough to account for the difference between wages and marginal products, and is there evidence of diminishing returns to recruiting expenditure?

To begin, the ratio of wages to marginal product implied by a model with recruiting can be much closer to 1 than would be implied by a dynamic monopsony model without a recruiting margin. As a consequence, the level of recruiting costs does not need to be large in order for the recruiting cost-adjusted markdown to be equal to 1. To see this, consider the evidence that the sum of recruiting and separation elasticities $\mathcal{E} \equiv \varepsilon_{R,w} - \varepsilon_{S,w}$ is typically estimated to be around 4 (Bassier et al., 2022; Datta, 2023; Sokolova and Sorensen, 2021). In a model without a recruiting margin ($\sigma = \chi$, $\sigma, \chi \rightarrow \infty$), this would imply a wage share of marginal product equal to $\frac{\mathcal{E}}{1+\mathcal{E}} = 0.8$. Consider instead a different benchmark of a constant returns to scale recruiting function ($\sigma = 0$) but firms face convexity in recruiting costs as a function of vacancies per incumbent, with $\chi = 2$ as proximately the value that we estimate for χ . In that case, holding fixed the recruiting minus separation elasticities $\mathcal{E} \equiv \varepsilon_{R,w} - \varepsilon_{S,w} = 4$, the ratio of the wage to marginal product becomes $\frac{\mathcal{E}(1+\chi)}{1+\mathcal{E}(1+\chi)+0} \approx 0.92$. Thus, a higher value of χ pushes wages closer to marginal product, holding $\varepsilon_{R,w} - \varepsilon_{S,w}$ fixed.

When interpreting evidence on turnover or hiring costs, it is important to distinguish between average and marginal hiring costs. In our setting, the marginal hiring cost equals $(1 + \chi)$ times the average hiring cost. Evidence on average hiring costs suggests they are moderate: Manning (2011) surveys the literature and Dube et al. (2010) document average hiring costs of around 2-5% of the wage bill. Blatter et al. (2012) find the average cost of replacing skilled workers to be 10–17 weeks of salary. Michailat and Saez (2024) argue that servicing a vacancy requires approximately one full-time worker. In terms of marginal hiring costs, Muehlemann and Pfeifer (2016) find that marginal hiring costs are increasing, consistent with $\chi = 1.3$. Bloesch and Weber (2023) find evidence of

¹⁸While spending on amenities may increase the desirability of a job, better amenities does not help a firm match with more workers. With exogenous separations, the firm must have a finite and positive long-run labor supply elasticity, as worker inflows are bounded by the number of matches each period.

congestion in onboarding software developers, consistent with convex costs in the rate of training new workers. Evidence on separation costs, and by implication the marginal cost of a hire, tends to be larger: Kline et al. (2019), Jäger and Heining (2022), Bertheau et al. (2021), and Bloesch et al. (2022) all point to separation costs of 1–3 years of a workers' wages using worker deaths, though deaths may be significantly more costly than replacing a worker after a typical quit.¹⁹ Models that calibrate an object like χ in order to match empirical relationships between vacancy rates and hiring rates imply values of χ between 1 and 8 (Davis et al., 2013; Kaas and Kircher, 2015; Gavazza et al., 2018; Mongey and Violante, 2025).

Our model has strict implications for the present value of a match to a firm and the level of marginal hiring costs in steady state. Given the expressions for wages and recruiting costs in Table 2.1, the recruiting expenditure per incumbent each period in steady state is $\frac{w}{\mathcal{E}(1+\chi)}$. In steady state, the number of hires equals the number of separations SN , so the recruiting expenditure per hire is $\frac{w}{S\mathcal{E}(1+\chi)}$. Given the convexity parameter χ , the marginal hiring cost equals $(1+\chi)$ times the average cost per hire, so the marginal hiring cost is $\frac{w}{S\mathcal{E}}$. Therefore, with $\mathcal{E} = 4$ and a monthly separation rate between 0.04 and 0.06,²⁰ this would imply a marginal hiring cost per hire of around 4–6 months of wages. To be consistent with evidence on the level of hiring costs, this would imply a value of χ roughly between 1 and 3, which is consistent with our estimates in the next section.

Wage and Employment Responses to Product Demand Shocks

In this section, we derive the response of wages to idiosyncratic firm demand shocks, which we use in Section 2.3 to estimate the recruiting cost parameters χ and σ , and by extension estimate the inverse labor supply elasticity $\varepsilon_{w,N}$ and the share of marginal product retained as profits.

We consider how a firm's optimal wage responds to an idiosyncratic shock to $z_{j,t}$, i.e., an idiosyncratic product demand shock. Define $h_{j,t}$ as the firm's hiring rate (hires per incumbent): $h_{j,t} \equiv V_{j,t}R(\mathcal{V}_{j,t,t})/N_{j,t-1}$. We log-linearizing the firm's problem around each firm's steady state, denoting percent deviation from steady state using tildes. The firm's optimal deviation of new hire wages from steady state $\tilde{w}_t$ is

$$\begin{aligned} \tilde{w}_t = \frac{1}{\Omega} & \left(\frac{\varepsilon_{R,w}}{\varepsilon_{R,w} - \varepsilon_{S,w}} (\chi \tilde{h}_t + \sigma \tilde{N}_{t-1}) \right. \\ & \left. + \frac{-\varepsilon_{S,w}}{\varepsilon_{R,w} - \varepsilon_{S,w}} \left(S \sum_{\kappa=0}^{\infty} (1-S)^{\kappa} (-(1+\chi)\varepsilon_{R,w}\tilde{w}_{t+\kappa+1} + \chi \tilde{h}_{t+\kappa+1} + \sigma \tilde{N}_{t+\kappa}) \right) \right) \end{aligned} \quad (2.13)$$

where $\Omega = 1 + \frac{\varepsilon_{R,w}}{\varepsilon_{R,w} - \varepsilon_{S,w}} (-\xi_{R,w} - \varepsilon_{S,w} + (1+\chi)\varepsilon_{R,w}) + \frac{-\varepsilon_{S,w}}{\varepsilon_{R,w} - \varepsilon_{S,w}} (-\xi_{S,w} - \varepsilon_{S,w})$. The ξ terms are super-elasticities: $\xi_{R,w} \equiv \frac{\partial \varepsilon_{R,w}}{\partial w} \frac{w}{\varepsilon_{R,w}}$ and $\xi_{S,w} \equiv \frac{\partial \varepsilon_{S,w}}{\partial w} \frac{w}{\varepsilon_{S,w}}$, which capture how the recruiting and

¹⁹We consider hiring costs that do not vary with wages or labor market tightness, such as training costs, in Appendix B.9. For these costs, higher wages save turnover costs only by preventing separations, resulting in more weight on the separation elasticity and less weight on the recruiting elasticity when setting wages.

²⁰The rate of worker flows across firms is similar Denmark and in the US. Caldwell and Harmon (2019). Bagger et al. (2022) find that the monthly separation rate in Denmark is around 6.6 percent. The total monthly job separation rate from the Job Opening and Labor Turnover Survey (JOLTS) in the US averages between 3 and 4 percent.

separation elasticities change with the firm's choice of wage. These super-elasticities $\xi_{R,w}$ and $\xi_{S,w}$ are not quantitatively important, and the model can be calibrated such that they are of opposite sign and similar in magnitude, canceling each other out.

The coefficient on the hiring rate is increasing in χ : when firms face diminishing returns to recruiting expenditure, matching with an additional worker becomes more costly. Given a limited pool of matched workers, firms raise wages to accelerate hiring, and this incentive is stronger when the marginal match is more expensive. Thus, a higher χ amplifies the firm's optimal wage response for a given hiring target. The coefficient on firm size is increasing in σ : a higher value of σ means that hiring costs increase with firm size, when holding the vacancy rate fixed; thus, leading to a higher optimal wage with larger firm size.

The intuition for identification is that in response to an idiosyncratic demand shock, the firm will respond with a path of wages, hiring rates, and employment levels. Identification of χ and σ will come from whether the wages of new hires co-vary with the hiring rate or the firm's employment level. That is, if wages are related to (recent and future) firm size, then that is evidence that $\sigma > 0$. If instead the wage of new hires co-moves with the hiring rate but not firm size, then that is evidence that $\chi > 0$ but $\sigma = 0$. Lastly, if wages of new hires do not change at all as the firm hiring rate and employment level move around, then that is evidence that $\chi = 0$ and $\sigma = 0$. Differing variation in hiring rates and employment arises because hiring is a flow, which responds quickly to a shock, while firm size is a stock, which lags. In the next section, we will estimate the empirical impulse responses of switcher wage growth, hiring rates, and employment levels at exporting firms to idiosyncratic firm export demand shocks in Denmark, and we will use these results to estimate χ and σ .

Discussion: Switcher vs. Stayer Wages Equation (2.13) notably does not include stayer wages. This is in contrast with much of the existing monopsony literature, which focuses on the wages of stayers. In our model, the response of stayer wages to demand shocks is not informative for estimating χ and σ . To illustrate this, we will discuss our assumptions on stayer wages in the model. In equation (2.9), we specified that firms pay at least the posted wage. This means that in response to a shock, stayer wages may rise, but stayer wages may also return back to their prior level if firms hire new workers and bring the marginal product of labor back down. We specified the constraint this way because it will be consistent with the temporary response of stayer wages to demand shocks in Section 2.3.

However, we could have specified incumbent wages to be downwardly rigid, so that stayers have permanent wage changes in response to demand shocks. Even under this different assumption, the response of stayer wages would not be informative for estimating χ and σ . Consider the following scenario: in response to a demand shock, incumbents receive permanently higher wages. With higher wages, these incumbents turn over more slowly, decreasing the number of hires the firm needs to make. Once the demand shock fades, the firm has a below steady-state separation rate, allowing the firm maintain its size with a below steady-state recruiting and hiring rate. This would mean that the firm could offer new hires lower wages, and with marginal workers being relatively cheap, the firm would find it optimal to have above steady-state employment and below steady-state marginal product of labor. In total, this example would result in elevated wages for stayers, even while the marginal product of labor falls *below* its steady state level, with wages potentially above marginal product. All the while, equation (2.13) still holds, and the path of stayer wages is

irrelevant for estimating χ and σ , conditional on the path of total employment, the hiring rate, and new hire wages. Thus, our model assumption on stayer wages does not influence our estimation of χ and σ . Moreover, the *realized* path of stayer wages and their retention rates do not add any information for the purpose of estimating χ and σ , since the retention behavior of stayers is implied by the firm's employment level and hiring rate.

Special Case: Isoelastic Labor Supply Curves

We now highlight a highly tractable special case that closely relates our model to existing models of monopsony.

Proposition 1. *If workers have log utility over consumption ($\eta \rightarrow 1$), there are no exogenous separations ($s_0 = 0$), workers fully discount the future ($\beta_w \rightarrow 0$), and σ is not too negative, then there exists a unique stationary equilibrium such that:*

1. *firms' optimal wage is constant in steady state,*
2. *ex-ante identical firms pay a common wage, and*
3. *the recruiting elasticity minus the separation elasticity $\mathcal{E}_j \equiv \varepsilon_{R,w_j} - \varepsilon_{S,w_j}$ is constant and is equal to γ for any w_j .*

Proof: See Appendix B.7. This case is particularly useful because when we shut down the recruiting margin ($\sigma = \chi$ and $\sigma, \chi \rightarrow \infty$), the model in steady state is isomorphic to static models where firms face an isoelastic residual labor supply curve, as is the case in Lamadon et al. (2022) without amenities, in Card et al. (2018) when $b = 0$ (in their notation), in Berger et al. (2022) when firms are atomistic, and in Manning (2003) when all firms have the same labor supply elasticity. By nesting this common special case, we can more clearly isolate how introducing a recruiting margin generates different implications relative to this important benchmark for markdowns and profits from wage setting power. We will also use this special case of log utility, myopic workers, and no exogenous separations in our exercise with concentration in Section 2.4, as the closed-form results with labor market concentration will depend on firms having no exogenous separations.

2.3 Empirical Evidence and Parameter Estimation

In this section, we estimate the effect of persistent firm-level export demand shocks on Danish manufacturing firms' employment, hiring rates, and the wage growth of job switchers and stayers. This exercise allows us to estimate the key parameters of our model: the recruiting cost function parameters, χ and σ , and the inverse labor supply elasticity, $\varepsilon_{w,N}$. We then use our estimates of σ and χ in our section on the labor market oligopsony and the profit puzzle in Section 2.4.

Data

We construct an annual employer-employee matched dataset using Danish administrative data for the period 2001-2019. The administrative datasets we use are maintained by Statistics Denmark.

For workers' employment history and wages, we use the registers in the Integrated Database for Labor Market Research (IDA), which reports workers' earnings, hours worked, occupation, employer, and demographic information. We add information about the workers' education from the education register (UDDA). We link the labor market data with firm characteristics from the Danish business register (FIRM), including revenue, employment, and NACE industry classification. We measure employment in full-time equivalents (FTEs) throughout, and we use two-digit industry codes in our baseline analysis. We use customs and trade data from two sources to construct our firm-level export demand shocks, in two steps. First, we obtain data on the goods exports of Danish firms by product and destination country from the customs data register (UHDM). We use 4-digit Harmonized System (HS) product codes. Second, we use data on international trade flows from the United Nations Comtrade database to measure each country's imports by product and partner source country.

Sample Our main estimation sample is a panel of exporting manufacturing firms. We drop observations with less than 5 employees or revenue below DKK 1 million. In our baseline regression specification, we control for lags of the export demand shock. Including lagged values imposes a balance requirement on our sample: firms must be observed in the data for at least five consecutive years: the year of the shock and the four preceding years. To maintain a consistent sample throughout the paper, we restrict further slightly by including only firms that have non-missing export demand shocks each of the 4 years prior to the shock. Our baseline sample thus includes shocks in years 2006–2019.

The estimation sample consists of 22,453 firm-year observations with 2,961 unique firms. We report descriptive statistics in Table 2.2. Denmark is a small and very open economy, so most manufacturing firms are exporters, and our sample includes many small- and medium-sized firms. The median firm size is 38 employees, but the distribution is heavily right-skewed with mean employment at 114 employees and several firms of considerable size at the very top. The firms in our sample produce multiple products and export to multiple countries. On average, firms export 13 different products (at the level 4-digit HS codes) and export to 8 destination countries.

Table 2.3 reports summary statistics for the sample of workers who switch into our sample of export manufacturing firms. The sample is 70% male, and the typical worker has 14 years of education and is 42 years old. The average hourly wage is 208.9 DKK/ hour (approximately \$30/hour), and the average worker sees an hourly wage change of 4.7 log points when changing jobs. There are 115,228 unique individuals switching jobs in our sample and 131,973 switcher-year events.

Construction of Export Demand Shocks

Our goal is to estimate the recruitment cost parameters χ and σ , which, in turn, will be informative about the inverse labor supply elasticity $\varepsilon_{w,N}$. For this purpose, we need a firm-specific product demand shock that is persistent and plausibly exogenous.

We construct shift-share export demand shocks to Danish exporting firms drawing on Hummels et al. (2014). The idea is that Danish exporters are differentially exposed to changes in foreign demand depending on the products they export and the countries they export to. The shift-share shock variable consists of exposure shares and foreign demand shifters. We construct the demand

Table 2.2: Firm Summary Statistics

	Mean	Median	10th Percentile	90th Percentile
Employment (FTE)	114.1	37.6	10.0	200.6
Revenue (DKK Mn)	333.8	61.0	13.9	465.0
Exports (DKK Mn)	130.6	20.3	0.9	211.5
Export destinations	17.4	13.0	2.0	39.4
Exported products (4-digit HS)	13.0	8.0	2.0	28.5
Hiring Rate	18.9	15.5	2.9	37.8
Switcher Wage Growth	3.9	3.7	-15.6	23.2
Observations	22,453			
No. of firms	2,961			

Notes: Sample is exporting manufacturing firms observed in 2006–2019. Employment is measured in full-time equivalents (FTE). Revenue and exports are in DKK millions. “Export destinations” counts distinct destination countries; “Exported products” counts distinct 4-digit HS product lines. “Hiring rate” is $H_{j,t}/N_{j,t-1}$, where $H_{j,t}$ are new hires in year t (workers on the firm’s November payroll in t , not employed by the firm in $t-1$ or $t-2$, and with ≥ 50 hours in t). “Switcher wage growth” is the firm-year mean two-year log wage change for new hires who worked at another firm two years earlier: $\Delta \log w_{j,t}^{\text{switcher}} = \frac{1}{|\text{switchers}_{j,t}|} \sum_{i \in \text{switchers}_{j,t}} (\log w_{i,j,t} - \log w_{i,k,t-2})$. Percentiles are across firm-year observations; “Observations” counts firm-years and “No. of firms” counts unique firms in the estimation sample.

Table 2.3: Switching Worker Summary Statistics

	Mean	Median	10th Percentile	90th Percentile
2 Year Log Wage Growth	5.9	4.7	-21.4	33.3
Wage (kr/hr)	208.9	190.6	134.1	299.6
Educ Years	14.4	14.0	11.0	18.0
Age	41.8	42.0	29.0	54.6
Share Male	0.7			
Observations	131973	No. of workers		115228

Notes: Sample includes job switchers (workers moving from employer k to j) in 2006–2019 who are full-time ($\geq 1,400$ annual hours) at both the prior and new employer. “2-year wage growth” is $\log w_{i,j,t} - \log w_{i,k,t-2}$. “Wage (kr/hr)” is hourly earnings in Danish kroner. “Educ Years” comes from the education registry; “Age” is in years; “Share Male” is the fraction male. Percentiles are across worker-level switch events; “Observations” counts worker-year switches and “No. of workers” counts unique individuals.

shifters as follows. For each country-product pair, we compute the world import demand $WID_{c,p,t}$, which is country c 's imports of product p from all countries excluding Denmark. We then compute the one-year change in log world import demand $d \log WID_{c,p,t} \equiv \log WID_{c,p,t} - \log WID_{c,p,t-1}$, for each country-product-year cell.

A Danish firm j is exposed to a shift in foreign demand $d \log WID_{c,p,t}$ if they export product p to country c . We define the exposure shares as

$$\omega_{j,c,p,t-1} = \frac{\text{exports}_{j,c,p,t-1}}{\text{revenue}_{j,t-1}},$$

which captures how exposed firm j 's total revenue is to changes in foreign demand for a given country-product pair (c, p) . We measure the shares in the previous year. Finally, we combine the shares and the shifters to construct our firm-level shift-share export demand shock:

$$z_{j,t} = \sum_{c,p} \omega_{j,c,p,t-1} d \log WID_{c,p,t}.$$

If firms have domestic sales, then the shares do not add up to one when we sum across countries and products. Instead, they add up to the (lagged) export share of revenue,

$$\Omega_{j,t-1} = \sum_{c,p} \omega_{j,c,p,t-1}.$$


If the export share of revenue varies across firms in a way that is correlated with unobservables, then our estimates will be biased. To deal with this problem, we follow Borusyak et al. (2022) and control for the sum of the export shares. Specifically, since our main specification includes industry-by-year fixed effects, we control for the interaction between the industry-year fixed effects and the lagged export share $\Omega_{j,t-1}$.

Idiosyncratic Export Demand Shocks A potential concern is that the export demand shocks may not be completely idiosyncratic due to global product shocks. Imagine there is a global increase in demand for a specific product p' . All Danish firms that export this product will be affected, and if these firms compete for workers in the same local labor market, then wages may rise more than if only one firm was hit by a truly idiosyncratic shock. If our industry classification is a good approximation of the local labor markets that firms hire in, then industry-by-year fixed effects can absorb the common component of the shocks. However, it is difficult to know whether this is the case in practice. To isolate the idiosyncratic component of the export demand shock, we carry out a decomposition of the demand shifters proposed by Garin and Silv rio (2024). We can write the shifters as

$$d \log WID_{c,p,t} = \underbrace{\overline{d \log WID}_{p,t}}_{\text{common component}} + \underbrace{(d \log WID_{c,p,t} - \overline{d \log WID}_{p,t})}_{\text{idiosyncratic component}},$$

where $\overline{d \log WID}_{p,t} \equiv \frac{1}{|C|} \sum_{c \in C} d \log WID_{c,p,t}$ is the mean change in log world import demand across countries (with C denoting the set of destination countries, excluding Denmark). In other words, we decompose the foreign demand shifter in year t into a product fixed effect and a residual. The

Figure 2.2: Export Demand Shock Autocorrelation and Revenue Response



Notes: Panel (a) shows the coefficients from a regression of a trade shock in period $t + \nu$ on a trade shock in period t , controlling for industry-by-year fixed effects. The coefficient on the current shock z_t is 1 by construction. Standard errors are clustered at the firm level. Panel (b) plots the effect of 100 log point predicted increase in revenue on a firm's log revenue. A predicted 100 log point export demand shock increases revenue by 30 log points on impact, with a peak response of 40 log points the year after the shock. The effect on revenue mostly fades three years after the shock.

product fixed effect is the “common component”, which reflects global changes in demand for p . The residual term is the “idiosyncratic component” of the shifter, which captures changes in the demand for a product p from a destination country c *relative* to the global average change in demand for p . The decomposition is

$$z_{j,t} = \underbrace{\sum_p \bar{\omega}_{j,p,t-1} \overline{d \log WID}_{p,t}}_{\equiv z_{j,t}^{\text{common}}} + \underbrace{\sum_{c,p} \omega_{j,c,p,t-1} (d \log WID_{c,p,t} - \overline{d \log WID}_{p,t})}_{\equiv z_{j,t}^{\text{idiosyncratic}}},$$

where $\bar{\omega}_{j,p,t-1} \equiv \sum_{c \in C} \omega_{j,c,p,t-1}$ are product exposure shares. We then winsorize the idiosyncratic export demand shocks at the 1st and 99th percentiles in each year.

Shock Diagnostics Figure 2.2 shows the diagnostics of our idiosyncratic export demand shock. Figure 2.2a reports the results of regressing the leads and lags of the shock on the time t shock, including industry-by-year fixed effects. These shocks exhibit minimal autocorrelation (Dhyne et al., 2025).

Figure 2.2b estimates the effect of an idiosyncratic export demand shock on firm log revenue using local projections (see equation Section 2.3 and equation (2.14) for a discussion of the specification). The export demand shocks predict increases in actual revenue: a 100 log point predicted

increase in revenue on average results in a 40 log point increase in actual revenue as shown in Figure 2.2b. On average, export demand shocks are relatively small: the standard deviation of shocks is 0.046. This means that in a typical year, an exporting firm experiencing a one-standard deviation shock in export demand should expect actual revenue to increase by around 2 percent.

Firm Responses to Export Demand Shocks

Outcome Variables The firm-level outcome variables we consider are log revenue, log employment, the hiring rate, and the wage growth of job switchers and job stayers. We denote revenue by $\mathcal{R}_{j,t}$. Our measure of firm-level employment $N_{j,t}$ is the number of full-time equivalent (FTE) employees. We define the hiring rate as

$$h_{j,t} \equiv \frac{H_{j,t}}{N_{j,t-1}},$$

where $H_{j,t}$ is the number of new hires by firm j in year t . A new hire is defined as a worker who is employed at the firm in November of year t but not in either of the two preceding years. We also require that the new hire worked at least 50 hours at the firm in year t .

A job switcher at firm j at time t is a new hire who was employed at a different firm in the previous two years. We define switcher wage growth for firm j in year t as the average wage growth of the arriving job switchers:

$$\Delta \log w_{j,t}^{switcher} \equiv \frac{1}{|switchers_{j,t}|} \sum_{i \in switchers_{j,t}} (\log w_{i,j,t} - \log w_{i,k,t-2}),$$

where k denotes the prior employer in $t-2$ of the job switcher. We measure wage growth using two-year differences because the annual frequency of the labor market data makes it difficult to know exactly when job transitions are made within a year.

Empirical Specification We estimate the dynamic response to a firm-level export demand shock using reduced-form panel local projections (Jordà, 2005b). The export demand shock occurs in year t , and we will measure the response of the firm-level outcome variables from $t-4$ to $t+3$. Our main regression specification is

$$y_{j,t+\nu} - y_{j,t-1} = \beta_{\nu}^y z_{j,t}^{idiosyncratic} + \delta_{\nu}^y X_{j,t} + u_{j,t+\nu}^y \quad (2.14)$$

for each horizon $\nu \in [-4, 3]$, where the long-differenced outcome variable $y_{j,t+\nu} - y_{j,t-1}$ is either log revenue ($\log \mathcal{R}_{j,t+\nu} - \log \mathcal{R}_{j,t-1}$), log employment ($\log N_{j,t+\nu} - \log N_{j,t-1}$), or the hiring rate ($h_{j,t+\nu} - h_{j,t-1}$). Because firms can have zero hires in a year, we leave the hiring variable in levels and convert to elasticities later. The control vector $X_{j,t}$ includes three lags of the export demand shock, the contemporaneous value and three lags of the common shock (as the idiosyncratic and common shocks are negatively correlated mechanically, see Garin and Silvério (2024)), firm fixed effects, industry-by-year fixed effects, and industry-by-year fixed effects interacted with the sum of the lagged export share (Borusyak et al., 2022). We estimate a similar regression for the wage growth of job switchers, except we do not difference the wage growth variable, as the panel of the

switcher wage growth is unbalanced and requiring a non-missing value for $\Delta \log w_{j,t-1}^{switcher}$ would greatly reduce sample size. We estimate

$$\Delta \log w_{j,t+\nu}^{switcher} = \beta_{\nu}^w z_{j,t}^{idiosyncratic} + \delta_{\nu}^w X_{j,t} + u_{j,t+\nu}^w \quad (2.15)$$

for each horizon $\nu \in [-4, 3]$. Standard errors are clustered at the firm level.

Results We plot the response of the firm employment level, hiring rate, and wage growth of switchers in Figure 2.3.²¹ Panel (a) shows the response of employment and switcher wage growth. The solid circles show the effect of this shock on employment, and the hollow triangles report the effect on the wage growth of switchers, where each point represents a different cohort of new hires each year. Conditional on the sequence of shocks from $t - 3$ to $t - 1$, firms with different values of shocks in period t have similar pre-trends in both the level of employment and in the wage growth of workers who switched into the firm. On impact, an export demand shock that predicts a 100 log point increase in predicted revenue increases a firm's employment level by approximately 15 log points. The effect peaks a year later, with a response of 22 log points. After this, the effect on employment decreases, with employment approximately returning to its pre-shock level three years after the shock. For switcher wages, in the year of the shock, job switchers who move into firms that have a 100 log point export demand shock receive 11 log points higher wage growth than workers who arrive at a firm with a 0 percent export demand shock. However, in the years after the shock, when the firm is larger but is no longer growing, the wage growth of newly hired switchers is similar to or lower than what switcher wage growth was before the shock occurred.


Panel (b) of Figure 2.3 reports the path of firms' hiring rate and reprises the path of switcher wage growth. On impact, the hiring rate increases by 10 percentage points. Since the average hiring rate is 18.9 percentage points, a 10 percentage point increase corresponds to a 53 percent increase on impact.

Importantly, these figures show that firms pay higher wages to new hires when firms are growing, but not necessarily when firms are large. At $t = 0$, when firms are initially shocked, firms pay high wages when they are growing and rapidly hiring. This is consistent with diminishing returns in hiring ($\chi > 0$). By $t = 1$, firms are larger than before the shock, but the hiring rate and the switcher wage growth are similar to pre-shock levels. These results are also consistent with an elastic labor supply curve ($\sigma = 0$): as long as firms are not rapidly hiring, large firms do not pay higher wages. Our results differ dramatically from static models of monopsony, which would predict that wages would remain above its steady state level as long as the level of the firm's employment is above its steady state.

In Appendix B.2, we estimate the effect of export demand shocks on the raw average log wage of newly hired switchers, both in the year that the workers arrive and workers' wage two years prior at their previous firm. Estimates are noisy, suggesting that differencing switcher wages removes highly variable worker fixed effects. The evidence in Appendix B.2 suggests that the composition of switchers, based on their wage prior to switching firms, did not significantly change in response to the export demand shocks.

²¹We align the samples so that the coefficient on firm employment level and switcher wage growth within each time horizon reflects the same sample. However, this means that not every horizon has exactly the same sample, as not every firm makes a full time hire of a switcher each year. The sample for each horizon has between 10,157 and 15,539 firm-year observations.

Figure 2.3: Response of Firm Employment, Hiring Rate, and Switcher Wage Growth to Export Demand Shocks



Notes: The figure plots the path of the firm employment level, firm hiring rate, and wage growth of switchers in response to a 100 log point predicted increase in revenue due to an export demand shock in period t . The blue circles in Panel (a) show that firms increase employment by approximately 15 log points on impact, with a peak effect of 22 log points the following year. The hollow red triangles show the effect of the period t trade shock on the wage growth of different cohorts of arriving switchers. Workers who switch into the firm in period t experience wage growth that is 11 log points greater than workers who switch into a firm that is not shocked. However, workers who switch into the shocked firm in subsequent years receive similar or smaller wage increases than workers who arrived prior to the shock. Panel (b) shows that hiring increases by 10 percentage points on impact, while future hiring rates are similar to or below the hiring rates prior to the shock. With an average hiring rate of 18.9 percentage points, a 10 percentage point increase in the hiring rate is equivalent to a 53 percent increase in the hiring rate in the year of the shock.

Robustness We now address a number of potential concerns with our empirical strategy. First, the composition of new hires may change in response to the shock. In particular, when hiring accelerates, firms may shift toward younger workers, who typically experience faster wage growth than older workers. Such compositional shifts can mechanically raise estimated wage growth for switchers. In Appendix B.3, we address this by residualizing switchers' individual wage growth using standard Mincerian controls and then averaging the residuals at the firm-year level. This adjustment has only modest effects, as shown in Figure B.3a.

Second, firms may selectively hire from high- or low-wage firms when expanding. We address this by estimating firm wage effects using the standard two-way fixed effects (AKM) method (Abowd et al., 1999), and include the estimated firm effects when residualizing workers' wages. This adjustment has only very small effects, as shown in Figure B.3b in Appendix B.3.²²

²²In Appendix B.14, we allow for heterogeneity in worker types and let firms set wages per efficiency unit. Our model can thus be consistent with firms having time-varying, log-additive wage effects for workers of different types, justifying the inclusion of AKM firm fixed effect controls when residualizing switchers' wages.

Third, rapidly expanding firms may relax hiring standards and hire workers who are less qualified than their usual hires. This can bias measured wage changes upward during expansion, conflating a wage premium with occupational upgrading. To address this, we compute market-wide average wages by occupation. Then, when constructing residualized wage changes for workers, we include as a control the difference between the average wages of the worker's prior and new occupation. This adjustment has only very small effects, as shown in Figure B.3c in Appendix B.3.

Fourth, an idiosyncratic shock to a large firm may affect the local labor market, violating the Stable Unit Treatment Value Assumption (SUTVA)—no interference across units and no multiple versions of treatment. In general, we are not concerned that this materially affects our results: the standard deviation of the shock is 0.046 log points, implying approximately a 1 percent increase in employment in the second year. The median firm in our sample has fewer than 40 employees, so a typical shock generates fewer than one additional full-time-equivalent employee after two years. This is a very small shock in the local labor market, so we assume SUTVA holds. However, as a robustness check, we estimate responses of employment and wages for small firms with log full-time-equivalent employment less than four (approximately 55 full-time-equivalent workers) and report the results in Figure B.3d in Appendix B.3. We find broadly similar results: employment rises on impact and remains elevated for a few periods, while switchers' wage growth increases on impact but declines once the firm is no longer expanding.

Lastly, firms may be hiring from other firms that are also affected by the shock, particularly when typical hires come from competing firms or similar industries. If the shock does not perfectly isolate idiosyncratic demand shifts, expanding firms may hire from similarly affected firms. To address this, we reconstruct the wage-growth measure using only hires from non-exporting firms and report the paths of employment, hiring, and switcher wage growth in Figure B.3e and B.3f. We find a nearly identical pattern: switchers' wage growth rises and falls with the hiring rate, while employment remains elevated even after switchers' wage growth returns to pre-shock levels.

Firm-by-Cohort AKM Interpretation In the model, the wage variable $\tilde{w}$ is the percent deviation from steady state wage of new hires in response to shocks to A_j , and we abstracted from heterogeneity in worker types. However, in the data, job switchers are of course heterogeneous. We generalize our model in Appendix B.14 so that workers have heterogeneous productivity (efficiency units). If workers have log flow utility, then firms choose to pay all worker types a common wage per efficiency unit of labor. In response to demand shocks, firms will raise the wage per efficiency unit of labor to all workers proportionately. $\tilde{w}$ can therefore be interpreted as the percent deviation from steady state of the firm's wage per efficiency unit of labor, which is log additive with heterogeneous worker fixed effects. This is like an AKM decomposition, except that the firm wage effects are cohort specific.

The empirical analogue to $\tilde{w}$ will be regression coefficients β_ν^w , which represent how firm deviations in new hire wage rates vary with the size of export demand shocks. In addition to the standard strict exogeneity assumption that worker-firm match effects are uncorrelated with mobility decisions, interpreting β_ν^w as an unbiased estimate of $\tilde{w}$ also requires that worker-firm match effects are uncorrelated with export demand shocks. Additionally, the wage effects of firms that workers leave must either be uncorrelated with export demand shocks or controlled for, which we do by residualizing for prior firm AKM effects, as discussed in the previous section. For more details on extending the model to have worker heterogeneity, see Appendix B.14. We discuss the mapping

of the estimated empirical coefficients β_ν^w to model objects $\tilde{w}$ in the next section and in Appendix B.1.

Stacked Regression Next we estimate a truncated version of the impulse responses in a stacked regression. We use these estimates in the next subsection when we estimate our structural parameters χ and σ , residualizing for worker demographics as described in the first robustness check in this section. Specifically, we stack the dependent variables in equations (2.14) and (2.15) and estimate the regression jointly. This imposes that all variables are non-missing, which particularly binds for the switcher wage growth variable, shrinking the sample size. The results are reported in Table 2.4. The estimates are similar to those in Figure 2.3.

Table 2.4: Estimated Coefficients from Stacked Regression

$\hat{\beta}_0^w$	$\hat{\beta}_1^w$	$\hat{\beta}_2^w$	$\hat{\beta}_0^h$	$\hat{\beta}_1^h$	$\hat{\beta}_2^h$	$\hat{\beta}_0^N$	$\hat{\beta}_1^N$	$\hat{\beta}_2^N$
0.141	0.016	-0.002	0.146	0.025	-0.042	0.158	0.251	0.176
Observations: 8,384								

Notes: This table reports the estimated coefficients from a stacked regression of equations (2.14) and (2.15) for horizons $\nu \in [0, 2]$, which jointly estimate the responses of switchers' wage growth, the hiring rate, and firms' log employment level to idiosyncratic export demand shocks.

These regression coefficients $\{\beta_\nu^h\}$, $\{\beta_\nu^w\}$ and $\{\beta_\nu^N\}$ will be used as the empirical analogues of the model's log-linearized firm responses to shocks $\{\tilde{w}_t\}$, $\{\tilde{h}_t\}$, $\{\tilde{N}_t\}$ from equation (2.13). See Appendix B.1 for how these objects map into each other.

Estimating the Recruitment Cost Function Parameters

In the previous section, we estimated the causal effect of an idiosyncratic export demand shock on the path of firm employment levels, hiring rates, and switcher wage growth. In this section, we use the results to estimate our model parameters of interest χ and σ and to obtain estimates for the inverse labor supply elasticity and profit share of marginal product.

We estimate the recruitment parameters χ and σ using minimum distance estimation. Our estimation strategy involves plugging empirical estimates into a model-derived equation that describes how the variables are related, and then choosing parameter values to minimize the sum of squared errors of that expression.²³ Define π as the sequence of the firm's log linearized choice variables: $\pi \equiv [\tilde{w}_t \ \tilde{w}_{t+1} \ \dots \ \tilde{h}_t \ \tilde{h}_{t+1} \ \dots \ \tilde{N}_t \ \tilde{N}_{t+1} \ \dots]'$. These variables are the model objects that describe the response of the firm's choice variables to a product demand shock in period t in terms percentage deviations from steady state. The empirical analogue $\hat{\pi} \equiv [\hat{\beta}_0^w \ \hat{\beta}_1^w \ \dots \ \hat{\beta}_0^h/\bar{h} \ \hat{\beta}_1^h/\bar{h} \ \dots \ \hat{\beta}_0^N \ \hat{\beta}_1^N \ \dots]$

²³Estimating the parameters using standard indirect inference would be challenging, as the large state variable of past committed wages for each cohort of workers makes simulating the model difficult.

is the estimated response of the corresponding variables in the data to an identified idiosyncratic demand shock (scaling the hiring rate by $\bar{h}$ to express it in percentage deviations).

Equation (2.13) describes how the endogenous variables are related to each other in response to a shock in the model. This equation provides infinitely many conditions, as w_t depends on N_{t-1} , period t variables, and future variables; w_{t+1} depends on N_t , period $t+1$ variables, and future variables, and so on. We truncate this infinite sequence of equations to T conditions and express them as

$$f(\pi, \chi, \sigma, \Theta) \equiv \mathcal{A}(\sigma, \chi, \Theta)\pi = 0, \quad (2.16)$$

where $\mathcal{A}(\sigma, \chi, \Theta)$ is a matrix with elements composed of σ , χ , and endogenous model objects $\Theta \equiv \{S, \varepsilon_{R,w}, \varepsilon_{S,w}, \xi_{R,w}, \xi_{S,w}\}$ that turns the conditions of equation (2.13) into a T equation linear system.

The estimation of χ and σ depends on the endogenous model objects in Θ , which potentially creates a fixed point problem: χ and σ are parameters that, in equilibrium, partially determine the endogenous variables in Θ needed to estimate χ and σ . Fortunately, this problem is easily addressed in two steps. First, we choose common values for the objects in Θ directly, setting $\varepsilon_{R,w} = -\varepsilon_{S,w} = \xi_{S,w} = -\xi_{R,w} = 2$ and $S = .05$. Given these values, we can estimate χ and σ . Then we calibrate the remaining parameters so that the model in equilibrium delivers our chosen values of endogenous outcomes Θ while also using our estimated values of χ and σ . There are many calibrations that satisfy this procedure, including the calibration in Section 2.4 as well as in Appendix B.12, so we will not have to take a stance on a specific calibration here.

Using our chosen values of Θ , we plug in for π with $\hat{\pi}$ from Table 2.4 and choose σ and χ to minimize a quadratic loss function

$$\min_{\sigma, \chi} f(\hat{\pi}, \sigma, \chi, \Theta)' f(\hat{\pi}, \sigma, \chi, \Theta). \quad (2.17)$$

The parameters are just identified if $T = 2$. Intuitively, identification of χ and σ comes from the different timing of the responses of hiring and employment. χ is primarily influenced by the comovement of wages with the hiring rate, and identification of σ comes from the residual variation of wages with employment that hiring cannot explain. We provide detailed steps in Appendix B.1.

One technical issue is that our model is best calibrated at a monthly frequency (to limit time aggregation bias), while our empirical results are annual. We discuss how we map our annual empirical estimates into the monthly model in Appendix B.1. Broadly speaking, we assume that hires are evenly distributed within a year, and we assign each annual estimate of $\hat{\pi}$ into 12 periods of corresponding model “months”. We use the first three years of estimated responses (years 0-2) and truncate the infinite sum in equation (2.13) such that switcher wages depend on future variables at most one year in advance. Converting to a monthly linear system gives 24 “months” of equations that depend on 9 estimated annual variables in $\hat{\pi}$, making $\mathcal{A}(\sigma, \chi, \Theta)$ a 24×9 matrix. $\mathcal{A}(\sigma, \chi, \Theta)$ is thus constructed such that the wages of a worker hired earlier within a year (suppose year 0) place greater weight on $\hat{\beta}_0^h$ and less weight on $\hat{\beta}_1^h$ than the wages of a worker hired later in year 0.

Table 2.5 reports the point estimates for σ , χ , and the inverse labor supply elasticity $\varepsilon_{w,N}$. We estimate $\sigma = 0.009$, $\chi = 2.21$, and the inverse labor supply elasticity $\varepsilon_{w,N} = 0.001$. To characterize the uncertainty of our estimates, we construct a distribution of estimates by sampling from the variance-covariance matrix of coefficients, fitting σ and χ to minimize (2.17) for each draw. The 10th and 90th percentiles of estimates for χ are 0.79 and 3.95, respectively. The 10th and 90th

Table 2.5: Fitted Parameters χ and σ and Inverse Labor Supply Elasticity $\varepsilon_{w,N}$

Parameter	χ	σ	$\varepsilon_{w,N}$
Estimate	2.21	0.009	0.001
10th percentile	0.79	-2.63	-0.18
90th percentile	3.95	3.33	0.27

Notes: This table reports point estimates and the 10th/90th percentiles for χ , σ , and the implied inverse labor supply elasticity $\varepsilon_{w,N}$ obtained via minimum-distance estimation.

percentiles of estimates for the inverse labor supply elasticity are -0.18 and 0.27 , respectively. We report distributions of the estimates in Appendix B.1.²⁴

Worker-Level Regressions for Switchers and Stayers


In Section 2.3, we focused on the path of new hire wages across arrival cohorts of workers. Three natural questions arise. First, are workers who switch into expanding firms on different wage trajectories? Second, are the wage effects for new hires persistent after moving firms? Third, how do the wages of job stayers respond? To answer these questions, we estimate the path of individual worker wages for both new hires and stayers in an event study specification. First, we find that workers who switch into positively shocked firms do not have significantly different pre-trends than comparable workers who switch into firms that receive a zero shock. Second, wage effects for workers hired in the shock year are persistent, while workers hired the following year show no effect. These results support our model assumptions of firm committing to posted wages and cohort-specific wage setting. Third, stayers' wages increase modestly in response to the export demand shocks, consistent with the literature's estimates of firm-specific labor supply elasticities of around 4 when estimated using stayer wages.

