

ESSAYS ON FIRM-LEVEL DISTORTIONS
AND AGGREGATE PRODUCTIVITY

Gorkem Bostanci

A DISSERTATION

in

Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2021

Supervisor of Dissertation

Co-Supervisor of Dissertation

Guillermo L. Ordoñez

Professor of Economics and Finance

Harold L. Cole

James Joo-Jin Kim

Professor of Economics

Graduate Group Chairperson

Jesús Fernández-Villaverde

Professor of Economics

Dissertation Committee

Aviv Nevo, George A. Weiss and Lydia Bravo Weiss University Professor

ESSAYS ON FIRM-LEVEL DISTORTIONS
AND AGGREGATE PRODUCTIVITY

Copyright ©

Gorkem Bostanci

2021

ACKNOWLEDGEMENTS

I am deeply grateful to my advisors Guillermo Ordoñez and Harold Cole and my committee member Aviv Nevo for their continued support, guidance, and trust. I would also like to thank Iourii Manovskii, Guido Menzio, José-Victor Ríos-Rull, and the rest of the Penn Econ faculty for their invaluable input on my development as a researcher.

I would like to acknowledge the support of Gina Conway, Kelly Quinn, Sandra Venne, and Cortni White, who made the Ph.D. as easy as possible for my friends and me. I also want to thank Andre Victor, Gokhan, Harun, Minsu, and Paolo for their help and advice from the very beginning.

I am grateful for the company of Gabby, Joao, Sergio, and Tomas who added great joy to my six years. I always felt at home with Can, Ipek, Omer, and Umut. I am thankful for having met Camille, Hanna, and Montse, whose friendships have become invaluable.

I thank Merve for being by my side any way she could. I am grateful to my parents Aytekin and Turkan and my brother Goktan for their continued support.

ABSTRACT

ESSAYS ON FIRM-LEVEL DISTORTIONS AND AGGREGATE PRODUCTIVITY

Gorkem Bostanci

Guillermo Ordoñez

Harold Cole

This thesis focuses on how the frictions at the firm-level production decisions impact aggregate productivity.

The first chapter quantifies the impact of trade secret protection on labor outsourcing, and consequently, on aggregate productivity. First, using event studies and differences-in-differences around the staggered adoption of the Uniform Trade Secrets Act, I show that better trade secret protection leads to increased outsourcing. Second, to quantify the resulting gains in productivity, I build a structural model of outsourcing and multi-industry dynamics and estimate it with data from the U.S. manufacturing sector. I decompose the cross-state differences in labor outsourcing into differences in firing cost, industry composition, demand volatility, and trade secret protection. Strengthening trade secret protection for all states to match the state with the strictest protection would increase the outsourcing employment by 29% and aggregate output by 0.8%.

The second chapter studies the role of information frictions by measuring how the informativeness of the stock prices changes with business cycles. We first build a stock market model in which both the information content and the noise in prices respond to changes in economic activity, affecting how well those prices reflect firm's performance. Then we incorporate this module in a dynamic model with heterogeneous firms to characterize how stock price informativeness and capital misallocation interact with one another. We find that an increase in liquidity concerns can simultaneously boost information production, decrease stock price informativeness, and increase capital misallocation.

The third chapter examines the strong positive correlation between job-to-job transition rates and nominal wage growth in the U.S. First, using time series regressions, structural monetary policy shocks, and survey data on search effort we provide evidence that inflationary shocks cause higher job-to-job transitions in the subsequent years. Second, we build a model with aggregate shocks and competitive on-the-job search in which wages react sluggishly to inflation. Third, we calibrate the model to the U.S. economy and find that the output response to inflation shock is non-monotonic. The monetary authority can stimulate productivity with an inflationary shock through job-to-job transitions.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iii
ABSTRACT	iv
LIST OF TABLES	xi
LIST OF ILLUSTRATIONS	xiii
1 Productivity Gains from Labor Outsourcing: The Role of Trade Secrets	1
1.1 Introduction	1
1.2 Background	8
1.2.1 Trade Secrets	8
1.2.2 Trade Secret Protection in the U.S.	9
1.2.3 Trade Secret Protection and Labor Outsourcing	10
1.3 Empirical Analysis	12
1.3.1 Facts on Domestic Labor Outsourcing	13
1.3.2 Evidence on the Effect of Trade Secret Laws	17
1.4 A Model of Outsourcing and Trade Secret Protection	25
1.4.1 Environment	26
1.4.2 Intermediate Firm's Dynamic Problem	33

1.4.3	Equilibrium	34
1.4.4	Discussion of the Model Elements	34
1.4.5	Outsourcing Choice	36
1.4.6	Model Extensions	37
1.5	Estimation	38
1.5.1	Data and Estimation Method	38
1.5.2	Estimation Results	43
1.5.3	Model Validation	46
1.5.4	Decomposition of the Outsourcing Heterogeneity	47
1.6	Productivity Gains from Better Trade Secret Protection	49
1.7	Conclusion	53
	References	54
1.A	Proofs	66
1.B	Data Sources	70
1.B.1	Measures of Labor Outsourcing	70
1.C	Cross-Country Evidence	82
1.D	Estimation Details	89
1.E	Trade Secret Protection	94
1.E.1	Trade Secret Protection under Common Law	95
1.E.2	Non-Disclosure Agreements	96
1.E.3	Governing Law in Trade Secret Disputes	98
1.F	Generalized Differences-in-Differences Methods	100
1.G	Outsourcing and Trade Secrets	102
1.G.1	Government Regulations	102
1.G.2	Self-regulation	103
1.H	Additional Figures and Tables	106

2	Price Informativeness and Business Cycle Misallocation	114
2.1	Introduction	114
2.2	Static Model	119
2.2.1	Environment	120
2.2.2	Portfolio Choice Problem	121
2.2.3	Solution	122
2.2.4	Equilibrium Characterization	124
2.2.5	Comparative Statics	126
2.3	Dynamic Model	130
2.3.1	Environment	131
2.3.2	Recursive Competitive Equilibrium	133
2.3.3	Discussion of the Model Ingredients	136
2.3.4	Liquidity Needs and Misallocation	137
2.4	Conclusion	139
	References	139
2.A	Proofs	143
2.B	Comparative Statics	149
3	Changing Jobs to Fight Inflation: Labor Market Reactions to Inflationary Shocks	150
3.1	Introduction	150
3.2	Empirical Analysis	158
3.2.1	Predictive Regressions	158
3.2.2	Structural Monetary Policy Shocks	162
3.2.3	Survey Evidence on Search Effort	163
3.3	The Model	168

3.3.1	Environment	168
3.3.2	Equilibrium	172
3.3.3	Effect of a Decrease in Real Wage	175
3.4	Quantitative Analysis	177
3.4.1	Calibration Strategy	177
3.4.2	Unexpected Inflation Shock	180
3.5	Conclusion	181
	References	182
3.A	Data Sources	187
3.B	Evidence on the Extent of Wage Indexation	188
3.C	Proofs	190
3.D	Solution Method	192
3.E	A Random Search Model with Effort	193
3.E.1	Preferences	193
3.E.2	Production Technology	193
3.E.3	Meeting Technology	194
3.E.4	Wage Setting	195
3.E.5	Market Tightness	199
3.E.6	Problem of the Firm	200
3.E.7	Problem of the Worker	201
3.E.8	Distribution Accounting	203
3.E.9	Equilibrium	204
3.F	Additional Analysis	205

LIST OF TABLES

Productivity Gains from Labor Outsourcing: The Role of Trade Secrets	1
1.1 Two-way Fixed Effects Estimation	23
1.2 Overall Treatment Effect via Group-Time ATT Estimates	24
1.3 Placebo Regressions	25
1.4 The Main Parameters and the Moments Used in the Estimation	44
1.5 The Counterfactual Results After an Improvement in Trade Secret Protection	50
1.6 Labor Outsourcing Sector in NAICS Classification	72
1.7 Labor Outsourcing Sector in Census 1990 Classification	74
1.8 Manufacturing Industry Groups (Chan, 2017) for 2-digit NACE and 3-digit NAICS Classifications	80
1.9 Cross-Country Panel Regressions	88
1.10 Externally Calibrated Industry-level Parameters	106
1.11 First-stage Estimation Results (Michigan)	107
1.12 State-Level Estimates for Trade Secret Protection	108
1.13 The Baseline and the Counterfactual Outsourcing to Payroll Ratios for States of the U.S.	112
1.14 The State-Level Counterfactual Results After an Improvement in Trade Secret Protection	113

Price Informativeness and Business Cycle Misallocation	114
2.1 Comparative Statics Benchmarks	128
2.2 Information Acquisition Possibilities	148
Changing Jobs to Fight Inflation: Labor Market Reactions to Inflation- ary Shocks	150
3.1 The Effect of Inflation on Search Effort	167
3.2 Calibrated Parameters	179

LIST OF ILLUSTRATIONS

Productivity Gains from Labor Outsourcing: The Role of Trade Secrets	1
1.1 Employment Trends in Multiple Industry Groups	14
1.2 The Cross-state Supply of and Demand for Labor Outsourcing	15
1.3 Event Study Estimates for the UTSA Adoption	20
1.4 The Task Allocation Problem of a Software Design Firm	28
1.5 The Estimation Results from the 1st Stage	45
1.6 The Estimation Results from the 2nd Stage	46
1.7 The Untargeted Moments	48
1.8 The State-Industry Level Gains	52
1.9 Employment Protection Laws and the UTSA	77
1.10 The Employment Share of the Labor Outsourcing Sector in 1970 and 2005	83
1.11 The Employment Share of the Labor Outsourcing Sector in 2017	84
1.12 The Labor Outsourcing Sector and Legal Protection	87
1.13 Trade Secret Protection and GDP per capita	89
1.14 The Distribution and the Coefficient of Variation for Outsourcing to Payroll Ratios	109
1.15 First Stage Estimates and Revenue Payroll Ratios	110
1.16 The Number of States that Adopted the UTSA (1980-2016)	111

Price Informativeness and Business Cycle Misallocation	114
2.1 Comparative Statics for the Benchmark Model with respect to γ . . .	127
2.2 Comparative Statics for the Benchmark Model and ‘Noise Trader’ Model with Respect to γ	130
2.3 Comparative Statics for the Model with Many Risky Assets	131
2.4 Comparative Statics for the Full Model	138
2.5 Comparative Statics for the Benchmark Model with Good Shocks . .	149
Changing Jobs to Fight Inflation: Labor Market Reactions to Inflation- ary Shocks	150
3.1 The Discrepancy Between the SPF Forecast and Realized Inflation . .	152
3.2 National Predictive Regressions	160
3.3 State-Level Predictive Regressions	162
3.4 Monetary Policy Shocks and Acceptance Rate	164
3.5 Histogram of ι_i Values	167
3.6 Impulse Responses to Inflation Shocks	180
3.7 Worker x in Firm y Matches with Firm $y' > y$	197
3.8 Worker x in Firm y Matches with Firm $y' \in [\tilde{y}, y]$	198
3.9 Worker x in Firm y Matches with Firm $y' < \tilde{y}$	198
3.10 The Discrepancy Between the MCS Forecast and Realized Inflation .	205
3.11 Predictive Regressions for Inflation Mistake and Acceptance Rate . .	206

Chapter 1

Productivity Gains from Labor Outsourcing: The Role of Trade Secrets

BY GORKEM BOSTANCI[†]

1.1 Introduction

Producers' demand for workers changes over time due to fluctuating demand for goods and the presence of tasks that are not performed frequently. Labor outsourcing allows producers to make quick adjustments to their workforce, bypassing hiring and firing costs. However, many jobs, which could be outsourced, also provide access to sensitive information. For example, accountants might see financial documents, machine operators might see product designs, and security guards might see visitor lists. Sharing such information with outsiders can be problematic if the legal environment does not

[†]University of Pennsylvania.

provide adequate protection for intellectual property. In such cases, producers will be reluctant to use outsourced workers, leading to an inefficiently small outsourcing sector, slower reallocation of workers, and reduced aggregate productivity.¹

In this paper, I quantify the impact that trade secret protection has on aggregate productivity by affecting the extent of outsourcing in the economy. To show that the legal environment impacts labor outsourcing, I first use the staggered adoption of the Uniform Trade Secrets Act (UTSA) among states of the U.S. Next, I develop and estimate a structural model of industry dynamics in which firms choose whether to use outsourced workers in each task. I use the estimated model to measure the impact of distorted outsourcing decisions on aggregate productivity. I find that if all states of the U.S. could protect trade secrets as well as the state with the strictest protection, the fraction of outsourced workers would increase by 29%, and aggregate output would increase by 0.8%.

The U.S. provides a good laboratory to study this question because it features considerable variation in both trade secret protection and the extent of outsourcing. First, for reasons that were exogenous to outsourcing, the switch to statutory law via the UTSA happened in different years for different states, creating the heterogeneity in protection. Second, the extent of outsourcing varies substantially, both over time and across states. The firms that provide labor-intensive services, which were historically done in-house, employed 11% of the U.S. labor force in 2018, yet this share was just over 3% in 1971. In 2018, these firms had an employment share of 14.3% in California (90th percentile) but only 7.6% in Wisconsin (10th percentile).

I start by documenting three main stylized facts on the patterns of labor outsourcing in the U.S. First, I show that the growth in outsourcing was not an artifact of

¹See Decker et al., 2020 for an exercise and an overview of the literature on the relation between input reallocation and aggregate productivity.

growth in industries that demand outsourcing more than others. Second, the growth in labor outsourcing is not accompanied by a similar growth in the outsourcing of physical goods. Third, the cross-state heterogeneity in demand for outsourced workers does not diminish once I compare the demand from more disaggregated industry groups. These facts motivate a state- and time-specific factor that determines the extent of labor outsourcing for all industries.

To understand the role of trade secret protection, I use the staggered adoption of the UTSA across U.S. states. First, using historical anecdotes and event studies, I argue that timing of the adoptions was exogenous to outsourcing patterns. Second, using difference-in-differences, I show that stronger trade secret protection has a positive and significant impact on the size of the labor outsourcing sector. Quantitatively, improvements in trade secret law explain 14% of the outsourcing share growth from 1971 to 1997, translating to 0.7 million new jobs in the outsourcing sector. Third, I supplement the relevance of shared information by showing that the impact was not significant for tasks that are (1) unlikely to involve sensitive information or (2) already subject to auxiliary enforcement through professional organizations.

To quantify the aggregate productivity gains, I develop and estimate a structural model of industry dynamics that is based on Hopenhayn, [1992](#). I augment the model in two dimensions. First, I incorporate a task-based production framework in which firms decide whether to use their employees or outsourced workers for each task. Unlike employees, the number of outsourced workers can be adjusted freely, but their productivity is limited by how much sensitive information is shared. The extent of trade secret protection determines which information can be shared without risking leaks and, thus, the tasks that can be feasibly outsourced. Second, I extend the model to accommodate multiple industries that use different technologies, including different tastes for outsourced labor. In total, the extent of outsourcing can differ *across states*

due to differences in four components: (1) strength of employment protection; (2) within-industry firm characteristics; (3) industry compositions; and (4) strength of trade secret protection.

I estimate the model using state-industry-level data from the U.S. manufacturing sector in 2007. I use establishment size distributions and job flows among others to identify the magnitude of firing costs and the parameters of the production technologies (components (1) and (2)). The fundamental identifying assumption for distinguishing (3) and (4) is that the comparative advantage of outsourced workers (e.g., specialized knowledge) depends on the industry but not on the state. In contrast, the extent of trade secret protection depends on the state, but not on the industry. My identification relies on parameters that are constant across states; hence, it requires estimating all state-industry pairs simultaneously. To make the estimation feasible, I continue in two stages. In the first stage, I use the method of moments to estimate the full model separately for each state under assumptions where the task-based production function simplifies to a CES aggregate of employees and outsourced workers. In the second stage, I treat the estimated CES factor shares as data and estimate the trade secret protection and outsourcing efficiency parameters separately using non-linear least squares. The estimated trade secret protection parameters are highly correlated with the UTSA adoption dates. I find the impact of differences in trade secret protection to be considerable. If all states had the same (average) level of trade secret protection, the cross-state dispersion of outsourcing would decline by 19%.

Using the model estimates, I ask how the extent of outsourcing and aggregate productivity would change if all states enforced trade secret protection as well as the state with the strictest protection. I find that the ratio of purchased outsourcing to payroll expenses would increase by 4.5 pp (from 13.6% to 18.1%), while the aggregate output would go up by 0.8% (\$165B in 2018). A large portion of the output growth

would come through the entry of new firms, while the size-productivity correlation in the economy would also improve. Since the only productive input in the economy, labor, is fixed, all productivity gains essentially stem from the improved allocation of workers between producers. The wage levels would increase more than the increase in output, implying an increase in the labor share. There would also be modest gains in business dynamism through increased job reallocation and entry rates in the steady-state.

My paper is closely related to others that use estimated distortions in firm decisions to analyze the importance of contract enforcement and trust for aggregate productivity. Bloom, Sadun, and Van Reenen, [2012](#) find that the regions that have lower trust measures have firms with more centralized structures, slower worker reallocation, and lower productivity. Akcigit, Alp, and Peters, [2021](#), who quantify the impact of lack of enforcement and the resulting lack of delegation, find that the differences in enforcement can explain 11% of the productivity difference between India and the U.S. Grobovšek, [2020](#) finds similar quantitative effects from lack of enforcement using data from France. The closest paper to mine is Boehm and Oberfield, [2020](#). They study the impact of weak contract enforcement on aggregate productivity through distortions in the choice of intermediate inputs. In particular, in Indian states where courts are more congested, firms substitute away from specialized intermediate inputs towards generic ones to avoid hold-up problems. My empirical strategy is similar to theirs in that I use cross-state variation in wedges to structurally identify distortions. However, there are methodological differences beyond the differences in our questions. Boehm and Oberfield, [2020](#) use firm-level data on intermediate input use, which allows them to control for a larger set of differences across states than mine. At the same time, their model is static, which does not permit analysis of the dynamic flexibility gains from labor outsourcing. While their measure of court

congestion is constant over time, I can use state-level changes in laws to control for many state-specific covariates through state fixed effects.

My paper also contributes to the literature on the cost of employment protection. The patterns and implications of labor flows have been studied extensively,² but especially more recently after Restuccia and Rogerson, 2008 and Hsieh and Klenow, 2009 who showed that input misallocation can explain a large part of cross-country differences in aggregate TFP. Hopenhayn and Rogerson, 1993, using a general equilibrium setting, found that a firing cost equal to 1 year of wages can decrease employment by as much as 2.5%.³ Focusing largely on the fixed-term contracts commonly used in Europe, a branch of the literature asked whether alternative forms of employment can help (Bentolila and Saint-Paul, 1992, Cahuc and Postel-Vinay, 2002, Caggese and Cuñat, 2008, Katz and Krueger, 2019). My contribution here is two-fold. First, I study the importance of a wide range of labor outsourcing practices instead of the fixed-term workers that tend to work in lower-skilled occupations. Second, I allow outsourced workers to be imperfect substitutes to permanent workers and evaluate distortions that limit their utilization.

My paper is also related to the literature that examines the determinants and consequences of labor outsourcing. The large growth in labor outsourcing practices brought nationwide surveys, as in Harrison and Kelley, 1993, K. Abraham and Taylor, 1996 and Houseman, 2001. The three biggest reasons managers list for outsourcing are higher flexibility, access to specialized labor, and cost savings. Autor, 2001, Houseman, Kalleberg, and Erickcek, 2003 and Autor and Houseman, 2010 analyze

²See Davis and J. Haltiwanger, 1992, Caballero and Hammour, 1994, Bartelsman and Doms, 2000, Foster, J. C. Haltiwanger, and Krizan, 2001, Autor, Kerr, and Kugler, 2007.

³Bento and Restuccia, 2017 and Da-Rocha, Restuccia, and Tavares, 2019 have found the impact of firing costs on employment and productivity becomes even larger once the life-cycle productivity growth of firms is endogenized. The impact of employment protection laws on labor allocation had been an active area, following the early contributions by Lazear, 1990 and Bentolila and Bertola, 1990.

how outsourcing allows employers to screen potential hires. Bidwell, 2012, using data on outsourcing projects within a single firm, suggests that personal interests of managers play a role in outsourcing decisions. More recently, Goldschmidt and Schmieder, 2017 and Drenik et al., 2020 use microdata on both the employer and client of outsourced workers to confirm the cost saved by outsourcing instead of hiring.⁴ Adding to the literature, I propose and quantify the trade secret protection as a concern in labor outsourcing decisions. My model incorporates an examination of how outsourcing impacts flexibility, access to specialized talent, and cost savings in a simplified way. However, it does not incorporate the potential benefits through screening or an organizational conflict within the firm. Lastly, Bloom, Eifert, et al., 2013 and Bruhn, Karlan, and Schoar, 2018, using RCTs, document large sustained gains from receiving free management consulting services. I confirm their findings in a macroeconomic setting.

Last, my paper is related to studies of firm boundaries. Following Coase, 1937, Williamson, 1975, and Grossman and Hart, 1986, the literature analyzes how imperfect contract enforcement impacts the organization of production. The empirical literature has broadly focused on either the make-or-buy decisions for physical inputs by multinationals or the competitive effects of vertical integration.⁵ I contribute by showing that intellectual property protection is specifically important for the make-or-buy decision for services.⁶

⁴For papers that analyze the macroeconomic implications of growing labor outsourcing, see Berlingieri, 2013 for the structural transformation in the U.S., Giannoni and Mertens, 2019 for the trends in labor share, and Bilal and Lhuillier, 2020 and Bergeaud et al., 2020 for wage inequality.

⁵See Antras, 2003, Nunn, 2007, Corcos et al., 2013, Antras, Fort, and Tintelnot, 2017, and Boehm, Forthcoming for discussion on multinational organizations. See Alfaro et al., 2016, Crawford et al., 2018, Hansman et al., 2020 for research on competitive effects and Lafontaine and Slade, 2007 for a broad review of this literature.

⁶The idea that firms provide a structure that protects secrets has been proposed as early as Alchian and Demsetz, 1972 and Liebeskind, 1996. See Rajan and Zingales, 2001 and Henry and Ruiz-Aliseda, 2016b for theoretical analyses and Ethier and Markusen, 1996, Fosfuri, Motta, and Rønde, 2001, Bolatto et al., 2020, and Kukharsky, 2020 for the make or buy decision of multinationals in

The rest of the paper is structured as follows. Section 1.2 summarizes trade secret protection in the U.S. and how it matters for labor outsourcing in particular. Section 1.3 documents new facts on outsourcing as well as a causal link from trade secret protection that motivates the structural model. Section 1.4 presents the structural model, while Section 1.5 presents the estimation strategy and results. Section 1.6 presents the counterfactual exercise and Section 1.7 concludes.

1.2 Background

I start this section by defining trade secrets and discussing their significance for businesses. Second, I discuss the historical development of the trade secret law in the U.S., emphasizing the Uniform Trade Secrets Act (UTSA). Third, I discuss how trade secret law impacts employees and outsourced workers differently.

1.2.1 Trade Secrets

The USPTO defines trade secrets as “information that has either actual or potential independent economic value by virtue of not being generally known, has value to others who cannot legitimately obtain the information, and is subject to reasonable efforts to maintain its secrecy”. Business information such as customer lists and pricing strategy as well as R&D related information such as manufacturing techniques and designs can be trade secrets.

Trade secrets are arguably the most important form of IP for most businesses. Protecting information on clients and suppliers, pricing strategies, and long-term growth plans have historically been essential for firms. On the other hand, only a fraction of firms engage in formal R&D, and among those that do, a small fraction

countries with weak IP protection.

holds patents. Moreover, trade secrets are still a fundamental part of R&D, even when the ultimate goal is to get a patent.

Trade secrets are understudied compared to other forms of IP, as assigning a dollar value to secrets is hard with the absence of an explicit market.⁷ The lack of legal uniformity has also limited statistical research on trade secret protection, even though they are the most litigated form of intellectual property (Lerner, 2006).

1.2.2 Trade Secret Protection in the U.S.

Before 1979, trade secrets were protected exclusively under common law.⁸ This created two main problems. First, as no two cases are the same, there was uncertainty regarding the law's extent.⁹ Second, three standard requirements -to declare the act as a trade secret violation- were unfit for outsourcing practices: (1) information had to be illegally appropriated, (2) the accused party had to be in direct competition with the plaintiff, and (3) those who have paid an amount in good faith to purchase the information from the accused were not prevented from further use (Lao, 1998). Because the outsourced worker would usually receive the information legally and act only as an intermediary between the client and its competitor, the law did not provide adequate protection for outsourcing relationships.

The Uniform Law Commission has drafted the Uniform Trade Secrets Act (UTSA) in 1979. The UTSA statutes defined which information constitutes a trade secret, which acts constitute misappropriation, and which are the associated remedies. It

⁷The 2017 report by the Commission on the Theft of American Intellectual Property estimates the total cost of trade secret theft to the U.S. economy to be between 1% to 3% of GDP, which is somewhere between the yearly outlays to the Dept. of Education and Dept. of Defense.

⁸Common law, as opposed to statutory law, does not rely on a codified set of rules. Instead, it uses previous court decisions to reach new ones.

⁹"... even in states in which there has been significant litigation, there is undue uncertainty concerning the parameters of trade secret protection, and the appropriate remedies for misappropriation of a trade secret.", UTSA Prefatory Note (1985). See Appendix 1.E for details on the legal environment under common law.

also broadened the law’s scope, e.g., by making misappropriation itself a crime, without the information being used or disclosed. Most importantly, it made third parties liable if they receive this information with a reasonable expectation that it is misappropriated. Each state had to opt-in for the UTSA to be effective in its courts. Minnesota, Idaho, Arkansas, Kansas, and Louisiana were the first states to adopt it in 1980. By 1988, 26 states had already adopted it, and by 2019, all states did.¹⁰

1.2.3 Trade Secret Protection and Labor Outsourcing

There are two main reasons why trade secret law is crucial for labor outsourcing. First, although its extent varies, all outsourced workers are exposed to some trade secrets. Second, it is harder to prevent outsourced workers from disclosing secrets to third parties compared to employees.

High-skill outsourcing generally provides a personalized solution to the client’s problem; hence it is straightforward how an outsourced R&D expert or an accountant would be exposed to secret information. Albeit to a lesser degree, trade secrets are also relevant for the low-skilled. An outsourced machine operator would be exposed to product designs and daily production volumes. An outsourced personal assistant would have access to manager’s daily activities, including meetings with other branches and business partners. Furthermore, having access to facilities may enable overhearing the managers’ discussions and the rumors circulating among other workers¹¹. In short, outsourced workers’ regular activities inherently create exposure

¹⁰There have been two other main developments in trade secrets protection. Economic Espionage Act of 1996 made trade secrets misappropriation that is either interstate or benefits a ‘foreign power’ a federal crime. The Defend Trade Secrets Act of 2016 (DTSA) allowed any trade secret misappropriation case to be seen in federal courts. Although both are significant developments, they happened at the national level, making it harder to measure their impact.

¹¹In *SEC v. Steffes*, No. 01 Civ. 06266 (N.D. Ill. Sept. 30, 2010), the SEC alleged railroad workers “traded and tipped on observations made on the job, including seeing people in suits tour the rail yards, hearing coworkers discuss the possible sale of their company, and being asked to prepare asset valuations.”Cohen and Dunning, [2010](#)

to firm secrets unless the firm explicitly limits their access, which would reasonably reduce their value.

The data from trade secret litigation confirm the intuition. First, limiting access to certain ‘labs’ does not protect the business from trade secret misappropriation. Almeling, Snyder, and Sapoznikow, 2009 shows, in their sample of U.S. federal district court cases in 2008, only 35% involved any technical information or know-how. 31% involved customer lists, and 35% involved non-technical business information. Second, the misappropriator is almost always someone who has physical access to the secret: an employee or a business partner in 90% and 93% of the cases for the cases in federal and state appellate courts, respectively (Almeling, Snyder, Sapoznikow, and McCollum, 2010). Similarly, the defendant was either a former, current, or an outsourced worker in 76% of the cases tried under the Economic Espionage Act (Searle, 2012).

Employees are less susceptible to these concerns than outsourced workers for two main reasons. First, voluntary disclosure of secrets is less likely for employees. Because the employment relationship is generally of longer-term¹², it allows the design of better incentives for the employee to work in the best interest of the employer (Liebeskind, 1996, Gibbons, Roberts, et al., 2013). Second, inevitable disclosure is less likely for employees. While covenant not to compete (CNC) agreements¹³ are ubiquitous among employees that work with sensitive data (Jeffers, 2018, Shi, 2020), they are not common in outsourcing agreements, being directly at odds with the business model of most outsourcing firms.¹⁴ Lastly, signing a non-disclosure agree-

¹²There is no legal constraint on how long an outsourcing relationship lasts. However, longer relationships make it more likely that the courts will interpret it as a de facto employment relationship in case of a dispute, especially upon termination. See *Amarnare v. Merrill Lynch, Pierce, Fenner & Smith Inc.*, (611 F. Supp. 344 S.D.N.Y. 1984).

¹³CNC agreements designate a period for which the employee cannot work in the same industry with the previous employer upon termination of the employment contract.

¹⁴“Firms regularly hire consultants to advise on sensitive business problems, and one of the im-

ment helps, but how it is enforced is largely determined by the trade secret law (See Appendix 1.E.2).

In short, firms have reason to avoid labor outsourcing to limit the risks of losing trade secrets. The next section tests and confirms this hypothesis using the cross-state legal variation across the U.S. The modeling choices in Section 1.4 are based on the frictions discussed here.

1.3 Empirical Analysis

In the first half of this section, I document two broad facts on domestic¹⁵ labor outsourcing in the U.S, focusing on its growth and cross-state heterogeneity. In the second half, I argue the trade secret laws in the U.S. help explain the two facts.

I define labor outsourcing as the purchase of labor-intensive services that can otherwise be done in-house. My definition is far from being arbitrary. The businesses that provide outsourcing as I define it are conveniently classified into two 2-digit NAICS sectors.¹⁶ NAICS 54 (The Professional, Scientific, and Technical Services) principally employs high-skill occupations such as consultants, accountants, and data analysts. NAICS 56 (The Administrative and Support and Waste Management and Remediation Services) principally employs lower-skilled occupations such as machine operators, security guards, and janitors. The output of both sectors is mainly used as an intermediate input by other sectors. The set of industries in this definition is similar to Berlingieri, 2013, but more extensive than Autor, 2003 and Katz and

portant qualifications of the consultants seems to be that they know the industry well—they have offered similar consulting services to the competitors.” Kitch, 1980

¹⁵I abstract from foreign outsourcing (e.g., call centers abroad) because it constitutes a relatively small fraction (3.5% in 2004) of total labor outsourcing practices (Amiti et al., 2005). See Appendix 1.C for the cross-country evidence on the relationship between outsourcing and trade secret protection. The cross-country evidence broadly supports the analysis within the U.S.

¹⁶See Appendix 1.B for the few exceptions, the details of the selection of industries, and how I map different classifications to one another.

Krueger, 2019 who prioritize temp agencies.

Throughout the paper, I refer to the firms and the industries that supply labor outsourcing services as the outsourcing firms and the outsourcing sector for brevity.

1.3.1 Facts on Domestic Labor Outsourcing

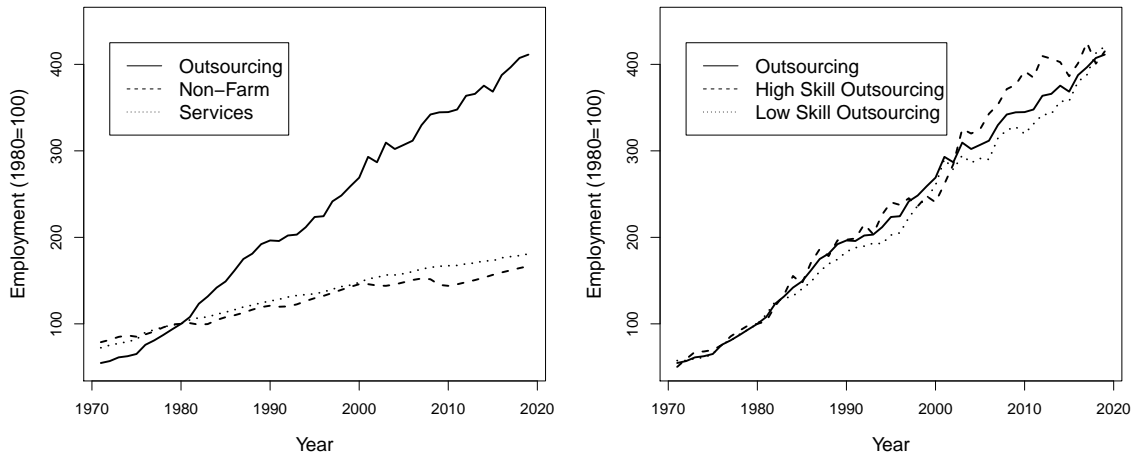
Here, I present two sets of facts that shows a large heterogeneity in labor outsourcing across states and over time in the U.S. Furthermore, the heterogeneity is not explained by differences in skill levels, industries, and occupations.

Fact 1: The outsourcing sector's employment share has tripled since the 70s.

The outsourcing sector's employment share increased from 3% in 1971 to 11% in 2019. The left-hand side panel in Figure 1.1 depicts the normalized non-farm employment, service employment, and employment in the outsourcing sector. The average growth in the outsourcing sector far exceeds the US non-farm and services employment. The right-hand side panel shows the large growth was evident for both skill groups. So, the underlying reasons cannot be exclusively based on the skill level.

The growth in outsourcing was also not an artifact of (1) the growth in industries that historically had above-average demand for outsourcing or (2) the growth in demand for occupations that historically had been outsourced more than others. I use the BEA Integrated Production Account and find the aggregate ratio of purchased services to value-added has increased from 0.25 in 1963 to 0.44 in 2018. Using the time series for 63 industries, I compute the counterfactual growth if each industry's purchased services ratio remained constant while the output shares changed as they did (between-industry), and if the output shares remained constant while the purchased services ratios changed as did (within-industry). I find that 84% of the growth


Figure 1.1: Employment Trends in Multiple Industry Groups (1971-2019) Notes: See Appendix 1.B and Table 1.7 for details on how I pick and classify sectors into low and high skill outsourcing. Service Employment in the left panel consists of all U.S. Census 1990 3-digit industry groups from 400 to 892. Sector level employment is from the Annual Social and Economic Supplement (ASEC) of IPUMS-CPS. Total Non-farm employment is published by the Bureau of Labor Statistics (BLS).



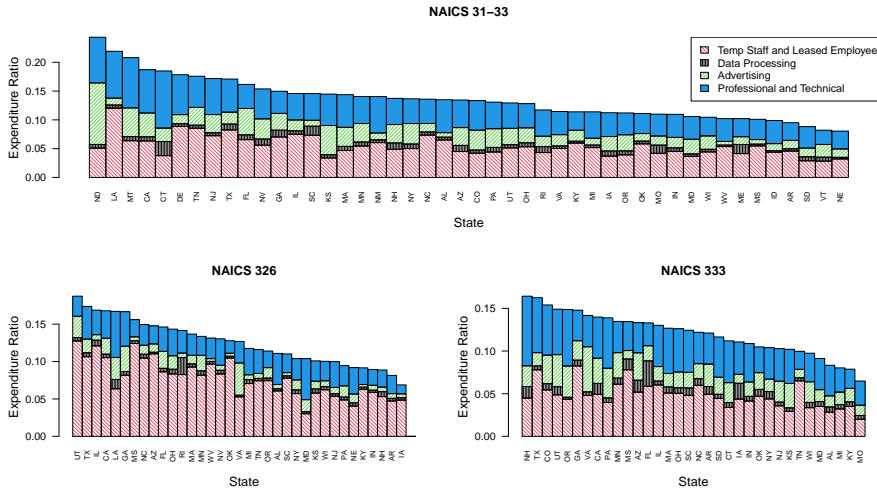
is within-industry, i.e., would still happen with no structural change.

I further check whether the growth in services outsourcing is part of a broader trend of shrinking firm boundaries. On the contrary, the ratio of all intermediate inputs to value-added has decreased from 0.83 to 0.76 during the same period. Although each industry uses more intermediate inputs on average, the structural shift from manufacturing to services more than canceled the growth.¹⁷

Figure 1.2: The Cross-state Supply of and Demand for Labor Outsourcing Notes: The details on the data sources and the state abbreviations are available in Appendix 1.B.



(a) Employment Share of Outsourcing Sectors (2017) Notes: The full length of the bar designates the employment share of outsourcing, while the shaded length (in red) designates the portion that is in high skill outsourcing sectors. The data is from IPUMS USA. See Figure 1.1 for details on which industries are included.



(b) Ratio of Outsourcing Expenses to Annual Payroll in Manufacturing Sectors (2017) Notes: The top panel provides estimates for all NAICS manufacturing sectors (31-33), the bottom left panel for Plastics and Rubber Products Manufacturing (326), and the bottom right panel for Machinery Manufacturing (333). In each panel, only the states with complete data on each of the four outsourcing expenses are included. All panels use data from the 2017 Census of Manufactures.

Fact 2: The supply of and demand for outsourcing is heterogeneous across states.

I define a state's 'supply' of outsourcing as how much outsourcing services it provides, and its 'demand' as how much outsourcing services is used there.¹⁸ To measure the supply of outsourcing, I use the American Community Survey from the IPUMS USA database to get employment shares for outsourcing providing sectors. Figure 1.2a presents the shares across the states of the U.S. First, there is considerable heterogeneity: the state at the 90th percentile has a share of 14.3% while the 10th has 7.6%. Second, a large part of the heterogeneity comes from high-skill outsourcing: the outsourcing employment share and high skill ratio have a correlation of 0.6.

To measure the demand for outsourcing, I use the 2017 Census of Manufactures in Figure 1.2b, which provides estimates of expense items for employer establishments. Specifically, it gives expense estimates for Temporary Staff and Employee, Data Processing Services, Advertising and Promotional Services, and Professional and Technical Services among others. For each state, I plot the ratio of their sum to the Annual Payroll. First, the state-level heterogeneity is comparable to the heterogeneity in supply. The state in the 90th percentile has a ratio of 0.18, while the 10th has 0.1. Second, heterogeneity does not concentrate on one of the four types of outsourcing expenses. Third, it does not disappear at more disaggregated levels. For example, both the Plastics and Rubber Products Manufacturing and the Machinery Manufacturing exhibit similar degrees of heterogeneity in outsourcing expenses, although their composition is very different.¹⁹ Fourth, states with higher outsourcing

¹⁷Berlingieri, 2013 does a similar test for occupations. He picks occupations that are predominantly employed in outsourcing sectors and tracks their employment share over time. He finds that this share shows no trend after 1970, where most of the outsourcing growth happens.

¹⁸The two need not equal as outsourcing services provided by a firm in one state can be used by a firm in another state.

¹⁹The degree of heterogeneity also persists at the 6-digit industry level; however, the data is censored for most state-industry pairs to ensure the confidentiality of firm data. For example, the

ratios are also the ones that have a larger share of their outsourcing in high-skill tasks, with a correlation of 0.32.

1.3.2 Evidence on the Effect of Trade Secret Laws

The previous facts presented a considerable heterogeneity in labor outsourcing both across states and over time that was not explained by differences in skill levels, industries, and occupations. Here, I test whether the differences in trade secret protection over time and across states play a role.

Data and the Estimation Method

Testing the impact of trade secret protection is not straightforward for a few reasons. First, the legal frameworks differ across states in clarity and scope, which are hard to quantify. I use two measures in this section, namely, adoption of the Uniform Trade Secrets Act (UTSA) and the trade secret protection index (TSP index henceforth) constructed by Png, 2017a and Png, 2017b. The adoption of the UTSA was essential both for reducing the uncertainty about the trade secret protection and extending its coverage, particularly for labor outsourcing relationships. The TSP index evaluates whether states had certain types of protections in a given year and assigns a score ranging from 0 to 1 (See Appendix 1.B for details).

Second, I need a measure of the extent of outsourcing. Unfortunately, comprehensive data on demand for labor outsourcing does not exist before 2007. Thus, I use the supply of labor outsourcing as my measure.²⁰ I use the state-year level employ-

10th and the 90th percentiles are 9% and 18% in the Plastics Pipe and Pipe Fitting Manufacturing (NAICS 326122).

²⁰Although there was no definitive procedure, the governing law was of the state where the misappropriation happened in a large majority of cases (See Appendix 1.E.3). This state would generally be the one where the client operates, especially in the 80s and 90s. As long as outsourcing firms are more likely to serve clients in their states, my mechanism predicts a positive relationship between the strength of trade secret protection and the employment share of the outsourcing sector in that

ment shares of the outsourcing sector from the ASEC samples. In total, I have an unbalanced panel of 50 states and the District of Columbia from 1970 to 1997.

Last, to measure the causal link, I need exogenous variation in protection. The UTSA provides precisely that. After being drafted, each state had to opt-in to start using it. The adoption times differed significantly (See Figure 1.16), creating cross-sectional variation in trade secret protection on top of the time-series variation. After arguing its exogeneity, I use the staggered adoption of the UTSA as my exogenous variation for trade secret protection.

The staggered adoption of the UTSA allows aggregating the information from difference-in-differences (DiD) comparisons across multiple pairs of states over many periods. The Two-Way Fixed Effects (TWFE) estimator provides an intuitive tool and is widely used in studies with staggered adoptions. However, the recent work following Goodman-Bacon, 2018 has shown TWFE may fail to give (1) consistent test statistics for pre-trends and (2) intuitive measures of treatment effects without strong assumptions (Appendix 1.F for details). In my analysis, I primarily yield to the historical setting to argue for the exogeneity of the UTSA adoptions, together with statistical tests for pre-trends. I then provide estimates from both the TWFE estimator and the estimator proposed by Callaway and P. Sant’Anna, 2020, which remains consistent under multiple dimensions of treatment heterogeneity and selection into treatment based on covariates.

Exogeneity of the UTSA Adoption

I start by confirming that the adoption of the UTSA did not coincide with the adoption of other major state-level laws. The adoption time of the UTSA has a weak correlation with the adoption of other commercial uniform laws (<0.13) and employ-

state.

ment protection laws (<0.04) across states.

The adoptions' history suggests the timing choices of states were less about economic concerns and more about differences in legal structures and opinions. First, Ribstein and Kobayashi, 1996 show the basic economic characteristics like size, population density, and state expenditures were irrelevant in explaining the adoption of any uniform law. The structure of the state legislatures (e.g., size of chambers), on the other hand, had predictive power on the adoption dates. Second, Sandeen, 2010 documents, many states postponed their adoption of UTSA to after 1985 due to the opposition organized by a single attorney who argued certain clauses could be misinterpreted. Last, Png, 2017a discusses how UTSA was adopted in California only when proposed a second time and rejected in New York for reasons unrelated to the intended coverage of the UTSA. The opposition came from farmworkers in California and trial lawyers in New York. They were concerned that the law can be used to hide information about pesticides and trial evidence, respectively.²¹ The convergence also supports the argument for differences in legal opinions: all states adopted a version of the UTSA eventually.

The quantitative tests do not suggest the presence of pre-trends either.²² First, I run the classical event study regression with the leads and lags of the treatment in a TWFE setting

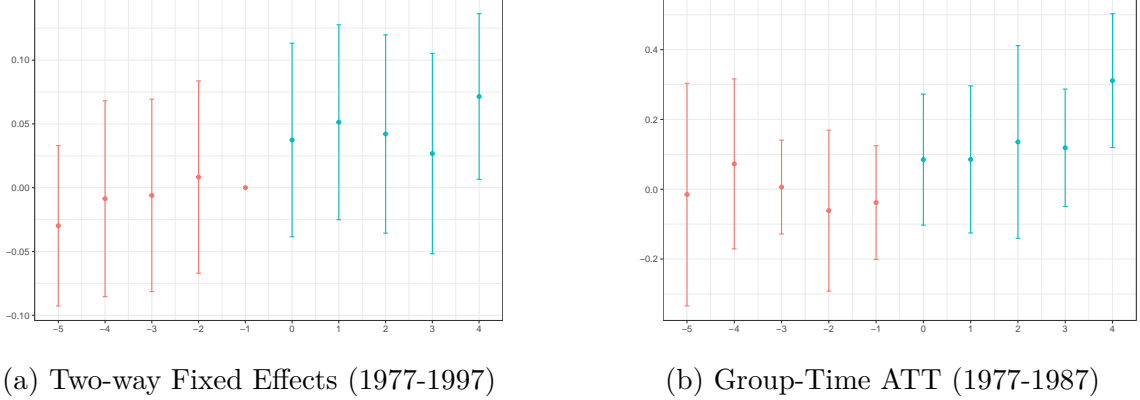
$$y_{it} = \sum_{l \in \{-4, -3, -2, 0, 1, 2, 3\}} \delta_l A_{itl} + \delta_4 A_{it, l \geq 4} + \delta_{-5} A_{it, l \leq -5} + \beta x_{it} + \alpha_i + \gamma_t + \epsilon_{it} \quad (1.1)$$

where y_{it} is the log employment share of outsourcing sectors, A_{itl} is equal to 1 if for

²¹Similarly, during the United Kingdom's implementation of the Trade Secrets Directive in 2018, the opposition centered around whether the law would be used against journalists and whistle-blowers (IPO, 2018).

²²Png, 2017a and Klasa et al., 2018 provide several tests and conclude variables used in their analysis including R&D expenditures and capital structures of firms do not predict the adoption of the UTSA.

