Essays on the U.S. labor market: Evidence from immigrants' wages and the downward nominal wage rigidity constraint

by

Sidonia Lucrisha McKenzie

M.S., University of the West Indies, Mona, Jamaica, 2013

M.A., Kansas State University, 2018

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2024

Abstract

This dissertation consists of three essays that focus on the empirical regularities of the U.S. labor market and their consequences for economic opportunities available to native and immigrant workers. In particular, I reexamine two crucial aspects of the U.S. labor market: the disparity in the labor market outcomes of immigrants and the sluggish adjustment of nominal wages over the business cycle.

The first essay addresses the question of whether various features of an occupation affect the wages of workers in that occupation. I leverage information on the variation in an occupation's English language skill requirement and the ethnic and nativity concentration in a given occupation and year to explain the native-immigrant wage differences within occupations. I show that the fraction of immigrants in an occupation plays an essential role in explaining this wage gap in the U.S. Furthermore, the nature of this fraction, whether immigrants are co-nationals or from a different birthplace, has a salient effect on the magnitude of the gap. Notably, across state labor markets, a one percent increase in the fraction of immigrant men in the workforce of a given occupation leads to an increase in the wage gap in that occupation by about 20 percentage points on average. On the other hand, if this fraction is from the same birthplace, the within-occupation wage gap narrows substantially by nearly 32 percentage points.

In the second essay, I revisit the impact of access to the Supplemental Nutrition Assistance Program (SNAP) on immigrants' labor market outcomes. SNAP stands as one of the largest safety net programs in the United States. Given its significance, the welfare implications of SNAP have been extensively studied from various angles. In particular, applied researchers often focus on the 1996 federal welfare reform, which introduced significant restrictions to SNAP eligibility, creating a clear distinction between citizens (typically the control group) and legal non-citizens (usually the treated group). Subsequent policies that

reinstated eligibility for legal non-citizens at different times across states offer a natural setting for a staggered treatment adoption, which is well-suited for a difference-in-differences (DiD) analysis.

As previous researchers have done, East (2018) employs this methodology using a two-way fixed effects (TWFE) model to examine the effects of SNAP eligibility. However, recent advancements in empirical estimation, often referred to as the "DiD credibility revolution," have highlighted inherent biases in TWFE models, including issues related to "negative weighting" that can lead to misleading and oppositely signed estimates. These more robust methods cast significant doubt on the validity of the effects reported by earlier studies that employed staggered rollout designs.

In this essay, I revisit the effects of SNAP access on immigrants' labor market outcomes, building off the work of East (2018). By employing more recent econometric techniques, I confirm that while SNAP eligibility does indeed reduce immigrants' labor supply—affecting the likelihood of employment—the magnitude of these effects is considerably smaller than previously reported. This finding implies that policy changes impacting SNAP access for vulnerable groups, such as immigrants, should account for the fact that labor supply disincentive effects are relatively modest.

The third essay, with Lance Bachmeier and Benjamin Keen, investigates the response of the labor market to substantial increases in oil prices. Our analysis utilizes microdata from the Current Population Survey (CPS) spanning from 1983 to 2018 to construct a linked sample of hourly paid workers who did not change jobs over the interview cycle. The most important finding, which contributes new insights to the literature, is that overtime compensation serves as a crucial margin for firms to adjust to significant increases in oil prices. Specifically, the likelihood of receiving overtime pay declines by nearly 45% across both blue-collar and white-collar workers due to a reduction in overtime hours. We argue that explanations for the macroeconomic response to oil price shocks may benefit from emphasizing the central role of overtime compensation as a key margin where firms adjust production costs. A second clear takeaway from our analysis is that hours worked fall in most industries in response to large oil price increases. This pattern aligns with the notion that firms are hesitant to reduce wages, opting instead to decrease working hours as a means of cost adjustment.

Essays on the U.S. labor market: Evidence from immigrants' wages and the downward nominal wage rigidity constraint

by

Sidonia Lucrisha McKenzie

M.S., University of the West Indies, Mona, Jamaica, 2013

M.A., Kansas State University, 2018

A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2024

Approved by:

Co-Major Professor
Hugh Cassidy

Approved by:

Co-Major Professor
Steven P. Cassou

Copyright

Abstract

This dissertation consists of three essays that focus on the empirical regularities of the U.S. labor market and their consequences for economic opportunities available to native and immigrant workers. In particular, I reexamine two crucial aspects of the U.S. labor market: the disparity in the labor market outcomes of immigrants and the sluggish adjustment of nominal wages over the business cycle.

The first essay addresses the question of whether various features of an occupation affect the wages of workers in that occupation. I leverage information on the variation in an occupation's English language skill requirement and the ethnic and nativity concentration in a given occupation and year to explain the native-immigrant wage differences within occupations. I show that the fraction of immigrants in an occupation plays an essential role in explaining this wage gap in the U.S. Furthermore, the nature of this fraction, whether immigrants are co-nationals or from a different birthplace, has a salient effect on the magnitude of the gap. Notably, across state labor markets, a one percent increase in the fraction of immigrant men in the workforce of a given occupation leads to an increase in the wage gap in that occupation by about 20 percentage points on average. On the other hand, if this fraction is from the same birthplace, the within-occupation wage gap narrows substantially by nearly 32 percentage points.

In the second essay, I revisit the impact of access to the Supplemental Nutrition Assistance Program (SNAP) on immigrants' labor market outcomes. SNAP stands as one of the largest safety net programs in the United States. Given its significance, the welfare implications of SNAP have been extensively studied from various angles. In particular, applied researchers often focus on the 1996 federal welfare reform, which introduced significant restrictions to SNAP eligibility, creating a clear distinction between citizens (typically the control group) and legal non-citizens (usually the treated group). Subsequent policies that

reinstated eligibility for legal non-citizens at different times across states offer a natural setting for a staggered treatment adoption, which is well-suited for a difference-in-differences (DiD) analysis.

As previous researchers have done, East (2018) employs this methodology using a two-way fixed effects (TWFE) model to examine the effects of SNAP eligibility. However, recent advancements in empirical estimation, often referred to as the "DiD credibility revolution," have highlighted inherent biases in TWFE models, including issues related to "negative weighting" that can lead to misleading and oppositely signed estimates. These more robust methods cast significant doubt on the validity of the effects reported by earlier studies that employed staggered rollout designs.