Job Switchers In this section, we estimate an event study showing the path of individual workers' wages as they change jobs, using worker-level local projections. As before, we focus on full-time workers who also worked full-time in their previous job. Therefore, job switchers are defined as workers in their "main job" who worked full-time (at least 1,400 hours) at firm j in year t and were full-time in their "main job" at some firm $k \neq j$ in year $t-2$. We use two-year changes because of how the data is constructed: employment is reported as of a worker's employer in November. Workers who are full-time in year t may have switched jobs in year $t-1$, but would not yet have completed a full year with the new employer by November of year $t-1$.

Equation (2.18) estimates the effect of a period t shock on a worker's wage growth from $t-2$ to $t+\nu$. The sequence $\{\beta^\nu\}$ traces out the effect of an export demand shock hitting in year t on the path

²⁴We report percentiles and distributions because our estimates of $\hat{\chi}$ are not normally distributed.

Figure 2.4: Path of Switcher Wage Level by Trade Shock Timing Relative to Job Switch



Notes: This figure reports the estimated coefficients from equations (2.18) and (2.19). Panel (a) plots the path of an individual worker's wage level for a worker who switches in year t into a firm that experiences a 100 log-point positive shock to revenue in year t . Panel (b) plots the path of an individual worker's wage level for a worker who switches in year t into a firm that experiences a shock of the same size in year $t - 1$.

of wages for workers who switch into the firm in year t . Define $\Delta \log w_{i,t+\nu} \equiv \log w_{i,t+\nu} - \log w_{i,t-2}$.

$$\text{Arriving in year of shock: } \Delta \log w_{i,t+\nu} = \beta^\nu z_{j(i,t),t} + \zeta x_{i,t} + \delta^\nu X_{j(i,t),t} + u_{i,t+\nu}, \quad (2.18)$$

$$\text{Arriving year after shock: } \Delta \log w_{i,t+\nu} = \tilde{\beta}^\nu z_{j(i,t),t-1} + \tilde{\zeta} x_{i,t} + \tilde{\delta}^\nu X_{j(i,t),t} + \tilde{u}_{i,t+\nu}, \quad (2.19)$$

where $j(i,t)$ is the firm that worker i is employed at in year t , $z_{j(i,t),t}$ is the firm-level export demand shock, and controls $x_{i,t}$ include gender, years of education, potential experience, its square, and the interaction of potential experience with education, and $X_{j(i,t),t}$ includes firm fixed effects, industry-by-year fixed effects, industry-by-year fixed effects interacted with firm's export share of revenue in $t - 1$, three years of the lagged idiosyncratic shock, and the contemporaneous and three years of lags of the common shock. Hence, (2.18) addresses the question: How does an export demand shock in year t affect the wages of job switchers that arrive at the firm in year t ?

As an additional test, we also estimate (2.19), which is the same regression but excludes the time t shock, and the controls include industry-by-year fixed effects and industry-by-year effects interacted with the export share of revenue in year $t - 2$. That way we can get a consistent estimate of the coefficient β_{-1} , which estimates the effect of an export demand shock in period $t - 1$ on the path of wages for workers who arrive in year t . Controls of lagged shocks include shocks through year $t - 4$.

Figure 2.4 shows the estimates from equations (2.18) and (2.19). Figure 2.4(a) shows the coefficient of the time t shock on the wage growth of workers from $t - 2$ to $t + \nu$. Visual inspection of pre-trends indicates that workers moving into differently shocked firms have similar wage paths prior to switching. Consistent with evidence in Figure 2.3a, wages of new hires are approximately

10% higher at firms facing a 100 log-point shock to predicted revenue relative to workers who arrive at firms with a zero shock. The wage premium for workers hired in the shock year is persistent.

Figure 2.4(b) reports the coefficient on the lagged shock, $z_{j(i,t),t-1}$, on the wage changes of switchers who arrive in period t from a local-projection specification analogous to equation (2.18), omitting the trade shock in period t .²⁵ This figure shows that workers who arrive in firms that had larger export demand shocks in period $t - 1$ receive no wage premium upon arrival in period t and exhibit no subsequent catch-up in wage growth. This provides further evidence that firms do not need to pay new hires higher wages once the firm has already expanded.

Job Stayers In addition to estimating the path of wages for switchers around export demand shocks, we also estimate the wage growth of job stayers. We define a stayer as a worker who was full-time ($\geq 1,400$ hours/year) at the same firm from in years $t - 1$ through $t + \nu$. We estimate a local projection of the change in workers' wages relative to year $t - 1$, with the firm export-demand shock in year t as the independent variable. We regress

$$\log w_{i,t+\nu} - \log w_{i,t-1} = \bar{\beta}^\nu z_{j(i,t),t} + \bar{\delta}^\nu X_{i,t} + \tilde{u}_{i,t+\nu} \quad (2.20)$$

for $\nu \in [-4, 3]$, where $X_{i,t}$ includes firm fixed effects, three lags of the idiosyncratic shock $z_{j(i,t),t}$, the current value and three lags of the common shock, industry-by-year fixed effects, and industry-by-year fixed effects interacted with $t - 1$ export shares. Standard errors are clustered at the firm level.

Figure 2.5 reports the response of stayer wages to an export-demand shock. The peak response to a 100 log-point predicted increase in revenue is the year after the shock, when stayer wages rise by 5.2 log points. As a back-of-the-envelope calculation, the implied inverse labor supply elasticity estimated using the response of stayer wages one year after the shock is $d \log w^{\text{stayers}} / d \log N = (d \log w^{\text{stayers}} / d \log Y) / (d \log N / d \log Y) = 0.052 / 0.22 = 0.236$, so the implied labor supply elasticity is $1 / 0.236 \approx 4.2$. Thus, we replicate the literature's finding that using stayer wages to calculate labor supply elasticities yields a value of around 4, implying much less elastic labor-supply curves than the (effectively) perfectly elastic long-run labor-supply curves we find using switchers in Section 2.3. Since firms' long-run marginal cost of labor is determined by the wages of new hires, the wages of new hires determine markdowns in steady state, not the wages of stayers. Accordingly, using the pass-through of demand shocks to stayer wages will understate firms' labor supply elasticity.


Our estimated response of stayer wages relative to the effect on firm revenue are meaningfully smaller than the response of wages in Garin and Silv rio (2024). This may be because our shocks are comparatively small and have less persistent effects on revenue.

2.4 Profit Puzzle with Oligopsonistic Competition

In this section, we extend the equilibrium analysis in Section 2.2 to allow for an oligopsonistic labor market with strategic behavior. Using our evidence for the recruiting cost parameters from Section 2.3, we characterize the aggregate income shares for labor, capital, and profits as a function of

²⁵We correspondingly include the industry-by-year fixed effects interacted with export shares for period $t - 2$ instead of for year $t - 1$.

Figure 2.5: Response of Stayer Wages to Export Demand Shock



Notes: This figure reports coefficients from equation (2.20). The dependent variable is $\log w_{i,t+\nu} - \log w_{i,t-1}$. The regressor $z_{j(i,t),t}$ is the firm-level idiosyncratic export-demand shock. Sample: full-time stayers ($\geq 1,400$ hours/year) who remain at the same firm from $t-1$ through $t+\nu$. Controls $X_{i,t}$ include three lags of the idiosyncratic shock, the current value and three lags of the common shock, industry-by-year fixed effects, and their interactions with $t-1$ export shares. Standard errors clustered at the firm level. Coefficients are scaled to a 100 log-point shock in predicted revenue; entries are in log points.

parameters and the level of labor market concentration. We show that in a calibration that shuts down the recruiting margin and assumes a labor-supply elasticity of 4, the implied labor share is at least 10 percentage points too low when price markups and capital income are calibrated to realistic levels. We then show that with constant-returns-to-scale recruiting function and a long-run elastic labor supply curve, the model matches a labor share in line with national accounts data in rich countries. Furthermore, we find that moving from an equilibrium with no labor market concentration to an empirically realistic level of concentration (an employment Herfindahl–Hirschman Index (HHI) of 0.1) reduces wages relative to marginal product by only 3.4%, an order of magnitude smaller than what is implied by models with Cournot competition and no recruiting margin (Berger et al., 2022).

Setting We now enrich the model from Section 2.2 by assuming that firms are non-atomistic and internalize that their actions affect labor market aggregates, as well as that workers may encounter their own firm’s vacancies when searching. We restrict to no exogenous separations ($s_0 = 0$). Define $\phi_{j,s,t-1}^n$ as the share of employed workers who were hired in period s by firm j and are still employed by firm j in period $t-1$, define $\phi_{j,t}^n \equiv \frac{N_{j,t}}{\sum_k N_{k,t}}$ as firm j ’s share of employment in period t with $\phi_{k,t-1}^n = \sum_{s=-\infty}^{-1} \phi_{k,s,t-1}^n$, and define $\phi_{j,t}^v = \frac{V_{j,t}}{\sum_k V_{k,t}}$ as firm j ’s share of vacancies posted in period t . If a worker encounters their own firm’s vacancy, the worker chooses between two jobs as if they

were any two jobs.²⁶ The recruiting and cohort-specific separation functions take the form

$$R(\mathcal{V}_{j,t,t}, \mathcal{V}_{-j,t}, V_{j,t}, V_{-j,t}, N_{j,t-1}, N_{-j,t-1}) = g(\theta_t) \left[(1 - \Phi_{t-1}^U) \sum_k \sum_{s=-\infty}^{t-1} \phi_{k,s,t-1}^n \frac{\exp(\gamma \mathcal{V}_{j,t,t})}{\exp(\gamma \mathcal{V}_{j,t,t}) + \exp(\gamma \mathcal{V}_{k,s,t})} + \Phi_{t-1}^U \mathbb{1}\{\mathcal{V}_{j,t,t} > \mathcal{V}_U\} \right] \quad (2.21)$$

$$S(\mathcal{V}_{j,s,t}, \mathcal{V}_{-j,t}, V_{j,t}, V_{-j,t}) = \lambda_{EE} f(\theta_t) \sum_k \phi_{k,t}^v \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})} + (1 - \lambda_{EE} f(\theta_t)) \mathbb{1}\{\mathcal{V}_U > \mathcal{V}_{j,s,t}\} \quad (2.22)$$

where $\mathcal{V}_{-j,t}$ denotes the value of all workers not at firm j , we define $\Phi_{t-1}^U \equiv \frac{U_{t-1}}{U_{t-1} + \lambda_{EE}(1 - U_{t-1})}$, $\theta_t \equiv \frac{\sum_k V_{k,t}}{U_{t-1} + \lambda_{EE}(1 - U_{t-1})}$ denotes labor-market tightness, and $U_{t-1} = 1 - \sum_k N_{k,t-1}$. We assume that firms operate in an extended version of Bertrand competition, where firms solve their problems taking as given the wage and vacancy policies of their competitors.

Firm's Problem The firm's problem differs from the firm's problem in Section 2.2 in two ways. First, the firm knows that its own workers will not quit if the worker encounters their own firm's vacancy. Second, the firm internalizes that their choice of vacancies affects labor market tightness, and that additional vacancies congest the firm's own recruiting, as well as congest the labor market for the firm's current workers. Firms solve (2.23) taking $\{V_{-j,t}\}_{t=0}^\infty$ and $\{w_{-j,t}\}_{t=0}^\infty$ as given.

$$\begin{aligned} & \max_{\{V_{j,t}, K_{j,t}, N_{j,t}, \{N_{j,s,t}\}_{s \leq t}, \{w_{j,s,t}\}_{s \leq t}\}_{t=0}^\infty} \\ & \sum_{t=0}^\infty \beta^t \left[A_{j,t} \left(K_{j,t}^\alpha N_{j,t}^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - \sum_{s \leq t} w_{j,s,t} N_{j,s,t} - r^K K_{j,t} - c_j \left(\frac{V_{j,t}}{N_{j,t-1}} \right)^\chi N_{j,t-1}^\sigma V_{j,t} \right] \end{aligned} \quad (2.23)$$

subject to (2.6), (2.7), (2.8), and (2.9), where the recruiting and separation functions are (2.21) and (2.22), and labor-market clearing condition is $U_t + \sum_k N_{k,t} = 1$, and tightness is $\theta_t = \frac{\sum_k V_{k,t}}{U_{t-1} + \lambda_{EE}(1 - U_{t-1})}$. As before, we will focus on the limiting case where $\beta_f \rightarrow 1$.

Equilibrium A stationary equilibrium is defined as in Section 2.2, except that each firm maximizes profits, taking competitors' vacancy and wage strategies as given.

Markdowns We now characterize recruiting cost-adjusted markdowns for the economy with oligopsony in the labor market. We can set b sufficiently low so that workers never voluntarily quit into unemployment, implying zero unemployment in equilibrium. Since all workers are employed in equilibrium, and under the normalization $\sum_k N_{k,t} = 1$, firm j 's share of employment is equal to its employment level, $\phi_{j,t}^n = N_{j,t}$. We now highlight a set of special cases that allows us to characterize markdowns in closed form.

²⁶This assumption is needed for a closed-form solution in section in equation 2.24, and can be thought of as workers applying for another job internally or applying to a different establishment within the firm.

Proposition 2. *If the unemployment benefit b is sufficiently low, if firms have common recruiting cost shifters $c_j = c \forall j$, and if either $\sigma = 0$ or $A_j = A \forall j$, then the recruiting cost-adjusted markdown takes the form*

$$\mu_j \equiv \frac{w_j + c \left(\frac{V_j}{N_j} \right)^{1+\chi} N_j^\sigma}{MRPL_j} = \frac{1 + (1 + \chi)\mathcal{E}_j}{1 + (1 + \chi)\mathcal{E}_j + \sigma + (1 + \chi) \frac{\phi_j^n}{1 - \phi_j^n}}. \quad (2.24)$$

Proof: See Appendix B.10. Using our estimates from Section 2.3 that σ is functionally equal to 0, the above expression holds even if firms have heterogeneous productivity A_j .

Labor and Profit Shares in Equilibrium Now we characterize the share of income that goes to labor, capital, and profits. We will assume that turnover costs accrue to labor income.²⁷ In equilibrium, the share of income that goes to wage income and turnover costs (which are allocated to wages) as a share of gross output is

$$\text{Gross Labor Share} \equiv \sum_j \frac{w_j N_j + c \left(\frac{V_j}{N_j} \right)^{1+\chi} N_j^\sigma}{P_j Y_j} = \frac{\epsilon - 1}{\epsilon} (1 - \alpha) \frac{1}{\sum_j (\phi_j^n / \mu_j)}.$$

We can define the aggregate wage markdown as $\mathcal{M} = 1 / (\sum_j (\phi_j^n / \mu_j))$, and the Herfindahl-Hirschman index is $\sum_j (\phi_j^n)^2$. As we will discuss momentarily, we will be interested in the labor share of net value added, which subtracts depreciation of capital from GDP. Subtracting out depreciation from the denominator, the net labor share and profit share are

$$\text{Net Labor Share} = \frac{\frac{\epsilon-1}{\epsilon}(1-\alpha)\mathcal{M}}{1 - \frac{\epsilon-1}{\epsilon}\alpha\frac{\delta}{\delta+rK}}, \quad \text{Net Profit Share} = \frac{\frac{1}{\epsilon} + \frac{\epsilon-1}{\epsilon}(1-\alpha)(1-\mathcal{M})}{1 - \frac{\epsilon-1}{\epsilon}\alpha\frac{\delta}{\delta+rK}},$$

where the profit share of gross value added from product market power is $\frac{1}{\epsilon}$, and the profit share of gross value added from wage markdowns is $\frac{\epsilon-1}{\epsilon}(1-\alpha)(1-\mathcal{M})$.

Evidence on Profit and Labor Shares The labor share measure that we focus on is the net labor share of domestic corporate value added, which excludes housing and proprietor's income. Our preferred measure of the labor share excludes housing because housing is not traditionally a factor of production, and focusing on the corporate share allows us to avoid taking a stance on how to allocate proprietor's income, which can be challenging (Elsby et al., 2013). The "net" part of net domestic corporate value added means that depreciation is subtracted from value added, since payments to capital that pay off depreciation are not available for consumption (Koh et al., 2020). Rognlie (2015) shows for advanced economies, the labor share of net value added for the domestic corporate sector is between 70-80%.²⁸ The allocation of non-labor income between capital income and pure profit is usually not measured directly, so we will use additional evidence and the model to infer a capital and profit share of income.

²⁷As before, our results are not sensitive to this choice. For more details, see Appendix B.8.

²⁸For the United States, data on value added of nonfinancial domestic corporate business can be found in Table 1.15 of the National Income and Product Accounts.

Table 2.6: Calibration

Parameter	Description	Value	Rationale
ϵ	Demand elasticity	7	Price markup
r^K	Rental rate of capital	0.125	Return on capital = 0.05
δ	Depreciation rate	0.075	Match depreciation/GDP
α	Cobb–Douglas capital elasticity	0.3	Match $K/Y \approx 2$
η	CRRA parameter	1	Simplicity: $\mathcal{E} = \gamma$
β^w	Workers’ discount factor	0	Simplicity: $\mathcal{E} = \gamma$
γ	Inverse scale parameter of non-wage preferences	4	Match $\varepsilon_R - \varepsilon_S \equiv \mathcal{E} = 4$
s_0	Exogenous separation rate	0	Closed-form markdowns

Calibration and Labor Share We now calibrate our model and compute the labor and profit shares of income under three scenarios: (i) a standard atomistic monopsony case, (ii) our preferred calibration with a constant returns to scale (CRS) recruiting function, and (iii) an extension with labor market concentration.

Table 2.6 presents our choice of parameter values that are constant across our three scenarios. We set the rate of exogenous separations at $s_0 = 0$, let workers be myopic $\beta^w = 0$, and we assume log utility $\eta \rightarrow 1$. We set the inverse scale parameter of workers’ preferences over non-wage amenities $\gamma = 4$ so $\varepsilon_{R,w} - \varepsilon_{S,w} = 4$. We set the product demand elasticity $\epsilon = 7$ to generate a price markup over marginal cost of 16.7%, which is relatively small but consistent with empirical evidence (Kline et al., 2019). We assume an ad hoc risk premium on capital of $r = 0.05$ (Jordà et al., 2019), a risk-free real rate of zero, and a depreciation rate of $\delta = 0.075$. This yields a rental rate of capital $r^K = 0.125$. Together with a Cobb–Douglas capital share $\alpha = 0.3$, the parameters jointly deliver a ratio of capital to output $K/Y \approx 2$, consistent with Rognlie (2018), a capital-depreciation share of GDP of 15%,²⁹ and a gross capital share of income of 25%, consistent with Barkai (2020). As can be seen in the adjusted markdown expression in equation (2.24), all other parameters that govern the search aspect of the model drop out.

Table 2.7 reports the share of income that goes to labor, capital, and profits under three scenarios. In the first column, we return to the assumption of atomistic firms, so $\phi_j^n = 0$ for all firms, but the recruiting margin is shut down ($\sigma = \chi \rightarrow \infty$). In this case, wages are 20% below marginal product. Profits from price markups over marginal cost account for 17% of net value added, and net capital income (excluding depreciation) accounts for 12% of net value added. In total, this leaves a profit share of net value added equal to 31%, while the labor share of net value added is only 57%, significantly below the empirical range of 70-80%.

In the second column, we assume a constant returns to scale recruiting function ($\sigma = 0$) with $\chi = 2$ and only atomistic firms ($\text{HHI} = 0$). Product market profits and net capital shares are the same as in column 1, so the only difference is that the labor share of net value added is 71%, i.e., 14 percentage points higher than in the no recruiting case. As a consequence, the profit share of

²⁹See National Income and Product Accounts (NIPA) Table 1.11.

Table 2.7: Distribution of Income: Standard Monopsony, CRS Recruiting, and Concentration

	Standard Monopsony $\sigma = \chi, \sigma, \chi \rightarrow \infty$	CRS Recruiting $\sigma = 0, \chi = 2$	CRS + Oligopsony HHI = 0.1	National Accounts Data
Recruiting cost-adjusted markdown	0.80	1.00	0.97	
Product Profits/Net Output	0.17	0.17	0.17	
Total Profit/Net Output	0.31	0.17	0.19	
Net Capital Income/Net Output	0.12	0.12	0.12	
Non Labor Income/Net Output	0.43	0.29	0.31	0.20–0.30
Labor Income/Net Output	0.57	0.71	0.69	0.70–0.80

Notes: This table reports the recruiting cost-adjusted markdown and the shares of aggregate net value added: product market profits, total profit, net capital income, non labor income, and labor income. For each model, we assume the labor supply elasticity conditional on vacancies $\mathcal{E} \equiv \varepsilon_{R,w} - \varepsilon_{S,w} = 4$. The three economies are (i) a standard monopsony model with no recruiting margin; (ii) our dynamic monopsony model with a constant returns to scale recruiting function; and (iii) one large firm with employment share $\sqrt{0.1} \approx 0.316$ (yielding labor market HHI = 0.1) while all other firms are small.

net value added is 14 percentage points lower. This case matches national accounts data on labor and non-labor shares of national income in rich countries.

In the third column, we assume there is one large firm that accounts for 31.6% of employment, yielding a labor market HHI of 0.1. We show in Appendix B.11 that for a given HHI, the widest markdown is achieved by concentrating employment in only one firm. Compared to the case with no concentration, the recruiting cost-adjusted markdown is 3.4% lower, lowering the aggregate labor share of net income by 2 percentage points, putting the labor share just outside the empirical range. For labor markets that are less concentrated, as found by Schubert et al. (2025) (median HHI of 0.018 on a 0–1 scale), the labor share of net value added would rise back above 70%.

Could the puzzle be resolved by eliminating price markups? As shown in Column 1 of Table 2.7, profits from pricing power in the product market as a share of net value added (17%) are roughly equal to profits from wage setting power in the labor market ($31\% - 17\% = 14\%$). What happens to the profit share if we shut down product market rents? To calculate this, we recalibrate the capital elasticity to $\alpha = 0.25$ to get a capital to output ratio of 2. Holding other parameters fixed with $\gamma = 4$, $\sigma = \chi$, $\sigma, \chi \rightarrow \infty$ (no recruiting margin), and letting $\epsilon \rightarrow \infty$ (no price markups), we obtain a net labor share of 71%. Thus, under a labor supply elasticity of 4 and no recruiting margin, matching the observed labor share of net output requires eliminating price markups.

Misallocation We now consider misallocation. Because large firms strategically under-hire, there may be dispersion in marginal products. Using our result on markdowns from equation (2.24), and using the calibration with $\sigma = 0$ and $\chi = 2$, we numerically assess the output losses from misallocation due to large firms exercising power in the labor market. We reverse engineer firm-specific demand shifters that ensure that the large firm employs a share 31.6% of employment in the market, so the labor market HHI is 0.1. We then compute the steady-state local market output

$Y_m = \left(\sum_j Y_{j,m}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$ and the capital consumption $\delta K_m = \delta \sum_j K_{j,m}$. We then reallocate labor within the local labor market to equalize marginal products and calculate the reallocated output Y_m^* as well as capital consumption δK_m^* given re-optimized choices for capital. The net output gain from reallocating labor is $(Y_m^* - \delta K_m^*) / (Y_m - \delta K_m) = 1.006$, so output increases by 0.6%. At comparable levels of labor market concentration (HHI), our estimated misallocation-induced output losses are about an order of magnitude smaller than those reported by Berger et al. (2022).

2.5 Conclusion

In this article, we provide a model of dynamic monopsony where workers have heterogeneous preferences over firms and can search on the job, and firms can attract workers with higher wages and recruiting expenditures. We use this model to analytically decompose the share of marginal product into wages, recruiting costs, and profits. We provide empirical evidence consistent with constant returns to scale recruiting functions, implying that firms' long-run labor supply curves are very elastic, and that the rents that firms get from exploiting their wage setting power are spent on recruiting. We show that, in equilibrium, elastic labor supply curves help match aggregate profit, capital, and labor shares of income if firms also have pricing power in the product market, even taking into account labor market concentration and strategic interactions. Our setting tractably combines numerous sources of monopsony, including preference heterogeneity, search frictions, and labor market concentration, in a way that is easily disciplined by new and existing empirical evidence and nests many existing monopsony models.

The lack of pure wage dispersion may be useful for researchers who wish to model a labor market with wage posting and on-the-job search but do not want the added complexity of an equilibrium wage distribution. For example, Bloesch et al. (2024) use a similar setting in the labor block of a New Keynesian model to study how wages evolve in response to cost-of-living shocks if firms post wages and workers search on the job. In general, our model may be useful for tractably modeling other search environments such as spending on advertising, where buyers or sellers remain matched while searching for a new match.

Chapter 3

The Incidence and Efficiency of Land Value Taxation

with Ulf Nielsson and Caleb Wroblewski

We thank Alan Auerbach, Sarah Baker, Antoine Levy, Ethan McClure, Mathilde Muñoz, Emi Nakamura, Jón Steinsson, Emmanuel Saez, Benjamin Schoefer, and Danny Yagan, as well as seminar and conference participants at U.C. Berkeley, the Urban Economics Association, the 2024 DAEiNA Meeting, and the National Tax Association for valuable comments and suggestions. We gratefully acknowledge financial support from the Opportunity Lab Place-Based Policy Initiative and the Fisher Center for Real Estate and Urban Economics.

We do need to provide for certain essential government functions... So the question is, which are the least bad taxes? In my opinion the least bad tax is the property tax on the unimproved value of land, the Henry George argument of many, many years ago.

Friedman (1978)

3.1 Introduction

Governments need tax revenue to provide public services, but commodity taxes induce costly economic distortions. It is well understood that taxes on inelastic commodities reduce the deadweight loss from taxation because distortions in *relative prices* cause smaller changes in *relative quantities*. As a result, economists and other social scientists have long recognized that taxes on the value of land can be uniquely efficient. Because the supply of land is thought to be fixed, taxes on land should not reduce the quantity of land; thus, there should be zero efficiency cost of taxation.¹

Taxes on land also have a unique place in intellectual history, having been a subject of active debate by philosophers and social scientists for centuries. The Physiocrats discussed taxes on land rents before modern economics emerged as a discipline. The efficiency arguments for land taxes were well understood by Smith (1776), and the desirability of land taxes has been the subject of discussion by Marx, Hayek, Friedman, Vickrey, and many others. Land taxes today are most prominently associated with George (1880), who argued for a single land tax in lieu of other commodity taxes, and whose work inspired political movements in both the United States and around the world.² Though they are less common in practice than traditional property taxes, land taxes are also the subject of an active policy debate. Detroit is considering a proposal to transition the property tax to a land tax, and California and Colorado are considering or studying similar proposals.

Despite this historical interest and contemporary policy debate, we have relatively little modern empirical evidence about the effects of land taxes. In this paper, we use a natural experiment in Denmark to study individual and aggregate responses to land taxes. We exploit a municipal reform that forced nearby municipalities with distinct land tax rates to set new, unified rates. These new rates varied across municipalities but were constrained by a formula created by the central government. As a result, bordering municipalities experienced very different changes in their land tax rate. Based on this formula, we construct an instrument for the change in the tax rate experienced by these neighboring municipalities using historical data from a decade before the reform. Importantly, because these bordering municipalities were constrained to set unified rates going forward, the differential changes experienced by these bordering municipalities as a result of the reform were permanent. Because this reform happened almost two decades ago, we can investigate both the short- and long-term responses to tax changes. We compare market and aggregate outcomes across these bordering municipalities to identify the effects of land taxes using an instrumental variables (IV) difference-in-differences strategy.

¹Though recent efforts to reclaim land in places like the Netherlands and Denmark suggest the supply curve is not perfectly inelastic over the long term.

²Denmark, which is the setting for this paper, has a political party centered on the work of Henry George, which held seats in Parliament as recently as the 1950s.

We first build a neoclassical model with land taxes to guide our empirical analysis. Our model highlights a critical assumption which drives how individuals respond to land taxes: how the government spends tax revenue. When tax revenue is fully refunded to incumbent homeowners, a land tax does not affect any outcome except the asset price of land. A decrease in the purchase price for land offsets land tax payments. The full response to the tax comes through price changes, and the incidence falls entirely on the incumbent landowner who is then fully compensated by the government. In contrast, if households are not refunded their tax payments, behavioral responses emerge due to income effects. Because the household is poorer due to unrebated tax payments, they demand less of all normal goods, including leisure, buildings, and non-durable consumption.

Guided by the theoretical model, we focus on two main sets of empirical results. First, we study the effects of land taxes on the prices of residential property. The model states that taxes must be capitalized into property prices. Because quantities cannot adjust to clear the market, prices must change instead. We use two sources of variation to study price capitalization. First, we compare homes in nearby bordering municipalities that experienced distinct changes in land value tax rates due to municipal reform (which we refer to as the “policy shock”). Second, we compare homes that are differentially exposed to changes in land tax rates because they have different levels of land intensity (which we refer to as “exposure”). While land taxes are levied on land value alone, we only observe prices for the entire housing unit, including land and structures. However, because different housing units have different land intensities, they should experience different price changes in response to the same policy shock. For example, for an empty plot of land, 100% of the value of the property comes from land, and thus, the land tax will be equivalent to a tax on the value of the entire property. In contrast, in the case where the land and the structure are owned separately, such as a ground lease, the value of the structure will be completely distinct from the value of the land it sits on. Thus, the land tax should have no effect on the structure price. While there are relatively few of these extreme cases in the data, there is significant variation in the land intensity of different housing units. Because land intensity depends on endogenous development decisions, we instrument for exposure using historical land intensity at the plot level.

In stark contrast to the predictions of the theoretical model, we estimate precise null effects of land taxes on housing prices. We use a repeat-sales index, where we compare the transacted prices of the same home over time, before and after the tax change. In our baseline, IV specification, our point estimate is very close to zero and statistically insignificant. The degree to which taxes should affect housing prices depends on the interest rate households use to discount housing cash flows (or, alternatively, the rent-price ratio). In public finance, it is standard to assume a discount rate between 3 - 5%. Recent work has also used institutional features of housing markets to estimate discount rates: Giglio et al. (2014) and Bäcker-Peral et al. (2024) estimate discount rates of 1.9% and 2.8 - 5.3% respectively in the U.K. housing market. He et al. (2024) estimate housing discount rates of 2.4%, while Koster and Pinchbeck (2022) use variation in property taxes to estimate discount rates between 3-4%. Finally, in Denmark, the setting of our study, Jordá et al. (2019), estimates that rental yields in the most recent period are approximately 3.4%. Our baseline estimates rule out all of these estimates with a high degree of confidence. The lower bound of our 95% confidence interval rules out discount rates below 8%.

Our findings have significant implications for the relative desirability of taxing land as opposed to other commodities. A key feature underlying the efficiency properties of land taxes is that the land tax burden is unaffected by behavioral responses. If taxes are fully capitalized into land prices,

landowners cannot avoid land tax increases by selling their land: future purchasers of the land are fully compensated for their higher tax burden with a lower purchase price today. In contrast, when land taxes are not capitalized into prices, as in our setting, incumbent homeowners can avoid taxes by selling to new purchasers, which affects the allocation and efficiency of housing. Because we find land taxes are not capitalized into home prices, our results suggest that taxes on land may be less efficient than in theoretical models. Similarly, when taxes are not capitalized into land prices, the tax burden does not entirely fall on incumbent landowners. Instead, a lack of capitalization implies that the economic incidence of the tax is equivalent to the statutory incidence. This means that the tax burden is shared between current landowners and future purchasers.

Beyond the focus on land taxes directly, our setting is ideal for studying the capitalization of taxes into housing more broadly. First, price capitalization depends not just on contemporaneous prices but rather on expectations of future tax changes. Thus, naive regressions of prices on taxes will be biased downwards if tax changes are *expected* because prices will have already adjusted to the tax change. Similarly, if the market expects the tax change to be transitory, it will have smaller effects on prices, which can be interpreted as a lack of capitalization. We study a single unexpected reform that generated heterogeneous tax changes across space. Thus, we can carefully follow the narrative record and test whether prices adjusted during periods before the statutory tax rate changed when the reform was announced and debated. We find no response in either the period where the reform was being discussed or after the tax rate changed. In addition, because of the nature of the municipal reform we study, the quasi-experimental changes in tax rates we used were permanent. The effect of the policy shock instrument on effective tax rates is identical for each year in our sample period from 2007, when tax rates changed, through 2019.

Second, if households value the municipal services purchased with higher tax revenue, the effect of taxes on prices will be confounded by the effect of changes in municipal service levels on prices. Two aspects of our setting allay this concern. First, in our setting, land taxes are a small share of the municipal budget, and the reform we study was explicitly formulated not to change service levels. Second, in a robustness check, we use only within-area variation from differential exposure to land tax changes. Thus, we compare price changes among homes that are entitled to the exact same level of municipal services but differentially affected by land taxes because they have different land shares. Though less precisely estimated, we once again estimate zero effects that imply less capitalization than the baseline estimates, suggesting municipal service changes do not confound our estimates. Finally, the degree of price capitalization depends on the slope of the housing supply curve. As a result, disagreements about the extent of price capitalization may arise if there are imprecise estimates of the housing supply elasticity. Because land is more inelastically supplied than housing, uncertainty about the housing supply elasticity is less relevant for interpreting price capitalization.

Theoretically, land taxes should generate few behavioral responses. In the model, land taxes generate no substitution effects but can generate income effects if households are not compensated for the taxes that they pay. In our setting, however, behavioral responses may be more relevant because the lack of price capitalization implies households can avoid taxes by moving or selling their property. In addition, the behavioral responses to land taxes have been a key part of the ongoing policy debate about the desirability of land taxes. In particular, land taxes are viewed as a potential substitute for traditional property taxes because land taxes, unlike property taxes, do not disincentivize increased development. Some public commentators in the Georgist tradition even

argue that higher land taxes alone can increase development if property owners need to generate additional cash flow to cover higher tax payments or by reducing the purchase price of land, which might improve land allocation in the presence of financial frictions.

We use the variation caused by the policy shock to study whether land taxes have long-run, aggregate effects on property development, mobility, and housing allocation. Because these variables are slow-moving, the long-term, persistent nature of our shock is particularly valuable. We estimate precise null effects on aggregate residential development. Our estimates imply, consistent with the theoretical benchmark, that land taxes do not disincentivize property development. This means that substituting land taxes for property taxes might be an effective way to increase aggregate density, though the quantitative magnitude depends on the slope of the housing supply curve. We also focus on the development of empty plots, as much of the public discussion of land taxes has focused on whether the shift to land taxes might increase the development of empty plots. Again, we find null effects on the rate of development of undeveloped land plots. Overall, we find that higher land taxes do not affect aggregate development. This suggests that the move to land taxes and away from more distortionary taxes may encourage property development. However, we also provide evidence against the optimistic viewpoints expressed by George (1880) and public commentators that higher land taxes alone might incentivize development.

Finally, we study the effect of higher taxes on housing allocation. Theoretically, land taxes should not affect homeownership unless incumbent landowners respond to income losses from higher taxes by selling their homes. When taxes are not capitalized into land prices, however, homeowners can move to avoid higher tax payments. In addition, a perennial concern about property taxes is that liquidity-constrained homeowners may have to move if they cannot meet their annual tax obligations. Publicly, much of this fear is especially focused on elderly retirees that are not in the labor force. We estimate precise zero effects of land taxes on the probability that homeowners sell their existing property. However, we find that land taxes change the composition of homeowners in an area. In particular, we find that areas that experienced tax increases had fewer older homeowners in the years after the tax change. Taken together, we interpret this evidence as suggesting that higher land taxes do not push incumbent homeowners out of their homes but that older individuals are less likely to move to high tax areas.

In the final part of the paper, we discuss potential drivers of the deviations we document from the neoclassical benchmark. One potential explanation for our result is that landlords pass through increases in land taxes to tenants. While we do not have direct evidence on rent increases during this period, pass-through of *property taxes* to rental prices has been documented in a variety of different settings (Löffler and Siegloch, 2021; Baker, 2024; Watson and Ziv, 2024). In addition, rents in Denmark are regulated and the most common types of rent control specifically allow tax increases to be passed through to rents. We show in an extension of our model that when rents are capped but taxes can be passed through to tenants, the price of land no longer depends on the tax rate. However, this implies that the incidence of the tax falls on tenants, rather than landowners. Because landowners earn higher incomes and have more wealth than tenants and future home purchasers, the tax is significantly more regressive than implied by standard models. Overall, our results suggest that land taxes are less desirable on both efficiency and equity grounds than the neoclassical benchmark. In addition, our work provides evidence in the context of property and capital taxes that demand and supply elasticities are not sufficient statistics for tax incidence (Auerbach, 2006;

Benzarti, 2024). Instead, our work highlights how institutional factors and interactions with other regulations may shape tax incidence.

Altogether, our work sheds light on the real effects of land taxes. We find evidence that contrasts with the stark predictions of neoclassical models. In particular, we find strong evidence that tax increases are not capitalized into housing prices. However, we also see that higher land taxes do not disincentivize development but do affect the allocation of housing within a local area. Overall, we interpret our evidence as suggesting that land taxes have limited efficiency consequences but that the tax incidence does not fall on incumbent landowners. Instead, tax payments fall on tenants and are shared between incumbent and future homeowners.

Related Literature

Our project is related to two main literatures. First, our project contributes to the literature on property taxation (Oates, 1969; Rosen, 1982; Löffler and Siegloch, 2021; Koster and Pinchbeck, 2022; Brockmeyer et al., 2024; Wong, 2023; Baker, 2024; Horton, 2024; Kopplin, 2024; Wong, 2024) and land taxation (Bonnet et al., 2021; Schwerhoff et al., 2022). We add to this literature by studying a unique natural experiment that led to persistent changes in land tax rates across space and time to study the effects of tax rates over the long run on prices and aggregate outcomes. In contemporaneous work, Murphy and Seegert (2024) create a measure of “implicit” land taxes based on differences between assessed values and market valuations to show that higher implicit taxes are associated with a variety of economic outcomes, including density. Coven et al. (2024) use a theoretical model to argue that property taxes can change housing allocation by increasing homeownership, especially for financially constrained households. In contrast, we use a quasi-experimental variation in statutory tax rates to estimate the causal effect of land taxes, providing new evidence on these important questions of homeownership, density, and housing allocation.

Two papers study the short-term effects of the same reform on home prices (Nielsen and Rzeźnik, 2014; Høj et al., 2018). Relative to both papers, we use a different empirical strategy over a longer time horizon, which results in substantial differences between our estimates. We discuss differences between our estimates and previous work in Section 3.7. We also expand on this previous work by looking at longer-term effects, heterogeneity in the price results, and real effects on aggregate outcomes.

Second, we relate to a long literature on tax incidence. Benzarti (2024) reviews the literature on the incidence of consumption and labor taxes and argues that, in contrast to theoretical models, the demand and supply elasticities are not sufficient statistics for tax incidence. We show that similar results hold in the case of taxes on real property, in this case for land. Benzarti (2024) argues that non-standard factors, like fairness norms, shape the incidence of labor taxes in particular. In the final part of the paper, we argue that institutional regulations related to Danish rental markets may provide an explanation for the limited price capitalization we estimate. As a result, we also show how regulations and institutions can be important drivers of tax incidence.

The rest of the paper is organized as follows. Sections 3.2 and 3.3 lay out our theoretical model and present a series of empirical predictions. Section 3.4 presents institutional context about the reform we study and Danish land taxes. Section 3.5 outlines the data we use, while Section 3.6 describes our instrument and identification strategy. Sections 3.7 present the results of taxes on

home prices. Section 3.8 shows the effect on aggregate development and housing allocation. Section 3.9 discusses potential drivers of the main results and concludes.

3.2 Model

We build a neoclassical model of land-value taxation to guide the empirical analysis and highlight the unique theoretical properties of land taxes.

Households

Households maximize utility derived from nondurable consumption (C_t), housing services (H_t), and leisure over an infinite horizon. Housing services are an aggregate of land and buildings owned by the household ($H_t = H(L_t, B_t)$). Individuals also invest in a real financial asset (A_t) and receive lump-sum transfers from the government (T_t). The government levies a tax (τ^L) that is paid each period based on the price of land. The stock of buildings depreciates at a rate δ . Formally, the household problem can be written as:

$$U_0 = \sum_{t=0}^{\infty} \beta^t u(C_t, N_t, H_t).$$

The household optimizes while respecting a series of flow budget constraints:

$$C_t + P_t^L(L_t - L_{t-1}) + P_t^B(B_t - (1 - \delta)B_{t-1}) + (A_t - A_{t-1}) = W_t N_t + r_{t-1} A_{t-1} - \tau^L P_t^L L_{t-1} + T_t \quad (3.1)$$

for $t = 0, 1, \dots$, and given initial values L_{-1} , B_{-1} , A_{-1} . Consumption is the numeraire. P_t^L and P_t^B are the asset prices of land and buildings, respectively. The first order conditions for labor and consumption are standard. The first order conditions for the land and building holdings of the household are given by the following expression:

$$\frac{u_{L,t}}{u_{C,t}} = P_t^L - \frac{(1 - \tau^L)P_{t+1}^L}{1 + r_t} \quad (3.2)$$

$$\frac{u_{B,t}}{u_{C,t}} = P_t^B - \frac{(1 - \delta)P_{t+1}^B}{1 + r_t}. \quad (3.3)$$

Equation 3.2 shows that the household will purchase land until the marginal utility of land consumption (left-hand side) equals the **user cost of land** (right-hand side). Because land is a durable good, the allocative price of land is not the asset price (P_L), but rather the cost of holding a unit of land for one period and then selling it. This cost is increasing in the land tax rate and in the interest rate (which represents the opportunity cost of investing in financial assets relative to land) and is decreasing in any capital gains.