Figure 1.3: Event Study Estimates for the UTSA Adoption Notes: The X-axis refers to l in (1.1) for the left panel and $t - g$ in (1.3) for the right panel. Y-axis provides the corresponding estimates with 95% confidence intervals constructed from standard errors clustered at the state level. I use the doubly-robust balancing procedure in the right panel. The outsourcing shares and employment series are from the IPUMS-CPS database. The controls are population, GDP, manufacturing GDP, manufacturing employment, unionization rate, high school and college shares, and adoption of exceptions to at-will employment. Since the CS estimator relies on propensity score matching, the control group must be sufficiently large for estimation. Hence, the estimation only runs for adoptions until 1987. See Figure 1.1 for details on included industries.



state i , year t is l years after the adoption of the UTSA. The coefficient estimates are in Figure 1.3a. There are no signs of a pre-trend, i.e., the states that are closer to adoption have comparable outsourcing shares to others. However, the plot also hints at dynamic treatment effects: it takes a few years for the treatment to have full effect. Thus, the pre-trend test likely suffers from the bias suggested by Sun and S. Abraham, 2020. Thus, I supplement the analysis by using the estimator by Callaway and P. Sant’Anna, 2020 (CS henceforth).

CS starts with the concept of group-time average treatment effects on the treated:

$$ATT(g, t) = E[Y_t(g) - Y_t(0) | G = g] \quad (1.2)$$

where g denotes group index (the adoption time), G_i denotes the group of unit i , $Y_t(g)$ ($Y_t(0)$) denotes the outcome variable at time t conditional on being treated at time g (never being treated). Thus, $ATT(g, t)$ denotes the effect of being treated at time

g that is measured in time t , thus allows heterogeneity across groups and dynamic treatment effects. Furthermore, by conditioning on being treated, it controls for selection into treatment.²³ After identifying $ATT(g, t)$, CS aggregates them over t to get average dynamic effects:

$$\theta_D(e) := \sum_{g=2}^{\mathcal{T}} \mathbf{1}\{g + e \leq \mathcal{T}\} ATT(g, g + e) P(G = g | G + e \leq \mathcal{T}) \quad (1.3)$$

where e denotes the exposure time and $\theta_D(e)$ are the counterparts of the event study estimates of the classical DiD under homogenous treatment. Lastly, $ATT(g, t)$ can be aggregated over both g and t to get an overall treatment effect:

$$\theta_S^O := \sum_{g=2}^{\mathcal{T}} \theta_S(g) P(G = g) \quad (1.4)$$

Figure 1.3b plots the event study estimates from (1.3), which confirm the findings with the TWFE: there are no apparent pre-trends, and the full effect is realized only a few years after the adoption.

The Impact of Trade Secrets Laws

Having established a case for the exogeneity of the UTSA adoption, I use the variation it created to estimate the impact on outsourcing employment.

I have so far ignored that trade secret protection may have differed both pre- and post-adoption across states. I use the TSP index as the regressor in the main specification below, instrumented by the adoption dummy in a TWFE model. There-

²³CS identifies $ATT(g,t)$ under the assumptions of parallel trends (conditional on observables) and absorbing treatment. In particular, to avoid the bias generated by dynamic treatment effects, CS only uses units that are not yet treated in the control group and uses propensity score matching to balance the two groups on relevant observables to take potential selection into treatment into account.

fore, I measure the impact through an index that quantifies this heterogeneity while restricting attention to changes through the UTSA. To test the results' robustness, I also use the CS estimator to take selection into treatment, dynamic treatment effects, and treatment heterogeneity over time of adoption into account. In the main specification, I estimate a TWFE-IV model of the form:

$$y_{it} = \beta tsp_{it} + \tilde{\beta} x_{it} + \alpha_i + \gamma_t + \epsilon_{it} \quad (1.5)$$

where y_{it} is the log employment share of outsourcing sectors, tsp_{it} is the TSP index, x_{it} is the vector of controls, α_i and γ_t are the state and year fixed-effects. α_i helps control for state-specific factors that remain constant over time, such as persistent differences in state subsidies and the availability of natural resources. γ_t provides a non-parametric time trend, controlling for broad trends in the economy, such as the growth in information technology and changes in the federal subsidies. I instrument the TSP index with the adoption dummy for the UTSA and use White standard errors clustered at the state level.

Table 1.1 presents the regression results. Trade secret protection has a positive and statistically significant effect at 5% level, in line with my hypothesis. Moreover, the quantitative estimates are similar across specifications without controls or instrumentation. Using the estimates, I find the outsourcing sector would be 14% smaller in 1997 had all the controls changed as they did, but the TSP indices remained the same as the 1971 levels, translating to 0.7M jobs.

I also use the CS estimator's overall treatment effect in Equation (1.4), which gives comparable results.²⁴ The CS estimates are qualitatively in line with the TWFE

²⁴Since CS takes potential selection into treatment based on observable covariates into account, it requires a large enough control group for balancing the treatment and the control groups. A larger estimation period allows using more pairwise DiD estimates to incorporate in the estimation. However, as the estimation horizon grows, the control group's size gets smaller, and the balancing

Table 1.1: Two-way Fixed Effects Estimation

	Adoption	Index	IV	Adoption	Index	IV
	(1)	(2)	(3)	(4)	(5)	(6)
TS Protection	0.05* (0.03)	0.12* (0.06)	0.12* (0.06)	0.06** (0.03)	0.13** (0.05)	0.13** (0.06)
Demographics				Yes	Yes	Yes
Ind Composition				Yes	Yes	Yes
Union				Yes	Yes	Yes
WDL				Yes	Yes	Yes
State & Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Range	'70-'97	'70-'97	'70-'97	'70-'97	'70-'97	'70-'97
Observations	1,180	1,180	1,180	1,180	1,180	1,180

Notes: The dep. variable is the log outsourcing sector share of employment. The employment series are from IPUMS-CPS. See Figure 1.1 for details on included industries. The main variable of interest is the UTSA adoption dummy in columns (1) and (4), and the TSP index in others. Columns (2) and (4) present OLS estimates while (3) and (6) present IV estimates. Columns (4)-(6) controls for unionization rate, the share of college and high school graduates, the exceptions (good faith, implied contract, public policy) to the at-will employment as well as logged population, GDP, manufacturing GDP, and manufacturing employment. See Appendix 1.B for details on how each variable is constructed. I cluster the standard errors at the state level. *p<0.1; **p<0.05; ***p<0.01

estimates, although their magnitude is larger. The difference in magnitudes may indicate large dynamic treatment effects, as suggested by the event study estimates in Figure 1.3.

Two additional concerns bias the estimates towards 0 and cannot be resolved without additional data. First, the treatment also impacts the control group. Once a state adopts the UTSA, its subsequent decisions may affect others that are yet to adopt. As the extent of cross-state citations increases, my estimates' bias would be greater. Second, the data available for this period is on the supply side, while becomes less precise and eventually infeasible. Lastly, the CS estimator requires a balanced panel; hence the longest estimation period I can use is from 1977 to 1987. I also present results from smaller horizons where the balancing is more precise.

Table 1.2: Overall Treatment Effect via Group-Time ATT Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
UTSA Adoption	0.20 (0.077)	0.13 (0.076)	0.16 (0.065)	0.16 (0.057)	0.11 (0.047)	0.16 (0.051)
Range	'77-'82	'77-'83	'77-'84	'77-'85	'77-'86	'77-'87
Number of Adopted States	6	9	11	11	13	19

Notes: The estimates correspond to Callaway and P. Sant'Anna, 2020 group-time att estimates integrated over time of adoption and the length of exposure to treatment using (1.4). The dependent variable is the log outsourcing sector share of employment, and the treatment is the adoption of the UTSA. The control group consists of states that are not-yet-treated, and the balancing is done via the doubly-robust estimation method by P. H. Sant'Anna and Zhao, 2020. See the notes for Table 1.1 for a list of control variables included and details on the variables.

the adoption reasonably impacts the demand. As the extent of cross-state trade of outsourcing services increases, my estimates' bias would be greater. The structural model in Section 1.4 uses demand-side data to circumvent the second problem, while the first problem requires measuring the extent of cross-state legal influence.

Placebo Regressions

If trade secret protection is indeed important, the effect of laws should be greater for high-skill outsourcing, where the exposure to trade secrets is arguably higher. In columns 2 and 3 of Table 1.1, I estimate Equation (1.5) for high-skill and low-skill outsourcing sectors separately. In line with the theory, the impact on high skill outsourcing is greater. In column 4, I address 3-digit sectors 841 and 890, which mainly employ lawyers and accountants subject to client privilege codes: her association would disbar an accountant or lawyer that discloses her client's information to 3rd parties.²⁵ Hence, these two sectors should be affected to a lesser extent. The estimate confirms this, where the estimate is both quantitatively smaller and not different from

²⁵See the American Institute of Certified Public Accountants' Trust Services Criteria and the American Bar Association's Model Rules of Professional Conduct.

0 at a 10% significance level. Lastly, in column (5), I re-run column (1) excluding sub-sector 732 (Computer and data processing services) and confirm that the concurrent growth of the role of computers in businesses does not drive the results.

Table 1.3: Placebo Regressions

	Outsourcing Share	High-Skill	Low-Skill	Leg-Acct	Except Comp
	(1)	(2)	(3)	(4)	(5)
TSP Index	0.13** (0.06)	0.18** (0.09)	0.12 (0.08)	0.13 (0.11)	0.13* (0.07)
Range	'70-'97	'70-'97	'70-'97	'70-'97	'70-'97
Observations	1,180	1,174	1,175	1,177	1,180

Notes: The outsourcing shares and employment series are from the IPUMS-CPS database. See Figure 1.1 for details on included industries and their assignment into skill bins. The fourth column is the total employment in 3-digit 1990 U.S. Census sectors 841 (Legal services) and 890 (Accounting, auditing, and bookkeeping services). The fifth column is all 3-digit high skill outsourcing sectors except for 732 (Computer and data processing services). Standard errors are clustered at the state level. See Table 1.1 for details on the controls. *p<0.1; **p<0.05; ***p<0.01

1.4 A Model of Outsourcing and Trade Secret Protection

In this section, I construct a multi-industry firm dynamics model based on Hopenhayn, 1992, where firms decide whether to use in-house or outsourced workers for various tasks. Outsourced workers are more productive in certain tasks and are easier to adjust, but need firm-specific information to perform. The effective trade secret protection determines what amount is safe to share, i.e., the size of the *enforcement friction*.

The model provides three main inputs that allow quantifying the output cost of enforcement frictions using the observed cross-state heterogeneity in outsourcing use.

First, it provides a mapping between observables such as firm size distribution and job destruction rates and structural parameters such as demand persistence and labor adjustment costs. Second, it incorporates an intuitive restriction: the productivity advantage of outsourced workers depends on the industry but not on the state. In contrast, the strength of trade secret protection depends on the state but not on the industry. Third, it maps estimated firm-level distortions to aggregate productivity by taking general equilibrium effects through product and labor markets into account, providing the final piece.

1.4.1 Environment

Agents and Preferences

The economy consists of (1) a decreasing returns-to-scale (DRS) intermediate goods sector with K industries, (2) a constant returns-to-scale (CRS) final good sector, (3) a CRS outsourcing sector, and (4) a unit measure of workers. Each K industries in the intermediate sector have a continuum of firms and a large pool of potential entrants. All firms maximize expected discounted profits. Each worker inelastically supplies one unit of labor and is indifferent between being a permanent or outsourced worker.

Technology

The Final Good and Outsourcing Sectors

All the action in the model is in the intermediate goods sector, so I quickly discuss the other two sectors here. The final goods sector produces the final good by

aggregating the intermediate goods, solving:

$$\max_{\{Y_k\}_{k=1}^K} P \left(\sum_{k=1}^K Y_k^\omega \right)^{\frac{1}{\omega}} - \sum_{k=1}^K p_k Y_k \quad (1.6)$$

where $1/(1 - \omega)$ is the elasticity of substitution across intermediate goods.²⁶ The outsourcing sector transforms each worker into an outsourced worker. Since both sectors make 0 profits, firms' ownership and size are irrelevant.

The Intermediate Goods Sector

The intermediate goods sector consists of K industries. To simplify the notation, I avoid the industry subscript whenever possible. The structure of the environment is the same across all industries; only the parameter values potentially differ.

I use a task-based production technology similar to Zeira, 1998 and Acemoglu and Restrepo, 2018. The production of each firm is a CES aggregate of production in individual tasks that are indexed by $i \in [0, 1]$:

$$y = \left(\int_0^1 y(i)^\gamma di \right)^{\frac{\theta}{\gamma}} \quad (1.7)$$


where $\theta < 1$ controls returns to scale and $1/1 - \gamma$ is the elasticity of substitution across tasks. Each task i can be done with **permanent** or **outsourced** workers:

$$y(i) = g(i)n(i) + 1_{\{z \geq \zeta(i)\}} \delta r(i) \quad (1.8)$$

where $n(i)$ and $r(i)$ denote the number of permanent and outsourced workers assigned to task i , $g(i)$ denotes the marginal product of permanent workers in task i , δ denotes the marginal product of rented workers, and z denotes the amount of firm-specific

²⁶I do not model demand shares for intermediate goods explicitly, since it is not possible to distinguish them from intermediate goods prices without data on quantities.

Figure 1.4: The Task Allocation Problem of a Software Design Firm Notes: The dashed line pieces denote the upper-envelope of the two lines.



knowledge shared with each outsourced worker. $\zeta(i)$ denotes the minimum amount of information that must be shared to outsource task i . The relative sizes of $g(i)$ and δ determine gains from outsourcing a task, while $\zeta(i)$ puts a hard constraint on which tasks are feasible to be outsourced.²⁷

I assume $g(i)$ is strictly increasing, i.e. (1) the tasks are ordered by how suitable they are to outsourcing, and (2) there is a strict ordering of their suitability. The next assumption is less innocuous.

Assumption 1. $\zeta(i)$ is strictly increasing.

Assumption 1 implies that the gains from outsourcing ($g(i)$) strictly decreases with the required amount of information for the task to be outsourced. This assumption can be micro-founded with a model with communication costs. Relaxing it requires a two-dimensional task space, which is mathematically straightforward but also harder to interpret and complicates the notation. Nevertheless, this assumption is rather conservative for the impact of strengthening trade secret laws. The tasks that would provide the highest marginal gain once outsourced are assumed to be the ones that are already outsourced.

²⁷I abstract from capital as an additional input in the production process. Veracierto, 2001 has previously shown that explicitly modeling capital does not impact the quantitative inference on steady-state labor flows in industry dynamics models.

To make the structure more concrete, imagine SD, a software design firm whose tasks can be grouped into office security, testing, and design. The left-hand side panel in Figure 1.4 places the tasks in the x axis, where the increasing and flat lines represent the marginal product of permanent and outsourced agents respectively in each task i . Design tasks are the firm’s core functions and require knowing the specifications of clients, how the data is organized, etc. The extent of information required would make it more efficient to use a permanent worker. On the other hand, office security requires little firm-specific knowledge; it could be even more productive once outsourced from a security company with better training material. Testing would be in the middle, requiring some firm-specific knowledge, such as the designed software’s potential flaws, but not as much as required by the designers. First, suppose the information-sharing constraint ($z > \zeta(i)$) was not present. Assuming the marginal costs are constant and equal, SD would choose to use permanent workers for design and some testing functions and outsource the rest as in the middle panel of 1.4. However, when the information-sharing constraint is binding, as in the right-hand side panel, effective marginal product becomes zero for the outsourced in tasks that do not satisfy the constraint. Hence, SD would be forced to outsource a smaller set of tasks.

Why does SD not share as much information as possible then, i.e., maximize z ? If SD shares too much, the outsourced would find it more profitable to steal the knowledge, risking a potential lawsuit. Instead of explicitly modeling the ‘trade secret theft’ and its aftermath, which is not the focus of the current paper (See Section 1.4.4), I simplify it into a hard constraint: the firm only shares an amount that does not induce the outsourced worker to steal. How much information is ‘too much’ is determined by π , which I introduce next, which represents the trade secret protection provided by the courts.

The Intermediate Firm’s Static Allocation Problem

Before completing the description of the environment, I first characterize the firm's static task allocation problem with a given number of workers. I then use the solution to this problem later, which simplifies describing the rest of the environment. The firm with n permanent and r outsourced workers chooses how many to allocate each task $(n(i), r(i))$, and how much information to share with the outsourced (z) to solve:

$$F(n, r) = \max_{n(i), r(i), z} \left(\int_0^1 y(i)^\gamma di \right)^{\frac{\theta}{\gamma}}$$

$$\text{(Task production)} \quad y(i) = g(i)n(i) + 1_{\{z \geq \zeta(i)\}} \delta r(i) \tag{1.9}$$

$$\text{(Resource Constraints)} \quad \int_0^1 r(i) di = r, \quad \int_0^1 n(i) di = n$$

$$\text{(Information-Sharing)} \quad z \leq \pi$$

The last constraint represents the legal friction: with perfect enforcement, π would equal one and the information-sharing constraint would be redundant. Given the assumptions on $g(i)$ and $\zeta(i)$, the problem simplifies substantially:

Lemma 1. *Let $n, r, \pi > 0$, $\gamma < 1$. For $g(i), \zeta(i)$ strictly increasing, \exists a unique $0 \leq \bar{z} \leq \zeta^{-1}(z)$ s.t. tasks $i \leq \bar{z}$ only use outsourced and tasks $i > \bar{z}$ only use permanent workers.*

Proof. See Appendix [1.A](#) for all proofs. ■

Thus, the problem of choosing $n(i), r(i)$ boils down to choosing the threshold \bar{z} . The model does not allow identifying the level of $g(i)$ from δ . Although the shape of the $g(i)$ is still important, it matters mainly for counterfactuals that extrapolate from the range of data. Since I do not have task-level data that helps me identify its shape, I go ahead and assume $g(i) = i$ and stick to counterfactuals within the

range of my data. Lastly, it is neither possible nor necessary to identify $\zeta(\cdot)$ and π separately. Thus, I normalize $\zeta^{-1}(\pi) = \pi$. These provide a simple characterization of $F(n, r)$, the maximum production that can be achieved with n and r :

Proposition 1. *The solution to (1.9) can be written as*

$$F(n, r) = \left(\underbrace{\left((1 - \gamma)(1 - \bar{z}^{\frac{1}{1-\gamma}}) \right)^{1-\gamma}}_{\alpha_n(n,r)} n^\gamma + \underbrace{\bar{z}^{1-\gamma} \delta^\gamma}_{\alpha_r(n,r)} r^\gamma \right)^{\frac{\theta}{\gamma}} \quad (1.10)$$

where \bar{z} is a known function of π, n , and r .

Although (1.10) looks like a Constant Elasticity of Substitution (CES) function in permanent and outsourced workers, \bar{z} being a function of n and r complicates things. The next assumption is not required for solving the model but makes the estimation procedure feasible.²⁸

Assumption 2. *The information-sharing constraint is binding.*

Corollary 1. *Under Assumption 2, $\bar{z} = \pi$. Thus, the solution to (1.9) can be written as*

$$F(n, r) = A(\pi, \delta) \left(\alpha(\pi, \delta) n^\gamma + (1 - \alpha(\pi, \delta)) r^\gamma \right)^{\frac{\theta}{\gamma}} \quad (1.11)$$

where $A(\pi, \delta)$ is strictly increasing and $\alpha(\pi, \delta)$ is strictly decreasing in π .

To sum up, under certain assumptions, the solution to the task allocation problem boils down to a CES function, where the factor shares are determined both by the marginal product of outsourced workers (δ) and the strength of trade secret protection

²⁸Specifically, it allows estimating the model for each state of the U.S. separately. Assumption 2 is not on parameters, but on equilibrium outcomes. After estimation, I confirm that this assumption is satisfied for the vast majority of the firms under the estimated parameters. I discuss its benefits and caveats in detail in Section 1.5.

(π). Stronger protection has two effects on F : (1) the factor share of permanent workers $\alpha(\pi, \delta)$ go down, and (2) the productivity multiplier $A(\pi, \delta)$ goes up. The first effect derives since a smaller share of tasks use permanent workers while the second effect follows from a larger choice set. Lastly, the parameter that determines the substitution elasticity across tasks (γ) is inherited in the CES form to determine the elasticity of substitution between permanent and outsourced workers.

Intermediate Goods Sector - Dynamic Elements

The firms are ex-ante identical, but they are subject to idiosyncratic productivity shocks s that follow an AR(1) process $s' = \rho s + \epsilon$ where $\epsilon \sim N(0, \sigma^2)$ and shocks are independent across firms.²⁹ Adjusting the stock of permanent workers has a cost of $\tau \max\{0, n_- - n\}$, where n_- is the stock of workers that were under contract, n is the new stock of workers, and τ is a per-worker firing cost. The incumbent firms have to pay a fixed cost of operating c every period or exit and pay a one-time cost of firing all workers (τn_-).³⁰ The entrants have to pay a cost of entry c^E before drawing a shock from the distribution $\phi(\cdot)$. Both the fixed cost of operating and the entry cost are paid in the units of final goods.

Timing

The timing of events in a given period is as follows:

1. Entry decisions are made
2. Intermediate firms learn their productivity shocks and decide whether to stay or exit.

²⁹I use revenues to discipline the production function; hence s may represent fluctuations in both prices and quantities. I will call s demand shocks for brevity.

³⁰I use the specification here following the empirical evidence in Bottasso, Conti, and Sulis, 2017 that countries with higher firing costs also have lower firm exit rates. If I modeled the exit cost as a fixed number, my model would generate the opposite pattern. I do not model a separate hiring cost, since its implications are indistinguishable from those of firing costs in this model.

3. Intermediate firms make hiring/firing and outsourcing decisions and produce
4. Final good sector produces

1.4.2 Intermediate Firm's Dynamic Problem

I restrict attention to the steady-state, where firms' distribution across state variables stays constant for all industries. I denote the steady-state value function of the intermediate firm with V :

$$V(s, n_-) = \max \left\{ \max_{n, r} p_k s F(n, r) - n - r - \tau \max\{0, n_- - n\} - Pc + \beta EV(s', n), -\tau n_- \right\} \quad (1.12)$$

where $F(n, r)$ is given in (1.11). p_k and P refer to the intermediate and final good prices, and the wage is normalized to 1. There is a single market wage for the hired and outsourced since outsourcing is provided competitively, and workers are indifferent.³¹ The firm compares the exit cost to the expected discounted value of profits to decide whether to stay in business. The decision to use permanent versus outsourced workers depends both on the structure of $F(n, r)$, and the firing cost τ (See Section 1.4.5). Lastly, potential entrants compare the cost of entry to the expected future discounted profits to decide whether to enter or not. Since the product prices are determined in equilibrium, increased entry moves prices down, depressing the profits firms make, thus feeding back to slow entry.

³¹I only have data on outsourcing expenditures, instead of the number of outsourced workers. Hence, the differences in input prices and factor shares are not separately identified. The model captures any cost savings or markups attached to outsourced workers with the factor share (α).

1.4.3 Equilibrium

A steady-state equilibrium consists of the final good producer's demand for intermediate goods $\{Y_k\}_{k=1}^K$, value and policy functions of the intermediate firms $\{V_k, n_k, r_k\}_{k=1}^K$, the intermediate good prices $\{p_k\}_{k=1}^K$, the final good price P , the measure of entrants in each industry $\{\mu_k\}_{k=1}^K$, and the steady-state distribution of intermediate firms $\{\psi_k\}_{k=1}^K$ that solve

1. $V_k(s, n_-)$ solves (1.12) $\forall k \in K$ (Intermediate Problem)
2. $EV_k(s, 0) = Pc_k^E \quad \forall k \in K$ (Free Entry)
3. $\sum_k \int [n_k(s, n_-) + r_k(s, n_-)] d\psi_k(s, n_-) = L^s$ (Labor Market Clearing)
4. $\psi_k(s, n_-) = T(\psi_k(s, n_-), \mu_k) \quad \forall k \in K$ (Stationary Dist)
5. $\frac{Y_k}{Y_j} = \left(\frac{P_k}{P_j}\right)^{\frac{1}{\omega-1}} \quad \forall k, j \in K$ (Intermediate Good Demand)
6. $P = \left(\sum_k p_k^{\frac{\omega}{\omega-1}}\right)^{\frac{\omega-1}{\omega}}$ (Final Good Price)

1.4.4 Discussion of the Model Elements

The equilibrium defined in 1.4.3 describes the economy of a single state. The model allows four possible channels to explain the state-level differences in outsourcing use: differences in (1) cost of firing, (2) within-industry firm dynamics, (3) industry compositions, and (4) trade secret protection. In this subsection, I discuss how the model generates and quantitatively disciplines each channel.

Since each state recognizes different exceptions to at-will employment, effective firing costs potentially differ across states. The firing costs only apply to the permanent workers in the model, thus, incentivize outsourcing. The model allows industries

to differ in almost all dimensions, including the relative average productivity of outsourcing δ_k . Since industry compositions are available in the data, the model allows controlling for ‘industry fixed-effects’ that would lead to different outsourcing choices across industries.

When the same industry has different outsourcing levels across states, the model does not automatically assign the differences to state policies. Instead, it takes into account that firms that belong to the same industry may be fundamentally different across states and face different operating costs or productivity fluctuations. Only when firms in the same industry have different outsourcing behavior across states that cannot be explained by differences in firm characteristics or the firing costs, the model will assign this to differences in the extent of enforcement friction. Thus, there is a natural link from the enforcement frictions to labor allocation and aggregate output.

Lastly, I conceptualize trade secret theft only as a threat, which never happens in equilibrium. Thus, the model assumes a lack of trade secret protection is unequivocally inefficient, which does not have to be true. The unregulated transmission of secrets in the economy can theoretically be welfare improving. On top of reduced incentives to innovate (Samaniego, 2013), there are two additional barriers against this free flow of ideas. First, when the legal protection is lacking, companies invest in costly physical barriers to prevent theft³² Second, in business partnerships, the sides become more hesitant to share information, which is the main idea of this paper.³³ I

³²Risch, 2007 documents how a client boasted about introducing to the workplace “fingerprint scanners, almost no Internet access, expensive network filtering appliances to scan outgoing email, special locks on the computers, disabled CD-ROM drives, and portable drives, extensive physical security, and so forth.” to avoid trade secret theft. See Henry and Ruiz-Aliseda, 2016a for a theoretical analysis of deterring access to secrets.

³³Increasing collaboration in innovative activities was one of the main aims behind the EU legislation that introduced a uniform trade secret law across the EU in 2016 (Directive on the Protection of Trade Secrets). <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=SWD:2013:0471:FIN:EN:PDF>

assume these effects dominate the gains from the chaotic flow of ideas through theft; i.e., the current level of trade secret protection is below the socially optimal level. The strong correlation between trade secret protection and GDP per capita across countries is consistent with this idea.³⁴

1.4.5 Outsourcing Choice

Characterizing firms' policy functions is difficult in the full model due to discrete exit choice and non-convex adjustment cost. Ignoring entry and exit, assuming a differentiable adjustment cost function $\Phi(n_-, n)$, and a binding information-sharing constraint gives a formula that carries the full model's intuition and allows a simple characterization of the forces at work. The problem of the firm in industry k would simplify to

$$V_k(s, n_-) = \max_{n, r} p_k s A(\pi, \delta_k) (\alpha(\pi, \delta_k) n^{\gamma_k} + (1 - \alpha(\pi, \delta_k)) r^{\gamma_k})^{\frac{\theta_k}{\gamma_k}} - n - r - \Phi(n_-, n) + \beta E V_k(s', n) \quad (1.13)$$

where the ratio of outsourcing expenditures over payroll expenses would become

$$\frac{r}{n} = \left[\frac{1 - \alpha(\pi, \delta_k)}{\alpha(\pi, \delta_k)} \left(1 + \Phi'_2(n_-, n) + \beta E \Phi'_1(n, n') \right) \right]^{\frac{1}{1 - \gamma_k}} \quad (1.14)$$

where Φ'_j is the first derivative of $\Phi()$ according to its j th element. The expenditure share on outsourced workers would increase if adjusting permanent workers is more costly, i.e., Φ has a larger slope. The importance of adjustment costs is further amplified if the expected future adjustments are larger: σ_k^2 is higher or ρ_k is lower. The outsourced share also goes down as it becomes easier to substitute permanent

³⁴See Figure 1.13 in Appendix 1.C. See Ottoz and Cugno, 2011 and Acemoglu and Akcigit, 2012 for theoretical analyses of the optimal scope of trade secret protection.

workers for rented workers, that is, when γ_k is higher. Lastly, firms outsource more when the factor share of outsourcing is larger, i.e., $\alpha(\pi, \delta_k)$ is lower. $\alpha(\pi, \delta_k)$ is low when either the relative marginal product of outsourcing (δ) or the strength of trade secret protection (π) is high.

Although this simplified analysis helps tease out some of the model’s central mechanisms, I estimate the full model in the next section. The estimation confirms that the general equilibrium effects have a significant impact on aggregate outsourcing.

1.4.6 Model Extensions

I solve the model numerically, using grid-search on the value functions and forward iterations to compute firms’ stationary distributions. I make a couple of adjustments before estimating the model. These do not affect the primary mechanism but simplify the computation and the estimation of the model.

First, I discretize the idiosyncratic productivity process to 10 grid points using Rouwenhorst, 1995’s method. Second, I add Type 1 Extreme Value (T1EV) shocks to the exit decision, ensuring the equilibrium moments change smoothly with parameter values which simplifies the estimation procedure. Each period, to continue operating, firms need to pay $c^F + \nu_1$, or they exit and pay $\tau n_+ \nu_2$ where ν_1, ν_2 are identically distributed T1EV shocks with shape parameter η . I assume the ν_1, ν_2 are independent over time, across firms, from productivity shocks, and one another. The difference of two T1EV shocks has a logistic distribution, which allows the analytical characterization of the probability that a firm with state (s, n_-) chooses to exit. Last, as in Boedo and Mukoyama, 2012, incumbents receive an ‘offer they cannot refuse’ after production ends with probability κ_j and have to exit. This shock helps generate realistic exit patterns in the model for large establishments.

1.5 Estimation

In this section, I estimate the model to make quantitative statements. Section 1.5.1 describes the data, the estimation procedure and the identification strategy. The estimation results are in 1.5.2. Section 1.5.3 evaluates the ability of the model to match untargeted moments. Section 1.5.4 provides the quantitative decomposition of state-level outsourcing heterogeneity while productivity gains from better trade secret protection are discussed in Section 1.6.

1.5.1 Data and Estimation Method

I use establishment-level moments for each state-industry pair in the manufacturing sector (NAICS 31-33) from 2007 to estimate the model. I use three primary data sources to compute the moments. The Census of Manufactures (CMF) provides state-industry level revenue shares, revenue to payroll ratios, and outsourcing expenditures. The Statistics of U.S. Businesses (SUSB) provides state-industry level moments on establishment size distribution. Lastly, the Business Dynamics Statistics (BDS) provides state-level moments on job flows, which are only available at the manufacturing sector level.

The model has parameters that are global, industry-specific, state-specific, and state-industry specific. I use subscript j to denote that the parameter varies across states and k to denote it varies across industries. The full set of parameters necessary to compute the extended model is the vector³⁵:

$$\Omega = \{\beta, \omega, \gamma_k, \sigma_k^2, \kappa_j, \tau_j, c_{jk}^F, c_{jk}^E, \rho_{jk}, \theta_{jk}, \pi_j, \delta_k\} \quad (1.15)$$

³⁵I fix the productivity distribution of entrants (ϕ) and the shape parameter for the T1EV shocks for now.

I set β and ω to standard values, and γ_k and σ_k^2 to previous estimates in the literature. I estimate the rest of the parameters $(\kappa_j, \tau_j, c_{jk}^F, c_{jk}^E, \rho_{jk}, \theta_{jk}, \pi_j, \delta_k)$ in two stages. The first stage assumes the information sharing constraint binds and treats $\alpha(\pi_j, \delta_k)$ in (1.11) as a state-industry level parameter α_{jk} . This assumption allows the first stage to be estimated separately for each state. This substantially relieves the computational burden since the stationary distribution of the firms has to be solved numerically. The second stage treats α_{jk} as data generated by $\alpha(\pi_j, \delta_k) + \epsilon_\alpha$ where ϵ_α are zero-mean iid shocks and uses non-linear least squares to estimate $\{\pi_j\}_{j=1}^J$ and $\{\delta_k\}_{k=1}^K$.

Externally Set Parameters

I set the discount factor $\beta = 0.94$ and the parameter governing the demand substitution between intermediate goods to $\omega = -0.5$. Two sets of parameters are hard to identify with the available data. The first is the elasticity of substitution parameter between permanent and outsourced workers. Identifying it either requires wage data with an exogenous wage shifter or an establishment-level panel with information on dynamic inputs. Neither data is available, so I take the estimates of Chan, 2017 directly, who uses an establishment panel from Denmark to do the latter³⁶ for four manufacturing industry groups. The second is the variance of the productivity process. It is not possible to nonparametrically identify both the persistence and the variance of an AR(1) process from cross-sectional data. I take the industry-level estimates from Bloom, Floetotto, et al., 2018, who use the Annual Survey of Manufacturers to estimate an AR(1) process for the log TFP estimates for each manufacturing industry.³⁷

³⁶Both the relative size of the outsourcing sector, and its skill composition are remarkably similar between Denmark and the U.S.

³⁷Unlike this paper, Bloom, Floetotto, et al., 2018 estimate value-added production functions and include capital and materials. However, for a Cobb-Douglas production function between materials,

Method of Moments Estimation and Identification Idea

I estimate $\Omega_E = \{\kappa_j, \tau_j, c_{jk}^F, c_{jk}^E, \rho_{jk}, \theta_{jk}, \alpha_{jk}\}$ via method of moments, minimizing the weighted distance between the model $M(\Omega_E)$ and data M^D moments:

$$\hat{\Omega}_E = \arg \min_{\Omega_E} \left(M^D - M(\Omega_E) \right)' W \left(M^D - M(\Omega_E) \right) \quad (1.16)$$

where W is a weighting matrix. The estimator is consistent for any choice of W , but the efficient estimator has $W = V^{-1}$, i.e., the inverse covariance matrix of the data moments. Estimating the covariance matrix requires micro-data. I instead use a diagonal matrix where $W_{nn} = (M_n^D)^{-2}$, which transforms the objective function into one that minimizes total squared percent deviations.

The model admits a general equilibrium where common labor and product markets connect all establishments in a state. The steady-state distribution of firms does not have a closed-form solution either; thus, I can only provide intuitive arguments on why the selected moments inform the structural parameters. I suppress the state subscript j as all the parameters here are state-specific. The only parameter that maps one-to-one to a moment is the exogenous exit probability κ . The model generates essentially no endogenous exit for the largest firms; thus, κ becomes equal to the exit probability of large establishments (more than 250 employees).

The aggregate entry rate, average establishment size, and the revenue shares of industries jointly inform c_k , the fixed cost of operating, and c_k^E , the entry cost. Both a small c_k and a small c_k^E incentivize entry and are associated with a large industry. Thus, a decrease in either cost would increase the revenue share of an industry. On the other hand, the average establishment size moves in opposite directions when c_k

capital, labor services (CES of permanent and outsourced workers), and competitive input markets, their variance estimates can be applied to my setting up to a constant multiplier. The multiplier is not identified in my model; hence, its value is irrelevant for the estimation. See Table 1.10 for the calibrated values of γ_k and σ_k .

and c_k^E increases. A large average establishment size is associated with a large c_k because establishments would not find it profitable to pay a high operating cost at a small scale and exit instead. On the other hand, a small cost of entry c_k^E would result in a large average establishment size, as the competitive pressure through new entrants would lead small unproductive firms to exit. Thus the two moments provide a single crossing condition for the two parameters. Lastly, the economy's overall scale is not pinned down; therefore, there are only $K - 1$ linearly independent revenue shares. The aggregate entry rate helps pin down the average level of entry costs across industries.

While an increase in the returns to scale parameter θ_k increases both the average establishment size and the revenue share of an industry, the ratio of revenues to payroll expenses allows distinguishing it from c_k and c_k^E . The two costs have no direct influence on this ratio, except through the firms' steady-state distributions. On the other hand, θ_k directly impacts the labor share of revenues by determining the elasticity of revenues to the labor inputs.

It is relatively easier to distinguish the persistence of the idiosyncratic shocks ρ_k and the firing cost τ from the parameters I discussed so far (c_k , c_k^E , and θ_k): while the latter parameters have first-order effects only on the first moments of the firm distribution, ρ_k and τ are crucial for the second moments and the flows.³⁸ On the other hand, it is notoriously difficult to separately identify adjustment costs and the parameters of the idiosyncratic shock process (Bloom, 2009, Decker et al., 2020). I use the share of small establishments (less than 20 employees) and the aggregate job destruction rate. Both a high persistence and a high firing cost reduce the rate of job destruction. If shocks' persistence is high, establishments face the need to

³⁸The only exception to this is the impact on the entry rate, which directly affects the job destruction rate. In model validation, I specifically check whether the estimated model does a good job matching the fraction of job flows through exits.

change their workforce less frequently while under high firing costs, establishments choose to operate at a sub-optimal scale instead of having to fire workers later. The two parameters also impact the share of small establishments in the same direction. If persistence is high, entrants stay small for a long time until their productivity increases. High firing costs also discourage establishments from increasing the number of workers anticipating the possibility of having to fire them later. On the other hand, for a wide range of reasonable firing costs (0 to 4 years of wages), the impact on the share of small establishments is modest (less than 1%). Thus, a local single crossing condition is satisfied. The intuition for the modest impact of firing costs relies on the firm size distribution's long right tail. Given the high fixed costs of operating and low returns to scale parameters, the return from hiring workers is very high for small productive firms.³⁹

Last but not least, the ratio of outsourcing expenses to payroll expenses helps identify α , the factor share of permanent workers. As discussed in Section 1.4.5, the parameters that have a direct effect on the ratio of outsourcing expenses are γ , σ^2 , ρ , τ and α . I externally calibrate γ and σ^2 with structural estimates from the literature. The share of small establishments again helps distinguish ρ from α , as the impact of α is negligible once the average size of establishments is held constant. Finally, although both a low α and a high τ increase the ratio, the large response of job destruction rate and the small response of the outsourcing ratio to τ allows distinguishing the two.

³⁹One moment that would allow a global identification would be the 'job destruction' rate for outsourced workers, i.e., the average decline in outsourcing expenses for firms that decrease their outsourcing. Because outsourcing is not subject to firing costs, its flow helps discipline the fluctuations in the idiosyncratic shock process. Unfortunately, there are no public estimates for this moment.

Nonlinear Least Squares

In the second stage, I minimize the sum of squared residuals between the model implied $\alpha(\pi_j, \delta_k)$ as derived in (1.11) and $\hat{\alpha}_{jk}$ estimates from the first stage (1.16):

$$\{\hat{\pi}_j, \hat{\delta}_k\} = \arg \min_{\{\pi_j, \delta_k\}} \sum_{j,k} (\hat{\alpha}_{jk} - \alpha(\pi_j, \delta_k))^2 \quad (1.17)$$

This procedure is similar in spirit to a fixed effects regression; once the factor shares are estimated, the ‘state fixed effects’ give the π_j and the ‘industry fixed effects’ give the δ_k . Similar to a two-way fixed-effects regression, it is impossible to separately identify the level of π_j from the level of δ_k . Therefore, in the counterfactuals, I do a normalization a la Hsieh and Klenow, 2009 and consider the state with the largest π_j as unconstrained and use it as the baseline for comparisons based on enforcement frictions. Table 1.4 summarizes the full calibration/estimation strategy, together with data sources. The first four rows of parameters are externally calibrated. The ones in the middle are jointly estimated to match the moments in the first stage. The ones in the last two rows are jointly estimated to match the α_{jk} estimates from the first stage.

1.5.2 Estimation Results

I have estimated the model for 28 states so far, where I divide the manufacturing sector into $K = 4$ industry groups: Food Products ($k = 1$), Wood and Paper Products ($k = 2$), Heavy Industry and Extraction ($k = 3$), and Tools, Machinery and Consumer Goods ($k = 4$). Figure 1.5a presents the estimated factor shares for all industry-state groups.⁴⁰

⁴⁰I follow the same grouping as in Chan, 2017 to have a one-to-one match with his γ_k estimates. The details of how I match the U.S. NAICS 3-digit sectors with the Danish NACE 2-digit sectors are in Appendix 1.B. The first-stage in-sample results are in Table 1.11, where I provide the results

Table 1.4: The Main Parameters and the Moments Used in the Estimation Notes: The details of the data sources and how the moments are calculated can be found in Appendix 1.B.


Par	Role	Moment	Source
β	Discount Factor	External	0.94
ω	Int. Good Subst.	External	-0.5
γ_k	Permanent/Outsourced Subst.	External	Chan 2017
σ_k^2	Idio. Shock Variance	External	Bloom et al. 2018
κ_j	Exog Exit Prob	Exit Rate>250	BDS
τ_j	Firing Cost	Job Destruc. Rate	BDS
c_{jk}	Fixed Cost of Operating	Avg. Estb Size	SUSB
c_{jk}^E	Entry Cost	Ind. Output Shares	CMF
ρ_{jk}	Idio. Shock Persistence	Share of Estb Size<20	SUSB
θ_{jk}	Returns to Scale	Receipts/Payroll	CMF
α_{jk}	Permanent Factor Share	Outsourcing/Payroll	CMF
π_j	Trade Secret Enforcement	Agg. Entry Rate	BDS
δ_k	Outsourcing Suitability	$\hat{\alpha}_{jk}$	1st Stage

Figure 1.5b summarizes how the estimated factor share parameters relate to the observed outsourcing ratios. In a model with no adjustment costs, the outsourcing ratios would only depend on γ_k and α_{jk} because there would be no flexibility gains from outsourcing. The cross-state patterns are as expected within each industry. However, the estimates suggest the factor share of outsourcing is considerably lower in food manufacturing, even though it outsources as much as the other industry groups. Also, the estimates for heavy manufacturing are broadly similar to wood manufacturing, even though heavy manufacturing has a considerably higher outsourcing to payroll ratio.

Two channels mainly drive these results. First, permanent and outsourced workers are easier to substitute in food and heavy manufacturing, according to the externally calibrated γ_k values (Table 1.10). This implies a larger outsourcing ratio for a fixed $\alpha_{jk} > 0.5$ (see (1.14)). Second, in the data, food and heavy manufacturing establishments have a larger revenue to payroll ratio, even though their average size is not

for Michigan for brevity.

Figure 1.5: The Estimation Results from the 1st Stage Notes: Each shape refers to a state-industry pair. See Table 1.11 for details on the first and second stage estimation results and Appendix 1.B for details on the outsourcing to payroll ratios.




(a) Estimated Outsourcing Factor Shares ($1 - \alpha_{jk}$)

(b) Outsourcing to Payroll Ratios vs Estimated Outsourcing Factor Shares ($1 - \alpha_{jk}$)

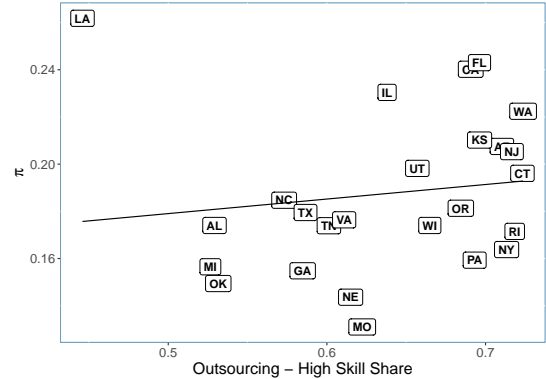
significantly different than the other two groups. Hence, they are estimated to have low θ_{jk} and c_{jk}^E and high c_{jk} (See Figure 1.15). The low returns to scale together with high fixed costs create a fat-tailed size distribution, and the low c_E ensures the total size of these industries is as large as in the data. In the model, larger firms outsource a bigger fraction of their workforce, fearing mass layoffs in the future. The very large firms in the food and heavy manufacturing hence outsource a large fraction of their workforce, generating the pattern in Figure 1.5a. Lastly, these two effects are large enough to offset the lower-variance productivity shocks for food and heavy manufacturing, given the externally calibrated σ_k values.

Table 1.12 presents the results from the second stage; hence the main estimation results. I find, without enforcement frictions, the industry that would benefit the most from outsourcing is heavy manufacturing, and the one that would benefit the least is food manufacturing. The average productivity of an outsourced worker (δ_k) is estimated to be twice as large in the former than the latter (0.36 vs 0.16). Louisiana is the state with the strongest secret protection, and Missouri is the one with the weakest. Most importantly, as Figure 1.6a shows, the results from the struc-

Figure 1.6: The Estimation Results from the 2nd Stage Notes: Figures only presents states who adopted the UTSA by 2007.



(a) Estimated Strength of Protection (π) vs Number of Years Since the Adoption of the UTSA (Correlation 0.37)



(b) Estimated Strength of Protection (π) vs the Share of Outsourcing in High-Skilled Tasks (Correlation 0.14, 0.51 without LA)

natural estimation align with the adoption date of the UTSA. The states that adopted the UTSA earlier are the ones that have better trade secret protection on average. Figure 1.6b further shows that states with better protection spend a larger fraction of their labor outsourcing budget on high-skill tasks. The two figures provide an important first step for validating the model: the estimation results are consistent with (1) the actual legal environment of the states and (2) laws being more important for information-sensitive tasks, even though neither pattern was targeted in the estimation.

1.5.3 Model Validation

I validate the model through its ability to match the share of job destruction that happens through establishment exits, establishment shares of industry groups, and the share of employment in small establishments.