In this essay, I revisit the effects of SNAP access on immigrants' labor market outcomes, building off the work of East (2018). By employing more recent econometric techniques, I confirm that while SNAP eligibility does indeed reduce immigrants' labor supply—affecting the likelihood of employment—the magnitude of these effects is considerably smaller than previously reported. This finding implies that policy changes impacting SNAP access for vulnerable groups, such as immigrants, should account for the fact that labor supply disincentive effects are relatively modest.

The third essay, with Lance Bachmeier and Benjamin Keen, investigates the response of the labor market to substantial increases in oil prices. Our analysis utilizes microdata from the Current Population Survey (CPS) spanning from 1983 to 2018 to construct a linked sample of hourly paid workers who did not change jobs over the interview cycle. The most important finding, which contributes new insights to the literature, is that overtime compensation serves as a crucial margin for firms to adjust to significant increases in oil prices. Specifically, the likelihood of receiving overtime pay declines by nearly 45% across both blue-collar and white-collar workers due to a reduction in overtime hours. We argue that explanations for the macroeconomic response to oil price shocks may benefit from emphasizing the central role of overtime compensation as a key margin where firms adjust production costs. A second clear takeaway from our analysis is that hours worked fall in most industries in response to large oil price increases. This pattern aligns with the notion that firms are hesitant to reduce wages, opting instead to decrease working hours as a means of cost adjustment.

Table of Contents

List of Figures

List of Tables

Acknowledgments

This dissertation reflects the diligence, thoughtful guidance, encouragement, and dedication of several individuals and groups; without their support, this work would not have been possible.

Firstly, I am extraordinarily grateful to my dissertation committee for their patience and support, allowing me to acclimate to a new job while completing my dissertation. In particular, I am deeply indebted to my committee co-chair, Hugh Cassidy, for his insightful feedback, encouragement, and unwavering support throughout my Ph.D. journey. He consistently made time to meet with me, even during the pandemic, and offered detailed explanations to enhance my technical skills. I am a better researcher because of his guidance. I also thank Steve Cassou for his invaluable advice during the early stages of my dissertation and for being an outstanding instructor and researcher. His high standards and passion for understanding macroeconomic fluctuations have profoundly influenced my research and teaching practices.

I extend special thanks to the other members of my committee, who contributed significantly to this work. Valerie Bostwick, whom I have admired since her job market presentation at Kansas State, has significantly enhanced my empirical skills through her thoughtful guidance and support. I also gratefully acknowledge Lance Bachmeier for his invaluable guidance in developing my skills as a scholar and for showing me how to mentor other students. Our collaborations have been immensely rewarding, and I deeply appreciate his ongoing support. I look forward to future opportunities to work together. Benjamin Keen, who has become my unofficial external committee member, deserves recognition for our many discussions on the broader implications of macroeconomic models. His insights have been both intellectually stimulating and encouraging.

In addition to my committee, I have benefited from the expert guidance of and learned from several faculty members including Philip Gayle, William Blankenau, Yoon-Jin Lee,

Yang-Ming Chang, Ross Milton, Dan Kuester, Vincent Amanor-Boadu, Jason Bergtold, and Nathan Hendricks. Thank you for creating stimulating learning environments and for always finding time to address my numerous questions. To the extent that I am regarded as a good instructor, I owe it to the exceptional examples you have provided. To the administrative staff at K-State, notably Susan Koch, Crystal Strauss, Judy Duryee, Deana Foster, and Angie Pfizenmaier, thank you for your encouraging words and for consistently going beyond the call of duty to assist students.

I owe a debt of gratitude to Jim and Kathy Haymaker, generous supporters of the Wayne Nafziger Scholarship and champions of advancing economic knowledge at Kansas State University. I have not met a kinder, more thoughtful couple than the Haymakers. I deeply appreciated our correspondence and their support throughout my academic journey.

My graduate school experience was much more meaningful because of the tremendous support of my friends and family. I am immensely blessed by the steadfast love, prayers, and support of my parents, siblings, friends, and Church community, especially my mother, my youngest sister Paige McKenzie, Pastor Sterling Hudgins, Pastor Jason Anderson, and Rose Ndegwa. I am profoundly grateful for your advice and encouragement, and I cherish our many fond memories, adventures, and, of course, the potlucks.

Finally, I must thank those who played a crucial role at the outset of this academic journey: Abdullahi Abdulkadri, Peter-John Gordon, Nadine McCloud, Christine Clarke, Patrice Whitely, and Calvin McDonald. Your encouragement ignited my pursuit of a PhD in Economics and reminded me that intelligence plus character is the true goal of education.

Dedication

To my family, may this inspire you to never give up on your dreams.

To Adonai, the shield around me, my glory, and the One who lifts my head high, you get the glory from this.

Chapter 1

Explaining the intra-occupational native-immigrant wage gap in the U.S.

1.1 Introduction

How immigrants perform in the U.S. labor market has long been a concern for policymakers. The economic assimilation of immigrants is typically measured by the rate at which their wages converge to those of native-born citizens. Previous studies have largely attributed the extent of this convergence to various human capital-driven factors, including cohort quality (Borjas, 2015, 1985); countries of origin (LaLonde and Topel, 1992); job-education mismatch (Chiswick and Miller, 2009b); and age-at-arrival effects (Bleakley and Chin, 2004; Cassidy, 2019). However, they have not adequately addressed the role of occupation-specific factors.

To the best of my knowledge, this paper provides the first quantitative assessment of how specific features of an occupation contribute to understanding the dynamics between the structure of the occupation and the wages of immigrant workers within that occupation. A related study by Goldin (2014) explores a similar issue by examining the degree of (non)linearity in wages concerning the time flexibility of an occupation to explain the gender wage gap. In a comparable approach, I utilize information on the variation in the level of English proficiency required for each occupation, as well as the fraction of immigrants in a given occupation and year, and the role of immigrant networks, to explain the observed intra-occupational native-immigrant wage gap.

I begin by explaining the rationale behind selecting these specific occupational features, drawing on insights from the existing literature. English language proficiency is widely recognized as a crucial form of human capital integral to the assimilation of immigrants (Chiswick and Miller, 1995, 2012; Ferrer et al., 2006; Imai et al., 2019; Lazear, 1999). Indeed, Chiswick and Miller (2012) highlight that for immigrants, language skills and wages are closely interconnected. To assess this, I use skill data from the Occupational Information Network (O*NET) to construct a measure of English language skill requirements for each occupation. Additionally, I create a measure of the concentration of immigrants within each occupation to consider the idea that immigrants may compete more directly with fellow immigrants than with natives.