Suppose that $r_t = r$ is constant (for expositional purposes). Then, based on the first order condition of the household, we can solve for the asset price of land. Denote the marginal rate of substitution for land: $MRS_t^L = \frac{u_{L,t}}{u_{C,t}}$. Then, we have that the asset price of land P_t^L is given by

$P_t^L = MRS_t^L + \frac{(1-\tau^L)}{1+r} P_{t+1}^L$. Solving the equation forward results in:

$$P_t^L = \sum_{s=0}^{\infty} \left(\frac{1-\tau^L}{1+r} \right)^s MRS_{t+s}^L. \quad (3.4)$$

Thus, the asset price of land is the sum of future expected marginal rates of substitution, discounted by the interest rate and the land value tax rate. Note that this is a standard asset pricing condition, where the price is the discounted sum of future cash flows. In a model with renters, this would be equivalent to the price of land equalling the discounted sum of future rents generated by the land.

Production

Consumption goods are produced by a firm that hires labor and produces according to the following production technology:

$$Y_t = F(N_{Y,t}). \quad (3.5)$$

Similarly, there is a construction firm that produces new buildings I_t by hiring labor and producing using the following production function:

$$I_t = G(N_{I,t}). \quad (3.6)$$

Both firms are price takers in both input and output markets and maximize profits with the following equations for the consumption ($\Pi_t^Y = Y_t - W_t N_{Y,t}$) and construction firms, ($\Pi_t^I = P_t^B I_t - W_t N_{I,t}$). Their first order conditions for production are given by:

$$F_N(N_{Y,t}) = W_t \quad (3.7)$$

and

$$P_t^B G_N(N_{I,t}) = W_t. \quad (3.8)$$

Government

The government collects tax revenue $\tau^L P_t^L L_{t-1}$ at the beginning of each period and pays lump-sum transfers T_t to households. Many of the model's key results are driven by how the government spends its collected tax revenue. As a result, we consider two cases. In the first case, the government fully rebates all the tax revenue back to landowners as a lump-sum transfer: $T_t = \tau^L P_t^L L_{t-1}$. We refer to this as the "full compensation" case. In the second case, the government does not return tax revenue to the landowners. To nest both of these cases, we write the lump-sum transfer as:

$$T_t = \tau^L P_t^L L_{t-1} \times \mathbb{I}(\text{Comp.}) \quad (3.9)$$

where $\mathbb{I}(\text{Comp.}) = 1$ if the government compensates landowners and zero otherwise.

Market Clearing

We impose four market-clearing conditions:

$$N_t = N_{Y,t} + N_{I,t} \quad (3.10)$$

$$A_t = 0 \quad (3.11)$$

$$L_t = \bar{L} \quad (3.12)$$

$$B_t = I_t + (1 - \delta)B_{t-1}. \quad (3.13)$$

Equation 3.10 is the labor-market clearing condition. Labor supply equals the combined demand for labor from the consumption and construction firms. Equation 3.11 is an asset-market clearing condition, which states that the financial asset is in zero net supply. Equation 3.12 is the land-market clearing condition. Our key assumption is that the supply of land is inelastic and endowed by nature. Thus, demand for land from the household has to equal $\bar{L}$, the fixed supply of land. Equation 3.13 states that building demand equals the supply of new construction produced by the construction firm, plus the non-depreciated stock of buildings from the previous period. The market for non-durable goods clears by Walras' law.

Equilibrium

Given an initial allocation for (L_{-1}, B_{-1}, A_{-1}) , an equilibrium is prices (P_t^L, P_t^B, W_t, r_t) and quantities $(C_t, Y_t, N_t, N_{Y,t}, N_{I,t}, L_t, B_t, I_t, A_t, T_t)$ such that the household optimizes, the buildings and consumption firms optimize, the government collects taxes and rebates transfers, and markets clear.

Special Case

To highlight the analytical predictions of the model, we make a series of functional-form restrictions. Assume utility is given by:

$$u(C_t, N_t, L_t, B_t) = \log C_t - \vartheta \frac{N_t^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}} + \theta \log L_t + \varphi \log B_t.$$

In addition, assume the production functions for non-durable consumption and buildings are given by $F(N_{Y,t}) = Z_Y N_{Y,t}^\gamma$ and $G(N_{I,t}) = Z_I N_{I,t}^\eta$, respectively.

3.3 Theoretical Predictions

We use the model outlined in Section 3.2 to highlight how land taxes affect equilibrium prices and quantities.

The Effect of Land Taxes on Land Prices

Higher land taxes make holding land less desirable. However, because land is in fixed supply, the equilibrium quantity of land cannot fall as a result of higher taxes. As a result, all of the adjustment


to higher land taxes **must come through changes in the price of land**. In steady state, the asset price of land is given by the following expression:

$$P^L = \frac{1+r}{r+\tau^L} \theta \bar{L}^{-1} C. \quad (3.14)$$

The asset price of land is decreasing in the tax rate τ_L . This effect is highlighted in Figure 3.1, which shows supply and demand in the market for land after a land tax increase. The line labeled D shows the demand for land prior to a tax increase, while D' shows the demand for land after a tax increase. As tax rates go up, higher taxes are *capitalized* in the price of land. That is, higher taxes reduce the price of land such that future landowners are fully compensated for future higher tax payments with a lower purchase price today.

Since the supply curve for land is vertical, the change in the *asset price* of land does not change the quantity of land in equilibrium. As a result, there is no deadweight loss of higher taxes, implying no efficiency cost of land taxes. In fact, when the tax is fully compensated, the tax rate has no affect on any aggregate quantity or price except for the asset price of land.

Figure 3.1: A Model of the Market for Land



Income Effects

When taxes are not fully rebated to households, income effects cause behavioral responses. We focus on the effects of land taxes on labor supply. The following expression gives the steady-state level of labor:

$$N = \vartheta^{-\frac{1}{1+\frac{1}{\chi}}} \left(1 + \frac{1+r}{r+\delta} \varphi \delta + \frac{1+r}{r+\tau^L} \theta \tau^L (1 - \mathbb{I}(\text{Comp.})) \right)^{\frac{1}{1+\frac{1}{\chi}}}. \quad (3.15)$$

Equation 3.15 highlights how the government's redistributive preferences drives the effects of land taxes on equilibrium prices and quantities. Consider the case where the government fully compensates landowners by returning the land tax payments as a lump sum. In that case, the steady-state labor supply is given by:

$$N = \vartheta^{-\frac{1}{1+\frac{1}{\chi}}} \left(1 + \frac{1+r}{r+\delta} \varphi \delta \right)^{\frac{1}{1+\frac{1}{\chi}}}.$$

When land taxes are fully rebated, the tax rate has **no effect on equilibrium labor supply**. Labor supply depends only on the parameters that determine household disutility of labor and household demand for buildings. The expression highlights the tradeoff faced by households in the model: the optimal quantity of labor balances household demand for consumption goods and buildings with leisure. Household preferences over buildings matter for period-by-period labor choices because buildings depreciate over time, and households must cover the cost of maintaining the housing stock. In contrast, consider the case where the government spends the land tax revenue on services that do not enter the household utility function. Then, the labor allocation is given by:

$$N = \vartheta^{-\frac{1}{1+\frac{1}{\chi}}} \left(1 + \frac{1+r}{r+\delta} \varphi \delta + \frac{1+r}{r+\tau^L} \theta \tau^L \right)^{\frac{1}{1+\frac{1}{\chi}}}.$$

When taxes are not remitted back to the household, labor supply is increasing in the tax rate. This is a wealth effect on labor supply. Households' work decisions now **also** reflect the tradeoff between leisure and covering land tax payments. Similar tradeoffs affect the allocations of other normal goods in the model: higher land taxes reduce the demand for consumption of non-durable goods, the demand for housing, and the demand for land. Because the supply of buildings and consumption goods is upward-sloping, this also reduces equilibrium quantities of both goods.

In contrast, because land supply is fixed, the quantity of land is invariant to the tax rate. Because land is a normal good, like leisure or consumption, household demand for land decreases due to wealth effects. However, when taxes are not returned to owners, income effects can reduce land demand, further depressing land prices. This effect is highlighted by Equation 3.14: when there is no income effect, the **user cost of land** does not change in response to a change in taxes $UC^L = \theta \bar{L}^{-1} C$ because C is invariant to the tax rate. In contrast, when taxes have income effects, households reduce C in response to higher taxes, which results in the user cost falling in response to higher tax rates. In practice, in the presence of income effects, prices must fall more to make the household indifferent between land consumption and non-durable consumption, which now has a higher marginal utility. This puts further downward pressure on P_t^L in response to a land tax increase.

Taken together, our model highlights two main effects of land taxes. First, higher tax rates should reduce the price of land (and thus housing) as the inelastic land supply implies that taxes should be fully capitalized into land prices. Second, if taxes are not rebated to landowners, or if the government spends tax revenue on services that do not enter the household utility function, there will be behavioral responses due to income effects. In particular, households should reduce their consumption of normal goods: buildings, non-durable goods, and leisure. In addition, income effects can also reduce demand for land, which will put further downward pressure on land prices.

3.4 Institutional Details

We study a municipal reform in Denmark that generated permanent, quasi-experimental variation across towns in land tax rates.

Municipal Reform

Before 2007, Denmark had three primary levels of government: the central government, the counties, and municipalities. Municipalities are the smallest level of government and are responsible for a variety of public services. Most prominently, the municipalities provide primary schooling and elder care, though they also have some responsibility for infrastructure and other social and child care services. At the time of the municipal reform, the thirteen counties were primarily responsible for healthcare services.

In 2002, discussion of a potential reform to consolidate the subnational governments into larger entities began. Both Christiansen and Klitgaard (2008) and Christiansen (2012) date the beginning of the discussion of the reform to the Summer of 2002, when newspaper articles indicated support for reform among both Denmark's largest business organization and younger members of a governing party at the time. The primary motivations for the reform were to improve public sector service provision through economies of scale, especially in the hospital sector, where it was thought that the counties were too small to be effective. Crucially, the goal of the reform was not to change **the level** of municipal services. Before the increased discussion of the reform, a municipal reform of this type was considered extremely unlikely. Christiansen and Klitgaard (2008) write that:

"There is no shortage of arguments that, before it is actually implemented, the structural reform must indeed be considered almost unthinkable...The unthinkable, as is well known, happened."

Appendix Figure C.1 displays the time series of references to structural reform in major Danish newspapers. Before 2002, there were very few mentions of structural reforms. This changed in the summer of 2002 when the number spiked. Structural reform was also not part of the debate during the general election in 2001 (Bundgaard and Vrangbæk, 2007).

In August of 2002, the Prime Minister announced that the government intended to appoint a commission to study the administrative structure of the Danish state (Bundgaard and Vrangbæk, 2007). In January 2004, the commission presented its report. The report included proposals similar to the municipal reform that was eventually adopted, including proposals eliminating the counties and requiring larger municipalities. Between April and June 2004, the government presented a proposed reform, which was debated and negotiated between parties. In June 2004, a law implementing the structural reform was passed.

The structural reform required the creation of larger municipalities, which were constructed through municipal mergers. The counties were eliminated and replaced by five regions, primarily responsible for the healthcare sector but also for some tasks that would benefit from less decentralization (such as regional development and critical infrastructure). The larger municipalities retained most of their responsibilities but took over some that were previously the purview of the counties.

Once the reform passed, individual municipalities had leeway in determining which other municipalities to merge with. The government required a baseline population level for new municipalities

and that merging municipalities bordered each other. Even before the agreement by the central government, several municipalities had announced their merger plans, and most of the remaining municipalities were very close to merger agreements (Lassen and Serritzlew, 2011). By 2005, the mergers were finalized, and local residents had voted for representatives for the new merging municipalities. The reform took effect in January of 2007.

The Reform Changes Land Tax Rates

Denmark has a land tax set at the municipal level.³ The land tax is one relatively small source of financing for municipal spending. In practice, most local tax revenue comes from income taxes. Municipalities also receive significant financing from the central government through grants and an equalization scheme. Before the reform, the overall land taxes paid by households ranged from around 1.6% to 3.4% of the assessed value of land annually. Of this, 1% was used to fund services provided by the county, while the rest was used to fund municipal services. After the reform, the municipalities kept all of the tax revenue generated by the land tax.

The municipal reform generated substantial variation in land taxes because merging municipalities that previously set distinct land tax rates were now required to set a unified one. The central government also constrained how municipalities could set this new unified tax rate. The government set an upper bound for the new tax rate that was a weighted average of the existing land tax rates in the merging areas (Høj et al., 2018). In practice, most of the new merging municipalities set their new tax rate very close to the ceiling provided by the government (Høj et al., 2018).

Land assessments are the responsibility of the central government and are based on sales of unimproved land and hedonic regressions. The amount that land tax assessments can increase each year is capped. From 2011 until the 2020s, land assessments remained fixed as the government developed a new system for land valuations. However, **taxable** land values continued to increase because the cap had resulted in taxable valuations being well below market values after the increase in home prices in the early 2000s (Klein et al., 2016). There are some exemptions from land taxes, such as for private schools and non-profits. In addition, agriculture is entitled to lower land tax rates and benefits from low property valuations.

There are three primary benefits to using the variation in tax rates induced by the municipal reform. First, by restricting our comparison to nearby municipalities that merged, we can isolate a permanent change in tax rates. Consider two hypothetical towns that merged into a single municipality, in which one town had previously set a 1% tax rate, and the other had previously set a 2% tax rate. The government constrained the newly merged municipality to set their tax rate at 1.5%. Households in one town experienced a 50 basis point (b.p.) tax increase, while the other experienced a 50 b.p. decrease. The 100 b.p. difference between the two towns is held fixed **even if the newly merged municipality chooses to change its tax rate later** because the towns are constrained by the reform to set a unitary tax rate going forward. Suppose that later on, the new municipality chooses to increase its tax rate to 2%. Now, one town is at its old tax rate, while the other has experienced a 100 b.p. increase, but the difference remains at 100 b.p. Because price and behavioral responses to capital tax changes depend not just on contemporaneous policy but also on expectations of future policy, our permanent variation ensures our results are not driven by

³There is also a national-level property tax, which was not affected by the reform.

expectations of future policy changes, that may have occurred subsequent to the reform, and allows us to interpret our effects as the effect of a permanent tax change.

Second, the ceiling set by the central government generated formula-based changes in tax rates. Naive regressions of the effect of property tax changes may be confounded if local economic conditions or financing needs drive both the stance of policy and behavioral responses. We create an instrument for the tax change experienced by an individual municipality based on historical data and the central government formula. This allows us to isolate variation driven by the municipal reform and ensures our results are not confounded by municipal expenditure needs or other local changes in economic conditions or demographics. Third, because asset prices are forward-looking, behavioral responses to tax changes should reflect expectations of future policy. Because we focus on a single reform with a detailed narrative record that generated substantial variation in tax rates, we can carefully test whether expectations of future policy affect asset prices and individual behavior. Finally, municipal reform generated large tax changes in towns that bordered each other. Our empirical strategy studies the effect of land taxes by comparing nearby towns that experienced these differing changes in land tax rates. These nearby towns serve as a natural control group that allows us to carefully control for other cyclical and regional factors that might influence aggregate and individual outcomes.

3.5 Data

Our analysis combines a variety of Danish administrative registers covering land use and the housing market, as well as administrative data from tax records on individual households.

Housing Data

We combine a series of data on residential housing in Denmark. We start with data on land and property taxation, which provides information for each property on the assessed value of land and property taxes owed in each year. We calculate effective tax rates for each property by dividing taxes owed in each year by the property's assessed value. We combine this with property-level data from the Danish Housing Census, which contains comprehensive property-level characteristics. We merge this with data on property sales, which provide us with information on the sale date and price of sold properties. To create a consistent panel of properties before and after the reform, we rely on a key that links properties over time. We sum individual property level outcomes, including property development and residential square footage, to the treatment area level to study aggregate effects of the policy. When studying the effect of taxes on prices, we focus on residential properties and drop vacation homes.

Individual Data

We create an annual panel of individual-level income and wealth using administrative tax data. The data contain detailed information about the amount and composition of income and wealth holdings. Nearly all components of income and wealth are third-party reported (Jakobsen et al., 2024). We use a property ownership registry to link individuals to their properties, which we use to

assign treatment. We use the individual data to study whether land taxes change the number and composition of homeowners along observable dimensions like wealth, income, and age.

Municipal Data

We collect annual, municipality-level data from Statistics Denmark. The data include demographic and economic variables and information about local government services and spending. We use the statutory land tax rates to construct our instrument and define treatment. We also include pre-treatment data on demographic characteristics as controls. Note that because our treatment areas are smaller than municipalities after the reform, we cannot use *municipal-level* variables after 2007.

3.6 Instrument and First Stage

Our empirical strategy regresses various property-level and aggregate outcomes on tax rates. However, because the change in the tax rate was formula-based, transitory changes in tax policy or population before the reform may have mechanically affected the area's post-reform tax rate. Thus, a potential concern with a simple OLS regression of outcomes on tax rate changes is biased by the effects of other changes immediately before the reform. To account for this concern, we construct an instrument for the change in the tax rate based on historical data more than a decade before the reform. As a result, our empirical results are only based on variation from historical tax rates and populations, and not from changes that occur close to the reform.

Our instrument is based on the formula the central government used to constrain the new land value tax rates set by the merged municipalities. The central government set a ceiling for the new tax rate for each new municipality that was a weighted average of the land tax rates of the existing, merging municipalities (Høj et al., 2018). We construct a predicted value for the ceiling set by the central government for each new municipality (M), using the following formula:

$$\tilde{T}_M = \sum_{i \subset M}^M T_{i,1995} \frac{p_{i,1995}}{p_{M,1995}}.$$

Here, $\tilde{T}_M$ represents the predicted tax rate ceiling for new municipality M that is imposed by the central government. It is a weighted average of the tax rates in the existing municipalities i that would merge into municipality M , where the weights are municipality i 's share of the total population of M in 1995. We then define the instrument at the treatment area (a) level, which is a cross of each old municipality i with a new municipality M .⁴ For each treatment area, we define the instrument as the difference between the area's 1995 tax rate and the predicted tax rate ($\tilde{T}_M$), according to the following formula:

$$\Delta \tilde{\tau}_a = T_{i,1995} - \tilde{T}_M$$

Thus, the instrument represents the predicted change in land tax rate experienced by treatment area a , based on historical data from 1995. Figure 3.2 tests whether the first-stage relationship is

⁴Note that the treatment areas are not exactly equal to old municipalities i because some old municipalities split apart during the reform.


strong. Figure 3.2 plots the following regression specification:

$$\Delta\tau_a = \beta\Delta\tilde{\tau}_a + \eta_{m(a)} + \delta'X_a + e_a.$$

In this specification, $\Delta\tau_a$ is the actual change in the statutory land tax rate experienced by treatment area a between 2007 and 2005. The primary coefficient of interest is the tax change instrument ($\Delta\tilde{\tau}_a$). The regression also includes controls at the treatment area level, as well as new municipality fixed effects. Thus, the regression compares bordering towns that merged together, and as a result of the merger experienced heterogeneous changes in land tax rates, mimicking the variation used in the paper.

Figure 3.2 shows that the first stage is strong: the points are closely bunched around the line of best fit, and the regression coefficient is highly statistically significant. Overall, we find evidence that our instrument is highly predictive of changes in statutory land tax rates.

Figure 3.2: First Stage: The Effect of Tax Change Instrument on Statutory Tax Rate



Notes: This figure plots the first stage regression of the land tax instrument on the actual change in statutory tax rates from 2005 to 2007. For both the x- and y-axes, the units are in percentage points: thus, a value of 1 represents a 1 p.p. change in the land value tax rate. The blue line represents the 45-degree line. Beta coefficients represent and 95% confidence intervals, with robust standard errors are plotted.


Figure 3.3 shows that this change in statutory tax rates led to changes in effective tax rates at the property level. We use the following event study specification:

$$y_{j,a,t} = \sum_{h \neq -1} \beta_h \Delta \tilde{\tau}_a \mathbf{1}_{\{t-h=2005\}} + \gamma_j + \eta_{m(a) \times t} + e_{j,i,t}. \quad (3.16)$$

In this specification, j represents a property, which is assigned to a treatment area a . The outcome variable is the property-level effective tax rate, defined as the taxes owed in year t divided by the property's assessed land value. We plot β_h coefficients, representing the interaction between the treatment instrument and year fixed effects. The regression includes a property fixed effect (γ_j) and new municipality by year fixed effects.⁵

Figure 3.3 plots the β_h coefficients. For the decade before the reform, trends in land tax rates are relatively flat. In 2007, the year the reform took effect, there is a strong increase in effective tax rates for properties in treatment areas where statutory tax rates increased. The magnitude of the effect is similar to the effect on statutory tax rates, with a 1 p.p. predicted increase in the tax rate corresponding to an actual increase in effective tax rates of around 80 basis points.

Figure 3.3: The Effect of Statutory Tax Changes on Property-Level Effective Tax Rates



Notes: This figure plots the first stage regression of the land tax instrument on effective tax rates at the property level. 95% confidence intervals are shown, with standard errors clustered at the treatment area level. We estimate the regression without controls for computational tractability.

⁵For computational tractability, we omit the controls interacted with year fixed effects.

Crucially, this effect is persistent throughout the entire sample period. The point estimates and standard errors in the years after 2007 are nearly identical to those in 2007. This highlights that our empirical strategy successfully isolates changes in land taxes that have a strong first stage and are very persistent. This provides a clean interpretation of our price and behavioral results as reflecting the effect of persistent tax changes.

To study the effect on price capitalization, we combine this policy-induced variation with variation in exposure to the land tax change. While all residential properties in a treatment area face the same effective tax rate, the extent to which this passes through to prices depends on the share of the property's value derived from land. Consider a simplified Gordon growth model, with a land tax, that is only levied on the share of the value of the property that is assessed to derive from the land. Then, the effect of land taxes on the housing prices is equal to:

$$\frac{\partial V}{\partial \tau} = -\frac{\theta}{r}$$

where τ is the land tax rate and θ is the land share of the property. Thus, property's with higher land shares, will have larger price changes in response to the same change in land tax rates. This is the second source of variation we use to study tax capitalization.

A property's land share is an endogenous variable. Thus, we instrument for the land share with the property's land share in 1996, a decade before the reform. We run the following regression:

$$y_{j,a} = \beta \Delta \tilde{\tau}_a \theta_j + \zeta_a + \psi_j + e_a. \quad (3.17)$$

In this specification, $y_{j,a}$ represents total land taxes owed, normalized by the total assessed value of the property, at the property level. In this specification the main coefficient of interest is the policy shock $\Delta \tilde{\tau}_a$ interacted with the properties pre-existing land share ω_j . ζ_a is a treatment area fixed effect. Thus, all of the variation in land tax exposure comes from properties that are in the same treatment area, exposed to the same tax rate and municipal services, but with differential exposure to land taxes. ψ_j is a land-share decile fixed effect.


Figure 3.4 shows the binscatter plot at the property level of total land taxes paid normalized by the total assessed value of the property. Figure 3.4 shows that there is substantial variation, even within the same treatment area, in exposure to land tax changes. The coefficient is close to 0.5 and highly statistically significant.

Overall, we find substantial variation, both from the policy shock and exposure to the land tax changes, with which to study the effect of land taxes on residential property prices. Figure 3.5 shows the instrument combining both the policy shock variation and the exposure variation, mimicking the combined instrument from our baseline price capitalization specification. In this case, the regression takes the following form:

$$y_{j,a} = \beta \Delta \tilde{\tau}_a \theta_j + \eta_{m(a)} + \delta' X_a + \psi_j + e_a.$$

This regression has the same dependent variable and instrument as Equation 3.17. However, here, we replace the treatment area fixed effect with a new municipality fixed effect. As a result, we use both variation from the exposure shares, as well as variation in tax rates from the policy shock. Figure 3.5 shows that, as with the instruments used separately, the first stage is strong, with the instrument being highly predictive of changes in taxes owed normalized by the total valuation. Overall, both the policy shock and exposure shares induced strong changes in land taxes paid as a result of the reform.

Figure 3.4: First Stage: The Effect of Land-Tax Exposure Instrument Within A Treatment Area




Notes: This figure plots the first stage regression of the land tax instrument on the actual change in effective taxes paid, normalized by the assessed value of the property, from 2005 to 2007. For both the x- and y-axes, the units are in percentage points: thus, a value of 1 represents a 1 p.p. change in the tax rate, normalized by the assessed value. The blue line represents the 45-degree line. Beta coefficients and 95% confidence intervals, with standard errors clustered at the treatment area by land share decile level, are plotted.

3.7 Capitalization in Housing Prices

The efficiency of land value taxes rests on the assumption the supply curve for land is vertical. While a vertical supply curve implies that the economy does not adjust to the tax with changes in total quantities of land, it also implies that all adjustment to the tax comes through changes in prices. Without price adjustments, incumbent landowners can shift the tax burden to future purchasers of the land. Similarly, limited price effects may reflect the ability of landowners to avoid the tax by shifting tax payments onto tenants.

We use our empirical setting to investigate the effect of land taxes on prices in the residential housing market. We use a repeat-sales methodology to study the effect of land taxes on the prices of homes. Our outcome variable is the log sale price of transacted homes. We run the following

Figure 3.5: First Stage: The Effect of Land-Tax Exposure Instrument Combining Policy Shock and Exposure Measures



Notes: This figure plots the first stage regression of the land tax instrument on the actual change in effective taxes paid, normalized by the assessed value of the property, from 2005 to 2007. For both the x- and y-axes, the units are in percentage points: thus, a value of 1 represents a 1 p.p. change in the tax rate, normalized by the assessed value. The blue line represents the 45-degree line. Beta coefficients and 95% confidence intervals, with standard errors clustered at the treatment area by land share decile level, are plotted.

regression for our baseline specification:

$$y_{j,a,t} = \sum_{h \neq -1} \beta_h \Delta \tilde{\tau}_a \theta_j \mathbf{1}_{\{t-h=2001\}} + \gamma_j + \eta_{m(a) \times t} + \delta' X_a \mathbf{1}_{\{t-h=2005\}} + \psi_{j \times t} + e_{j,i,t}. \quad (3.18)$$

Our baseline specifications include a property fixed effect and new municipality by year fixed effects. Thus, our estimator identifies the impact of land taxes by comparing houses in neighboring towns with different tax rate changes pre- and post-reform, and by comparing homes with differing exposure to the land tax through different land shares. By including property fixed effects, we hold fixed time-invariant property characteristics and ensure that the results are not driven by the land tax changing the composition of transacted homes across time. We also include $\psi_{j \times t}$ land share decile by year fixed effects, to ensure our results are not affected by time-varying demand for more vs.

less land intensive properties. We normalize the year coefficients relative to 2001, the year before discussion of the reform began.

Figure 3.6 shows that higher land taxes do not affect residential home prices. The coefficients show the impact of a one percentage point increase in the land tax (normalized by the assessed home price) on sale prices. Between 1995 and 2001, prices of homes in areas that faced increases in their land tax rate experienced very similar trends to areas where land tax rates would eventually fall. All of the coefficients are close to zero and statistically insignificant. In 2002, discussion of the reform began. However, between 2002 and 2006, when details of the reform are finalized, there is no stark change in home prices. All coefficients are close to zero, and only one is (marginally) significant. After 2007, when the tax rates actually changed, there is still no change in home prices. Across all post-period years, the coefficients are close to zero and insignificant.


The line in red plots the post-period coefficient, pooling all post-period years together and comparing it to 2001, the last year before the discussion of the reform began. The post-period coefficient lies directly on top of the zero line. Overall, we estimate precise null effects of higher land tax rates on residential property prices. This suggests that, in contrast to the predictions of the theoretical model, land taxes are not capitalized into home prices. This is consistent with the slope of the demand and supply curves not acting as sufficient statistics for tax incidence, as discussed in the context of labor and consumption taxes by Benzarti (2024).

Figure 3.7 shows that the results are consistent using either source of identification separately. Panel (A) plots the beta coefficients, using only the policy shock to instrument for the change in the tax rate. The results look nearly identical to the baseline. All but one of the post-period coefficients are statistically insignificant and close to zero. As a result, the overall post-period coefficient is very close to zero. Panel (B) shows the effect after including treatment area by year fixed effects. Thus, this estimate compares homes that are in the same area and have the same municipal services but have different land intensities. Though much less precisely estimated, the effect is again close to zero, and all post-period coefficients are statistically insignificant. If anything, the point estimate is larger, implying less pass-through than in the baseline specification. This provides suggestive evidence, consistent with our empirical setting, that the estimates are not affected by changes in municipal services.

Finally, Figure 3.8 shows the effect in a sample of single-family homes. Single-family homes are a particularly important sub-sample because they are more land-intensive than most apartments or condominiums and because the vast majority of single-family homes are owner-occupied as opposed to rented. If landlords can pass on cost increases to tenants, this will mute the price pass through. However, because for owner-occupied homes the owner is both the landlord and the tenant, reactions to tax rates may be higher. However, Figure 3.8 shows that the results are remarkably consistent in the single-family sample. Once again, all post-period coefficients are statistically insignificant and close to zero. Across all samples and identification strategies, we find consistent evidence of zero pass through of taxes to home prices.

Because the right-hand side variable in our event study is the instrument for the change in the tax rate, Figures 3.6, 3.7, and 3.8 show the intent-to-treat effect of land taxes. To quantify the magnitude of the tax effect, we instrument directly for the property-level change in taxes owed divided by assessed land value. The results are shown in Table 3.1. In the IV specification, we pool together years into an “Anticipation” period (2002 - 2006), which corresponds to the period when discussion of the reform was moving expectations around prices, and a “Post” period (2007 -

Figure 3.6: Higher Taxes Have No Effect on Sales Prices



Notes: This figure plots the effect of land taxes on sales prices. 95% confidence intervals are shown for the event-study coefficients, with standard errors clustered at the treatment area by land share decile level. The red line shows the difference in difference estimate that pools all post-2007 years compared to the price in 2001, and the red shading shows the 95% confidence intervals from that regression.

2019). Table 3.1, Column (1) shows the OLS coefficients. The results indicate that a 1 p.p. increase in land taxes (relative to total assessed values) lowers home prices by 3% during the anticipation period and 7% during the post period, though neither coefficient is statistically significant.

Table 3.1, Column (2) shows the baseline IV specification. The bottom row of the table reports that the IV has an F-Statistic of over 200, implying that the first stage is strong and the instrument meets the relevance requirement for a valid IV. Home prices are unchanged in the anticipation period. In the post period, the coefficient is positive, implying higher taxes increase home prices, though it is very close to zero and statistically insignificant.


Table 3.1 reports the discount rate consistent with full capitalization using the lower bound of the 95% confidence intervals we estimate. This represents the lowest possible discount rate that lies within our 95% confidence interval. In our baseline specification, we can rule out full capitalization at discount rates that are lower than 8%. This is substantially higher than most recent estimates of the discount rate for housing cash flows. Recent work has carefully estimated this parameter in a variety of settings. Giglio et al. (2014) estimate net discount rates of 1.9%, He et al. (2024)

Figure 3.7: The Effect of Land Taxes Using Separate Identification Strategies



Notes: This figure plots the effect of land taxes on sales prices. 95% confidence intervals are shown for the event-study coefficients, with standard errors clustered at the treatment area by land share decile level. The red line shows the difference in difference estimate that pools all post-2007 years compared to the price in 2001, and the red shading shows the 95% confidence intervals from that regression. Panel (A) shows the effect using the policy shock instrument only, while Panel (B) shows the effect after residualizing on treatment area fixed effects. Thus, Panel (B) just uses variation from within the same treatment area, coming from heterogeneous exposure to land tax rates.

Figure 3.8: Higher Taxes Have No Effect on Sales Prices of Single-Family Homes



Notes: This figure plots the effect of land taxes on sales prices in a subsample of single-family homes. 95% confidence intervals are shown for the event-study coefficients, with standard errors clustered at the treatment area by land share decile level. The red line shows the difference in difference estimate that pools all post-2007 years compared to the price in 2001, and the red shading shows the 95% confidence intervals from that regression.

estimate discount rates of 2.4%, while Bäcker-Peral et al. (2024) provide estimates between 5.3% and 2.8%. Most relevant to our work, Jordá et al. (2019) estimate housing returns of 3.4% in Denmark during this period. All of these estimates, except for the earliest estimates from Bäcker-Peral et al. (2024), are outside of our 99.9% confidence intervals. Thus, our estimates rule out price movements consistent with all of these estimates with a high degree of confidence.

Table 3.1, Columns (3) through (7) show robustness checks for our baseline estimate. Columns (2) and (3) trim outlier observations in the data, with minimal effects on the coefficients or standard errors. Column (4) restricts attention to the single-family home sample. The coefficients are even closer to zero, and all of the previous estimates lie outside our 95% confidence interval. Columns (6) and (7) show the results when using the policy shock and exposure instruments separately. Though somewhat less precisely estimated (especially for the exposure instrument alone), when using the policy instrument alone, we can rule out the estimates from Jordá et al. (2019), Giglio et al. (2014), He et al. (2024), as well as the more recent estimates from Bäcker-Peral et al. (2024). Using the

Table 3.1: Effect on Home Prices and Implied Discount Rates

	OLS	IV Baseline	IV Trim 1%	IV Trim 5%	IV SF	IV Policy Shock	IV Exposure
Ant. x Treat	-3.72 (2.80)	-0.90 (6.67)	-1.24 (6.64)	-4.06 (6.37)	0.75 (6.15)	-8.56 (6.17)	15.95 (20.02)
Post x Treat	-7.17 (4.22)	2.66 (7.76)	3.85 (7.67)	0.61 (6.94)	0.33 (7.86)	-7.39 (7.90)	13.20 (23.02)
Obs.	1015414	1015414	1007798	942671	314045	616419	1015385
Discount rate l.b.	6.5	8.0	8.9	7.7	6.6	4.4	3.1
F-Stat		209.2	208.2	213.6	215.3	197.2	161.1

Notes: This table shows IV estimates of the effects of land tax changes on home prices and computes the smallest discount rate that is consistent with full capitalization. Specifically, the row discount rate l.b. (lower bound) denotes the discount rate implied by our results assuming full capitalization of taxes into home prices. The discount rate is based on the lower bound of the calculated 95% confidence interval. Standard errors in parentheses.

exposure instrument alone leads to larger coefficients, implying less pass-through and providing suggestive evidence that service changes are unlikely to affect our estimates.

However, the literature on property tax capitalization provides less of a consensus on whether changes in taxes are passed through to home prices. In canonical work, Rosen (1982) estimates price effects that are consistent with full capitalization of taxes into prices with discount rates of approximately 14%. Recent work finds varying results: Horton (2024) finds zero effects on home prices, while Koster and Pinchbeck (2022) finds nominal discount rates between 3 and 4%. The persistence of our policy shock, variation from the exposure instrument within municipal boundaries, and the fact that our setting studies land, as opposed to property taxes, makes our setting uniquely suited to study this question.

More relevant to our work is previous work studying the same reform. Høj et al. (2018) estimate significant effects of land taxes on home prices. In Appendix C.3, we provide a detailed comparison of our results. While we use distinct empirical strategies and different sources of variation, we show that when using a specification closer to the one that they use, we can replicate the significant effects of land taxes on prices. However, we show that this effect is confounded by a correlation between the change in the tax rate and the number of vacation homes in a treatment area, as well as diverging price trends in those areas that coincided with the reform. Once we drop vacation homes from the sample and include a control for the share of vacation homes in a treatment area, the price effect is eliminated.

The reform we study provides a unique window to study tax capitalization. We document precise null effects of land taxes on residential home prices. Our estimates allow us to easily rule out full capitalization at standard discount rates. Overall, we provide crucial new evidence on how land taxes affect residential property markets. Our results provide new evidence that the slope of the

supply and demand curves is insufficient for understanding and predicting tax incidence for taxes on residential property. Our results suggest that households can avoid paying land taxes by selling their property or moving. This contradicts the basic efficiency argument for land taxes, which is that an individual's tax burden is unaffected by their behavioral responses. In the next section, we empirically test for behavioral responses by studying the effect of land taxes on development, mobility, and housing allocation.

3.8 Aggregate Effects of Land Taxes

Much of the policy debate about land taxes has focused on their efficiency. Unlike the conventional property tax, land taxes do not contain any **explicit** disincentive to develop property. However, we find no effect of land taxes on prices, suggesting that land taxes operate differently in practice than in theoretical models. In addition, policymakers and the public often express concern about whether or not property taxes may push older retirees out of their homes because they do not have the liquidity to pay taxes on (potentially rising) home values. We use our setting to provide new evidence on these policy debates by investigating the long-term effect of land taxes on development, mobility, and housing allocation.

The Effect of Land Taxes on Property Development

The effects of land taxes on development is uncertain. In our model, when landowners are fully compensated, taxes do not affect the demand or supply of buildings. However, without full compensation, income effects reduce demand for buildings and, thus, the total amount of developed property. In addition, if taxes are capitalized into land values, and land is used as collateral for development or if developers have short-term liquidity constraints, higher land taxes might serve as a negative net wealth or liquidity shock that can reduce the ability of developers to increase the housing supply. In contrast, a Georgist viewpoint suggests that **higher taxes** can actually **increase development**. George (1880) argued that higher taxes would improve land usage. This could come about either because higher taxes would reduce the asset price of land, which would improve land allocation and reduce speculation in the land market, or because the need to pay higher land taxes would lead to higher investment in land.

We provide new empirical estimates on the effects of land taxes on property development. Because property development is a slow-moving stock variable, we estimate the long-term effects of changes in land taxes using a long differences specification. We estimate the following regression at the treatment area level:

$$y_{a,2015} - y_{a,2006} = \beta \Delta \tilde{\tau}_a + \eta_{m(a)} + \delta X_a + e_a. \quad (3.19)$$

In this specification, the dependent variable is the change in our outcome variables from the year before the tax change occurred until the last year of our data.⁶ Our primary coefficient of interest is β , which represents the ITT effect of the policy shock instrument. We only use variation from

⁶We end most of the property development specifications in 2015 because of changes in the Housing Census.

the policy shock to study the effects of land taxes on aggregate outcomes, because only the policy shock varies at the treatment area level. We include treatment area controls, including detailed fixed effects for the share of vacation homes given the correlation documented in Appendix C.3, and a new municipality fixed effect.

We first investigate the effect of land taxes on total residential square footage. We sum up the total area dedicated to residential housing in each treatment area, investigating whether the total amount of housing increased or decreased more in areas with higher land tax increases after the reform.⁷ Panel (A) of Figure 3.9 shows that we estimate a precise null effect of land taxes had on aggregate housing development. The β coefficient is 0.001 and statistically insignificant. We can rule out decreases or increases of more than 1 p.p. in total residential development in response to a 1 p.p. increase in (predicted) land value tax rates over 9 years. Figure 3.9, Panel (B) shows no pre-trends in residential development in nearby areas that were eventually exposed to tax increases vs. decreases before the reform. Our results suggest that higher land taxes do not affect aggregate residential housing. This is consistent with our baseline neoclassical model, where land taxes do not affect the incentive to develop housing.

Much of the public debate on land taxes has focused on the effects on undeveloped plots. A revenue-neutral reform that moves from a property tax to a land tax (as proposed in Detroit) shifts the tax burden from highly developed properties (low θ) to land-intensive, empty plots (high θ). While the tax should not change the **marginal** incentive to develop undeveloped plots, many commentators have suggested this could improve the usage of undeveloped plots. To investigate the effects of land taxes on undeveloped plots, we focus on a set of undeveloped plots in 1995 and test whether the propensity to develop those plots depends upon the land tax rate.⁸ We say a property is developed if it is recorded as having a business use by the property tax authorities or if the housing census records residential development on that plot. In this case, our outcome variable is the log change in the number of undeveloped properties.

Figure 3.10, Panel (A) shows the effect of land taxes on the change in the number of undeveloped properties in each treatment area that **remain undeveloped** in 2015 relative to 2006. Our results show that higher taxes do not have an effect on the development of empty plots. The point estimate is positive, suggesting that higher land taxes lead to less development of empty plots. However, the point estimate is smaller than the standard error, indicating the effect is highly statistically insignificant. The results suggest that a 1 p.p. change in the tax change instrument increases the number of properties that remain undeveloped by around 1.3 p.p. Panel (B) of Figure 3.10 shows that there is no evidence of pre-trends. Before the reform, undeveloped land was developed at a similar rate in areas that experienced tax decreases and increases. Overall, our results suggest that land taxes have limited effects on the development of empty plots. As a result, our results suggest that moving from a property tax to a land tax might increase housing development if property taxes are a significant impediment to residential growth. However, our results also allow us to rule out that land taxes alone might increase development. The 95% confidence interval for our ITT estimates can rule out that a 1 p.p. increase in land taxes increases the number of developed plots by more than 1.5 p.p. over nine years.

⁷We once again drop vacation homes.

⁸We focus on a set of properties for which we have a balanced panel and can observe the use of the property in each year.