Although the estimation targets the rates of exit and job destruction, the share of job destruction through exits can be anywhere between 0 and 1 depending on the

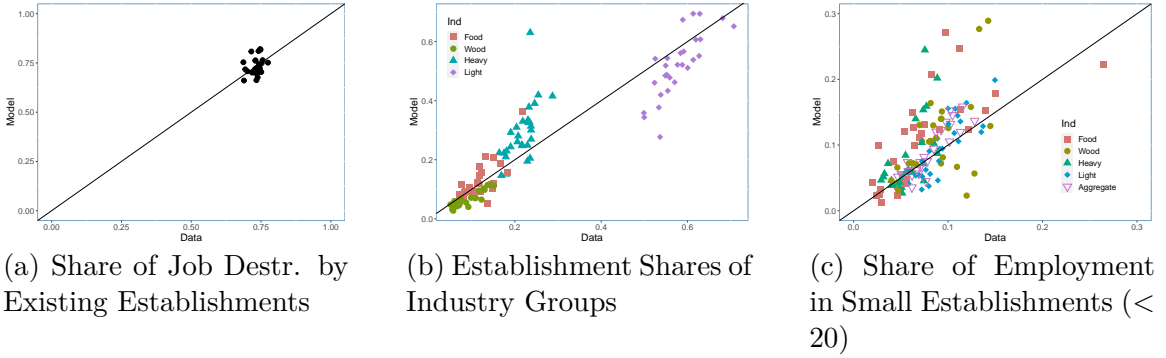
exiting establishments' average size. The model does an excellent job of predicting the share (Figure 1.7a), hence the average size of exiting establishments. The estimation targets the revenue share, the revenue payroll ratio, and the average establishment size for each industry group. If workers' average wages across industries differed significantly, the model would do a bad job predicting the fraction of establishments that belong to each industry. Figure 1.7b suggests the model still does a good job. The only exceptions are the wages at California's Light and Heavy industries, where the model undervalues the former and overvalues the latter. Lastly, the model targets the share of establishments with less than 20 employees but does not target the size distribution below 20. If the model did a bad job at matching that distribution, it would make a bad prediction of the expected size of establishment conditional on less than 20. Figure 1.7c suggests the model does an okay job, except for food manufacturing, which is a relatively smaller part of the manufacturing sector. In particular, the model cannot account for the states with small food manufacturing establishments.

The model does a poor job predicting the size distribution's right-tail, generating too few very-large establishments (larger than 250, 500). The model's inability to match both tails is partly due to the assumption of normal shocks to the productivity process. A shock distribution that has fatter tails would help the model generate more large establishments.

1.5.4 Decomposition of the Outsourcing Heterogeneity

In this section, I ask how the cross-state heterogeneity in labor outsourcing would change if all states had the same (1) firing cost, (2) industry composition, (3) within-industry firm characteristics, and (4) trade secret protection. According to the model, these four objects constitute a mutually exclusive and exhaustive list of the differences

Figure 1.7: The Untargeted Moments



between states. However, they might interact with one another and amplify/dampen each other's effects. Notably, the industry composition and the within-industry firm characteristics are equilibrium objects, making the decomposition non-trivial.

To equate the labor protection and the trade secret protection across states, I replace the values of τ and π with the average estimates. To 'equate' the industry compositions, I take simple weighted averages of industry-level outsourcing shares for each state, weights being average industry share of employment across states. To find the impact of equating within-industry firm characteristics, I take the average values of the other three (τ , π , and industry shares) for each state and compute the remaining dispersion (See Figure 1.14). Now I can answer one of the main questions I have started with: what generates the cross-state dispersion in outsourcing use? I use the coefficient of variation (standard deviation divided by the average) as my measure of dispersion. The cross-state dispersion would be

- 19% less with average trade secret protection
- 14% less with average industry composition
- 14% more with average firing cost
- 90% less with average within-industry firm characteristics

The differences in within-industry firm characteristics create the lion’s share of the observed dispersion across states. While equating industry shares would reduce the heterogeneity, equating firing costs would amplify it. The counter-intuitive implication is that the states with the higher estimated firing costs outsource less than others on average due to the other three channels’ counteracting force.

Equating the strength of trade secret protection decreases the cross-state dispersion by 19%. This result, however, is built on considerable heterogeneity across states. In particular, there are states with weak trade secret protection that still outsource a significant amount of their workforce. Bringing the strength of trade secret protection up to the average level increases outsourcing shares for these states, pushing for increased dispersion. For example, Tennessee is a state with an above-average outsourcing ratio of 0.2, and improving its trade secret protection up to the average level would bring the ratio up to 0.22.

1.6 Productivity Gains from Better Trade Secret Protection

In this section, I answer the question I started with: how large are the productivity gains from better trade secret protection? Specifically, I calculate the counterfactual outcomes when every state has the same trade secret protection (π) as the ‘best state,’ which is Louisiana, according to my estimates.

Table 1.5 presents the main results. The median state increases its outsourcing to payroll ratio from 0.12 to 0.17. While both the gross and the net output (net of all costs) of the median state grows by 0.9%, the state that benefits the most has a net output growth as large as 2%. The growth is mostly through the entry channel: the number of firms increases by 0.8% in the median state. Lastly, wages

also reflect productivity growth, increasing by as much as 1.4% for the median state. I compute the aggregate gains as the weighted average of the net output gains in each state, where the weights are equal to each state’s manufacturing output in 2007. The aggregate output grows by 0.7%. In the remainder of the section, I quantify individual channels that lead to output gains.

Table 1.5: The Counterfactual Results After an Improvement in Trade Secret Protection Notes: The first and second rows give the result for the median and maximum value across states. The third row gives the aggregate response, which is an output-weighted average of the responses of states. The values for columns 4 to 7 are relative to a baseline value of 1. Base and Best TSP refer to the outsourcing to payroll ratio in the baseline estimation and the counterfactual where each state’s π is equal to the state with the highest π . Gross Output is the aggregate amount of final goods produced, and the net output is gross output net of all entry, operating, and firing costs. The number of firms is aggregated over industries. See Table 1.14 for state-by-state details.

	Base	Best TSP	Gross Out	Net Out	# of Firms	Wage
Median	0.12	0.17	1.009	1.009	1.008	1.014
Max	0.20	0.26	1.019	1.020	1.020	1.029
Aggregate	0.14	0.18	1.007	1.008	1.006	1.014

The Role of Labor Adjustment Costs

Improved trade secret protection decreases the job destruction rate, i.e., increased outsourcing leads to more job stability for permanent manufacturing employees. Yet, the aggregate decline is relatively small, from 10.80% to 10.77%. Although the job destruction rate remains relatively constant, the total amount of job destruction declines substantially because the fraction of workers under employment goes down. These lead to savings through avoided firing costs: even though the number of firms increases by 0.6%, the aggregate firing cost paid declines by 2.7%. The magnitude of the savings is small on the macroeconomic scale (4 basis points of GDP).

On the other hand, the gains from better allocation of workers are significant. The dispersion of the marginal product of labor across firms (0 at a frictionless equilibrium) declines by 1.4%. The correlation between size and productivity, a commonly used


measure of labor (mis)allocation between firms (1 at a frictionless equilibrium), would also have a modest increase for both in-house employees and outsourced workers from 0.822 to 0.824 and from 0.854 to 0.856 respectively. In other words, the reduction in the firms that have excess and too little employed workers leads to a better allocation of outsourced workers across firms as well.

Entry and Exit

The entry/exit channel impacts the aggregate gains both through the number of firms that operate in the steady-state and through the rate of entry/exit as a force that generates steady creative destruction. Although the aggregate rate of entry/exit goes up, it is quantitatively small: the change is 4 basis points relative to a baseline level of 7.54%. On the other hand, the number of firms in the steady-state increases substantially by 0.6%. This increase is reflected by the economically significant growth in aggregate entry costs and operating costs paid by 0.7% and 0.6% (0.1% and 0.2% of GDP).

The increase in the number of firms is accompanied by a 0.4 p.p. increase in small firms' share (less than 20 employees). This increase is not surprising since the total number of employees employed by the manufacturing firms decreases while the total number of firms increases, i.e., the average firm size must be decreasing. A decrease in the fraction of very large firms accompanies the increase in small firms' fraction. While small firms find it easier to grow in size with the added flexibility provided by outsourcing, they also face more intense competition for workers due to the increased number of firms. For the large firms, flexibility and competition work in the same direction: they find it easier to decrease their size after bad shocks. Hence, firms *hoard labor* to a lesser extent when the outsourcing sector is larger.

Figure 1.8: The State-Industry Level Gains



(a) The Growth in Outsourcing to Payroll Ratio

(b) The % Growth in Net Output

The Role of Industries

The industries differ in δ_k ; therefore, the importance of trade secret protection is potentially different across industries, which helps explain why some states enjoy more significant gains from improved protection than the others.

Figure 1.8 shows the industry that changes its workforce composition the most is heavy manufacturing, followed by light manufacturing. Both industries heavily rely on secrecy for comparative advantage. The secret formulas and processes are integral parts of light and chemical manufacturing. The negative information on R&D, which cannot be patented, is critical for pharmaceuticals. Similarly, the information on the location of raw materials and manufacturing processes is essential for oil and metals industries.

On the other hand, the industry-level output growth rates are much more similar to one another than the outsourcing growth. This similarity is largely driven by the value of the parameter that controls the demand elasticity of the final good producer ($\omega = -0.5$).⁴¹ Since intermediate goods are gross complements, an increase in

⁴¹I choose an elasticity that implies gross complementarity between intermediate goods because I estimate the model using a revenue (instead of value-added) production function. Since I do not

one intermediate industry's productivity increases the demand for other intermediate industries. This complementarity aligns the output of different industries together; hence all industries benefit from a productivity gain in one industry.

1.7 Conclusion

I study the impact of trade secret protection on producers' willingness to use outsourced workers, and consequently, aggregate output. Through an analysis of this channel in the U.S. I make two main points. First, better legal protection for trade secrets can induce managers to use outsourced workers for a larger number of tasks. Second, the consequent expansion in outsourcing use generates a better allocation of workers across firms and a quantitatively significant increase in aggregate output.

To make the first point, I rely on the Uniform Trade Secrets Act and utilize the variation in adoption times across states. My analysis shows that adopters enjoyed a higher pace of subsequent growth in outsourcing employment relative to non-adopters. Also, the effect was more pronounced for tasks that provide greater access to sensitive information. Quantitatively, the improvements in trade secret law explain 13% of the growth in outsourcing employment in the U.S. from 1977 to 1997.

I build and estimate a structural model of industry dynamics to make the second point. The model teases out the part of cross-state heterogeneity in outsourcing that is attributable to variation in trade secret protection and maps it to aggregate productivity measures. Estimating it with data from the U.S. manufacturing sector shows that the gains from better trade secret protection are sizeable. If all states could protect trade secrets as adequately as the 'best state,' the aggregate output would increase by 0.8%.

model an explicit production network between manufacturing industries, I introduce a reduced-form supply chain through complementarity in final good production.

These findings suggest large gains for the U.S., a country that is at the forefront of trade secret protection (See Figure 1.13). The gains might be even larger for countries where the statutory law is still missing, common law is underdeveloped, or the enforcement of existing laws is lacking. Improving legal protection requires trained judges, lawyers, expert witnesses, and functioning audit and appeals systems that supervise the legal system. None of these come easy or cheap, but neither do tax breaks or R&D subsidies.

References

- Abraham, Katharine and Susan K Taylor (1996). “Firms’ use of outside contractors: Theory and evidence”. In: *Journal of labor economics* 14.3, pp. 394–424.
- Acemoglu, Daron and Ufuk Akcigit (2012). “Intellectual property rights policy, competition and innovation”. In: *Journal of the European Economic Association* 10.1, pp. 1–42.
- Acemoglu, Daron and Pascual Restrepo (2018). “The race between man and machine: Implications of technology for growth, factor shares, and employment”. In: *American Economic Review* 108.6, pp. 1488–1542.
- Adler, Paul S (2003). “Making the HR outsourcing decision”. In: *MIT Sloan Management Review* 45.1, p. 53.
- Akcigit, Ufuk, Harun Alp, and Michael Peters (Jan. 2021). “Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries”. In: *American Economic Review* 111.1, pp. 231–75. DOI: [10.1257/aer.20180555](https://doi.org/10.1257/aer.20180555). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20180555>.
- Alchian, Armen A and Harold Demsetz (1972). “Production, information costs, and economic organization”. In: *The American economic review* 62.5, pp. 777–795.

- Alfaro, Laura, Paola Conconi, Harald Fadinger, and Andrew F Newman (2016). “Do prices determine vertical integration?” In: *The Review of Economic Studies* 83.3, pp. 855–888.
- Almeling, David S, Darin W Snyder, and Michael Sapoznikow (2009). “A statistical analysis of trade secret litigation in federal courts”. In: *Gonz. L. Rev.* 45, p. 291.
- Almeling, David S, Darin W Snyder, Michael Sapoznikow, and Whitney E McCollum (2010). “A statistical analysis of trade secret litigation in state courts”. In: *Gonz. L. Rev.* 46, p. 57.
- Amiti, Mary, Shang-Jin Wei, Jonathan Haskel, and Emmanuelle Auriol (2005). “Fear of Service Outsourcing: Is It Justified?” In: *Economic Policy* 20.42, pp. 307–347. ISSN: 02664658, 14680327. URL: <http://www.jstor.org/stable/3601072>.
- Antras, Pol (2003). “Firms, contracts, and trade structure”. In: *The Quarterly Journal of Economics* 118.4, pp. 1375–1418.
- Antras, Pol, Teresa C Fort, and Felix Tintelnot (2017). “The margins of global sourcing: Theory and evidence from us firms”. In: *American Economic Review* 107.9, pp. 2514–64.
- Autor, David H (2001). “Why do temporary help firms provide free general skills training?” In: *The Quarterly Journal of Economics* 116.4, pp. 1409–1448.
- (2003). “Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing”. In: *Journal of labor economics* 21.1, pp. 1–42.
- Autor, David H and Susan N Houseman (2010). “Do temporary-help jobs improve labor market outcomes for low-skilled workers? Evidence from “Work First””. In: *American economic journal: applied economics* 2.3, pp. 96–128.
- Autor, David H, William R Kerr, and Adriana D Kugler (2007). “Does employment protection reduce productivity? Evidence from US states”. In: *The Economic Journal* 117.521, F189–F217.

- Bartelsman, Eric J and Mark Doms (2000). “Understanding productivity: Lessons from longitudinal microdata”. In: *Journal of Economic literature* 38.3, pp. 569–594.
- Bento, Pedro and Diego Restuccia (2017). “Misallocation, establishment size, and productivity”. In: *American Economic Journal: Macroeconomics* 9.3, pp. 267–303.
- Bentolila, Samuel and Giuseppe Bertola (1990). “Firing costs and labour demand: how bad is eurosclerosis?” In: *The Review of Economic Studies* 57.3, pp. 381–402.
- Bentolila, Samuel and Gilles Saint-Paul (1992). “The macroeconomic impact of flexible labor contracts, with an application to Spain”. In: *European Economic Review* 36.5, pp. 1013–1047.
- Bergeaud, Antonin, Clement Malgouyres, Clement Mazet-Sonilhac, and Sara Signorelli (2020). “Technological Change and Domestic Outsourcing”. Working Paper.
- Berlingieri, Giuseppe (Apr. 2013). “Outsourcing and the Rise in Services”. CEP Discussion Papers. URL: <https://ideas.repec.org/p/cep/cepdps/dp1199.html>.
- Beulen, Erik and Pieter Ribbers (2010). *Managing IT outsourcing*. Routledge.
- Bidwell, Matthew J (2012). “Politics and firm boundaries: How organizational structure, group interests, and resources affect outsourcing”. In: *Organization Science* 23.6, pp. 1622–1642.
- Bilal, Adrien and Hugo Lhuillier (2020). “Outsourcing, Inequality and Aggregate Output”. University of Chicago, Becker Friedman Institute for Economics Working Paper.
- Bloom, Nicholas (2009). “The impact of uncertainty shocks”. In: *Econometrica* 77.3, pp. 623–685.

- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (2013). “Does management matter? Evidence from India”. In: *The Quarterly Journal of Economics* 128.1, pp. 1–51.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry (2018). “Really uncertain business cycles”. In: *Econometrica* 86.3, pp. 1031–1065.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012). “The organization of firms across countries”. In: *The Quarterly Journal of Economics* 127.4, pp. 1663–1705.
- Boedo, Hernan J Moscoso and Toshihiko Mukoyama (2012). “Evaluating the effects of entry regulations and firing costs on international income differences”. In: *Journal of Economic Growth* 17.2, pp. 143–170.
- Boehm, Johannes (Forthcoming). “The Impact of Contract Enforcement Costs on Value Chains and Aggregate Productivity”. In: *The Review of Economics and Statistics*. eprint: https://doi.org/10.1162/rest_a_00940. URL: https://doi.org/10.1162/rest_a_00940.
- Boehm, Johannes and Ezra Oberfield (2020). “Misallocation in the Market for Inputs: Enforcement and the Organization of Production”. In: *The Quarterly Journal of Economics* 135.4, pp. 2007–2058.
- Bolatto, Stefano, Alireza Naghavi, Gianmarco Ottaviano, and Katja Zajc Kejzar (Jan. 2020). “Intellectual property and the organization of the global value chain”. CEP Discussion Papers. URL: <https://ideas.repec.org/p/cep/cepdps/dp1673.html>.
- Bottasso, Anna, Maurizio Conti, and Giovanni Sulis (2017). “Firm dynamics and employment protection: Evidence from sectoral data”. In: *Labour Economics* 48, pp. 35–53.

- Bradsher, Keith (Jan. 2020). “How China Obtains American Trade Secrets”. In: *The New York Times*. URL: <https://www.nytimes.com/2020/01/15/business/china-technology-transfer.html>.
- Bruhn, Miriam, Dean Karlan, and Antoinette Schoar (2018). “The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico”. In: *Journal of Political Economy* 126.2, pp. 635–687. DOI: [10.1086/696154](https://doi.org/10.1086/696154). eprint: <https://doi.org/10.1086/696154>. URL: <https://doi.org/10.1086/696154>.
- Caballero, Ricardo J and Mohamad L Hammour (Dec. 1994). “The Cleansing Effect of Recessions”. In: *American Economic Review* 84.5, pp. 1350–1368.
- Caggese, Andrea and Vicente Cuñat (Oct. 2008). “Financing Constraints and Fixed-term Employment Contracts”. In: *The Economic Journal* 118.533, pp. 2013–2046. ISSN: 0013-0133.
- Cahuc, Pierre and Fabien Postel-Vinay (2002). “Temporary jobs, employment protection and labor market performance”. In: *Labour economics* 9.1, pp. 63–91.
- Callaway, Brantly and Pedro Sant’Anna (2020). “Difference-in-Differences with multiple time periods”. In: *Journal of Econometrics*. ISSN: 0304-4076.
- Chaisemartin, Clément de and Xavier D’Haultfœuille (Sept. 2020). “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects”. In: *American Economic Review* 110.9, pp. 2964–96. DOI: [10.1257/aer.20181169](https://www.aeaweb.org/articles?id=10.1257/aer.20181169). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20181169>.
- Chan, Mons (2017). “How Substitutable are Labor and Intermediates?” In: *Unpublished working paper. University of Minnesota*.
- Coase, Ronald H. (1937). “The Nature of the Firm”. In: *Economica* 4.4, pp. 386–405.
- Cohen, Joel M. and Mary Kay Dunning (Dec. 2010). “Insider Trading: It’s Not Just for Suits”. In: *Business Crimes Bulletin* 18.4.

- Corcos, Gregory, Delphine M Irac, Giordano Mion, and Thierry Verdier (2013). “The determinants of intrafirm trade: Evidence from French firms”. In: *Review of economics and statistics* 95.3, pp. 825–838.
- Covey, Anne E and Michael S Morris (1983). “The Enforceability of Agreements Providing for Forum and Choice of Law Selection”. In: *Denv. LJ* 61, p. 837.
- Crawford, Gregory S, Robin S Lee, Michael D Whinston, and Ali Yurukoglu (2018). “The welfare effects of vertical integration in multichannel television markets”. In: *Econometrica* 86.3, pp. 891–954.
- Davis, Steven J and John Haltiwanger (1992). “Gross job creation, gross job destruction, and employment reallocation”. In: *The Quarterly Journal of Economics* 107.3, pp. 819–863.
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda (Dec. 2020). “Changing Business Dynamism and Productivity: Shocks versus Responsiveness”. In: *American Economic Review* 110.12, pp. 3952–90. DOI: [10.1257/aer.20190680](https://doi.org/10.1257/aer.20190680). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20190680>.
- Drenik, Andres, Simon Jäger, Miguel Pascual Plotkin, and Benjamin Schoefer (2020). “Paying outsourced labor: Direct evidence from linked temp agency-worker-client data”.
- Eckert, Fabian, Teresa C Fort, Peter K Schott, and Natalie J Yang (2020). “Imputing Missing Values in the US Census Bureau’s County Business Patterns”.
- Elzankaly, Marwa (May 2018). “The Not-So-Powerful Non-Disclosure Agreement (NDA)”. In: URL: <https://www.mcmanislaw.com/blog/2018/the-not-so-powerful-non-disclosure-agreement-nda>.
- Ethier, Wilfred J and James R Markusen (1996). “Multinational firms, technology diffusion and trade”. In: *Journal of International Economics* 41.1-2, pp. 1–28.

- Fisk, Catherine L (2000). “Working knowledge: Trade secrets, restrictive covenants in employment, and the rise of corporate intellectual property, 1800-1920”. In: *Hastings LJ* 52, p. 441.
- Fort, Teresa C, Shawn D Klimek, et al. (2016). “The effects of industry classification changes on US employment composition”. In: *Tuck School at Dartmouth*.
- Fosfuri, Andrea, Massimo Motta, and Thomas Rønde (2001). “Foreign direct investment and spillovers through workers’ mobility”. In: *Journal of international economics* 53.1, pp. 205–222.
- Foster, Lucia, John C Haltiwanger, and Cornell John Krizan (2001). “Aggregate productivity growth: lessons from microeconomic evidence”. In: *New developments in productivity analysis*. University of Chicago Press, pp. 303–372.
- Fujita, Shigeru and Makoto Nakajima (2016). “Worker flows and job flows: A quantitative investigation”. In: *Review of Economic Dynamics* 22, pp. 1–20.
- Giannoni, Marc and Karel Mertens (2019). “Outsourcing, Markups and the Labor Share”.
- Gibbons, Robert, John Roberts, et al. (2013). “Economic theories of incentives in organizations”. In: *Handbook of Organizational Economics*, pp. 56–99.
- Goldschmidt, Deborah and Johannes F. Schmieder (2017). “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure*”. In: *The Quarterly Journal of Economics* 132.3, pp. 1165–1217. DOI: [10.1093/qje/qjx008](https://doi.org/10.1093/qje/qjx008). eprint: [/oup/backfile/content_public/journal/qje/132/3/10.1093_qje_qjx008/1/qjx008.pdf](https://oup/backfile/content_public/journal/qje/132/3/10.1093_qje_qjx008/1/qjx008.pdf). URL: <http://dx.doi.org/10.1093/qje/qjx008>.
- Goodman-Bacon, Andrew (2018). “Difference-in-differences with variation in treatment timing”.
- Grobovšek, Jan (2020). “Managerial delegation, law enforcement, and aggregate productivity”. In: *The Review of Economic Studies* 87.5, pp. 2256–2289.

- Grossman, Sanford J and Oliver D Hart (1986). “The costs and benefits of ownership: A theory of vertical and lateral integration”. In: *Journal of political economy* 94.4, pp. 691–719.
- Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, and Matthieu Teachout (2020). “Vertical integration, supplier behavior, and quality upgrading among exporters”. In: *Journal of Political Economy* 128.9, pp. 000–000.
- Harrison, Bennett and Maryellen R Kelley (1993). “Outsourcing and the Search for Flexibility”. In: *Work, Employment and Society* 7.2, pp. 213–235.
- Henry, Emeric and Francisco Ruiz-Aliseda (Aug. 2016a). “Keeping Secrets: The Economics of Access Deterrence”. In: *American Economic Journal: Microeconomics* 8.3, pp. 95–118. DOI: [10.1257/mic.20140190](https://doi.org/10.1257/mic.20140190). URL: <https://www.aeaweb.org/articles?id=10.1257/mic.20140190>.
- (2016b). “Keeping secrets: the economics of access deterrence”. In: *American Economic Journal: Microeconomics* 8.3, pp. 95–118.
- Hirsch, Barry T, David A Macpherson, and Wayne G Vroman (2001). “Estimates of union density by state”. In: *Monthly Labor Review* 124.7, pp. 51–55.
- Hopenhayn, Hugo (1992). “Entry, exit, and firm dynamics in long run equilibrium”. In: *Econometrica: Journal of the Econometric Society*, pp. 1127–1150.
- Hopenhayn, Hugo and Richard Rogerson (1993). “Job turnover and policy evaluation: A general equilibrium analysis”. In: *Journal of political Economy* 101.5, pp. 915–938.
- Houseman, Susan N (2001). “Why employers use flexible staffing arrangements: Evidence from an establishment survey”. In: *ILR Review* 55.1, pp. 149–170.
- Houseman, Susan N, Arne L Kalleberg, and George A Erickcek (2003). “The role of temporary agency employment in tight labor markets”. In: *ILR Review* 57.1, pp. 105–127.

- Hsieh, Chang-Tai and Peter J Klenow (2009). “Misallocation and manufacturing TFP in China and India”. In: *The Quarterly journal of economics* 124.4, pp. 1403–1448.
- IPO (May 2018). “Government response to technical consultation on trade secrets regulations”.
- Jefferis, Jessica (2018). “The impact of restricting labor mobility on corporate investment and entrepreneurship”. In: *Available at SSRN 3040393*.
- Katz, Lawrence F and Alan B Krueger (2019). “The rise and nature of alternative work arrangements in the United States, 1995–2015”. In: *ILR Review* 72.2, pp. 382–416.
- Kendrick, Rupert (2009). *Outsourcing IT: A governance guide*. It Governance Ltd.
- Kitch, Edmund W (1980). “The law and economics of rights in valuable information”. In: *The Journal of Legal Studies* 9.4, pp. 683–723.
- Klasa, Sandy, Hernan Ortiz-Molina, Matthew Serfling, and Shweta Srinivasan (2018). “Protection of trade secrets and capital structure decisions”. In: *Journal of Financial Economics* 128.2, pp. 266–286.
- Kukharsky, Bohdan (2020). “A tale of two property rights: Knowledge, physical assets, and multinational firm boundaries”. In: *Journal of International Economics* 122, p. 103262.
- Lafontaine, Francine and Margaret Slade (2007). “Vertical integration and firm boundaries: The evidence”. In: *Journal of Economic literature* 45.3, pp. 629–685.
- Lao, Marina (1998). “Federalizing Trade Secrets Law in an Information Economy”. In: *Ohio St. LJ* 59, p. 1633.
- Lazear, Edward P (1990). “Job security provisions and employment”. In: *The Quarterly Journal of Economics* 105.3, pp. 699–726.
- Lerner, Josh (2006). “Using Litigation to Understand Trade Secrets: A Preliminary Exploration”. Working Paper, Available at SSRN.

- Liebeskind, Julia Porter (1996). “Knowledge, strategy, and the theory of the firm”. In: *Strategic management journal* 17.S2, pp. 93–107.
- Lippoldt, Douglas C. and Mark F. Schultz (2014). “Uncovering Trade Secrets - An Empirical Assessment of Economic Implications of Protection for Undisclosed Data”. In: *OECD Trade Policy Papers* 167. DOI: <https://doi.org/https://doi.org/10.1787/5jxzl5w3j3s6-en>. URL: <https://www.oecd-ilibrary.org/content/paper/5jxzl5w3j3s6-en>.
- Martinis, Lorenzo de, Francesca Gaudino, and Thomas S Respress III (2013). “Study on Trade Secrets and Confidential Business Information in the Internal Market”.
- Mukoyama, Toshihiko (2014). “The cyclicalness of job-to-job transitions and its implications for aggregate productivity”. In: *Journal of Economic Dynamics and Control* 39, pp. 1–17.
- Nunn, Nathan (2007). “Relationship-specificity, incomplete contracts, and the pattern of trade”. In: *The Quarterly Journal of Economics* 122.2, pp. 569–600.
- Ottoz, Elisabetta and Franco Cugno (2011). “Choosing the scope of trade secret law when secrets complement patents”. In: *International Review of Law and Economics* 31.4, pp. 219–227.
- Png, Ivan (2017a). “Law and Innovation: Evidence from State Trade Secrets Laws”. In: *The Review of Economics and Statistics* 99.1, pp. 167–179. DOI: [10.1162/REST_a_00532](https://doi.org/10.1162/REST_a_00532). eprint: https://doi.org/10.1162/REST_a_00532. URL: https://doi.org/10.1162/REST_a_00532.
- (2017b). “Secrecy and patents: Theory and evidence from the Uniform Trade Secrets Act”. In: *Strategy Science* 2.3, pp. 176–193.
- Pooley, James (1989). *Trade secrets: A guide to protecting proprietary business information*. American Management Association New York.

- Pooley, James (Aug. 2020). “Don’t Focus on the Fight: When it Comes to Trade Secrets, it’s the Transaction that Counts”. In: URL: <https://www.ipwatchdog.com/2020/08/30/dont-focus-fight-trade-secrets-transaction-counts/id=124700/>.
- Quinto, David W and Stuart H Singer (2012). *Trade secrets: law and practice*. Lexis Nexis.
- Rajan, Raghuram G and Luigi Zingales (2001). “The firm as a dedicated hierarchy: A theory of the origins and growth of firms”. In: *The Quarterly Journal of Economics* 116.3, pp. 805–851.
- Restuccia, Diego and Richard Rogerson (2008). “Policy distortions and aggregate productivity with heterogeneous establishments”. In: *Review of Economic dynamics* 11.4, pp. 707–720.
- Ribstein, Larry E and Bruce H Kobayashi (1996). “An economic analysis of uniform state laws”. In: *The Journal of Legal Studies* 25.1, pp. 131–199.
- Risch, Michael (2007). “Why do we have trade secrets”. In: *Marq. Intell. Prop. L. Rev.* 11, p. 1.
- Da-Rocha, José-María, Diego Restuccia, and Marina Mendes Tavares (2019). “Firing costs, misallocation, and aggregate productivity”. In: *Journal of Economic Dynamics and Control* 98, pp. 60–81.
- Rogers, Edward S (1964). “New Concepts of Unfair Competition under the Lanham Act”. In: *Trademark Rep.* 54, p. 752.
- Roth, Jonathan (2018). “Pre-test with caution: Event-study estimates after testing for parallel trends”.
- Rouwenhorst, K. Geert (1995). “Asset Pricing Implications of Equilibrium Business Cycle Models”. In: *Frontiers of Business Cycle Research*. Princeton University

- Press, pp. 294–330. ISBN: 9780691043234. URL: <http://www.jstor.org/stable/j.ctv14163jx.16>.
- Samaniego, Roberto M (2013). “Knowledge spillovers and intellectual property rights”. In: *International Journal of Industrial Organization* 31.1, pp. 50–63.
- Sandeen, Sharon K (2010). “The Evolution of Trade Secret Law and Why Courts Commit Error When They Do Not Follow the Uniform Trade Secrets Act”. In: *Hamline L. Rev.* 33, p. 493.
- Sant’Anna, Pedro HC and Jun Zhao (2020). “Doubly robust difference-in-differences estimators”. In: *Journal of Econometrics*.
- Schaller, William Lynch (2009). “Secrets of the trade: Tactical and legal considerations from the trade secret plaintiff’s perspective”. In: *Rev. Litig.* 29, p. 729.
- Schniederjans, Marc J, Ashlyn M Schniederjans, and Dara G Schniederjans (2015). *Outsourcing and insourcing in an international context*. Routledge.
- Searle, Nicola (2012). “The Criminalization of the Theft of Trade Secrets: An Analysis of the Economic Espionage Act”. In: *IP Theory* 2.2, p. 2.
- Sherwood, Robert M (1990). *Intellectual property and economic development*. Routledge.
- Shi, Liyan (2020). “The Macro Impact of Noncompete Contracts”.
- Siegel, Donald and Zvi Griliches (July 1992). “Purchased Services, Outsourcing, Computers, and Productivity in Manufacturing”. In: *Output Measurement in the Service Sectors*. NBER Chapters. National Bureau of Economic Research, Inc, pp. 429–460.
- Sun, Liyang and Sarah Abraham (2020). “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects”. In: *Journal of Econometrics*.

- Timmer, Marcel P, Mary O Mahony, Bart Van Ark, et al. (2007). “EU KLEMS growth and productivity accounts: an overview”. In: *International Productivity Monitor* 14, p. 71.
- USITC, US, International Trade Commission, et al. (2011). “China: Effects of intellectual property infringement and indigenous innovation policies on the US economy”. In: *United States International Trade Commission, Investigation 332-519*.
- Veracierto, Marcelo (2001). “Employment flows, capital mobility, and policy analysis”. In: *International Economic Review* 42.3, pp. 571–596.
- Williamson, Oliver E (1975). “Markets and hierarchies”. In: *New York* 2630.
- Zeira, Joseph (1998). “Workers, machines, and economic growth”. In: *The Quarterly Journal of Economics* 113.4, pp. 1091–1117.

1.A Proofs

Proof of Lemma 1. I will first show that if a unique \bar{z} exists, it has to satisfy $0 \leq \bar{z} < \zeta^{-1}(z)$. Second, I show the task-level production $y(i)$ is increasing in i . Last, I will show that a unique \bar{z} exists s.t. tasks $i \leq \bar{z}$ only use outsourced and tasks $i > \bar{z}$ only use hired labor in the optimal solution.

First, the manager would not assign any outsourced workers to tasks $i \geq \zeta^{-1}(z)$ because (1) outsourced workers assigned to tasks above $\zeta^{-1}(z)$ do not generate any output while their output would be strictly positive in tasks $i < \zeta^{-1}(z)$ and (2) the marginal contribution of each task’s output approaches infinity as the output in that task approaches 0.⁴² Hence, the manager would assign a positive measure of permanent workers and no outsourced workers to all tasks $i \geq \zeta^{-1}(z)$.

⁴²Because $\zeta(i)$ is strictly increasing, $\zeta^{-1}(z)$ exists, and is strictly increasing.

Second, $y(i)$ should be weakly increasing in i . Assume towards a contradiction that $y(i_1) > y(i_2)$ for $i_2 > i_1$. Let the total number of permanent and outsourced workers assigned to these tasks be $n(i_1), r(i_1)$ and $n(i_2), r(i_2)$. Then, the marginal product of an outsourced worker in these tasks would be

$$MP_r(i) = \theta Y^{\frac{\theta-\gamma}{\gamma}} y(i)^{\gamma-1} \delta$$

For $y(i_1) > y(i_2)$, the manager could increase Y by reassigning an infinitesimal measure of outsourced workers from task i_1 to i_2 . Similarly, the marginal product of a permanent worker in these tasks would be

$$MP_n(i) = \theta Y^{\frac{\theta-\gamma}{\gamma}} y(i_1)^{\gamma-1} g(i)$$

For $y(i_1) \geq y(i_2)$, the manager could increase Y by reassigning an infinitesimal measure of permanent workers from task i_1 to i_2 because $g(i)$ is strictly increasing. Hence $y(i)$ has to be weakly increasing in i .

Last, for tasks $i \leq \zeta^{-1}(z)$, assume towards a contradiction that a permanent worker is assigned to task i_1 and an outsourced worker is assigned to task $i_2 > i_1$ in the optimal solution. Let the total number of permanent and outsourced workers assigned to these tasks be $n(i_1), r(i_1)$ and $n(i_2), r(i_2)$. Then, the manager could increase its output by switching the permanent and the outsourced worker in these tasks because, the strictly increasing $g(i)$ and weakly increasing $y(i)$ imply the last inequality

$$\begin{aligned}
MP_n(i_1) + MP_r(i_2) &> MP_n(i_2) + MP_r(i_1) \Leftrightarrow \\
\theta Y^{\frac{\theta-\gamma}{\gamma}} (y(i_1)^{\gamma-1} g(i_1) + y(i_2)^{\gamma-1} \delta) &> \theta Y^{\frac{\theta-\gamma}{\gamma}} (y(i_2)^{\gamma-1} g(i_2) + y(i_1)^{\gamma-1} \delta) \Leftrightarrow \\
y(i_1)^{\gamma-1} (g(i_1) - \delta) &> y(i_2)^{\gamma-1} (g(i_2) - \delta)
\end{aligned}$$

Hence, if a permanent worker is assigned to task i_1 , no outsourced worker would be assigned to a task $i_2 > i_1$ in the optimal solution. This guarantees that a unique \bar{z} exists s.t. tasks $i \leq \bar{z}$ only use outsourced and tasks $i > \bar{z}$ only use hired labor in the optimal solution. ■

Proof of Proposition 1. I will first characterize the assignment of workers across tasks for a given \bar{z} and then characterize the optimal choice of \bar{z} . The idea is that, hired (rented) workers should be allocated across tasks $i > \bar{z}$ ($i \leq \bar{z}$) in a way to equalize marginal products across those tasks. Second, if the threshold task is interior, i.e. $\exists \bar{z} < z$, then the firm should be indifferent between using hired or rented labor for that task. If not, then the firm should strictly prefer renting to hiring at the threshold task $\exists \bar{z} = z$. First, since the productivity of outsourced workers in tasks does not depend on the identity of the task i , the CES aggregation of the tasks together with the budget constraint for rented workers imply

$$r(i) = \frac{r}{\bar{z}} \tag{1.18}$$

For permanent workers, the equalization of the marginal product across tasks requires:

$$\gamma g(i)^\gamma n(i)^{\gamma-1} = \bar{n}$$

Using $g(i) = i$ gives

$$n(i) = \left(\frac{\gamma}{\bar{n}} g(i)^\gamma \right)^{\frac{1}{1-\gamma}} \quad (1.19)$$

where \bar{n} is a constant. The budget constraint for the permanent workers gives

$$\left(\frac{\gamma}{\bar{n}} \right)^{\frac{1}{1-\gamma}} \int_{\bar{z}}^1 g(i)^{\frac{\gamma}{1-\gamma}} di = n$$

which pins down the constant term:

$$\bar{n} = \gamma \left(\frac{(1-\gamma)(1-\bar{z}^{\frac{1}{1-\gamma}})}{n} \right)^{1-\gamma} \quad (1.20)$$

(1.19) and (1.20) allow writing $n(i)$ as a function of n and \bar{z} :

$$n(i) = \frac{ni^{\frac{\gamma}{1-\gamma}}}{(1-\gamma)(1-\bar{z}^{\frac{1}{1-\gamma}})}$$

Denote with \tilde{z} the threshold task in an unconstrained (by z) allocation of workers across tasks. At task \tilde{z} , manager should be indifferent between using permanent or outsourced workers:

$$r\delta = \frac{\tilde{z}^{\frac{2-\gamma}{1-\gamma}} n}{(1-\gamma)(1-\tilde{z}^{\frac{1}{1-\gamma}})}$$

This condition does not give an analytical solution for \tilde{z} . The right-hand side is a continuous and strictly increasing function of \tilde{z} that is equal to 0 when $\tilde{z} = 0$ and is unbounded above as \tilde{z} approaches 1. The left hand side is a positive constant. Hence, there exists a unique \tilde{z} that satisfies the condition. If $\tilde{z} > z$, then $\bar{z} = z$. Otherwise, $\bar{z} = \tilde{z}$.

Using the derived formulas for $r(i)$ and $n(i)$, I can write down the total firm output as a function of n , r , and $\bar{z}(n, r)$:

$$\begin{aligned}
F(n, r) &= \left(\int_{\bar{z}}^1 \left(\frac{ni^{\frac{1}{1-\gamma}}}{(1-\gamma)(1-\bar{z}^{\frac{1}{1-\gamma}})} \right)^\gamma di + \int_0^{\bar{z}} \left(\frac{r\delta}{\bar{z}} \right)^\gamma di \right)^{\frac{\theta}{\gamma}} \\
&= \left(\underbrace{\left((1-\gamma)(1-\bar{z}^{\frac{1}{1-\gamma}}) \right)^{1-\gamma} n^\gamma}_{\alpha_n(n,r)} + \underbrace{\bar{z}^{1-\gamma} \delta^\gamma r^\gamma}_{\alpha_r(n,r)} \right)^{\frac{\theta}{\gamma}}
\end{aligned}$$

■

Proof of Corollary 1. Once the IC constraint binds, i.e., $\bar{z} = \pi$:

$$Y(n, r) = s \left(\underbrace{\left((1-\gamma)(1-\pi^{\frac{1}{1-\gamma}}) \right)^{1-\gamma} n^\gamma}_{\alpha_n} + \underbrace{\delta^\gamma r^\gamma}_{\alpha_r} \right)^{\frac{\theta}{\gamma}}$$

Defining $A = \alpha_n + \alpha_r$ and $\alpha = \alpha_n/A$ allows rewriting this in the classical CES form:

$$Y(n, r) = sA(\pi, \delta) \left(\alpha(\pi, \delta) n^\gamma + (1 - \alpha(\pi, \delta)) r^\gamma \right)^{\frac{\theta}{\gamma}}$$

■

1.B Data Sources

In this section, I describe the data sources and sample construction procedures.

1.B.1 Measures of Labor Outsourcing

I conduct analyses with data from different time periods and geographical levels, hence the best available data changes according to the question at hand. Throughout the paper, I use employment data that uses NAICS, SIC, and 1990 Census classifications and outsourcing expenditures data from Census of Manufactures (CMF). I carefully

designate which industries in NAICS classification provide labor outsourcing services. Then, for other classifications, I choose the industries that correspond the best to the designated NAICS industries.

Definition of Labor Outsourcing

I define labor outsourcing as the purchase of business services that are labor intensive and can potentially be done in-house. First, I restrict attention to business services, because the main decision (hire vs outsource) I analyze in this paper is not relevant for households. I operationalize this criterion by restricting attention to 4-digit NAICS services industries who earn more than 70% of their revenues from serving businesses and government according to the 2017 Services Annual Survey (SAS). Second, I restrict attention to labor intensive services because the decision to outsource capital-intensive services may rely on financial concerns that I abstract from in this paper. I operationalize this criterion by restricting attention to services industries who have less than 5% of their expenditures as depreciation in the 2017 Services Annual Survey (SAS) conducted by the U.S. Census Bureau. Last, I restrict attention to purchase of services where there is a meaningful make or buy decision. I use this criterion intuitively, and rule out the information technology (IT) industry (NAICS 51)⁴³, finance providing industries (NAICS 52, 53) and central offices of holding companies (NAICS 55).

This definition roughly translates to two 2-digit industries: NAICS 54 (The Professional, Scientific, and Technical Services) and NAICS 56 (The Administrative and Support and Waste Management and Remediation Services) with the following exceptions. I exclude 4-digit subsectors 5419 (Other Professional, Scientific, and Technical Services, roughly employs 8% of the total employment in NAICS54, consists mainly of

⁴³The portion of the IT sector that provides personalized services to each client firm will still be in my sample as NAICS 5415 Computer Systems Design and Related Services.

veterinary and photographic services) and 5615 (Travel Arrangement and Reservation Services, roughly employs 3% of the total employment in NAICS56) because 46% and 68% of their revenues come from households respectively. I also exclude the 3-digit subsector 562 (Waste Management and Remediation Services, roughly employs 5% of the total employment in NAICS56) because depreciation roughly corresponds to 10% of its expenses.

Table 1.6: Labor Outsourcing Sector in NAICS Classification Notes: Employment (1000s) figures are from the 2018 Current Employment Statistics. Total revenues (\$B) and the ratio of depreciation expenditures to total expenditures is from are from the 2017 Services Annual Survey (SAS). The share of revenues from households are from the 2019 Q3 Quarterly Services Survey (QSS). The fraction of employment with Bachelor’s degree (or more) is from 2019 IPUMS CPS. The SAS and QSS do not have full breakdowns by 4-digit sectors of NAICS 561, the last row provides the aggregate values.

Industry	NAICS	Emp.	Rev.	HH Share	Deprec.	College
Scientific R&D	5417	710	166	0.05	0.04	0.79
Comput. Sys. Design and Rel.	5415	2,154	304	0.00	0.03	0.73
Manag., Sci., and Tech. Consult.	5416	1,501	210	0.06	0.02	0.72
Advertising and Related	5418	493	72	0.07	0.04	0.70
Legal	5411	1,142	203	0.29	0.01	0.69
Architect., Eng., and Rel.	5413	1,493	253	0.03	0.02	0.67
Specialized Design	5414	142	15	0.30	0.02	0.64
Account., Tax, Book., Payroll	5412	1,009	136	0.15	0.02	0.61
Office Admin.	5611	517				0.38
Facilities Support	5612	160				0.38
Other Support	5619	331				0.38
Employment	5613	3,669				0.31
Business Support	5614	890				0.26
Investigation and Security	5616	951				0.19
Serv. to Buildings	5617	2,158				0.09
Admin. and Support	561		632	0.15	0.03	

Table 1.6 presents the list of 4-digit NAICS industries that fall into my definition of labor outsourcing sectors, ordered according to the share of employment with a Bachelor’s degree. The total employment in these industries is around 17 million workers, where the employment shares of NAICS 54 and 56 are almost equal with 8.5 million workers each.

Overview of Data Availability on Labor Outsourcing

I use data on both the demand for outsourcing and the supply of outsourcing. Unfortunately, historical data on demand for outsourcing has many problems. The U.S. Census first started collecting establishment-level data on outsourcing use in 1977 with Annual Survey of Manufactures (ASM) and Census of Manufactures (CMF), but restricted attention to purchase of capital-intensive services: repair and communication services.⁴⁴ Furthermore, the treatment of transactions with the establishments' Central Administrative Offices (CAO) or other auxiliary establishments of the same firm has changed in 1997. Within SIC classification, these auxiliary establishments were classified according to the primary activity of the establishment they are serving. On the other hand, NAICS classifies these establishments according to their own activity, thus these transactions show up as purchased services for the main establishment after 1997. See the discussions in Siegel and Griliches, 1992, Berlingieri, 2013, and Fort, Klimek, et al., 2016 for more details. The U.S. Census only started to collect relevant information on purchase of labor outsourcing services 1992 through ASM and CMF, while the measurement of expenditures on temporary workers only started in 2007.