Peri and Sparber (2009) demonstrate that immigrants tend to sort into occupations where they possess a comparative advantage, which in turn affects their wages.[1] This comparative advantage extends beyond skills and location decisions to include the level of competition among immigrants. The degree of substitutability between the skills of immigrant workers within the same occupation is significantly higher compared to that between observationally similar natives in identical roles. In short, immigrants are closer substitutes for each other than for natives within the same occupations, indicating substantial labor market competition among immigrants within any given occupation. Consequently, when an exogenous factor increases the share of immigrants in a particular occupation (excluding the self-selection effect), it tends to drive down wages for immigrants in that occupation. This dynamic, in turn, widens the wage gap between natives and immigrants within that occupation.

Furthermore, non-pecuniary compensation costs may also depress immigrant wages relative to observationally equivalent natives. Immigrants might accept lower wages in occu-

[1] For instance, low-skilled immigrants often pursue more manual or physically intensive occupations, whereas observationally similar natives may choose more interactive or language-intensive roles.

pations where they are surrounded by individuals from the same country of origin or ethnic network, even if they are equally skilled for higher-paying positions elsewhere. Conversely, ethnic networks can provide valuable information about occupations, potentially enhancing the match quality between workers and jobs, which could increase wages (Damm, 2009). To account for these effects, I include a measure of the concentration of immigrants from the same birthplace within each occupation.

I exploit variations in the English language skill requirements and the share of immigrants (both co-nationals and those from different birthplaces) within occupations over time to identify the impact of these occupation-specific characteristics on wages. The two Herfindahl-type immigration indices gauge the concentration of immigrants within an occupation, which is crucial considering the endogenous occupational adjustment of natives to immigration (Llull, 2017), which impacts the magnitude of bias arising when such labor market adjustments are overlooked. To account for endogeneity bias, I use the residual wage difference approach to estimate the intra-occupation wage gap between natives and immigrants. This approach is implemented in two stages: the first stage regression controls for human capital characteristics, geographic, and time fixed effects, while the second stage incorporates occupation-specific controls. This first stage follows the residual approach in the wage decomposition literature, controlling for observable human capital characteristics associated with an occupation (Goldin, 2014; Sharpe and Bollinger, 2020). In essence, the assumption is that individuals sharing identical occupation codes are expected to possess comparable skills.[2] Therefore, any observed wage differences are potentially attributed to occupation-specific features or other unobserved factors, or both. I present results separately for each year, as well as within year-state and year-metropolitan area cells.

Using data from the decennial Census and the American Community Survey, the results reveal compelling heterogeneous effects of the intra-occupation wage penalty in response to changes in occupation-specific characteristics across different genders and labor markets. Notably, the wage penalty faced by immigrant men(women) increases(decreases) with the share

[2] This assumption is not exclusive to the gender wage gap or immigration literature; sociologists have long regarded occupations as collections of diverse roles, each defined by tasks, necessary skills, and responsibilities that remain consistent across the spectrum (Hauser and Warren, 1997; Xie, 2016).

of immigrant workers in a given occupation and year, with the effects being more pronounced in state and metropolitan labor markets. The concentration of co-national immigrants in identical occupations significantly reduces the wage penalty, with a more substantial effect for immigrant men compared to women. These results suggest that the wage differential between natives and immigrants in the same occupations within a given year is influenced by the ethnic and nativity composition of the workforce. This finding contrasts with the gender wage gap literature, where Goldin (2014) asserts that changing the gender distribution of workers within occupations is insufficient to "close" the wage gap. Nevertheless, the occupation-specific features examined in this paper seem to more consistently explain the wage differential over time for immigrant men compared to women. This opens the possibility for further exploration of other gender-specific characteristics of occupations.

The paper is organized as follows: Section 2 reviews the relevant literature on the native-immigrant wage gap in the U.S. and highlights the contributions of this study. In Section 3, I describe the data sources, introduce the immigrant diversity indices, and present descriptive statistics for selected human capital variables among both natives and immigrants. Section 4 outlines the empirical strategy and details the alternative wage function specifications. The main results of the analysis are presented in Section 5, and the paper concludes with Section 6.

1.2 Background

1.2.1 Literature Review

The native-immigrant wage gap in the U.S. has been widely studied across many different dimensions. In the early literature on immigrants' post-migration economic adjustment, the human-capital wage framework was the leading paradigm (Borjas, 1995, 1985; Chiswick, 1978; Gabriel and Schmitz, 1987). In this strand of literature, immigrants' labor market outcomes were largely attributed to an assimilation process where differences in immigrant cohort quality (Borjas, 1985), immigrants' countries of origin (LaLonde and Topel, 1992), educational quality and returns to foreign schooling (Bratsberg and Terrell, 2002; Dustmann et al., 2016) and differences in language ability (Borjas, 1995; Chiswick and Miller, 1995) play a significant role in explaining wage differentials.[3] Recently, it turns out that selective out-migration (Lubotsky, 2007) and age-at-arrival effects (Bleakley and Chin, 2004; Cassidy, 2019; Friedberg, 1992) also explain the wage progress of immigrants.

Another strand of literature moves away from the human-capital paradigm and focuses on the occupational choice of immigrants. Arguably, this shift is partly influenced by studies in the gender wage gap literature, which largely document that human capital variables such as education and experience have played a diminished role in explaining the gender wage differential over time.[4,5] Indeed, more recent studies in the immigration literature have challenged the human capital approach, presenting evidence that education is not always a perfect proxy for skill when evaluating the labor market outcomes of immigrants. These studies have found that both the quality of education and the returns to foreign schooling for immigrants are diminished in the labor markets of host countries (Bratsberg and Terrell, 2002; Dustmann et al., 2016). Additionally, newly arrived immigrants are more likely to be overeducated compared to natives in the jobs they hold (Chiswick and Miller, 2009c).

It turns out that using occupations as a proxy for skill provides a more robust approach for analyzing wage differentials. Research on the gender wage gap (Bizopoulou, 2019; Fortin and Huberman, 2002; Goldin, 2014) has revisited the taxonomy proposed by Dolton and Kidd (1994), emphasizing the relative importance of inter-occupation versus intra-occupation effects. These studies find that gender wage differentials *within* occupations are greater than those observed *across* occupations. My paper builds on this literature, arguing that a similar distinction—*within* versus *across* occupations—is particularly relevant for understanding wage differentials between natives and immigrants in the U.S.

[3]Studies have shown that discrimination, a topic that I do not explore in this paper, plays a significant role in immigrants' integration in the host country labor market (Dechief and Oreopoulos, 2012; Oreopoulos, 2011).

[4]See (Blau and Kahn, 2017) for a comprehensive overview of this literature.