Figure 3.9: The Effect of Land Taxes on Aggregate Residential Development

(a) Long Difference: Effect of Land Taxes on Residential Development




(b) Placebo Test: Effect of Land Taxes on Residential Development in Pre-Period




Notes: Panel (A) plots the long difference between 2015 and 2006 in total residential square footage on the land tax change instrument. Panel (B) plots the same effect, but from 1995 to 2001, before the tax change, as a test for pre-trends. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

Figure 3.10: The Effect of Land Taxes on Development of Empty Lots

(a) Long Difference: Effect of Land Taxes on Development of Empty Plots



(b) Placebo Test: Effect of Land Taxes on Development of Empty Plots in Pre-Period



Notes: Panel (A) plots the long difference between 2015 and 2006 in log count of undeveloped plots from 1995 that remain undeveloped on the land tax change instrument. Panel (B) plots the same effect, but from 1995 to 2001, before the tax change, as a test for pre-trends. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

Figure 3.11 investigates heterogeneity in the effects of land taxes on the usage of undeveloped properties. Panels (A) and (B) show the effect split by whether the plot is on agricultural or non-agricultural land. Public discussion of land taxes has focused chiefly on empty plots in urban land. In addition, no response on agricultural land might mask significant results in urban areas because agricultural land is subject to lower tax rates and advantageous land assessments, and the propensity to develop on agricultural land is much lower. Figure 3.11, however, shows that we estimate null effects on both agricultural and non-agricultural land. The estimate for building on non-agricultural plots is exactly zero, while the estimate for non-agricultural plots is close to zero and statistically insignificant.

Overall, our results suggest that higher land taxes have precisely estimated zero effects on property development. Our findings hold when looking specifically at the development of empty plots of land, or when looking at aggregate residential development. Our results imply that higher land taxes do not reduce the total housing stock or residential building in an area. As a result, a move from property taxes to land taxes might lead to higher development. However, our results suggest that land taxes alone do not increase growth over the medium- to long term.

The Effect of Land Taxes on Mobility


An especially strong concern in policy circles is that higher property taxes may push individuals out of their homes. This concern is especially prevalent for retirees on fixed incomes, who may be unable to keep up with tax payments if property values increase. This concern is somewhat mitigated in our setting because retirees in Denmark are eligible for a subsidized loan to cover their land tax payments. However, in our setting, these mobility concerns may be amplified because, in contrast to the predictions of theoretical models, we find null effects of taxes on home prices. As a result, the statutory incidence is identical to the economic incidence, and households can avoid the tax by moving to low-tax areas. This has important implications for the efficiency of land value taxes. If households resort where they live in response to higher tax rates, this reduces the efficiency case for land taxes.

We study this empirically by investigating whether property owners exposed to higher vs. lower land taxes become more likely to sell their home in subsequent years. We focus on owners of primary residences and run the following regression:

$$y_{a,t} = \beta \Delta \tilde{\tau}_a + \eta_{m(a)} + \delta X_a + e_a. \quad (3.20)$$


This specification mirrors the long difference specification, but our outcome variable is the (log) share of incumbent homeowners (from 1995) who retain ownership of their home in year t . Figure 3.12, Panel (A) shows the effect on ownership in 2019. We estimate a precise zero effect of land taxes on the probability an individual still owns their home in 2019. This indicates that the tax shock does not affect individual decisions to sell or remain in their house. Panel (B) shows the effect in 2001, before the shock, and serves as a placebo test. As with the long-term effect, we estimate a small, statistically insignificant effect of (eventual) tax changes on the propensity of incumbents to move. Our results show that higher taxes do not push incumbent homeowners out of their homes. This suggests that, although homeowners could theoretically avoid the land tax by moving, the tax shock does not prompt households to sell their house.

Figure 3.11: The Effect of Land Taxes on Development of Empty Lots



Notes: Panel (A) plots the long difference between 2015 and 2006 in log count of undeveloped non-agricultural plots from 1995 that remain undeveloped on the land tax change instrument. Panel (B) plots the same effect, but for agricultural plots. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

Figure 3.12: Land Taxes Do Not Cause Homeowners to Sell Their Home



Notes: Panel (A) plots the share of primary, residential homeowners from 1995 who still own their home in 2019 on the land tax change instrument. Panel (B) plots the same effect, but in 2001, to test for pre-trends. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

The Effect of Land Taxes on the Allocation of Housing

Though our model abstracts from the homeownership decision, in a standard neoclassical model, land taxes should not affect the decision about whether or not to purchase a home. Theoretically, higher land taxes should be offset by lower purchase prices, and households should be indifferent between purchasing a home and paying the rental price (the user cost of housing) each period. However, in models with financial frictions, individuals who are down-payment constrained may be induced to purchase homes if higher taxes reduce the asset price of housing, thus allowing some households that were previously down-payment constrained to purchase homes. Coven et al. (2024) argue that higher property taxes would increase homeownership, especially among financially constrained households. However, since we estimate null effects on purchase prices, this channel is unlikely to be operative in our setting.

We use our experiment to study whether higher land taxes increased homeownership rates. We re-estimate our long-difference specification but replace the outcome variable with the (log) change in the number of residential homeowners. Since we find null effects on development, this specification will capture whether higher land taxes affect the decision of whether to own or rent a home. Figure 3.13, Panel (A) plots the long difference between 2019 and 2006. While the estimate is positive, suggesting that homeownership numbers rose, the estimate is close to zero and statistically insignificant. Our results suggest that higher land taxes did not induce higher homeownership rates in the medium-to-long run. Panel (B) serves as a placebo test and shows that in the pre-period, the slope is nearly identical to the post-period slope and statistically insignificant (though only marginally). Overall, we find little evidence that higher taxes increase homeownership by relaxing financial constraints.

We also investigate whether higher land taxes changed the composition of homeowners. We repeat the long-difference specification but replace the dependent variable with the log share of homeowners in a treatment area with particular characteristics. Figure 3.14, Panel (A) shows the change in the share of young homeowners, which we define as under 35. We find no change in the share of young homeowners in areas that experienced tax increases relative to areas that experienced tax decreases. This provides additional evidence against the arguments made in Coven et al. (2024) that higher property taxes can increase home purchases by young, financially constrained households. If the price of housing does not fall in response to higher taxes, then there is no change households ability to pay for a home.


In contrast, Panel (B) finds a significant decrease in the share of homeowners over 60 years old. The effect is highly statistically significant ($p < 0.001$). Appendix Figure C.2 also shows no significant differences between areas that eventually experienced tax increases versus decreases in the pre-period. As a result, we find that while higher land taxes did not lead to increased home sales in high-tax areas among incumbent homeowners, older individuals sorted away from high-tax areas to low-tax regions.

In contrast, Panels (C) and (D) show that higher taxes have no effect on the share of lower-income or very high-income homeowners. Panel (C) shows the effect of taxes on the share of working households with below-median income.⁹ We find a precise zero effect of higher taxes on the share of lower-income homeowners. Panel (D) shows the effects of higher taxes on the share

⁹We focus on working households to avoid the results being driven by retirees, who have low reported incomes.

Figure 3.13: Land Taxes Do Not Increase Homeownership

(a) Long Difference: Change in the Number of Homeowners



(b) Placebo Test: Change in the Number of Homeowners in the Pre-Period



Notes: This figure plots binned scatterplots of the change in the number of primary homeowners in a given treatment area. Panel (A) shows the long difference between 2006 and 2019, while Panel (B) shows the long difference between 1995 and 2001, as a placebo test. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

of rich households, which we define as households in the top 10% of the income distribution. The coefficient is negative but has a relatively small magnitude and is statistically insignificant.¹⁰

Taken together, we estimate that higher land taxes did not affect the number of homeowners in a treatment area. This contradicts the argument that higher property taxes reduce financial constraints by lowering sales prices, leading to increased homeownership (Coven et al., 2024). However, we find that higher taxes reallocated the existing housing stock: higher taxes led to a decrease in the share of older homeowners. This is consistent with older homeowners sorting away from areas with higher land taxes during retirement years. Because we find no effect of higher taxes on the probability of remaining in a home, the decrease in the share of older homeowners is unlikely to be driven by older homeowners selling their homes in response to higher tax payments. This suggests that land taxes affect the composition of homeownership within an area, and represents an efficiency cost of land taxation not present in standard models. This also shows that differentiated local taxes affect the allocation of homeowners across space.

3.9 Discussion

Discussion

What can explain our null effect on home prices? A natural explanation in the Danish context comes from rent control for housing. In Denmark, approximately 50% of households are renters, and nearly all private rental housing is subject to rent controls. In addition, for the major types of rent control regulations, rent can explicitly be raised when taxes are increased (Whitehead, 2012). In Appendix C.2, we present a version of our model where land is owned by landlords who rent to the household. In this model, the highest rent a landlord can charge for a unit of land is:

$$R = \bar{R} + \tau P^L.$$

When the rent cap is binding, the price of land is given by the following expression:

$$P^L = \frac{1+r}{r} \bar{R}.$$

Thus, when landlords can pass through the tax increases, the price no longer depends on taxes: $V = \frac{\bar{R}}{r}$. In contrast, in this model higher taxes lead to higher rent payments:


$$R = \bar{R} + \tau P^L = \bar{R} + \tau \frac{1+r}{r} \bar{R}.$$

As a result, a key implication of this setting and explanation is that because the price does not adjust, but the rents do, renters pay the tax rather than landlords. After the land tax increase, the landlord's net profits are the same (as higher rents exactly offset higher tax payments), but the rent owed by the tenant has increased.

Given this potential explanation, it is particularly interesting that we also estimated precise null effects of land taxes in the single-family home market. In this market, more than 80% of households

¹⁰Appendix Figure C.2 shows that none of these outcomes have significant pre-trends.

Figure 3.14: Land Taxes and the Allocation of Housing



Notes: The binscatters in this figure plots the long difference between 2019 and 2006 in log share of homeowners in a given treatment area on the land tax change instrument. Panel (A) plots the effect for homeowners who are under 35 years old. Panel (B) plots the same effect, but for homeowners who are over 60 years old. Panel (C) plots the effect for below-median income working homeowners, while Panel (D) plots the effect individuals in the top 10% of the income distribution. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

are both landowners and tenants (to themselves). As a result, even in markets dominated by owner-occupied housing, higher taxes do not lead to higher prices. It may be the case that prices are set by the outside option of renting, in which case rental restrictions will still affect market prices, even if they bind no tenants in equilibrium. However, given the substantial evidence of non-standard pricing behavior in the housing market (Genesove and Mayer, 2001; Andersen et al., 2022), this channel is speculative.

Similarly, given the recent evidence of pass-through of property taxes to rents documented across a variety of institutional contexts (Löffler and Siegloch, 2021; Baker, 2024; Watson and Ziv, 2024), it is entirely plausible that pass-through of land taxes to renters could hold in markets less bound by rent control. This study also provides new evidence of how non-neoclassical forces may affect tax incidence. Benzarti (2024) emphasizes how fairness norms play a role in determining incidence of the payroll tax. Our work suggests that other institutional regulations may be important for understanding the incidence of property taxation. More work on land and property taxes across a variety of institutional settings is vital for understanding how regulations may interact with tax policy and affect incidence.


Conclusion

Despite a long and rich intellectual history, we have surprisingly little evidence on the effects of land-value taxation. In this paper, we exploit a unique institutional reform that generated large, persistent, quasi-experimental variation in land-value tax rates across towns to study how individuals and markets respond to land taxes. We have two main sets of results. First, we show that land taxes are not capitalized into home prices. We estimate precise null effects of taxes on residential property prices using a repeat-sales index. This result contradicts the standard view that incumbent owners pay for higher taxes. This result is also inconsistent with the efficiency argument for land taxes: incumbents can sell their homes to avoid paying taxes, implying taxes may generate an inefficient allocation of housing. Finally, it suggests that taxes are paid by renters and future purchasers of land, as well as incumbent homeowners.

Figure 3.15 shows the CDF of land ownership in Denmark in 2001, prior to the reform. Unsurprisingly, land holdings are concentrated at the top of the distribution. The top 3 deciles own nearly 50% of the value of the land. In a neoclassical world, it is these landowners who would pay a land tax. In contrast, our results show that the tax burden is shared with renters and future purchasers. The bars in pink show the share of individuals in each income decile who own no real property. In contrast, non-property owners are concentrated at the bottom of the income distribution: in each decile in the bottom half of the income distribution, more than 60% of individuals are renters. In the bottom decile, the number is more than 80%. Our results show that land tax payments are shifted to these individuals, implying a much more regressive tax than in standard models.

Second, we study how higher taxes affect aggregate development, mobility, and housing allocation. Consistent with the neoclassical benchmark, but inconsistent with the most optimistic Georgist predictions, we find null effects of higher land taxes on aggregate development. This suggests that moving from traditional property taxes to land taxes may spur new development, though higher land taxes alone do not lead to more intensive land use. We find this effect when looking at aggregate residential development or specifically at the development of empty plots, which has been the subject of an active policy debate. We also find that higher taxes do not affect a region's

Figure 3.15: Land Ownings Across the Income Distribution



Notes: This figure plots CDFs of holdings of land across the income distribution. The bars in pink show the share of individuals in each income decile who own no real property.

total quantity of homeownership. Taken together, these results speak against recent work that has suggested higher land taxes may spur increased density or homeownership (Murphy and Seegert, 2024; Coven et al., 2024).

However, we find evidence that higher land taxes may have reallocated housing away from older homeowners. Because we find null effects on the probability that incumbent homeowners sell their homes, we interpret these results as suggesting that older households sort away from areas with high taxes. This work also provides evidence against the view that land taxes are perfectly efficient because taxes affect where households choose to locate.

Our findings have important implications for policymakers choosing from a range of possible tax bases. Land taxes have traditionally been viewed favorably because they do not induce substitution effects. In contrast, our results suggest that land taxes have consequences for efficiency. However, these effects are somewhat limited as land taxes do not disincentivize housing production or land usage. However, land taxes are paid broadly across the income and wealth distribution. Our findings suggest that the land tax is paid by renters, future landowners, and current owner-occupied homeowners, who are both landlord and tenant. The only people who do not pay any tax are absentee landlords, who can pass off the tax to tenants. This imperils the political viability of

land and property taxes and provides a potential explanation for the limited popular support for property taxation (Elmendorf et al., 2024). In ongoing work, we study how individual renters and residents of owner-occupied housing responded to this persistent change in their land tax burden.

Bibliography

- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67, 251–333.
- ACCONCIA, A., G. CORSETTI, AND S. SIMONELLI (2014): “Mafia and Public Spending: Evidence on the Fiscal Multiplier from a Quasi-experiment,” *American Economic Review*, 104, 2185–2209.
- ACEMOGLU, D. AND W. B. HAWKINS (2014): “Search with Multi-Worker Firms,” *Theoretical Economics*, 9, 583–628.
- AIYAGARI, S. R., L. J. CHRISTIANO, AND M. EICHENBAUM (1992): “The output, employment, and interest rate effects of government consumption,” *Journal of Monetary Economics*, 30, 73–86.
- ALBRECHT, J., C. CARRILLO-TUDELA, AND S. VROMAN (2018): “On-the-Job Search with Match-Specific Amenities,” *Economics Letters*, 162, 15–17.
- AMODEO, F. AND E. BRIGANTI (2025): “High-Frequency Cross-Sectional Identification of Military News Shocks,” Working paper.
- ANDERSEN, S., C. BADARINZA, L. LIU, J. MARX, AND T. RAMADORAI (2022): “Reference Dependence in the Housing Market,” *American Economic Review*, 112, 3398–3440.
- ANGELETOS, G.-M., C. LIAN, AND C. K. WOLF (2024): “Can Deficits Finance Themselves?” *Econometrica*, 92, 1351–1390.
- ANTOLIN-DIAZ, J. AND P. SURICO (2025): “The Long-Run Effects of Government Spending,” *American Economic Review*, 115, 2376–2413.
- ANTONOVA, A., R. LUETTICKE, AND G. J. MÜLLER (2025): “The Military Multiplier,” Working paper.
- AUCLERT, A., B. BARDÓCZY, AND M. ROGNLIE (2023): “MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models,” *The Review of Economics and Statistics*, 105, 700–712.
- AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models,” *Econometrica*, 89, 2375–2408.
- AUCLERT, A. AND M. ROGNLIE (2018): “Inequality and Aggregate Demand,” Tech. Rep. w24280, National Bureau of Economic Research.

- AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2020): “Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model,” Working paper.
- (2024): “The Intertemporal Keynesian Cross,” *Journal of Political Economy*, 132, 4068–4121.
- AUERBACH, A., Y. GORODNICHENKO, AND D. MURPHY (2020): “Local Fiscal Multipliers and Fiscal Spillovers in the USA,” *IMF Economic Review*, 68, 195–229.
- AUERBACH, A. J. (2006): “Who Bears the Corporate Tax? A Review of What We Know,” *Tax Policy and the Economy*, 20, 1–40.
- AUERBACH, A. J. AND Y. GORODNICHENKO (2012): “Measuring the Output Responses to Fiscal Policy,” *American Economic Journal: Economic Policy*, 4, 1–27.
- (2016): “Effects of Fiscal Shocks in a Globalized World,” *IMF Economic Review*, 64, 177–215.
- AUERBACH, A. J., Y. GORODNICHENKO, AND D. MURPHY (2024): “Macroeconomic Frameworks: Reconciling Evidence and Model Predictions from Demand Shocks,” *American Economic Journal: Macroeconomics*, 16, 190–229.
- AUERBACH, A. J. AND D. YAGAN (2025): “Robust Fiscal Stabilization,” Tech. Rep. w33374, National Bureau of Economic Research.
- BAGGER, J., F. FONTAINE, M. GALENIANOS, AND I. TRAPEZNIKOVA (2022): “Vacancies, Employment Outcomes and Firm Growth: Evidence from Denmark,” *Labour Economics*, 75, 102103.
- BAKER, S. (2024): “Property Tax Pass-Through to Renters: A Quasi-Experimental Approach,” *Working Paper*.
- BARKAI, S. (2020): “Declining Labor and Capital Shares,” *The Journal of Finance*, 75, 2421–2463.
- BARRO, R. J. (1981): “Output Effects of Government Purchases,” *Journal of Political Economy*, 89, 1086–1121.
- BARRO, R. J. AND C. J. REDLICK (2011): “Macroeconomic Effects From Government Purchases and Taxes,” *The Quarterly Journal of Economics*, 126, 51–102.
- BASSIER, I., A. DUBE, AND S. NAIDU (2022): “Monopsony in Movers The Elasticity of Labor Supply to Firm Wage Policies,” *Journal of Human Resources*, 57, S50–s86.
- BASU, S. AND M. KIMBALL (2003): “Investment planning costs and the effects of monetary and fiscal policy,” Working paper.
- BAXTER, M. AND R. G. KING (1993): “Fiscal Policy in General Equilibrium,” *American Economic Review*, 83, 315–334.
- BAYER, C., B. BORN, AND R. LUETTICKE (2024): “Shocks, Frictions, and Inequality in US Business Cycles,” *American Economic Review*, 114, 1211–1247.

- BEN ZEEV, N. AND E. PAPPA (2017): “Chronicle of A War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks,” *The Economic Journal*, 127, 1568–1597.
- BEN ZEEV, N., V. A. RAMEY, AND S. ZUBAIRY (2023): “Do Government Spending Multipliers Depend on the Sign of the Shock?” *AEA Papers and Proceedings*, 113, 382–87.
- BENIGNO, P. AND G. B. EGGERTSSON (2024): “The Slanted-L Phillips Curve,” *AEA Papers and Proceedings*, 114, 84–89.
- BENZARTI, Y. (2024): “Tax Incidence Anomalies,” NBER Working Paper w32819, National Bureau of Economic Research.
- BERGER, D., K. HERKENHOFF, A. R. KOSTØL, AND S. MONGEY (2024): “An Anatomy of Monopsony: Search Frictions, Amenities, and Bargaining in Concentrated Markets,” *NBER Macroeconomics Annual*, 38, 1–47.
- BERGER, D., K. HERKENHOFF, AND S. MONGEY (2022): “Labor Market Power,” *American Economic Review*, 112, 1147–1193.
- BERTHEAU, A., P. CAHUC, S. JAGER, AND R. VEJLIN (2021): “Turnover Costs: Evidence from Unexpected Worker Separations,” Working Paper.
- BI, H., M. PHILLOT, AND S. ZUBAIRY (2025): “Treasury Supply Shocks: Propagation through Debt Expansion and Maturity Adjustment,” Working paper.
- BILBIIE, F. O. (2008): “Limited asset markets participation, monetary policy and (inverted) aggregate demand logic,” *Journal of Economic Theory*, 140, 162–196.
- (2011): “Nonseparable Preferences, Frisch Labor Supply, and the Consumption Multiplier of Government Spending: One Solution to a Fiscal Policy Puzzle,” *Journal of Money, Credit and Banking*, 43, 221–251.
- BIOLSI, C. (2019): “Local effects of a military spending shock: Evidence from shipbuilding in the 1930s,” *Review of Economic Dynamics*, 32, 227–248.
- BLANCHARD, O. AND R. PEROTTI (2002): “An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output,” *The Quarterly Journal of Economics*, 117, 1329–1368.
- BLATTER, M., S. MUEHLEMANN, AND S. SCHENKER (2012): “The Costs of Hiring Skilled Workers,” *European Economic Review*, 56, 20–35.
- BLOESCH, J., B. LARSEN, AND B. TASKA (2022): “Which Workers Earn More at Productive Firms? Position Specific Skills and Individual Worker Hold-up Power,” Working Paper.
- BLOESCH, J., S. J. LEE, AND J. WEBER (2024): “Firm Wage Setting and On-the-Job Search Limit Wage-Price Spirals,” *Available at SSRN 4734451*.

- BLOESCH, J. AND J. WEBER (2023): “Congestion in Onboarding Workers and Sticky R&D,” *FRB of New York Staff Report*.
- BLOOM, N., F. GUVENEN, B. S. SMITH, J. SONG, AND T. VON WACHTER (2018): “The Disappearing Large-Firm Wage Premium,” in *AEA Papers and Proceedings*, vol. 108, 317–22.
- BOEHM, C. E. AND T. N. KRONER (2025): “The U.S., Economic News, and the Global Financial Cycle,” *The Review of Economic Studies*, rda020.
- BONNET, O., G. CHAPPELLE, A. TRANNOY, AND E. WASMER (2021): “Land is back, it should be taxed, it can be taxed,” *European Economic Review*, 134, 103696.
- BORUSYAK, K., P. HULL, AND X. JARAVEL (2022): “Quasi-Experimental Shift-Share Research Designs,” *The Review of Economic Studies*, 89, 181–213.
- BRIGANTI, E., G. BRUNET, AND V. SELLEMI (2025): “When Does Government Spending Matter? It’s All in the Measurement,” Working paper.
- BROCKMEYER, A., A. ESTEFAN, J. C. SUÁREZ SERRATO, AND K. RAMÍREZ (2024): “Taxing Property in Developing Countries: Theory and Evidence from Mexico,” Working paper, Duke University.
- BROER, T., P. KRUSELL, AND E. ÖBERG (2023): “Fiscal multipliers: A heterogeneous-agent perspective,” *Quantitative Economics*, 14, 799–816.
- BRUNET, G. (2024): “Stimulus on the Home Front: The State-Level Effects of WWII Spending,” *The Review of Economics and Statistics*, 1–46.
- BUNDGAARD, U. AND K. VRANGBÆK (2007): “Reform by Coincidence? Explaining the Policy Process of Structural Reform in Denmark,” *Scandinavian Political Studies*, 30, 491–520.
- BURDETT, K. AND D. T. MORTENSEN (1998): “Wage Differentials, Employer Size, and Unemployment,” *International Economic Review*, 39, 257–273.
- BURDETT, K. AND T. VISHWANATH (1988): “Balanced Matching and Labor Market Equilibrium,” *Journal of Political Economy*, 96, 1048–1065.
- BURKS, S. V., B. COWGILL, M. HOFFMAN, AND M. HOUSMAN (2015): “The Value of Hiring through Employee Referrals,” *The Quarterly Journal of Economics*, 130, 805–839.
- BURNSIDE, C., M. EICHENBAUM, AND J. D. M. FISHER (2004): “Fiscal shocks and their consequences,” *Journal of Economic Theory*, 115, 89–117.
- BÄCKER-PERAL, V., J. HAZELL, AND A. MIAN (2024): “Dynamics of the Long Term Housing Yield: Evidence from Natural Experiments,” Working paper.
- CALDARA, D. AND C. KAMPS (2017): “The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers,” *The Review of Economic Studies*, 84, 1015–1040.

- CALDWELL, S. AND N. HARMON (2019): “Outside Options, Bargaining, and Wages: Evidence from Coworker Networks,” Working paper.
- CARD, D. (2022): “Who Set Your Wage?” *American Economic Review*, 112, 1075–1090.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 36, S13–S70.
- CARRILLO-TUDELA, C., H. GARTNER, AND L. KAAS (2023): “Recruitment Policies, Job-Filling Rates, and Matching Efficiency,” *Journal of the European Economic Association*, 21, 2413–2459.
- CARVALHO, M., J. GALINDO DA FONSECA, AND R. SANTARROSA (2024): “How Are Wages Determined?: A Quasi-Experimental Test of Wage Determination Theories,” Working paper.
- CHAN, M., S. SALGADO, AND M. XU (2023): “Heterogeneous Passthrough from TFP to Wages,” *Available at SSRN 3538503*.
- CHODOROW-REICH, G. (2019): “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?” *American Economic Journal: Economic Policy*, 11, 1–34.
- CHODOROW-REICH, G., L. FEIVESON, Z. LISCOW, AND W. G. WOOLSTON (2012): “Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act,” *American Economic Journal: Economic Policy*, 4, 118–145.
- CHRISTIANO, L., M. EICHENBAUM, AND S. REBELO (2011): “When Is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119, 78–121.
- CHRISTIANSEN, P. M. AND M. KLITGAARD (2008): *Den Utænkelige Reform: Strukturreformens Tilblivelse 2002-2005*, University of Southern Denmark Studies in History and Social Sciences, Princeton University Press.
- CHRISTIANSEN, T. (2012): “Ten Years of Structural Reforms in Danish Healthcare,” *Health Policy*, 106, 114–119.
- COIBION, O. AND Y. GORODNICHENKO (2012): “Why Are Target Interest Rate Changes So Persistent?” *American Economic Journal: Macroeconomics*, 4, 126–162.
- COLES, M. G. (2001): “Equilibrium Wage Dispersion, Firm Size, and Growth,” *Review of Economic Dynamics*, 4, 159–187.
- COLES, M. G. AND D. T. MORTENSEN (2016): “Equilibrium Labor Turnover, Firm Growth, and Unemployment,” *Econometrica*, 84, 347–363.
- CORSETTI, G., A. MEIER, AND G. J. MÜLLER (2012): “What determines government spending multipliers?” *Economic Policy*, 27, 521–565.
- COTTON, C. D. (2024): “Debt, deficits and interest rates,” *Economica*, 91, 911–943.
- COVEN, J., S. GOLDER, A. GUPTA, AND A. NDIAYE (2024): “Property Taxes and Housing Allocation Under Financial Constraints,” *Working Paper*.

- COX, L., G. J. MÜLLER, E. PASTÉN, R. SCHOENLE, AND M. WEBER (2024): “Big G ,” *Journal of Political Economy*, 132, 3260–3297.
- DAGGETT, S. AND R. O’ROURKE (2007): “Defense Procurement: Full Funding Policy—Background, Issues, and Options for Congress,” Report RL31404, Congressional Research Service.
- DATTA, N. (2023): “The Measure of Monopsony: The Labour Supply Elasticity to the Firm and its Constituents,” Working Paper.
- DAVIS, S. J., R. J. FABERMAN, AND J. C. HALTIWANGER (2013): “The Establishment-Level Behavior of Vacancies and Hiring,” *The Quarterly Journal of Economics*, 128, 581–622.
- DEBORTOLI, D. AND J. GALÍ (2025): “Heterogeneity and Aggregate Fluctuations: Insights from TANK Models,” *NBER Macroeconomics Annual*, 39, 307–357.
- DHYNE, E., A. K. KIKKAWA, T. KOMATSU, M. MOGSTAD, AND F. TINTELNOT (2025): “Firm Responses and Wage Effects of Foreign Demand Shocks with Fixed Labor Costs and Monopsony,” *American Economic Review*, 115, 4328–4368.
- DRECHSEL, T. AND K. MIURA (2025): “The macroeconomic effects of bank regulation: New evidence from a high-frequency approach,” Working paper.
- DUBE, A., E. FREEMAN, AND M. REICH (2010): “Employee Replacement Costs,” Working paper.
- DUPOR, B. AND R. GUERRERO (2017): “Local and aggregate fiscal policy multipliers,” *Journal of Monetary Economics*, 92, 16–30.
- DUSTMANN, C., A. GLITZ, U. SCHÖNBERG, AND H. BRÜCKER (2016): “Referral-Based Job Search Networks,” *The Review of Economic Studies*, 83, 514–546.
- EDELBERG, W., M. EICHENBAUM, AND J. D. M. FISHER (1999): “Understanding the Effects of a Shock to Government Purchases,” *Review of Economic Dynamics*, 2, 166–206.
- ELMENDORF, C. S., C. NALL, AND S. OKLOBDZIJA (2024): “What State Housing Policies Do Voters Want? Evidence from a Platform-Choice Experiment,” *Journal of Political Institutions and Political Economy*.
- ELSBY, M. W., B. HOBIJN, AND A. ŞAHİN (2013): “The Decline of the US Labor Share,” *Brookings Papers on Economic Activity*, 2013, 1–63.
- ENGBOM, N., C. MOSER, AND J. SAUERMANN (2023): “Firm Pay Dynamics,” *Journal of Econometrics*, 233, 396–423.
- ERCEG, C. J., D. W. HENDERSON, AND A. T. LEVIN (2000): “Optimal monetary policy with staggered wage and price contracts,” *Journal of Monetary Economics*, 46, 281–313.
- FARHI, E. AND I. WERNING (2016): “Fiscal Multipliers: Liquidity Traps and Currency Unions,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 2417–2492.

- FISHER, J. D. AND R. PETERS (2010): “Using Stock Returns to Identify Government Spending Shocks,” *The Economic Journal*, 120, 414–436.
- FLODEN, M. AND J. LINDÉ (2001): “Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?” *Review of Economic Dynamics*, 4, 406–437.
- FRIEDMAN, M. (1978): “Milton Friedman Gives the Answers!” Collected Works of Milton Friedman Project Records, Hoover Institution Archives.
- GAGLIARDONE, L. AND M. GERTLER (2023): “Oil Prices, Monetary Policy and Inflation Surges,” Working Paper 31263, National Bureau of Economic Research.
- GALÍ, J. (2011): “The Return of the Wage Phillips Curve,” *Journal of the European Economic Association*, 9, 436–461.
- GALÍ, J., J. D. LÓPEZ-SALIDO, AND J. VALLÉS (2007): “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 5, 227–270.
- GARIN, A. AND F. SILVÉRIO (2024): “How Responsive Are Wages to Firm-Specific Changes in Labour Demand? Evidence from Idiosyncratic Export Demand Shocks,” *The Review of Economic Studies*, 91, 1671–1710.
- GAVAZZA, A., S. MONGEY, AND G. L. VIOLANTE (2018): “Aggregate Recruiting Intensity,” *American Economic Review*, 108, 2088–2127.
- GENESOVE, D. AND C. MAYER (2001): “Loss Aversion and Seller Behavior: Evidence from the Housing Market*,” *The Quarterly Journal of Economics*, 116, 1233–1260.
- GEORGE, H. (1880): *Progress and Poverty: An Inquiry into the Cause of Industrial Depressions, and of Increase of Want with Increase of Wealth. The Remedy.*
- GERTLER, M. AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7, 44–76.
- GIGLIO, S., M. MAGGIORI, AND J. STROEBEL (2014): “Very Long-Run Discount Rates,” *The Quarterly Journal of Economics*, 130, 1–53.
- GOLDIN, C. D. (1980): “War,” in *Encyclopedia of american economic history: Studies of the principal movements and ideas*, ed. by G. Porter, New York: C. Scribner’s Sons, vol. 1, 935–957.
- GOMEZ CRAM, R., H. KUNG, AND H. LUSTIG (2025): “Can Treasury Markets Add and Subtract?” Working Paper 33604, National Bureau of Economic Research.
- GOTTFRIES, A. AND G. JAROSCH (2023): “Dynamic monopsony with large firms and noncompetes,” NBER Working Paper No. 31965.
- GOUIN-BONENFANT, É. (2022): “Productivity Dispersion, Between-Firm Competition, and the Labor Share,” *Econometrica*, 90, 2755–2793.

- GRIGSBY, J., E. HURST, AND A. YILDIRMAZ (2021): “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data,” *American Economic Review*, 111, 428–471.
- GÜRKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 1, 55–93.
- HALL, R. E. (1986): “The Role of Consumption in Economic Fluctuations,” in *The American Business Cycle: Continuity and Change*, ed. by R. J. Gordon, University of Chicago Press, 237–266.
- (2009): “By How Much Does Gdp Rise If the Government Buys More Output?” *Brookings Papers on Economic Activity*, 40, 183–249.
- HALL, R. E. AND A. I. MUELLER (2018): “Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values,” *Journal of Political Economy*, 126, 1594–1637.
- HAZELL, J. AND S. HOBLER (2025): “Do Deficits Cause Inflation? A High Frequency Narrative Approach,” Working paper.
- HE, Z., M. HU, Z. WANG, AND V. YAO (2024): “Valuing Long-Term Property Rights with Anticipated Political Regime Shifts,” *American Economic Review*, 114, 2701–47.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): “Optimal Tax Progressivity: An Analytical Framework*,” *The Quarterly Journal of Economics*, 132, 1693–1754.
- HEBOUS, S. AND T. ZIMMERMANN (2021): “Can government demand stimulate private investment? Evidence from U.S. federal procurement,” *Journal of Monetary Economics*, 118, 178–194.
- HEISE, S. AND T. PORZIO (2022): “Labor Misallocation Across Firms and Regions,” Working Paper.
- HERBST, E. P. AND B. K. JOHANSEN (2024): “Bias in Local Projections,” *Journal of Econometrics*, 240, 105655.
- HIRSCH, B., E. J. JAHN, A. MANNING, AND M. OBERFICHTNER (2022): “The Wage Elasticity of Recruitment,” Working Paper.
- HOOKE, M. A. AND M. M. KNETTER (1997): “The Effects of Military Spending on Economic Activity: Evidence from State Procurement Spending,” *Journal of Money, Credit and Banking*, 29, 400–421.
- HORTON, E. (2024): “Are Tax Benefits Priced into Home Values? Evidence from the Tax Cuts and Jobs Act,” *Working Paper*.
- HUMMELS, D., R. JØRGENSEN, J. MUNCH, AND C. XIANG (2014): “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data,” *American Economic Review*, 104, 1597–1629.

- HØJ, A. K., M. R. JØRGENSEN, AND P. SCHOU (2018): “Land Tax Changes and Full Capitalisation,” *Fiscal Studies*, 39, 365–380.
- ILZETZKI, E. (2025): “Guns & Growth: The Economic Consequences of Defense Buildups,” Kiel Report 2, Kiel Institute for the World Economy.
- ILZETZKI, E., E. G. MENDOZA, AND C. A. VÉGH (2013): “How big (small?) are fiscal multipliers?” *Journal of Monetary Economics*, 60, 239–254.
- JÄGER, S. AND J. HEINING (2022): “How Substitutable are Workers? Evidence from Worker Deaths,” NBER Working Paper No. w30629.
- JAKOBSEN, K., H. KLEVEN, J. KOLSRUD, C. LANDAIS, AND M. MUNÓZ (2024): “Taxing Top Wealth: Migration Responses and Their Aggregate Economic Implications,” *Working Paper*.
- JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96, 1589–1610.
- JORDÀ, Ò. (2005a): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- (2005b): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, AND A. M. TAYLOR (2019): “The Rate of Return on Everything, 1870–2015,” *The Quarterly Journal of Economics*, 134, 1225–1298.
- JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, AND A. M. TAYLOR (2019): “The Rate of Return on Everything, 1870–2015*,” *The Quarterly Journal of Economics*, 134, 1225–1298.
- JORDA, O. AND A. M. TAYLOR (2025): “Local Projections,” *Journal of Economic Literature*, 63, 59–110.
- JØRGENSEN, P. L. AND S. H. RAVN (2022): “The inflation response to government spending shocks: A fiscal price puzzle?” *European Economic Review*, 141.
- KAAS, L. AND P. KIRCHER (2015): “Efficient Firm Dynamics in a Frictional Labor Market,” *American Economic Review*, 105, 3030–3060.
- KÄNZIG, D. R. (2021): “The Macroeconomic Effects of Oil Supply News: Evidence from Opec Announcements,” *American Economic Review*, 111, 1092–1125.
- (2023): “The Unequal Economic Consequences of Carbon Pricing,” Working Paper 31221, National Bureau of Economic Research.
- KÄNZIG, D. R. AND R. RAGHAVAN (2025): “The Macroeconomic Effects of Supply Chain Shocks: Evidence from Global Shipping Disruptions,” Working paper.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary Policy According to HANK,” *American Economic Review*, 108, 697–743.

- KAPLAN, G. AND G. L. VIOLANTE (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82, 1199–1239.
- (2022): “The Marginal Propensity to Consume in Heterogeneous Agent Models,” *Annual Review of Economics*, 14, 747–775.
- KEKRE, R. AND M. LENEL (2025): “The High-Frequency Effects of Dollar Swap Lines,” *American Economic Review: Insights*, 7, 107–123.
- KLEIN, A., S. J. HVIID, T. S. HVOLBØL, P. L. KRAMP, AND E. H. PEDERSEN (2016): “House Price Bubbles and the Advantages of Stabilising Housing Taxation,” *Danmarks Nationalbank Monetary Review*.
- KLINE, P., N. PETCOVA, H. WILLIAMS, AND O. ZIDAR (2019): “Who Profits from Patents? Rent-Sharing at Innovative Firms,” *The Quarterly Journal of Economics*, 1343, 1404.
- KOH, D., R. SANTAEULÀLIA-LLOPIS, AND Y. ZHENG (2020): “Labor Share Decline and Intellectual Property Products Capital,” *Econometrica*, 88, 2609–2628.
- KOPPLIN, K. (2024): “County Property Tax Capitalization in U.S. Cities,” *Working Paper*.
- KOSTER, H. R. AND E. W. PINCHBECK (2022): “How Do Households Value the Future? Evidence from Property Taxes,” *American Economic Journal: Economic Policy*, 14, 207–39.
- KUHN, P. (2004): “Is Monopsony the Right Way to Model Labor Markets? A Review of Alan Manning’s Monopsony in Motion,” *International Journal of the Economics of Business*, 11, 369–378.
- LAMADON, T., M. MOGSTAD, AND B. SETZLER (2022): “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market,” *American Economic Review*, 112, 169–212.
- LANGELLA, M. AND A. MANNING (2021): “Marshall lecture 2020: The Measure of Monopsony,” *Journal of the European Economic Association*, 19, 2929–2957.
- LASSEN, D. D. AND S. SERRITZLEW (2011): “Jurisdiction Size and Local Democracy: Evidence on Internal Political Efficacy from Large-scale Municipal Reform,” *American Political Science Review*, 105, 238–258.
- LEEPER, E. M. (1991): “Equilibria under ‘active’ and ‘passive’ monetary and fiscal policies,” *Journal of Monetary Economics*, 27, 129–147.
- LEEPER, E. M., T. B. WALKER, AND S.-C. S. YANG (2013): “Fiscal Foresight and Information Flows,” *Econometrica*, 81, 1115–1145.
- LÖFFLER, M. AND S. SIEGLOCH (2021): “Welfare Effects of Property Taxation,” *Working Paper*.
- MANNING, A. (2003): *Monopsony in Motion*, Princeton University Press.

- (2006): “A Generalised Model of Monopsony,” *The Economic Journal*, 116, 84–100.
- (2011): “Imperfect Competition in the Labor Market,” in *Handbook of Labor Economics*, Elsevier, vol. 4, 973–1041.
- (2021): “Monopsony in Labor Markets: A Review,” *ILR Review*, 74, 3–26.
- (2025): “The Immobile Incumbent Problem in a Model of Short-Term Wage-Posting,” *German Economic Review*.
- MATSUDAIRA (2014): “Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations,” *The Review of Economics and Statistics*, 96, 92–102.
- MCKAY, A. AND R. REIS (2016): “The Role of Automatic Stabilizers in the U.S. Business Cycle,” *Econometrica*, 84, 141–194.
- MELOSI, L., H. MORITA, A. ROGANTINI PICCO, AND F. ZANETTI (2025): “The Signaling Effects of Fiscal Announcements,” Working paper.
- MERTENS, K. AND M. O. RAVN (2012): “Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks,” *American Economic Journal: Economic Policy*, 4, 145–181.
- MICHAILLAT, P. AND E. SAEZ (2024): “ $u^* = \sqrt{uv}$: The Full-Employment Rate of Unemployment in the United States,” Brookings Papers on Economic Activity Conference Draft, Fall.
- MONACELLI, T. AND R. PEROTTI (2008): “Fiscal Policy, Wealth Effects, and Markups,” Working Paper 14584, National Bureau of Economic Research.
- MONGEY, S. AND G. L. VIOLANTE (2025): “Macro Recruiting Intensity from Micro Data,” *American Economic Journal: Macroeconomics*, *Forthcoming*.
- MONTIEL OLEA, J. L. AND C. PFLUEGER (2013): “A Robust Test for Weak Instruments,” *Journal of Business & Economic Statistics*, 31, 358–369.
- MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): “Local Projection Inference Is Simpler and More Robust Than You Think,” *Econometrica*, 89, 1789–1823.
- MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, ERIC QIAN, AND C. K. WOLF (forthcoming): “Local Projections or VARs? A Primer for Macroeconomists,” *NBER Macroeconomics Annual*.
- MOUNTFORD, A. AND H. UHLIG (2009): “What are the effects of fiscal policy shocks?” *Journal of Applied Econometrics*, 24, 960–992.
- MUEHLEMANN, S. AND H. PFEIFER (2016): “The Structure of Hiring Costs in Germany: Evidence from Firm-Level Data,” *Industrial Relations: A Journal of Economy and Society*, 55, 193–218.
- MURPHY, D. AND N. SEEGER (2024): “Implicit Land Taxes and Their Effect on the Real Economy,” *Working Paper*.

- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models*,” *The Quarterly Journal of Economics*, 123, 1415–1464.
- (2014): “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review*, 104, 753–792.
- (2018): “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect*,” *The Quarterly Journal of Economics*, 133, 1283–1330.
- NEENAN, A. G. (2024a): “Defense Primer: Department of Defense Contractors,” In Focus IF10600, Congressional Research Service.
- (2024b): “Defense Primer: Procurement,” In Focus IF10599, Congressional Research Service.
- NIELSEN, M. G. AND A. RZEŹNIK (2014): “House Prices and Taxes,” *Working Paper*.
- NORMAN, C. (1988): “B-1’s Problems Stir Debate on Bomber’s Role,” *Science*, 241, 1282–1283.
- OATES, W. E. (1969): “The Effects of Property Taxes and Local Public Spending on Property Values: An Empirical Study of Tax Capitalization and the Tiebout Hypothesis,” *Journal of Political Economy*, 77, 957–971.
- OHANIAN, L. E. (1997): “The Macroeconomic Effects of War Finance in the United States: World War II and the Korean War,” *American Economic Review*, 87, 23–40.
- ORCHARD, J. D., V. A. RAMEY, AND J. F. WIELAND (2025): “Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates*,” *The Quarterly Journal of Economics*, 140, 2001–2052.
- PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 103, 2530–2553.
- PEROTTI, R. (2007): “In Search of the Transmission Mechanism of Fiscal Policy,” *NBER Macroeconomics Annual*, 22, 169–249.
- RAMEY, V. A. (2011a): “Can Government Purchases Stimulate the Economy?” *Journal of Economic Literature*, 49, 673–685.
- (2011b): “Identifying Government Spending Shocks: It’s all in the Timing*,” *The Quarterly Journal of Economics*, 126, 1–50.
- (2016): “Macroeconomic Shocks and Their Propagation,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 71–162.
- (2019): “Ten Years After the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?” *Journal of Economic Perspectives*, 33, 89–114.
- RAMEY, V. A. AND M. D. SHAPIRO (1998): “Costly capital reallocation and the effects of government spending,” *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–194.

- RAMEY, V. A. AND S. ZUBAIRY (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126, 850–901.
- ROGNLIE, M. (2015): "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?" *Brookings Papers on Economic Activity*, 2015, 1–69.
- (2018): "Comment on "Accounting for Factorless Income"," in *NBER Macroeconomics Annual 2018, volume 33*, University of Chicago Press, 235–248.
- ROSEN, K. T. (1982): "The Impact of Proposition 13 on House Prices in Northern California: A Test of the Interjurisdictional Capitalization Hypothesis," *Journal of Political Economy*, 90, 191–200.
- ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.
- ROTEMBERG, J. J. AND M. WOODFORD (1992): "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, 100, 1153–1207.
- SATURNO, J. V. AND M. S. LYNCH (2023): "The Appropriations Process: A Brief Overview," Report R47106, Congressional Research Service.
- (2024): "Defense Primer: Defense Appropriations Process," In Focus IF10514, Congressional Research Service.
- SCHMIEDER, J. F. (2023): "Establishment Age and Wages," *Journal of Econometrics*, 233, 424–442.
- SCHUBERT, G., A. STANSBURY, AND B. TASKA (2025): "Employer Concentration and Outside Options," Working Paper.
- SCHWERHOFF, G., O. EDENHOFER, AND M. FLEURBAEY (2022): "Equity and Efficiency Effects of Land Value Taxation," *Working Paper*.
- SEEGMILLER, B. (2021): "Valuing Labor Market Power: The Role of Productivity Advantages," Available at SSRN 4412667.
- SHOAG, D. (2013): "Using State Pension Shocks to Estimate Fiscal Multipliers since the Great Recession," *American Economic Review*, 103, 121–124.
- SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.
- SMITH, A. (1776): *An Inquiry into the Nature and Causes of the Wealth of Nations*.
- SOKOLOVA, A. AND T. SORENSEN (2021): "Monopsony in Labor Markets: A Meta-Analysis," *ILR Review*, 74, 27–55.
- SORKIN, I. (2018): "Ranking Firms using Revealed Preference," *The Quarterly Journal of Economics*, 133, 1331–1393.

- STEVENS, M. (2004): “Wage-tenure Contracts in a Frictional Labour Market: Firms’ Strategies for Recruitment and Retention,” *The Review of Economic Studies*, 71, 535–551.
- STOCK, J. H. AND M. W. WATSON (2018): “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *The Economic Journal*, 128, 917–948.
- SUÁREZ SERRATO, J. C. AND P. WINGENDER (2016): “Estimating Local Fiscal Multipliers,” Working Paper 22425, National Bureau of Economic Research.
- TANAKA, S., L. WARREN, AND D. WICZER (2023): “Earnings Growth, Job Flows and Churn,” *Journal of Monetary Economics*, 135, 86–98.
- THE WALL STREET JOURNAL (2010): “Gates Talks of Tough Choices Ahead,” *The Wall Street Journal*, yochi J. Dreazen. May 10, 2010.
- UNITED STATES SENATE (1951): “Concentration of Defense Contracts: Report of the Select Committee on Small Business,” Report 551, U.S. Senate, Washington D.C.
- WATSON, C. L. AND O. ZIV (2024): “A Test for Pricing Power in Urban Housing Markets,” Unpublished Manuscript.
- WHITEHEAD, C. (2012): “The Private Rented Sector in the New Century: A Comparative Approach,” Tech. rep., Boligøkonomisk Videncenter.
- WIEGAND, C. (2025): “The Effect of Fiscal Policy Shocks on Asset Prices,” Working paper.
- WONG, F. (2023): “Mad as Hell: Property Taxes and Financial Distress,” *Working Paper*.
- (2024): “Taxing Homeowners Who Won’t Borrow,” *Working Paper*.
- WOODFORD, M. (2011): “Simple Analytics of the Government Expenditure Multiplier,” *American Economic Journal: Macroeconomics*, 3, 1–35.

Appendix A

Appendix to The Macroeconomic Effects of Defense Spending News

A.1 Campbell-Shiller Decomposition

Let P_t be the price of a stock index of defense contractors, and let D_t denote dividends. The gross return from period t to $t + 1$ is $R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t}$. We can express the log price $p_t \equiv \log(P_t)$ as follows using the Campbell-Shiller approximation:

$$p_t - d_t = \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{t+1+j}) + \frac{\kappa}{1 - \rho}$$

where lowercase variables are in logs and $\rho = \frac{1}{1 + \frac{D}{P}} \approx 0.99$. We can solve this for the return r_{t+1} :

$$r_{t+1} = d_t - p_t + \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \sum_{j=1}^{\infty} \rho^j r_{t+1+j} + \frac{\kappa}{1 - \rho}.$$

We can do the same exercise for a broader market index with price P_t^m and subtract the market return to get the excess return of defense contractors relative to the market:

$$r_{t+1} - r_{t+1}^m = (d_t - p_t) - (d_t^m - p_t^m) + \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - \Delta d_{t+1+j}^m) - \sum_{j=1}^{\infty} \rho^j (r_{t+1+j} - r_{t+1+j}^m).$$

This equation holds ex post, so it also holds ex ante for rational expectations and for non-rational expectations that respect identities (see Campbell's textbook). We assume rational expectations and take the expectation conditional on information available at time t :

$$\mathbb{E}_t(r_{t+1} - r_{t+1}^m) = (d_t - p_t) - (d_t^m - p_t^m) + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - \Delta d_{t+1+j}^m) - \mathbb{E}_t \sum_{j=1}^{\infty} \rho^j (r_{t+1+j} - r_{t+1+j}^m).$$

Now consider an announcement that affects the expected the path of future dividends and discount rates. Let $\mathbb{E}_t$ be the expectation immediately *before* the announcement that does not take the new

information into account, and let $\tilde{\mathbb{E}}_t$ denote the expectation immediately *after* the announcement, but before $t + 1$. The change in excess return when the news arrive is

$$\underbrace{(\tilde{\mathbb{E}}_t - \mathbb{E}_t)(r_{t+1} - r_{t+1}^m)}_{\text{Our shock } z_t} = \underbrace{(\tilde{\mathbb{E}}_t - \mathbb{E}_t) \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - \Delta d_{t+1+j}^m)}_{\text{Cash flow news}} - \underbrace{(\tilde{\mathbb{E}}_t - \mathbb{E}_t) \sum_{j=1}^{\infty} \rho^j (r_{t+1+j} - r_{t+1+j}^m)}_{\text{Discount rate news}}.$$

Consider the following set of assumptions:

1. The announcement contains news about the path of future defense spending $\{G_{t+k}\}_{k \geq 1}$
2. Defense cash flows depend on defense spending through the simple relation:

$$\Delta d_t = \alpha + \gamma G_t + e_t$$

where $\gamma > 0$ and $e_t \sim iid(0, \sigma_e^2)$.

3. Market cash flows are unaffected by the announcement. (This is stronger than necessary. It is enough to assume that they are affected less than defense stocks.)
4. The announcement affects the expected future discount rates r_t and r_t^m in the same way (so the discount rate news term drops out).

Under these assumptions, we get

$$(\tilde{\mathbb{E}}_t - \mathbb{E}_t)(r_{t+1} - r_{t+1}^m) = (\tilde{\mathbb{E}}_t - \mathbb{E}_t) \sum_{j=0}^{\infty} \rho^j \gamma \Delta G_{t+1+j},$$

so in this case our shocks reflect news about future defense spending.

What could go wrong? If the announcements contains no or little news, or if the cash flows don't depend on spending ($\gamma = 0$), then we would have no first stage (but this is testable and we find that our shocks are highly predictive of future defense spending). Another possibility is that markets are irrational and the assumption that asset prices reflect all available information is a not a good approximation. In this case, stock prices may not react to news about defense spending. If they consistently over- or under-react, then our shock series is still a valid instrument, since the shocks are only identified up to scale. Nonetheless, we will simply assume expectations are rational.

Another possibility is that assumption is violated and market cash flows rise (future earnings expectations go up). Defense stocks constitute about 2–3% of total public market capitalization in the U.S., so direct effect of higher defense cash flows is small. However, as we emphasize in the paper, there may be multiplier effects: An increase in defense spending can potentially stimulate economic activity more broadly. Assume that defense spending news affect market cash flows *less* than defense contractor cash flows:

$$\Delta d_t^m = \alpha^m + \lambda \gamma G_t + e_t^m$$

where $\lambda \in [0, 1)$ determines how much the market is affected relative to the defense industry. Maintaining the other assumption, we get

$$(\tilde{\mathbb{E}}_t - \mathbb{E}_t)(r_{t+1} - r_{t+1}^m) = (\tilde{\mathbb{E}}_t - \mathbb{E}_t) \sum_{j=0}^{\infty} \rho^j (1 - \lambda) \gamma \Delta G_{t+1+j},$$

i.e. this simply re-scales the instrument.

Finally, discount rates may move differentially in response to the announcement. The discount rates reflect time preference and risk premia. For instance, if the announcement affects the risk premium of the defense stock index relative to the market, then our shock will also reflect discount rate news.


Why does the expected path of defense spending change when a defense bill passes in Congress? We can consider two edge cases:

- Investors expected (with high confidence) a different level of spending than the one that was announced. For example, they knew the bill was going to pass, but in the last hours additional spending was added.
- Investors were uncertain about future levels of spending and the announcement resolved this uncertainty. For example, the bill is set in stone, but it is a coin flip whether it is going to pass. The expected spending increase is $G^e = \frac{1}{2}G^{pass} + \frac{1}{2}G^{fail}$, and if the bill passes the surprise is $G^{pass} - G^e$ and defense stocks go up.

Our reading of the Congressional Record and the news coverage around is that late changes to the bill do occur, but the second case is more common.


A.2 Additional Figures

Figure A.1: Real Government Spending per Capita in the United States




Notes: This figure plots defense spending and total government spending in real 2017-chained dollars quarterly from 1947 to 2020.

Figure A.2: Defense Stock Index Coverage: Share of Top 100 Defense Contracts




Notes: This figure plots the coverage of our index, as well as the top 25 defense contractors, as a percentage of the top 100 defense contracts in dollar terms quarterly from 1958 to 2023.

Figure A.3: Quarterly Defense Spending Surprises




Notes: This figure plots our defense spending surprise series quarterly from 1947 to 2010. We convert our daily series to quarterly by summing all shocks in a given quarter.

Figure A.4: Quarterly Defense Spending Surprises




Notes: This figure plots our defense spending surprise series against (Ramey, 2011a)'s narrative shocks quarterly from 1947 to 2010. We convert our daily series to quarterly by summing all shocks in a given quarter. The defense spending series is measured in excess returns while the narrative series is measured as a percentage of GDP.

Figure A.5: Shock Autocorrelation




Notes: This figure plots autocorrelation tests of our defense spending surprise series. The left panel plots the sample autocorrelation and the p-value of a Ljung-Box test of autocorrelation. The right panel plots the partial sample autocorrelation and the p-value of an Engle test.

Figure A.6: Shock Predictability (Granger Causality Tests)




Notes: This figure plots the p-values of a series of Granger-Causality tests. Each column reports the p-value from a test of whether that variable Granger-Causes the defense spending surprise series. The last column reports the p-value from a joint test of all variables in the proceeding columns.

Figure A.7: Defense and Market Returns in September 2001




Notes: This figure plots the returns of defense contractors and the market around the September 11 attacks. The left panel plots the cumulative excess returns (with September 11th normalized to zero). The right panel plots the corresponding returns of defense contractors in blue circles and the market in red triangles.

Figure A.8: Spending Responses to a Defense Spending Surprise (Baseline, Long Horizon)




Notes: This figure plots the response of defense spending and total government spending to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars). The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons $h = 0, 1, 2, \dots, 48$. The dark and light shaded areas are 68% and 95% confidence intervals.

Figure A.9: Spending Responses to a Defense Spending Surprise (Post 1955, Long Horizon)




Notes: This figure plots the response of defense spending and total government spending to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars), restricting the sample to post-1955. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons $h = 0, 1, 2, \dots, 48$. The dark and light shaded areas are 68% and 95% confidence intervals.

Figure A.10: Spending Responses to a Defense Spending Surprise (Excluding September 11 Attacks, Long Horizon)




Notes: This figure plots the response of defense spending and total government spending to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars), excluding the September 11th attacks from the defense spending surprise series. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons $h = 0, 1, 2, \dots, 48$. The dark and light shaded areas are 68% and 95% confidence intervals.

Figure A.11: Spending Responses to a Defense Spending Surprise (Excluding Supplemental Appropriations, Long Horizon)




Notes: This figure plots the response of defense spending and total government spending to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars), excluding supplemental appropriations from the defense spending surprises. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons $h = 0, 1, 2, \dots, 48$. The dark and light shaded areas are 68% and 95% confidence intervals.

Figure A.12: Response of Taxes and Deficits to a Defense Spending Surprise (Pre-2000 Subsample)



Notes: This figure plots the response of defense spending and total spend to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars), using the pre-2000 sample. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.

Figure A.13: Response of Taxes and Deficits to a Defense Spending Surprise (Post-1954 Subsample)



Notes: This figure plots the response of defense spending, nondefense spending, federal taxes and transfers, and the federal deficits to a defense spending surprise, normalized so the defense spending response peaks at one dollar per capita (chained 2017-dollars), using the post 1954 sample. The solid blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are 68% and 95% confidence intervals.

A.3 Model

The construction and presentation of this model closely resembles the two-account HANK models of Auclert et al. (2024) and Kaplan et al. (2018).

Environment

Time is discrete, and all agents have perfect foresight over aggregate shocks. The economy is populated by final good producers, intermediate good producers, households, labor packers, individual labor unions, banks, a monetary authority, and a fiscal authority.

Firms

We adopt the common two-tier production model with a continuum of monopolistically competitive intermediate goods producers that sell their output to a final goods producer.

Final Good Producer The final good Y_t is a composite, made from combining intermediate inputs $Y_{j,t}$ using a CES technology: $Y_t = \left(\int_0^1 Y_{j,t}^{\frac{\nu^p-1}{\nu^p}} dj \right)^{\frac{\nu^p}{\nu^p-1}}$. The price of intermediate input j is $P_{j,t}$. The final good producer sells the final good Y_t to consumers at price P_t . The final good producer maximizes $P_t Y_t - \int_0^1 P_{j,t} Y_{j,t} dj$. Profit maximization by the final goods producer generates a demand curve for each intermediate good:

$$Y_{j,t} = \left(\frac{P_{j,t}}{P_t} \right)^{-\nu^p} Y_t \quad (\text{A.1})$$

where $P_t = \left(\int P_{j,t}^{1-\nu^p} dj \right)^{\frac{1}{1-\nu^p}}$. In symmetric equilibrium, all firms set the same price $P_{j,t} = P_t$, so the final good producer makes no profit from production.

Intermediate Goods Producers Intermediate goods producers combine labor $N_{j,t}$ and capital $K_{j,t-1}$ to produce the intermediate good $Y_{j,t}$. The production function is Cobb-Douglas, i.e., $Y_{j,t} = K_{j,t-1}^\alpha N_{j,t}^{1-\alpha}$. They pay nominal wage W_t for workers, investment costs $X_{j,t}$, and sell their product for price $P_{j,t}$. Each producer faces both nominal and real rigidities. Specifically, each producer faces Rotemberg price adjustment costs with partial indexation $\Psi_{j,t}^p = \frac{\psi^p}{2} \left(\frac{P_{j,t}}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p} P_{j,t-1}} - 1 \right)^2$ and investment adjustment costs $\Xi_{j,t} = \frac{1}{2\delta\xi} \left(\frac{X_{j,t}}{X_{j,t-1}} - 1 \right)^2 X_{t-1}$. Here, $\Pi_t = \frac{P_t}{P_{t-1}}$ is the gross nominal inflation rate, Π is the steady-state inflation rate, ι^p is the degree of price indexation, and X_{t-1} is the aggregate investment last period which intermediate producers take as given. Intermediate

goods producers choose their nominal price, labor, and capital to maximize their value

$$J_t(S_{j,t-1}) = \max_{N_{j,t}, S_{j,t}} \left\{ \frac{P_{j,t}}{P_t} \left(\frac{P_{j,t}}{P_t} \right)^{-\nu^p} Y_t - \frac{W_t}{P_t} N_{j,t} - X_{j,t} - \mu_t \left[\left(\frac{P_{j,t}}{P_t} \right)^{-\nu^p} Y_t - F(K_{j,t-1}, N_{j,t}) \right] \right. \\ \left. - Q_t [K_{j,t} - (1 - \delta)K_{j,t-1} - X_{j,t}] - \Xi_{j,t} - \Psi_{j,t}^p + \frac{J_{t+1}(S_{j,t})}{1 + r_t} \right\} \quad (\text{A.2})$$

where $S_{j,t} = \{K_{j,t}, P_{j,t}, X_{j,t}\}$ are the state variables, and $1 + r_t$ is the risk-free rate. The envelope conditions of the problem are

$$J_{K,t} = \mu_t F_{K,t} + (1 - \delta)Q_t \\ J_{X,t} = \frac{1}{\delta\xi} \left(\frac{X_{j,t}}{X_{j,t-1}} - 1 \right) \frac{X_{j,t} X_{t-1}}{X_{j,t-1}^2} \\ J_{P,t} = \psi^p \left(\frac{P_{j,t}}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p} P_{j,t-1}} - 1 \right) \frac{P_{j,t}}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p} P_{j,t-1}^2}$$

and the optimality conditions for state variables $S_{j,t}$ are

$$0 = -Q_t + \frac{J_{K,t+1}}{1 + r_t} \\ 0 = (Q_t - 1) - \frac{1}{\delta\xi} \left(\frac{X_{j,t}}{X_{j,t-1}} - 1 \right) \frac{X_{t-1}}{X_{j,t-1}} + \frac{J_{X,t+1}}{1 + r_t} \\ 0 = (1 - \nu^p) \left(\frac{P_{j,t}}{P_t} \right)^{1-\nu^w} \frac{Y_t}{P_{j,t}} + \mu_t \nu^p \left(\frac{P_{j,t}}{P_t} \right)^{-\nu^w} \frac{Y_t}{P_{j,t}} \\ - \psi^p \left(\frac{P_{j,t}}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p} P_{j,t-1}} - 1 \right) \frac{1}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p} P_{j,t-1}} + \frac{1}{1 + r_t} J_{P,t+1}^y.$$

In symmetric equilibrium, all firms set the same prices ($P_{j,t} = P_t$), labor demand ($N_{j,t} = N_t$), marginal cost ($\mu_{j,t} = \mu_t$), capital demand ($K_{j,t} = K_t$), and investment ($X_{j,t} = X_t$).

Taking first-order conditions with respect to labor, we see that at the symmetric equilibrium

$$(1 - \alpha)\mu_t Y_t = \frac{W_t}{P_t} N_t \quad (\text{A.3})$$

where μ_t is the real marginal cost.

Q-Theory of Investment The above conditions imply two standard Q-theory equations for capital demand and investment:

$$Q_t = \frac{1}{1 + r_t} [\mu_{t+1} F_{K,t+1} + (1 - \delta)Q_{t+1}] \\ (g_{X,t} - 1) = \delta\xi[Q_t - 1] + \frac{g_{X,t+1}}{1 + r_t} (g_{X,t+1} - 1) \quad (\text{A.4})$$

where $g_{X,t} \equiv \frac{X_t}{X_{t-1}}$ is the (gross) growth rate of investment.

New Keynesian Phillips Curve Combining the envelope condition for past prices, and the optimal pricing equation yields the nonlinear New Keynesian Phillips Curve

$$\begin{aligned} \left(\frac{\Pi_t}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p}} - 1 \right) \frac{\Pi_t}{(\Pi_{t-1})^{\iota^p} (\Pi)^{1-\iota^p}} &= Y_t(\mu_t - \mu) \\ &+ \frac{1}{1+r_t} \left(\frac{\Pi_{t+1}}{(\Pi_t)^{\iota^p} (\Pi)^{1-\iota^p}} - 1 \right) \frac{\Pi_{t+1}}{(\Pi_t)^{\iota^p} (\Pi)^{1-\iota^p}} \end{aligned} \quad (\text{A.5})$$

where $\mu = \frac{\nu^p-1}{\nu^p}$ is the optimal markup in steady state and $\kappa^p = \frac{\nu^p}{\psi^p}$ is the slope of the Phillips curve.

Production Dividends The per-period dividend earned by intermediate production is

$$D_t = Y_t - \frac{W_t}{P_t} N_t - X_t - \Xi_t - \Psi_t^p. \quad (\text{A.6})$$

The ex-dividend price of the firm satisfies

$$P_t^y = J_t(S_{t-1}) - D_t. \quad (\text{A.7})$$

The intermediate firms issue equity, which is held by households.

Household

The economy is populated by a unit mass of households, facing income uncertainty. Each household i in period t faces an idiosyncratic ability shock e_{it} , which follows a Markov process. The proportion of households with ability e in a given period is $\omega(e)$, the stationary probability of transition matrix Ω . The average ability is normalized to 1, such that $\int \omega(e)de = 1$. Household i consumes $C_{i,t}$ and receive flow utility $u(C_{it})$. Household i 's post-income is $Z_{it} = \frac{e_{it}^{1-\theta}}{\int e_{it}(h)^{1-\theta} dh} Z_t$, where $Z_t = \frac{W_t}{P_t} N_t - T_t$ is aggregate post-tax income. Households hold liquid A_{it} and illiquid $\tilde{A}_{it}$ assets. Holding the liquid asset incurs flow cost $\zeta(1+r_t)A_{it-1}$. Let Λ_{it} be an independent Bernoulli distribution, which equals 1 with probability λ and 0 with probability $1-\lambda$. If $\Lambda_{it} = 1$, households can able to transfer funds H_{it} between these accounts, otherwise, households cannot transfer funds. Households face borrowing constraints for both liquid and illiquid assets $A_{it}, \tilde{A}_{it} \geq 0$. The households budget constraints are

$$\begin{aligned} C_{it} + A_{it} &= Z_{it} + (1+r_{t-1})(1-\zeta)A_{it-1} - H_{it}\Lambda_{it} \\ \tilde{A}_{it} &= (1+r_{t-1})\tilde{A}_{it-1} + H_{it}\Lambda_{it} \end{aligned}$$

For ease of notation, we drop the i from the following analysis of households optimality conditions. Households solve the following optimization problem

$$\begin{aligned} V_t(A, \tilde{A}, \epsilon, \Lambda) &= \max_{C, A, \tilde{A}, T} u(C) - v(N) + \beta \mathbf{E}[V_{t+1}(A', \tilde{A}', \epsilon', \Lambda')] \\ \text{s.t. } C + A' &= \epsilon Z_t + (1+r_{t-1})(1-\zeta)A - T\Lambda \\ \tilde{A}' &= (1+r_{t-1})\tilde{A} + T\Lambda \\ A', \tilde{A}' &\geq 0 \end{aligned}$$

where $\epsilon = \frac{e^{1-\theta}}{\int e(i)^{1-\theta} di}$ is defined for simplicity and assumed to follow a Markov process with the same transition matrix as e . While there is perfect foresight for aggregate shocks, households do not have perfect foresight over idiosyncratic shocks. We assume that labor unions make labor supply choices N_{it} for households.

Liquid Assets The first order condition for liquid savings is:

$$u'(C) \geq \beta \mathbf{E}[V_{At+1}(A', \tilde{A}', \epsilon', \Lambda') | \epsilon, \Lambda].$$

When the borrowing constraints both bind the consumption function becomes $C_t(A, \tilde{A}, \epsilon, \Lambda) = \epsilon Z_t + (1+r_{t-1})(1-\zeta)(A+\Lambda\tilde{A})$. Notice that the inability to access illiquid assets (i.e. $\Lambda = 0$) amplifies the severity of hitting the borrowing constraint. Moreover, this allows for wealthy hand-to-mouth households. The envelope condition for liquid assets is:

$$V_{At}(A, \tilde{A}, \epsilon, \Lambda) = (1+r_{t-1})(1-\zeta)u'(C_t(A, \tilde{A}, \epsilon, \Lambda)).$$

Combining the envelope condition with the first condition yields an Euler equation for liquid assets

$$u'(C_t(A, \tilde{A}, \epsilon, \Lambda)) \geq (1-\zeta)\beta(1+r_t)\mathbf{E}[u'(C_{t+1}(A', \tilde{A}', \epsilon', \Lambda')) | \epsilon, \Lambda]. \quad (\text{A.8})$$

which features excessive discounting of the future, due to the flow costs of holding liquid assets embodied in the parameter ζ .

Illiquid Assets Repeating this exercise for illiquid assets, is complicated by the ability or inability to access assets. Without access (i.e. $\Lambda = 0$), the next period illiquid assets are given by the budget constraint $\tilde{A}' = (1+r_t)\tilde{A}$. With access (i.e. $\Lambda = 1$), an Euler equation for illiquid asset emerges

$$\begin{aligned} u'(C_t(A, \tilde{A}, \epsilon, 1)) \geq & \lambda\beta(1+r_t)\mathbf{E}[u'(C_{t+1}(A', \tilde{A}', \epsilon', 1)) | \epsilon] \\ & + (1-\lambda)\beta(1+r_t)\mathbf{E}[V_{\tilde{A}t+1}(A', \tilde{A}', \epsilon', 0) | \epsilon] \end{aligned} \quad (\text{A.9})$$

where $V_{\tilde{A}t}(A, \tilde{A}, \epsilon, 0)$ is the marginal value of illiquid assets without assets. By the envelope condition, this marginal value equals

$$V_{\tilde{A}t}(A, \tilde{A}, \epsilon, 0) = \beta(1+r_t)\mathbf{E}[V_{\tilde{A}t+1}(A', \tilde{A}', \epsilon', \Lambda') | \epsilon, \Lambda] \quad (\text{A.10})$$

Labor Union

We adopt the usual two-tier union structure with a continuum of monopolistically competitive differentiated labor unions that sell their labor services to a representative competitive labor packer.

Labor Packer A competitive labor packer combines tasks N_{kt} into an aggregate employment services using a CES technology: $N_t = [\int N_{kt}^{\frac{\nu^w-1}{\nu^w}}]^{-\frac{\nu^w}{\nu^w-1}} dk$. Workers in each union earn a task-specific wage W_{kt} . The labor packer maximizes profit $W_t N_t - \int W_{kt} N_{kt} dk$. Maximization produces demand curve for each task:

$$N_{kt} = \left(\frac{W_{kt}}{W_t} \right)^{-\nu^w} N_t \quad (\text{A.11})$$

where $W_t = (\int W_{kt}^{1-\nu^w} dk)^{\frac{1}{1-\nu^w}}$ is the aggregate nominal wage. At the symmetric equilibrium, all unions set the same wage $W_{kt} = W_t$, so the labor packer earns no profit.

Individual Union Problem Households i rent their labor services to individual labor unions k . Each union k aggregates the effective services of its members into a union specific task $N_{kt} = \int e_{it} N_{ikt} di$, where N_{ikt} represents the hours supplied by household i to union k at time t and e_{it} is the labor efficiency. Hence, $\int e_{it} N_{ikt} di$ represents the efficiency-weighted hours worked. Households spend a portion, possibly zero, of their labor hours working at each union. Each task is paid a wage W_{kt} . Each individual union faces utility costs for adjusting the union-specific nominal wage W_{kt} . Namely, individual labor unions face Rotemberg adjustment costs with partial indexing. At any time t , union k sets its wage W_{kt} to maximize welfare for all workers it employs,

$$\sum_{\tau=0}^{\infty} \beta^{t+\tau} \left[\int (u(C_{it}) - v(N_{it})) d\Upsilon_{i,t+\tau} - \frac{\psi^w}{2} \left(\frac{W_{k,t+\tau}}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w} W_{k,t+\tau-1}} - 1 \right)^2 \right] \quad (\text{A.12})$$

where $\Pi_t^w = \frac{W_t}{W_{t-1}}$ is the gross aggregate nominal wage inflation at time t , $\pi_w = \Pi_t^w - 1$ is wage inflation in equilibrium, and ι^w is the degree of wage indexation. Taking as given, the initial distribution of households Ω_{it} and the union specific demand curve.

The first-order condition of the union k with respect to W_{kt} is therefore

$$\begin{aligned} 0 = & \int \left(\frac{\partial C_{it}}{\partial W_{kt}} u'(C_{it}) - \frac{\partial N_{it}}{\partial W_{kt}} v'(N_{it}) \right) d\Upsilon_{it} \\ & - \psi^w \left(\frac{W_{kt}}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w} W_{kt-1}} - 1 \right) \frac{1}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w} W_{kt-1}} \\ & + \beta \psi^w \left(\frac{W_{kt+1}}{(\Pi_t^w)^{\iota^w} (\Pi^w)^{1-\iota^w} W_{kt}} - 1 \right) \left(\frac{W_{kt+1}}{(\Pi_t^w)^{\iota^w} (\Pi^w)^{1-\iota^w} W_{kt}} \right) \frac{1}{W_{kt}} \end{aligned}$$

Given the labor demand curve from the representative labor packer, the marginal impact of a task-specific wage increase on labor hours is

$$\frac{\partial N_{it}}{\partial W_{kt}} = -\nu^w \left(\frac{W_{kt}}{W_t} \right)^{-\nu^w-1} \frac{N_t}{W_t}$$

Each household i , working for union k earns total real earnings,

$$Z_{it} = \tau_t \left(\frac{W_{kt}}{P_t} e_{it} N_{kt} \right)^{1-\theta} = \tau_t \left(e_{it} W_{kt}^{1-\nu^w} \frac{N_t}{P_t W_t^{-\nu^w}} \right)^{1-\theta}.$$

By the envelope theorem, we know that if households consume all income from the wage increase that $\frac{\partial C_{it}}{\partial W_{kt}} = \frac{\partial Z_{it}}{\partial W_{kt}}$. Therefore, the marginal impact of a wage increase on household i consumption is

$$\frac{\partial C_{it}}{\partial W_{kt}} = (1 - \nu^w)(1 - MTR_{it})e_{it} \left(\frac{W_{kt}}{W_t} \right)^{-\nu^w} \frac{N_{kt}}{P_t}$$

where $MTR_{it} = 1 - (1 - \theta)\tau_t \left(\frac{W_{kt}}{P_t} e_{it} N_{kt} \right)^{-\theta}$ is household i 's marginal tax rate at time t .

Wage Phillips Curve We exploit the fact that at a symmetric equilibrium, wages across unions are equal $W_{kt} = W_t$, and thus labor supply is equalized across unions $N_{kt} = N_t$ and thus across households $N_{it} = N_t$. Therefore we can rewrite the optimal wage formula:

$$\begin{aligned} & \left(\frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} - 1 \right) \frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} \\ &= \frac{\nu^w}{\psi^w} N_t \int \left(v'(N_{it}) - \frac{\nu^w - 1}{\nu^w} (1 - MTR_{it}) e_{it} \frac{W_t}{P_t} u'(C_{it}) \right) di \\ &+ \beta \left(\frac{\Pi_{t+1}^w}{(\Pi_t^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} - 1 \right) \frac{\Pi_{t+1}^w}{(\Pi_t^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} \end{aligned}$$

Note that in equilibrium,

$$(1 - MTR_{it}) e_{it} \frac{W_t}{P_t} = (1 - \theta) \tau_t e_{it}^{1-\theta} \left(\frac{W_t}{P_t} \right)^{1-\theta} N_t^{-\theta} = (1 - \theta) \frac{e_{it}^{1-\theta}}{\int e_{it}^{1-\theta} di} \frac{Z_t}{N_t}$$

where Z_t is average post-tax income. Replacing consumption with the “virtual aggregate consumption” term C_t^* , we arrive at a non-linear New Keynesian Wage Phillips Curve:

$$\begin{aligned} & \left(\frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} - 1 \right) \frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} = \frac{\nu^w}{\psi^w} \left(N_t v'(N_t) - \frac{\nu^w - 1}{\nu^w} (1 - \theta) Z_t u'(C_t^*) \right) \\ &+ \beta \left(\frac{\Pi_{t+1}^w}{(\Pi_t^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} - 1 \right) \frac{\Pi_{t+1}^w}{(\Pi_t^w)^{\iota^w} (\Pi^w)^{1-\iota^w}} \end{aligned} \tag{A.13}$$

where $u'(C_t^*) = \int \frac{e_{it}^{1-\theta} u'(C_{it})}{\int e_{it}^{1-\theta} di} di$. Where $\kappa^w = \frac{\nu^w}{\psi^w} N v'(N) = \frac{\nu^w}{\psi^w}$ because $N = 1$ in steady-state, and $\mu^w = \frac{\nu^w - 1}{\nu^w}$.

Financial Markets

The financial system includes a representative bank and production firms that issue equity.

Banks There is a single representative competitive bank. Let B_t real bonds, and B_t^n be a nominal bond in zero net-supply, that can only be held by the bank. Banks' real profits each period are

$$D_t^b = (1 + r_{t-1})B_{t-1} + (1 + i_{t-1})\frac{B_{t-1}^n}{P_t} - B_t - \frac{B_t^n}{P_t}$$

where r_t is the interest rate earned on real bonds, and i_t is the interest rate earned on nominal bonds. Profit maximization by banks and a no arbitrage conditions implies that the bank discounts at the real rate $\tilde{r}_t = r_t$, and that the Fisher equation

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}} \quad (\text{A.14})$$

holds for all $t \geq 0$. In equilibrium, these banks make no profits from buying and selling bonds. Since nominal bonds are in zero net supply, $B_t^n = 0$ for all banks. As in Auclert et al. (2024), we also assume that the banks start with zero assets, so we have $B_t = 0$ at all times.

Liquid and Illiquid Accounts For simplicity, we assumed that households could only trade liquid A_{it} or illiquid $\tilde{A}_{it}$ assets. Now we show how these accounts are constructed. Suppose households can now hold assets in either real bonds or shares of production. Under this assumption, household i 's budget constraints become

$$\begin{aligned} C_{it} + P_t^y V_{it} + B_{it} &= Z_{it} + (1 - \zeta)[(P_t^y + D_t)V_{it-1} + (1 + r_{t-1})B_{it-1}] - H_{it}\Lambda_{it} \\ P_t^y \tilde{V}_{it} + \tilde{B}_{it} &= (P_t^y + D_t)\tilde{V}_{it-1} + (1 + r_{t-1})\tilde{B}_{it-1} + H_{it}\Lambda_{it} \end{aligned}$$

where $\{B_{it}, \tilde{B}_{it}\}$ are household i 's liquid and illiquid holdings of bonds, $\{V_{it}, \tilde{V}_{it}\}$ are household i 's liquid and illiquid shares of equity, and $\{P_t^y, D_t\}$ are the price and dividend of the firm share. The no arbitrage conditions implies that these prices must equal

$$P_t^y + D_t = (1 + r_{t-1})P_{t-1}^y$$

By defining the liquid and illiquid asset position to be $A_{it} \equiv P_t^y V_{it} + B_{it}$ and $\tilde{A}_{it} \equiv P_t^y \tilde{V}_{it} + \tilde{B}_{it}$ we can rewrite the household budget constraints as

$$\begin{aligned} C_{it} + A_{it} &= Z_{it} + (1 - \zeta)(1 + r_{t-1})A_{it-1} - H_{it}\Lambda_{it} \\ \tilde{A}_{it} &= (1 + r_{t-1})\tilde{A}_{it-1} + H_{it}\Lambda_{it} \end{aligned} \quad (\text{A.15})$$

Monetary and Fiscal Policy

Monetary Policy The central bank implements monetary policy by directly setting the nominal interest rate i_t on the nominal bonds (i.e. "reserves") held by banks. The monetary authority adjusts the nominal interest rate in response to inflation deviations, while aiming to smooth interest rates across time. The central bank obeys a Taylor rule with interest-rate smoothing

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)[r + \phi^\pi \pi_t] \quad (\text{A.16})$$

where ϕ^p is the responsiveness of monetary policy, and ρ_i is the degree of interest rate smoothing.

Fiscal Policy The flow budget constraint of the government is

$$G_t + (1 + r_{t-1})B_{t-1} = T_t + B_t \quad (\text{A.17})$$

The fiscal authority adjusts taxes in response to government spending and debt deviations. Specifically, the fiscal authority follows a Leeper rule with tax smoothing

$$dT_t = \phi^g dG_t + \phi^b dB_t \quad (\text{A.18})$$

where dG_t , dT_t , and dB_t are the change in government expenditure, taxes, and the real debt, $\phi^g > 0$ is the responsiveness of taxes to government spending, and $\phi^b > 0$ is the responsiveness of taxes to debt. In the baseline calibration we set this smoothing parameter to zero, but in alternative calibrations we allow it to be non-zero.

Market Clearing

Market clearing in the goods market implies that

$$C_t + G_t + X_t + \Psi_t^p + \Xi_t + \zeta(1 + r_{t-1})A_{t-1} = Y_t \quad (\text{A.19})$$

where $C_t = \int C_{it} di$ is aggregate consumption, Ψ_t^p are the aggregate nominal adjustment costs, Ξ_t are the aggregate real adjustment costs, and $\zeta(1 + r_{t-1})A_{t-1}$ are the costs of holding liquid assets, where we have used that $A_t = \int A_{it} di$. Market clearing in financial markets implies

$$\begin{aligned} \int (V_{it} + \tilde{V}_{it}) di &= 1 \\ \int B_{it}^n di &= 0 \\ \int (B_{it} + \tilde{B}_{it}) di &= B_t \\ \int K_{jt} dj &= K_t \end{aligned} \quad (\text{A.20})$$

Walras's Law Aggregating budget constraints across households, we have

$$\begin{aligned} \int C_{it} di + \int (B_{it} + \tilde{B}_{it}) di + P_t^y \int (V_{it} + \tilde{V}_{it}) di &= \int Z_{it} di \\ + (P_t^Y + D_t) \int (V_{it-1} + \tilde{V}_{it-1}) di + (1 + r_{t-1}) \int (B_{it-1} + \tilde{B}_{it-1}) di \\ - \zeta(1 + r_{t-1}) \int A_{it-1} di \end{aligned}$$

which, given market clearing, means

$$C_t + B_t = Z_t + D_t + (1 + r_{t-1})B_{t-1} - \zeta(1 + r_{t-1})A_{t-1} \quad (\text{A.21})$$

Combining this expression with the government budget constraint, and the definition of post tax labor income $Z_t = \frac{W_t}{P_t} N_t - T_t$, and the definition of the dividend D_t , we arrive at the goods market clearing condition $C_t + X_t + G_t + \Psi_t^p + \Xi_t + \zeta(1 + r_{t-1})A_{t-1} = Y_t$.