The historical data on the supply of labor outsourcing (employment and value-added) is available through multiple sources, each with their own issues. The Bureau of Economic Analysis (BEA) publishes historical employment and output figures for some state-industry pairs based on 1987 SIC classification (SA25, SA25N, SAEMP25), but does not provide a clear separation of labor outsourcing sector from other sectors. In particular, it uses two-digit SIC industry 73 Business Services which combines labor outsourcing with many other capital intensive services such as equipment rental.

⁴⁴Siegel and Griliches, 1992 documents that even for the manufacturing sector, these services constituted only 28% of total service purchases once compared with Input-Output (I-O) tables for 1977.

County Business Patterns (CBP) collects very detailed industry level employment and number of establishment figures at the county level from the universe of employer establishments. However, (1) industry classifications change several times from its start with no clear bridge, and (2) it uses extensive censoring and imputation on employment values.⁴⁵ The decennial Census provides a large sample size together with a consistent industry definition provided by IPUMS USA, but the data frequency does not allow observing the impact of changes in laws. For historical data analysis, I rely on the March Current Population Survey (CPS) together with the historically consistent industry definition (1990 Census industry classification) provided by the IPUMS CPS. The CPS has a smaller sample size than the other data sources and suffers from small sample size in some state-industry bins, which does not necessarily create bias in diff-and-diff estimates.

Table 1.7: Labor Outsourcing Sector in Census 1990 Classification Notes: Employment figures are from the 2018 American Community Survey through IPUMS USA. The fraction of employment with Bachelor’s degree (or more) is from 2019 IPUMS CPS and the skill classification is based on how the industry compares to the U.S. average of 0.34.

Code	Subsector	Emp (1000s)	College	Skill Classification
20	Landscape and horticultural	1,731	0.10	Low-Skill
721	Advertising	672	0.70	High-Skill
722	Services to dwellings and other buildings	1,944	0.09	Low-Skill
731	Personnel supply	1,464	0.31	-
732	Computer and data processing	3,541	0.72	High-Skill
740	Detective and protective	1,051	0.19	Low-Skill
841	Legal	1,903	0.69	High-Skill
882	Engineering, architectural, and surveying	1,855	0.67	High-Skill
890	Accounting, auditing, and bookkeeping	1,397	0.61	High-Skill
891	Research, development, and testing	791	0.79	High-Skill
892	Management and public relations	2,103	0.72	High-Skill

Data Sources for the Panel Data Analysis

The Current Population Survey: I use the CPS mainly for state-industry level employment figures for labor outsourcing industries and education controls. I use the

⁴⁵See Eckert et al., [2020](#) for an ongoing project on making CBP available for historical comparisons.

Annual Social and Economic Supplement (ASEC) samples of CPS through IPUMS CPS. The IPUMS database provides an industry classification system ‘ind990’ that is based on the classification system used in 1990 Census and provides comparability over time. See Table 1.7 for the list of included industries. I also construct state-level manufacturing employment measures using Census 1990 industries with codes between 100 to 392 and total employment measures using employment status variable being at work (empstat=10). Lastly, IPUMS censors state-industry level employment estimates when the data quality is too low, hence the final sample becomes an unbalanced panel ranging from 1970 to 2019. I construct the state and industry level educational attainment measures from the ASEC samples, restricting attention to individuals of age 25 to 65. I use the ‘educ’ variable and classify values 71 to 100 as high school and above, and 110 and above as 4-year college and above. When necessary, I classify the industries that have educational attainment levels significantly above the U.S. average as high-skill labor outsourcing industries and those with significantly below as low-skill labor outsourcing industries.

The Trade Secret Protection Index: ”The index is constructed as a simple average of scores for three items of substantive law (i to iii), one item of civil procedure (iv), and two items of remedies (v to vi): (i) Whether a trade secret must be in continuous business use; (ii) Whether the owner must take reasonable efforts to protect the secret; (iii) Whether mere acquisition of the secret constitutes misappropriation; (iv) The limitation on the time for the owner to take legal action for misappropriation; (v) Whether an injunction is limited to eliminating the advantage from misappropriation; and (vi) The multiple of actual damages available in punitive damages. The index is the sum of the scores for each of the six items divided by six, so it is scaled between 0 and 1. For each item, a higher score represents stronger legal protection of trade secrets based on milestones including both common law (decisions

in cases that set legal precedent) and the UTSA taking effect.” (Png, 2017a). Png, 2017b extends this measure further until 2010.

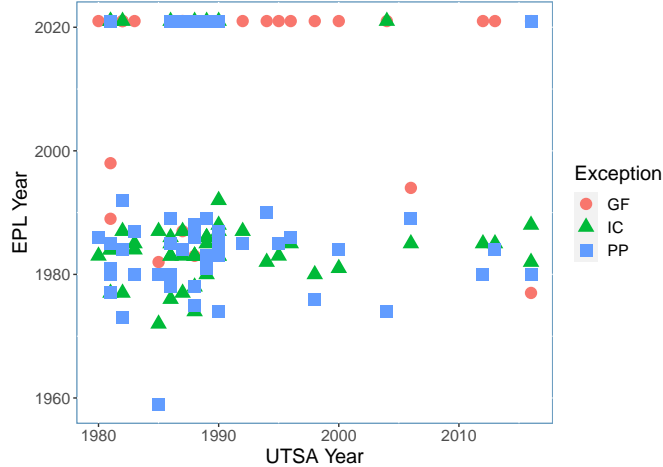
The Control Variables: I use data from the BEA to construct state level employment, population and gross domestic product (GDP) measures to serve as controls. The population measures are from the Table SA30, the employment measures are from SA25, and the inflation-adjusted GDP measures from SAGDP2S. The BEA/BLS Account covers 1987-2018 period while the BEA publishes another table for 1963-1997 period with the same industry definitions. I merge the two and compare the series in the period they coincide. The differences are very small compared to the trends I document. The decomposition results in Section 1.3.1 are broadly similar when I only use 1963-1997 or the 1987-2018 periods. I use the state-level union membership density estimates from Hirsch, Macpherson, and Vroman, 2001 who uses the CPS Outgoing Rotation Group earning files. I use the data on the state-level Wrongful Discharge Laws (WDL) from Autor, 2003 who provides public access to the data sample through his website. Figure 1.9 plots how the adoption dates of the WDL across states compare against the adoption of the UTSA.

I use the adoption data presented in Ribstein and Kobayashi, 1996 and Autor, 2003 which document the state-level adoption for 103 uniform laws and the exceptions to the at-will employment respectively to argue the UTSA adoption dates do not coincide with other laws. See also Figure 1.9.

Data Sources for the Cross-Sectional Analysis

The Census of Manufactures: The CMF collects information from the universe of manufacturing establishments as part of the Economic Census. The public data from CMF provides state and industry level data on revenues and detailed expenses, including expenses related to purchase of labor outsourcing services. I construct the

Figure 1.9: Employment Protection Laws and the UTSA The Figure has the adoption year for the Uniform Trade Secrets Act on the x-axis and for the exceptions to the at-will employment (Good Faith, Implied Contract, and Public Policy) on the y-axis. For the states that did not adopt the UTSA, the adoption year has been set to 2016 for the adoption of the DTSA. For the states that did not adopt the exceptions, the adoption year has been set to 2021.



labor outsourcing expenses by combining expenses on ‘Temporary staff and leased employee expenses’ (PCHTEMP), ‘Data processing and other purchased computer services’ (PCHADPR)⁴⁶, ‘Purchased professional and technical services’ (PCHPRTE), and ‘Advertising and promotional services’ (PCHADVT). I use the ‘Annual Payroll’ (PAYANN) as total expenses on employees on payroll, ‘Total value of shipments’ (RCPTOT) as total revenues, and ‘Value Added’ (VALADD) as value added. I use the 2007 CMF for the structural model estimation and the 2017 CMF for documenting cross-state heterogeneity in the use of labor outsourcing.

The public tables for 2007 Economic Census have state-industry level estimates for payroll, revenues, and value added but outsourcing expenses are only tabulated separately at the state and industry level. The identification only requires the state and industry level aggregates for identification. However, the two-stage estimation method I use requires state-industry level estimates for outsourcing, even though the

⁴⁶This expense does not include ‘Expensed computer hardware and other equipment’ and ‘Expensed purchases of software’, hence only documents the purchase of IT services. See Appendix 1.B for how I define labor outsourcing.

extra information is not used to identify the parameters. I construct synthetic state-industry estimates that are consistent with the state and industry level estimates and use these in the first-stage estimation⁴⁷

The Statistics of U.S. Businesses: The SUSB uses data from the universe of employer establishments and publishes statistics on establishment size distributions. I use it to construct and estimate the fraction of establishments with fewer than 20 employees and the average establishment size in each state-industry pair. To estimate the average establishment size, I compute a weighted average of average establishment sizes in each bin by weighting the bins by the listed number of establishments.

The Business Dynamics Statistics: The BDS is created from the Longitudinal Business Database and provides information on the universe of the U.S. establishments. Unfortunately, the state-level data the BDS provides is only available at the level of major industry sector. Hence, I use the BDS information to discipline state-level parameters only. In particular, I construct establishment-level job destruction and exit rates for the manufacturing sector in each state. I also use the exit rate of establishments with more than 250 employees to discipline the exogenous exit rate parameter.

The job destruction rate is very widely used as an estimator for the total separations subject to a firing cost (Boedo and Mukoyama, 2012, Decker et al., 2020), due to its standard definition and widespread availability. Yet, it is subject to two sources of bias, which act in opposite directions. First, it is subject to a time aggregation bias: because it is based on measures of establishments at certain points in time, it doesn't account for the separations in the middle that were replaced with a hire before the

⁴⁷The 2017 tables do report estimates for outsourcing expenses at the state-industry level. I use the same synthetic construction for 2017 as if only the state and industry level estimates are observed. The correlation between the actual and the synthetic estimates is 0.6. Considering the frequent censoring applied at the state-industry level, the synthetic data should closely follow the actual data.

next observation. Hence it underestimates the number of total separations. The bias becomes larger as the frequency of observations gets lower. Second, it overestimates the separations that are subject to a regulatory firing cost, as some job destruction is due to voluntary quits or retirement instead of layoffs⁴⁸.

I use the Job Openings and Labor Turnover Survey (JOLTS) by BLS to get a rough estimate of the direction and the size of the total bias. JOLTS provides estimates for the total count of separations in a time period, hence it is not subject to the time-aggregation bias. Furthermore, it distinguishes the separations as quits and layoffs. The (nationwide) approximate yearly rate of quits equals 14.5% relative to the job destruction rate of 11.4% for the manufacturing sector in 2007 (JOLTS doesn't publish state-level estimates). Since the discrepancy is not very large, I follow the literature and use the job destruction rate as the primary moment to target.

Data Conversions

The Elasticity of Substitution: I use the estimates from Chan, 2017 as elasticity of substitution parameters (between permanent and outsourced workers) in the structural model. Chan, 2017 groups 3-digit manufacturing industries in the second revision of The Statistical Classification of Economic Activities in the European Community (NACE) industry classification into four broad manufacturing industry groups: Food Products, Wood and Paper Products, Heavy Industry and Extraction, and Tools, Machinery and Consumer Goods. I match the NACE 2-digit sectors to 2007 NAICS 3-digit sectors using the official correspondence table from the Eurostat.⁴⁹ I leave NAICS industries out of my analysis if they do not strongly match to one of the 2-digit NACE industries. Table 1.8 lists both the NACE and NAICS

⁴⁸See Mukoyama, 2014 for a more detailed description of the first bias and Fujita and Nakajima, 2016 for the second bias.

⁴⁹See https://ec.europa.eu/eurostat/ramon/miscellaneous/index.cfm?TargetUrl=DSP_NACE_2_US_NAICS_2007.

industries included in this classification.

Food	Wood	Heavy	Machinery	
10	2	6	25	
11	16	9	26	
12	17	19	27	
		20	28	
		21	29	
		22	30	
		23	31	
		24	32	
Food	Wood	Heavy	Machinery	Left Out
311	321	324	332	313
312	322	325	333	314
		326	334	315
		327	335	316
		331	336	323
			337	339

Table 1.8: Manufacturing Industry Groups (Chan, 2017) for 2-digit NACE and 3-digit NAICS Classifications

The TFP Process: I use the estimates from Bloom, Floetotto, et al., 2018 to discipline the industry-level estimates of the variance of the productivity process. It is impossible to reach at the variance estimates at the group level without the micro-data, so I equate variance of the group equal to the weighted average of the variances. Since the average level of the TFP/demand shock is not identified in my model, I

only need the relative variances of different industries. In addition, since I model the TFP/demand as a log-normal process, errors in the parametrization of the variance process are partially corrected through the estimation of the persistence parameters. Bloom, Floetotto, et al., 2018 provides the estimates with the 4-digit 1987 SIC classification. Using the conversion table by Eckert et al., 2020, I first construct weights to compute variance estimates at the NAICS level and take a weighted average to get group level variance estimates.

Data Sources for the Cross-Country Analysis

The EU KLEMS Accounts: The EU KLEMS Growth and Productivity Accounts aims to provide data on industry level employment, output, and productivity estimates. The accounts include several updates that extend the coverage of countries, include more detailed industries, and make changes and corrections to the previous releases. I use the March 2008 release (Timmer, O Mahony, Van Ark, et al., 2007) which has a smaller coverage of countries relative to more recent releases, but goes back as early as 1970. In particular, I use the ‘Number of Persons Engaged’ (EMP) variable and use industry code 74 (Other business activities) as labor outsourcing. Although this industry code is not as precise as the definitions I have used with the Census and NAICS classifications, the implied labor outsourcing share is remarkable similar to the one I have derived for the U.S. through the 1990 Census classification.

The OECD Structural Analysis Database: The STAN collects and estimates data on industry level input and output from the countries’ own national accounts, using a harmonized industry definition in the process. I use the industry codes M-N (Professional, scientific and technical activities; administrative and support service activities) as labor outsourcing, which roughly corresponds to NAICS 54 and 56 but also includes equipment rental and leasing activities.

The OECD Employment Protection Index: The OECD have information on several types of employment protection, “...compiled using the Secretariat’s own reading of statutory laws, collective bargaining agreements and case law as well as contributions from officials from OECD member countries and advice from country experts.” The index has four versions that improves the method and increases the scope of the previous one. I restrict attention to the first version because it provides the longest panel of data. I use the strictness of employment protection (individual and collective dismissals) as a measure of firing cost consistent with the cross-state analysis I do in the main text. The index ranges from 0 to 5 from the weakest to strongest protection and is available yearly from 1985 to 2019.


The OECD Trade Secret Protection Index: I use two cross-country measures of trade-secret protection. The first one is an index constructed by Lippoldt and Schultz, [2014](#) for the OECD, which combines information on whether 26 criteria were satisfied in the trade secret law of 37 between 1985 and 2010. It ranges from 0 to 5 from the weakest to strongest protection. The index is only available for years ending in 0 and 5.

The Global IP Trade Secret Protection Index: The second index I use is constructed by the Global Innovation Policy Center of the U.S. Chamber of Commerce. It ranges from 0 to 3 from the weakest to strongest protection. Its country coverage is much larger than the OECD index with 50 countries but it only goes back as far as 2012.

1.C Cross-Country Evidence

In this section, I analyze the cross-country patterns of labor outsourcing and trade secret laws and discuss four more facts on (1) the growth of outsourcing, (2) the cross-

Figure 1.10: The Employment Share of the Labor Outsourcing Sector in 1970 and 2005 The total height of the bar denotes the size of the employment share of the labor outsourcing sector in 2005 while the shaded height denotes the share in 1970. The employment data is from the 2008 Revision of the EU KLEMS Accounts. I define the labor outsourcing sector as the industry code 74 (Other business activities). See Appendix 1.B for details.

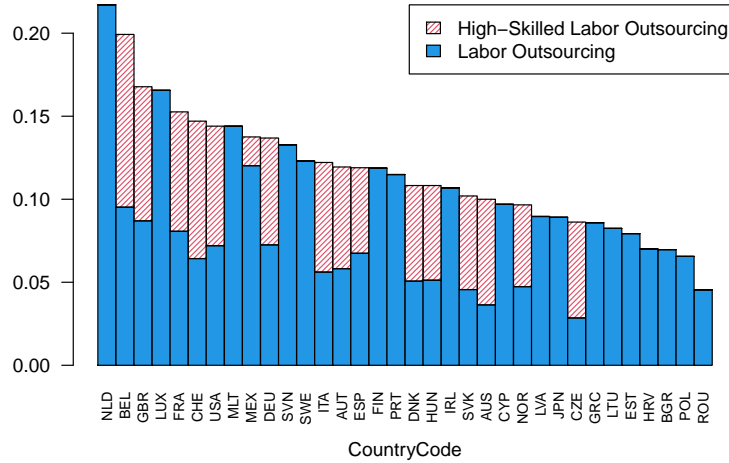


country heterogeneity in outsourcing, (3) the cross-country heterogeneity in trade secret protection and (4) how these patterns relate to the trade secret laws. I restrict attention to the analysis of the supply of labor outsourcing through employment data, because there is no available data for the demand side that allows cross-country comparisons. Hence, the scope of my analysis is determined by the availability of industry level employment data that allows cross-country comparisons.

Fact 3: The employment share of the labor outsourcing sector has grown globally since the 1970s.

The large growth in the employment share of the labor outsourcing sector was not specific to the U.S. I use the EU KLEMS Accounts (2008 Rev.) to construct measures of employment in labor outsourcing sectors for 14 countries in 1970 and 2005. Figure 1.10 presents how the employment share of the labor outsourcing sector has changed from 1970 to 2005. The sector has grown dramatically across all the countries in my sample and the growth in the U.S. is not an anomaly.

Figure 1.11: The Employment Share of the Labor Outsourcing Sector in 2017 The total height of the bar denotes the size of the employment share of the labor outsourcing sector. I depict the share of the high-skill labor outsourcing sector with the shaded height of the bar for countries where the data is separately available. The employment data is from the 2017 OECD STAN Accounts. I define the labor outsourcing sector as the industry codes M-N (Professional, scientific and technical activities; administrative and support service activities). See Appendix 1.B for details.



Fact 4: There is a large cross-country heterogeneity in the intensity of labor outsourcing.

The employment share of the labor outsourcing sector differs significantly across countries, similar to the heterogeneity present across the states of the U.S. I use the Organisation for Economic Co-operation and Development (OECD) STAN Accounts to construct measures of employment in labor outsourcing sectors for 34 countries in 2017. Figure 1.11 presents how the employment share of the labor outsourcing sector differs across countries. The employment share for the country in the 90th percentile (France, 15%) is twice of the country in the 10th percentile (Croatia, 7%) .

Fact 5: There is large variation in trade secret protection globally.

There have been many developments in the protection of trade secrets globally since 1970. The World Trade Organisation proposed the TRIPS Agreement (Agreement on Trade-Related Aspects of Intellectual Property Rights) in 1994. The Article 39 of the TRIPS Agreement is specifically dedicated to trade secrets and describes broadly what is protected under the definition. The member countries promise to enforce the protection of trade secrets, yet there is substantial heterogeneity in both the form and the enforcement of the laws across countries.

China has been at the center of trade secret violation discussions for some time (Bradsher, 2020). China provides protection for trade secrets under the Anti-Unfair Competition Law (AUCL) which was enacted as early as 1993, and amended in 2017 and 2019. Yet, foreign firms operating in China frequently complain about the lack of enforcement. The U.S. International Trade Commission conducted a survey of firms (USITC, Commission, et al., 2011) that are in IP-intensive sectors and are “particularly susceptible to IPR (intellectual property rights) violations in China.” According to their report, “Firms that provided quantitative responses estimated that improved IPR protection and enforcement in China could result in as much as a 10–20 percent increase in sales, royalties, and license fees earned in China, and a 2–5 percent increase in employment in their U.S. operations. These employment gains could translate into approximately 922,588 new U.S. jobs among IP-intensive firms.” More importantly, even though firms were suffering from trade secret theft, “Only 0.6 percent of those firms that reported material losses due to trade secret misappropriation during 2007–09 stated that they had pursued any trade secret misappropriation proceedings in China.”

Sherwood, 1990 reports the results of a survey on 1800 Brazilian firms in 1989. In the survey, although half of the firms have had ‘trade secret losses’, in 86% of those

cases, there was no attempt for a legal procedure. The firms reported as the main reasons they did not take legal action were “...lack of sufficient proof, a gap in the law on which to base a legal action, or the expectation that litigation would be too expensive or that enforcement would be poor even if the case were won.”


The European Union has enacted the Directive on the Protection of Trade Secrets (EUTSD) in 2016 after a lengthy process of drafting and consultations “to harmonise the existing diverging national laws [within the EU] on the protection against the misappropriation of trade secrets, so that companies can exploit and share their trade secrets with privileged business partners across the Internal Market, turning their innovative ideas into growth and jobs”. Before 2016, even the provision that guided the trade secret protection changed across countries. A large majority used their criminal code or an unfair competition law and the only country that had a specific trade secret law was Sweden. Furthermore, the countries differed in which types of damages were granted and on what conditions injunctive reliefs were issued.⁵⁰ According to an industry survey on 537 firms in 13 countries ran by Baker and McKenzie for the EU, “40% of EU companies would refrain from sharing trade secrets with other parties because of fear of losing the confidentiality of the information through misuse or release without their authorisation” and among 110 firms who had at least one case of misappropriation “only 57 (40.7% of responses) sought remedies in EU courts”.

Fact 6: The strength of trade secret protection and the size of the labor outsourcing sector are positively correlated across countries.

In this section, I ask whether there is any evidence of a link between the protection of trade secrets and labor outsourcing decisions across countries. Since there are

⁵⁰See Figure 4, Table A9, and Table A2.2 in <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=SWD:2013:0471:FIN:EN:PDF>

Figure 1.12: The Labor Outsourcing Sector and Legal Protection The x-axis is the OECD Trade Secret Protection Index in the left panel and the OECD Employment Protection Index in the right-panel. Each box refers to one country-year observation where the boxes with darker colors refer to earlier years. The employment data is from the 2008 Revision of the EU KLEMS Accounts. I define the labor outsourcing sector as the industry code 74 (Other business activities). Both indices range from 0 to 5 with 5 being the strongest protection. See Appendix 1.B for details.



large unobserved differences across countries beyond the intellectual property law, I treat the evidence here more descriptive rather than causal. I use a panel data on the employment shares of labor outsourcing sector through 2008 EU KLEMS and the trade secret protection index constructed by Lippoldt and Schultz, 2014. The final sample has quintennial observations for 12 countries between 1985 and 2005. The left panel of Figure 1.12 presents the patterns of trade secret protection and the extent of labor outsourcing. There is overall a positive correlation, with countries improving in both dimensions (e.g. Korea) and others that do not really increase the extent of outsourcing even though the law has improved (e.g. Lithuania). I do a similar analysis using the OECD employment protection index as shown in the right panel of Figure 1.12, and no real pattern emerges having in mind the little time-series variation present in employment protection laws.


Table 1.9: Cross-Country Panel Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
TSP	0.52** (0.21)	0.68*** (0.10)	0.22*** (0.08)	0.24** (0.09)	0.16*** (0.06)	0.18*** (0.06)
log(ManufShare)				0.38 (0.34)		0.35 (0.34)
EPL					-0.22 (0.13)	-0.21 (0.13)
FE	Year	Country	Both	Both	Both	Both
Range	'85-'05	'85-'05	'85-'05	'85-'05	'85-'05	'85-'05
Nobs	49	49	49	49	49	49

Notes: The dependent variable is the log outsourcing sector share of employment. TSP refers to the OECD Trade Secret Protection index and the EPL refers to the OECD Strictness of Employment Protection index. There are country and year fixed effects. Standard errors are clustered at the state level. The employment shares of the outsourcing sector and the manufacturing sector are computed from the 2008 EU KLEMS Accounts. See Appendix 1.B for details on sample construction and included industries.

To dig deeper, I run simple panel regression, controlling for country and year fixed effects. The country fixed effects allow controlling for important country-specific variables that are important for outsourcing but does not change much over time, such as the degree of corruption and trust. The time fixed effects allow controlling

Figure 1.13: Trade Secret Protection and GDP per capita The x-axis is the Global IP Trade Secret Protection Index and y-axis is the GDP per capita. Each box refers to one country observation. The GDP per capita data is from the 2008 Revision of the EU KLEMS Accounts. Global IP Trade Secret Protection Index ranges from 0 to 3 with 3 being the strongest protection. The line is the ordinary least squares estimate together with a 95% confidence interval. See Appendix 1.B for details.



for global trends in outsourcing, for example due to increasing use of information technology. I also use the share of manufacturing employment in each country to control for country-specific structural change. Table 1.9 presents the results of the panel regressions. The trade secret protection index has a statistically significant correlation with the outsourcing shares, after controlling for country and year specific variables.

Even though the trade secret protection and the extent of outsourcing tend to evolve together across countries, my analysis here does not rely on an exogenous variation in trade secret laws. Hence, it is important not to derive causal implications from this analysis.

1.D Estimation Details

The estimation of the structural model requires solving for the distribution of firms across the number of permanent workers and idiosyncratic shocks. Since I do the

estimation for multiple industries and multiple states of the U.S., even solving for the equilibrium can quickly become infeasible. I do several tricks to decrease the computational burden. I describe these tricks in three levels: the design of the model environment, the assumptions that allow approximating the equilibrium, and the estimation algorithm.

The Design of the Model Environment

I design the model environment in a way that allows estimating each state separately. This requires each state to have separate product and labor markets. Since neither the aggregate size of the workforce nor aggregate output is identified for states in the model, these restrictions do not play a role in the estimation. In other words, one can do the estimation ignoring cross-state interactions, then appropriately weight the states according to their size to compute nation-level aggregates. However, these restrictions do play a role in the counterfactual exercises. In particular, I assume the policies do not change the extent of cross-border activities: when one state improves its trade secret law, increased productivity does not attract workers or businesses from other states. Although this assumption is restrictive, it is necessary to keep the problem feasible. Another alternative would be to allow cross-state interactions, but decrease the cross-industry and cross-state heterogeneity across firms substantially. I anticipate the bias in policy evaluations that would arise from assigning the heterogeneity from other factors to trade secret protection would be larger than the bias from ignoring cross-state interactions. I leave the formal assessment for future projects.

Approximating Assumptions

The main identification assumption, i.e. the benefits to outsourcing varies across industries but not over time, implies a parameter that is constant across states. This parameter does not preclude separately computing the equilibria for each state, but requires the estimation to be done simultaneously for all states. Estimating all states simultaneously would necessitate the estimation of 1050 parameters altogether, which is computationally infeasible. To avoid this issue, I do the estimation under Assumption 2, where the parameters for the trade secret protection (π_j) and the outsourcing efficiency δ reduce to a factor share in a CES production function. Then, I treat the estimated factor shares ($\hat{\alpha}_{jk}$) as the sum of the model implied factor shares ($\alpha(\pi_j, \delta_k)$) and a symmetric zero-mean error term. This allows separately estimating each state, collecting the factor shares, and estimating the trade secret protection parameters (π_j) in the second stage.

At the estimated parameters, the assumption does not impact the vast majority of firms and does not have a large impact on the model implications.⁵¹ I do not impose Assumption 2 when I compute the counterfactuals, i.e., firms are not forced to use more outsourced workers when the trade secret protection improves.

Estimation Algorithm

Computing the stationary equilibrium requires two computationally intensive steps: (1) computing the value function of firms for each industry, and (2) computing the equilibrium rate of entry in each industry that ensures market clearing under the

⁵¹This does not preclude the possibility that it significantly impacts the estimated parameters, i.e., imposing the assumption at the ‘correct’ parameters would impact a significant portion of the firms. A complete verification requires simulating data from the model under different parameter sets and assessing the ability of the model to estimate those parameters accurately when the assumption is imposed.

implied steady state distribution of firms. I use Value Function Iteration (VFI) for the first step and a forward iteration with an exact transition function for the second step. It is possible to compute the equilibrium under a second with 200 grids points for permanent workers and 10 grid points for the idiosyncratic shock process with the classical algorithm by Hopenhayn and Rogerson, 1993.⁵² My model has two added levels of complexity on top of the classical version. First, due to the non-convex adjustment cost for permanent workers together with the task-based production function, the choice of outsourced workers requires the use of a non-linear solver for each choice of the number of permanent workers. Second, my model requires computing K (number of industries) prices, stationary distributions, and entry rates and the computation time does not scale linearly in K . I estimate the model efficiently without adding an extra layer of approximation. The classical algorithm (for one industry) prescribes

1. Use the free entry condition to determine the price of output
2. Find the mass of entrants that clears the labor market in the stationary distribution

When there are K industries that source from the same labor market, I need additional conditions to pin down the relative sizes of each industry. The final good industry provides K intermediate good demand conditions on top of the labor market clearing condition that help pin down the final good price and the K entry masses for each industry. Normally, for each guess of the parameters, solving the equilibrium requires simultaneously finding K prices that satisfy K free entry conditions, where each guess for the price requires running the VFI again to find the implied value of

⁵²I utilize the monotonicity and the concavity of the policy function in the stock of permanent workers, and the Howard's improvement algorithm. All three generate significant gains in computation speed.

entry. I use two tricks to ensure that I only need to run the VFI once for each industry for each guess of the parameters.

First, instead of finding the equilibrium intermediate good price for a given entry cost parameter, I treat the price as the parameter and the entry cost as the equilibrium object in the estimation. Hence, I only need to evaluate the VFI once for the given price, and the associated value of entry gives the ‘equilibrium’ entry cost. This uses the fact that the demand shares for the intermediate goods, intermediate good prices, and the level of productivity/demand shocks across industries are not separately identified. Hence, I can assume any K product prices, compute the associated entry cost, and set the demand shares to equate the relative size of each industry to data.

Second, although I model the entry cost and the fixed cost in the units of the final good, I measure them in units of the market wage which I normalize to unity. Hence, each firm’s value function only requires knowing the intermediate good price of its own industry and not the prices of the other intermediate goods. This allows computing the intermediate good prices separately. This trick uses the fact that the full equilibrium does not need to be computed for the estimation. When I compute the counterfactuals, I revert to measuring these costs in units of the final good price, hence computing the full equilibrium.

To sum up, for each set of (remaining) parameters, I use $K - 1$ relative industry sizes from the data, $K - 1$ conditions that ensure that the industry sizes are consistent with the equilibrium, K free entry conditions, one labor market condition and one aggregate entry rate to pin down K entry costs, K masses for entrants, K intermediate good price.⁵³ The gains in speed come from using the parameters to ensure equilibrium conditions while using the equilibrium objects to match moments.

⁵³The use of the entry rate to pin down the price level happens over the whole estimation, rather than for each set of parameters.

So my algorithm is to do the following steps for each set of ‘parameters’, where the parameters have the equilibrium price level but do not have the entry costs.

1. Use the revenue shares of industries from the data to pin down the price ratios⁵⁴, hence p_k (since the price level is a parameter)
2. Use the free entry condition to pin down the associated entry costs c_k^E
3. Choose the mass of entrants for each industry m_k to ensure the equilibrium distribution of firms in each industry is consistent with the revenue shares of industries from the data and the labor market clearing conditions

These tricks significantly speed up the computation of the equilibrium moments for each set of parameters without relying on any approximation. However, they also distort how the moments respond to changes in parameters. In particular, it reduces the efficiency of gradient based solvers, because once the parameters change, the normalization also changes. Since my model already has non-convexities due to adjustment costs and exit decisions, I prefer the gains in the speed of evaluating moments over the lost gains in efficiently searching the parameter space.

1.E Trade Secret Protection

In this section, I analyze some of the legal concepts and issues that relate to trade secret protection in more detail. Section 1.E.1 discusses the problems with trade secret protection under common law, Section 1.E.2 discusses why non-disclosure agreements are not sufficient to ensure trade secret protection, and Section 1.E.3 discusses how the courts determine which state’s law should govern a trade secret dispute.

⁵⁴This step practically puts infinite weight on the revenue share moments, forcing the estimation to match revenue shares exactly. I can always run my estimation algorithm to get a very good starting point, and let the usual procedure run without imposing this condition before finalizing the estimation.

1.E.1 Trade Secret Protection under Common Law

Before 1979, protection of trade secrets was established exclusively through common law. In addition, trade secret protection varied substantially across U.S. states. This created further uncertainty: to understand the legal practice, one had to analyze a separate set of cases for each state.

This problem was further amplified when the Supreme Court has ruled that state courts cannot use decisions made by federal courts as common law in *Erie Railroad Co. v. Tompkins* 304 U.S. 64 (1938). This landmark decision led to each state relying on the decisions made by their own courts, removing the only unifying body from the picture. Edward S. Rogers, who was chairman of the board of executives of Sterling Drug Co. and a member of Lawyers' Advisory Committee of U.S. Trademark Association would later say "Soon there was built up by decisions of the Federal Court a great body of Federal Law dealing with trademarks and unfair competition. It was a great convenience to the bar because lawyers knew or could easily learn what the decisions were and there were enough of them to give a comprehensive picture. Then came Erie ... which required Federal Courts to apply the law of the State in which they sit, and there was chaos. There were 48 different sovereignties, the decisions of whose courts were the only law. The body of Federal decision which was 50 years evolving was not binding either on the State or the Federal Courts. Nobody knew what the law was. It was frequently found that there were no applicable State decisions or that the decisions in the States comprising the same circuit were not uniform." (Rogers, 1964). Justice Joseph Story explained what creates this uncertainty as early as 1837: "One great advantage, therefore, of a code, an advantage which in a practical view can scarcely be over-estimated, is that it supersedes the necessity, in ordinary cases at least, of very elaborate researches into other books; and indeed, it often supersedes in all cases, but those of rare and extraordinary occurrence, the necessity of consulting

an immense mass of learned collections and digests of 243 antecedent decisions.” (Sandeen, 2010)

To resolve these issues, the American Law Institute has published several ‘Restatements of Torts’ before 1979, which summarized the theme of the previous decisions. However, the statements had no legal binding and were necessarily vague where uncertainty was the highest.

1.E.2 Non-Disclosure Agreements

A natural solution to prevent trade secrets from reaching the competitors would be to sign a non-disclosure agreements (NDA)⁵⁵, which are common practice today in outsourcing. However, the majority of cases do not involve a spy with malicious intent who steals obvious secrets hoping not to get caught. Instead, the issue either arises from a disagreement between the parties on what is secret and what would constitute a misappropriation, or an otherwise legitimate actor who sees a loophole in the agreement and tries to make quick profits.⁵⁶ In these scenarios, the NDA is far from being sufficient to ensure protection. First, to be enforceable, an NDA should explicitly designate what pieces of information are secret, which is very hard in practice (Elzankaly, 2018). The agreements that try to make an exhaustive list tend to fail, hence, the majority define secrets as broad and vague as possible to leave room for potential litigation. Pooley, 1989 prescribes “Overnarrow definitions of your trade secrets may restrict available protection.” and

As a practical matter, many experienced consultants will require you to

⁵⁵See Footnote 47 in Martinis, Gaudino, and Respass III, 2013 for example of a standard NDA.

⁵⁶According to the analysis of trade secret cases in federal courts in 2008 by Almeling, Snyder, and Sapoznikow, 2009, of cases where plaintiff eventually lost, 61% were because the plaintiff could not validate the information was a trade secret, 30% were because plaintiff could not prove information was misappropriated and 30% were because plaintiff could not prove it took reasonable measures to protect the secret. The percentages do not necessarily add up to 100% due to multiple issue being present in some cases.

define and describe your trade secrets in some detail. After all, consultants make their living by hopping from one firm to another in the same industry. They may justifiably insist on a strict limitation of their obligations not to use what you consider to be your trade secrets.

A word to consultants: do not sign a general nondisclosure clause if you can avoid it. Remember more than one person can possess the same trade secret, discovered independently. If you have to sign, insist on a precise definition and clarify your other consulting relationships.

Second, an NDA is only enforceable on information that is not readily available elsewhere. For example, if the secret is previously presented in a public fair, or if it is not clear what portion of the secret is already known in the industry, the NDA may not be enforced. Third, enforcement of the NDA requires taking proper precautions to protect the information, where the definition of proper is purposefully vague. While verbally discussing a document which is explicitly classified to be secret, additional information the firm gives may not be protected (Pooley, 2020). Fourth, the NDA can assign damages to violations, but cannot prevent further use or the disclosure of the secret once it is revealed. Fifth, although the NDA may designate a monetary transfer in case of a violation, it is rarely enforced and the court tends to update the number according to its own estimate of the actual damages. Last, but not least, small and inexperienced companies may not be able to draft a functional NDA. The trade secret law still provides protection if there is an implied confidentiality in the agreement when the NDA is missing or invalid (Smith v. Dravo Corp., 203 F.2d 369 (7th Cir. 1953)). Since the NDA fail to ensure a common understanding in most cases, the details of the trade secret law becomes important in how well the secrets are protected.

1.E.3 Governing Law in Trade Secret Disputes

If the governing law is important for trade secret disputes, can the sides benefit from the non-uniformity of laws across the U.S. by designating their favorite choice-of-law? The answer is largely no.

In transactions where both sides operate in the same state, the laws of that state govern the trade secret disputes.⁵⁷ In multi-state transactions, the U.S. law permits the sides to put a choice-of-law clause in their contract, designating which state-law should govern the disputes over it. There is no definitive rule that determines the enforceability of these clauses, but two legal principles favor the state the client is based.

First, either the disputed action or one of the sides should have an organic connection to the state that will handle the case. Designating a ‘choice of law’ in a contract (e.g. a non-disclosure agreement) is neither necessary nor sufficient to ensure the designated state court will handle the dispute. Either side can file a lawsuit in a state court that is different from the one designated on the contract and the state court designated on the contract can reject handling the dispute if it feels there is no organic connection between the state and the dispute. The organic connection requirement also prevents the sides to use simple loopholes in the legal system: a firm that operates in Florida cannot request the laws of Delaware to be applied in disputes just because it is officially established there. On the contrary, the courts tend to reject attempts to pick a ‘favorite state law’ in disputes. Schaller, 2009 summarizes the procedure for trade secret disputes:

The choice of law can be complex in trade secret cases. There is no federal choice-of-law code that dictates the application of governing law in

⁵⁷The discussion in this section is largely based on Covey and Morris, 1983.

state law diversity cases. Instead, in diversity jurisdiction cases, absent an enforceable contractual choice of law clause, a district court must apply the choice of law rules of the state in which it sits... For trade secret purposes, the applicable law might be that of the place where the secrets were stolen, the place where the secrets were disclosed or used, the place where the economic effects of misappropriation were felt, or possibly the place where products incorporating the secrets were ultimately sold. The test employed usually focuses upon which jurisdiction has the greatest “interest” or “governmental interest” in the litigation, upon which jurisdiction has the most significant relationship to the dispute, or some combination of these rules. Other jurisdictions follow the *lex loci delicti* rule, meaning they apply the law of the place where the misappropriation actually took place. At times, however, courts seem to follow no specific standard at all...This costly, confusing and uncertain inquiry can be bypassed in some jurisdictions if an enforceable choice-of-law clause exists in a nondisclosure or similar contract between the parties. The chosen law will be honored if the contract bears some reasonable relationship to the designated jurisdiction and does not offend any public policy of the state in which the court is sitting. Thus, designating the law of plaintiff’s state of incorporation will not carry the day if plaintiff and defendant have their relationship centered elsewhere...See e.g. *Curtis 1000, Inc. v Suess*, 24 F.3d 941, 943-44 (7th Cir. 1994) (holding that the designated law of Delaware lacked sufficient connection to trade secret and non-compete dispute between plaintiff headquartered in Georgia and defendant working in Illinois).

Second, when the outsourcing firm signs multiple contracts with multiple clients with the same choice of law clause, the courts may interpret these non-disclosure

agreements as one of adhesion. In other words, the choice-of-law clause could be perceived as one dictated by the outsourcing firm to the client, resulting from inequality of bargaining power. In the case where the choice of law favors the outsourcing firm over the client, the court may not enforce the choice-of-law clause.

There is another fundamental force that steers the choice of law towards the client's state: if a dispute ends up in a court, the client will have to be physically present in the courtroom. Hence, the clients have an intrinsic motive to designate the home state as the governing law.

This is also supported in Almeling, Snyder, and Sapoznikow, [2009](#) and Almeling, Snyder, Sapoznikow, and McCollum, [2010](#) for trade secret disputes. Although their data do not include the location of the sides or the dispute, they find the applied law differed substantially in cases, indicating that there was no convergence to the law of a particular state.

1.F Generalized Differences-in-Differences Methods

In a setting with two time periods and two groups (treatment and control), the differences-in-differences (DiD) estimator gives a consistent estimate of the average treatment effect for the treated (att) under the parallel trends assumption. Furthermore, one can test the parallel trends assumption using pre-treatment trends under additional assumptions.

The staggered adoption setting allows aggregating the information from DiD comparisons across multiple pairs of units over many periods. One simple counterpart of the DiD estimator with multiple periods and staggered adoption is the Two-Way Fixed Effects (TWFE) estimator and it is widely used in empirical studies. This

estimator corresponds to a regression with both time and unit fixed effects where the main regressor is a dummy D_{it} that equals 1 if unit i is under the effect of the treatment at time t . The TWFE does not adopt the nice properties of the DiD estimator due to two reasons. First, Goodman-Bacon, 2018 and Chaisemartin and D’Haultfoeuille, 2020 have recently shown TWFE estimate does not have a clear economic interpretation when the treatment effect is heterogeneous across units. The estimate can even be outside the convex hull of the pairwise DiD estimates of individual adoptions. Second, Sun and S. Abraham, 2020 pointed out that the TWFE estimator estimates the treatment effect by comparing units whose treatment has changed to those whose treatment remained constant. Thus, the control group includes units who have recently received treatment. In the presence of dynamic treatment effects, this introduces a bias in the estimates as well as tainting the tests for pre-treatment trends.⁵⁸

My setting is likely subject to both dimensions of heterogeneity. First, the effect of the UTSA can be smaller or larger for the states who adopted it later. It can be smaller if there are treatment spillovers to the control states, e.g. through the interstate provision of these services. It can also be larger if the UTSA becomes more effective as states that already adopted it accumulate decisions based on it to be used as a reference for future decisions. Second, the adoption potentially has dynamic effects, i.e., its effect on outsourcing may depend on how much time has passed since adoption. It is reasonable to think the effect may take a few years to fully realize since (1) it takes time for the clients to understand the law changes and demand more outsourcing and (2) it takes time for the outsourcing sector to grow to meet the growing demand.

⁵⁸See Roth, 2018 for further issues with statistical tests for pre-trends, even in the classical DiD settings.

1.G Outsourcing and Trade Secrets

In this section, I provide some direct evidence on how the concerns over protecting trade secrets indeed impact the outsourcing decisions of the firms. First, I discuss the government regulations that limit the form and extent of outsourcing due to concerns over loss of trade secrets.⁵⁹ Second, I provide anecdotes from experts and practitioners that emphasize the importance of trade secrets in outsourcing relationships.

1.G.1 Government Regulations

Financial Institutions

The Federal Financial Institutions Examination Council⁶⁰ publishes the Outsourcing Technology Services Booklet that regulates whether and how financial institutions can outsource a variety of IT functions “... to help ensure financial institutions operate in a safe and sound manner.”

Health Providers

Health Insurance Portability and Accountability Act (HIPAA) regulates the use of outsourcing by health institutions through the Omnibus Rule which requires the ‘business associates’ of health providers to also comply with the HIPAA Rules (Breach Notification Rule, HIPAA Security Rule, HIPAA Privacy Rule, etc.) and holds the health provider responsible for any loss of private information that happens through the business associate.

⁵⁹I do not model regulation explicitly, but the information sharing constraint can easily be interpreted as such.

⁶⁰The FFIEC consists of five banking regulators—the Federal Reserve Board of Governors (FRB), the Federal Deposit Insurance Corporation (FDIC), the National Credit Union Administration (NCUA), the Office of the Comptroller of the Currency (OCC), and the Consumer Financial Protection Bureau (CFPB).

Governmental Agencies

The Privacy Act of 1974 regulates the extent to which governmental agencies can share information that pertains to an individual: “No agency shall disclose any record which is contained in a system of records by any means of communication to any person, or to another agency, except pursuant to a written request by, or with the prior written consent of, the individual to whom the record pertains [subject to 12 exceptions]” 5U.S.C. § 552a(b). The first of these 12 exceptions, namely “need to know within agency”, makes it easier to communicate this information within the agency relative to third-party agencies such as outsourcing firms.⁶¹ There are also supplemental clauses through other regulations, such as the Protection of Privacy and Freedom of Information chapter of Federal Acquisition Regulation. Specific governmental agencies also have additional regulations restricting the use of contractors. For example, Department of Defense Privacy Program of 2007, C1.3.1.4. requires that for any contracted job, an internal system of contractor performance review to be established and special training to be given on the privacy programs.

1.G.2 Self-regulation

I restrict attention to either self-reports of firms and managers, first-hand documentation of these practices by observers, or recommendations from experts. Some of the evidence here explicitly mention outsourcing decision, while some imply it through emphasizing the importance of the length of a relationship to build trust.

“Because consultants have many of the privileges of a regular employee, though for a shorter period of time, they must be subject to nondisclosure obligations as well.

⁶¹In certain instances, the courts allow treating the employees of contractors as the employees of the agency, e.g. *Mount v. USPS*, 79 F.3d 531, 532-34 (6th Cir. 1996), in some others they do not, e.g. *Minshew v. Donley*, 911 F. Supp. 2d 1043, 1072 (D. Nev. 2012).