[5]This partly reflects the significant advancement women have made in their educational attainment relative to men as well as increased job experience due to higher labor force participation (Goldin, 2014). In fact, (Blau et al., 2013) found that reductions in occupational segregation were the largest among college-educated individuals.

Studies that utilize this dichotomy are relatively sparse in the immigration literature. Moreover, they primarily document the relative extent of the *intra*- and *inter*-occupational wage gaps without investigating the underlying determinants of these differences. Notably, Liu et al. (2004) was among the first to provide evidence on intra- and inter-occupational earnings differentials between immigrants and natives using data from the Hong Kong 1996 Census. They employed the methodology developed by Brown et al. (1980), which facilitates the direct estimation of occupational choices and the measurement of earnings differences within occupations. Their findings indicate that the intra-occupational earnings differential is slightly more significant than the inter-occupational earnings differential.

Similarly, Demoussis et al. (2010) applied the Brown et al. (1980) methodology to examine *inter*- and *intra*-occupational wage differentials using data from the Greek Household Budget Survey (GHBS) between 2004 and 2005. The authors found that less than half of the average wage gap could not be explained by differences in observed characteristics. Furthermore, they noted that the larger component of this unexplained gap was attributed to *inter*-occupational wage differences. This implies that asymmetric occupational access, rather than unequal treatment within occupations, is a key factor in explaining the wage differential between native and immigrant workers in Greece.

My analysis extends the existing literature in two ways. First, my empirical approach enables me to decompose the native-immigrant wage differential while maintaining a detailed classification of occupations. A limitation of the methodology employed by Brown et al. (1980) is its aggregation of occupations into relatively broad categories, which obscures valuable occupational information. Natives and immigrants often exhibit significant differences in their occupational distributions within these broad categories. Consequently, aggregating occupations at such broad levels may overlook crucial patterns that could shed light on the wage gap. By using a residual approach and pooling data from several Census years, I address this issue effectively. This allows me to examine *intra*-occupational wage differences across over 300 occupational codes. Additionally, I improve upon previous studies by exploring the factors that influence the relative magnitudes of the intra-occupational wage gap.

Goldin (2014) utilizes this approach by highlighting that the remaining gender gap is related to the time demands of an occupation and how firms choose to reward individuals based on these demands. Using data from O*NET, Goldin (2014) constructs a composite index of characteristics reflecting the time flexibility of an occupation.[6] Interestingly, occupations with less time flexibility, such as those in business and law, exhibit a larger gender wage gap. Methodologically, my paper closely follows Goldin (2014). I apply her approach to decompose wage differentials between natives and immigrants, examining the extent to which these differentials can be attributed to variations within occupations versus those across occupations.

To understand the extent of wage gaps over time, Figure 1.1 displays the log weekly wage differential between immigrants and observationally equivalent natives in a given year. The height of the bars represents the unconditional wage gap, which can be adjusted for differences in workers' human capital endowments such as educational attainment, English proficiency, and occupational choice. This adjustment, conventionally referred to as the "inter-occupation" wage gap (Chiswick and Miller, 2009a), highlights differences in wages attributed to discrete occupational preferences. The residual fraction of the bars, conversely, shows the wage difference within specific occupations, the "intra-occupation" wage gap. In cases where this differential is negative, it reflects the wage disadvantage (penalty) faced by an immigrant within the same occupational domain as their native counterparts.

Over the sample period, the average unconditional gap in the weekly log wage of immigrants is substantial. On average, male and female immigrant workers experience wages that are 23% and 13% lower than their native-born counterparts, respectively. Human capital characteristics and occupational choice contribute to explaining approximately 44% and 18% of this wage gap for immigrant men and women, respectively. [7] It is important to note that, while these factors account for a substantial portion of the wage gap, a significant differen-

[6] The characteristics include *time pressure*, *contact with others*, *establishing and maintaining interpersonal relationships*, *structured vs. unstructured work*, and *freedom to make decisions*.

[7] More specifically, the unadjusted native-immigrant gap in weekly wages fluctuated between 18%(1%) in 1980 and 32%(16%) in 2000 before declining(increasing) to 21%(22%) in 2018 for immigrant men(women). Human capital factors accounted for 53%(34%) of the unadjusted gap in 1980 and just 37%(30%) of the gap in 2018 for immigrant men (women).

Figure 1.1: Trend in native-immigrant wage gap, 1980-2018

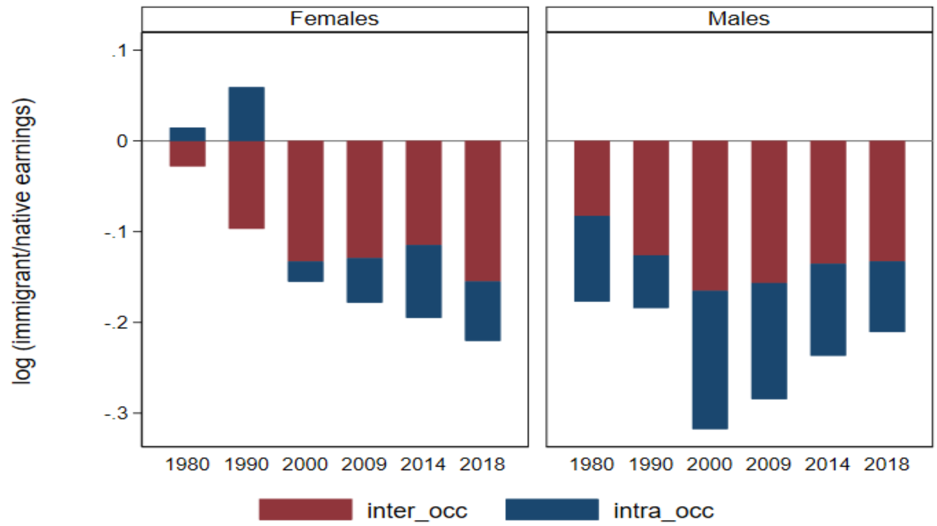

Sources: Data are from the 1980, 1990, and 2000 IPUMS U.S. Census and ACS five-year pooled samples: 2005-2009 (for 2009), 2010-2014 (for 2014), and 2014-2018 (for 2018).

Notes: The figure shows the log weekly wage differential between natives and immigrants calculated separately by each Census year using a Mincerian wage regression and a "basic" regression. The Mincer regression includes controls for a cubic function of age, dummy variables for education and occupation levels, and an interaction term of immigrant and occupation variables. The coefficient on this interaction term represents the intra-occupational wage gap, depicted by the red bars. The "basic" regression includes the immigrant dummy variable as the only covariate and represents the unconditional wage penalty of immigrants, as shown by the height of the bars. The blue portion illustrates the wage difference between natives and immigrants employed in different occupations. The sample is employed adults aged 25-64. All regressions are weighted by the Census-provided weights.

tial remains. Even after controlling for human capital traits and occupational variables, a notable wage gap persists between natives and immigrants within the same occupation.