Equilibrium

Table A.1: Variable Names and Symbols

Name	Symbol
Output	Y_t
Consumption	C_t
Virtual consumption	C_t^*
Government expenditure	G_t
Aggregate taxes	T_t
Post-tax income	Z_t
Capital stock	K_t
Investment	X_t
Liquid assets	A_t
Illiquid assets	$\tilde{A}_t$
Nominal wage	W_t
Price level	P_t
Gross price inflation	Π_t
Gross wage inflation	Π_t^w
Real interest on capital	R_t^k
Real interest rate	R_t
Nominal interest rate	I_t
Rental rate of capital	Q_t
Marginal cost	μ_t
Capital adjustment cost	Ξ_t^k
Investment adjustment cost	Ξ_t^x
Price adjustment cost	Ψ_t^p
Wage adjustment cost	Ψ_t^w

Variables

Equations Production

$$Y_t = \Theta K_{t-1}^\alpha N_t^{1-\alpha} \quad (\text{A.22})$$

Evolution of marginal cost

$$(1 - \alpha)\mu_t Y_t = \frac{W_t}{P_t} N_t \quad (\text{A.23})$$

New Keynesian Phillips Curve

$$\begin{aligned} \left(\frac{\Pi_t}{(\Pi_{t-1})^{\iota^p}(\Pi)^{1-\iota^p}} - 1\right) \frac{\Pi_t}{(\Pi_{t-1})^{\iota^p}(\Pi)^{1-\iota^p}} &= \frac{\nu^p}{\psi^p} Y_t [\mu_t - \mu] \\ &+ \frac{1}{1+r_t} \left(\frac{\Pi_{t+1}}{(\Pi_t)^{\iota^p}(\Pi)^{1-\iota^p}} - 1\right) \frac{\Pi_{t+1}}{(\Pi_t)^{\iota^p}(\Pi)^{1-\iota^p}} \end{aligned} \quad (\text{A.24})$$

The Q-theory of investment

$$\begin{aligned} Q_t &= \frac{1}{1+r_t} [\mu_{t+1} F_{Kt+1} + (1-\delta) Q_{t+1}] \\ (g_{Xt} - 1) g_{Xt} &= \frac{X_t}{\xi} [Q_t - 1] + \frac{1}{1+r_t} (g_{Xt+1} - 1) g_{Xt+1} \end{aligned} \quad (\text{A.25})$$

Adjustment costs

$$\begin{aligned} \Xi_t &= \frac{\xi}{2} \left(\frac{X_t}{X_{t-1}} - 1\right)^2 \\ \Psi_t^p &= \frac{\psi^p}{2} \left(\frac{\Pi_t}{(\Pi_{t-1})^{\iota^p}(\Pi)^{1-\iota^p}} - 1\right)^2 \\ \Psi_t^w &= \frac{\psi^w}{2} \left(\frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w}(\Pi^w)^{1-\iota^w}} - 1\right)^2 \end{aligned} \quad (\text{A.26})$$

Dividend

$$D_t^y = Y_t - \frac{W_t}{P_t} N_t - X_t - \Xi_t - \Psi_t^p. \quad (\text{A.27})$$

Asset market clearing and no arbitrage conditions

$$\begin{aligned} A_t &= B_t + P_t^y \\ P_t^y &= \frac{1}{1+r_t} (D_t + P_{t+1}^y) \end{aligned} \quad (\text{A.28})$$

Euler equation

$$u'(C_t(\epsilon, A_-)) \geq \beta(1+r_t) \mathbf{E}[u'(C_{t+1}(\epsilon', A)) | \epsilon]. \quad (\text{A.29})$$

The Wage Phillips curve

$$\begin{aligned} \left(\frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w}(\Pi^w)^{1-\iota^w}} - 1\right) \frac{\Pi_t^w}{(\Pi_{t-1}^w)^{\iota^w}(\Pi^w)^{1-\iota^w}} &= \frac{\nu^w}{\psi^w} [N_t v'(N_t) - \frac{\nu^w - 1}{\nu^w} (1-\theta) Z_t u'(C_t^*)] \\ &+ \beta \left(\frac{\Pi_{t+1}^w}{(\Pi_t^w)^{\iota^w}(\Pi^w)^{1-\iota^w}} - 1\right) \frac{\Pi_{t+1}^w}{(\Pi_t^w)^{\iota^w}(\Pi^w)^{1-\iota^w}} \end{aligned} \quad (\text{A.30})$$

Wage consistency condition

$$\Pi_t^w = \Pi_t \frac{W_t/P_t}{W_{t-1}/P_{t-1}} \quad (\text{A.31})$$

Goods market clearing

$$C_t + G_t + X_t + \Psi_t + \Xi_t = Y_t \quad (\text{A.32})$$

where $\Psi_t = \Psi_t^p + \Psi_t^w$ are the aggregate nominal adjustment costs, and Ξ_t are the aggregate real adjustment costs.

Fisher Equation

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}} \quad (\text{A.33})$$

Taylor Rule

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)[r + \phi^\pi \pi_t] \quad (\text{A.34})$$

Government budget constraint

$$B_t = G_t + (1 + r_{t-1})B_{t-1} - T_t \quad (\text{A.35})$$

Fiscal rule

$$dT_t = \phi^g dG_t + \phi^b dB_t \quad (\text{A.36})$$

Steady State

$$Y = \Theta K^\alpha N^{1-\alpha} \quad (\text{A.37})$$

$$\mu = \frac{\nu^p - 1}{\nu^p} \quad (\text{A.38})$$

$$\mu = (1 - \alpha) \frac{WN}{PY} \quad (\text{A.39})$$

$$r + \delta = \alpha \mu \frac{Y}{K} \quad (\text{A.40})$$

$$Q = 1 \quad (\text{A.41})$$

$$X = \delta K \quad (\text{A.42})$$

$$D = (1 - \mu)Y \quad (\text{A.43})$$

$$1 \geq (1 + r)\beta \quad (\text{A.44})$$

$$Nv'(N) = \mu^w(1 - \theta)(Y - T)u'(C^*) \quad (\text{A.45})$$

$$\Pi^w = \Pi \quad (\text{A.46})$$

$$1 + r = \frac{1 + i}{1 + \pi} \quad (\text{A.47})$$

$$T = G + rB \quad (\text{A.48})$$

$$Z = \frac{W}{P}N - T \quad (\text{A.49})$$

$$C + G + X + \zeta(1 + r)A = Y \quad (\text{A.50})$$

$$A = B + P^y \quad (\text{A.51})$$

$$P^y = \frac{(1 - \mu)}{r}Y \quad (\text{A.52})$$

$$C = \mathcal{C}(\{Z, r\}, P^y + D; \beta) \quad (\text{A.53})$$

$$C^* = \mathcal{C}^*(\{Z, r\}, P^y + D; \beta) \quad (\text{A.54})$$


$$g_X = 0 \quad (\text{A.55})$$

$$\Xi = 0 \quad (\text{A.56})$$

A.4 Additional Model Results

Role of Monetary Policy


Figure A.14: Role of Monetary Policy



Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the hawkish monetary policy ($\phi_\pi = 1.5$) model's response to the government spending shock in the first panel.

Role of Fiscal Policy

Figure A.15: Role of Fiscal Policy



Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the tax-financed fiscal policy ($\phi_g = 1$) model's response to the government spending shock in the first panel.

Alternative Models

Analytical Models The representative agent New Keynesian (RANK) model and the two-agent bonds-in-the-utility function New Keynesian (TABU) model are different calibrations of the same underlying analytical model. There are two types of agents: constrained and unconstrained. The $1 - \mu$ mass of unconstrained agents earn income $Z_t = \frac{W_t}{P_t} N_t - T_t$ and have access to a one-period risk-free bond, which pays interest r_t . Given their initial asset position A_{-1} the unconstrained households solve

$$\max_{\{C_{ut}, N_{ut}, A_{ut}\}_{t=0}^{\infty}} \frac{C_{ut}^{1-1/\sigma} - 1}{1 - 1/\sigma} - \frac{N_{ut}^{1+1/\eta}}{1 + 1/\eta} + w(A_{ut})$$

$$C_{ut} + A_{ut} = Z_t + (1 + r_{t-1})A_{ut-1}$$

where C_{ut} , N_{ut} and A_{ut} are the unconstrained consumption, labor supply, and asset position, σ is the intertemporal elasticity of substitution, η is the Frisch elasticity, and $w(\cdot)$ is the household

demand for liquidity. This demand for liquidity, $w(\cdot)$, is weakly concave. The remaining μ portion of the population is hand-to-mouth, with no access to financial markets. These constrained households maximize:

$$\max_{C_{ct}, N_{ct}} \frac{C_{ct}^{1-1/\sigma} - 1}{1 - 1/\sigma} - \frac{N_{ct}^{1+1/\eta}}{1 + 1/\eta}$$

$$C_{ct} = Z_t$$

each period t . Aggregate consumption is $C_t = (1 - \mu)C_{ut} + \mu C_{ct}$, aggregate labor supply is $N_t = (1 - \mu)N_{ut} + \mu N_{ct}$, and the aggregate assets are $A_t = (1 - \mu)A_{ut}$.

The Euler equation of this model, $C_{ut}^{-1/\sigma} = \beta(1 + r_t)C_{ut+1}^{-1/\sigma} + w'(A_{ut})$, implies that in steady-state

$$\beta(1 + r) = 1 - \frac{w'(A_{ut})}{(C_u)^{-1/\sigma}} \quad (\text{A.57})$$

the interest rate is below the inverse of the discount factor, i.e., $r < \beta^{-1} - 1$, if the marginal value of liquidity $w'(A_{ut})$ is positive. Moreover, Auclert et al. (2024) (section 4) show that the intertemporal marginal propensity to consume at the steady-state is

$$\mathbf{M} = \mu \mathbf{I} + (1 - \mu) \mathbf{M}^{BU}$$

$$\text{s.t. } M_{t0}^{BU} = (1 - \frac{\lambda}{1 + r}) \lambda^t, \forall t \geq 0 \quad (\text{A.58})$$

where λ is a parameter that maps to the shape of the liquidity demand function $w(\cdot)$, and M_{t0}^{BU} is the first column of $\mathbf{M}^{BU}$. See Appendix D.2 of Auclert et al. (2024) for more details.


Beyond the usual household preferences, i.e., σ, η, β , this model is characterized by the proportion of hand-to-mouth agents μ and the degree of liquidity demand λ . When there are no hand-to-mouth agents ($\mu = 0$) and no demand for liquidity ($\lambda = 0$), this model collapses to the RANK formulation with marginal propensity to consume

$$M_{t0}^{RA} = (1 - \beta)\beta^t, \forall t \geq 0 \quad (\text{A.59})$$

in the first column of the iMPC matrix $\mathbf{M}$. On the other hand, if there were hand-to-mouth households ($\mu > 0$) but not liquidity demand ($\lambda = 0$) the model would collapse to standard TANK model (e.g. Galí et al. (2007)).


In the main text we compare our baseline model to a RANK ($\mu = 0, \lambda = 0$) variation and a TA BU ($\mu > 0, \lambda > 0$) variation. To calibrate the RANK model, we use the discount factor β to discipline the iMPC matrix. To calibrate the TABU model we match the first two iMPCs, $M_{00}^{TABU}, M_{10}^{TABU}$ to match the first and second iMPC from a one-asset HANK model, i.e. $M_{00}^{HA1}, M_{10}^{HA1}$. This sets $\beta = .965, \mu = .138, \lambda = .946$.

Figure A.16: Alternative Model: RANK




Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the RANK model's response to the government spending shock in the first panel.

Figure A.17: Alternative Model: TABU




Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the TABU model's response to the government spending shock in the first panel.

Figure A.18: Alternative Model: Flexible Prices




Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the flexible price model's response.

Figure A.19: Alternative Model: Flexible Wages




Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the flexible wage model's response to the government spending shock in the first panel.

Figure A.20: Alternative Model: Capital Adjustment




Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the capital adjustment model's response to the government spending shock in the first panel.

Figure A.21: Alternative Model: No Price Indexation




Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the no indexation model's response to the government spending shock in the first panel.

Figure A.22: Alternative Model: Policy Smoothing



Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the no policy smoothing model's response to the government spending shock in the first panel.

Figure A.23: Alternative Model: larger Frisch elasticity ($\eta^{-1} = 1$)



Notes: This figure plots the estimated and model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to a defense spending surprise, normalized so the defense response peaks at one percent of GDP. The sold blue lines are our estimates of β_h in Equation (1.1) for different horizons h . The dark and light shaded areas are the 68% and 95% confidence intervals. The red dashed line is the baseline model's response to the government spending shock in the first panel. The green dot-dashed line is the larger Frisch elasticity model's response to the government spending shock in the first panel.

Additional Results




Figure A.24: Effect of an AR(1) Defense Spending Shock

Notes: This figure plots the model implied responses out output, consumption, investment, taxes, nominal interest rates, inflation, and real wages to an AR(1) government spending surprise. The red dashed line is the baseline model's response to the government spending shock in the first panel.

Table A.2: Alternative Calibrations.

Model	Hawkishness (ϕ^π)	Interest-smoothing (ρ_i)	Fiscal reaction to G (ϕ^g)	Fiscal reaction to B (ϕ^b)	Investment elasticity (ξ^I)
Baseline	1.206	0.164	.511	0.0	1.694
RANK	1.020	0.0	.392	0.0	1.151
TABU	1.158	0.358	.429	0.0	.970
Flexible price	1.046	0.036	.147	0.0	1.313
Flexible wage	1.045	0.972	1.00	0.0	.389
Capital adjustment	1.106	0.069	.099	0.0	0.0
No indexation	1.423	0.173	.033	0.0	1.273
Without policy smoothing	1.200	0.000	.524	0.0	.800
Larger Frisch	1.322	0.152	.465	0.0	1.673

Appendix B

Appendix to Monopsony with Recruiting

B.1 Estimation Details

We begin by explaining our estimation strategy using minimum distance. Define Z_t to be a vector of a finite sample of $z_{j,t}$ shocks in the model.

In response to a shock, the firm's path for the new hire wages, hiring rate, and employment are deterministic and satisfy (2.13). We assume we are log-linearizing around a point where $\varepsilon_{R,w} = -\varepsilon_{S,w} = 2$. We also know that in the linearized model, each variable is a linear function of shocks. For example, $\tilde{w}_{j,t} = \beta_0^w z_{j,t}$. Thus we can substitute firm-level endogenous variables with the firm j -specific shock $z_{j,t}$ times a variable-specific slope:

$$\beta_0^w z_{j,t} = \frac{1}{2\Omega} \left((\chi\beta_0^h z_{j,t} + \sigma\beta_{t-1}^N z_{j,t}) + S \sum_{\kappa=0}^{\infty} (1-S)^\kappa \chi\beta_{\kappa+1}^h z_{j,t} + \sigma\beta_\kappa^N z_{j,t} - (1+\chi)2\beta_{\kappa+1}^w z_{j,t} \right),$$

If this equation holds with equality for any individual firm j , summing over j on both sides of the equation also holds with equality, so we can cancel the sum of shocks, $\sum_j z_{j,t}$:

$$\begin{aligned} \sum_j \beta_0^w z_{j,t} &= \sum_j \frac{1}{2\Omega} \left((\chi\beta_0^h z_{j,t} + \sigma\beta_{t-1}^N z_{j,t}) + S \sum_{\kappa=0}^{\infty} (1-S)^\kappa \chi\beta_{\kappa+1}^h z_{j,t} + \sigma\beta_\kappa^N z_{j,t} - (1+\chi)2\beta_{\kappa+1}^w z_{j,t} \right) \\ \beta_0^w &= \frac{1}{2\Omega} \left((\chi\beta_0^h + \sigma\beta_{t-1}^N) + S \sum_{\kappa=0}^{\infty} (1-S)^\kappa \chi\beta_{\kappa+1}^h + \sigma\beta_\kappa^N - (1+\chi)2\beta_{\kappa+1}^w \right). \end{aligned} \quad (\text{B.1})$$

Suppose we did not know χ and σ , but we had simulated data from the linearized model. Define $\tilde{W}_t$ as the vector of values $\tilde{w}_{j,t}$ for the shocked firms, and $\tilde{W}_{t+1}$ as the vector of values $\tilde{w}_{j,t+1}$ for the shocked firms, $\tilde{\mathcal{H}}_t$ as the vector of log deviations of hiring rates of firms shocked in period t , and $\tilde{\mathcal{N}}_t$ as the vector of log deviations of employment levels of firms shocked in period t , and so on. Let $\tilde{\mathcal{Y}}_t \equiv [\tilde{W}_t \ \tilde{W}_{t+1} \dots \ \tilde{\mathcal{H}}_t \ \tilde{\mathcal{H}}_{t+1} \dots \ \tilde{\mathcal{N}}_t \ \tilde{\mathcal{N}}_{t+1} \dots]$ be a matrix of log deviations of steady state of firms shocked in period t . The slope coefficients $\pi \equiv [\beta_0^w \ \beta_1^w \dots \ \beta_0^h \ \beta_1^h \dots \ \beta_0^N \ \beta_1^N \dots]$ could be recovered from regressing firm outcomes $\tilde{\mathcal{Y}}_t$ on Z_t :

$$\pi = (Z_t' Z_t)^{-1} (Z_t' \tilde{\mathcal{Y}}_t).$$

Next we can use these coefficients to uncover the true values of χ and σ . Equation (B.1) can be rolled forward so that we have two conditions that must hold:

$$\begin{aligned}\beta_0^w + (1 + \chi) \frac{1}{\Omega} S \sum_{\kappa=0}^{\infty} (1 - S)^\kappa \beta_{\kappa+1}^w &= \frac{\chi}{2\Omega} \left(\beta_0^h + S \sum_{\kappa=0}^{\infty} (1 - S)^\kappa \beta_{\kappa+1}^h \right) + \frac{\sigma}{2\Omega} \left(\beta_{-1}^N + S \sum_{\kappa=0}^{\infty} (1 - S)^\kappa \beta_\kappa^N \right) \\ \beta_1^w + (1 + \chi) \frac{1}{\Omega} S \sum_{\kappa=1}^{\infty} (1 - S)^\kappa \beta_{\kappa+1}^w &= \frac{\chi}{2\Omega} \left(\beta_1^h + S \sum_{\kappa=1}^{\infty} (1 - S)^\kappa \beta_{\kappa+1}^h \right) + \frac{\sigma}{2\Omega} \left(\beta_0^N + S \sum_{\kappa=1}^{\infty} (1 - S)^\kappa \beta_\kappa^N \right).\end{aligned}$$

We have two equations and two unknowns, and χ and σ can be solved for exactly.

Next, we estimate (2.14) and (2.15) in a stacked regression and recover the vector of regression coefficients $\hat{\pi} \equiv [\hat{\beta}_0^w \ \hat{\beta}_1^w \ \hat{\beta}_2^w \ \hat{\beta}_0^h/\bar{h} \ \hat{\beta}_1^h/\bar{h} \ \hat{\beta}_2^h/\bar{h} \ \hat{\beta}_0^N \ \hat{\beta}_1^N \ \hat{\beta}_2^N]$ from Table 2.4 of the main text and the following variance-covariance matrix of regression coefficients in Table B.1, truncating the regression to have three years of horizon.

Table B.1: Variance Covariance Matrix $\hat{\Sigma}$

$\hat{\beta}_0^h$	$\hat{\beta}_1^h$	$\hat{\beta}_2^h$	$\hat{\beta}_0^N$	$\hat{\beta}_1^N$	$\hat{\beta}_2^N$	$\hat{\beta}_0^w$	$\hat{\beta}_1^w$	$\hat{\beta}_2^w$
0.0026	0.0015	0.0013	0.0006	0.0010	0.0010	0.0000	0.0000	0.0000
0.0015	0.0028	0.0015	-0.0003	0.0007	0.0010	0.0000	0.0000	0.0000
0.0013	0.0015	0.0025	-0.0003	-0.0001	0.0006	-0.0001	0.0000	0.0000
0.0006	-0.0003	-0.0003	0.0013	0.0013	0.0011	0.0000	0.0000	0.0000
0.0010	0.0007	-0.0001	0.0013	0.0024	0.0023	0.0001	0.0000	0.0000
0.0010	0.0010	0.0006	0.0011	0.0023	0.0032	0.0000	0.0001	0.0000
0.0000	0.0000	-0.0001	0.0000	0.0001	0.0000	0.0016	-0.0001	-0.0001
0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	-0.0001	0.0016	-0.0002
0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	-0.0001	-0.0002	0.0015

Notes: This table reports the variance-covariance matrix of regression coefficients for the joint regressions of (2.14) and (2.15) where the dependent variables are switcher wage growth, firm hiring rate, and log firm employment, in years 0, 1, 2 in response to an export demand shock.

Next, we construct f from equation (2.17) in the main text. To map our annual regression estimates into a monthly model, we assume that hires occur evenly across months within the year, and plug in annual estimates for each year's corresponding 12 months. The formula for this can be seen in equation (B.2). We then plug in our point estimates $\hat{\pi}$ into f , and search over values of σ and χ to minimize


$$\min_{\sigma, \chi} f(\hat{\pi}, \sigma, \chi, \Theta)' f(\hat{\pi}, \sigma, \chi, \Theta).$$

with $\Theta = \{S, \varepsilon_{R,w}, \varepsilon_{S,w}, \xi_{R,w}, \xi_{S,w}\}$, and where f is written out in equation (B.2) with $\Omega = 1 + \frac{\varepsilon_{R,w}}{\varepsilon_{R,w} - \varepsilon_{S,w}} (-\xi_{R,w} - \varepsilon_{S,w} + (1 + \chi)\varepsilon_{R,w}) + \frac{-\varepsilon_{S,w}}{\varepsilon_{R,w} - \varepsilon_{S,w}} (-\xi_{S,w} - \varepsilon_{S,w})$. When minimizing (2.17), we additionally include a small quadratic penalty term: $(12 \times .0025 \times (\chi - 0))^2 + (12 \times .0025 \times (\sigma - 0))^2$,

which helps rule out economically implausible estimates. We set $\varepsilon_{R,w} = -\varepsilon_{S,w} = -\xi_{R,w} = \xi_{S,w} = 2$ and choose $S = 0.05$. We use the Matlab function `lsqnonlin`. Our estimate for the inverse labor supply elasticity is then $\hat{\varepsilon}_{w,N} = \frac{\hat{\sigma}}{1+\mathcal{E}(1+\hat{\chi})}$.

$$f(\hat{\pi}, \sigma, \chi, \Theta) = \underbrace{\begin{bmatrix} \frac{-(1+\chi)S \sum_{\kappa=0}^{10} (1-S)^\kappa}{\Omega} - 1 & \frac{-(1+\chi)S \sum_{\kappa=0}^9 (1-S)^\kappa}{\Omega} - 1 & \dots & 0 \\ \frac{-(1+\chi)S(1-S)^{11}}{\Omega} & \frac{-(1+\chi)S \sum_{\kappa=10}^{11} (1-S)^\kappa}{\Omega} & \dots & -1 \\ 0 & 0 & \dots & \frac{-(1+\chi)S \sum_{\kappa=0}^{11} (1-S)^\kappa}{\Omega} \\ \frac{\chi(1+S \sum_{\kappa=0}^{10} (1-S)^\kappa)}{2\Omega} & \frac{\chi(1+S \sum_{\kappa=0}^9 (1-S)^\kappa)}{2\Omega} & \dots & 0 \\ \frac{\chi S(1-S)^{11}}{2\Omega} & \frac{\chi S \sum_{\kappa=10}^{11} (1-S)^\kappa}{2\Omega} & \dots & \frac{\chi}{2\Omega} \\ 0 & 0 & \dots & \frac{\chi S \sum_{\kappa=0}^{11} (1-S)^\kappa}{2\Omega} \\ \frac{\sigma(S \sum_{\kappa=0}^{11} (1-S)^\kappa)}{2\Omega} & \frac{\sigma(1+S \sum_{\kappa=0}^{10} (1-S)^\kappa)}{2\Omega} & \dots & 0 \\ 0 & \frac{\sigma S \sum_{\kappa=11}^{11} (1-S)^\kappa}{2\Omega} & \dots & \frac{\sigma}{2\Omega} \\ 0 & 0 & \dots & \frac{\sigma S \sum_{\kappa=0}^{11} (1-S)^\kappa}{2\Omega} \end{bmatrix}'}_{A_{24 \times 9}} \times \underbrace{\begin{bmatrix} \hat{\beta}_0^w \\ \hat{\beta}_1^w \\ \hat{\beta}_2^w \\ \hat{\beta}_0^h / \bar{h} \\ \hat{\beta}_1^h / \bar{h} \\ \hat{\beta}_2^h / \bar{h} \\ \hat{\beta}_0^N \\ \hat{\beta}_1^N \\ \hat{\beta}_2^N \end{bmatrix}}_{\hat{\pi}_{9 \times 1}} \quad (\text{B.2})$$

Lastly, we characterize the uncertainty of our estimates by sampling from a joint normal distribution $\sim \mathcal{N}(\hat{\pi}, \hat{\Sigma})$. For each draw, we re-estimate σ and χ using the same procedure and compute an inverse labor supply elasticity $\hat{\varepsilon}_{w,N} = \frac{\hat{\sigma}}{1+\mathcal{E}(1+\hat{\chi})}$ for each draw. We plot the distribution of $\hat{\chi}$ and $\hat{\varepsilon}_{w,N}$ in Figure B.1 and report the 10th and 90th percentiles of $\hat{\chi}$, $\hat{\sigma}$, and $\hat{\varepsilon}_{w,N}$ in Table 2.5 in the main text.

Figure B.1: Distribution Estimates of χ and Inverse Labor Supply Elasticity $\varepsilon_{w,N}$ 

Notes: This figure shows the distribution of estimates for the parameter χ and firm specific inverse labor supply elasticity $\varepsilon_{w,N}$. Estimates are obtained by resampling from a multivariate normal distributed $\mathcal{N}(\hat{\pi}, \hat{\Sigma})$, and choosing σ and χ to minimize equation (2.17), where f is defined in equation (B.2).

B.2 Switcher Wage Levels


We report how the raw log wage level of job switchers changes in response to a period 0 export demand shock. That is, we run the regression in equation (2.15), except the outcome variable is the average log wage of arriving workers, both in the year they arrive and their wages prior to switching. Let $\log w_{j,t+\nu}^{switcher} \equiv \frac{1}{|switchers_{j,t}|} \sum_{i \in switchers_{j,t}} \log w_{i,j,t+\nu}$ be the average log wage of workers who arrive in firm j in period $t+\nu$, and let $\log w_{j,t+\nu-2}^{switcher} \equiv \frac{1}{|switchers_{j,t}|} \sum_{i \in switchers_{j,t}} \log w_{i,k,t+\nu-2}$ be the average prior log wage of workers who arrive in firm j in period $t+\nu$. We estimate the following regressions:

$$\log w_{j,t+\nu-2}^{switcher} = \beta_{\nu}^w z_{j,t}^{idiosyncratic} + \delta_{\nu}^w X_{j,t} + u_{j,t+\nu}^w \quad (\text{B.3})$$

$$\log w_{j,t+\nu}^{switcher} = \beta_{\nu}^w z_{j,t}^{idiosyncratic} + \delta_{\nu}^w X_{j,t} + u_{j,t+\nu}^w \quad (\text{B.4})$$

Figure B.2a shows that prior log wage of workers did not change much around the timing of export demand shocks, suggesting that the composition of workers that firms did not change much in response to export demand shocks. While more positively shocked firms may have hired workers with slightly lower prior wages on impact, the estimate is very noisy. Figure B.2b shows how the raw log wage of newly hired job switchers changes around an export demand shock. Wages are higher in on impact, but not statistically significantly. The coefficient estimates are noisy, suggesting that differencing switcher wages removes highly variable person effects.


Figure B.2: Raw Log Wage of Switchers



Notes: This figure reports the coefficient of a period t export demand shock on the raw two-year prior log wage level and the raw log wage level and of different cohorts of job switchers arriving in period $t + \nu$. Each point corresponds to a different cohort of hired workers. For example, in horizon $\nu = 0$, these are workers who arrive in the firm in the same period t as the export demand shock, and their prior wage reflects the log wage the worker earned two years prior to the export demand shock ($t - 2$).

B.3 Robustness: Measures of Switcher Wage Growth

Figure B.3: Robustness: Employment, Hiring, and Switcher Wage Growth Response



Notes: This figure reports robustness to Figure 2.3a. In panel (a), switcher wages are constructed by first residualized switcher wage growth on mincer controls, and then taking the average residualized switcher wage at the firm-year level. In panel (b), switcher wage growth residuals are constructed taking into the AKM fixed effect of the firm that that workers leave, in addition to Mincer controls. In panel (c), switcher wage growth residuals are constructed taking in account Mincer Controls, the prior firm AKM fixed effect, and a variable that estimates the average wage change when a worker switches between two occupations. Panel (d) uses the same measure as Panel (b), but using only firms whose log full time equivalent employment was less than 4 (approximately 55 full-time equivalent workers) in the year prior to the shock. Panels (e) and (f) construct new hire wage growth using only workers who switch from non-exporting firms.

B.4 The Firm Size Wage Premium

In the main text, we showed that when manufacturing firms are subject to export demand shocks, these firms do not pay any higher wages to new hires than the firms did prior to the shock. In this section, we estimate the effect of firm size on workers' wages for Danish firms economy-wide between 2008-2019. We regress workers' log wages on the log of a firm's average employment level in full-time equivalent (FTE) workers, averaged over 2008-2019. We do this for two separate samples of firms: first restricting to firms with 5 or more full-time equivalent workers, and second, restricting to firms with 25 or more FTEs:

$$\log w_{ijt} = \alpha_i^{\bar{N}} + \beta^{\bar{N}} \overline{\log N_j} + \xi_{j,t}^{\bar{N}} + e_{ijt}^{\bar{N}},$$

where $\alpha_i^{\bar{N}}$ is a worker fixed effect, $\xi_{j,t}^{\bar{N}}$ is a firm-level industry-by-year fixed effect. We cluster the standard errors at the firm level.

Table B.2: Firm Size Wage Premium

	(1)	(2)
	$\log w_{ijt}$	$\log w_{ijt}$
$\overline{\log N_j}$	0.00531 (0.000543)	0.00146 (0.000706)
Sample	$\exp(\overline{\log N_j}) \geq 5$	$\exp(\overline{\log N_j}) \geq 25$
Observations	17,389,575	14,475,115
No. Firms	53,200	9,885

Notes: This table reports regressions estimating the effect of firm size on workers' log wages that includes worker fixed effects and industry-by-year fixed effects. For firms with an average employment of at least 5 full time equivalent workers, the elasticity of worker wages to firm size is 0.005. For firms with at least 25 workers, the elasticity of wages to firm size is 0.001. Standard errors are clustered at the firm level.

Table B.2 reports the results. For firms with 5 or more workers, the elasticity of workers' wages with respect to firm size is 0.005, or 0.5%. Focusing only on firms with on average 25 workers or more, the elasticity drops to .001, or around 0.1%. These estimates are similar to the elasticities of firm wage effects on log firm size estimated in the United States (Bloom et al., 2018), particularly for medium-sized and large firms, suggesting that our results are not solely a feature of the high union density in Denmark.

B.5 Worker's Problem

Worker's Value Function The worker gets utility from wages and draws a utility shock when given a choice between two firms. Let ν_k be the share of vacancies in the market posting wage $w = w_k$.

$$\begin{aligned} \mathcal{V}_{j,s,t} = & \frac{\eta}{\eta-1} \left(w_{j,s,t}^{\frac{\eta-1}{\eta}} - 1 \right) + \beta_w s_0 \mathcal{V}_U + \\ & \beta_w (1-s_0) \mathbb{E}_t \left[\lambda_{EE} f(\theta) \int_k \nu_k \gamma^{-1} \log \left(\exp(\gamma \mathcal{V}_{j,s,t+1}) + \exp(\gamma \mathcal{V}_{k,t+1}) \right) dk + \right. \\ & \left. (1 - \lambda_{EE} f(\theta)) \max\{\mathcal{V}_{j,s,t+1}, \mathcal{V}_U\} \right], \end{aligned} \quad (\text{B.5})$$

How Worker's Value Function Changes with a Firm's Wage When deriving how a change in the firm's wage policy affects worker values, we will assume that workers believe that the wage change is permanent. We will consider the case of no wage dispersion and $\mathcal{V}_U < \mathcal{V}_j$:

$$\frac{\partial \mathcal{V}_j}{\partial w_j} = w_j^{-\frac{1}{\eta}} + \beta_w (1-s_0) \left(f(\theta) \lambda_{EE} \int v_k \left(\frac{\exp(\gamma \mathcal{V}_j)}{\exp(\gamma \mathcal{V}_j) + \exp(\gamma \mathcal{V}_k)} \right) dk + (1 - f(\theta) \lambda_{EE}) \right) \frac{\partial \mathcal{V}_j}{\partial w_j}.$$

We can substitute in the expression for the separation rate from equation (2.3) (under the assumption that workers are not quitting into unemployment or exogenously separating). Under symmetry, values are equal across firms $\mathcal{V}_j = \mathcal{V}_k$, firm employment and firm vacancy shares are equal $\phi_k = v_k$, and if there are no exogenous separations, then $S = \lambda_{EE} f(\theta)/2$.

$$\begin{aligned} \frac{\partial \mathcal{V}_j}{\partial w_j} = & w_j^{-\frac{1}{\eta}} + \beta_w \left(S(w_j, w_{-j}, \theta) + (1 - f(\theta) \lambda_{EE}) \right) \frac{\partial \mathcal{V}_j}{\partial w_j} \\ = & w_j^{-\frac{1}{\eta}} + \beta_w \left(S(w_j, w_{-j}, \theta) + (1 - 2S(w_j, w_{-j}, \theta)) \right) \frac{\partial \mathcal{V}_j}{\partial w_j} \\ = & w_j^{-\frac{1}{\eta}} + \beta_w (1 - S(w_j, w_{-j}, \theta)) \frac{\partial \mathcal{V}_j}{\partial w_j}, \end{aligned}$$

resulting in

$$\frac{\partial \mathcal{V}_j}{\partial w_j} = \frac{w_j^{-\frac{1}{\eta}}}{(1 - \beta_w (1 - S))}. \quad (\text{B.6})$$

B.6 Firms' Dynamic Problem

This appendix section derives the decomposition of marginal product in Section 2.2, quantifies the effect of non-zero discounting, and derives the inverse labor supply elasticity in Section 2.2.

If the firm is operating within a stationary environment, (i.e., $S(w)$ and $R(w)$ are not changing), the firm maximizes the present discounted value of profits. Ignoring a firm subscript, the firm's problem is:

$$\begin{aligned} & \max_{\{V_t, K_t, N_t, \{N_{s,t}\}_{s \leq t}, \{w_{s,t}\}_{s \leq t}\}} \left\{ \sum_{t=0}^{\infty} \beta_f^t \left(A_t (K_t^\alpha N_t^{1-\alpha})^{\frac{\epsilon-1}{\epsilon}} - \sum_{s=-\infty}^t w_{s,t} N_{s,t} - r^K K_t - c \left(\frac{V_t}{N_{t-1}} \right)^\chi N_{t-1}^\sigma V_t \right) \right\}_{t=0}^{\infty} \quad (\text{B.7}) \end{aligned}$$

subject to

$$\begin{aligned} \text{Recruiting:} \quad & N_{t,t} = R(w_{t,t}) V_t \quad \text{for all } t \\ \text{Retention:} \quad & N_{s,t} = (1 - S(w_{s,t})) N_{s,t-1} \quad \text{for all } t \text{ and } s < t \\ \text{Employment is sum of cohorts:} \quad & N_t = \sum_{s=-\infty}^t N_{s,t} \quad \text{for all } t \\ \text{Wage at least posted wage:} \quad & w_{s,t} \geq w_{s,s} \quad \text{for all } t \text{ and } s < t \end{aligned}$$

where the total number of workers in the firm N_t is the sum of all workers from each past cohort still remaining in the firm: $N_t = \sum_{s=-\infty}^t N_{s,t}$ which we substitute into the equation and obtain the Lagrangian and check for the fourth restriction in Appendix B.7:

$$\begin{aligned} \mathcal{L} = \sum_{t=0}^{\infty} \beta_f^t & \left[A_t \left(K_t^\alpha \sum_{s=-\infty}^t N_{s,t}^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - \sum_{s=-\infty}^t w_{s,t} N_{s,t} - c V_t^{1+\chi} \left(\sum_{s=-\infty}^{t-1} N_{s,t-1} \right)^{\sigma-\chi} \right. \\ & + \lambda_t (R(w_{t,t}) V_t - N_{t,t}) \\ & \left. + \sum_{s < t} \mu_{s,t} ((1 - S(w_{s,t})) N_{s,t-1} - N_{s,t}) \right] \end{aligned}$$

Assuming that the wage posting constraint binds, the first order conditions are:

$$\begin{aligned} \mathcal{L}_{N_{s,t}} : & MRPL_t - w_{s,t} - \lambda_t - \beta_f c (\sigma - \chi) \left(\frac{V_{t+1}}{N_t} \right)^\chi N_t^{\sigma-1} V_{t+1} + \beta_f \mu_{s,t+1} (1 - S(w_{s,t+1})) = 0 \\ \mathcal{L}_{w_{0,0}} : & - \sum_{t=0}^{\infty} \beta_f^t N_{0,t} - \sum_{t=1}^{\infty} \beta_f^t \mu_{0,t} N_{0,t-1} S'(w_{0,0}) + \lambda_0 R'(w_{0,0}) V_0 = 0 \\ \mathcal{L}_{V_t} : & -c(1 + \chi) \left(\frac{V_t}{N_{t-1}} \right)^\chi N_{t-1}^\sigma + \lambda_t R(w_{t,t}) = 0. \end{aligned}$$

Rearranging the first order condition on V_t yields, imposing steady state, and using that $\frac{V}{N} = \frac{S(w)}{R(w)}$ whereby also $\lambda = \mu$:

$$\lambda = \mu = \frac{c(1 + \chi) \left(\frac{S(w)}{R(w)} \right)^\chi}{R(w)} N^\sigma.$$

With some algebra, the first order condition on wages yields

$$w = c(1 + \chi)N^\sigma \left(\frac{S(w)}{R(w)} \right)^{1+\chi} \left(\frac{1 - \beta_f(1 - S(w))}{S} \varepsilon_{R,w} - \beta_f \varepsilon_{S,w} \right). \quad (\text{B.8})$$

We can now solve for what shares of marginal product go to wages, turnover costs, and profits in steady state. The wage share of marginal product is then

$$\frac{w}{MRPL} = \frac{(1 + \chi) \left(\frac{1 - \beta_f(1 - S(w))}{S} \varepsilon_{R,w} - \beta_f \varepsilon_{S,w} \right)}{(1 + \chi) \left(\frac{1 - \beta_f(1 - S(w))}{S} \varepsilon_{R,w} - \beta_f \varepsilon_{S,w} \right) + \left((1 + \chi) \frac{1 - \beta_f}{S(w)} + \beta_f(1 + \sigma) \right)}.$$

Recruiting costs per worker are given by $c(V/N)^{1+\chi}N^\sigma$ in steady state, so as a share of marginal product we get

$$\frac{\text{Recruiting costs per worker}}{MRPL} = \frac{1}{(1 + \chi) \left(\frac{1 - \beta_f(1 - S(w))}{S} \varepsilon_{R,w} - \beta_f \varepsilon_{S,w} \right) + \left((1 + \chi) \frac{1 - \beta_f}{S(w)} + \beta_f(1 + \sigma) \right)}$$

The labor market profit per worker is the gap between marginal product and the sum of wages and per-incumbent recruiting costs.

$$\frac{\text{Labor market profits per worker}}{MRPL} = \frac{\left((1 + \chi) \frac{1 - \beta_f}{S(w)} + \beta_f(1 + \sigma) \right) - 1}{(1 + \chi) \left(\frac{1 - \beta_f(1 - S(w))}{S} \varepsilon_{R,w} - \beta_f \varepsilon_{S,w} \right) + \left((1 + \chi) \frac{1 - \beta_f}{S(w)} + \beta_f(1 + \sigma) \right)}$$

How large are the additional terms due to discounting? At a monthly frequency, total monthly separation rates are approximately $S = 0.05$, and given an annual discount rate, β_f can be approximated to being equal to 0.995. Setting $\sigma = 0$ and $\chi = 2$, $\varepsilon_{R,w} = 2$, and $\varepsilon_{S,w} = -2$, then

$$\frac{\text{Labor market profits per worker}}{MRPL} \approx 0.021.$$

This calculation implies that under standard parameters, the additional profits collected on the margin are 2.1% of marginal product. Economically, standard time preferences are quantitatively unimportant because the discounting from the time preference rate is an order of magnitude smaller than the separation rate. Because matching with a worker requires an upfront investment with a future flow of payoffs, but the life of that flow of payoffs is affected by the separation rate, the relevant discount rate to the firm is the sum of the separation rate and time preference parameter.

Deriving the Inverse Labor Supply Elasticity in General Form To derive the inverse labor supply elasticity, we start with the optimal wage equation:

$$w = c(1 + \chi) \left(\frac{S(w)}{R(w)} \right)^{1+\chi} (\varepsilon_{R,w} - \varepsilon_{S,w}) N^\sigma.$$

Taking logs (and) reinstating the (w) yields

$$\log(w) = \log(c(1 + \chi)) + (1 + \chi) \log\left(\frac{S(w)}{R(w)}\right) + \log \mathcal{E} + \sigma \log(N).$$

where $\mathcal{E} = \varepsilon_{R,w} - \varepsilon_{S,w}$. Taking the total derivative with respect to $\log(w)$, we get

$$\begin{aligned} 1 &= (1 + \chi) (\varepsilon_{S,w} - \varepsilon_{R,w}) + \varepsilon_{\mathcal{E},w} + \sigma \varepsilon_{N,w} \\ &= -(1 + \chi) \mathcal{E} + \varepsilon_{\mathcal{E},w} + \sigma \varepsilon_{N,w}. \end{aligned}$$

Rearranging and using $\varepsilon_{w,N} = \varepsilon_{N,w}^{-1}$ yields

$$\varepsilon_{w,N} = \frac{\sigma}{1 + (1 + \chi) \mathcal{E} - \varepsilon_{\mathcal{E},w}}.$$

B.7 Proof of Proposition 1

This appendix section proves Proposition 1 in three steps. First, in Section B.7, assuming a stationary equilibrium exists, we show that the sum of recruiting and separation elasticities is equal to γ at any wage level. In Section B.7, we characterize firms' optimal wages and employment levels as functions of parameters and aggregate variables. Then in Section B.7, we prove that firms do not find it optimal to backload wages. Lastly, in Section B.7, we show that the equilibrium is unique.