Indeed, it is essential to secure such agreements from consultants: the nature of their work suggests they will work later for a competitor, or may compete with you directly. In fact, the consultant may be serving other masters at the same time as working for you. The consultant presents all the problems of the ‘peripatetic employee’ magnified several times. Therefore, you must be extremely cautious and clear in establishing and managing your relationship. ” Pooley, [1989](#)

“Limit the consultant’s access to that portion of your facilities, records, and staff that is necessary to complete the work. Closely supervise what is done. At termination of the relationship, get additional reassurances of what the consultant will do to protect the integrity of your data, including the results of this project.” Pooley, [1989](#)

“Contracting with a supplier can expose a company to the possibility that confidential information might leak, perhaps even to competitors. The risk is heightened when the out-sourced activity involves technology that is novel in some competitively significant way and when the protection for it (for example, patent laws) is weak or unclear and the innovation is easy to imitate. Interdependencies are also of concern: Spillover risks are exacerbated when the interface between the outsourced activity and other internal functions is complex, requiring a company to reveal proprietary information to ensure a good fit between the two.

To protect against dependency and spillover risks, a company can rely on detailed legal contracts with vendors. But such documents are time-consuming and expensive to negotiate, and enforcement is uncertain and costly, thus discouraging outsourcing. Instead, outsourcing is greatly facilitated by trust between the two parties, particularly when both organizations are keen on maintaining their reputations as trustworthy partners. However, given the possibility of divergent business interests, trust between independent firms is, by nature, conditional. Note too that the trustworthiness of external partners should be compared with that of internal suppliers,

which sometimes rate poorly.” Adler, [2003](#)

“Referred to by Adler (2003) as spillover risk, outsourcing firms are exposed to the possibility that confidential or critical information might leak to competitors or be used by the outsourcing firm to eventually take over the client firm’s business.” M. J. Schniederjans, A. M. Schniederjans, and D. G. Schniederjans, [2015](#)

“Much essential company information, including strategic plans, is stored in computers. Under no circumstances should such information fall into the hands of competitors. The security risks involved in outsourcing are therefore frequently cited as a reason for not contracting out one’s information services delivery; these companies prefer to keep their internal IT departments (Willcocks and Fitzgerald 1994; Klepper and Jones 1998; Miller and Anderson, 2004). The IT procurement manager of Case III explains:

Our primary processes of producing coatings, fibres, chemicals and pharmaceuticals are supported by IT, which consequently has very much added value. Contracting out activities so close to our primary processes is not desirable. The risk of production secrets falling into the hands of our competitor by way of external suppliers is far too great.” Beulen and Ribbers, [2010](#)

“Outsourcing the IT function is likely to involve the supplier processing the organisation’s data in some form. The organisation remains responsible for compliant handling of its data even if this is under the control of a supplier. Risks may arise over the confidentiality of the organisation’s data and intellectual property. For instance, there may be misuse of confidential data relating to the organisation, its employees and customers; and inadequate security measures implemented by the supplier.” Kendrick, [2009](#)

“It is not unusual, however, for confidentiality orders to require that all experts or consultants, whether testifying or not, be disclosed before they receive access to

confidential documents produced by the other side. Such provisions reflect legitimate concerns that the disclosure of trade secret information to a consultant who has other clients in the industry or who may participate in the industry in other capacities, creates the risk of competitive injury.” Quinto and Singer, [2012](#)

“The principal issue at the start of the Du Pont-Masland litigation was whether Masland was using Du Pont’s trade secrets in manufacturing artificial leather, or whether he was using methods that were common knowledge among chemists in that line of business. The district court initially denied a preliminary injunction because Masland insisted that he was not using Du Pont trade secrets. During the litigation, Masland proposed to get expert testimony to establish that the processes that Du Pont claimed as trade secrets were in fact common knowledge among chemists. Fearing that litigation would reveal their secrets to their competitors, Du Pont wanted to prevent Masland from drawing his experts from the ranks of their competitors, preferring that he serve as his own expert or that he use experts drawn from the Government or academia.” Fisk, [2000](#)

1.H Additional Figures and Tables

Table 1.10: Externally Calibrated Industry-level Parameters Notes: The σ_k values are computed from Bloom, Floetotto, et al., [2018](#) by taking weighted averages of ‘Uncert_tfp’ estimates for 4-digit SIC sectors. The γ_k values are from Table 9 in Chan, [2017](#).

Industry Group	σ_k	γ_k
Food	0.555	0.417
Wood	0.407	0.653
Heavy	0.516	0.568
Light	0.427	0.62

Table 1.11: First-stage Estimation Results (Michigan)

Parameter		Moment	Model	Data
τ	0.2	% Job Dest. Rate	9.24	8.74
c_1^E	176	% Exit Rate	6.15	6.29
c_2^E	463	Revenue Share 1	0.314	0.323
c_3^E	280	Revenue Share 2	0.0424	0.0461
c_4^E	490	Revenue Share 3	0.264	0.287
c_1^F	194	Avg Estb. Size 1	35.5	35
c_2^F	93	Avg Estb. Size 2	44.7	44
c_3^F	172	Avg Estb. Size 3	54	53.5
c_4^F	101	Avg Estb. Size 4	62.6	60.5
θ_1	0.092	Revenue/Payroll 1	12	12
θ_2	0.14	Revenue/Payroll 2	7.62	7.26
θ_3	0.14	Revenue/Payroll 3	8.14	8.07
θ_4	0.14	Revenue/Payroll 4	8.22	7.81
ρ_1	0.95	Size < 20 Share 1	0.762	0.715
ρ_2	0.97	Size < 20 Share 2	0.505	0.591
ρ_3	0.95	Size < 20 Share 3	0.46	0.454
ρ_4	0.98	Size < 20 Share 4	0.538	0.608
α_1	0.81	Outsourcing/Payroll 1	0.0868	0.0855
α_2	0.73	Outsourcing/Payroll 2	0.0604	0.0596
α_3	0.72	Outsourcing/Payroll 3	0.115	0.109
α_4	0.75	Outsourcing/Payroll 4	0.0836	0.0847

Table 1.12: State-Level Estimates for Trade Secret Protection The first-stage estimation results for α_{jk} and the associated second-stage estimation results for π_j

State	π	α_{Food}	α_{Wood}	α_{Heavy}	α_{Light}
AL	0.17	0.19	0.27	0.28	0.27
AZ	0.21	0.21	0.28	0.29	0.30
CA	0.24	0.23	0.30	0.30	0.31
CT	0.20	0.21	0.28	0.28	0.29
FL	0.24	0.22	0.30	0.31	0.31
GA	0.15	0.16	0.26	0.29	0.28
IL	0.23	0.22	0.29	0.30	0.31
IN	0.18	0.20	0.28	0.28	0.28
KS	0.21	0.21	0.29	0.30	0.30
LA	0.26	0.25	0.28	0.31	0.32
MA	0.21	0.21	0.29	0.29	0.30
MI	0.16	0.19	0.27	0.28	0.25
MO	0.13	0.17	0.25	0.25	0.25
NE	0.14	0.17	0.26	0.26	0.26
NJ	0.21	0.21	0.28	0.29	0.29
NY	0.16	0.19	0.27	0.27	0.27
NC	0.18	0.16	0.30	0.31	0.29
OH	0.16	0.19	0.27	0.27	0.27
OK	0.15	0.19	0.23	0.27	0.27
OR	0.18	0.19	0.27	0.28	0.29
PA	0.16	0.19	0.26	0.27	0.27
RI	0.17	0.18	0.28	0.28	0.28
TN	0.17	0.19	0.29	0.31	0.23
TX	0.18	0.23	0.30	0.19	0.31
UT	0.20	0.21	0.28	0.28	0.29
VA	0.18	0.18	0.28	0.29	0.29
WA	0.22	0.22	0.29	0.30	0.30
WI	0.17	0.20	0.26	0.28	0.27

Figure 1.14: The Distribution and the Coefficient of Variation for Outsourcing to Payroll Ratios Under Baseline and the Counterfactual Scenarios Notes: Base refers to the baseline, Avg Ind refers to the counterfactual with the average composition of industries, Avg τ (π) refers to counterfactual with the average level of τ (π). The last two refers to counterfactuals where multiple objects are equal to their average values across states. See Table 1.13 for state-by-state details.




Figure 1.15: First Stage Estimates and Revenue Payroll Ratios Notes: The first three panels plot the first stage estimation results for returns to scale parameters (θ_{jk}), entry costs c_{jk}^E , and fixed operating costs c_{jk} respectively. The bottom right panel has the revenue payroll ratios from the CMF for each state-industry pair. See Appendix 1.B for details on the data sources.

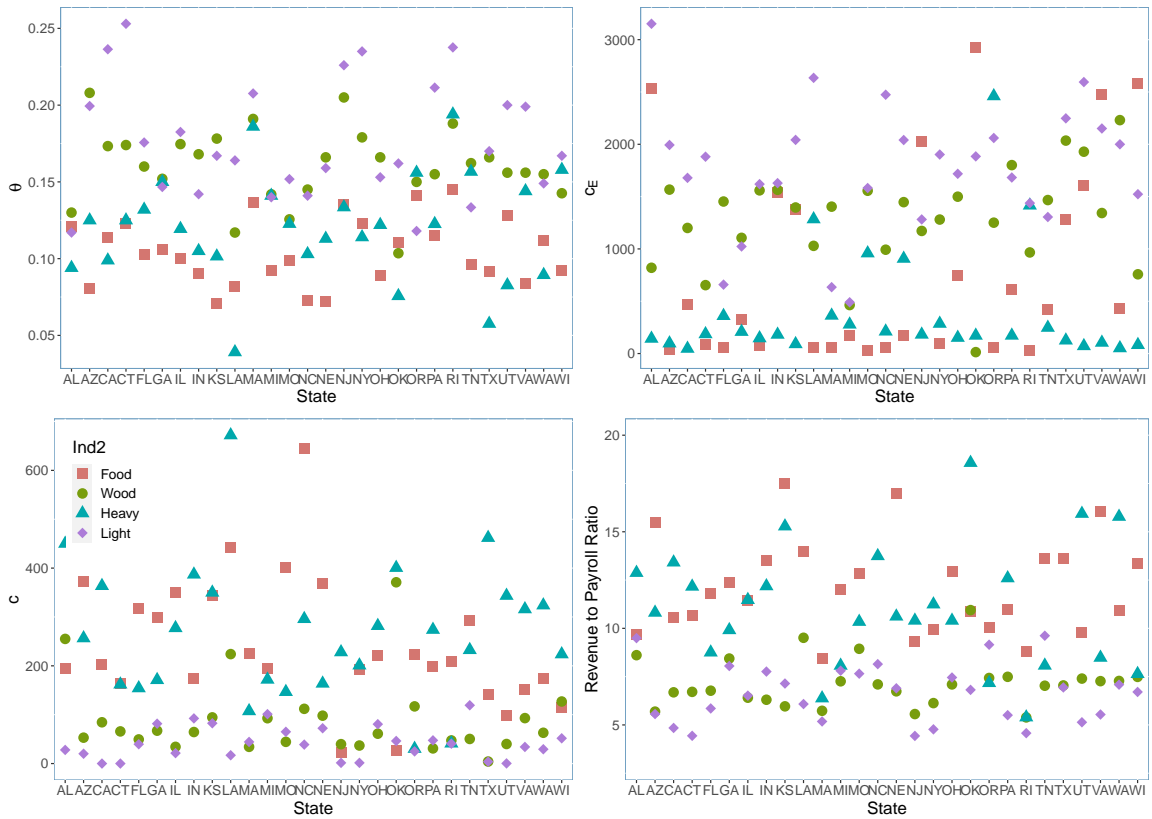


Figure 1.16: The Number of States that Adopted the UTSA (1980-2016) Notes: EEA refers to the Economic Espionage Act of 1996 and DTSA refers to the Defend Trade Secrets Act of 2016. The figures combines adoption years in Png, [2017b](#) with public announcements.




Table 1.13: The Baseline and the Counterfactual Outsourcing to Payroll Ratios for States of the U.S. The last row reports the coefficient of variation.

State	Base	Avg τ	Avg π	Avg Ind	Avg τ and π	Avg Ind, τ and π
AL	0.19	0.17	0.22	0.19	0.17	0.22
AZ	0.16	0.15	0.15	0.16	0.15	0.15
CA	0.22	0.19	0.17	0.22	0.19	0.17
CT	0.15	0.16	0.15	0.15	0.16	0.15
FL	0.17	0.22	0.19	0.17	0.22	0.19
GA	0.15	0.15	0.16	0.15	0.15	0.16
IL	0.19	0.17	0.22	0.19	0.17	0.22
IN	0.16	0.15	0.15	0.16	0.15	0.15
KS	0.22	0.19	0.17	0.22	0.19	0.17
LA	0.15	0.16	0.15	0.15	0.16	0.15
MA	0.17	0.22	0.19	0.17	0.22	0.19
MI	0.15	0.15	0.16	0.15	0.15	0.16
MO	0.19	0.17	0.22	0.19	0.17	0.22
NE	0.16	0.15	0.15	0.16	0.15	0.15
NJ	0.22	0.19	0.17	0.22	0.19	0.17
NY	0.15	0.16	0.15	0.15	0.16	0.15
NC	0.17	0.22	0.19	0.17	0.22	0.19
OH	0.15	0.15	0.16	0.15	0.15	0.16
OK	0.19	0.17	0.22	0.19	0.17	0.22
OR	0.16	0.15	0.15	0.16	0.15	0.15
PA	0.22	0.19	0.17	0.22	0.19	0.17
RI	0.15	0.16	0.15	0.15	0.16	0.15
TN	0.17	0.22	0.19	0.17	0.22	0.19
TX	0.15	0.15	0.16	0.15	0.15	0.16
UT	0.19	0.17	0.22	0.19	0.17	0.22
VA	0.16	0.15	0.15	0.16	0.15	0.15
WA	0.22	0.19	0.17	0.22	0.19	0.17
WI	0.15	0.16	0.15	0.15	0.16	0.15
CoV	0.19	0.22	0.15	0.16	0.17	0.15

Table 1.14: The State-Level Counterfactual Results After an Improvement in Trade Secret Protection The values for columns 4 to 7 are relative to a baseline value of 1.

State	Base	Best TSP	Gross Out	Net Out	# of Firms	Wage
AL	0.11	0.16	1.006	1.006	1.004	1.015
AZ	0.14	0.17	1.007	1.008	1.007	1.010
CA	0.16	0.17	1.000	1.000	0.999	1.004
CT	0.13	0.18	1.010	1.011	1.010	1.014
FL	0.13	0.14	1.003	1.003	1.003	1.002
GA	0.11	0.17	1.010	1.011	1.009	1.020
IL	0.15	0.17	1.001	1.001	1.000	1.005
IN	0.12	0.17	1.009	1.010	1.009	1.014
KS	0.13	0.15	1.004	1.004	1.003	1.008
LA	0.17	0.17	1.000	1.000	1.000	1.000
MA	0.14	0.17	1.006	1.007	1.006	1.009
MI	0.11	0.17	1.012	1.013	1.011	1.020
MO	0.09	0.18	1.017	1.019	1.016	1.029
NE	0.09	0.15	1.008	1.009	1.007	1.017
NJ	0.15	0.18	1.009	1.010	1.009	1.011
NY	0.12	0.19	1.019	1.020	1.020	1.022
NC	0.13	0.17	1.007	1.007	1.006	1.014
OH	0.12	0.18	1.013	1.015	1.013	1.022
OK	0.10	0.17	1.015	1.017	1.015	1.024
OR	0.11	0.15	1.010	1.010	1.010	1.012
PA	0.12	0.20	1.015	1.018	1.016	1.025
RI	0.12	0.19	1.017	1.019	1.017	1.021
TN	0.16	0.22	1.005	1.006	1.002	1.019
TX	0.20	0.26	1.007	1.007	1.005	1.023
UT	0.13	0.17	1.010	1.011	1.011	1.012
VA	0.11	0.16	1.008	1.009	1.007	1.014
WA	0.12	0.14	1.005	1.005	1.005	1.005
WI	0.11	0.16	1.010	1.010	1.010	1.013
Median	0.12	0.17	1.009	1.009	1.008	1.014
Max	0.20	0.26	1.019	1.020	1.020	1.029

Chapter 2

Price Informativeness and Business Cycle Misallocation

BY GORKEM BOSTANCI AND GUILLERMO ORDOÑEZ[†]

2.1 Introduction

Economic downturns are associated with a decrease in measures of aggregate productivity. Two consistent findings so far are that misallocation of inputs increases during downturns, and capital reallocation decreases.¹ These findings contradict the Schumpeterian growth theory, which suggests decreased demand makes recessions ideal for reallocation. One suspect is the counter-cyclical information quality on investment opportunities. In particular, stock prices reflect the actions of traders who spend time, money, and effort to evaluate firms' future potential and are commonly used to guide investment.² Hence, understanding the cyclical behavior of stock price infor-

[†]University of Pennsylvania.

¹See Eisfeldt and Rampini, 2006.

²Baker, Stein, and Wurgler, 2003 shows stock prices are important for the corresponding firm's investment when the firm is dependent on equity with external financing needs. Chen, Goldstein,

mativeness, i.e., the amount of information revealed by stock prices might shed some light on the capital (re)allocation puzzle.

Stock price informativeness crucially depends on (1) how informed traders are in the first place and (2) how well prices reflect their information. In this paper, we theoretically analyze how asset liquidity interacts with these two channels. We first build a model of information acquisition and trading where the noise in prices is endogenously determined. Second, we incorporate it into a neoclassical growth model where stocks are claims on real assets, and the real sector learns from the trades in the stock markets. Then, we study how the informativeness of these prices interacts with shocks to the economy.

A comprehensive analysis of price informativeness requires relaxing the commonplace assumption of exogenous noise in prices, which followed Grossman and Stiglitz, 1980. Exogenous noise prevents prices from perfectly revealing and maintains an equilibrium with incentives to acquire costly information. However, the assumptions made about the noise's exogenous characteristics also dictate how it responds to shocks. We start by building a model with different trading motives across investors. Specifically, we model two types of traders -day and night- interested in different asset properties -liquidity risk and fundamental payoff-. This structure creates endogenous noise in prices: a high price may indicate a low liquidity risk or a high fundamental payoff.

The stock trading model offers several insights by itself. First, when a larger share of traders are worried about the liquidity risk, the price informativeness on the fundamental payoff goes down. Second, traders may acquire more information about

and W. Jiang, 2006 and Bennett, R. Stulz, and Z. Wang, 2019 show that the sensitivity of the firm's investment and CEO turnover on its stock price increases as empirical measures of price informativeness increase. Edmans, Goldstein, and W. Jiang, 2012 shows a decrease in stock prices increases the likelihood that the corresponding firm will be subject to a takeover. Feldman and Schmidt, 2003 describes how regulators use stock prices in their decision making.

the firm's performance even when nothing changes about its profitability outlook. In particular, information acquisition intensifies when (1) more traders care about its stock's liquidity, (2) the quality of public information about its liquidity decreases, and (3) the quality of costly information about its liquidity increases. Third, in an otherwise symmetrical setting, the total resources spent on information acquisition are the largest when there is an equal number of day and night traders.

Next, we incorporate this stock market module inside a neoclassical growth model with heterogeneous firms. In the absence of information frictions, the investors would allocate capital across firms based on their true idiosyncratic productivity. In our model, the allocation is based on the investors' best guess, given the stock prices. While the stock traders acquire costly information and trade based on their information, the economy's real side uses the prevailing stock prices as signals of true productivity. Hence, the stock markets affect the resource allocation in the real sector through prevailing prices. On the other hand, real shocks affect the stock markets by changing the profitability of firms. This structure allows the amplification of small shocks through feedback loops between the real and financial sectors.

The main mechanism of the model works as follows. A shock that increases liquidity traders' share in the economy masks the information about fundamentals in prices. This mechanism raises two real distortions: (1) increased spending on acquiring costly information by traders and (2) worse allocation of capital across firms by the investors. The first effect allocates resources away from productive capital and towards information producing, while the second causes a misallocation of capital across firms and lead to lower aggregate investment. Hence a financial shock that increases the liquidity needs increases misallocation, decreases aggregate investment, and increases the resources spent on information production, consistent

with the patterns observed during recessions.³

The model also allows separately accounting for the efficiency losses through different channels. We look at the efficiency loss due to capital allocated to information production and the loss due to misallocation. We find the efficiency loss due to information production disappears when the share of liquidity traders is either close to 0 or 1.

Our paper lies at the intersection of the literature on price informativeness and the literature on input misallocation. The vast majority of the theoretical literature on price informativeness assumes an exogenous source of noise to prevent prices from being perfectly informative, following the impossibility theorem of Grossman and Stiglitz, 1980. We endogenize the information/noise ratio by assuming two dimensions of information condensed in a single price. The closest papers to ours here are Stein, 1987 and Vives, 2014. The former uses heterogeneity in market access while the latter uses heterogeneity in preferences to generate imperfectly informative prices without exogenous noise. Both papers are theoretical and restrict attention to implications for the stock markets. We extend their framework and allow the heterogeneity itself to change over time. The empirical literature focuses on measuring price informativeness. Dávila and Parlatore, 2018 are the first to map empirical regression estimates to parameters in a Grossman and Stiglitz, 1980 structure to infer price informativeness. They use time-series regressions to measure price informativeness for each stock, which requires them to make assumptions on how model parameters change over time to keep the cross-sectional variation flexible. We, on the other hand, use cross-sectional regressions to measure price informativeness over time. Thus, we make assumptions on the extent of heterogeneity across stocks to allow parameters to

³See H. Jiang, Habib, and Gong, 2015 and Loh and R. M. Stulz, 2018 for evidence on counter-cyclical information production by investors.

change flexibly over time. J. Bai, Philippon, and Savov, [2016](#), similar to us, analyzes the long-run trend in price informativeness using cross-sectional regressions. However, they are interested in the ability of prices to predict future stock performance, which is determined jointly by the availability of information on future prices and the ability of stock markets to communicate such information. Our model allows disentangling the two components.

One strand of the literature on input misallocation analyzes the patterns of input misallocation across firms in recessions. Foster, Grim, and Haltiwanger, [2016](#) find the extent of reallocation across the U.S. firms has declined during the great recession. Kehrig, [2015](#) finds dispersion of productivity distribution in the U.S. is larger in recessions than booms.⁴ Furthermore, Eisfeldt and Rampini, [2006](#) shows the amount of capital reallocation is procyclical and large countercyclical reallocation costs are required to justify it. Tighter financial constraints, counter-cyclical adverse selection in the market for used-capital, managers' incentives to hide reallocation needs from owners during recessions have been proposed as potential mechanisms for the counter-cyclical misallocation.⁵ On the other hand, others take increased uncertainty/misallocation as a primitive shock and analyze its effects to understand business cycles.⁶ We contribute to this literature by first proposing a novel mechanism that creates misallocation and analyzing how the misallocation caused by the drop in productivity can feedback to the price informativeness and amplify the initial shock.

⁴The increased misallocation is not specific to the U.S. and has been documented for other countries in economic crises as well. See Oberfield, [2013](#) for Chile, Sandleris and Wright, [2014](#) for Argentina, Dias, Marques, and Richmond, [2016](#) for Portugal and Di Nola, [2016](#) for the U.S.

⁵See Ordonez, [2013](#), Khan and Thomas, [2013](#), Fajgelbaum, Schaal, and Taschereau-Dumouchel, [2017](#) and Straub and Ulbricht, [2017](#) for tighter financial constraints, Fuchs, Green, and Papanikolaou, [2016](#) for adverse selection, and Eisfeldt and Rampini, [2008](#) for managerial incentives.

⁶See Christiano, Motto, and Rostagno, [2014](#), Arellano, Y. Bai, and Kehoe, [2016](#), and Bloom et al., [2018](#).

There is a recent, but growing, literature at the intersection of these two strands. Benhabib, Liu, and P. Wang, [2019](#) does a theoretical analysis similar to ours with two-way learning between the real and the financial sectors. Their model exhibits complementarity in information acquisition. Thus, a shock that reduces incentives to acquire information in one sector induces the other to reduce information acquisitions, creating equilibrium switches that amplify the initial shock. David, Hopenhayn, and Venkateswaran, [2016](#) and David and Venkateswaran, [2019](#) are the closest papers to our study. The former focuses on the role of informational frictions in resource allocation and measures how much each source of information contributes to productivity gaps. The latter has a larger scope and incorporates many potential frictions that can distort resource allocation on top of informational frictions. Both analyses provide static measures; thus, they are silent about cyclicalities. While our framework restricts attention to stock markets as the main source of information, we introduce endogenous noise, time-varying model parameters, and a two-way interaction between real and financial sectors.

The remainder of the paper is structured as follows. Section [2.2](#) introduces the stock market model with endogenous noise, and Section [2.3](#) incorporates it into an RBC model. Section [2.4](#) concludes.

2.2 Static Model

In this section, we present a simple static environment with a single risky asset. The setting is designed to be symmetric to flesh out the main mechanics of our mechanism. The price functions as a signal for two properties, which are valued differently by different traders. The trading behavior of one type of trader masks the information for the other type.

2.2.1 Environment

Preferences There is a measure one of traders with CARA period utility functions. That is, utility from consuming an amount W is given by $V(W) = -e^{-aW}$. A $\gamma \in (0, 1)$ fraction of traders live on a sunny island (sunny traders) and $1 - \gamma$ fraction live on a cloudy island (cloudy traders).

Technology There is a safe asset (money) with return R regardless of where it is consumed. There is also a risky asset (orchid) which gives a random return u_1 when planted in the sunny island and u_2 when planted in the cloudy island. Returns of this risky asset consist of two parts:

$$u_1 = \theta_1 + \epsilon_1$$

$$u_2 = \theta_2 + \epsilon_2$$

where θ_k can be privately observed at a cost $c_k(\cdot)$ while ϵ_k are unobservable. Both θ_k and ϵ_k are random variables. The cost of acquiring information $c_k(\cdot)$ is assumed to be an increasing function of the number of informed traders.⁷

Endowments Trader j is assumed to be endowed with \bar{M}_j of safe asset and \bar{X}_j of the risky asset. The total supply of risky asset in the economy is assumed to be 0 and it is common knowledge. We denote the price of the risky asset with P where the price of safe asset is normalized to 1. Trader j 's starting wealth becomes $\bar{M}_j + P\bar{X}_j$.

Distributional Assumptions We assume $\theta_k \sim \mathcal{N}(\bar{\theta}_k, \sigma_{\theta_k}^2)$ and $\epsilon_k \sim \mathcal{N}(0, \sigma_{\epsilon_k}^2)$. We also assume $\theta_1, \theta_2, \epsilon_1, \epsilon_2$ are jointly independent.

⁷This rules out any complementarity in information acquisition and prevents multiple equilibria. This assumption can be derived from a model where the cost of acquiring information is heterogeneous across traders and those with the smallest cost acquire the information first.

2.2.2 Portfolio Choice Problem

Both the sunny and cloudy traders first decide whether to pay the cost $c_k(\lambda_k)$ to get information and then decide on a portfolio of assets. After the informational decision, the portfolio choice problem of trader j from island k becomes

$$\begin{aligned} \max_{M_j, X_j} E \left[- \exp \left[- a [RM_j + u_k X_j] \right] \right] \\ \text{s.t. } M_j + P X_j = \bar{M}_j + P \bar{X}_j \end{aligned} \quad (2.1)$$

where M_j and X_j are safe and risky asset demands of trader j . Let's define $W_{0j} \equiv \bar{M}_j + P \bar{X}_j$. Since u_k is normally distributed, the end-of-period wealth of trader j is also normally distributed. Therefore, we can re-write the problem as

$$\max_{M_j, X_j} - \exp \left[- a \left[RW_{0j} + X_j [E[u_k] - RP] - \frac{a}{2} X_j^2 \text{Var}[u_k] \right] \right] \quad (2.2)$$

which yields

$$X_j^* = \frac{E[u_k] - RP}{a \text{Var}[u_k]}$$

Informed traders know the relevant θ . Therefore, they form their expectations on u_k based on θ and do not need to use the market price P . Uninformed traders look at the market price to form their expectations. For these traders, the price is not perfectly informative. The reason is, agents only care about the asset's payoff on the island they plant it but the market price is a function of payoffs in both periods. Therefore, a high price might be the outcome of a high payoff on island 1 as well as in island 2.

Demand for the risky asset for an informed trader j that lives in island k becomes

$$X_j^* = \frac{\theta_k - RP}{a \sigma_{\epsilon_k}^2}$$

Demand for the risky asset for an uninformed trader that lives in island k becomes

$$X_j^* = \frac{E[u_k|P] - RP}{aVar[u_k|P]}$$

2.2.3 Solution

Let's define λ_1 and λ_2 as the fraction of traders that pay the cost to be informed on island 1 and 2 respectively. We denote $\lambda \equiv \{\lambda_1, \lambda_2\}$. Market clearing for the risky asset requires

$$\begin{aligned} & \gamma \left[\lambda_1 \left[\frac{\theta_1 - RP}{a\sigma_{\epsilon_1}^2} \right] + (1 - \lambda_1) \left[\frac{E[u_1|P] - RP}{aVar[u_1|P]} \right] \right] \\ & + (1 - \gamma) \left[\lambda_2 \left[\frac{\theta_2 - RP}{a\sigma_{\epsilon_2}^2} \right] + (1 - \lambda_2) \left[\frac{E[u_2|P] - RP}{aVar[u_2|P]} \right] \right] = 0 \end{aligned} \quad (2.3)$$

Lemma 2. *Given the distributional assumptions, there exists a market price for a given λ with the form*

$$P_\lambda(\theta) = \alpha_{0\lambda} + \alpha_{0\lambda} \left(\theta_2 + \frac{\gamma\lambda_1\sigma_{\epsilon_2}^2}{(1-\gamma)\lambda_2\sigma_{\epsilon_1}^2} \theta_1 \right) \quad (2.4)$$

where $\alpha_{0\lambda}$ and $\alpha_{0\lambda}$ are real numbers that possibly depend on λ but not on θ_1 or θ_2 .

Proof. Appendix 2.A. ■

Corollary 2. *For a given λ , price becomes more informative about θ_1 when*

- (i) a larger fraction of traders are from island 1
- (ii) a larger fraction of island 1 traders are informed compared to island 2 traders
- (iii) the payoff at island 1 is less volatile conditional on information on θ compared to payoff at island 2.

Proof. Immediately follows from Equation 2. ■

Let's denote the end-of-period wealth for trader j from island k for a given λ as $W_j^{\lambda k}$. That is

$$W_{Ij}^{\lambda k} = R(W_{0j} - c^k(\lambda_k)) + [u^k - RP_\lambda(\theta)]X_I^j(P_\lambda(\theta), \theta) \quad (2.5)$$

$$W_{Uj}^{\lambda k} = RW_{0j} + [u^k - RP_\lambda(\theta)]X_U^j(P_\lambda(\theta), P_\lambda^*) \quad (2.6)$$

where subscripts I, U refer to being informed and uninformed and P_λ^* is the realized market price. Trader j would be willing to pay to be informed if and only if $E[V(W_{Ij}^{\lambda k})|P_\lambda] \geq E[V(W_{Uj}^{\lambda k})|P_\lambda]$.

Lemma 3. *Under the distributional assumptions,*

$$\frac{E[V(W_{Ij}^{\lambda k})|P_\lambda]}{E[V(W_{Uj}^{\lambda k})|P_\lambda]} = e^{ac_k(\lambda_k)} \sqrt{\frac{\text{Var}[u_k|\theta_k]}{\text{Var}[u_k|P_\lambda]}} \equiv \psi^k(\lambda) \quad (2.7)$$

for $k \in \{1, 2\}$ and $j \in [0, 1]$.

Proof. Appendix 2.A. ■

Corollary 3. $\psi^k(\lambda)$ is monotone in each λ_k . Therefore,

- (i) If $\psi^k(\lambda) > 1 \forall \lambda_k \in [0, 1]$, all island k traders become informed, i.e. $\lambda_k^* = 1$.
- (ii) If $\psi^k(\lambda) < 1 \forall \lambda_k \in [0, 1]$, no island k traders become informed, i.e. $\lambda_k^* = 0$.
- (iii) Otherwise, λ_k^* as a function of λ_{-k} is given by $\psi^k(\lambda) = 1$.

Definition 1. $P(\theta)^*, \lambda_1^*, \lambda_2^*$ constitutes a stochastic Rational Expectations Equilibrium (REE) such that

- (i) λ_1^*, λ_2^* are given by Corollary 3, given $P(\theta)^*$.
- (ii) $P(\theta)^*$ satisfies the market clearing condition given in Equation 2.3, given λ_1^*, λ_2^* .

Proposition 2. *Let $c(\lambda_k)$ be strictly increasing and strictly concave. There exists a unique linear REE where price has the form given in Lemma 2.*

Proof. There exists a $P(\theta)^*$ for any $\lambda_1^*, \lambda_2^* \in [0, 1]$ since market clearing condition is unbounded below and above for $P(\theta) \in R^+$. Therefore, showing $\lambda_1^*, \lambda_2^* \in [0, 1]$ exists that satisfies Corollary 3 is sufficient. Given the structure of Corollary 3, $\lambda_1^*, \lambda_2^* \in [0, 1]$ has to exist. Thus an equilibrium, which is not guaranteed to be interior, exists.

There exists a unique $P(\theta)^*$ for any $\lambda_1^*, \lambda_2^* \in [0, 1]$ since market clearing condition is strictly monotonic in price. By Lemma 3, λ_k is unique for a given λ_{-k} since $\psi^k(\lambda)$ is strictly monotonic in λ_k . Therefore, there cannot exist multiple equilibria that are not interior. For interior equilibria, $\lambda_k^*(\lambda_{-k})$ is a proper function and is strictly increasing and concave in λ_{-k} for $k = 1, 2$. Therefore, $\lambda_1^*(\lambda_2)$ and $\lambda_2^*(\lambda_1)$ can only have a single crossing inside $[0, 1] \times [0, 1]$. Thus the equilibrium is unique. ■

Corollary 4. *Let $\lim_{\lambda_k \rightarrow 0} c(\lambda_k) = 0$ and $\lim_{\lambda_k \rightarrow 1} c(\lambda_k) = \bar{C}$ where \bar{C} is large enough. Then the unique linear REE is interior.*

2.2.4 Equilibrium Characterization

Here, we focus on interior equilibria and follow Goldstein, Li, and Yang, 2014 to define price informativeness as the reduction in payoff variance faced after observing the price. However, we normalize it with the reduction in payoff variance after acquiring the costly signal.

Definition 2. *Price informativeness measure for island 1 traders is defined as*

$$PI_1 = \frac{1}{1 + \frac{\sigma_{\theta_2}^2}{\sigma_{\theta_1}^2} \left(\frac{(1-\gamma)\lambda_2\sigma_{\epsilon_1}^2}{\gamma\lambda_1\sigma_{\epsilon_2}^2} \right)^2} \quad (2.8)$$

and for island 2 traders is defined as

$$PI_2 = \frac{1}{1 + \frac{\sigma_{\theta_1}^2}{\sigma_{\theta_2}^2} \left(\frac{\gamma \lambda_1 \sigma_{\epsilon_2}^2}{(1-\gamma) \lambda_2 \sigma_{\epsilon_1}^2} \right)^2} \quad (2.9)$$

An immediate implication is whatever increases the informativeness of the price for the island 1 traders decreases it for the island 2 traders and vice versa.

Using Lemma 3, in interior equilibria, we can write price informativeness for island i trader (to be denoted as PI_i) as

$$PI_i = \sigma_{\theta_i}^2 + \sigma_{\epsilon_i}^2 (1 - e^{2ac_i(\lambda_i)}) \quad (2.10)$$

In Grossman and Stiglitz, 1980, equilibrium objects do not appear on the RHS, thus comparative statics become trivial. In our case, combining Definition 2 and (2.10) we have two equations to solve for two equilibrium objects λ_1 and λ_2 . We further assume a cost function form $c(\lambda_k) = C_k \mathcal{C}(\lambda_k)$ where C_k is a parameter.

Proposition 3. *In interior equilibria, fraction of informed island 2 traders λ_2 increases as*

- (1) a , C_1 , C_2 and $\sigma_{\epsilon_1}^2$ decreases
- (2) $\sigma_{\theta_1}^2$ and γ increases
- (3) $\sigma_{\epsilon_2}^2$ decreases and $\sigma_{\theta_2}^2$ increases.

Symmetric results hold for island 1 traders.

Proof. Appendix 2.A. ■

In Proposition 3, the comparative statics w.r.t. the parameters of the island 2 payoffs are similar to Grossman and Stiglitz, 1980. A larger fraction of island 2 traders acquire information when they are less risk averse, information acquisition

is cheaper, pre-acquisition uncertainty is higher, and post-acquisition uncertainty is lower.

What is new to our setting is the comparative statics w.r.t. to the payoffs in the other island. In particular, Corollary 2 shows the informed trades by one island masks the information about the other island in prices. Proposition 3 adds to it by describing how traders respond by changing their information acquisition. In particular, a larger fraction island 2 traders acquire information when for island 1 traders information acquisition is cheaper, pre-acquisition uncertainty is higher, and post-acquisition uncertainty is lower.

2.2.5 Comparative Statics

The model does not give closed-form solutions for all equilibrium objects, therefore, we rely on numerical solutions for comparative statics. To isolate the role of price informativeness, we focus on a fully symmetric structure, where beliefs and realizations of each island are identical. The benchmark values for parameters are given in Table 2.1. We analyze a case where both island 1 and island 2 payoffs turn out to be lower than expected.

Figure 2.1 presents comparative statics with respect to fraction of island 1 traders: γ . Here, we interpret the comparative statics as a response to a liquidity shock⁸. The equilibrium response crucially depends on the initial and final values of γ . Panel 1 shows an increase in γ increases λ_2 and decreases λ_1 as predicted by Proposition 3. Meanwhile, the fraction of all the traders that acquire information can increase or decrease based on γ_0 . In a situation where γ_0 is low, however, an increase is more likely to increase information acquisition than decrease.

⁸This can be justified in an infinitely played market game where θ realizations are i.i.d. over time, so static equilibria are played in each game. Therefore, transition dynamics can be ignored.

Figure 2.1: Comparative Statics for the Benchmark Model with respect to γ

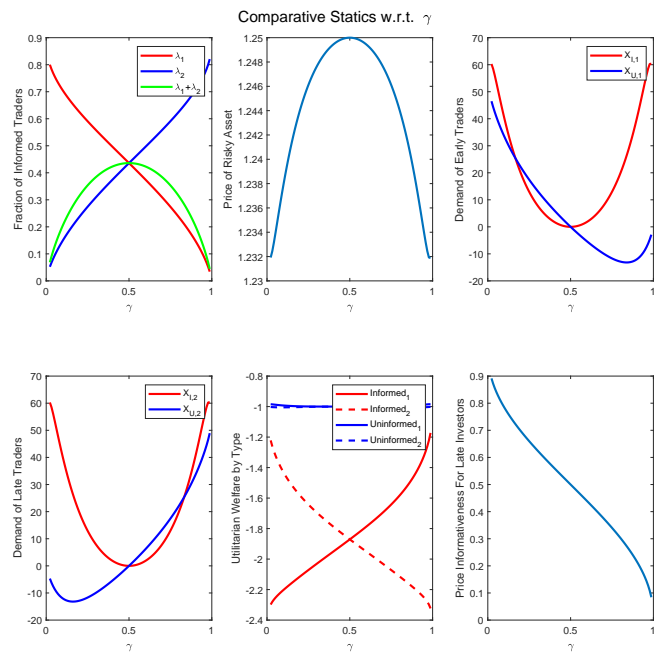


Table 2.1: Comparative Statics Benchmarks

Parameter		Value
<i>Technology and Preferences</i>		
a	Risk Aversion Parameter	0.03
γ	Fraction of Island 1 Traders	0.5
R	Risk Free Return	1
C_1	Cost of Acquiring Information for Island 1 Traders	0.5
C_2	Cost of Acquiring Information for Island 2 Traders	0.5
<i>Beliefs</i>		
$\bar{\theta}_1$	Mean of Island 1 Payoff Signal	1.35
$\sigma_{\theta_1}^2$	Variance of Island 1 Payoff Signal	0.01
$\sigma_{\epsilon_1}^2$	Quality of Island 1 Payoff Signal	0.05
$\bar{\theta}_2$	Mean of Island 2 Payoff Signal	1.35
$\sigma_{\theta_2}^2$	Variance of Island 2 Payoff Signal	0.01
$\sigma_{\epsilon_2}^2$	Quality of Island 2 Payoff Signal	0.05
<i>Realized Values</i>		
θ_1	Island 1 Payoff Signal	1.25
θ_2	Island 2 Payoff Signal	1.25

The second panel shows the price response. As shown in Table 2.1, the signal is 1.25 for both island 1 and island 2 traders. In other words, the expected payoff is equal for both islands conditional on the only piece of information available. Therefore, when price informativeness is ignored, there is no reason for the price to depend on γ . However, the price actually responds to changes in γ , since traders change how they perceive prices and their information acquisition behavior.

Panels 3 and 4 show how trader demands change. As expected, informed traders are on the buyers' side as the price is always below the signal. When γ is low, the price carries more information about θ_2 . Island 2 uninformed traders interpret the low price as low θ_2 and sell the risky asset. Island 1 uninformed traders rely mostly on their prior, since price doesn't carry much information on θ_1 , and buy the asset.

Panel 5 shows utilitarian welfare is lowest when the number of early and late traders is similar. Since trades are only transfers between agents, welfare is mainly

determined by the resources spent on acquiring information. Indeed, welfare function closely resembles the total fraction of informed traders in Panel 1. Lastly, Panel 6 confirms Corollary 2: information about θ_2 stored in prices decreases with γ ⁹.

Comparison with a Noise Trader Economy

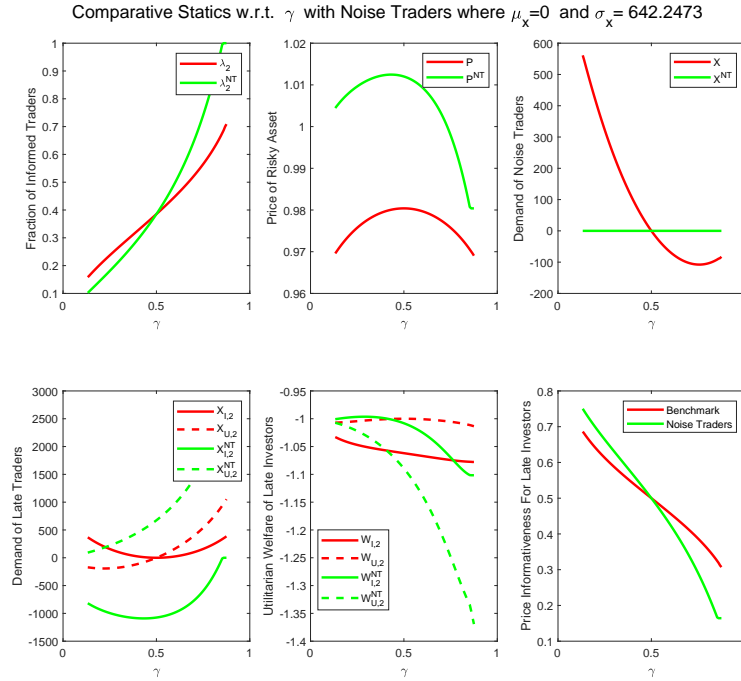
Here, we compare the properties of the model to a classical noise trader model as in Grossman and Stiglitz, 1980. Specifically, we ask how does the equilibrium price and behavior of island 2 traders would change if island 1 traders were replaced by noise traders. We equate the ex-ante distribution of noise trader demand to the ex-ante realized distribution of island 1 trader demand in the full model at $\gamma = 0.5$ as well as the realized demand. Then, changing the parameters of the model, we ask how these two models differ in their responses. Figure 2 summarizes the differences in comparative statics with γ for both models.

Multiple Assets

The model generalizes naturally to settings with multiple risky assets. CARA utility functions imply that equilibrium objects (demand, price etc.) related to the risky asset i only depend on its own payoff, not the availability of other risky assets. Therefore, the solution to the multi-asset problem is identical to the single-asset problem. This setting allows us to do comparative statics on market variables, such as portfolio returns and cross-sectional price distributions. Since these variables are directly observable, the results here are crucial for testing the model. Here, we focus on a situation where each risky asset i is characterized by a $\theta_i = \theta_{1i} = \theta_{2i}$, although prior beliefs about the risky assets are identical, realized information θ_i is uniformly distributed between 1 and 1.5.

⁹The comparative statics of the model where the signal turns out to be better than expected can be found in Figure 2.5 in Appendix.


Figure 2.2: Comparative Statics for the Benchmark Model and ‘Noise Trader’ Model with Respect to γ



2.3 Dynamic Model

In this section, we integrate a close variant of the static model in Section 2.2 into a neoclassical growth model with firm heterogeneity. We analyze how the degree of misallocation (of inputs) depends on the information production in the stock markets.

Figure 2.3: Comparative Statics for the Model with Many Risky Assets with respect to γ , where assets differ on their payoffs



2.3.1 Environment

Preferences There is a measure one of stock traders who live one period and a measure one of infinitely lived households.

At the start of each period, newborn stock traders receive a liquidity shock and a $\gamma \in (0, 1)$ fraction of them ('day traders') need to consume early while the rest ('night traders') consume after the production is finalized. Stock traders have CARA period utility functions, that is, utility from consuming an amount W is given by

$$\nu(W) = -e^{-aW}.$$

The representative household has CRRA utility function with inter-temporal elasticity of substitution parameter η and discounts the future with β .

Technology There is a measure one of firms (indexed by i) with the production function $z_{in}f(K_i)$ where z_{in} is the idiosyncratic productivity of firm i and K_i is the capital used. z_{in} are assumed to be *iid* across assets.