The insights drawn from Figure 1.1 indicate a widening trend in within-occupation wage disparities over time, particularly for men, despite adjusting for occupational factors. This prompts an exploration into the causes of this growing wage gap within occupations and the potential relevance of various occupational characteristics. While Goldin (2014) has explored the impact of (non)linearity in wages due to time flexibility as an explanation for the gender wage gap, I exploit variations in the level of English proficiency required for each occupation, the fraction of immigrants in each occupation and year, and the role of immigrant networks to explain the observed residual intra-occupational native-immigrant wage gap.

A select number of related studies employ occupations as a proxy for skill and leverage the variation within occupation groups to examine the impact of immigration on wages in the U.S. (Camarota, 1997; Card, 2001; Orrenius and Zavodny, 2007; Sharpe and Bollinger, 2020). Empirically, this branch of research can be broadly categorized into two groups: studies utilizing the "*national skill-cell approach*" and those adopting the "*mixture approach*", as termed by Dustmann et al. (2016). The fundamental distinction between these methodologies lies in the unit of analysis and, consequently, the source of variability. Sharpe and Bollinger (2020) utilize the national skill-cell approach, demonstrating that the partial equilibrium impact of immigration on native wages exhibits greater magnitude when the labor market is stratified by occupation-experience cells, in contrast to the impact when education-experience is employed, as in the case of Borjas (2003). Card (2001) and Orrenius and Zavodny (2007) adopt the mixture approach, leveraging variation within occupational groups and across sub-labor markets, specifically cities and metropolitan areas, respectively. As discussed below, a common theme across these studies is that occupational characteristics offer a more nuanced perspective for assessing the relative wage effects of immigration. However, the question of how these occupational features influence wages within specific occupations remains underexplored. This paper contributes to the literature by leveraging both specifications to present a multi-dimensional framework that explains the intra-occupational wage differentials between natives and immigrants.

1.2.2 Expected Effects

In modeling immigrant network effects, I employ the conceptual framework of Patel and Vella (2013) regarding the impact of ethnic clustering in an occupation on the intra-occupational wage gap. The prediction of their model is twofold: first, an immigrant's probability to sort in an occupation is a function of the share of their ethnic network in that occupation; and second, the accepted wage by an immigrant is positively associated with the share of their ethnic network in the occupation. These assumptions imply that three possible effects might be at play to explicate the impact of ethnic networks on the intra-occupational wage differential. They are given as follows:

1. *Competition of immigrants.* Immigrants in identical occupations are more substitutable with other immigrants than natives in a given occupation. If this assumption holds among immigrants within an occupation, then the fraction of immigrants captures the competition among immigrants in an occupation/year cell. An increase in this fraction, due to larger inflows of immigrants in a given occupation (or area), raises the relative labor supply of immigrants, which should lower wages in that occupation (or area) (Orrenius and Zavodny, 2007). To capture this effect, I construct the variable, `frac_immig`, as a ratio of the immigrant size in a given occupation and year to the total employment in that occupation and year.

2. *Immigrant Occupation Selection.* It is well-documented that the selection of immigrants into specific occupations, their geographical location, and their occupational choices are not random. Immigrants tend to sort into occupations where they have a comparative advantage in terms of required skills or compensation relative to other occupations. Specifically, Peri and Sparber (2009) demonstrates that immigrants are more likely to pursue manual or physically intensive occupations, while natives tend to occupy roles that are more interactive or language-intensive. In a wage regression, this effect would be reflected by a positive coefficient on the fraction of immigrants, indicating a wage "premium" that immigrants might earn by excelling in occupations suited to their skills. Controlling for occupation should largely account for this effect, assuming that the distribution of occupations remains relatively stable across years.

3. *Ethnic enclaves.* Relative to the national level analysis, the effect of ethnic clustering may be more apparent in the local labor market model because of geographic proximity. Patel and Vella (2013) find that foreign-born workers from the same birthplace and located within the same area (metropolitan area or state, by extension) may have a greater probability of forming networks. These networks provide information regarding occupations that can increase occupation-worker match quality, improving wage

(Damm, 2009). In this respect, the effect of ethnic network size on immigrants' wages can be positive. On the other hand, the effect could be harmful if immigrants desire to work more closely in the same occupation as members of their ethnic enclaves and, therefore, are willing to forgo higher wages in other occupations. Indeed, Bentolila et al. (2010) find that social networks can create a mismatch between workers' occupational choices and their abilities. Notably, some workers may be encouraged to undertake common occupations in their social networks but do not fully exploit their productive skills. Overall, the arguments presented here show that the coefficient on the measure of ethnic networks cannot be unambiguously signed *a priori*.

I show that the concentration of immigrants in an occupation is a central feature of the occupation that is not adequately addressed in the literature. In this respect, my paper complements the sparse strand of literature examining immigrants' occupational allocation. Notably, relative to Patel and Vella (2013), I explore whether the concentration of immigrants from the same birthplace matter above the overall fraction of immigrants in an occupation. Finally, it turns out that the inclusion of the overall fraction of immigrants in an occupation explains the wage differential within occupations and influences the effect of ethnic networks on the observed intra-occupational wage gap.

1.3 Data

1.3.1 Decennial Census and American Community Survey

For my analysis, I utilize the U.S. decennial Census data for the years 1980, 1990, and 2000 as well as the ACS five-year pooled samples: 2005-2009 (2009), 2010-2014 (2014), and 2014-2018 (2018) (Ruggles et al., 2020).[8] The Census dataset is nationally representative and provides a comprehensive source of information on the U.S. population, which includes a

[8]Data was retrieved from the Integrated Public Use Microdata Series (IPUMS) website. The 5% state samples (1-in-20 random sample of the national population) were used for each dataset. The English proficiency variable, `speakeng` is not available in the 1970 decennial Census and therefore excluded from my analysis.

suite of household and individual demographic characteristics such as citizenship status and labor market variables, which are central to my analysis. I focus only on the prime-aged working population. I separately examine men and women aged 25-64 who report positive weekly wages, hours worked, and weeks of work during the survey year.[9] I further restrict the sample to exclude students, military personnel, self-employed, and individuals living in group quarters.