Constant Sum of Recruiting and Separation Elasticities

This section proves the claim from Section 2.2 that if there are no exogenous separations $s_0 = 0$, workers completely discount $\beta_w = 0$, and workers have log indirect utility from wages $\eta \rightarrow 1$, and if a stationary equilibrium exists where firms are in steady state, then the sum of recruiting and separation elasticities $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$ for any choice of w in steady state.

Suppose that a steady state equilibrium exists, characterized by tightness θ , the aggregate wage index $\tilde{w}$, the distribution of posted wages $\Upsilon(w)$, the distribution of employed wages $\Phi(w)$, and wage, employment, and vacancy policies w_j^* , N_j^* , and V_j^* . Suppose there are K wage levels. Given $w_1, \dots, w_K$ and $V_1, \dots, V_K$, we want to solve for steady state employment shares $\phi_1, \dots, \phi_K$. We also want to show that $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$ in this steady state. Let us assume that:

$$\frac{\phi_i}{\phi_k} = \frac{v_i}{v_k} \left(\frac{w_i}{w_k} \right)^\gamma$$

for any two wage levels i and k . Then we need to (1) show that the employment share in any wage level is in a steady state, (2) solve for the level of the ϕ_k 's, and (3) show that $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$ for any firm j 's wage policy w_j .

First let's show that inflows are equal to outflows for any sector i . Inflows to sector i are:

$$\begin{aligned} & V_i g(\theta) \sum_{k \neq i} \phi_k \frac{w_i^\gamma}{w_i^\gamma + w_k^\gamma} \\ &= V_i \frac{g(\theta)}{f(\theta)} f(\theta) \sum_{k \neq i} \phi_i \frac{v_k}{v_i} \left(\frac{w_k}{w_i} \right)^\gamma \frac{w_i^\gamma}{w_i^\gamma + w_k^\gamma} \quad \left(\text{since } \frac{\phi_k}{\phi_i} = \frac{v_k}{v_i} \left(\frac{w_k}{w_i} \right)^\gamma \right) \end{aligned}$$

$$\begin{aligned}
&= V_i \frac{1}{\theta} f(\theta) \phi_i \frac{1}{v_i} \sum_{k \neq i} v_k \frac{w_k^\gamma}{w_i^\gamma + w_k^\gamma} \\
&= V_i \frac{S}{V} f(\theta) \phi_i \frac{V}{V_i} \sum_{k \neq i} v_k \frac{w_k^\gamma}{w_i^\gamma + w_k^\gamma} \\
&= S N f(\theta) \phi_i \sum_{k \neq i} v_k \frac{w_k^\gamma}{w_i^\gamma + w_k^\gamma} \\
&= N_i \lambda_{EE} f(\theta) \sum_{k \neq i} v_k \frac{w_k^\gamma}{w_i^\gamma + w_k^\gamma},
\end{aligned}$$

which is the formula for outflows from sector i (using $f(\theta)/g(\theta) = \theta$, $\theta = V/S$, $S = \lambda_{EE}N$, $\phi_i = N_i/N$). Thus, when $\frac{\phi_i}{\phi_k} = \frac{v_i}{v_k} \left(\frac{w_i}{w_k} \right)^\gamma$, this labor market is in steady state.

Next we solve for the values of ϕ_i . First, define some constant C such that

$$\frac{v_i w_i^\gamma}{\phi_i} = \frac{v_k w_k^\gamma}{\phi_k} = C \rightarrow \phi_k = \frac{v_k w_k^\gamma}{C}, \quad \forall k.$$

We also know that $\sum_{k=1}^K \phi_k = 1$. Thus

$$\sum_{k=1}^K \frac{w_k^\gamma v_k}{C} = 1 \rightarrow \sum_{k=1}^K w_k^\gamma v_k = C.$$

Thus

$$\phi_i = \frac{v_i w_i^\gamma}{C} = \frac{v_i w_i^\gamma}{\sum_{k=1}^K v_k w_k^\gamma}.$$

Lastly, we show that for any firm j , the firm faces a constant $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$ at any value of w_j . The firm's ratio of separation rate to recruiting rate is:

$$\begin{aligned}
\frac{S(w_j)}{R(w_j)} &= \frac{f(\theta) \lambda_{EE} \left(v_1 \frac{w_1^\gamma}{w_1^\gamma + w_j^\gamma} + \dots + v_K \frac{w_K^\gamma}{w_K^\gamma + w_j^\gamma} \right)}{g(\theta) \left(\phi_1 \frac{w_j^\gamma}{w_1^\gamma + w_j^\gamma} + \dots + \phi_K \frac{w_j^\gamma}{w_K^\gamma + w_j^\gamma} \right)} \\
&= \lambda_{EE} \theta \frac{\left(v_1 \frac{w_1^\gamma}{w_1^\gamma + w_j^\gamma} + \dots + v_K \frac{w_K^\gamma}{w_K^\gamma + w_j^\gamma} \right)}{\left(\phi_1 \frac{w_j^\gamma}{w_1^\gamma + w_j^\gamma} + \dots + \phi_K \frac{w_j^\gamma}{w_K^\gamma + w_j^\gamma} \right)}
\end{aligned}$$

Using that $v_i w_i^\gamma = \phi_i \sum_{k=1}^K w_k^\gamma v_k$, we can rewrite the above equation as:

$$\begin{aligned}
\frac{S(w_j)}{R(w_j)} &= \lambda_{EE} \theta \frac{\sum_{k=1}^K v_k w_k^\gamma}{w_j^\gamma} \frac{\left(\phi_1 \frac{1}{w_1^\gamma + w_j^\gamma} + \dots + \phi_K \frac{1}{w_K^\gamma + w_j^\gamma} \right)}{\left(\phi_1 \frac{1}{w_1^\gamma + w_j^\gamma} + \dots + \phi_K \frac{1}{w_K^\gamma + w_j^\gamma} \right)} \\
&= \lambda_{EE} \theta w_j^{-\gamma} \sum_{k=1}^K v_k w_k^\gamma = \lambda_{EE} \theta \left(\frac{\tilde{w}}{w_j} \right)^\gamma,
\end{aligned}$$

with $\tilde{w} = (\sum_{k=1}^K v_k w_k^\gamma)^\frac{1}{\gamma}$, where the elasticity of this function with respect to w_j is $-\gamma$.

Characterization of Constant Elasticity of Equilibrium

Next, we characterize firms' optimal employment and wages, focusing on the case of $\eta = 1$, $s_0 = 0$, $\beta_f = 1$, and $\beta_w = 0$. By solving out for capital, we can define auxiliary parameters

$$\tilde{A}_j = A_j^{\frac{1}{1-\alpha_j \frac{\epsilon-1}{\epsilon}}} \left(\frac{\alpha_j \frac{\epsilon-1}{\epsilon}}{r^K} \right)^{\frac{\alpha_j \frac{\epsilon-1}{\epsilon}}{1-\alpha_j \frac{\epsilon-1}{\epsilon}}} \quad \tilde{\alpha}_j = \frac{(1-\alpha_j) \frac{\epsilon-1}{\epsilon}}{1-\alpha_j \frac{\epsilon-1}{\epsilon}}.$$

Given parameters $\tilde{\alpha}_j$, $\tilde{A}_j$, χ , σ , and c_j , aggregate labor market tightness θ , and aggregate wage index $\tilde{w}$, a firm j 's optimal wage is

$$w_j^* = \left(c_j \gamma (1+\chi) (\lambda_{EE} \theta)^{1+\chi} \tilde{w}^{\gamma(1+\chi)} \left(\frac{\tilde{\alpha}_j \tilde{A}_j \gamma (1+\chi)}{1+\gamma(1+\chi)+\sigma} \right)^{\frac{\sigma}{1-\tilde{\alpha}_j}} \right)^{\frac{1-\tilde{\alpha}_j}{(1-\tilde{\alpha}_j)(1+\gamma(1+\chi))+\sigma}}, \quad (\text{B.9})$$

where $\tilde{w}$ is a vacancy-weighted index of the distribution of wages: $\tilde{w} = (\int_{w_k} \nu(w_k) w_k^\gamma dw_k)^\frac{1}{\gamma}$. Taking $\beta_f \rightarrow 1$, optimal employment at firm j is

$$N_j^* = \left(c_j \gamma (1+\chi) (\lambda_{EE} \theta)^{1+\chi} \tilde{w}^{\gamma(1+\chi)} \right)^{\frac{-1}{(1-\tilde{\alpha}_j)(1+\gamma(1+\chi))+\sigma}} \left(\tilde{\alpha}_j \tilde{A}_j \frac{\gamma(1+\chi)}{(1+\gamma(1+\chi))+\sigma} \right)^{\frac{1+\gamma(1+\chi)}{(1-\tilde{\alpha}_j)(1+\gamma(1+\chi))+\sigma}} \quad (\text{B.10})$$

To have an interior solution, we must have that $\sigma > -(1-\tilde{\alpha}_j)(1+\gamma(1+\chi))$, i.e., σ cannot be too negative. The optimal vacancy policy is simply $V_j^* = \bar{V}(\tilde{w}/w_j)^\gamma N_j^*$, and labor market tightness is $\theta = (\int_{k \in J} V_k dk) / (\lambda_{EE} N)$.

No Backloading of Wages

We want to show that in our equilibrium with myopic workers and no exogenous separations, no firm has an incentive to backload wages. In an equilibrium with log utility $\eta \rightarrow 0$, no exogenous separations $s_0 = 0$, and myopic workers $\beta_w = 0$, we have the result that $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$ at any wage level. We now want to show it in general for K wages. That is, wages are $w_1, w_2, \dots, w_K$, employment shares, $\{\phi_k\}_{k=1}^K$, vacancy shares, $\{v_k\}_{k=1}^K$ where $\sum_{k=1}^K \phi_k = 1$ and $\sum_{k=1}^K v_k = 1$.

In equilibrium, the firm's first order condition on wages satisfies

$$N^* = \lambda(R'(w^*)V^* - S'(w^*)N^*),$$

where stars denote the optimal choice variables. Rearranging and plugging in steady-state condition $\frac{V}{N} = \frac{S}{R}$, this yields

$$w^* = S(w^*) \lambda(\varepsilon_{R,w^*} - \varepsilon_{S,w^*}). \quad (\text{B.11})$$

Now consider a firm that is setting wages period by period. The first order condition yields

$$\begin{aligned} -S'(w_t)N_{t-1}\lambda &= N_t \\ -S'(w_t)N_{t-1}\lambda &= (1 - S(w_t))N_{t-1} \\ -S'(w_t)\lambda &= (1 - S(w_t)). \end{aligned}$$

Let's define $w^\#$ as the optimal wage if the firm were setting wages period by period. Then the optimal wage would satisfy

$$w^\# = \lambda(-\varepsilon_{S,w^\#}) \frac{S(w^\#)}{1 - S(w^\#)}. \quad (\text{B.12})$$

We know that the left hand side and right hand side of equation (B.11) cross only once because when we plug in the value of λ into equation (B.11), we get

$$w = c(1 + \chi) \left(\frac{S(w)}{R(w)} \right)^{1+\chi} (\varepsilon_{R,w} - \varepsilon_{S,w}) N^\sigma.$$

For any firm size N , since $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$ and $S(w)/R(w) \propto w^{-\gamma}$, the left and right hand side cross only once.

Thus, we can show that the optimal discretionary wage $w^\#$ is below the commitment wage w^* if the right hand side of equation (B.12) lies below the right hand side of equation (B.9) for all w , that is if

$$\lambda S(w)(\varepsilon_{R,w} - \varepsilon_{S,w}) + \lambda \varepsilon_{S,w} \frac{S(w)}{1 - S(w)} > 0$$

for all w . With some rearranging, this expression becomes

$$\begin{aligned} \lambda S(w)\varepsilon_{R,w} - \lambda S(w)\varepsilon_{S,w} + \lambda \varepsilon_{S,w} \frac{S(w)}{1 - S(w)} &> 0 \\ S(w)(\varepsilon_{R,w} - \varepsilon_{S,w})(1 - S(w)) + \varepsilon_{S,w}S(w) &> 0 \\ (\varepsilon_{R,w} - \varepsilon_{S,w})(1 - S(w)) + \varepsilon_{S,w} &> 0 \\ (\varepsilon_{R,w} - \varepsilon_{S,w}) + \varepsilon_{S,w} &> (\varepsilon_{R,w} - \varepsilon_{S,w})S(w) \end{aligned}$$

With $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$, this becomes

$$\varepsilon_{R,w} > \gamma S(w).$$

So now we need to prove this condition. Define the separation function for a firm paying wage w :

$$S(w) = \lambda_{EE} f(\theta) \sum_{k=1}^K \frac{\bar{w}_k^\gamma}{w^\gamma + \bar{w}_k^\gamma} v_k.$$

The recruiting function is

$$R(w) = g(\theta) \sum_{k=1}^K \frac{\bar{w}_k^\gamma}{w^\gamma + \bar{w}_k^\gamma} \phi_k.$$

the elasticity is then: The recruiting elasticity is:

$$\varepsilon_{R,w} = \gamma \cdot \frac{\sum_{k=1}^K \frac{\bar{w}_k^\gamma w^\gamma}{(w^\gamma + \bar{w}_k^\gamma)^2} \phi_k}{\sum_{k=1}^K \frac{w^\gamma}{w^\gamma + \bar{w}_k^\gamma} \phi_k}.$$

The condition is then

$$\varepsilon_{R,w} > \gamma S(w).$$

If and only if

$$\gamma \frac{\sum_{k=1}^K \frac{\bar{w}_k^\gamma w^\gamma}{(w^\gamma + \bar{w}_k^\gamma)^2} \phi_k}{\sum_{k=1}^K \frac{w^\gamma}{w^\gamma + \bar{w}_k^\gamma} \phi_k} > \gamma \lambda_{EE} f(\theta) \sum_{k=1}^K \frac{\bar{w}_k^\gamma}{w^\gamma + \bar{w}_k^\gamma} v_k,$$

where $\lambda_{EE} f(\theta) < 1$. Let

$$\frac{\phi_i}{\phi_k} = \frac{v_i}{v_k} \left(\frac{w_i}{w_k} \right)^\gamma$$

for any two wage levels i and k .

First, define some constant C such that

$$\frac{v_i w_i^\gamma}{\phi_i} = \frac{v_k w_k^\gamma}{\phi_k} = C,$$

Using the condition that $\sum_{k=1}^K \phi_k = 1$, we find:

$$\sum_{k=1}^K \frac{v_k w_k^\gamma}{C} = 1 \Rightarrow C = \sum_{k=1}^K v_k w_k^\gamma$$

Therefore we obtain

$$\phi_k = \frac{v_k w_k^\gamma}{\sum_{j=1}^K v_j w_j^\gamma}.$$

Define $D \equiv \sum_{j=1}^K v_j w_j^\gamma$. Then:

$$\varepsilon_{R,w} = \gamma \cdot \frac{\sum_{k=1}^K \frac{\bar{w}_k^\gamma w^\gamma v_k \bar{w}_k^\gamma}{(w^\gamma + \bar{w}_k^\gamma)^2 D}}{\sum_{k=1}^K \frac{w^\gamma v_k \bar{w}_k^\gamma}{(w^\gamma + \bar{w}_k^\gamma) D}}$$

which is equal to

$$\varepsilon_{R,w} = \gamma \cdot \frac{\sum_{k=1}^K \frac{\bar{w}_k^\gamma w^\gamma v_k \bar{w}_k^\gamma}{(w^\gamma + \bar{w}_k^\gamma)^2}}{\sum_{k=1}^K \frac{w^\gamma v_k \bar{w}_k^\gamma}{(w^\gamma + \bar{w}_k^\gamma)}}$$

Cancel w^γ as there is no k :

$$\varepsilon_{R,w} = \gamma \cdot \frac{\sum_{k=1}^K \frac{(\bar{w}_k^\gamma)^2 v_k}{(w^\gamma + \bar{w}_k^\gamma)^2}}{\sum_{k=1}^K \frac{v_k \bar{w}_k^\gamma}{(w^\gamma + \bar{w}_k^\gamma)}}$$

The condition then becomes

$$\gamma \cdot \frac{\sum_{k=1}^K \frac{(\bar{w}_k^\gamma)^2}{(w^\gamma + \bar{w}_k^\gamma)^2} v_k}{\sum_{k=1}^K \frac{\bar{w}_k^\gamma}{(w^\gamma + \bar{w}_k^\gamma)} v_k} > \gamma \lambda_{EE} f(\theta) \sum_{k=1}^K \frac{\bar{w}_k^\gamma}{w^\gamma + \bar{w}_k^\gamma} v_k.$$

Define the function:

$$f_k = \frac{\bar{w}_k^\gamma}{w^\gamma + \bar{w}_k^\gamma}.$$

Rewriting the condition to obtain

$$\frac{\sum_{k=1}^K f_k^2 v_k}{\sum_{k=1}^K f_k v_k} > \lambda_{EE} f(\theta) \sum_{k=1}^K f_k v_k.$$

which is equal to

$$\sum_{k=1}^K f_k^2 v_k > \lambda_{EE} f(\theta) \left(\sum_{k=1}^K f_k v_k \right)^2, \quad (\text{B.13})$$

which holds by Jensen's inequality as $\lambda_{EE} f(\theta) < 1$.

Proving Uniqueness

Finally, we show that all the endogenous outcomes θ , $\tilde{w}$, N_j , w_j , and V_j can be solved for as functions of parameters. Using the result from Appendix B.7 that in any equilibrium firms will have a constant value of $\varepsilon_{R,w} - \varepsilon_{S,w} = \gamma$, and as such optimal wages and employment will take the form of (B.9) and (B.10), we then confirm that the stationary equilibrium is unique.

Lemma 3. *Vacancy shares v_j are a function of only common parameters and do not depend on aggregate wages $\tilde{w}$ or labor market tightness θ .*

Proof: First we want to show that relative wages are not a function of $\tilde{w}$ or θ (assuming a common $\tilde{\alpha}$ across firms). For two firms j and k , relative wages are:

$$\frac{w_j}{w_k} = \left(\frac{c_j}{c_k} \right)^{\frac{1-\tilde{\alpha}}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\frac{\tilde{A}_j}{\tilde{A}_k} \right)^{\frac{\sigma}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}}$$

Next we show that relative employment is not a function of $\tilde{w}$ or θ . For two firms j and k , relative employment is:

$$\frac{N_j}{N_k} = \left(\frac{c_j}{c_k} \right)^{\frac{-1}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\frac{\tilde{A}_j}{\tilde{A}_k} \right)^{\frac{1+\gamma(1+\chi)}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}}$$

Thus, relative vacancies are a function of parameters only:

$$\frac{V_j}{V_k} = \frac{N_j}{N_k} \left(\frac{w_j}{w_k} \right)^{-\gamma} = \left(\frac{c_j}{c_k} \right)^{\frac{-1-\gamma(1-\tilde{\alpha})}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\frac{\tilde{A}_j}{\tilde{A}_k} \right)^{\frac{1+\gamma(1+\chi)-\gamma\sigma}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}}.$$

Let ω_j be the mass of firms with parameters c_j and $\tilde{A}_j$, with $\int_{\tilde{A}} \int_c \omega_j dc d\tilde{A} = 1$. The share of vacancies of firm type j is

$$v_j = \frac{\omega_j V_j}{\int_{k \in J} \omega_k V_k} = \frac{1}{\int_{k \in J} \frac{\omega_k}{\omega_j} \frac{V_k}{V_j}}.$$

Since V_j/V_k is a function of only parameters, and the share of each type of firm ω_j is a parameter, then the vacancy share v_j is a function of only parameters.

Lemma 4. *The ratio of a firm's optimal wage w_j to the index of aggregate wages $\tilde{w}$ is a function of parameters.*

Proof: (assuming a common $\tilde{\alpha}$ across firms)

$$\begin{aligned} \frac{\tilde{w}}{w_j} &= \frac{\left(\int_k v_k w_k^\gamma dk \right)^{1/\gamma}}{w_j} \\ &= \frac{\left(\int_k v_k \underbrace{\left[c_k \gamma (1+\chi) (\lambda_{EE} \theta)^{1+\chi} \tilde{w}^{\gamma(1+\chi)} \right]}_{=:X}^{\frac{\gamma(1-\tilde{\alpha})}{D}} \left(\frac{\tilde{\alpha} \tilde{A}_k \gamma (1+\chi)}{1+\gamma(1+\chi)+\sigma} \right)^{\frac{\gamma\sigma}{D}} dk \right)^{1/\gamma}}{\left[X^{\frac{1-\tilde{\alpha}}{D}} \left(\frac{\tilde{\alpha} \tilde{A}_j \gamma (1+\chi)}{1+\gamma(1+\chi)+\sigma} \right)^{\frac{\sigma}{D}} \right]}, \end{aligned}$$

where $D \equiv (1 - \tilde{\alpha})(1 + \gamma(1 + \chi)) + \sigma$. All factors in X (i.e., γ , χ , λ_{EE} , θ , and the common power of $\tilde{w}$) cancel between the numerator and denominator. Therefore,

$$\frac{\tilde{w}}{w_j} = \frac{\left(\int_k v_k c_k^{\frac{\gamma(1-\tilde{\alpha})}{D}} \tilde{A}_k^{\frac{\gamma\sigma}{D}} dk \right)^{1/\gamma}}{c_j^{\frac{1-\tilde{\alpha}}{D}} \tilde{A}_j^{\frac{\sigma}{D}}}.$$

Because the vacancy shares v_k are functions only of parameters (see Lemma above), the ratio $\tilde{w}/w_j$ is a function of parameters alone and does not depend on labor-market tightness θ (nor on λ_{EE} , γ , or χ).

Thus, we can write relative wages as a function of parameters only:

$$\frac{\tilde{w}}{w_j} = \frac{\tilde{w}}{w_j}(\mathfrak{c}, \tilde{\mathbf{A}}).$$

Lemma 5. *Given distributions of firm-specific parameters $\tilde{\mathbf{A}}$ and $\mathfrak{c}$, the wage index $\tilde{w}$ is an increasing function of tightness θ .*

Proof: Based on equation (B.9), individual firm wages w_j are increasing in aggregate wages $\tilde{w}$ and tightness θ . Solving out individual wages, and solving $\tilde{w}$ in terms of only θ and parameters yields:

$$\begin{aligned} \tilde{w} &= \left(\int_j v_j w_j^\gamma dj \right)^{\frac{1}{\gamma}} \\ \tilde{w} &= \left(\int_j v_j \left(\theta^{1+\chi} \tilde{w}^{\gamma(1+\chi)} \right)^{\frac{\gamma(1-\tilde{\alpha})}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} g(c_j, \tilde{A}_j) dj \right)^{\frac{1}{\gamma}} \end{aligned}$$

$$\tilde{w} = \left(\theta^{1+\chi} \tilde{w}^{\gamma(1+\chi)} \right)^{\frac{(1-\tilde{\alpha})}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\int_j v_j g(c_j, \tilde{A}_j) dj \right)^{\frac{1}{\gamma}}, \quad (\text{B.14})$$

with

$$g(c_j, \tilde{A}_j) = \left(c_j \gamma (1+\chi) \lambda_{EE}^{1+\chi} \right)^{\frac{\gamma(1-\tilde{\alpha})}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\frac{\tilde{\alpha} \tilde{A}_j \gamma (1+\chi)}{1+\gamma(1+\chi)+\sigma} \right)^{\frac{\gamma\sigma}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}}.$$

Grouping all $\tilde{w}$ terms on the left hand side yields:

$$\tilde{w}^{\frac{(1-\tilde{\alpha})+\sigma}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} = \theta^{\frac{(1+\chi)(1-\tilde{\alpha})}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\int_j v_j g(c_j, \tilde{A}_j) dj \right)^{\frac{1}{\gamma}}. \quad (\text{B.15})$$

Simplifying yields:

$$\tilde{w} = \theta^{\frac{(1+\chi)(1-\tilde{\alpha})}{(1-\tilde{\alpha})+\sigma}} \left(\int_j v_j g(c_j, \tilde{A}_j) dj \right)^{\frac{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}{\gamma(1-\tilde{\alpha}+\sigma)}}.$$

Proof of Proposition Let $\omega(\tilde{A}, c)$ be the joint density function of the mass of firm populations over parameters c and $\tilde{A}$. Then the aggregate mass of vacancies V is

$$V = \int_{\tilde{A}} \int_c V_j \omega(\tilde{A}, c) dc d\tilde{A}$$

$$V = \int_{\tilde{A}} \int_c \lambda_{EE} \theta \left(\frac{\tilde{w}}{w_j} \right)^{\gamma} N_j \omega(\tilde{A}, c) dc d\tilde{A}$$

We have already shown that $\tilde{w}/w_j$ is a function of parameters:

$$V = \int_{\tilde{A}} \int_c \lambda_{EE} \theta \left(\frac{\tilde{w}}{w_j}(\mathbb{c}, \tilde{A}) \right)^{\gamma} N_j \omega(\tilde{A}, c) dc d\tilde{A} \quad (\text{B.16})$$

We can abbreviate the expression for optimal employment as

$$N_j = \left(\theta^{1+\chi} \tilde{w}^{\gamma(1+\chi)} \right)^{\frac{-1}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} f(c_j, \tilde{A}_j), \quad (\text{B.17})$$

with

$$f(c_j, \tilde{A}_j) = \left(c_j \gamma (1+\chi) \lambda_{EE}^{1+\chi} \right)^{\frac{-1}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \left(\frac{\tilde{\alpha} \tilde{A}_j \gamma (1+\chi)}{1+\gamma(1+\chi)+\sigma} \right)^{\frac{1+\gamma(1+\chi)}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}}.$$

Plugging in (B.15) and (B.17) into equation (B.16) yields:

$$V = \int \int \theta \times \theta^{\frac{-1-\chi}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \times \theta^{\frac{(1+\chi)(1-\tilde{\alpha})}{(1-\tilde{\alpha})+\sigma} \frac{-\gamma(1+\chi)}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} H(\tilde{A}_j, c_j, \tilde{A}, \mathbb{c}) dc d\tilde{A},$$

with

$$H(\tilde{A}_j, c_j, \tilde{A}, \mathbb{c}) = \left(\frac{\tilde{w}}{w_j}(\mathbb{c}, \tilde{A}) \right)^{\gamma} f(c_j, \tilde{A}_j) \left(\int_j v_j g(c_j, \tilde{A}_j) dj \right)^{\frac{-1-\chi}{(1-\tilde{\alpha})+\sigma}} \omega(\tilde{A}_j, c_j).$$

In equilibrium where all workers are employed, and employed workers search with probability λ_{EE} , then (normalizing the worker population to 1) tightness θ is:

$$\theta = \frac{V}{\lambda_{EE}}.$$

The prior equation then becomes:

$$\lambda_{EE}\theta = \int \int \theta \times \theta^{\frac{-1-\chi}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} \times \theta^{\frac{(1+\chi)(1-\tilde{\alpha})}{(1-\tilde{\alpha})+\sigma} \frac{-\gamma(1+\chi)}{(1-\tilde{\alpha})(1+\gamma(1+\chi))+\sigma}} H(\tilde{A}_j, c_j, \tilde{\mathbb{A}}, \mathbb{C}) dc d\tilde{A}.$$

The θ terms on the right-hand side can be pulled out of the integral, and a θ term on both sides cancels. With some algebra, we have

$$\theta = \left(\frac{1}{\lambda_{EE}} \int \int H(\tilde{A}_j, c_j, \tilde{\mathbb{A}}, \mathbb{C}) dc d\tilde{A} \right)^{\frac{(1-\tilde{\alpha})+\sigma}{1+\chi}}.$$

Thus equilibrium labor market tightness can be calculated strictly as a function of parameters. Plugging in our value of θ in equation (B.15), we can solve for the aggregate wage $\tilde{w}$ as a function of parameters. Given $\tilde{w}$, we can compute individual firms' N_j , w_j , and V_j .

B.8 Recruiting Costs as Reallocating Labor

In this section, we show that counting recruiting costs as taking away from labor inputs generates identical formulas as treating recruiting costs as expenditure. Suppose that recruiting takes away from labor inputs, so a firm's labor input is

$$N_t - cV_t^{1+\chi}N_{t-1}^{\sigma-\chi}.$$

For simplicity, assume $\sigma = 0$. If firms do not discount ($\beta_f = 1$) and revenue is $PY = \left(K_t^\alpha \left(N_t - cV_t^{1+\chi}N_{t-1}^{-\chi} \right)^{1-\alpha} \right)^{\frac{\epsilon}{\epsilon-1}}$ then firm maximizes the flow rate of profits:

$$\left(K_t^\alpha \left(N_t - cV_t^{1+\chi}N_{t-1}^{-\chi} \right)^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - w_t N_t - r^K K_t.$$

Given that $V_t = \frac{S(w_t)}{R(w_t)}N_t$ in steady state, and dropping time subscripts, we have that the firm solves

$$\begin{aligned} \max_{K, N, w} & \left(K^\alpha \left(N - c \left(\frac{S(w)}{R(w)} \right)^{1+\chi} N \right)^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - wN - r^K K \\ & K^{\alpha \frac{\epsilon-1}{\epsilon}} N^{(1-\alpha) \frac{\epsilon-1}{\epsilon}} \left(1 - c \left(\frac{S(w)}{R(w)} \right)^{1+\chi} \right)^{(1-\alpha) \frac{\epsilon-1}{\epsilon}} - wN - r^K K. \end{aligned}$$

Taking the first order condition with respect to N , we have

$$FOC_N : (1 - \alpha) \frac{\epsilon - 1}{\epsilon} K^{\alpha \frac{\epsilon - 1}{\epsilon}} N^{(1 - \alpha) \frac{\epsilon - 1}{\epsilon} - 1} \left(1 - c \left(\frac{S(w)}{R(w)} \right)^{1 + \chi} \right)^{(1 - \alpha) \frac{\epsilon - 1}{\epsilon}} = w.$$

Multiplying both sides by N , we have

$$(1 - \alpha) \frac{\epsilon - 1}{\epsilon} K^{\alpha \frac{\epsilon - 1}{\epsilon}} N^{(1 - \alpha) \frac{\epsilon - 1}{\epsilon}} \left(1 - c \left(\frac{S(w)}{R(w)} \right)^{1 + \chi} \right)^{(1 - \alpha) \frac{\epsilon - 1}{\epsilon}} = wN \quad (\text{B.18})$$

Rearranging the first order condition on N once again, we come away with the simple formula

$$(1 - \alpha) \frac{\epsilon - 1}{\epsilon} PY = wN,$$

and consequently the labor share is the same as in the main text:

$$\frac{wN}{PY} = (1 - \alpha) \frac{\epsilon - 1}{\epsilon}.$$

B.9 Firm's Problem with Hiring Costs

In this section, we solve a general model where firms not only have to pay a vacancy cost, but firms also pay a hiring cost which may depend on firm size and the hiring rate. We abstract from capital and downward sloping product demand, assuming that output $= AN^{1 - \alpha}$. Hires each period are $H_t = V_t R(w_t)$. We also assume for simplicity that firms are committed to pay their posted wage, unlike in the main text where the posted wage is a floor. Firms maximize

$$\begin{aligned} \max_{\{N_{s,t}\}, \{w_{s,t}\}, \{V_t\}} \sum_{t=0}^{\infty} \beta_f^t & \left(A \left(\sum_{s=-\infty}^t N_{s,t} \right)^{\alpha} - \sum_{s=-\infty}^t w_{s,t} N_{s,t} \right. \\ & \left. - c \left(\frac{V_t}{N_{t-1}} \right)^{\chi} V_t N_{t-1}^{\sigma} - c_h \left(\frac{H_t}{N_{t-1}} \right)^{\chi_h} H_t N_{t-1}^{\sigma_h} \right) \end{aligned}$$

subject to

$$\begin{aligned} N_{s,t} &= (1 - S(w_s)) N_{s,t-1}, & N_{t,t} &= H_t = R(w_t) V_t, \\ w_{s,t+1} &\geq w_{s,t}, & N_t &= \sum_s N_{s,t}. \end{aligned}$$

Using the same separation and recruiting functions as in the main text, the Lagrangian is

$$\begin{aligned} \mathcal{L} : \sum_{t=0}^{\infty} \beta_f^t & \left(A \left(\sum_{s=-\infty}^t N_{s,t} \right)^{\alpha} - \sum_{s=-\infty}^t w_{s,t} N_{s,t} - c \left(\frac{V_t}{N_{t-1}} \right)^{\chi} V_t N_{t-1}^{\sigma} + \right. \\ & \left. - c_h \left(\frac{V_t}{N_{t-1}} \right)^{\chi_h} V_t R(w_{s,t})^{1 + \chi_h} N_{t-1}^{\sigma_h} + \sum_{s=-\infty}^t \mu_{s,t} \left[-N_{s,t} + N_{s,t-1} (1 - S(w_{s,t})) \right] \right. \\ & \left. + \lambda_t [-N_{t,t} + R(w_t) V_t] \right). \end{aligned}$$

Define $\frac{c_v}{N} \equiv c \left(\frac{V}{N}\right)^{1+\chi} N^\sigma$ and $\frac{c_h}{N} \equiv c \left(\frac{H}{N}\right)^{1+\chi_h} N^{\sigma_h}$. Taking first order conditions and solving, it can be shown that the recruiting and hiring cost adjust markdown is

$$\frac{w + \frac{c_v}{N} + \frac{c_h}{N}}{MRPL} = \frac{(1 + (1 + \chi)(\varepsilon_{Rw} - \varepsilon_{Sw})) \frac{c_v}{N} + (1 - \varepsilon_{Sw}(1 + \chi_h)) \frac{c_h}{N}}{(1 + (1 + \chi)(\varepsilon_{Rw} - \varepsilon_{Sw}) + \sigma) \frac{c_v}{N} + (1 - \varepsilon_{Sw}(1 + \chi_h) + \sigma_h) \frac{c_h}{N}}$$

which nests the case with only recruiting costs when $c_h = 0$. Introducing hiring costs puts more weight on the separation elasticity, as the value of higher wages becomes more concentrated on avoiding extra separations (and therefore hires), rather than filling vacancies more quickly.

B.10 Oligopsony

In this section, we derive the first order conditions for the oligopsonistic firm's problem. After deriving the first order conditions, we prove Proposition 2. We solve the firm's problem by first assuming that there is no unemployment, which can be achieved in equilibrium so long as the unemployment benefit b is sufficiently low (under this assumption, $U_t \equiv 0$ so $\Phi_{t-1}^U = 0$ and the unemployment terms in (B.26)–(B.27) drop out).

Workers The worker gets utility from wages and draws a utility shock when given a choice between two firms. Let $\phi_{k,t}^v$ be the share of vacancies posted by firm k in period t . If workers encounter their own firm's vacancies, the workers redraw preferences (as if they are choosing between two jobs), but the firm is not obligated to pay the same wage as its other hires in period t , and the worker continues to receive the wage of the cohort of workers they were initially hired with in period s .

$$\begin{aligned} \mathcal{V}_{j,s,t} = & \frac{\eta}{\eta-1} \left(w_{j,s,t}^{\frac{\eta-1}{\eta}} - 1 \right) + \beta_w s_0 \mathcal{V}_U + \\ & \beta_w (1 - s_0) \mathbb{E}_t \left[\lambda_{EE} f(\theta) \left(\sum_{k \neq j} \phi_{k,t}^v \gamma^{-1} \log \left(\exp(\gamma \mathcal{V}_{j,s,t+1}) + \exp(\gamma \mathcal{V}_{k,t+1,t+1}) \right) + \right. \right. \\ & \left. \left. \phi_{j,t}^v \gamma^{-1} \log \left(\exp(\gamma \mathcal{V}_{j,s,t+1}) + \exp(\gamma \mathcal{V}_{j,t+1,t+1}) \right) \right) + (1 - \lambda_{EE} f(\theta)) \max\{\mathcal{V}_{j,s,t+1}, \mathcal{V}_U\} \right], \end{aligned} \quad (\text{B.19})$$

where the second logarithm compares staying in the current job (value $\mathcal{V}_{j,s,t+1}$) with accepting a new vacancy at firm j (value $\mathcal{V}_{j,t+1,t+1}$).

We will now consider one-period deviations in period t for firms hiring workers in period t , under the assumption from workers that the posted wage will be equal to their wage for the duration of their match with the firm. Dropping the time subscripts, we have

$$\frac{\partial \mathcal{V}_j}{\partial w_j} = w_j^{-\frac{1}{\eta}} + \beta_w \left(f(\theta) \lambda_{EE} \left(\sum_k \phi_k^v \left(\frac{\exp(\gamma \mathcal{V}_j)}{\exp(\gamma \mathcal{V}_j) + \exp(\gamma \mathcal{V}_k)} \right) + \phi_j^v \right) + (1 - f(\theta) \lambda_{EE}) \right) \frac{\partial \mathcal{V}_j}{\partial w_j}.$$

under symmetry, wages are equal across firms ($w_j = w_k$), so $\mathcal{V}_j = \mathcal{V}_k$, and employment and vacancy shares coincide ($\phi_{k,t}^n = \phi_{k,t}^v$). If there are no exogenous separations ($s_0 = 0$), then $S_j = \lambda_{EE} f(\theta) (1 - \phi_{j,t}^v)/2$.

$$\begin{aligned} \frac{\partial \mathcal{V}_j}{\partial w_j} &= w_j^{-\frac{1}{\eta}} + \beta_w \left(\frac{f(\theta) \lambda_{EE}}{2} + (1 - f(\theta) \lambda_{EE}) \right) \frac{\partial \mathcal{V}_j}{\partial w_j} \\ &= w_j^{-\frac{1}{\eta}} + \beta_w \left(1 - \frac{\lambda_{EE} f(\theta)}{2} \right) \frac{\partial \mathcal{V}_j}{\partial w_j} \\ &= w_j^{-\frac{1}{\eta}} + \beta_w (1 - S_j) \frac{\partial \mathcal{V}_j}{\partial w_j}, \end{aligned}$$

where the last line uses that $S_j = \lambda_{EE} f(\theta)/2$ under symmetry. Thus

$$\frac{\partial \mathcal{V}_j}{\partial w_j} = \frac{w_j^{-\frac{1}{\eta}}}{(1 - \beta_w(1 - S_j))}. \quad (\text{B.20})$$

Firms Firms solve

$$\begin{aligned} \max_{\{V_{j,t}\}_{t=0}^\infty, \{N_{j,t}\}_{t=0}^\infty, \{N_{j,s,t}\}_{s \leq t}^\infty, \{w_{j,s,t}\}_{s \leq t}^\infty} \sum_{t=0}^\infty \beta_f^t \left[A_{j,t} \left(K_{j,t}^\alpha N_{j,t}^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - \sum_{s \leq t} w_{j,s,t} N_{j,s,t} \right. \\ \left. - r^K K_{j,t} - c_j \left(\frac{V_{j,t}}{N_{j,t-1}} \right)^\chi N_{j,t-1}^\sigma V_{j,t} \right] \end{aligned} \quad (\text{B.21})$$

subject to

$$\text{Recruiting:} \quad N_{j,t,t} = R(\mathcal{V}_{j,t,t}) V_{j,t} \quad \forall t \quad (\text{B.22})$$

$$\text{Retention:} \quad N_{j,s,t} = (1 - S(\mathcal{V}_{j,s,t})) N_{j,s,t-1} \quad \forall t \text{ and } s < t \quad (\text{B.23})$$

$$\text{Employment is sum of cohorts:} \quad N_{j,t} = \sum_{s=-\infty}^t N_{j,s,t} \quad \forall t \quad (\text{B.24})$$

$$\text{Wage at least posted wage:} \quad w_{j,s,t} \geq w_{j,s,s} \quad \forall t \text{ and } s < t, \quad (\text{B.25})$$

where R and S are

$$\begin{aligned} R(\mathcal{V}_{j,t,t}, \mathcal{V}_{-j,t}, V_{j,t}, V_{-j,t}, N_{j,t-1}, N_{-j,t-1}) = \\ g(\theta_t) \left[(1 - \Phi_{t-1}^U) \sum_k \sum_{s=-\infty}^{t-1} \phi_{k,s,t-1}^n \frac{\exp(\gamma \mathcal{V}_{j,t,t})}{\exp(\gamma \mathcal{V}_{j,t,t}) + \exp(\gamma \mathcal{V}_{k,s,t})} + \Phi_{t-1}^U \mathbb{1}\{\mathcal{V}_{j,t,t} > \mathcal{V}_U\} \right] \end{aligned} \quad (\text{B.26})$$

$$\begin{aligned} S(\mathcal{V}_{j,s,t}, \mathcal{V}_{-j,t}, V_{j,s,t}, V_{-j,t}) = \lambda_{EE} f(\theta_t) \sum_k \phi_{k,t}^v \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})} \\ + (1 - \lambda_{EE} f(\theta_t)) \mathbb{1}\{\mathcal{V}_U > \mathcal{V}_{j,s,t}\} \end{aligned} \quad (\text{B.27})$$

where $\phi_{k,s,t-1}^n$ is the share of workers who were hired in period s by firm k and still employed in period $t-1$, $\phi_{k,t}^v$ is firm k 's share of vacancies in period t , and $\Phi_{t-1}^U = \frac{U_{t-1}}{U_{t-1} + \lambda_{EE}(1-U_{t-1})}$ is the share of searchers who are unemployed.