Endowments The firms are owned by stock traders who can freely trade the shares of the firms among themselves. The outstanding share amount of each firm is normalized to 1. The households are not allowed to hold shares¹⁰. The day traders sell their stocks to incoming traders before the production is finished to be able to consume early. Each stock loses a z_{id}/z_{in} fraction of its value when sold prematurely. Thus, z_{id} , is a measure of how much resale risk is associated with stock i . We assume z_{id} are *iid* across assets.

Since the stock traders are short lived, there is no capital trade between the stock traders and the households. Total capital holdings of the stock traders are assumed to be 0. All the capital is held by the households and rented to a hedge fund for a fixed rental rate r . The hedge fund allocates the rented capital across firms to maximize its return¹¹.

Information Both z_{id} and the z_{in} consist of two parts:

$$\begin{aligned} z_{id} &= \theta_{id} + \epsilon_{id} \quad \theta_{id} \sim \mathcal{N}(\bar{\theta}_{id}, \sigma_{\theta_{id}}^2) \quad \& \quad \epsilon_{id} \sim \mathcal{N}(0, \sigma_{\epsilon_{id}}^2) \\ z_{in} &= \theta_{in} + \epsilon_{in} \quad \theta_{in} \sim \mathcal{N}(\bar{\theta}_{in}, \sigma_{\theta_{in}}^2) \quad \& \quad \epsilon_{in} \sim \mathcal{N}(0, \sigma_{\epsilon_{in}}^2) \end{aligned} \tag{2.11}$$

¹⁰The allocation of the stocks across stock traders is irrelevant for aggregate quantities, since CARA utility functions give rise to linear policy functions.

¹¹Here, the households face no risk since hedge fund offers a deterministic interest rate for the capital. The hedge fund faces the idiosyncratic risk of each firm. Because the productivity shocks are independent, the hedge fund can perfectly pool this risk and face no aggregate risk.

where

$$\theta_{ikt} = \tilde{\theta}_{ik}(1 - \rho_k) + \rho_k z_{ik(t-1)} + v_{ikt}$$

and $v_{idt}, v_{int}, \epsilon_{idt}, \epsilon_{int}$ are assumed to be jointly **independent**. Thus

$$\begin{aligned} \bar{\theta}_{ikt} &= \tilde{\theta}_{ik}(1 - \rho_k) + z_{ik(t-1)} \\ \sigma_{\theta_{ikt}}^2 &= \sigma v_{ikt}^2 \end{aligned} \tag{2.12}$$

For each stock i , traders can pay a cost $c_i(\lambda_{ik})$ to learn the realizations for θ_{id} and θ_{in} ¹² while ϵ_{id} and ϵ_{in} cannot be learned. We denote with λ_{ik} the fraction of $k \in \{d, n\}$ traders that are informed about stock i and assume $c_i(\lambda_{ik})$ is convex in λ_{ik} .

The hedge fund doesn't have access to the information technology. Thus, similar to the traders who chose not to pay the cost, the hedge fund infers each θ_i by observing the stock market prices.

2.3.2 Recursive Competitive Equilibrium

Portfolio Choice Problem

After receiving the liquidity shock, for each stock, the traders first decide whether to pay the cost $c_k(\lambda_{ik})$ to get information and then decide on a portfolio of assets. After the informational decision, the portfolio choice problem of a night trader who is endowed with \bar{k} amount of capital and \bar{X}_i amount of stock i becomes

$$\begin{aligned} \max_{k, \{X_{in}\}_{i \in (0,1)}} & E \left[- \exp \left[- a \left[(1+r)k + \int_i X_{in} (z_{inf}(K_i) - rK_i - p_i) \right] \right] \right] \\ \text{s.t. } & k + \int_i P_i X_{in} = \bar{k} + \int_i P_i \bar{X}_i \end{aligned} \tag{2.13}$$

¹²Since the only random part in θ_{id} and θ_{in} are v_{id} and v_{in} learning *thetas* is equivalent to learning *vs.*

and the portfolio choice problem of a day trader becomes

$$\begin{aligned} & \max_{k, \{X_{id}\}_{i \in (0,1)}} E \left[- \exp \left[- a \left[(1+r)k + \int_i X_{id} ((z_{in} - z_{id})f(K_i) - rK_i - p_i) \right] \right] \right] \\ & \text{s.t. } k + \int_i P_i X_{id} = \bar{k} + \int_i P_i \bar{X}_i \end{aligned} \quad (2.14)$$

where k and $\{X_i\}_{i \in (0,1)}$ are capital and stock demands of the traders.

Given the information structure, we can write the demand for risky asset by the informed and uninformed traders as follows:

$$\begin{aligned} X_{id}^{U*} &= \frac{E[z_{in} - z_{id}|p_i]f(K_i) - rK_i - (1+r)p_i}{aV[z_{in} - z_{id}|p_i]f^2(K_i)}, \\ X_{id}^{I*} &= \frac{(\theta_{in} - \theta_{id})f(K_i) - rK_i - (1+r)p_i}{a(\sigma_{\epsilon_{in}}^2 + \sigma_{\epsilon_{id}}^2)f^2(K_i)}, \\ X_{in}^{U*} &= \frac{E[z_{in}|p_i]f(K_i) - rK_i - (1+r)p_i}{aV[z_{in}|p_i]f^2(K_i)}, \\ X_{in}^{I*} &= \frac{(\theta_{in})f(K_i) - rK_i - (1+r)p_i}{a\sigma_{\epsilon_{in}}^2 f^2(K_i)} \end{aligned} \quad (2.15)$$

Information Acquisition Problem

Each trader decides whether to acquire information about each stock i , by comparing the cost of acquiring information with the decrease in the consumption variance. Define $\psi^k(\lambda_{ik})$ as

$$\psi^k(\lambda_{ik}) \equiv e^{ac_k(\lambda_{ik})} \sqrt{\frac{\text{Var}[z_{ik}|\theta_{ik}]}{\text{Var}[z_{ik}|p_i]}} \quad (2.16)$$

for $k \in \{d, n\}$.

Corollary 5. $\psi^k(\lambda)$ is monotone in each λ_{ik} . Therefore,

- (i) If $\psi^k(\lambda_i) > 1 \forall \lambda_{ik} \in [0, 1]$, all k traders become informed, i.e. $\lambda_{ik}^* = 1$.
- (ii) If $\psi^k(\lambda_i) < 1 \forall \lambda_{ik} \in [0, 1]$, no k traders become informed, i.e. $\lambda_{ik}^* = 0$.
- (iii) Otherwise, λ_{ik}^* as a function of λ_{i-k} is given by $\psi^k(\lambda_{ik}^*) = 1$.

Proof. Follows Lemma 3. ■

Problem of the Household and the Hedge Fund

The representative household solves

$$\begin{aligned} V(K, k) &= \max_{k'} u(k(1 + r(K)) - k') + \beta V(K', k') \\ K' &= G(K) \end{aligned} \tag{2.17}$$

where $V(\cdot)$ is the value function, β is the discount factor, k is the individual capital holdings and K is the aggregate capital holdings. $G(\cdot)$ determines how the household forms expectations over the future path of the aggregate capital. The household has the classical Euler Equation:

$$u'(k(1 + r(K)) - k') = \beta u'(k'(1 + r(K')) - k'') \tag{2.18}$$

The hedge fund allocates the capital to the firms to maximize expected profits, such that

$$E[z_{in}|p]f'(K_i) = r \quad \forall i \tag{2.19}$$

Market Clearing

Let's define λ_{id} as the fraction of day traders that pay the cost to be informed and λ_{in} as the fraction of night traders that pay the cost to be informed. We denote $\lambda_i \equiv \{\lambda_{id}, \lambda_{in}\}$. Market clearing condition for stock i is

$$\gamma \left[\lambda_{id} X_{id}^I + (1 - \lambda_{id}) X_{id}^U \right] + (1 - \gamma) \left[\lambda_{in} X_{in}^I + (1 - \lambda_{in}) X_{in}^U \right] = 0 \tag{2.20}$$

The capital market clearing condition is

$$\int_i K_i = K + \bar{K} \quad (2.21)$$

where K is the total capital held by households and \bar{K} by the traders.

Equilibrium

Definition 3. $V, r, k', \{K_i, X_{id}^I, X_{id}^U, X_{in}^I, X_{in}^U, \lambda_{id}, \lambda_{in}, \phi_{i0}, \phi_{id}, \phi_{in}, p_i\}_{i \in (0,1)}$ constitute a Recursive Linear Rational Expectations Equilibrium such that

(i) k' solves (2.18).

(ii) V solves (2.17).

(iii) $p_i = \phi_{i0} + \phi_{id}\theta_{id} + \phi_{in}\theta_{in}$

(iv) ϕ_{i0}, ϕ_{id} , and ϕ_{in} solve (2.20).

(v) $X_{id}^I, X_{id}^U, X_{in}^I$, and X_{in}^U solve (2.15).

(vi) r solves (2.21).

(vii) $\lambda_{id}, \lambda_{in}$ are given by Corollary 5.

2.3.3 Discussion of the Model Ingredients

The model is similar to a neoclassical growth model with heterogeneous firms. The main difference is in the problem of the hedge fund. In an environment with readily available information, the hedge fund would allocate the capital across firms according to

$$\theta_{in} f'(K_i) = r \quad \forall i \quad (2.22)$$

in which case, the sole ‘misallocation’ would be due to ϵ_{in} , which is inevitable. In our setting, the hedge fund can only rely on the stock prices p_i to infer θ_{in} . In terms of

allocation of capital across firms, the stock liquidity is irrelevant. However, a high price could stem from a high θ_{in} or a low θ_{id} . In other words, stocks may be priced higher due to higher long term value or higher short term liquidity. Thus, compared to the benchmark in (2.22), firms with lower (higher) than expected θ_{id} shocks are allocated more (less) capital. In summary, the existence of day traders prevent prices from perfectly revealing θ_{in} .

The discrepancy between the model's allocation and the benchmark allocation in (2.22) depends on how close $E[z_{in}|p]$ is to θ_{in} , that is, how informative the prices are for inferring θ_{in} . Given that we restrict attention to linear pricing equilibria where $p_i = \phi_{i0} + \phi_{id}\theta_{id} + \phi_{in}\theta_{in}$, informativeness solely depend on ϕ_{in}/ϕ_{id} ratios of each stock. This ratio depends on two components: (1) the fraction of informed night traders to informed day traders and (2) how aggressively informed night traders trade based off their information relative to informed day traders. The first component is simply $\frac{(1-\gamma)\lambda_{in}}{\gamma\lambda_{id}}$. Given the information, the sole uncertainty faced by the traders are the ϵ_i s. Thus, the second component is $\frac{\sigma_{\epsilon_{in}}^2 + \sigma_{\epsilon_{id}}^2}{\sigma_{\epsilon_{in}}^2}$ ¹³.

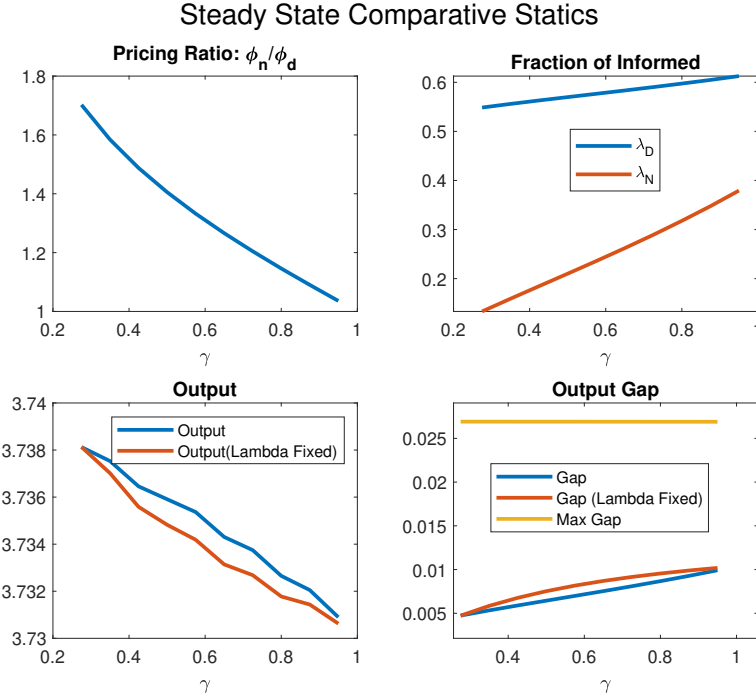
2.3.4 Liquidity Needs and Misallocation

We now analyze how the degree of misallocation changes with the fraction of day traders. Specifically, we look at (1) the ‘output gap’, which is the output difference between the perfect information case and our model of inference from prices and (2) the ‘maximum output gap’, which is the output difference between the perfect information case and a no-information case where capital is allocated across all firms equally.

Figure 2.4 presents the results of a comparative static of steady-state equilibria where we change the fraction of day traders in the economy. The pricing ratio declines

¹³See Definition 2 for more details.

Figure 2.4: Comparative Statics for the Full Model with respect to γ , where assets differ on their payoffs



as γ increases, similar to the result in Section 2.2.5. The top right panel shows that night traders respond by acquiring more information, as expected. Surprisingly, day traders also acquire more information in this parametrization. A larger fraction of day traders masks the information about the fundamental payoff which day traders also care about. In a way, increasing day trader activity can mask the information for themselves for some parameters. The input allocation depends on how much can be learned from the price regarding the fundamental payoff. Following the deteriorating information, the output gap increases with a decline in the output. Thus, the model with a larger number of day traders has larger misallocation even though agents spend more on information acquisition in steady-state.

The model structure also helps us disentangle the contribution of changing trading behavior and the changing information acquisition behavior. The bottom panels also

show what the output response would be if information acquisition did not change. Shutting this channel amplifies the output drop by not letting the traders respond to the deteriorating information.

2.4 Conclusion

Recessions are characterized by decreased investment and productivity, increased misallocation of capital, and information production activities. We suggest an increase in the concern for asset liquidity may be behind all these patterns. We build a model of stock trading with costly information acquisition where noise in prices is endogenously determined. Then we introduce this stock trading model into a neoclassical growth model where the real sector observes the stock markets to learn about investment opportunities. The model can simultaneously generate an increase in information acquisition and an increase in input misallocation by a shock to the number of traders that value asset liquidity. It also allows separating the decline in output due to costs of information acquisition and increased misallocation of capital across firms.

The model provides a mapping between the evolution of structural parameters and stock prices' behavior over time. A next step would be to use the model to quantify these channels' contribution to the observed recessions.

References

Arellano, Cristina, Yan Bai, and Patrick J Kehoe (2016). *Financial frictions and fluctuations in volatility*. Tech. rep. National Bureau of Economic Research.

- Bai, Jennie, Thomas Philippon, and Alexi Savov (2016). “Have financial markets become more informative?” In: *Journal of Financial Economics* 122.3, pp. 625–654.
- Baker, Malcolm, Jeremy C Stein, and Jeffrey Wurgler (2003). “When does the market matter? Stock prices and the investment of equity-dependent firms”. In: *The Quarterly Journal of Economics* 118.3, pp. 969–1005.
- Benhabib, Jess, Xuewen Liu, and Pengfei Wang (2019). “Financial markets, the real economy, and self-fulfilling uncertainties”. In: *The Journal of Finance* 74.3, pp. 1503–1557.
- Bennett, Benjamin, René Stulz, and Zexi Wang (2019). “Does the stock market make firms more productive?” In: *Journal of Financial Economics*.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry (2018). “Really uncertain business cycles”. In: *Econometrica* 86.3, pp. 1031–1065.
- Chen, Qi, Itay Goldstein, and Wei Jiang (2006). “Price informativeness and investment sensitivity to stock price”. In: *The Review of Financial Studies* 20.3, pp. 619–650.
- Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno (2014). “Risk shocks”. In: *American Economic Review* 104.1, pp. 27–65.
- David, Joel M, Hugo A Hopenhayn, and Venky Venkateswaran (2016). “Information, misallocation, and aggregate productivity”. In: *The Quarterly Journal of Economics* 131.2, pp. 943–1005.
- David, Joel M and Venky Venkateswaran (2019). “The sources of capital misallocation”. In: *American Economic Review* 109.7, pp. 2531–67.
- Dávila, Eduardo and Cecilia Parlatore (2018). *Identifying price informativeness*. Tech. rep. National Bureau of Economic Research.

- Di Nola, Alessandro (2016). “Capital Misallocation During the Great Recession”. Capital Misallocation During the Great Recession (August 1, 2016). Available at SSRN.
- Dias, Daniel A., Carlos Robalo Marques, and Christine Richmond (2016). “Misallocation and productivity in the lead up to the Eurozone crisis”. In: *Journal of Macroeconomics* 49, pp. 46–70. ISSN: 0164-0704. DOI: <https://doi.org/10.1016/j.jmacro.2016.04.009>. URL: <http://www.sciencedirect.com/science/article/pii/S0164070416300179>.
- Edmans, Alex, Itay Goldstein, and Wei Jiang (2012). “The real effects of financial markets: The impact of prices on takeovers”. In: *The Journal of Finance* 67.3, pp. 933–971.
- Eisfeldt, Andrea L. and Adriano A. Rampini (2006). “Capital reallocation and liquidity”. In: *Journal of Monetary Economics* 53.3, pp. 369–399. ISSN: 0304-3932. DOI: <https://doi.org/10.1016/j.jmoneco.2005.04.006>. URL: <http://www.sciencedirect.com/science/article/pii/S0304393206000298>.
- (2008). “Managerial incentives, capital reallocation, and the business cycle”. In: *Journal of Financial Economics* 87.1, pp. 177–199. ISSN: 0304-405X.
- Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel (2017). “Uncertainty traps”. In: *The Quarterly Journal of Economics* 132.4, pp. 1641–1692.
- Feldman, Ronald and Jason Schmidt (2003). “Supervisory use of market data in the federal reserve system”. In: *FRB Minneapolis, Manuscript*.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger (2016). “Reallocation in the great recession: cleansing or not?” In: *Journal of Labor Economics* 34.S1, S293–S331.
- Fuchs, William, Brett Green, and Dimitris Papanikolaou (2016). “Adverse selection, slow-moving capital, and misallocation”. In: *Journal of Financial Economics* 120.2,

- pp. 286–308. ISSN: 0304-405X. DOI: <https://doi.org/10.1016/j.jfineco.2016.01.001>. URL: <http://www.sciencedirect.com/science/article/pii/S0304405X16000027>.
- Goldstein, Itay, Yan Li, and Liyan Yang (2014). “Speculation and hedging in segmented markets”. In: *The Review of Financial Studies* 27.3, pp. 881–922.
- Grossman, Sanford J and Joseph E Stiglitz (1980). “On the impossibility of informationally efficient markets”. In: *The American economic review* 70.3, pp. 393–408.
- Jiang, Haiyan, Ahsan Habib, and Rong Gong (2015). “Business cycle and management earnings forecasts”. In: *Abacus* 51.2, pp. 279–310.
- Kehrig, Matthias (2015). “The Cyclical Nature of the Productivity Distribution”. Earlier version: US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15. Available at SSRN.
- Khan, Aubhik and Julia K Thomas (2013). “Credit shocks and aggregate fluctuations in an economy with production heterogeneity”. In: *Journal of Political Economy* 121.6, pp. 1055–1107.
- Loh, Roger K and René M Stulz (2018). “Is Sell-Side Research More Valuable in Bad Times?” In: *The Journal of Finance* 73.3, pp. 959–1013.
- Oberfield, Ezra (2013). “Productivity and misallocation during a crisis: Evidence from the Chilean crisis of 1982”. In: *Review of Economic Dynamics* 16.1. Special issue: Misallocation and Productivity, pp. 100–119. ISSN: 1094-2025. DOI: <https://doi.org/10.1016/j.red.2012.10.005>. URL: <http://www.sciencedirect.com/science/article/pii/S1094202512000634>.
- Ordonez, Guillermo (2013). “The asymmetric effects of financial frictions”. In: *Journal of Political Economy* 121.5, pp. 844–895.

- Sandleris, Guido and Mark L. J. Wright (2014). “The Costs of Financial Crises: Resource Misallocation, Productivity, and Welfare in the 2001 Argentine Crisis”. In: *The Scandinavian Journal of Economics* 116.1, pp. 87–127. DOI: [10.1111/sjoe.12050](https://doi.org/10.1111/sjoe.12050). eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/sjoe.12050>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/sjoe.12050>.
- Stein, Jeremy C (1987). “Informational externalities and welfare-reducing speculation”. In: *Journal of political economy* 95.6, pp. 1123–1145.
- Straub, Ludwig and Robert Ulbricht (2017). “Endogenous uncertainty and credit crunches”. In:
- Vives, Xavier (2014). “On the possibility of informationally efficient markets”. In: *Journal of the European Economic Association* 12.5, pp. 1200–1239.

2.A Proofs

Lemma 4. *Under the distributional assumptions,*

$$\frac{E[V(W_{I_j}^{\lambda_k})|P_\lambda]}{E[V(W_{U_j}^{\lambda_k})|P_\lambda]} = e^{ac_k(\lambda_k)} \sqrt{\frac{\text{Var}[u_k|\theta_k]}{\text{Var}[u_k|P_\lambda]}} \equiv \psi^k(\lambda) \quad (2.23)$$

for $k \in \{1, 2\}$ and $j \in [0, 1]$.

Proof. The proof adapts the corresponding proof in Grossman and Stiglitz, 1980 to this environment. Expected value of being informed for investor j that consumes in period k can be written as

$$E[V(W_{I_j}^{\lambda_k})|P_\lambda] = E[e^{-aW_{I_j}^{\lambda_k}}] = -\exp \left[-a \left[E[W_{I_j}^{\lambda_k}|\theta] - \frac{a}{2} \text{Var}[W_{I_j}^{\lambda_k}|\theta] \right] \right] \quad (2.24)$$

Using Equation 2.5, we can write

$$E[W_{I_j}^{\lambda_k}|\theta] = R(W_{0j} - c^k(\lambda_k)) + \frac{(E[u^k|\theta] - RP_\lambda)^2}{a\sigma_{\epsilon_k}^2} \quad (2.25)$$

$$Var[W_{I_j}^{\lambda_k}|\theta] = \frac{(E[u^k|\theta] - RP_\lambda)^2}{a^2\sigma_{\epsilon_k}^2} \quad (2.26)$$

since W_{0j} and P_λ are not random given θ . Thus, we can rewrite Equation 2.24 as

$$\begin{aligned} E[V(W_{I_j}^{\lambda_k})|P_\lambda] &= -\exp\left[-aR(W_{0j} - c^k(\lambda_k)) - \frac{(E[u^k|\theta] - RP_\lambda)^2}{2\sigma_{\epsilon_k}^2}\right] \\ &= -\exp\left[-aR(W_{0j} - c^k(\lambda_k))\right] E\left[\exp\left(\frac{-1}{2\sigma_{\epsilon_k}^2}(E[u^k|\theta] - RP_\lambda)^2\right)|P_\lambda\right] \end{aligned} \quad (2.27)$$

Now define

$$\begin{aligned} h_\lambda^k &:= Var(E[u^k|\theta]|P_\lambda) \\ z_\lambda^k &:= \frac{E[u^k|\theta] - RP_\lambda}{\sqrt{h_\lambda^k}} \end{aligned}$$

Now we can rewrite Equation 2.27 as

$$E[V(W_{I_j}^{\lambda_k})|P_\lambda] = e^{ac^k(\lambda_k)} V(RW_{0j}) E\left[\exp\left(\frac{-h_\lambda^k}{2\sigma_{\epsilon_k}^2}(z_\lambda^k)^2\right)|P_\lambda\right] \quad (2.28)$$

Since P_λ is a linear function of θ , conditional on P_λ , $E[u^k|\theta]$ is normally distributed. Therefore, $(z_\lambda^k)^2$ is distributed with Chi-squared. Hence, moment generating function of $(z_\lambda^k)^2$ has the form

$$E[e^{-t(z_\lambda^k)^2}|P_\lambda] = \frac{1}{\sqrt{1+2t}} \exp\left(\frac{-t(E[z_\lambda^k|P_\lambda])^2}{1+2t}\right) \quad (2.29)$$

Also, by definition, $Var[u_k|P_\lambda] = \sigma_{\epsilon_k}^2 + h_\lambda^k$. Thus we can simplify the second term in Equation 2.28 as

$$\begin{aligned} E\left[\exp\left(\frac{-h_\lambda^k}{2\sigma_{\epsilon_k}^2}(z_\lambda^k)^2\right)|P_\lambda\right] &= \frac{1}{\sqrt{1 + \frac{h_\lambda^k}{\sigma_{\epsilon_k}^2}}}\exp\left(\frac{-E[u^k|\theta] - RP_\lambda}{2(h_\lambda^k + \sigma_{\epsilon_k}^2)}\right) \\ &= \sqrt{\frac{Var[u_k|\theta]}{Var[u_k|P_\lambda]}}\exp\left(\frac{-E[u^k|\theta] - RP_\lambda}{2Var[u_k|P_\lambda]}\right) \end{aligned} \quad (2.30)$$

Using (2.30) we can rewrite (2.28) as

$$E[V(W_{I_j}^{\lambda k})|P_\lambda] = e^{ac^k(\lambda_k)}V(RW_{0j})\sqrt{\frac{Var[u_k|\theta]}{Var[u_k|P_\lambda]}}\exp\left(\frac{-E[u^k|\theta] - RP_\lambda}{2Var[u_k|P_\lambda]}\right) \quad (2.31)$$

Using similar steps, we can also write

$$E[V(W_{U_j}^{\lambda k})|P_\lambda] = V(RW_{0j})\exp\left(\frac{-E[u^k|\theta] - RP_\lambda}{2Var[u_k|P_\lambda]}\right) \quad (2.32)$$

The result immediately follows. ■

Lemma 5. *Given the distributional assumptions, there exists a market price for a given λ with the form*

$$P_\lambda(\theta) = \alpha_{0\lambda} + \alpha_{0\lambda}\left(\theta_2 + \frac{\gamma\lambda_1\sigma_{\epsilon_2}^2}{(1-\gamma)\lambda_2\sigma_{\epsilon_1}^2}\theta_1\right) \quad (2.33)$$

where $\alpha_{0\lambda}$ and $\alpha_{0\lambda}$ are real numbers that possibly depend on λ but not on θ_1 or θ_2 .

Proof. Conjecture a linear price function for a given λ :

$$P_\lambda(\theta) = \alpha_{0\lambda} + \alpha_{1\lambda}\theta_1 + \alpha_{0\lambda}\theta_2 \quad (2.34)$$

Then, the signal uninformed period i investor will use from observing the price can be drawn from

$$\theta_i = \frac{P_\lambda - \alpha_{0\lambda} - \alpha_{k\lambda}\theta_k}{\alpha_{i\lambda}} \quad (2.35)$$

where i and k denote opposite periods. Since prior distribution is normal and the signal is a linear function of a normally distributed random variable, posterior distribution for u_i will also be normal with mean and variance:

$$E[u_i|P_\lambda] = \frac{\frac{\bar{\theta}_i}{\sigma_{\theta_i}^2} + \frac{1}{\sigma_{\theta_k}^2} \left(\frac{\alpha_{i\lambda}}{\alpha_{k\lambda}} \right)^2 \frac{P_\lambda - \alpha_{0\lambda} - \alpha_{k\lambda}\bar{\theta}_k}{\alpha_{i\lambda}}}{\frac{1}{\sigma_{\theta_i}^2} + \frac{1}{\sigma_{\theta_k}^2} \left(\frac{\alpha_{i\lambda}}{\alpha_{k\lambda}} \right)^2} \quad (2.36)$$

$$Var[u_i|P_\lambda] = \sigma_{\epsilon_i}^2 + \frac{1}{\frac{1}{\sigma_{\theta_i}^2} + \frac{1}{\sigma_{\theta_k}^2} \left(\frac{\alpha_{i\lambda}}{\alpha_{k\lambda}} \right)^2} \quad (2.37)$$

Plugging Equations 2.36 and 2.37 into the market clearing condition in 2.3, we can rewrite as

$$\begin{aligned} & \gamma \left[\lambda_1 \left[\frac{\theta_1 - RP_\lambda}{a\sigma_{\epsilon_1}^2} \right] \right. \\ & \left. + (1 - \lambda_1) \left[\frac{\frac{\bar{\theta}_1}{\sigma_{\theta_1}^2} + \frac{1}{\sigma_{\theta_2}^2} \left(\frac{\alpha_{1\lambda}}{\alpha_{2\lambda}} \right)^2 \left(\frac{P_\lambda - \alpha_{0\lambda} - \alpha_{2\lambda}\bar{\theta}_2}{\alpha_{1\lambda}} \right) - RP_\lambda \left(\frac{1}{\sigma_{\theta_1}^2} + \frac{1}{\sigma_{\theta_2}^2} \left(\frac{\alpha_{1\lambda}}{\alpha_{2\lambda}} \right)^2 \right)}{a + a\sigma_{\epsilon_1}^2 \left(\frac{1}{\sigma_{\theta_1}^2} + \frac{1}{\sigma_{\theta_2}^2} \left(\frac{\alpha_{1\lambda}}{\alpha_{2\lambda}} \right)^2 \right)} \right] \right] \\ & + (1 - \gamma) \left[\lambda_2 \left[\frac{\theta_2 - RP_\lambda}{a\sigma_{\epsilon_2}^2} \right] \right. \\ & \left. + (1 - \lambda_2) \left[\frac{\frac{\bar{\theta}_2}{\sigma_{\theta_2}^2} + \frac{1}{\sigma_{\theta_1}^2} \left(\frac{\alpha_{2\lambda}}{\alpha_{1\lambda}} \right)^2 \left(\frac{P_\lambda - \alpha_{0\lambda} - \alpha_{1\lambda}\bar{\theta}_1}{\alpha_{2\lambda}} \right) - RP_\lambda \left(\frac{1}{\sigma_{\theta_2}^2} + \frac{1}{\sigma_{\theta_1}^2} \left(\frac{\alpha_{2\lambda}}{\alpha_{1\lambda}} \right)^2 \right)}{a + a\sigma_{\epsilon_2}^2 \left(\frac{1}{\sigma_{\theta_2}^2} + \frac{1}{\sigma_{\theta_1}^2} \left(\frac{\alpha_{2\lambda}}{\alpha_{1\lambda}} \right)^2 \right)} \right] \right] = 0 \end{aligned} \quad (2.38)$$

From Equation 2.38, after a basic but tedious algebra P_λ can be left alone. From

there, coefficient of θ_1 becomes $\alpha_{1\lambda}$, coefficient of θ_2 becomes $\alpha_{2\lambda}$ and the rest becomes $\alpha_{0\lambda}$.

Lastly, it is trivial to see from Equation 2.38 that

$$\alpha_{1\lambda} = \frac{\gamma\lambda_1\sigma_{\epsilon_2}^2}{(1-\gamma)\lambda_2\sigma_{\epsilon_1}^2}\alpha_{2\lambda} \quad (2.39)$$

Hence the expression in the Lemma follows. ■

Proposition 4. *In interior equilibria, fraction of informed period 1 investors λ_1 increases as*

- (1) a, C_1, C_2, γ and $\sigma_{\epsilon_2}^2$ decreases
- (2) $\sigma_{\theta_2}^2$ increases.

Proof. Equating $\psi^k(\lambda)$ and taking the square of both sides, we can write

$$e^{2ac_1(\lambda_1)}\sigma_{\epsilon_1}^2 = \sigma_{\epsilon_1}^2 + \frac{1}{\frac{1}{\sigma_{\theta_1}^2} + \frac{1}{\sigma_{\theta_2}^2} \left(\frac{\gamma\lambda_1\sigma_{\epsilon_2}^2}{(1-\gamma)\lambda_2\sigma_{\epsilon_1}^2} \right)^2} \quad (2.40)$$

and

$$e^{2ac_2(\lambda_2)}\sigma_{\epsilon_2}^2 = \sigma_{\epsilon_2}^2 + \frac{1}{\frac{1}{\sigma_{\theta_2}^2} + \frac{1}{\sigma_{\theta_1}^2} \left(\frac{\gamma\lambda_1\sigma_{\epsilon_2}^2}{(1-\gamma)\lambda_2\sigma_{\epsilon_1}^2} \right)^{-2}} \quad (2.41)$$

Further algebra yields $\lambda_1(\lambda_2)$ and $\lambda_2(\lambda_1)$:

$$\lambda_2 = \lambda_1 \frac{\gamma\sigma_{\epsilon_2}^2}{(1-\gamma)\sigma_{\epsilon_1}\sigma_{\theta_2}} \frac{1}{\sqrt{\frac{1}{e^{2ac_1(\lambda_1)}-1} - \frac{\sigma_{\epsilon_1}^2}{\sigma_{\theta_1}^2}}} \quad (2.42)$$

Table 2.2: Information Acquisition Possibilities The table shows which -if any- equations would not be satisfied if the parameter in the row were to increase and the equilibrium objects λ_1 and λ_2 were to move as designated in the column.

Parameter	Direction of λ_1, λ_2			
	+,+	-, -	+, -	-, +
γ	(2.44)	✓	(2.43)	(2.42)
c_1	(2.44)	✓	(2.42)	(2.43)
c_2	(2.44)	✓	(2.43)	(2.42)
$\sigma_{\epsilon_1}^2$	(2.44)	✓	✓	(2.43)
$\sigma_{\epsilon_2}^2$	(2.44)	✓	(2.42)	✓
$\sigma_{\theta_1}^2$	✓	(2.44)	(2.43)	✓
$\sigma_{\theta_2}^2$	✓	(2.44)	✓	(2.42)

and

$$\lambda_1 = \lambda_2 \frac{(1 - \gamma)\sigma_{\epsilon_1}^2}{\gamma\sigma_{\epsilon_2}\sigma_{\theta_1}} \frac{1}{\sqrt{\frac{1}{e^{2ac_2(\lambda_2)} - 1} - \frac{\sigma_{\epsilon_2}^2}{\sigma_{\theta_2}^2}}} \quad (2.43)$$

Lastly, using (2.42) and (2.43), we can derive

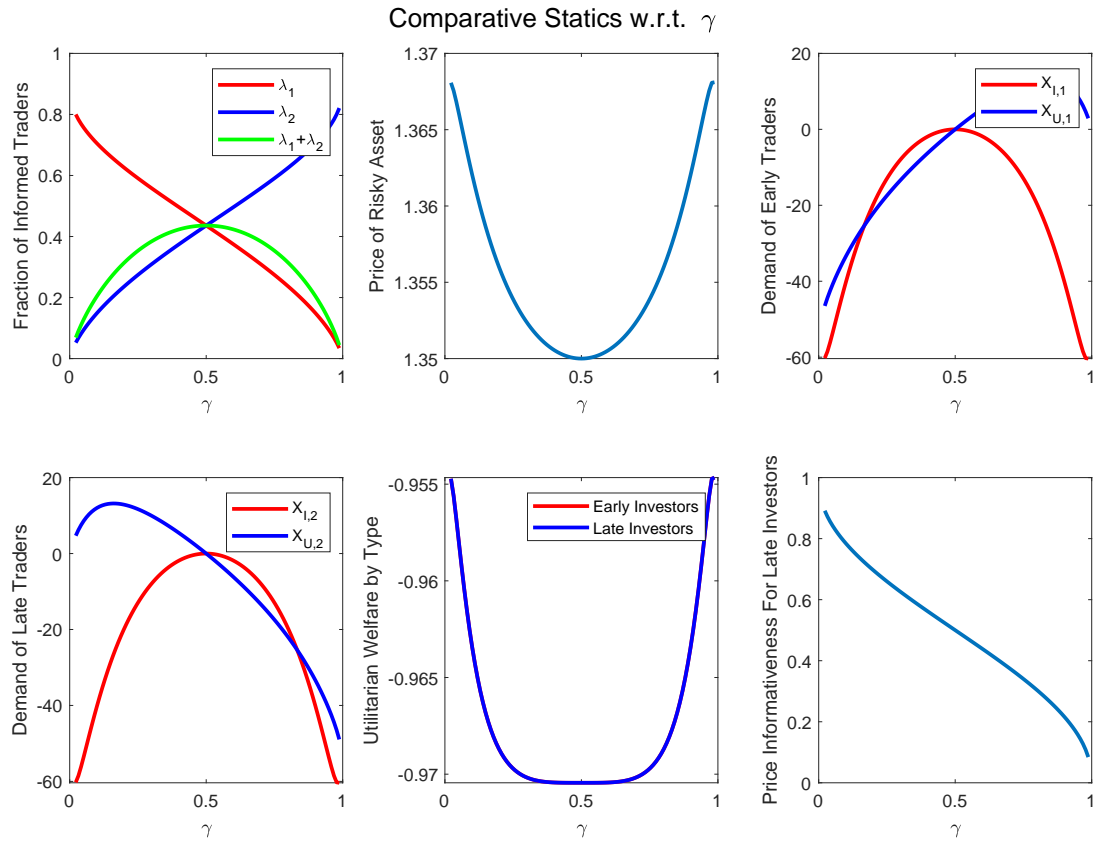
$$\frac{\sigma_{\epsilon_1}\sigma_{\epsilon_2}}{\sigma_{\theta_1}\sigma_{\theta_2}} = \sqrt{\frac{1}{e^{2ac_1(\lambda_1)} - 1} - \frac{\sigma_{\epsilon_1}^2}{\sigma_{\theta_1}^2}} \sqrt{\frac{1}{e^{2ac_2(\lambda_2)} - 1} - \frac{\sigma_{\epsilon_2}^2}{\sigma_{\theta_2}^2}} \quad (2.44)$$

Comparative static results can be derived from impossibility results, i.e. whether a certain direction of the movement in the parameters and the equilibrium object can be compatible with the equations (2.42), (2.43) and (2.44). Below is a table that summarizes which directions can be compatible with the equations and which equations the other directions violate.

The results follow from the table. ■

2.B Comparative Statics

Figure 2.5: Comparative Statics for the Benchmark Model with respect to γ , where $\theta_1 = \theta_2 = 1.25$ and $\theta_1 = \theta_2 = 1.35$.



Chapter 3

Changing Jobs to Fight Inflation: Labor Market Reactions to Inflationary Shocks

BY GORKEM BOSTANCI, OMER KORU AND SERGIO VILLALVAZO [†]

3.1 Introduction

Since job switches are usually associated with wage and productivity increases¹, the speed at which employees change employers is considered a measure of the health of the economy. Understanding what drives differences in job-to-job transitions across time and countries can be crucial for improving economic performance. In this paper, we identify a novel policy tool that affects the rate of job-to-job transitions: monetary

[†]University of Pennsylvania.

¹See e.g. Fallick and Fleischman, 2004, Christensen et al., 2005 and Jolivet, Postel-Vinay, and J.-M. Robin, 2006. Under a large variety of theoretical models, job changes come with changes in both wage and productivity (See Postel-Vinay and J.-M. Robin, 2002 and Menzio and Shi, 2011).

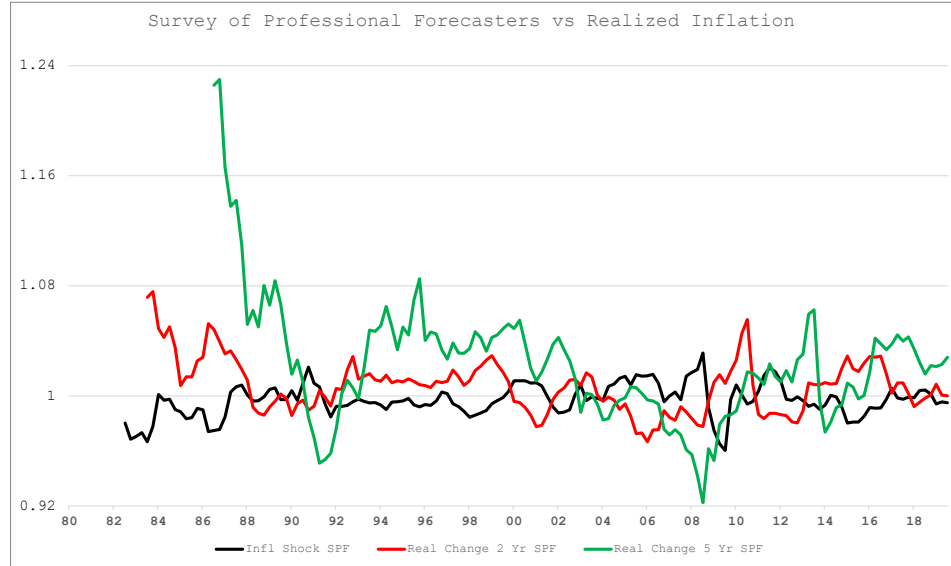
policy. When wages are not indexed to inflation², workers' real wages decrease at a faster rate in periods with unexpectedly high inflation. Therefore, potential gains from being able to renegotiate wages are higher for workers. Workers could respond to a positive inflationary shock by (1) increasing their search effort, thus, making it more likely that they will receive a job offer and (2) being less selective, i.e., accepting lower wage offers which lead to less productive matches. The first channel (search effort, henceforth) increases the number of job transitions, while the extent to which these transitions lead to more productive matches depends on the size of the second channel (selectivity, henceforth). Hence, the impact of inflation shocks on output is ambiguous and potentially depends on the size of the shock.

We measure how unexpected inflation affects aggregate productivity through its impact on the job search behavior of workers. We first utilize reduced-form causal inference to argue a quantitatively meaningful change in the rate of job switches following inflation shocks. We find that a 1% decline in real wages due to an unexpected inflation shock leads to a 7 percentage points increase in the probability of receiving a job offer in the following six-month period. To understand the resulting change in productivity, we build a model of directed on-the-job search with aggregate shocks. We calibrate the model to match the empirical job switching patterns and associated wage increases. The calibrated model suggests a non-monotonic output response following inflation shocks, suggesting both channels (search effort and selectivity) are quantitatively meaningful.

Although unexpected inflation movements have been relatively small for the U.S., they can imply a large drop in real wages once accumulated. Figure 3.1 summarizes this idea. The black line represents the discrepancy between Survey of Professional

²Existence of nominal frictions in wage setting has long been documented. See Appendix 3.B for a broad overview of the evidence regarding the extent of wage indexation.

Figure 3.1: The Discrepancy Between the SPF Forecast and Realized Inflation The x axis refers to the calendar year. The black line represents the difference between the 1-year ahead SPF forecast and the realized inflation. The values above 1 indicate inflation exceeded forecasts. The red line represents the cumulative real wage loss for a worker who signed his contract two years ago, based on SPF forecasts. The green line represents the cumulative real wage loss for a worker who signed his contract five years ago, based on SPF forecasts.



Forecasters (SPF) 1-year ahead inflation forecast versus the realized inflation. The red (green) line represents what fraction of the intended real wage is received by a worker who signed a contract 2 (5) years ago based on SPF forecasts. The real wage losses can be as high as 8%, and gains can be as high as 16% for some workers³ even though the surprise inflation never exceeds 6% and is mostly below 3% in magnitude. Hence, the output response can be large once small inflation shocks accumulate.

Our paper is motivated by the recent finding by Moscarini and Postel-Vinay, 2017 and Karahan et al., 2017 that once job-to-job transition rates are controlled for, unemployment-to-employment transition rates have little to no predictive power

³See Appendix Figure 3.10 for the same plot with the Michigan Consumer Survey inflation forecasts.

on nominal wage growth. On the other hand, the job-to-job transition rate and nominal wage growth have a large significant positive correlation. This is at odds with the classical Philips Curve idea where low unemployment strengthens workers' bargaining position and puts upward pressure on wages. It rather suggests the real threat point of the workers being switching to another job, that is, firms are more likely to increase wages when job-to-job transitions are more likely. Our analysis confirms the co-movement between job-to-job transitions and the inflation rate. Acknowledging that both objects are equilibrium outcomes, we try to unpack which shocks might be behind the positive correlation and the aggregate implications of the connection between the two.

In the first half of the paper, we provide three main pieces of empirical evidence that suggests the positive correlation between inflation and the job-to-job transition rate is driven by the positive effect of the former on the latter, rather than the other way around. First, we run Granger Causality tests on the aggregate data as well as panel regressions across U.S. regions and states. While inflation helps predict future job-to-job transition rates, job-to-job transitions do not help predict future inflation movements. Second, we use the previous estimates of structural monetary policy shocks instead of inflation in our regressions. This analysis allows us to look beyond the reverse causality argument, as these shocks are arguably exogenous to the economic conditions. Our results suggest that an unexpected one percent decrease in nominal interest rates can bring an increase in the job-to-job transition rates up to 0.4% percent. Third, we provide some direct evidence on the mechanism using individual-level survey data on on-the-job search behavior. We find that a cumulative wage loss of 1% due to unexpectedly high inflation increases the likelihood of receiving an offer by 7 percentage points and the expected number of offers by 0.17 in a six-month period.

In the second half, we build a model of competitive on-the-job search with endogenous search effort where the contract space is restricted to nominal wage contracts. The environment involves aggregate shocks to productivity where the agents form rational expectations. In the model, the agents respond to an unexpected positive inflation shock by increasing their search effort, as the option value of search increases⁴. Simultaneously, the agents also respond by searching in markets with lower posted wages as their current situation becomes more desperate. Hence, they trade a higher wage for a higher probability of finding a new job. The increased search effort leads to more frequent job-to-job transitions, which, by itself, would increase average productivity. However, the reduced asking wage makes these transitions less productivity-enhancing, therefore creates a force that decreases average productivity. In short, inflationary shocks unambiguously increase job-to-job transitions while their effect on productivity is undetermined. A preliminary calibration of the model to the U.S. economy confirms the non-monotone response of the output. When the unexpected increase in inflation is bigger than a threshold value, the selectivity channel starts to dominate, and the output decreases.