Immigrants and natives are categorized based on the survey's citizenship status question. I define *immigrants* as those who are either non-citizens or naturalized citizens and all other persons as *natives*. Following Borjas (2015), I consider those with at least one year of potential labor market experience but no more than 40 years.[10,11] In addition, I restrict the sample to immigrants who migrated as adults to avoid possible confounding effects of childhood immigrants (Bleakley and Chin, 2004).

The primary variable of interest is weekly log wages, which is obtained by taking the ratio of earned income to weeks worked.[12] It is instructive to note that on average, since 1980, 12 percent of earned income has been imputed. Although this ratio is relatively low, Hirsch and Schumacher (2004) and Bollinger and Hirsch (2006) established that imputation methods used to generate earnings data for non-respondents in surveys like the Census and CPS lead to biased coefficients as a result of imperfect matching of donors to non-respondents (match bias). In fact, Hersch and Shinall (2018) document that match characteristics for earnings imputations in the Census and ACS do not include citizenship status or nativity, which is particularly important for immigrants. As such, I exclude imputed wages from the

[9] The reference period refers to the past 12 months for decennial Census and ACS survey years.

[10] Labor market experience is approximately defined as the difference between a respondent's age in the survey year and their age of labor market entry. Following Borjas (2015), I determine the age of labor market entry based on years of schooling. In short, I assume that high school dropouts enter the labor market at age 17, high school graduates at 19, persons with some college education at 21, college graduates at 23, and workers with more than a college degree at 24.

[11] This cutoff is based on immigrant economic assimilation literature, which shows that immigrants with more than 40 years of U.S. labor market experience are almost identical to the U.S. born citizens in terms of wage growth (Borjas, 1994, 1985; Chiswick, 1978).

[12] The IPUMS variable `incwage` was used, which reports a worker's total pre-tax wage and salary income. Log weekly wage was also used as the dependent variable in Borjas (2015); Patel and Vella (2013); Sharpe and Bollinger (2020) while Borjas (1994); Cassidy (2020); Hersch and Shinall (2018) used log of weekly earnings. The results are qualitatively similar when a log of weekly earnings is used. These results are available upon request.

analysis.[13]

Given the wide scope of the Census survey years used for my analysis, attention must be paid to the occupational classification schemes. Since the occupation classification scheme has changed over time, I use the time-consistent occupational classification developed by Autor and Dorn (2013), which is similar to the `occ1990` variable available in the IPUMS. Specifically, I use their crosswalk file, which contains the 2000 Census-defined occupation codes, `occ`, and their consistent occupation codes, `occ1990dd`, to match the occupation codes in the Census data. I utilize the `occ1990dd` in all subsequent analysis.

1.3.2 Occupational Information Network (O*NET)

A distinguishing feature of this paper is that I control for occupational characteristics in the wage equation by focusing on three central attributes of the occupation. One of these attributes is required English language proficiency. Given the link between language skills and occupational attainment (Imai et al., 2019) and the advantages to dominant-language proficiency (Chiswick and Miller, 1995, 2012; Ferrer et al., 2006), I source skill data from the O*NET that summarizes the degree of knowledge of the English language required to perform each occupation. The O*NET produces a knowledge questionnaire about the work-related areas of knowledge in one's current job.[14] For each knowledge area, two questions are asked based on level and importance. For example, English language is a knowledge area, and the precise question asked on level is "What *level* of English language knowledge is needed to perform your current job", which is measured on a scale of 1-7.[15] I use this attribute as my main variable of interest to capture an occupation's English language skill requirement. I also control for English language proficiency from the Census, which is a self-reported rating on language fluency in English.[16]

[13] The results are relatively sensitive to the inclusion of imputed wages.

[14] The O*NET defines knowledge areas as the "sets of facts and principles needed to deal with problems and issues that are part of a job."

[15] The question related to knowledge asks "How important is knowledge of the English language to the performance of your current job?" and is rated from *Not important* to *Very important*.

[16] The Census asks respondents "How well does this person speak English?", based on ratings "Not at all [0]", "Not very well [1]", "Well [2]" and "Very well [3]".

1.3.3 Descriptive Statistics

To illustrate the extent of the native-immigrant wage gap in the U.S. over the sample period, one could employ the standard Mincerian wage-generating function, pioneered by Mincer (1974) and Chiswick (1978). This model typically relates the logarithm of annual (or weekly) wages to several human capital variables. In particular, years of schooling completed, age, year since migration, English proficiency, birthplace, potential labor market experience, weeks worked, and a set of indicator variables for marital status, race, location, and citizen status. Table 1.1 presents the sample statistics for a few of these covariates separately for the native and immigrant sub-samples and each survey year. The most apparent statistic is the average weekly wage difference between immigrants and natives over the sample period. On average, natives earn wages that are 18% (0.13 log points) higher than immigrants.

On average, natives' years of schooling exceed that of immigrants by roughly one year. This difference was much higher at the beginning of the sample period but declined sharply after the 2000 Census. This convergence in years of schooling between natives and immigrants coincides with the significant legislative increase in the annual Cap on H-1B visas.[17] In other socioeconomic characteristics, such as years of experience, age, and weeks worked, immigrants are slightly more comparable to natives, particularly after 2000. Finally, immigrants are more likely to marry than natives in the sample.

1.4 Empirical Strategy

Occupational choice is essential when explaining wage differentials. Previous empirical work has adopted various approaches when incorporating occupations in wage analyses. For example, the unit of analysis used in these studies is either at the national level (Borjas, 2003; Borjas et al., 1997; Sharpe and Bollinger, 2020) or metro area/city-level (Card, 2001; Orrenius and Zavodny, 2007; Patel and Vella, 2013). This is important to highlight due to the potential problems when the labor market is stratified. Notably, endogeneity concerns

[17]Notably, in 1998 and 2000, Congress instituted a piece of legislation that sharply raised the annual cap from 65,000 to 195,000 visas for five years (Kerr and Lincoln, 2010).

and occupational selectivity bias occur when the labor market is segregated geographically and by occupation. The main methods for addressing these concerns are the national labor market approach, first proposed by Borjas et al. (1997), and the Brown methodology, which explicitly models occupational choice.

Immigrants and natives are allocated differently across the distribution of occupations, which may be attributed to various reasons other than their observed characteristics.[18] Therefore, the occupational choice is likely not random (Demoussis et al., 2010; Liu et al., 2004; Sharpe and Bollinger, 2020). As such, the estimated wage differential will be biased if the observed difference in occupational distribution between immigrants and natives is influenced by unobserved factors and not accounted for in the regression. For this reason, the Blinder-Oaxaca decomposition, with occupational dummies included as covariates, is not appropriate because of the implicit assumption that occupational choice is exogenous (Demoussis et al., 2010). Liu et al. (2004) note that the Brown methodology avoids this problem by providing a direct approach for predicting the occupational structure for immigrants that would have occurred in the absence of discrimination.