Lagrangian. The Lagrangian is

$$\begin{aligned} \mathcal{L} = & \sum_{t=0}^{\infty} \beta_f^t \left(A_{j,t} \left(K_{j,t}^{\alpha} \left(\sum_{s=-\infty}^t N_{j,s,t} \right)^{1-\alpha} \right)^{\frac{\epsilon-1}{\epsilon}} - \sum_{s=-\infty}^t w_{j,s,t} N_{j,s,t} - c_j V_{j,t}^{1+\chi} \left(\sum_{s=-\infty}^{t-1} N_{j,s,t-1} \right)^{\sigma-\chi} \right. \\ & - r^K K_{j,t} + \sum_{s=-\infty}^{t-1} \mu_{s,t} \left[-N_{j,s,t} + N_{j,s,t-1} \left(1 - \lambda_{EE} f(\theta_t) \sum_k \phi_{k,t}^v \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})} \right) \right] \\ & \left. + \lambda_t \left[-N_{j,t,t} + V_{j,t} g(\theta_t) \sum_k \sum_{s=-\infty}^{t-1} \phi_{k,s,t-1}^n \frac{\exp(\gamma \mathcal{V}_{j,t,t})}{\exp(\gamma \mathcal{V}_{j,t,t}) + \exp(\gamma \mathcal{V}_{k,s,t})} \right] \right). \end{aligned}$$

We take the first-order condition with respect to $w_{j,t,t}$ under the assumption that condition (B.25) binds: that is, when the firm raises its posted wage, it raises the wage that the firm pays for as long as the cohort of hired workers is in the firm. For the cohort posted at $t = 0$,

$$\begin{aligned} \mathcal{L}_{w_{j,0,0}} : & - \sum_{t=0}^{\infty} \beta_f^t N_{j,0,t} + \lambda_0 V_{j,0} g(\theta_0) \sum_k \sum_{s=-\infty}^{-1} \phi_{k,s,-1}^n \gamma \frac{\exp(\gamma \mathcal{V}_{j,0,0}) \exp(\gamma \mathcal{V}_{k,s,0})}{(\exp(\gamma \mathcal{V}_{j,0,0}) + \exp(\gamma \mathcal{V}_{k,s,0}))^2} \frac{\partial \mathcal{V}_{j,0,0}}{\partial w_{j,0,0}} \\ & + \sum_{t=1}^{\infty} \beta_f^t \mu_{0,t} N_{j,0,t-1} \lambda_{EE} f(\theta_t) \sum_k \phi_{k,t}^v \gamma \frac{\exp(\gamma \mathcal{V}_{j,0,t}) \exp(\gamma \mathcal{V}_{k,t,t})}{(\exp(\gamma \mathcal{V}_{j,0,t}) + \exp(\gamma \mathcal{V}_{k,t,t}))^2} \frac{\partial \mathcal{V}_{j,0,t}}{\partial w_{j,0,0}} = 0. \end{aligned}$$

Notes. (i) Only the $t = 0$ recruiting constraint contributes in the first line (hence λ_0 , $V_{j,0}$, θ_0); for $t > 0$, $w_{j,0,0}$ does not enter $\mathcal{V}_{j,t,t}$. (ii) For $t \geq 1$, $w_{j,0,0}$ affects retention via $\mathcal{V}_{j,0,t}$ inside $S(\cdot)$, yielding the second-line terms. (iii) If you impose stationarity, you may replace θ_t with θ throughout.

$$\begin{aligned} \mathcal{L}_{V_{j,t}} : & - c_j (1 + \chi) V_{j,t}^{\chi} N_{j,t-1}^{\sigma-\chi} \\ & - \sum_{s=-\infty}^{t-1} \mu_{s,t} N_{j,s,t-1} \lambda_{EE} \frac{\partial f}{\partial \theta} \frac{\partial \theta_t}{\partial V_{j,t}} \sum_k \phi_{k,t}^v \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})} \\ & - \sum_{s=-\infty}^{t-1} \mu_{s,t} N_{j,s,t-1} \lambda_{EE} f(\theta_t) \sum_k \frac{\partial \phi_{k,t}^v}{\partial V_{j,t}} \frac{\exp(\gamma \mathcal{V}_{k,t,t})}{\exp(\gamma \mathcal{V}_{j,s,t}) + \exp(\gamma \mathcal{V}_{k,t,t})} \\ & + \lambda_t V_{j,t} \frac{\partial g}{\partial \theta} \frac{\partial \theta_t}{\partial V_{j,t}} \sum_k \sum_{s=-\infty}^{t-1} \phi_{k,s,t-1}^n \frac{\exp(\gamma \mathcal{V}_{j,t,t})}{\exp(\gamma \mathcal{V}_{j,t,t}) + \exp(\gamma \mathcal{V}_{k,s,t})} \\ & + \lambda_t g(\theta_t) \sum_k \sum_{s=-\infty}^{t-1} \phi_{k,s,t-1}^n \frac{\exp(\gamma \mathcal{V}_{j,t,t})}{\exp(\gamma \mathcal{V}_{j,t,t}) + \exp(\gamma \mathcal{V}_{k,s,t})} = 0. \end{aligned}$$

$$\begin{aligned} \mathcal{L}_{N_j,t,t} : \quad & \text{MRPL}_t - w_{j,t,t} + \beta_f c_j (\chi - \sigma) V_{j,t+1}^{1+\chi} N_{j,t}^{\sigma-\chi-1} - \lambda_t + \beta_f \mu_{t,t+1} (1 - S(\mathcal{V}_{j,t,t+1})) \\ & + \beta_f \lambda_{t+1} V_{j,t+1} g(\theta_{t+1}) \sum_k \sum_{s=-\infty}^t \frac{\partial \phi_{k,s,t}^n}{\partial N_{j,t,t}} \frac{\exp(\gamma \mathcal{V}_{j,t+1,t+1})}{\exp(\gamma \mathcal{V}_{j,t+1,t+1}) + \exp(\gamma \mathcal{V}_{k,t+1,t+1})} = 0. \end{aligned}$$

We will show that all firms choose the same wage when c_j is common and $\sigma = 0$ or A_j is common by positing symmetric wages $w_j = w_k = \bar{w}$ for all k and then verifying that a common wage is an equilibrium outcome.

In steady state we have $V/N = S/R$. From wage symmetry ($\mathcal{V}_j = \mathcal{V}_k \forall k$), we have that $R_j = \frac{g(\theta)}{2}(1 - \phi_j^n)$ and $S_j = \frac{\lambda_{EE} f(\theta)}{2}(1 - \phi_j^v)$. Further, under wage symmetry, $\phi_j^n = \phi_j^v \equiv \phi_j$, so S_j/R_j simplifies to :

$$\frac{S_j}{R_j} = \lambda_{EE} \theta$$

for all firms (using $f(\theta)/g(\theta) = \theta$).

With some algebra, it can be shown that the first order condition on wages simplifies to

$$\lambda \gamma \frac{1}{4} \frac{\partial \mathcal{V}_j}{\partial w_j} (1 - \phi_j) \left(\frac{V_j}{N_j} g(\theta) + \frac{\lambda_{EE} f(\theta) S_j}{1 - \beta_f (1 - S_j)} \right) = \frac{S_j}{1 - \beta_f (1 - S_j)}.$$

Using that $V_j/N_j = S_j/R_j = \lambda_{EE} \theta$, as well as that $g(\theta)\theta = f(\theta)$, this expression simplifies to

$$\lambda \gamma \frac{1}{4} \frac{\partial \mathcal{V}_j}{\partial w_j} (1 - \phi_j) \lambda_{EE} f(\theta) \left(1 + \frac{S_j}{1 - \beta_f (1 - S_j)} \right) = \frac{S_j}{1 - \beta_f (1 - S_j)}. \quad (\text{B.28})$$

If wages are equal at all firms, then the first order condition on vacancies becomes

$$\begin{aligned} c_j(1 + \chi) V_j^\chi N_j^{\sigma-\chi} = & -\lambda \frac{N}{V} \lambda_{EE} f(\theta) \varepsilon_{f,\theta} \varepsilon_{\theta,V_j} (1 - \phi_j) \frac{1}{2} - \lambda \frac{N_j}{V_j} \phi_j \lambda_{EE} f(\theta) (\phi_j - 1) \frac{1}{2} \\ & + \lambda g(\theta) \varepsilon_{g,\theta} \varepsilon_{\theta,V_j} (1 - \phi_j) \frac{1}{2} + \lambda g(\theta) (1 - \phi_j) \frac{1}{2} \end{aligned}$$

Combining terms and using that $V/N = \lambda_{EE} \theta$, we therefore have $\frac{N}{V} = \frac{R}{S} = \frac{1}{\lambda_{EE} \theta}$, and $\varepsilon_{f,\theta} = 1 + \varepsilon_{g,\theta}$ (from the matching function being constant returns to scale), the first order condition on vacancies simplifies down to

$$c_j(1 + \chi) V^\chi N^{\sigma-\chi} = \lambda \frac{1}{2} (1 - \phi_j) g(\theta) \quad (\text{B.29})$$

Plugging this in to (B.28) yields

$$c_j(1 + \chi) V^\chi N^{\sigma-\chi} \gamma \frac{1}{2} \frac{\partial \mathcal{V}_j}{\partial w_j} \lambda_{EE} \theta \left(1 + \frac{S_j}{1 - \beta_f (1 - S_j)} \right) = \frac{S_j}{1 - \beta_f (1 - S_j)} \quad (\text{B.30})$$

Combining Elements As $\beta_f \rightarrow 1$, this becomes

$$c_j(1 + \chi)V_j^\chi N_j^{\sigma-\chi} \gamma \lambda_{EE} \theta \frac{\partial V_j}{\partial w_j} = 1 \quad (\text{B.31})$$

Recalling that $\lambda_{EE} \theta = S/R = V/N$, this expression becomes

$$\gamma c_j(1 + \chi)(\lambda_{EE} \theta)^{1+\chi} N_j^\sigma \frac{\partial V_j}{\partial w_j} = 1. \quad (\text{B.32})$$

Thus if $\sigma = 0$ and c_j are common among firms, then the wage is identical for all firms, regardless of their value of A_j and their size N_j , proving the first claim in Proposition 2. Additionally, if firms are ex-ante identical (common values of c_j and A_j), they will choose identical wages. Note that under symmetry (common values of $\mathcal{V}$) and no unemployment,

$$\begin{aligned} \varepsilon_{R,w_j} &\equiv \frac{\partial R}{\partial w_j} = \frac{\gamma}{2} \frac{\partial \mathcal{V}_j}{\partial w_j} \\ \varepsilon_{S,w} &\equiv \frac{\partial S}{\partial w_j} = -\frac{\gamma}{2} \frac{\partial \mathcal{V}_j}{\partial w_j} \end{aligned}$$

and so $\varepsilon_{R,w_j} - \varepsilon_{S,w_j} = \gamma \frac{\partial \mathcal{V}_j}{\partial w_j} w_j$. Thus equation (B.32) turns into

$$c_j(1 + \chi)(\lambda_{EE} \theta)^{1+\chi} N_j^\sigma (\varepsilon_{R,w_j} - \varepsilon_{S,w_j}) = w_j.$$

Turning to employment, if all firms pay the same wage, and recognizing that $\sum_k \sum_{s=-\infty}^t \frac{\partial \phi_{k,s,t}^n}{\partial N_{j,t}} = 0$, the first order condition on employment simplifies to

$$MRPL_j - w_j + c_j(\chi - \sigma)V_j^{1+\chi} N_j^{\sigma-\chi-1} - \lambda S_j = 0. \quad (\text{B.33})$$

Combining (B.33) with (B.29) and (B.32), as well as using that $R_j = \frac{g(\theta)}{2}$ under wage symmetry, $V_j/N_j = \lambda_{EE} \theta \forall j$, and $\phi_j^n = \phi_j^v$, with some algebra yields our expression for the firm's recruiting cost-adjusted markdown

$$\frac{w_j + c_j \left(\frac{V_j}{N_j} \right)^{1+\chi} N_j^\sigma}{MRPL_j} = \frac{1 + (1 + \chi) (\varepsilon_{R,w_j} - \varepsilon_{S,w_j})}{1 + (1 + \chi) (\varepsilon_{R,w_j} - \varepsilon_{S,w_j}) + \sigma + (1 + \chi) \frac{\phi_j^n}{1 - \phi_j^n}}, \quad (\text{B.34})$$

establishing the second claim of Proposition 2.

B.11 Bounds on Markdowns with Oligopsony

In this section, we prove that under two allocations that deliver the same labor market HHI, the aggregate recruiting cost-adjusted markdown $\mathcal{M}$ is the lowest when that HHI is achieved with one large firm. Let ϕ_j^n be firm j 's share of employment. The employment HHI is $HHI = \sum_j (\phi_j^n)^2$.

Consider two hypothetical allocations: one in which a given HHI is achieved by one firm: $(\phi_1^n)^2 = \overline{HHI}$, where all other firms are sufficiently small such that they do not contribute to the HHI.

Consider a second allocation in which two firms achieve that same HHI : $(\phi_2^n)^2 + (\phi_3^n)^2 = \overline{HHI}$, and all other firms are sufficiently small. This gives $(\phi_1^n)^2 = (\phi_2^n)^2 + (\phi_3^n)^2$.

Without loss of generality, we can assume $\phi_3^n \leq \phi_2^n$. The aggregate markdown is

$$\mathcal{M} = \frac{1}{\sum_j \phi_j^n / \mu_j}, \quad \text{with} \quad \mu_j = \frac{1 + \gamma(1 + \chi)}{1 + \gamma(1 + \chi) + \sigma + (1 + \chi) \frac{\phi_j^n}{1 - \phi_j^n}}$$

Thus for a given $\overline{HHI}$ and ϕ_2^n , and focusing on the case of $\sigma = 0$, the aggregate markdown can be written as

$$\mathcal{M} = \frac{1 + \gamma(1 + \chi)}{1 + \gamma(1 + \chi) + (1 + \chi) \left(\frac{(\phi_2^n)^2}{1 - \phi_2^n} + \frac{(\phi_1^n)^2 - (\phi_2^n)^2}{1 - ((\phi_1^n)^2 - (\phi_2^n)^2)^{.5}} \right)}$$

The goal is to show that this expression is decreasing in ϕ_2^n . Define

$$A(\phi_2^n) = \frac{(\phi_2^n)^2}{1 - \phi_2^n} + \frac{(\phi_1^n)^2 - (\phi_2^n)^2}{1 - ((\phi_1^n)^2 - (\phi_2^n)^2)^{.5}}.$$

Then we just need to show that $\mathcal{M}$ is a decreasing function of ϕ_2^n if A is increasing in ϕ_2^n . With some algebra, we have

$$\frac{\partial A(\phi_2^n)}{\partial \phi_2^n} = \phi_2^n \left(\frac{(2 - \phi_2^n)}{(1 - \phi_2^n)^2} - \frac{(2 - \phi_3^n)}{(1 - \phi_3^n)^2} \right)$$

The function $\frac{2-x}{(1-x)^2}$ is increasing for $x \in [0, 1)$, so since $\phi_3^n \leq \phi_2^n$ we have $\frac{\partial A}{\partial \phi_2^n} > 0$. Thus $\mathcal{M}$ is decreasing in ϕ_2^n (for fixed $\overline{HHI}$), and is therefore minimized when ϕ_2^n is as large as possible—i.e., when the given HHI is achieved by a single large firm.

B.12 Calibration with Forward-Looking Workers

In Section 2.2, we assumed that workers are myopic $\beta_w = 0$ and that there are no exogenous separations $s_0 = 0$, as this delivered a constant sum of recruiting and separation elasticities. In this section, we examine how the model changes if we choose a more realistic calibration. We now set workers' discount factor $\beta_w = .99$, and we set the exogenous separation rate $s = 0.025$ to get a steady-state unemployment rate between 4 and 5 percent. We assume firms are ex-ante identical.

The equilibrium wage is 0.657. Figure B.4 plots what the recruiting and separation rates would be if a firm chose alternative wages. Consistent with evidence from Hirsch et al. (2022) and Langella and Manning (2021), the recruiting and separation elasticities are highest in the middle of the wage distribution. Because some separations are exogenous (and therefore are not wage sensitive), the separation elasticity declines at high wages because the only separations are exogenous and marginal changes in wages do not affect separation probabilities very much. At low wages, the separation elasticity is low because workers leave as soon as they find another job, and marginal changes in the wage have little effect on that separation decision. The recruiting elasticity is low at low wages for similar reasons: the only workers who accept jobs are the wage-insensitive unemployed workers.


Table B.3: Calibration with Forward Looking Workers

Parameter		Value	Reason
γ	Inverse scale parameter of non-wage preferences	1/4.225	Match $\varepsilon_R - \varepsilon_S = 4$
η	Worker CRRA parameter	.375	$\xi_{R,w} \approx \xi_{S,w}$
s_0	Exogenous separation rate	.025	Match $U = 4.3$
c	Hiring cost constant	892	Match $S = .05$
χ	Hiring cost-vacancy rate elasticity	2.2	Evidence in Section 2.3
σ	Hiring cost-size elasticity	0	Evidence in Section 2.3
λ_{EE}	On-the-job search probability	.08	Match J-J mobility
b	Unemployment benefit	.1	Workers accept low wage offers

At high wages, the recruiting elasticity is low because all workers will accept the job if they find it, and marginal changes in wages have little effect.

The key difference in calibration is the value of γ : as workers become more forward-looking, the variance of non-wage preference draws can be rescaled to maintain $\varepsilon_{R,w} - \varepsilon_{S,w}$ around 4. For this calibration, we set $\gamma = 0.2367$.

The other difference in calibration is η , the CRRA preference parameter in worker's flow utility. This is calibrated to offset the option value that would arise if workers have log utility and are forward-looking. With log utility, the feature of workers being forward-looking disproportionately lowers recruiting and separation elasticities at low wage levels. This is because of option value: workers know they will not be in low-wage jobs as long, and so marginal changes in wages have smaller effects on workers' value functions at low wages compared to at high wages. This feature makes recruiting and separation elasticities rise with wages, all else equal. This can be offset by making workers more risk averse: with a lower value of η , workers care about marginal increases of wages less as wages go up.

Figure B.4: Recruiting and Separation Elasticities for Different Wage Levels $\beta^w = .99$, $s_0 = .025$ 

Notes: This figure plots the recruiting elasticity $\varepsilon_{R,w}$ and separation elasticity $\varepsilon_{S,w}$ for different wages if all other firms pay the equilibrium wage of $w = 0.6322$. Both the recruiting and separation elasticities peak near the equilibrium wage and become small for high and low values of the wage.

In this calibration we have $\varepsilon_{R,w} = -\varepsilon_{S,w} = -\xi_{R,w} = \xi_{S,w} = 2$ and $S = 0.05$. Thus the values χ and σ in Section 2.3 would be internally consistent with an equilibrium from this calibrated model.

B.13 Amenities

In this section, we sketch a model with on-the-job search, costly amenities, but no recruiting margin. We argue that the inclusion of amenities can create an infinitely elastic labor supply curve only in the special case where there are no exogenous separations ($s_0 = 0$). This is in contrast with the model with a recruiting margin but no amenities in the main text. The key difference is that amenities do not help firms overcome matching frictions.

We modify the model in the main text so that (i) firms post constant wages for each cohort by assumption and (ii) workers now get utility from the level of amenities a_j when employed at firm j , with relative wage ξ . The value of a worker employed at firm j is

$$\mathcal{V}_j = \log w_j + \xi \log a_j + \beta_w \left(f(\theta') \lambda_{EE} \sum_{k \neq j}^{\mathcal{K}} \phi_k^{v'} \gamma^{-1} \phi_k^{v'} \log \left(\exp(\gamma \mathcal{V}_j') + \exp(\gamma \mathcal{V}_k') \right) + (1 - f(\theta') \lambda_{EE}) \mathcal{V}_j' \right),$$

where $\phi_k^{v'}$ now indicates the shares of vacancies associated with a $\{w, a\}$ bundle. The firm's problem is (suppressing the j):

$$\max_{\{\{N_{s,t}\}\}, \{w_t\}, \{a_t\}} \sum_{t=0}^{\infty} \beta^\tau \left(A \left(\sum_{s=-\infty}^t N_{s,t} \right)^\alpha - \sum_{s=-\infty}^t w_s N_{s,t} - c_a \frac{\sum_{s=-\infty}^t a_s N_{s,t}}{\left(\sum_{s=-\infty}^t N_{s,t} \right)^\psi} \right)$$

subject to

$$\begin{aligned} N_{s,t} &= (1 - S(w_s, a_s))N_{s,t-1} \\ N_{t,t} &= R(w_s, a_s)V_t. \end{aligned}$$

Here, $\psi \in [0, 1)$ allows for returns to scale in providing amenities: as the firm gets bigger, it may become cheaper to provide the same level of benefits per worker.

While a firm can increase the probability that workers accept a job and decrease endogenous separations with higher wages and better amenities, a greater level of amenities does not increase the number of matches, as search is random. Since random search limits worker inflows, any level of exogenous separations puts a hard limit on employment. Define V_0 as the exogenous amount of recruiting activity available to the firm. As in the main text, in steady state, inflows must equal outflows: $SN = VR$. As the firms offer better wages and amenities, $R \rightarrow 1$ and $S \rightarrow s_0$. The maximum employment of any firm is therefore $N^{max} = V_0/s_0$, and the marginal cost of employment becomes infinite because higher compensation (in the form of either wages or amenities) yields no additional workers.

B.14 Extension with Worker Heterogeneity

In this section, we show that we can extend the model in Section 2.2 where workers differ in ability and firms pay wage the same wage per efficiency unit of labor if workers have log flow utility.

Workers Worker value functions will be the same as before, except they will be indexed by worker type a . a is both an index for worker type and the worker's efficiency units of labor when employed. When unemployed, workers get flow income ba , where b is the unemployment benefit of workers type $a = 1$. Workers search and firms post vacancies in type-specific submarkets indexed by a . The value of an employed worker i who is type a and earns $w_{s,j,t}(a)$, and unemployed worker are, respectively,

$$\begin{aligned} \mathcal{V}_{j,s,t}(a) &= \log(w_{j,s,t}(a)) + \beta_w s_0 \mathcal{V}_U(a) \\ &\quad + \beta_w (1 - s_0) \mathbb{E}_t \left[\lambda_{EE} f(\theta^a) \int v_k \gamma^{-1} \log(\exp(\gamma \mathcal{V}_{j,s,t+1}(a)) + \exp(\gamma \mathcal{V}_{k,t+1,t+1}(a))) d\mathcal{V}_k \right. \\ &\quad \left. + (1 - \lambda_{EE} f(\theta^a)) \max\{\mathcal{V}_{j,s,t+1}(a), \mathcal{V}_U(a)\} \right], \\ \mathcal{V}_U(a) &= \log(b(a)) + \beta_w \mathbb{E}_t \left[f(\theta^a) \int_{\mathcal{V}_k \geq \mathcal{V}_U} v_k \mathcal{V}_{k,t+1,t+1}(a) d\mathcal{V}_k + (1 - f(\theta^a)) \mathcal{V}_U(a) \right]. \end{aligned}$$

Total Employment and Recruiting Costs Let total labor inputs and total recruiting effort for firm j be

$$N_{j,t} = \sum_{s=-\infty}^t \sum_a a N_{j,s,t}^a$$

$$V_{j,t} = \sum_a a V_{j,t}^a$$

Firms post vacancies $V_{j,t}^a$ in different submarkets a . Firm j 's total cost of recruiting effort is

$$C_{j,t} = c_j \left(\sum_a a V_{j,t}^a \right)^{1+\chi} \left(\sum_{s=-\infty}^t \sum_a a N_{j,s,t}^a \right)^{\sigma-\chi}.$$

Firm's Problem Dropping the j subscript to save space, the firm's problem is:

$$\begin{aligned} & \max_{\{V_t^a, K_t, N_t, \{N_{s,t}^a\}_{a,s \leq t}, \{w_{s,t}^a\}_{a,s \leq t}\}_{t=0}^{\infty}} \\ & \sum_{t=0}^{\infty} \beta_f^t \left(A_t (K_t^\alpha N_t^{1-\alpha})^{\frac{\epsilon-1}{\epsilon}} - \sum_{s=-\infty}^t \sum_a w_{s,t} N_{s,t}^a - r^K K_t - c \left(\frac{V_t}{N_{t-1}} \right)^\chi N_{t-1}^\sigma V_t \right) \end{aligned} \quad (\text{B.35})$$

subject to

$$\begin{aligned} \text{Recruiting:} \quad & N_{t,t}^a = R_a(w_{t,t}) V_t^a \quad \text{for all } t \text{ and } a \\ \text{Retention:} \quad & N_{s,t}^a = (1 - S(w_{s,t})) N_{s,t-1}^a \quad \text{for all } t \text{ and } s < t \text{ and } a \\ \text{Employment is sum of cohorts:} \quad & N_t = \sum_{s=-\infty}^t N_{s,t} \quad \text{for all } t \\ \text{Cohort Efficiency Units:} \quad & N_{s,t} = \sum_a a N_{s,t}^a \quad \text{for all } t \\ \text{Total Recruiting:} \quad & V_t = \sum_a a V_t^a \quad \text{for all } t \\ \text{Wage at least posted wage:} \quad & w_{s,t}^a \geq w_{s,s}^a \quad \text{for all } t \text{ and } s < t \text{ and } a \end{aligned}$$

Candidate equilibrium Define $\omega_j^a = \frac{w_j^a}{MRPL_j^a}$, where $MRPL_j^a$ is the marginal product of labor of worker type a at firm j .

Property 1: Firm j pays the same wage per efficiency unit to workers of all types: $w_{j,s,t}^a = \omega_{j,s,t}^a a$.

Property 2: The value function of a worker type a takes the following form:

$$\begin{aligned} (\bar{V}_{j,s,t} + \mathcal{A}) &= \log(\omega_{j,s,t}^a) + \log(a) + \beta_w s_0 (\bar{V}_U + \mathcal{A}) \\ &+ \beta_w (1 - s_0) \mathbb{E}_t \left[\lambda_{EE} f(\theta) \int v_k \gamma^{-1} \log(\exp(\gamma(\bar{V}_{j,s,t+1} + \mathcal{A})) + \exp(\gamma(\bar{V}_{k,t+1,t+1} + \mathcal{A}))) dk \right. \\ &\quad \left. + (1 - \lambda_{EE} f(\theta)) \max\{\bar{V}_{j,s,t+1} + \mathcal{A}, \bar{V}_U + \mathcal{A}\} \right] \end{aligned}$$

where $\mathcal{A} = \frac{\log(a)}{1-\beta^w}$ and $\bar{\mathcal{V}}_{j,s,t}$ is the value for a worker with $a = 1$.

Property 3: Worker's job acceptance probability between two jobs is a function of ω but not a . For a worker in firm j considering firm k , the probability of staying in firm j is

$$\begin{aligned} P[\text{accept } j \text{ over } k] &= \frac{\exp(\gamma \mathcal{V}_{j,s,t}(a))}{\exp(\gamma \mathcal{V}_{j,s,t}(a)) + \exp(\gamma \mathcal{V}_{k,t,t}(a))} \\ &= \frac{\exp(\gamma (\bar{\mathcal{V}}_{j,s,t} + \mathcal{A}))}{\exp(\gamma (\bar{\mathcal{V}}_{j,s,t} + \mathcal{A})) + \exp(\gamma (\bar{\mathcal{V}}_{k,t,t} + \mathcal{A}))} \\ &= \frac{\exp(\gamma \bar{\mathcal{V}}_{j,s,t})}{\exp(\gamma \bar{\mathcal{V}}_{j,s,t}) + \exp(\gamma \bar{\mathcal{V}}_{k,t,t})}. \end{aligned}$$

Because employed workers have the same outside options and expect a constant wage, the difference between values $\bar{\mathcal{V}}_{j,s,t}$ and $\bar{\mathcal{V}}_{k,t,t}$ comes only from different values of ω .

Property 4a: For a given firm j , recruiting and separation rates are identical for workers of different a .

Property 4b: Labor market tightness is identical across submarkets a .

Property 4c: Firms are indifferent between recruiting in different submarkets.

These properties come from the symmetry across submarkets. The cost of recruiting workers with a higher value of a scales linearly with a . Since their recruiting and separation elasticities are the same, the firm chooses the same markdown for all worker types. Thus the expected flow profit per efficiency unit of labor and the cost per efficiency unit of labor across worker types is identical. Thus firms are indifferent to recruiting workers of different types. Common labor market tightness is sustained because firms can costlessly switch the composition of their recruiting activity across submarkets. If one submarket were differentially slack, firms would instantly switch recruiting to that market because the cost of labor per efficiency unit of labor would be lower.

AKM Interpretation In such an equilibrium, both in steady state and in response to shocks, the wages of new hires can be described statistically as being decomposed into firm-by-cohort fixed effect $\log \omega_{j,t,t}$ and a time invariant worker fixed effect $\log a(i)$:

$$\log w_t(i, j) = \log \omega_{j,t,t}(i) + \log a(i) + \epsilon_{i,j,t}$$

The left-hand side variable in the local projection specification in equation (2.15) of the main text is the percent deviation of the firm's wage policy from steady state. If firm j were in steady state in period $t - 1$, this can be written as

$$\begin{aligned} \Delta w_t(i, j, k) - \Delta w_{t-1}(i', j, k') &= (\log \omega_{j,t,t}(i) - \log \omega_{k,s,t-2}(i)) - (\log \omega_{j,t-1,t-1}(i') - \log \omega_{k',s',t-3}(i')) \\ &\quad + (\epsilon_{i,j,t} - \epsilon_{i,k,t-2}) - (\epsilon_{i',j,t-1} - \epsilon_{i',k',t-3}). \end{aligned}$$

The coefficient β_ν^w captures how this percent deviation in cohort-specific wage policy changes in response to a shock z_j at horizon ν . Unbiased estimation of β_ν^w is satisfied under three assumptions.


The first is standard strict exogeneity assumption, which says that worker mobility does not depend on firm-by-worker match effects at any time horizon. Second, we also require that firm-by-worker match effects are uncorrelated with export demand shocks. Third, we must control for the steady-state wage effects of the firms that workers leave (k and k'), which we do by residualizing wage changes on the leaving firm's AKM fixed effect.

Appendix C

Appendix to The Incidence and Efficiency of Land Value Taxation


C.1 Appendix Figures

Figure C.1: Mentions of Structural Reform in Major Danish Newspapers



Notes: This figure plots mentions of structural reforms in major Danish newspapers across time. Figure is taken from Christiansen and Klitgaard (2008).

Figure C.2: Land Taxes and the Allocation of Housing: Placebo Tests



Notes: The bincatters in this figure plots the long difference between 2001 and 1995 in log share of homeowners in a given treatment area on the land tax change instrument. Panel (A) plots the effect for homeowners who are under 35 years old. Panel (B) plots the same effect, but for homeowners who are over 60 years old. Panel (C) plots the effect for below-median income working homeowners, while Panel (D) plots the effect individuals in the top 10% of the income distribution. Point estimates and standard errors are shown, with standard errors clustered at the new municipality level.

C.2 Model with Rent Ceiling

We build a neoclassical model of land-value taxation with a price ceiling on rents to provide a possible interpretation of our empirical results. The economy consists of households, landlords, producers, and government.

Households

Households consume non-durables C_t and supply labor N_t . They rent land L_t from landlords.¹ They have access to a real financial asset A_t . They receive lump sum transfers T_t and profits Π_t from the goods producer (both of which they take as given). They maximize utility

$$U_0 = \sum_{t=0}^{\infty} \beta^t u(C_t, N_t, L_t)$$

subject to

$$C_t + P_t^L L_t + A_t = W_t N_t + (1 + r_{t-1})A_{t-1} + T_t + \Pi_t.$$

The first-order conditions are

$$u_{C,t} = \beta(1 + r_t)u_{C,t+1} \quad (C.1)$$

$$-\frac{u_{N,t}}{u_{C,t}} = W_t \quad (C.2)$$

$$\frac{u_{L,t}}{u_{C,t}} = R_t \quad (C.3)$$

where $u_{X,t} \equiv \frac{\partial u}{\partial X}(C_t, N_t, L_t)$.

We note that equation (C.3) determines the demand for land services (i.e. renting land). With a binding rent ceiling land will be in excess demand. The equilibrium quantity of land will be determined by the intersection of the supply curve and the rent ceiling, so the household will be off their demand curve. However, we can still use (C.3) to learn the amount demanded at the prevailing price (the rent ceiling).

Landlords

Landlords own land L_t and rent it to households at rental price R_t . Landlords can trade land at price P_t^L and they pay land value tax at rate τ^L . They take prices as given and choose L_t to maximize

$$\sum_{t=0}^{\infty} q_t (R_t L_t - P_t^L (L_t - L_{t-1}) - \tau^L P_t^L L_t)$$

where $q_0 \equiv 1$ and $q_t = \prod_{s=1}^t \left(\frac{1}{1+r_s} \right)$ for $t \geq 1$. The first order condition is

$$R_t = (1 + \tau^L)P_t^L - \frac{P_{t+1}^L}{1 + r_t}. \quad (C.4)$$

This equation tell us that landlords will supply land services to households until the rental price R_t equals the user cost of land.

¹For simplicity, this model considers only land (not buildings).

Producers

Producers hire labor N_t and produce a non-durable output good Y_t with the technology

$$Y_t = F(N_t). \quad (\text{C.5})$$

They are price-takers in input and output markets, and they maximize profits $\Pi_t = Y_t - W_t N_t$. The first-order condition is

$$F_N(N_t) = W_t. \quad (\text{C.6})$$

Government

The government collects land taxes $\tau^L P_t^L L_{t-1}$ from landlords at the beginning of each period and pay lump sum taxes T_t to households. We assume

$$T_t = 0 \quad (\text{C.7})$$

i.e. the government “throws the revenue in the ocean”.

Market Clearing and Rent Ceiling

The goods and asset market must clear:

$$C_t + R_t \bar{L} = Y_t \quad (\text{C.8})$$

$$A_t = 0 \quad (\text{C.9})$$

The aggregate amount of land in the economy is fixed:

$$L_t = \bar{L}$$

There is a price ceiling on the rental price of land:

$$R_t \leq \bar{R} + \tau^L P_t^L. \quad (\text{C.10})$$

The rent ceiling increases with land taxes: When the landlords’ tax bill increases, the ceiling increases by the same amount. We will assume the rent ceiling binding in steady state. (We can verify this by checking that there excess demand for land services.)

Functional Forms

We assume

$$u(C, N, L) = \log C_t + \vartheta \frac{N^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}} + \varphi \log L_t,$$

$$F(N) = Z_Y N^\gamma.$$

Model Equations

The model features 11 variables ($C_t, N_t, L_t, A_t, Y_t, R_t, P_t^L, W_t, r_t, T_t, L_t^D$) and 11 equations:

$$\frac{1}{C_t} = \beta(1 + r_t) \frac{1}{C_{t+1}} \quad (\text{C.11})$$

$$\vartheta N_t^{\frac{1}{\chi}} C_t = W_t \quad (\text{C.12})$$

$$\varphi \frac{C_t}{L_t^D} = R_t \quad (\text{C.13})$$

$$R_t = (1 + \tau^L) P_t^L - \frac{P_{t+1}^L}{1 + r_t} \quad (\text{C.14})$$

$$Y_t = Z_Y N_t^\gamma \quad (\text{C.15})$$

$$\gamma \frac{Y_t}{N_t} = W_t \quad (\text{C.16})$$

$$T_t = 0 \quad (\text{C.17})$$

$$L_t = \bar{L} \quad (\text{C.18})$$

$$C_t + R_t L_t = Y_t \quad (\text{C.19})$$

$$A_t = 0 \quad (\text{C.20})$$

$$R_t = \bar{R} + \tau^L P_t^L \quad (\text{C.21})$$

Steady State

Since we are studying a permanent change in land taxes, we are primarily interested in the steady state. We denote steady values of the variable simply by removing the t -subscript. We are particularly interested in how land rents R and the price of land P^L changes after an increase in the land value tax rate τ^L . We solve for the steady state of the model as follows:

1. From the Euler equation, we have $r = \frac{1}{\beta} - 1$. We also immediately have $L = \bar{L}$, $A = 0$, $T = 0$.
2. Combine the landlord's FOC with the rent ceiling:

$$\bar{R} + \tau^L P^L = (1 + \tau^L) P^L - \frac{P^L}{1 + r} \Rightarrow P^L = \frac{1 + r}{r} \bar{R}.$$

3. The rental price of land is then

$$R = \bar{R} + \tau P^L = \bar{R} + \tau \frac{1 + r}{r} \bar{R} = \frac{(1 + \tau)(1 + r) - 1}{r} \bar{R} \approx \frac{r + \tau}{r} \bar{R}.$$

4. Labor market clearing implies $\vartheta N^{\frac{1}{\chi}} C = \gamma Z_Y N^{\gamma-1}$. Goods market clearing implies $C = Z_Y N^\gamma - R \bar{L}$. Combining these to eliminate C yields

$$\vartheta N^{\frac{1}{\chi}} (Z_Y N^\gamma - R \bar{L}) = \gamma Z_Y N^{\gamma-1}.$$

We solve this equation numerically.

5. Once we know N , it is straightforward to back out C , Y , W .
6. Finally, we can check the demand for land services and verify that there is excess demand $L^D > \bar{L}$ (so our initial assumption that the rent ceiling is binding is valid).

The Effect of Land Tax Increase on Land Prices and Rents

How does the steady state change when the land value tax rate τ^L increases? The price of land P^L does not change: It depends only on the interest rate r and the fixed part of the rent ceiling $\bar{R}$. When the tax rate increases, landlords pass through the tax increase to tenants (the households) through higher rents. This happens because the rent ceiling increases and there is excess demand for land.

C.3 Comparison with Høj et al. (2018)


In this appendix, we explain potential differences between our results and those in Høj et al. (2018). Høj et al. (2018) study the effect of the same reform on prices of single-family homes and find large price effects. They use a different regression than the one we do; for example, they do not instrument for the change in the land tax rate, do not use a repeat-sales index, do not restrict the treatment group to border towns, and only examine the years immediately surrounding the reform, from 2001 - 2008. However, we argue that these differences in the empirical strategy do not drive the differences between our results and those found by Høj et al. (2018). Instead, we argue that a confound likely drives the differences between their results and ours.

Appendix Figure C.3 shows a binscatter plot of a regression of the actual change in the land tax rate and the number of vacation homes in a treatment area, after residualizing on new municipality fixed effects. The share of vacation homes in an area is strongly predictive of the land tax change. Areas that experienced land tax decreases had significantly more vacation homes than areas that experienced tax increases. Due to this correlation, in our main specification, we control for the share of vacation homes in an area and drop vacation homes from our baseline sample.

Appendix Figure C.4 shows an event study that uses variation in land tax changes alone (the policy shock) on the prices of single-family homes. In this specification, we keep all of our baseline controls and primary sample, except we drop our control for the share of vacation homes in the treatment area and add vacation homes to the sample. In this specification, we find large, persistent effects of land taxes on home prices. Our baseline results are actually quite similar to those in Høj et al. (2018), multiplying our ITT effect by our first stage results in a discount rate of $\approx 2.4\%$. This is nearly identical to what Høj et al. (2018) estimate for 2007. However, their treatment effect falls by nearly half in 2008, relative to 2007. In contrast, our estimates are quite persistent through 2019.


Appendix Figure C.5 shows the results of the exact same specification as Appendix Figure C.4 after dropping vacation homes and adding in our control for the pre-existing vacation home share in a treatment area. These two choices are enough to give us the baseline result that land taxes have no effect on sales prices. We once again get a precise zero effect, and can rule out with a high degree of confidence the point estimate from Høj et al. (2018).

Figure C.3: Land Tax Changes Strongly Correlated With Share of Vacation Homes in Treatment Area




Notes: This figure plots the first stage regression of the land tax instrument on the actual change in effective taxes paid, normalized by the assessed value of the property, from 2005 to 2007. For both the x- and y-axes, the units are in percentage points: thus, a value of 1 represents a 1 p.p. change in the tax rate, normalized by the assessed value. The blue line represents the 45-degree line. Beta coefficients and 95% confidence intervals, with standard errors clustered at the treatment area by land share decile level, are plotted.

Figure C.4: ITT Effect of Land Taxes Including Vacation Homes and Without Control



Notes: This figure plots the effect of land taxes on sales prices, including vacation homes and without the treatment area vacation share control. 95% confidence intervals are shown for the event-study coefficients, with standard errors clustered at the treatment area by land share decile level. The red line shows the difference in difference estimate that pools all post-2007 years compared to the price in 2001, and the red shading shows the 95% confidence intervals from that regression.

Figure C.5: ITT Effect of Land Taxes Without Vacation Homes and Including Control



Notes: This figure plots the effect of land taxes on sales prices. 95% confidence intervals are shown for the event-study coefficients, with standard errors clustered at the treatment area by land share decile level. The red line shows the difference in difference estimate that pools all post-2007 years compared to the price in 2001, and the red shading shows the 95% confidence intervals from that regression.