The proposed mechanism has important implications. First, it explains how output response may not be monotonic in the size of the inflation shock. Thus, it provides a bridge between seemingly disparate estimates of the literature on the real effects of monetary policy shocks⁵. Second, it provides a novel mechanism on how monetary policy can affect the real economy in the short run. Through monetary policy shocks, the monetary authority can improve the allocation of labor in the economy, thus increase productivity. Third, it provides a novel channel that can explain why some recessions are associated with a more pronounced ‘cleansing’ effect than the others.

⁴See e.g., Christensen et al., 2005 and Mueller, 2010 for evidence on job search effort decreasing as workers move up the job ladder.

⁵See Wolf, 2019 for an overview of these findings.

In our model, the sign and the magnitude of the unexpected price movement can affect both the speed and the effectiveness of job reallocation during the recession.

This paper is closely related to the literature that analyzes the interaction between inflation and the efficiency of labor markets. In particular, the idea that inflation helps reduce labor market frictions and increase productivity was first proposed by Tobin, 1995 and tested by Card and Hyslop, 1997. In this channel, nominal downward wage rigidity can be made non-binding with a positive inflation rate that ensures nominal rigidity doesn't translate to a real rigidity⁶. Our model incorporates this benefit of inflation, on top of our novel channel, that it incentivizes job switches. The most closely related work to ours is by Moscarini and Postel-Vinay, 2019 (MPV henceforth), who incorporate a random on-the-job search framework into a New Keynesian DSGE model. When job-to-job transition rates are high, employees receive more offers, some of which are matched by the incumbent firm. Matched offers are essentially cost shocks to the firm and it responds by raising prices. Hence, a higher than average job-to-job transition rate brings higher than average price inflation. The mechanism in MPV and ours are complementary. MPV shows how a labor demand shock brings wage inflation and therefore price inflation. We show how a shock to price inflation increases job-to-job transitions. Thus, our contribution is three-fold. First, our mechanism, in combination with theirs, explains how labor demand shocks can be amplified through a combination of offer matching and changing search behavior. Second, shocks to price inflation can also trigger this cycle. Third, the monetary policy recommendations could change⁷. because the monetary authority needs to consider

⁶Lunnemann and Wintr, 2010 find real wage rigidity is indeed more substantial in Luxembourg where there is a state-imposed automatic wage indexation.

⁷Tom Fairless of the Wall Street Journal, in his article based on the results by MPV, argues "If workers are less willing to switch jobs, central banks could press harder on the gas pedal to stimulate the economy without worrying about inflation. And there may be little policy-makers can do to influence the job-switching rate except to watch it." (2019, Nov 17 <https://www.wsj.com/articles/one-explanation-for-weak-wage-growth-workers-reluctance->

the job switching response to predict the response of the real economy. MPV assumes on-the-job search effort is fixed, hence shuts down our channel by assumption⁸. The empirical evidence in Section 3.2 favors our channel if one or the other has to be picked.

This paper also contributes to the literature on the efficiency of job reallocation. This literature asks when reallocation is productivity-enhancing and when it is not. A broad finding is that U.S. recessions were accompanied with productivity-enhancing job reallocation until the great recession⁹ while the reallocation during the great recession was both slower and less productivity-enhancing (Mukoyama, 2014 and Foster, Grim, and J. Haltiwanger, 2016). J. C. Haltiwanger et al., 2018 asks whether the decline is due to a decreased number of transitions or a smaller productivity gain conditional on making a transition and find most of the decline comes from the latter. Caballero and Hammour, 1994 discusses potential frictions that may create inefficient job reallocation during recessions. Barlevy, 2003 emphasizes increased credit market frictions while Ouyang, 2009 suggests early exits as mechanisms large enough to reverse the ‘cleansing’ effect of the recessions¹⁰. Gautier, Teulings, and Van Vuuren, 2010, in a model with on-the-job search, analyzes which wage-setting mechanisms generate socially efficient job switches. They conclude, for social efficiency, the hiring

to-switch-jobs-11573999201?shareToken=st5a849d04f72440fca240048db4bad6d1). Our mechanism suggests there is a direct link from monetary policy shocks to job-to-job transition rates.

⁸Incorporating the search effort channel in their model is not trivial. In MPV, only the distribution of productivities across jobs is a state variable while adding the search effort makes the joint distribution of wages and productivities a state variable. The surplus function is not sufficient to characterize the transitions either because the search effort choice is not efficient due to the restricted contract space. Hence, the tricks in Lise and J.-M. Robin, 2017 cannot be used to simplify the problem. Our model avoids this issue by utilizing the block-recursivity of competitive search where the distributions are no longer state variables. We present a version of our model under random search in Appendix 3.E

⁹See e.g. Davis and J. Haltiwanger, 1992, Caballero and Hammour, 1994, Davis, Faberman, and J. Haltiwanger, 2006 and (Davis, Faberman, and J. Haltiwanger, 2012).

¹⁰Foster, J. Haltiwanger, and Syverson, 2008 shows if pricing decisions are not taken into account, the effect of demand and productivity shocks on profitability can be confounded. Thus, the reallocation that is only profitability enhancing can be mislabeled as productivity-enhancing.

premium (to induce the worker to undertake search) should equal the no-quit premium (to prevent the worker from making a job switch later) which happens in wage posting with commitment but not in wage bargaining or the sequential auctions of Postel-Vinay and J.-.-M. Robin, 2002. The competitive search framework we use also satisfies the efficiency requirement posited here; the inefficient switches in our setting are purely due to nominal frictions. The closest papers to ours in this literature are by Moscarini, 2001 and Barlevy, 2002. Moscarini, 2001 considers a trade-off similar to ours. In his model, similar to the competitive search models, workers decide between a good match with a long queue and a mediocre match with a short queue. Thus, in tight labor markets, the initial matches are of higher quality and the reallocation is slow. Barlevy, 2002 shows decreasing job-to-job transitions during recessions can generate an effect large enough to offset the ‘cleansing’ effect of recessions. In his model, after a bad productivity shock, firms post fewer vacancies, which reduces the rate of job-to-job transitions, thus the productive reallocation of workers in the economy. In contrast, our model focuses on the effect of the inflationary shocks and generates productivity drops even when the reallocation rate is higher.

Lastly, our mechanism is also related to the literature that analyzes how the extent of wage flexibility affects the output response to monetary policy shocks. Olivei and Tenreyro, 2007 shows that the effects of monetary policy shocks depend on their timing during the year, and it is consistent with the fact that a significant fraction of firms renegotiate wage contracts at the end of the year. Björklund, Carlsson, and Nordström Skans, 2019 find that the output response to monetary policy is bigger in periods where a larger fraction of wage contracts are nominally fixed, using a micro-level dataset on details of collective wage agreements in Sweden¹¹.

¹¹See also Benabou, 1992 and Diamond, 1993 for how inflation affects search effort in product markets.

We proceed with the description of the data used. Section 3.2 provides the empirical analysis. Section 3.3 lays down the model and provides the theoretical results. Quantitative results of the model are presented in Section 3.4.

3.2 Empirical Analysis

This section presents three types of evidence to argue that the positive correlation between inflation and job-to-job transitions stems from the causal effect of inflation on job-to-job transitions. First, subsection 3.2.1 uses the time-series structure of the data to show that a high inflation today predicts a high job-to-job transition rate in the future. In contrast, a high job-to-job transition rate today does not predict high inflation in the future. For this aim, both Vector Auto Regressions with aggregate data and panel regressions with state level data are used. Second, subsection 3.2.2 uses popular estimates of structural Monetary Policy shocks to get a causal estimate of the effect of inflation on job-to-job transitions and confirms that higher inflation causes higher job-to-job transitions. Third, subsection 3.2.3 provides direct evidence on how inflation increases the job search effort of the employed from survey data. We later use the estimates from this subsection to discipline the macro model.

3.2.1 Predictive Regressions

This subsection presents findings from three datasets: (1) national monthly job-to-job transition and inflation series between 1995-2018, (2) national yearly series between 1976-2018, and (3) quarterly state-level series between 2000-2018. All three analyses use different periods due to data limitations but support the same argument: higher inflation predicts higher job-to-job transitions in the future, while higher job-to-job transitions do not predict higher inflation.

Following Moscarini and Postel-Vinay, 2019, we will use a variable called ‘acceptance rate’ as introduced therein. This variable is the ratio of the job-to-job transition rate to the unemployment-to-employment transition rate. The division is to ensure that employees’ willingness to switch jobs is isolated from job availability, which moves both rates simultaneously. The ‘acceptance rate’ is a natural candidate for what our mechanism is about; higher inflation affects job-to-job transitions by changing the employees’ willingness to switch¹².

The other primary variable we construct is called ‘inflation mistake’ and defined as the discrepancy between the expected and the realized inflation for a one-year period. At a time t , this measures the accumulated unexpected prices moves since time $t - 1$.

Monthly Analysis, Nation Level

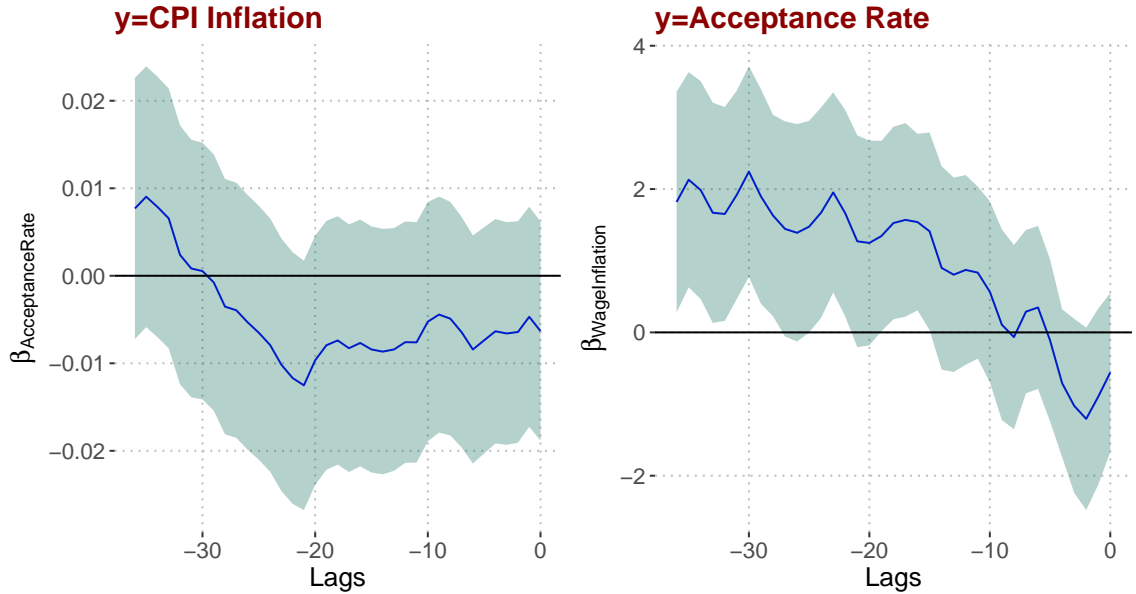
In this section, we use the series made available by Fujita, Moscarini, and Postel-Vinay, 2019¹³ that covers the period from September 1995 to December 2018 for the monthly job-to-job transition rates. Over-the-year log changes in the Consumer Price Index (CPI) provide price inflation. Inflation expectations are taken from the University of Michigan Survey of Consumers. We take logs of all labor market variables, and HP filter all variables with a smoothing parameter of 8.1×10^6 as in Moscarini and Postel-Vinay, 2019.

The Granger Causality test rejects if the lags of variable x help predict variable y above and beyond the lags of variable y . We find inflation Granger-causes job-to-job transition rates with 5% significance, while the other direction shows no predictive

¹²In MPV, the acceptance rate is primarily determined by the position of the workforce in the job ladder. Since the search effort of the employed is not a choice, and the switches are exogenous, no other model component can affect the acceptance rate once conditioned on the distribution of workers across jobs.

¹³See Appendix 3.A for details on the data sources used throughout the empirical analysis.

Figure 3.2: National Predictive Regressions The left panel presents the coefficient estimates and the associated 95% CI for β where price inflation is regressed on the ‘acceptance rate’ with the specification in Equation 3.1. Each point corresponds to an estimate where the associated lag is in the x-axis. The right panel provides the same plot where the ‘acceptance rate’ is regressed on the price inflation. See Appendix 3.A for details of the data sources.



relationship.

We continue by replicating the analysis in Moscarini and Postel-Vinay, 2019 that questions and rejects a price Philips curve. Specifically, we run OLS regressions of the form:

$$y_t = \beta x_{t-L} + \gamma Z_{t-L} + \epsilon_t \quad (3.1)$$

Firstly, we set the price inflation as y and the acceptance rate as x and then switch their places. The unemployment rate and unemployment-to-employment transition rate constitute Z in both types of regressions. We vary L from 0 to 36 months and analyze how β changes. Figure 3.2 presents the results of this analysis. The left panel indicates no significant relationship between the lags of the acceptance rate and the CPI inflation, where most of the estimates up to 2,5 years are negative. On the other

hand, as shown in the right panel, a higher CPI inflation predicts a higher ‘acceptance rate’ 15 to 36 months after¹⁴.

Quarterly Analysis, State Level

Here we utilize the Longitudinal Employer Household Dynamics (LEHD) data set by the U.S. Census. The LEHD provides publicly available job-to-job transition rates in quarterly frequency at the state level starting from 2000. This structure allows using the state-level variation in prices and job-to-job transitions¹⁵. Unfortunately, the state-level inflation data is only available in yearly frequency and starts from 2008¹⁶. Therefore, we use state-level wage inflation data from the Quarterly Census of Employment and Wages (QCEW) as a proxy. We take logs and four-quarter trailing moving averages of all labor market variables, and HP filter all variables with a smoothing parameter of 10^5 as in Moscarini and Postel-Vinay, 2019. We then combine all the data and run OLS regressions of the form:

$$y_{it} = \beta x_{it-L} + \gamma y_{it-L} + v_i + \nu_t \epsilon_{it} \quad (3.2)$$

where we analyze the lead-lag relationship between the wage inflation and the ‘acceptance rate’¹⁷. v_i and ν_t denote the state and time fixed effects. The results are in Figure 3.3. Our mechanism would be able to explain both panels. The search effort channel would suggest wage inflation be a positive predictor of the ‘acceptance


¹⁴Potential causal channels in either direction would take some time to show up in the data. In our mechanism, workers need to realize the real wage changes and manage to find a job after increasing their search effort before a change in job-to-job transition numbers can be observed. Similarly, under the classical menu cost assumptions, the mechanism argued by Moscarini and Postel-Vinay, 2019 requires firms to adjust their prices after their labor costs go up.

¹⁵CPS, which is monthly, provides information regarding the location of the participant. However, once the sample is divided into job switchers across states, the sample size becomes an issue.

¹⁶See Personal Consumption Expenditures (PCE) by State in <https://apps.bea.gov/regional/downloadzip.cfm>.

¹⁷The results are robust to removing the fixed-effects or y_{it-L} from the right-hand side of (3.2).

Figure 3.3: State-Level Predictive Regressions The left panel presents the coefficient estimates and the associated 95% CI for β where wage inflation is regressed on the ‘acceptance rate’ with the specification in Equation 3.2. Each point corresponds to an estimate where the associated lag is in the x-axis. The right panel provides the same plot where the ‘acceptance rate’ is regressed on the wage inflation. See Appendix 3.A for details of the data sources.



rate’ through its effect on price inflation. On the other hand, the job ladder channel, which is first proposed by MPV, would suggest the ‘acceptance rate’ be a negative predictor of wage inflation. If the ‘acceptance rate’ is high, workers are at the bottom of the ladder, and switches come with small wage improvements.

3.2.2 Structural Monetary Policy Shocks

Although the results in Section 3.2.1 are suggestive, they do not prove any causal relationship between inflation and job-to-job transitions. Here, we use structural estimates of monetary policy shocks as exogenous proxies for the inflation level. Our mechanism would imply a negative relationship with nominal interest rate shocks and job-to-job transitions.

We use several popular monetary policy shock estimates in the literature. The first

measure is computed from narrative records of FOMC meetings and internal forecasts of Federal Reserve by C. D. Romer and D. H. Romer, 2004, which is updated until 2007 by Wieland and Yang, 2016. The second measure is by Barakchian and Crowe, 2013 that uses Fed Funds futures to see exogenous changes in policy. The third measure is by Sims and Zha, 2006, who use structural VAR estimates to identify shocks to monetary policy. Fourth, fifth and sixth measures are by Gertler and Karadi, 2015 and Nakamura and Steinsson, 2018 who use high-frequency movements in financial series during FOMC announcements to identify monetary policy shocks¹⁸. The periods that match with the availability of job-to-job transitions data are all different across these measures, but results from regressions with all measures are consistent.

$$y_t = \beta x_{t-L} + \epsilon_t \tag{3.3}$$

Here, the majority of the coefficients are negative as expected. Furthermore, all but one of the significant coefficients are negative. These results further add to the evidence in support of our theory, that is, higher price inflation leads to higher job-to-job transitions.

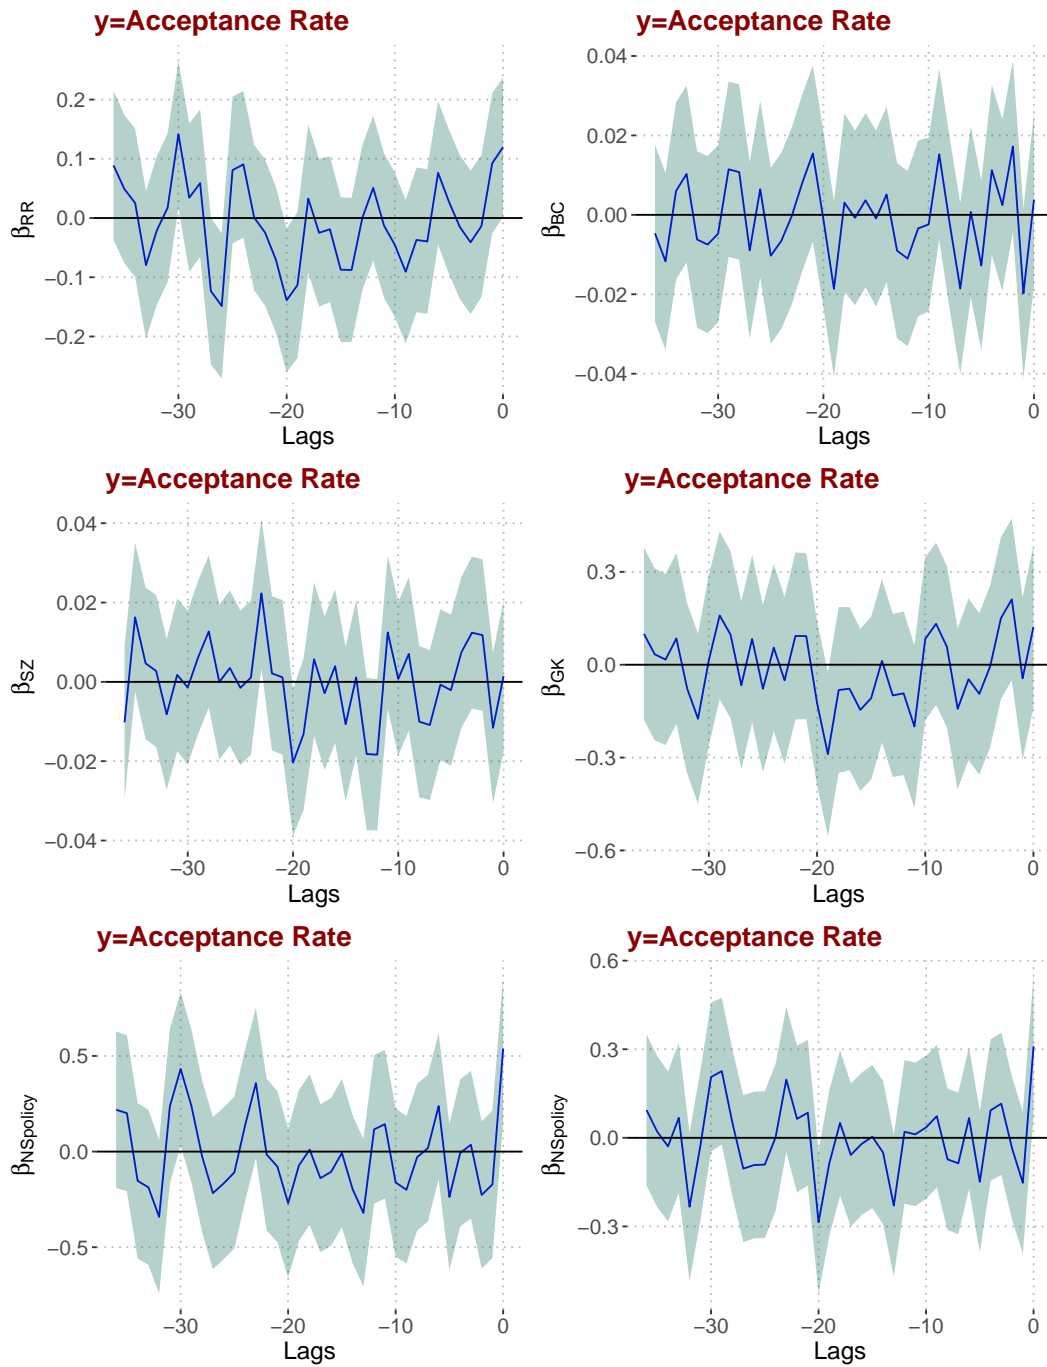
3.2.3 Survey Evidence on Search Effort

The analysis here utilizes the Job Search supplement of the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE)¹⁹. We use the publicly available data from 2013 to 2016. The Survey of Professional Forecasters (SPF), which is

¹⁸Readers should refer to Ramey, 2016 for an excellent review on these and other monetary policy shock estimation methods.

¹⁹The SCE is administered monthly as a rotating panel, and the Job Search supplement adds detailed questions on job search behavior in the October survey. Since no respondent stays in the SCE for more than a year, the supplement becomes a repeated cross-section.

Figure 3.4: Monetary Policy Shocks and Acceptance Rate Each panel presents the coefficient estimates and the associated 95% CI for β where 'acceptance rate' is regressed on a structural monetary policy shock estimate with the specification in Equation 3.3. Each point and the bar correspond to an estimate where the regressors are with the associated lag in the x-axis. See Appendix 3.A for details of the data sources.



administered quarterly, provides one-year ahead inflation expectations²⁰.

To understand the effect of inflation on job search effort, the ideal measure would be the accumulated real wage loss (or gain) the agent has due to unexpected price movements. This object, unfortunately, is not available at the individual level. We instead use another object denoted as ι_i :

$$\iota_i = \prod_{t=\tau_{0i}}^{\tau_{si}} \frac{1 + i_t}{1 + \hat{i}_{\tau_{0i}}} \quad (3.4)$$

where τ_{0i} and τ_{si} denote the dates individual i started her job and took the survey, respectively. i_t denotes the realized CPI inflation rate and \hat{i}_t denotes the SPF inflation expectations at date t . If the realized sequence of inflation rates is higher (lower) than the inflation expectations in the beginning, then the agent's real wage will be less (more) than intended, and ι_i will be larger. At the individual level, this measure only requires the job-start date of the worker, which is available in SCE.

The measure also has two main drawbacks. First, if the contract is renegotiated after the start date, the measure will break down. To alleviate this issue, we will focus on individuals who started their current job recently²¹. Second, SPF inflation expectations are only available at 1-year and 10-year horizons. Thus, we assume that inflation expectations n year ahead are the same as the 1-year ahead inflation repeating itself n times.

In the regressions below, we will restrict attention to full-time employees with a single job, who are (1) searching for another full-time job, (2) have been working for at least a year and (3) the reason for the search is not a firing notice or a non-work

²⁰Although the SCE provides the inflation expectations of each respondent, we believe the relevant inflation expectation that shapes a wage bargaining process is the one given by the firms and the policymakers.

²¹We will use workers with tenures for less than five years to have a compromise between guaranteeing that the start date is the last negotiation date and keeping a large sample.

related reason.

The empirical design we will use is of the form:

$$y_i = \beta_0 + \beta_1 \iota_i + \beta_2 \ln(\textit{tenure}) + \gamma_{\tau_{0i}} + \alpha_{\tau_{si}} + \beta X_i + \epsilon_i \quad (3.5)$$

where $\gamma_{\tau_{0i}}$ and $\alpha_{\tau_{si}}$ denote fixed effects for the job-start year and the survey year, respectively. X includes demographic controls for age, gender, education, and marital status. y_i denotes outcome variables measuring the extent of the search effort. In our exercise, we will use the number of offers received and a dummy variable for whether any offers were received in the past six months²².


The identification idea is built on the random sampling of the surveys. Conditional on a job-start date, the survey dates of individuals are randomly assigned barring survival bias. Once we control for the job start and survey years and the tenure of the worker, we can treat ι_i as randomly assigned²³. Figure 3.5 presents the histogram of ι_i values in the final sample.

Table 3.1 presents the results for an Ordinary Least Squares and a Linear Probability Model. The results indicate that unexpected inflation increases the likelihood of receiving an offer as well as the number of offers received at a 5% level. The ‘Wage Mistake’ variable is a ratio and is expected to be centered around 1. According to our estimates, a 1% positive inflation shock translates to 7.9% higher probability of receiving an offer and 0.17 more offers on average. Later in Section 3.4, we will use these coefficients to validate our model’s ability to assess the relationship between

²²SCE has other potential outcome variables such as the number of employers applied and the hours spent searching. However, any measure that quantifies effort through intermediate steps requires caution. The time spent searching or employers contacted are highly related to whether the specific type of effort translates into offers. For example, a high amount of time spent might indicate employee’s inefficient search strategies. Similarly, a large number of employers contacted might indicate a quantity/quality trade-off in the application strategy.

²³According to our identification argument, the demographic controls are also not strictly required. We include them only to reduce the regression variance. Including them has minimal effect on our quantitative results.

Figure 3.5: Histogram of ι_i Values



inflation and job-to-job transitions correctly.

Table 3.1: The Effect of Inflation on Search Effort Each column presents the coefficient estimate for β_1 and the associated standard error with the specification in Equation 3.5. The independent variable is ι_i as constructed in Equation 3.4. The dependent variables are the number offers received and whether an offer was received by the respondent in the past six months respectively. The controls whose estimates are excluded from the table are job-start date, survey date, tenure, age, gender, education, and marital status. See Appendix 3.A for details of the data sources.

	<i>Dependent variable:</i>	
	Number of Offers Received	Received (0-1)
	(1)	(2)
WageMistake	17.395** (7.293)	7.920** (3.991)
Observations	374	374
R ²	0.189	0.082
Adjusted R ²	0.155	0.043
Residual Std. Error (df = 358)	0.536	0.293
F Statistic (df = 15; 358)	5.555***	2.118***

*p<0.1; **p<0.05; ***p<0.01

3.3 The Model

3.3.1 Environment

The environment has two main frictions that are required to generate the monetary non-neutrality. First, firms and employees are not allowed to sign state-contingent contracts. Second, search frictions prevent perfect competition in the labor markets. Therefore, shocks to inflation introduce shifts in real wages of existing employees. Since employed also search on the job, the model exhibits monetary non-neutrality even though the wages of new hires are completely flexible. If all labor contracts were inflation-adjusted or labor markets were competitive, our model would exhibit monetary neutrality.

Here, we describe an environment where all variables are real. We then introduce shocks to the real wages of existing employees as inflation shocks and match these shocks to the discrepancy between the inflation forecasts and the realized inflation in the data. This will allow us to avoid nominal variables in our modeling which can be conceptualized as a limit of the classical New Keynesian model where pricing frictions go to zero²⁴.

Preferences

The economy consists of a continuum of individuals with measure one and a continuum of firms with positive measure. Both the workers and the firms are risk-neutral and maximize the expected discounted income/profits. Time is discrete, and firms and workers share the same discount factor, $\beta \in (0, 1)$.

²⁴We choose to avoid a full New Keynesian structure with pricing frictions. First, this allows us to isolate the effects of inflation through the labor market, without having to worry about other moving parts. Second, once included, pricing frictions require dynamically optimizing firms that break block-recursivity. Thus, we would be forced to use Taylor approximations to solve the model.

Production Technology

There is a single homogeneous consumption good in the economy. When a worker and a firm match, they produce $y + z$ units of output. The first component, y , is the aggregate productivity, and it is the same across firms. The second component, z , is match specific. Upon meeting, z is drawn from a distribution \mathcal{G} and remains the same until separation.

Unemployed workers produce b units of output.

Meeting Technology

Workers and firms need to find each other to produce. Search is directed, and markets are indexed by the value offered by a firm to a worker. We denote submarkets by $X \in \mathbb{R}$.

Both unemployed and employed workers can search for a job. After they choose in which submarket to search for a job, workers choose the search effort, e . The cost of exerting effort is denoted by $c(e)$ and it is a strictly increasing and convex function with the following properties: $c(0) = 0, c'(0) = 0$ ²⁵.

Firms also choose in which submarket to post their vacancies. The cost of opening a vacancy for one period is $\kappa > 0$.

In a submarket, firms and workers meet each other via a constant returns to scale matching function, M . Given v measure of vacancies and E unit of total search effort, there are $M(v, E)$ measure of matches. Constant returns to scale assumption implies that market tightness θ , i.e. vacancy-to-total search effort ratio, is sufficient to characterize the probability of matching. Specifically, a worker that exerts e unit of search effort finds a job with probability $ep(\theta)$, where $p : \mathbb{R} \rightarrow [0, 1]$ is a strictly increasing and concave function with following properties: $p(0) = 0, p(x) \rightarrow 1$ as

²⁵We consider the search cost as a utility cost, thus it doesn't appear in the output calculations.

$x \rightarrow \infty$. On the other hand, a vacancy meets a worker with probability $q(\theta)$, where $q : \mathbb{R} \rightarrow [0, 1]$ is a strictly decreasing function with the following property: $\theta q(\theta) = p(\theta)$.

After a firm and a worker meets, they draw match productivity z and decide whether to form a match or not.

Wage Setting

The contract space is limited to fixed-wage contracts. In other words, if a firm and a worker meet in a submarket X and decide to form a match, then firm offers a wage rate w that provides an expected lifetime utility of X to worker, taking into consideration the search effort cost and the separation risk (either exogenous or through the worker finding a better job). X and the aggregate state are sufficient to pin down the wage, since it depends on future lifetime utility X , not past outcomes. Also, the match productivity does not affect the lifetime value of the worker since it is constant throughout the firm-worker match. Let ψ be the aggregate state of the economy, which consists of aggregate productivity y and distribution of workers across jobs and wages $\Gamma(z, w)$ ²⁶. We denote the entry wage of a worker in submarket X when the aggregate state is ψ by $h(X, \psi)$.

Timeline

Each period is divided into five sub-periods. In the first sub-period, aggregate productivity y is drawn. In the second sub-period, exogenous separations occur with probability $\delta \in (0, 1)$. In the third sub-period, workers choose where to search and how much effort to exert. In this stage, workers who were separated from their job in the current period cannot search for a job; they remain unemployed with probability

²⁶Unlike Menzio and Shi, 2011, the wage distribution matters for determining future tightness because it determines the aggregate search effort.

one. In the fourth sub-period, workers and firms meet and decide whether to form a match. In the last sub-period, production takes place, and wages are paid.

Discussion of the Model Elements

While setting the environment, we make five main simplifications. Four of them are innocuous while the fifth is not.

First, we denote all the variables in real terms. Second, we avoid modeling an inflation process with rational expectations over it. Third, we assume fixed-wage contracts although all that is needed for the mechanism is that they are not state-contingent. In principle, we can focus on nominal wages, allow an inflation process that follows an $AR(\infty)$ and contracts that are functions of time. In that scenario, employees and firms could sign contracts that take the expected future inflation into account and designate an associated increase in nominal wages over time. Therefore, nominal wages would follow a path that leaves the real wages constant over time absent shocks to inflation and aggregate productivity. Using the real wages as the model element allows us to abstract from the expected paths of the nominal variables and focus on the shocks to the inflation process. None of these three simplifications have a bearing on the final results while they simplify the notation greatly.

Fourth, we don't allow firms to make counter offers for their poached employees. In theory, this might result in workers moving to jobs with lower productivity than their current jobs, which wouldn't happen if the incumbent firms could respond. We make the assumption for computational simplicity. More importantly, in our quantitative exercise, we don't observe this behavior with the calibrated parameters. Therefore, allowing the firms to respond should have no quantitative effect on our results.

The fourth simplification, namely, treating inflation as an exogenous process, is not completely innocuous. In a fully-fledged New Keynesian model, output shocks and

monetary shocks both contribute to determining the inflation. Therefore, treating the inflation shocks as completely independent from output shocks would not be entirely correct. On the other hand, introducing firms that price dynamically would break the block-recursivity of the equilibrium. Thus, whenever we draw conclusions from the past data, we will not only rely on the inflation series. Instead, we will focus on the discrepancy between the inflation expectations and the realized inflation while remaining agnostic on how these expectations are formed in the economy.

3.3.2 Equilibrium

Problem of a Firm

Since the production technology is constant returns to scale, the size of the firm is indeterminate. Hence, we consider single vacancy firms. Let $K(w, z, \psi)$ be the value function of a filled vacancy with match productivity z , wage rate w and aggregate state ψ . Observe that a firm is willing to form a match in submarket X if and only if the match productivity z satisfies $K(h(X, \psi), z, \psi) \geq 0$. Since the firm value is increasing in z , define \underline{z} such that $K(h(X, \psi), \underline{z}, \psi) = 0$. If such \underline{z} exists, the expected value of finding a worker is:

$$J(X, \psi) = \int_{z \geq \underline{z}} K(h(X, \psi), z, \psi) dG(z).$$

The free entry condition implies that

$$k \geq q(\theta)J(X, \psi), \tag{3.6}$$

where left-hand side is the cost of vacancy, and the right-hand side is the expected value of a vacancy, which is the product of the probability of finding a worker and

the expected value of a filled vacancy. This condition holds with equality whenever there is a positive mass of workers searching for a job in submarket X . Hence, there is a one-to-one relationship between market tightness θ and (X, ψ) . Hence, we can write $\theta(X, \psi)$ as the market tightness in active submarkets.

Let $\bar{p}(H(w, \psi), \psi)$ be the probability that a worker leaves the job when his lifetime value is $H(w, \psi)$ and the aggregate state is ψ . Then,

$$K(w, z, \psi) = y + z - w + \beta(1 - \delta)\mathbb{E}[(1 - \bar{p}(H(w, \psi'), \psi'))K(w, z, \psi)] \quad (3.7)$$

The model has endogenous separations, which affect the wage-setting problem in a non-trivial way. In a search model where job switches are efficient, a la Postel-Vinay and J.-.-M. Robin, 2002, the probability of losing a worker is completely exogenous. Thus, the sequential auctions protocol dictates firms to pay the minimum wage that will allow them to keep/attract the worker. Once search effort is introduced, firms may want to offer a wage that is more than absolutely needed to reduce the incentives of the worker to exert search effort and attract more offers. This kills the simple structure of the sequential auctions protocol. The additional complication is smaller in a directed search framework, however, results in a firm value function K that is not monotone in the wage (or value) offered.

Problem of an Unemployed Worker

Consider an unemployed worker. We write down the problem of the unemployed right before the production sub-period. The value function of an unemployed worker is

$$U(\psi) = b + \beta\mathbb{E} \left[\max_e eR(\psi', U) - c(e) + U(\psi') \right], \quad (3.8)$$

where $R(\psi, V)$ is return to searching in the optimal submarket for an agent with lifetime value of V :

$$R(\psi, V) = \max_x p(\theta(\psi, X))(X - V)(1 - G(\underline{z}(\psi, X)))$$

e does not appear in $R(\psi, X)$, because search effort is exerted after the choice of submarket²⁷. After the choice of submarket, worker chooses an effort level to maximize the term inside the brackets in (3.9).

Problem of an Employed Worker

Similarly, we can define the value function of an employed worker as:

$$H(w, \psi) = w + \beta \mathbb{E} \left[\delta U(\psi') + (1 - \delta) \max_e (eR(\psi', H(w, \psi')) - c(e) + H(w, \psi')) \right]. \quad (3.9)$$

Equilibrium Definition

Following Menzio and Shi, 2011, we consider block recursive equilibria. In a block-recursive equilibrium, policy functions do not depend on the distribution of workers across jobs. Hence, the only relevant aggregate variable is aggregate productivity y , i.e., $\psi = y$ ²⁸.

A block-recursive equilibrium consists of a market tightness function $\theta : Y \times \mathbb{R} \rightarrow \mathbb{R}_+$, a value function for the unemployed $U : Y \rightarrow \mathbb{R}$, a value function for the employed $H : \mathbb{R}_+ \times Y \rightarrow \mathbb{R}$, a value function for the firm $K : \mathbb{R}_+ \times Z \times Y \rightarrow \mathbb{R}$, optimal choice of submarket $m : \mathbb{R} \times Y \rightarrow \mathbb{R}$, optimal choice of search effort $e : \mathbb{R} \times Y \rightarrow \mathbb{R}_+$, entry

²⁷Since a worker is measure zero, his choice of e does not effect θ , hence it does not effect the choice of submarket.

²⁸Since the search effort choice is an innocuous extension of the framework in Menzio and Shi, 2010, we do not prove existence and uniqueness of the block-recursive equilibrium here. Schaal, 2017 provides a discussion of the possible scenarios where block-recursivity may fail.

wage $h : \mathbb{R} \times Y \rightarrow \mathbb{R}$ and the cutoff for match productivity $\underline{z} : \mathbb{R} \times Y \rightarrow \mathbb{R}_+$ such that:

1. $\underline{z}(X, \psi)$ satisfies $K(h(X, \psi), \underline{z}, \psi) = 0$,
2. entry wage $h(X, \psi)$ solves $H(h, \psi) = X$,
3. $H(w, \psi)$ satisfies (3.9), $U(\psi)$ satisfies (3.8), $K(w, z, \psi)$ satisfies (3.7) where probability that a worker finds a job is

$$\bar{p}(w, \psi) = e(m(H(w, \psi), \psi))p(\theta(\psi, m(H(w, \psi), \psi)))(1 - G(\underline{z})),$$
4. $e(V, \psi)$ and $m(V, \psi)$ solve worker's problem,
5. $\theta(\psi, X)$ satisfies the free entry condition (3.6).

3.3.3 Effect of a Decrease in Real Wage

What happens if a worker's real wage decreases for some exogenous reason, for example, inflation? In this section, we show that there are two competing mechanisms: a decrease in selectivity in on-the-job search and an increase in the search effort. First, we prove that when a worker's current lifetime utility decreases, she searches in a lower-valued submarket, which has a lower cutoff for match-specific productivity. Second, we prove that the worker increases the search effort.

Lemma 6. $\underline{z}(X, \psi)$ is increasing in promised lifetime utility X .

This lemma states that as the promised lifetime utility increases, to form a match, a better match specific productivity draw is needed. The intuition is clear: if a firm promises higher value, its lifetime value decreases. Hence, at the marginal match specific productivity, the firm starts making a loss. Therefore, the firm is more selective in high indexed markets.

Lemma 7. $m(V, \psi)$ is increasing in current lifetime utility V .

This lemma states that workers with low current lifetime utility searches in a market that promises lower lifetime utility compared to a worker with higher current lifetime utility. This mechanism implies a job ladder, workers start from the bottom and gets better lifetime utilities as they find new jobs and climb the job ladder.

Lemma 8. $R(\psi, V)$ is decreasing in current lifetime utility V .

As a worker's current lifetime utility increases, there is a lower gain from finding a better job. Hence, return to searching for a job increases. This mechanism also implies that search effort is decreasing with lifetime value.

Lemma 9. $e(V, \psi)$ is decreasing in current lifetime utility V .

Lemmas 6 and 7 show that a worker with a lower current lifetime utility search in a lower indexed submarket, in which cutoff for the match-specific productivity is lower. Hence, if a worker's wage decreases, the expected productivity of her next job is lower than the expected productivity in the market she previously searched in²⁹. On the other hand, Lemma 9 shows that the worker increases his search effort. Hence, the probability of moving to a better job increases.

At the micro-level, inflation has a direct impact on individual's lifetime utility. However, at the macro level, inflation does not have a direct effect, i.e., if workers do not change their behavior, there would be no change in the aggregate output. However, due to these two competing channels, the aggregate output might decrease or increase in the short-run due to inflation. One time inflation shock does not have an impact on the steady-state, thus, there are no long-run implications.

²⁹There might even be a probability that she ends up at a worse job than the current one she has. In the calibrated model, we don't observe this possibility.

If the first channel dominates, workers end up with lower match specific productivities, which leads to lower aggregate output. If the second channel dominates, workers increase their search effort and form new matches with higher match productivity. This mechanism leads to a higher aggregate output. Therefore, impact of an inflation shock is ambiguous. We proceed to quantify the importance of each channel in Section 3.4.

3.4 Quantitative Analysis

This section presents the preliminary calibration strategy and the quantitative results.

3.4.1 Calibration Strategy

For the output predictions to have a quantitative interpretation, two implied elasticities should be plausible: (1) the response of job-to-job transitions to an inflationary shock and (2) the response of aggregate output to job-to-job transitions. We measure the former elasticity from micro-data that documents how workers adjust their search behavior with inflationary shocks (see Section 3.A). The latter can be inferred from wage increases following job switches and a measure of how surplus is shared between firms and workers. Although matching these two elasticities is necessary for pinning down the output response, it is not sufficient. The response of the aggregate output to job-to-job transitions depends on the underlying reasons for these transitions. The output response following increased transitions due to a labor demand shock does not necessarily equal the response due to an inflationary shock. Thus, it is crucial to model these two together instead of stitching two elasticities that are computed separately.

We use a telephone-line matching function: $p(\theta) = \theta(1 + \theta^\gamma)^{-1/\gamma}$ ³⁰ and assume the match specific productivity distribution \mathcal{G} follows a Pareto distribution with location parameter z_{min} and shape parameter z_{shape} . Lastly, we assume a quadratic search cost function $c(e) = Ae^2$ where the level potentially differs for the employed A_e and the unemployed A_u .

The full set of parameters necessary to compute the model is the vector:

$$\Omega = \{\beta, \delta, \gamma, \kappa, A_e, A_u, b, z_{min}, z_{shape}, \rho_y, \sigma_y\} \quad (3.10)$$

The model period is taken to be a month. We normalize z_{min} to equal the unemployment benefit replacement rate, calibrate β and δ externally, and calibrate the remaining parameters internally. We calibrate the parameters to match the steady state moments, except for the parameters that determine the process of aggregate productivity process. Then, we calibrate the aggregate productivity process to match the business cycle statistics.

We set the monthly discount factor $\beta = 0.95^{1/12}$ and exogenous separation rate $\delta = 0.011$ consistent with the average EU rate in 2005 (Fallick and Fleischman, 2004).

Calibration Idea

The model doesn't admit an analytic expression for the steady state distribution of workers across jobs, hence we stick to discussing the broad intuition of how the moments inform the parameter values. The calibration uses all moments to discipline all parameters, since general equilibrium effects through market tightness prevents isolating the response of different moments.

The residual wage distribution informs the match productivity distribution z_{shape} ,

³⁰The telephone-line matching function, proposed by Stevens, 2007, is a flexible matching function that has the Cobb-Douglas as a special case.

Table 3.2: Calibrated Parameters All parameters in the table are jointly calibrated to match all the moments. The last column provides an intuitive mapping between the parameters and the moments that are most related. Avg. labor prod. is constructed by HP filtering the logged series with smoothing parameter 10^5 . In order to construct the residual wage distribution, we first construct an hourly wage measure through dividing the weekly wage by the usual hours worked. Then, we regress the hourly wage on age, age squared, gender, race, marital status, and education level in the cross-section for each month of 2005 in CPS. Lastly, we take the average of the quantiles of the distributions of residuals from each regression.

Parameter	Value	Moment	Data	Model	Source
γ Match. Funct. Elasticity	3	UE	0.30	0.30	Shimer '05
A_e for employed	0.066	EE	0.024	0.016	Shimer '05
A_u for unemployed	4.5	labor share	0.60	0.60	
b Unemployment Flow	5.50	residual log wage q75	0.54	0.32	CPS
z_{shape}	1.95	residual log wage q25	-0.32	-0.2679	CPS
κ Vacancy Cost	0.066	median tenure	48	31	CPS

and the flow benefit of unemployment b . The flow benefit disciplines the left tail because the wage bargaining between the firm and an unemployed worker depends on the outside option of the worker. The right tail depends on how large the match productivity can be, hence on z_{shape} .

The employment-to-employment (EE) and unemployment-to-employment (UE) transition rates inform the search effort cost level parameters for the employed A_e and the unemployed A_u respectively. A higher transition rate implies a lower cost.

The labor share disciplines the vacancy cost κ , hence the surplus sharing between the firm and the worker in the model. A higher labor share implies a low κ . Lastly, the median tenure helps discipline the matching function elasticity γ . As the elasticity gets larger, firms become more aggressive with the wage postings and the median tenure goes down.

Calibration Results

The calibrated parameters together with the matched moments are given in Table [3.2](#).

3.4.2 Unexpected Inflation Shock

This section presents how the economy responds to unexpected shocks to inflation of different sizes. In particular, the quantitative findings confirm the analytic results in Section 3.3.3. While small positive inflation shocks increase the output in the short run, large positive inflation shocks decrease it. Negative shocks to inflation uniformly decrease the output.

Figure 3.6: Impulse Responses to Inflation Shocks Each panel presents the impulse response functions with respect to an unexpected inflation shock of different sizes.