Despite the clear advantage of this approach, one weakness is that the model can only facilitate broadly defined occupational groups.[19] Therefore, if a researcher is interested in inference about the occupation wage gap at a more granular level, the Brown methodology would be limited in this sense.

Since immigrants are generally attracted to areas with favorable demand conditions, the national labor market approach provides a richer treatment of this selection problem by treating the U.S. as one national labor market.

[18] Discrimination is the most widely cited reason in the literature. However, other reasons such as licensing can affect the occupational distribution of natives and immigrants. Notably, Cassidy and Dacass (2019) document that immigrants are less likely to possess an occupational license relative to an observationally identical native.

[19] In other words, each occupation group needs a sufficient number of observations to be estimated.

Table 1.1: Summary statistics by Census year

	1980	1990	2000	2009	2014	2018
Panel A: Natives						
Log weekly wage	5.474	6.067	6.378	6.645	6.697	6.796
Years of schooling	15.155	16.100	16.272	16.526	16.666	16.738
Years of experience	19.002	17.962	19.727	21.703	22.082	21.788
Age	40.156	40.062	41.998	44.229	44.748	44.526
Weeks worked	47.081	47.360	47.189	46.765	47.272	47.855
Married	0.782	0.737	0.698	0.698	0.675	0.656
Observations	2,056,884	2,982,269	3,356,793	3,991,234	3,658,111	2,764,532
Panel B: Immigrants						
Log annual wage	5.355	5.985	6.222	6.497	6.559	6.645
Years of schooling	13.771	14.816	14.941	15.759	15.900	15.946
Years of experience	21.977	20.157	20.367	21.620	22.723	23.509
Age	41.749	40.973	41.308	43.378	44.622	45.455
Weeks worked	45.881	46.093	45.349	46.233	46.853	47.229
Married	0.817	0.786	0.797	0.791	0.781	0.779
Average English proficiency	2.214	2.247	2.092	2.124	2.172	2.199
Average years since migration	12.367	13.060	14.528	17.114	18.415	19.074
Observations	121,475	211,495	354,684	427,976	426,873	315,391

Notes: Data are from the 1980, 1990 and 2000 IPUMS U.S. Census and ACS five-year pooled samples: 2005-2009 (for 2009), 2010-2014 (for 2014) and 2014-2018 (for 2018). The sample is employed adults aged 25-64.

In the following empirical analyses, I strive to explain the within-occupation wage differences between natives and immigrants by examining various occupation characteristics. Following the methodologies of Borjas (2003) and Sharpe and Bollinger (2020), I use national-level data to derive a measure of the intra-occupational wage gap between natives and immigrants, which varies over time by pooling multiple Census samples. The advantage of using several survey years is twofold: first, I can track changes over time, and second, I can take advantage of the variation in the wage penalty within occupation but across years, which allows me to control for occupation fixed effects. One contribution of this paper is that the wage gap is finely partitioned across the full spectrum of occupational codes. As noted above, existing studies have only been able to stratify occupations into broad groups, which may undermine important intra-occupational wage trends that can only be observed at a disaggregated level.

1.4.1 Modelling intra-occupational wage

I adapt the methodology used by (Goldin, 2014) in the gender wage gap literature to compute the residual wage as follows:

$$ln(wage_{i,o,t}) = \lambda + \beta X_{i,t} + \tau occ_{o,t} + \omega_{o,t} occ * (immig_{i,t} = 1) + \varepsilon_{i,o,t} \tag{1.1}$$

where i denotes the individual, o indexes occupations and t indexes the survey years. $X_{i,t}$ is a vector of control variables that include the worker's age expressed as a third-order polynomial, state fixed effects, and dummies for levels of education. occ is a vector of over 300 three-digit occupations as defined by Autor and Dorn (2013). The interaction of occ with the indicator variable for immigrants, $immig$, allows the coefficient on the immigrant variable to vary systematically by occupation over time. Generally, the coefficient on the immigrant dummy measures the "wage penalty", which, if negative, reflects the wage disadvantage faced by an immigrant in the same occupation as an observationally identical native. As such, the coefficient ω is the amount of the wage gap not attributable to human capital characteristics or differences in occupational distribution. It is the intra-occupation wage penalty.

$$\hat{\omega}_{o,t} = \kappa + \gamma \mathbf{D} + \nu_{o,t} \tag{1.2}$$

The estimated coefficients from the vector, ω, become the dependent variable in the second stage. In short, $\hat{\omega}$ is the human capital-adjusted wage differential between natives and immigrants within each occupation/year cell.[20] Equation 1.2 is the workhorse model used to assess not only the extent of the intra-occupational wage gap but also to explore the systematic factors that drive the differences in wages of natives and immigrants in identical occupations over time.

A simple way to think about this specific measure of the residual wage gap is the unexplained portion that would arise from the typical Blinder-Oaxaca decomposition. For example, consider two workers, an American and a Jamaican, employed as an economist

[20]Since $\hat{\omega}$ is a generated dependent variable, future versions of this paper will bootstrap the standard errors.

(occupation code *166*) with identical years of schooling and age. Suppose we control for the unobserved wage determinants within a state and occupation through state and occupation fixed effects. In that case, any observed wage difference between the American and Jamaican economists over and above the difference explained by the observables is the unexplained portion. I argue that this unexplained or residual wage differential can be partly attributed to the specific features of the occupation.[21] Empirical studies from the gender wage gap literature and, to a lesser extent, the immigration literature have primarily attributed the systemic factors to discrimination (Bartolucci, 2014). However, the potential drivers of the wage gap between natives and immigrants within occupations remain unclear. In this regard, a worthwhile contribution to the immigration literature is considering the probable explanation of this gap. Accordingly, I examine three specific occupation characteristics that could explain the variation in the wage penalty for immigrants within a given occupation group and year. These characteristics include the O*NET's measure of English level, average English proficiency, and the fraction of immigrants in a given occupation, which are all captured in the vector D in equation 1.2. Therefore, it is essential to note that these variables in vector D are not logged transformed. Thus, the estimated coefficient on English level, for example, reflects the average percent change in the native-immigrant intra-occupation wage gap associated with a one percentage point increase in the level of English needed to perform the job.

Observations are weighted by the total number of observations used to calculate the native-immigrant wage gap in the occupation-year cell.