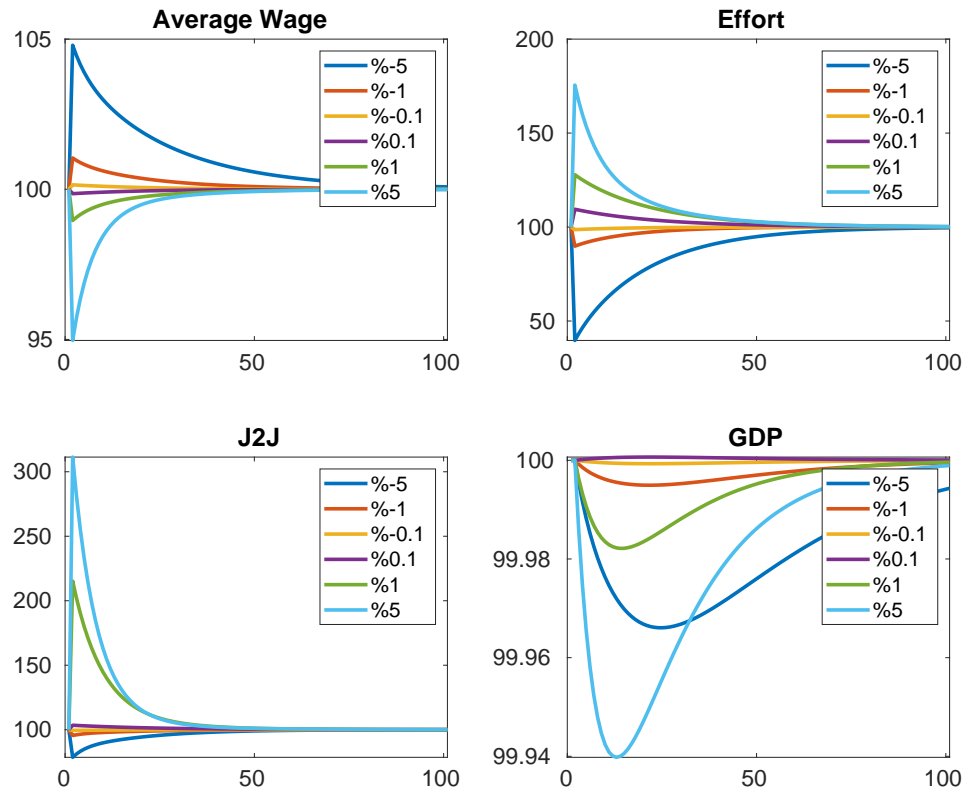


Figure 3.6 displays the impulse-responses for shocks to inflation of sizes 1 pp and 0.5 pp. The instantaneous change in average wages reflects the size of the inflation shock. The job-to-job transition rate increases following both shocks together with

the average on-the-job search effort. However, while the smaller shock brings a short-run boost to output, the larger shock causes a short-run decline. Here, one important implication of the counter-acting mechanisms manifests itself. The drop in real wages brings the search effort up, which results in an increase in output. On the other hand, the same drop causes the employed to be more nervous about finding a new job more quickly. Hence, they look for jobs in markets where it is easier to find a job, where wages and productivity are lower as well. When the shock is small enough, the increased number of switches dominates the fact that each switch is less productivity-enhancing. When the shock gets larger, the latter channel starts to dominate and we see a drop in output.

Since the wages of new hires are perfectly flexible, job switches undo the effects of the one-time inflation shocks. Therefore, the model exhibits money neutrality in the long run.

3.5 Conclusion

In this paper, we try to understand the positive correlation with inflation and job-to-job transitions in the economy. We first show reduced form and causal evidence suggesting higher inflation causes more job-to-job transitions. In time-series and panel structures, we find that shocks to inflation precede shocks to job-to-job transition rates: lags of inflation are consistently good predictors of job-to-job transitions. In addition, using several monetary policy shock estimates, we argue the relationship seems to be causal and economically significant: 1% increase in the nominal interest rate corresponds up to 6.5% decrease in job-to-job transition rates in the U.S. We proceed by constructing a model that can explain these observations. In settings with wage rigidities, higher than expected inflation rates increase the benefit of searching

on the job. As employees increase their search effort, more job-to-job transitions occur and allocation of labor across firms improves in the short run. The mechanism carries important implications for monetary policy: an expansionary monetary policy can improve the allocation of resources in the economy and increase productivity in the short run.

References

- Barakchian, S Mahdi and Christopher Crowe (2013). “Monetary policy matters: Evidence from new shocks data”. In: *Journal of Monetary Economics* 60.8, pp. 950–966.
- Barlevy, Gadi (2002). “The sullyng effect of recessions”. In: *The Review of Economic Studies* 69.1, pp. 65–96.
- (2003). “Credit market frictions and the allocation of resources over the business cycle”. In: *Journal of monetary Economics* 50.8, pp. 1795–1818.
- Benabou, Roland (1992). “Inflation and efficiency in search markets”. In: *The Review of Economic Studies* 59.2, pp. 299–329.
- Björklund, Maria, Mikael Carlsson, and Oskar Nordström Skans (2019). “Fixed-Wage Contracts and Monetary Non-neutrality”. In: *American Economic Journal: Macroeconomics* 11.2, pp. 171–92.
- Caballero, Ricardo J. and Mohamad L. Hammour (1994). “The Cleansing Effect of Recessions”. In: *The American Economic Review* 84.5, pp. 1350–1368. ISSN: 00028282. URL: <http://www.jstor.org/stable/2117776>.
- Card, David (1990). “Unexpected Inflation, Real Wages, and Employment Determination in Union Contracts”. In: *The American Economic Review* 80.4, pp. 669–688. ISSN: 00028282. URL: <http://www.jstor.org/stable/2006702>.

- Card, David and Dean Hyslop (1997). “Does inflation” grease the wheels of the labor market?” In: *Reducing inflation: Motivation and strategy*. University of Chicago Press, pp. 71–122.
- Christensen, Bent Jesper, Rasmus Lentz, Dale T Mortensen, George R Neumann, and Axel Werwatz (2005). “On-the-job search and the wage distribution”. In: *Journal of Labor Economics* 23.1, pp. 31–58.
- Davis, Steven J, R Jason Faberman, and John Haltiwanger (2006). “The flow approach to labor markets: New data sources and micro-macro links”. In: *Journal of Economic perspectives* 20.3, pp. 3–26.
- (2012). “Labor market flows in the cross section and over time”. In: *Journal of Monetary Economics* 59.1, pp. 1–18.
- Davis, Steven J and John Haltiwanger (1992). “Gross job creation, gross job destruction, and employment reallocation”. In: *The Quarterly Journal of Economics* 107.3, pp. 819–863.
- Diamond, Peter A (1993). “Search, sticky prices, and inflation”. In: *The Review of Economic Studies* 60.1, pp. 53–68.
- Druant, Martine, Silvia Fabiani, Gabor Kezdi, Ana Lamo, Fernando Martins, and Roberto Sabbatini (2012). “Firms’ price and wage adjustment in Europe: Survey evidence on nominal stickiness”. In: *Labour Economics* 19.5, pp. 772–782.
- Fallick, Bruce C. and Charles A. Fleischman (2004). *Employer-to-employer flows in the U.S. labor market: the complete picture of gross worker flows*. Finance and Economics Discussion Series 2004-34. Board of Governors of the Federal Reserve System (US). URL: <https://ideas.repec.org/p/fip/fedgfe/2004-34.html>.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger (2016). “Reallocation in the great recession: cleansing or not?” In: *Journal of Labor Economics* 34.S1, S293–S331.

- Foster, Lucia, John Haltiwanger, and Chad Syverson (2008). “Reallocation, firm turnover, and efficiency: selection on productivity or profitability?” In: *American Economic Review* 98.1, pp. 394–425.
- Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay (2019). *Measuring Employer-to-Employer Reallocation*. Tech. rep. Mimeo.
- Gautier, Pieter A, Coen N Teulings, and Aico Van Vuuren (2010). “On-the-job search, mismatch and efficiency”. In: *The Review of Economic Studies* 77.1, pp. 245–272.
- Gertler, Mark and Peter Karadi (2015). “Monetary policy surprises, credit costs, and economic activity”. In: *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.
- Grigsby, John, Erik Hurst, and Ahu Yildirmaz (2019). *Aggregate nominal wage adjustments: New evidence from administrative payroll data*. Tech. rep. National Bureau of Economic Research.
- Haltiwanger, John C, Henry R Hyatt, Lisa B Kahn, and Erika McEntarfer (2018). “Cyclical job ladders by firm size and firm wage”. In: *American Economic Journal: Macroeconomics* 10.2, pp. 52–85.
- Hofmann, Boris, Gert Peersman, and Roland Straub (2012). “Time variation in US wage dynamics”. In: *Journal of Monetary Economics* 59.8, pp. 769–783.
- Jolivet, Grégory, Fabien Postel-Vinay, and Jean-Marc Robin (2006). “The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US”. In: *European Economic Review* 50.4, pp. 877–907. ISSN: 0014-2921. DOI: <https://doi.org/10.1016/j.euroecorev.2006.02.005>. URL: <http://www.sciencedirect.com/science/article/pii/S0014292106000286>.
- Karahan, Fatih, Ryan Michaels, Benjamin Pugsley, Ayşegül Şahin, and Rachel Schuh (2017). “Do Job-to-Job Transitions Drive Wage Fluctuations over the Business Cycle?” In: *American Economic Review* 107.5, pp. 353–57.

- Lise, Jeremy and Jean-Marc Robin (2017). “The macrodynamics of sorting between workers and firms”. In: *American Economic Review* 107.4, pp. 1104–35.
- Lunnemann, Patrick and Ladislav Wintr (2010). “Downward wage rigidity and automatic wage indexation: Evidence from monthly micro wage data”. In:
- McLaughlin, Kenneth J (1994). “Rigid wages?” In: *Journal of Monetary Economics* 34.3, pp. 383–414.
- Menzio, Guido and Shouyong Shi (2010). “Block recursive equilibria for stochastic models of search on the job”. In: *Journal of Economic Theory* 145.4, pp. 1453–1494.
- (2011). “Efficient search on the job and the business cycle”. In: *Journal of Political Economy* 119.3, pp. 468–510.
- Moscarini, Giuseppe (2001). “Excess worker reallocation”. In: *The Review of Economic Studies* 68.3, pp. 593–612.
- Moscarini, Giuseppe and Fabien Postel-Vinay (2017). “The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth”. In: *American Economic Review* 107.5, pp. 364–68.
- (2019). “The job ladder: Inflation vs. reallocation”. In: *Unpublished draft, Yale University*.
- Mueller, Andreas (2010). “On-the-job search and wage dispersion: New evidence from time use data”. In: *Economics Letters* 109.2, pp. 124–127.
- Mukoyama, Toshihiko (2014). “The cyclical nature of job-to-job transitions and its implications for aggregate productivity”. In: *Journal of Economic Dynamics and Control* 39, pp. 1–17.
- Nakamura, Emi and Jón Steinsson (2018). “High frequency identification of monetary non-neutrality: The information effect”. In: *Quarterly Journal of Economics*.

- Olivei, Giovanni and Silvana Tenreyro (2007). “The timing of monetary policy shocks”. In: *American Economic Review* 97.3, pp. 636–663.
- Ouyang, Min (2009). “The scarring effect of recessions”. In: *Journal of Monetary Economics* 56.2, pp. 184–199.
- Postel-Vinay, Fabien and Jean-Marc Robin (2002). “Equilibrium wage dispersion with worker and employer heterogeneity”. In: *Econometrica* 70.6, pp. 2295–2350.
- Postel-Vinay, Fabien and Jean-Marc Robin (2004). “To match or not to match? Optimal wage policy with endogenous worker search intensity”. In: *Review of Economic Dynamics* 7.2, pp. 297–330.
- Ragan Jr, James F and Bernt Bratsberg (2000). “Un-COLA: Why have cost-of-living clauses disappeared from union contracts and will they return?” In: *Southern Economic Journal*, pp. 304–324.
- Ramey, Valerie A (2016). “Macroeconomic shocks and their propagation”. In: *Handbook of Macroeconomics*. Vol. 2. Elsevier, pp. 71–162.
- Romer, Christina D and David H Romer (2004). “A new measure of monetary shocks: Derivation and implications”. In: *American Economic Review* 94.4, pp. 1055–1084.
- Schaal, Edouard (2017). “Uncertainty and unemployment”. In: *Econometrica* 85.6, pp. 1675–1721.
- Sims, Christopher A and Tao Zha (2006). “Were there regime switches in US monetary policy?” In: *American Economic Review* 96.1, pp. 54–81.
- Stevens, Margaret (2007). “New microfoundations for the aggregate matching function”. In: *International Economic Review* 48.3, pp. 847–868.
- Tobin, James (1995). “Inflation and unemployment”. In: *Essential Readings in Economics*. Springer, pp. 232–254.
- Wieland, Johannes F and Mu-Jeung Yang (2016). *Financial dampening*. Tech. rep. National Bureau of Economic Research.

Wolf, Christian K (2019). “SVAR (Mis-) Identification and the Real Effects of Monetary Policy Shocks”. In: *Manuscript, Princeton University*.

3.A Data Sources

Monthly Data

For the job-to-job flows, we use the series made available by Fujita, Moscarini, and Postel-Vinay, 2019³¹ that is computed from the Current Population Survey (CPS). The unemployment-to-employment transition (UE) rates are from Fallick and Fleischman, 2004, similarly computed from the CPS³². The Consumer Price Index (CPI) inflation and the unemployment rate (U) series are from the U.S. Bureau of Labor Statistics³³. The wage inflation series is computed the same way using the ‘Average hourly earnings of production and nonsupervisory employees, total private, not seasonally adjusted’, from the Current Employment Series (CES)³⁴.

Yearly Data

CPS provides an approximate measure for yearly job-to-job transition rates starting at 1976. We use the methodology proposed by Mukoyama, 2014 to deal with the time aggregation bias introduced by low frequency of data³⁵.

³¹This series is based on the method introduced by Fallick and Fleischman, 2004 while corrects for an attrition bias that starts with the changes in the survey questions in 2007 (<https://sites.google.com/site/fabienpostelvinay/working-papers/EEProbability.xlsx?attredirects=0&d=1>). We repeat our empirical exercises using the original series by Fallick and Fleischman, 2004 as a robustness check. The results are BLANK and are presented in Appendix BLANK.

³²<https://www.federalreserve.gov/pubs/feds/2004/200434/200434abs.html>.

³³The analyses where CPI is replaced with Personal Consumption Expenditures (PCE, <https://fred.stlouisfed.org/series/PCEPI#0>) provide quantitatively and qualitatively similar results.

³⁴<https://beta.bls.gov/dataViewer/view/timeseries/CEU0500000008>

³⁵The monthly series is based on the question introduced to CPS at 1994, that asks whether there were any changes in the employment status of the worker since last month (“SAMEMP”). The yearly data asks whether the employee works for the same employer as last year. If the answer is no

Quarterly Data

For the job-to-job flows, we use the series J2JHireR and J2JSepR which are computed by dividing the number of hires (or separations) with no unemployment period in between to the total labor force. The two series closely follow each other and give very similar qualitative and quantitative results. In the main text, we focus on the analysis with J2JSepR. LEHD does not have information on unemployed-to-employed transition rates, therefore we use the variable NEHireR instead. This variable is computed by dividing the number of hires (or separations) from non-employment to the total labor force. We get state-level wage inflation data from the Quarterly Census of Employment and Wages (QCEW). Specifically, we use the percentage change in state-level average weekly wages between quarters t and $t-4$ in privately owned firms³⁶. We use Local Area Unemployment Statistics (LAUS) from the BLS for state-level unemployment and labor-force data³⁷.

3.B Evidence on the Extent of Wage Indexation

Explicit measures of what fraction of wage contracts are indexed to inflation are unavailable for the U.S. economy. The measures that are based on the actual contract terms are restricted to collective agreements in the U.S., which varies in coverage over the years and does not apply to a random sample of the workers. Measures based

and there hasn't been a long unemployment spell, there needs to be at least one switch. However, there is no way to confirm multiple switches within a year. Furthermore, recall becomes a bigger problem when the time period the respondent is expected to remember is further back. Although the data quality is potentially lower due to these issues, the yearly series goes back to high inflation periods in the U.S., thus provides important variation in inflation levels that is not present in recent years.

³⁶We exclude the public sector to isolate the market forces in the wage changes. Using data from all the firms has little impact on qualitative and quantitative outcomes.

³⁷See <https://www.j2jexplorer.ces.census.gov>, <https://www.bls.gov/cew/downloadable-data-files.htm>, and <https://www.bls.gov/lau/>.

on changes in the nominal wages are imperfect due to several other factors affecting the wage process. However, even the most conservative estimates imply a very low level of wage indexation (less than 25%) in developed countries. Here, we discuss the implications of prior research on the extent of wage indexation.

Evidence Based on Contract Terms

The main papers on the prevalence of ‘cost-of-living adjustment’ (COLA) terms in contracts are Card, 1990 for Canada and Ragan Jr and Bratsberg, 2000 for the U.S. Card, 1990 looks at the universe of manufacturing union contracts (with more than 500 employees) signed between 1968 and 1983. He finds that 26% of them have an ‘escalation clause’ on average while the explicit indexation is very rare. The fraction with ‘escalation clause’ peaks at 65% in a period where the inflation is over 10%. Ragan Jr and Bratsberg, 2000 use the U.S. Bureau of Labor Statistics data on collective bargaining settlements to see the prevalence of COLA provisions. They document that even though 61% of the settlements had COLA provisions back in 1976, it has fallen all the way to 22% in 1996 when the data is no longer available. The COLA provisions are known to be much less prevalent among non-union workers. With the decline in unionization, collective agreements cover a smaller fraction of the labor force in either country today. We consider these measures as an upper bound on the extent of wage indexation. Druant et al., 2012 utilize a firm-level survey conducted in 17 European countries regarding wage adjustment practices. Across 15,000 firms from all industries, they document that only 11.5 % of the firms employ any formal indexation clause in employment contracts while only 10.9% have any informal inflation considerations in wage setting³⁸. More importantly, the survey also asks about the frequency of wage adjustments. This gives us a back-of-the-envelope

³⁸There is still large variation across countries. In Belgium, 98.2% of the firms have automatic wage indexation while in Italy, only 5.8% of the firms have any form of wage indexation.

mapping between the degree of indexation and the frequency of wage adjustments. Wage adjustments happen either yearly or more frequently for 74.4% of the firms. Thus, even when firms adjust wages frequently, this does not imply an implicit wage indexation.

Evidence Based on Wage Movements

McLaughlin, 1994, using PSID data, finds that the effect of unanticipated inflation on nominal wage growth is consistent with 42% indexation between 1970 and 1986. Hofmann, Peersman, and Straub, 2012, using a DSGE model, infers the extent of wage indexation in the economy from the time variation in U.S. wage dynamics. They estimate the degree of wage indexation to be 0.17 in 2000, compared to 0.91 in 1974, which is roughly in line with the time path of COLA coverage in collective bargaining agreements³⁹. More recently, Grigsby, Hurst, and Yildirmaz, 2019, using data from a payroll processing company in the U.S., found that approximately 36% of job stayers experience no nominal wage changes in a one-year period. Once contrasted with the evidence in Druant et al., 2012, the implied wage indexation should be less than 11.5%.

3.C Proofs

Lemma 6. $\underline{z}(X, \psi)$ is increasing in promised lifetime utility X .

Proof. Recall that \underline{z} solves

$$K(h(X, \psi), \underline{z}, \psi) = 0.$$

³⁹A major implication from the paper is that wage indexation is a response to increasing monetary policy uncertainty. Thus, the level of indexation should be endogenous to run counter-factual exercises that change monetary policy. Since we focus on one-time shocks, we abstract from endogenous indexation.

Clearly, as promised lifetime utility X increases, value of the firm decreases. In order to satisfy equality, z must be increased. ■

Lemma 7. $m(V, \psi)$ is increasing in current lifetime utility V .

Proof. Let $V_h > V_\ell$. We want to show that $m(V_h, \psi) \geq m(V_\ell, \psi)$. For simplicity, we drop the aggregate state variable, since we are only considering the change in current lifetime utility V and denote the associated choices as m_h and m_ℓ and associated market tightness as θ_h and θ_ℓ . Suppose the contrary: $m_h < m_\ell$. This implies that $m_\ell - V_h > m_h - V_h$. Since m_h is the optimal choice for V_h

$$\begin{aligned} p(\theta_h)(m_h - V_h) &\geq p(\theta_\ell)(m_\ell - V_h) \\ \implies p(\theta_h) &> p(\theta_\ell). \end{aligned}$$

Rearranging the first line also gives us:

$$p(\theta_h)m_h - p(\theta_\ell)m_\ell \geq [p(\theta_h) - p(\theta_\ell)]V_h.$$

Similarly, since m_ℓ is the optimal choice for V_ℓ

$$\begin{aligned} p(\theta_\ell)(m_\ell - V_\ell) &\geq p(\theta_h)(m_h - V_\ell) \\ [p(\theta_h) - p(\theta_\ell)]V_\ell &\geq p(\theta_h)m_h - p(\theta_\ell)m_\ell. \end{aligned}$$

Using combining these two conditions:

$$[p(\theta_h) - p(\theta_\ell)]V_\ell \geq [p(\theta_h) - p(\theta_\ell)]V_h \implies V_\ell \geq V_h.$$

Which contradicts the assumption that $V_h > V_\ell$. ■

Lemma 8. $R(\psi, V)$ is decreasing in current lifetime utility V .

Proof. By envelope theorem:

$$R_V(\psi, V) = -p(\theta(m(V, \psi))) < 0.$$

Hence, R is decreasing in V . ■

3.D Solution Method

We use Value Function Iteration with 20 grid points for the distribution of z , 5 points for the distribution of y , 200 points for the grid for V , and 600 points for the grid for w . We define $\tilde{K}(V, y, z) = K(h(V, y), y, z)$ for convenience and start with an initial guess $\tilde{K}^0(V, y, z)$. The algorithm works sequentially. At step i , we compute

1. $J^i(V, y)$ given $\tilde{K}^{i-1}(V, y, z)$
2. $\underline{z}^i(V, y)$ and $\theta^i(V, y)$ given $J^i(V, y)$
3. $U^i(y)$, $e^i(V, y)$, $R^i(V, y)$, and $m^i(V, y)$ given $\underline{z}^i(V, y)$ and $\theta^i(V, y)$
4. $H^i(w, y)$ given $e^i(V, y)$, $R^i(V, y)$, $m^i(V, y)$, $\underline{z}^i(V, y)$, $\theta^i(V, y)$, and $U^i(y)$
5. $K^i(w, y, z)$ given $e^i(V, y)$, $m^i(V, y)$, $\underline{z}^i(V, y)$, and $\theta^i(V, y)$
6. $h^i(V, y)$ given $H^i(w, y)$
7. $\tilde{K}^i(V, y, z)$ given $K^i(w, y, z)$ and $h^i(V, y)$

We stop when $d_{max}(\tilde{K}^i(V, y, z), \tilde{K}^{i-1}(V, y, z)) < \epsilon$ where d_{max} gives the maximum distance between the two vectors.

3.E A Random Search Model with Effort

In this section, we present a random-search version of our model in Section 3.3. The random-search version here doesn't have the selectivity channel, since workers do not direct their search to particular types of firms.

3.E.1 Preferences

The discrete-time economy is populated by a continuum of infinitely-lived workers and firms. The total measures of workers and firms are fixed and normalized to one. Each worker has ability x , which is distributed with cumulative distribution function G , and each firm has productivity y , which is distributed by cumulative distribution function Γ . Time is discrete.

Both firms and workers are risk neutral and have the same discount factor, $\beta \in (0, 1)$.

3.E.2 Production Technology

There is only one consumption good in the economy. A worker-firm pair (x, y) can produce $f(x, y)$ output, where f is strictly increasing in both arguments and super-modular, i.e. $f_i(x, y) > 0$ for $i \in \{x, y\}$ and $f_{xy}(x, y) > 0$, where f_i is the derivative with respect to i .

Super-modularity of f implies that output maximizing allocation is to match high productivity workers with high productivity firms.

Each unemployed worker produces $b(x)$ unit of output by herself. Lastly, each worker-firm pair dissolves with probability δ in a given period.

3.E.3 Meeting Technology

In order to produce workers and firms need to find each other through random search.

Both unemployed and employed worker can search for a job. In order to find a vacancy, workers need to exert search effort. $c(e)$ denotes the utility cost of exerting e units of effort for the employed. For simplicity, we assume that search effort of unemployed worker is fixed to 1 and there is no cost attached to this effort. However, employed person chooses e optimally⁴⁰. We assume $c(e)$ is convex and strictly increasing in e , with $\lim_{e \rightarrow 1} c(e) \rightarrow \infty$ to simplify matching probabilities.

Firms, on the other hand, choose how many vacancies to open. In order to open v units of vacancies, a firm needs to pay $\kappa(v)$, where $\kappa(v)$ is convex and strictly increasing.

Let E and V be the total measure of search effort and vacancies, respectively. Total measure of matches to be formed is denoted with $M(E, V)$, for a given E and V . We assume that $M(E, V)$ is homogeneous of degree one. Then, measure of matches per unit of search effort is given by $M(1, E/V)$. Let λ denote the probability that one unit of search effort matches a vacancy: $\lambda = M(E, V)/E$. The probability that a vacancy meets with a worker is given by $\lambda^f = M(E, V)/V$.

We define market tightness to be the measure of vacancies available per unit search effort and denote it with $\theta = V/E$. Homogeneity of degree one implies that match probabilities of workers and vacancies only depend on aggregate quantities through tightness: $\lambda(\theta) = M(\theta, 1)$, $\lambda^f(\theta) = M(1, 1/\theta)$. This implies that $\lambda^f(\theta) = \theta\lambda(\theta)$.

⁴⁰It is assumed that production level does not depend on search behavior of the worker.

3.E.4 Wage Setting

Upon meeting, firm makes a take or leave it offer to worker. Firms can only propose constant nominal wage contracts to workers from which workers can walk away from anytime. Contracts can be re-negotiated without cost.

Contract space is not complete. Firms cannot make wage rate contingent on the state of the economy. Moreover, search effort of worker is not contractible. Hence, when a firm makes an offer, it needs to take into account the search effort of the worker.

When an employed worker meets with another vacancy, incumbent firm can make a counter-offer. As in Postel-Vinay and J.-.-M. Robin, 2002, this triggers Bertrand competition between incumbent firm and poaching firm.

Let $V_t(w, x, y)$ be the lifetime utility of a worker type x who is employed at firm y with a wage w and let $J_t(w, x, y)$ be the present discounted profits of a firm with productivity y that employs worker x at wage w ⁴¹. Consider two firms with productivity $y' > y$ that are bargaining over a worker with type x . In Bertrand competition, the maximum that a firm can offer as wage is the entire output. In such a situation, the lifetime utility of a worker type x would be $V_t(f(x, y), x, y)$ with firm y . Therefore, firm y' should solve the following problem:

$$\begin{aligned} \max_w J_t(w, x, y') \\ \text{s.t. } V_t(w, x, y') \geq V_t(f(x, y), x, y). \end{aligned}$$

where constraint ensures that firm y cannot outbid the offer.

When a firm makes an offer, it needs to take into account the search effort of

⁴¹For brevity, instead of writing aggregate states in the value function, we index value functions with the time subscript.

the worker. Even though an increase in w decreases the output share of firm, it discourages the worker from searching for a job and getting new offers, which is good for the firm. Depending on which effect dominates, value function J_t might be increasing or decreasing with w . To simplify the model, as in Postel-Vinay and J.-M. Robin, 2004, we assume that $J_t(w, x, y)$ is a decreasing function of w .

Assumption 3. $J_t(w, x, y)$ is a decreasing function of w .

This assumption implies that constraint must hold with equality, since $V_t(w, x, y)$ is increasing in w .

There are three possibilities for a worker employed at a firm with productivity y . First, she might match with a firm that has higher productivity, $y' > y$. In this case high productive firm wins the bargaining and worker changes his job. The worker's lifetime utility becomes $V(f(x, y'), x, y)$ ⁴². Let $\phi(x, y, y')$ be the wage that solves $V(\phi(x, y, y'), x, y') = V(f(x, y), x, y)$. In other words, $\pi(x, y, y')$ is the wage rate of worker type x when she moves from y to y' .

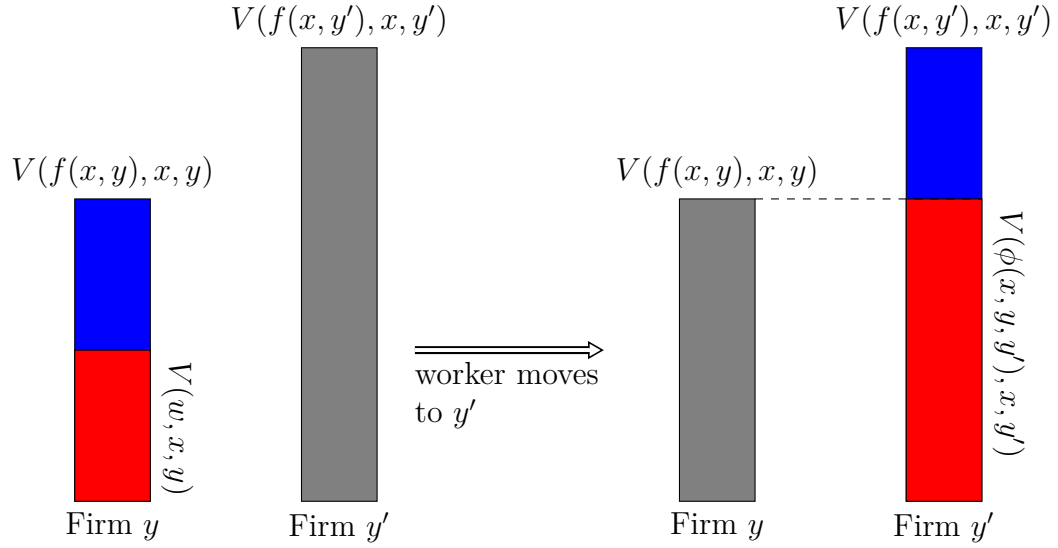
In the second case, the worker matches with a firm that has lower productivity, $y > y''$, however poaching firm can offer higher lifetime utility to worker than she currently has. In this case, poacher cannot win the bargaining, though bargaining increases the wage of the worker in the current firm. In this case, the worker's lifetime utility increases to $V(f(x, y''), x, y'')$ and her wage increases to $\pi(x, y'', y)$.

The second case can only happen if the current lifetime utility of the worker is lower than the maximum utility she could get from the poaching firm, i.e. $V(w, x, y) < V(f(x, y''), x, y'')$. In this situation, there is a room for firm y'' to make an offer.

In the third case, the poaching firm's productivity is so low that it cannot make any offer that triggers Bertrand competition. In this case there is no change in the

⁴²Here, since all comparisons happen at the same aggregate state, we suppress the time subscripts to reduce notation.

Figure 3.7: Worker x in Firm y Matches with Firm $y' > y$.



worker's wage and lifetime utility.

Let $\tilde{y}(w, x, y)$ be the minimum productivity level that a firm can trigger a bargaining. The following table summarizes the bargaining outcome between incumbent firm with productivity y and poaching firm with productivity y' :

- $y' > y$: Poaching firm offers $\phi(x, y, y')$, worker moves to firm y' and her lifetime utility becomes $V(f(x, y), x, y)$. See Figure 3.7.
- $\tilde{y}(w, x, y) \leq y' \leq y$: Incumbent firm offers $\pi(x, y', y)$, worker stays with the incumbent firm and her lifetime utility becomes $V(f(x, y'), x, y')$. See Figure 3.8.
- $y' < \tilde{y}(w, x, y)$: The worker ignores the poaching firm, stays with the incumbent firm and her lifetime utility remains $V(w, x, y)$. See Figure 3.9.

Now consider an unemployed worker. If she meets a vacancy, the firm has the all the bargaining power, since there is no other firm to make a counter offer. Hence, the firm offers the wage rate that makes the unemployed worker indifferent. Let

Figure 3.8: Worker x in Firm y Matches with Firm $y' \in [\tilde{y}, y]$.

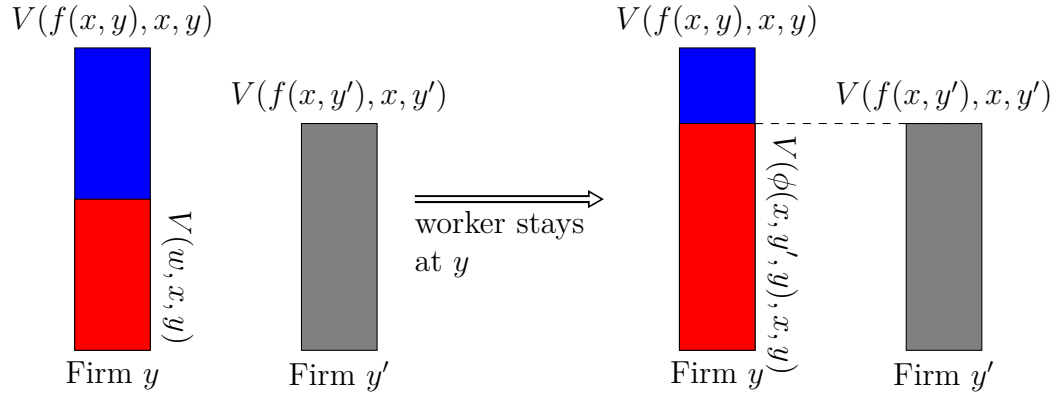
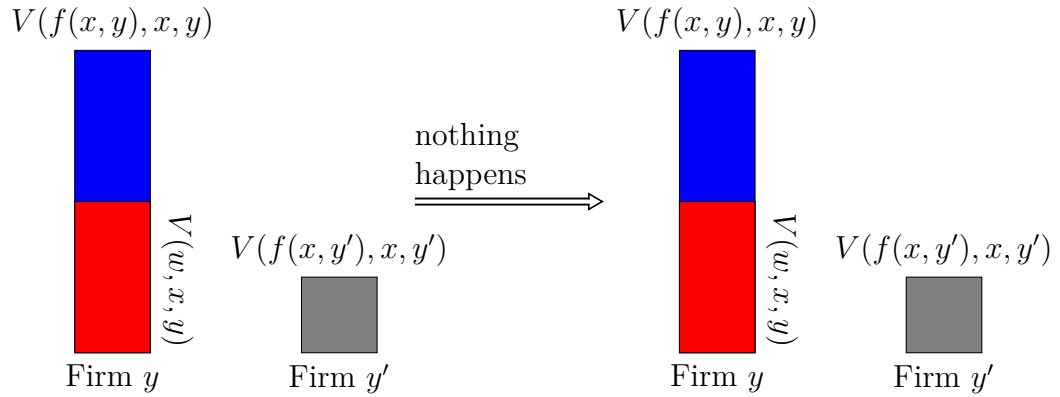


Figure 3.9: Worker x in Firm y Matches with Firm $y' < \tilde{y}$.



$\phi_t(x, 0, y')$ be the wage rate that firm y' offers to unemployed worker. $\phi_t(x, 0, y')$ solves $V_t(\phi_t(x, 0, y'), x, y') = U_t(x)$.

3.E.5 Market Tightness

Let $h(w, x, y)$ be the measure of workers with skill x employed at firm y earning wage w and let $e^*(w, x, y)$ be the optimal search effort. Let $u(x)$ be the measure of unemployed workers with skill x . Lastly, let $v(y)$ be the measure of vacancies posted by firms of type y .

Total search effort in the economy is given by

$$E_t = \int u_t(x)dx + \int \int \int e_t^*(w, x, y)h_t(w, x, y)dw dy dx.$$

Total measure of vacancies in the economy is given by

$$V_t = \int v_t^*(y)dG(y).$$

Then, market tightness is given by

$$\theta_t = V_t/E_t. \tag{3.11}$$

We define the distribution of vacancies as Γ :

$$\Gamma_t(y) \equiv \int^y \frac{v_t(y')}{V_t} dy'.$$

with $\gamma_t(y)$ is the associated density function.

3.E.6 Problem of the Firm

The present value of a filled vacancy by firm of productivity y that employs worker with skill x at wage w is

$$\begin{aligned}
J_t(w, x, y) &= f(x, y) - w \\
&+ \beta \left[(1 - \delta)(1 - e_{t+1}^*(w, x, y)\lambda(\theta_{t+1}))J_{t+1}(w, x, y) \right. \\
&+ (1 - \delta)e_{t+1}^*(w, x, y)\lambda(\theta_{t+1}) \left[\Gamma_{t+1}(\tilde{y}_{t+1}(w, x, y))J_{t+1}(w, x, y) \right. \\
&\left. \left. + \int_{\tilde{y}_{t+1}}^y J_{t+1}(\phi_{t+1}(x, y', y), x, y)\gamma_{t+1}(y')dy' \right] \right].
\end{aligned}$$

Using integration by parts we get

$$\begin{aligned}
J_t(w, x, y) &= f(x, y) - w + \beta(1 - \delta) \left[J_{t+1}(w, x, y) \right. \\
&\left. + e_{t+1}^*(w, x, y)\lambda(\theta_{t+1})J'_{t+1}(\phi_{t+1}(x, y', y), x, y)\gamma(y')dy' \right], \tag{3.12}
\end{aligned}$$

where J'_{t+1} is the derivative of $J_{t+1}(\phi_{t+1}(x, y', y), x, y)$ with respect to y' .

The main decision the firm gives is how many vacancies to post each period:

$$\begin{aligned}
\max_v v\lambda^f(\theta_t) &\left[\int u_t(x)J_t(\phi_t(x, 0, y), x, y)dx \right. \\
&\left. + \int^y \int \int J_t(\phi_t(x, y', y), x, y)h_t(w, x, y')dwdxdy' \right] - \kappa(v)
\end{aligned}$$

A vacancy can be filled by an unemployed worker or an employed worker. The first term inside the bracket is the expected return to vacancy that is filled by an unemployed worker while the second term is the expected return to vacancy filled by an employed worker. A firm with productivity y can hire any worker employed at a

firm with lower productivity $y' < y$ and pays the worker $\phi_t(x, y', y)$.

First order condition with respect to v is

$$\lambda^f(\theta_t) \left[\int u_t(x) J_t(\phi_t(x, 0, y), x, y) dx + \int^y \int \int J_t(\phi_t(x, y', y), x, y) h_t(w, x, y') dw dx dy' \right] = \kappa'(v). \quad (3.13)$$

At an interior optimum, firm equates the marginal cost of opening an extra vacancy to return to vacancy.

3.E.7 Problem of the Worker

Now, we are in a position to define the value function for a worker.

First consider a worker with skill level x employed at firm y and earning w . Suppose she searches for a job with effort level e . The worker gets flow utility of $w - c(e)$ this period. Next period, with probability δ she becomes unemployed and earns lifetime utility of an unemployed worker, $U_{t+1}(x)$. With probability $(1 - \delta)$ she remains employed and searches for a job. For a given effort level e , she does not meet with a firm with probability $1 - \lambda(\theta_{t+1})e$ and her lifetime utility becomes $W_{t+1}(w, x, y)$. With probability $\lambda(\theta_{t+1})e$ she meets with a firm. With probability $1 - \Gamma_{t+1}(y)$ the poaching firm has productivity $y' > y$. In this case the lifetime utility of the worker becomes $W_{t+1}(f(x, y), x, y)$. With probability $\Gamma(\tilde{y}_{t+1})$, the poaching firm has productivity $y' < \tilde{y}$. In this case, the lifetime utility of the worker remains as $W_{t+1}(w, x, y)$. If the poaching firm has productivity $y' \in [\tilde{y}_{t+1}, y]$, then his lifetime utility becomes $W_{t+1}(f(x, y'), x, y')$.

Hence, the lifetime utility of a worker with skill level x employed at firm y at wage

w is

$$\begin{aligned}
W_t(w, x, y) = \max_e w - c(e) + \beta & \left[\delta U_{t+1}(x) \right. \\
& + (1 - \delta)[1 - \lambda(\theta_{t+1})e]W_{t+1}(w, x, y) \\
& + (1 - \delta)\lambda(\theta_{t+1})e \left[(1 - \Gamma_{t+1}(y))W_{t+1}(f(x, y), x, y) \right. \\
& \left. \left. + \int_{\tilde{y}_{t+1}}^y W_{t+1}(f(x, y'), x, y')d\Gamma_{t+1}(y') + \Gamma_{t+1}(\tilde{y}_{t+1})W_{t+1}(w, x, y) \right] \right] \quad (3.14)
\end{aligned}$$

Consider a worker employed at a firm y with wage rate $f(x, y)$. Clearly, she has no gain from matching an outside firm, since no firm offers more than $W_{t+1}(f(x, y), x, y)$. In other words, optimal search effort for such worker is 0. This implies that the lifetime utility for her is

$$W_t(f(x, y), x, y) = f(x, y) + \beta\delta U_{t+1}(x) + \beta(1 - \delta)W_{t+1}(f(x, y), x, y) \quad (3.15)$$

Observe that state variables affect it through the value of unemployment. Since she does not search on the job, market tightness is irrelevant for on the job value. This implies that y only effects it through the production function. Hence, the derivative of $W_t(f(x, y), x, y)$ with respect to y is $f_y(x, y)/[1 - \beta(1 - \delta)]$.

Using integration by part and derivative of $W_{t+1}(f(x, y), x, y)$, the lifetime utility of an employed worker becomes

$$\begin{aligned}
W_t(w, x, y) = \max_e w - c(e) + \beta & \left[\delta U_{t+1}(x) + (1 - \delta) \left[W_{t+1}(w, x, y) \right. \right. \\
& \left. \left. + \lambda(\theta_{t+1})e \int_{\tilde{y}_{t+1}}^y \frac{f_y(x, y)}{1 - \beta(1 - \delta)} [1 - \Gamma_{t+1}(y')]dy' \right] \right] \quad (3.16)
\end{aligned}$$

Taking derivative with respect to e gives us

$$c'(e) = \beta(1 - \delta)\lambda(\theta_{t+1}) \int_{\tilde{y}_{t+1}}^y \frac{f_y(x, y)}{1 - \beta(1 - \delta)} [1 - \Gamma_{t+1}(y')] dy'. \quad (3.17)$$

where the left hand side is the marginal cost of effort. The right hand side is the marginal return to search effort. $(1 - \delta)\lambda(\theta)$ is the increase in the probability of meeting with a firm. The integral is the return to finding a match. β is the discount factor. At the optimal solution, the cost of increasing the search effort should be equal to benefit of increasing the search effort.

Now consider unemployed worker. Unemployed worker has no choice, she searches for a job with effort level 1. In the current period she gets flow utility of unemployment $b(x)$. In the next period, with probability $\lambda(\theta)$ she finds a job. Given the assumption that firms can make take-it-or-leave-it offers to unemployed, finding a job does not increase lifetime utility. Hence, the lifetime utility of an unemployed worker with skill x can be written as:

$$U_t(x) = b(x) + \beta U_{t+1}(x). \quad (3.18)$$

3.E.8 Distribution Accounting

In this section, we derive how the distribution of workers over employment status changes over time.

First, consider distribution of unemployed: $u_t(x)$. $\lambda(\theta_t)$ fraction find a job and leave unemployment. δ fraction of employed workers with skill level x separate from their job and become unemployed. Hence, unemployment distribution evolves ac-

ording to

$$u_{t+1}(x) = u_t(x) - \lambda(\theta_t)u_t(x) + \delta \int \int h(w, x, y)dw dy. \quad (3.19)$$

Similarly, employed distribution evolves according to

$$\begin{aligned} h_{t+1}(w, x, y) = & h_t(w, x, y) - \delta h_t(w, x, y)\lambda(\theta_t)[1 - \Gamma_t(\tilde{y})] \\ & + \int e_t(w', x, \hat{y}_t(x, w))\lambda(\theta_t)\gamma_t(y)h_t(w', x, \hat{y}_t(x, w))dw' \\ & + \int^w e_t(w', x, y)\lambda(\theta_t)\gamma_t(\hat{y}_t(x, w))h_t(w', x, y)dw' \\ & + 1\{w = \phi_t(x, 0, y)\}u_t(x)\gamma_t(y), \end{aligned} \quad (3.20)$$


where $\hat{y}_t(x, w)$ satisfies $\phi_t(x, \hat{y}_t(x, w), y) = w$.

3.E.9 Equilibrium

Definition 4. For given initial distributions $u_0(x)$ and $h_0(w, x, y)$, a competitive equilibrium is a set of value functions $\{U_t(x), W_t(w, x, y), J_t(w, x, y)\}_t$, policy functions $\{e_t^*(w, x, y), v_t^*(y)\}_t$ prices $\{\phi(x, y, y')\}_t$, market tightness $\{\theta_t\}_t$ and distributions $\{u_t(x), h_t(w, x, y)\}_t$ such that

- Value functions solve (3.18), (3.16), (3.12),
- policy functions solve (3.13), (3.17),
- $\phi(x, y, y')$ is the wage rate that solves $W(f(x, y), x, y) = W(w, x, y')$, and $\phi(x, 0, y)$ is the wage rate that solves $U(x) = W(w, x, y)$,
- market tightness is given by (3.11),
- distributions evolve according to (3.19) and (3.20).

Figure 3.10: The Discrepancy Between the MCS Forecast and Realized Inflation The x axis refers to the calendar year. The black line represents the difference between the 1-year ahead MCS forecast and the realized inflation. The values above 1 indicate inflation exceeded forecasts. The red line represents the cumulative real wage loss for a worker who signed his contract two years ago, based on MCS forecasts. The green line represents the cumulative real wage loss for a worker who signed his contract five years ago, based on MCS forecasts.



3.F Additional Analysis

Figure 3.11: Predictive regressions for inflation mistake and acceptance rate The left panel presents the coefficient estimates and the associated 95% CI for β where ‘inflation mistake’ is regressed on the ‘acceptance rate’ with the specification in Equation 3.1. Each point and the bar correspond to an estimate where the regressors are with the associated lag in the x-axis. The right panel provides the same plot where the ‘acceptance rate’ is regressed on the ‘inflation mistake’. See Appendix 3.A for details of the data sources.