[21] Following (Liu et al., 2004) and the gender wage gap literature, I take characteristics relating to "taste" as given and assume that various attributes of the occupation can explain the existing intra-occupation wage gap.

1.5 Results

1.5.1 Intra-occupational wage gap at the national level

I begin the analysis by examining the estimation results at the national level. This specification abstracts from immigrants' location decisions by treating the U.S. as a single labor market. The coefficients on the year-fixed effects reflect the time trend in the average intra-occupational wage gap relative to 1980 and are reported in each column; these coefficients are significantly different from zero in all specifications. The OLS estimation results are presented separately for males and females in Tables 1.2 and 1.3, respectively.

Each pair of columns displays results with and without occupation fixed effects. Given that immigrants do not randomly choose their occupations, and assuming that the characteristics of an occupation remain constant over time, I control for these characteristics by pooling multiple Census samples and including occupation fixed effects. These fixed effects account for time-invariant differences in occupational characteristics and mitigate endogeneity bias due to immigrant selection. Without occupation fixed effects, the estimated coefficients on the covariates may be biased.

Across both genders, the model specifications in columns 1 and 2 present the baseline results without incorporating any measures of occupational attributes, as specified in vector D in equation 1.2. Column 1 in Table 1.2 shows that residual wage differences are highly significant and have expanded over time for men. For instance, in 1980, the intra-occupational wage gap between natives and immigrants was 8.8%. By 1990, this gap had widened by 3.4 percentage points, reaching 12.2%. By 2000, the wage penalty had peaked at 16.6% before declining to 12.7%, relative to the 1980 wage gap. A similar analysis applies to immigrant women. Interestingly, while immigrant women earned an intra-occupation wage premium of 2.4% compared to observationally equivalent native women in 1980, the intra-occupation wage differential for women deteriorated more rapidly than the differential between native-born and immigrant men, especially after 1990. This trend is consistently observed across all other model specifications presented in the remaining columns. Schoeni (1998) also utilized decennial Census data to document that labor force participation rates for native-born and immigrant women were relatively similar in 1970 and 1980. However, the participation rates for immigrant women declined in 1990, which could account for the increased wage penalty within occupations for immigrant women if this trend continued over time. It is important to note that my model does not account for the labor supply decisions of either immigrants or natives. Consequently, it is challenging to determine whether the observed patterns in the intra-occupational wage gap are driven by self-selection biases or broader secular trends.[22]

In column 2, which includes a vector of occupation fixed effects as covariates, the intra-occupational wage gap has expanded even more than anticipated, indicating that the estimates in column 1 were likely understated. This notable finding suggests that, even after accounting for unobservable characteristics of occupations, immigrants' wages have deteriorated significantly. However, this does not rule out the possibility that structural aspects of occupations may contribute to the decline in wages over time. For instance, factors such as unionization and collective bargaining can influence wage composition. Although occupation fixed effects are included to control for these changes, this adjustment may not fully account for the sorting bias, raising the question of whether occupation-specific characteristics could explain the observed wage gap. To explore this possibility, I turn to column 3, where I introduce the fraction of immigrants within a given occupation-year cell.

On average, when period effects are accounted for, I find that the fraction of immigrants in a given occupation and year is positively associated with the residual wage gap between natives and immigrants in identical occupations and is statistically significant for both men and women. Interesting heterogeneous effects are revealed across genders when I estimate the exact specification but control for occupation fixed effects in column 4. For men, the coefficient on the fraction of immigrants remains statistically significant. However, it is now negative.[23] This result suggests that accounting for unobservable determinants of wages within an occupation may eliminate self-selection bias arising from immigrants sorting

[22]Additionally, I do not distinguish between part-time and full-time workers in any of the model specifications, which could also affect the wage gap.

[23]Also, in the other model specifications, whenever occupation fixed effects are adjusted for, the associated impact of a change in the fraction of immigrant men on the intra-occupational wage gap remains consistently negative.

Table 1.2: Determinants of intra-occupational wage gap for men

	(1)	(2)	(3)	(4)	(5)	(6)
1990	-0.034	-0.053	-0.046	-0.049	-0.049	-0.051
	(0.018)*	(0.014)***	(0.017)***	(0.014)***	(0.016)***	(0.016)***
2000	-0.078	-0.096	-0.118	-0.083	-0.119	-0.121
	(0.023)***	(0.015)***	(0.020)***	(0.017)***	(0.020)***	(0.020)***
2009	-0.059	-0.085	-0.099	-0.072	-0.093	-0.097
	(0.019)***	(0.012)***	(0.018)***	(0.014)***	(0.017)***	(0.017)***
2014	-0.041	-0.075	-0.082	-0.062	-0.072	-0.079
	(0.019)**	(0.012)***	(0.018)***	(0.014)***	(0.018)***	(0.018)***
2018	-0.039	-0.061	-0.075	-0.048	-0.063	-0.068
	(0.020)*	(0.012)***	(0.018)***	(0.014)***	(0.019)***	(0.019)***
Fraction of immigrants			0.371	-0.147	0.976	0.950
			(0.062)***	(0.089)*	(0.181)***	(0.187)***
English skill level (O*NET)					0.026	-0.115
					(0.008)***	(0.043)***
Avg. English ability						-0.148
						(0.063)**
constant	-0.088	-0.041	-0.121	-0.039	-0.209	0.163
	(0.012)***	(0.021)*	(0.013)***	(0.021)*	(0.028)***	(0.141)
R^2	0.02	0.68	0.11	0.68	0.16	0.18
N	1,723	1,723	1,723	1,723	1,723	1,723
Occ FE		X		X		X

Notes: Data are sourced from the 1980, 1990, and 2000 IPUMS U.S. Census and from the ACS five-year pooled samples: 2005-2009 (for 2009), 2010-2014 (for 2014), and 2014-2018 (for 2018). The sample consists of employed adult females aged 25-64. The dependent variable is the residual intra-occupation wage gap between natives and immigrants. Each column represents a different model specification and displays the weighted results from a pooled cross-sectional regression. The weights used correspond to the sample size of the dependent variable. Robust standard errors are reported in parentheses. *, ** and *** indicates significance at 10%, 5% and 1% significance level,respectively.

Table 1.3: Determinants of intra-occupational wage gap for women

	(1)	(2)	(3)	(4)	(5)	(6)
1990	-0.050	-0.044	-0.054	-0.052	-0.053	-0.054
	(0.026)**	(0.014)***	(0.028)*	(0.014)***	(0.028)*	(0.026)**
2000	-0.090	-0.101	-0.124	-0.121	-0.122	-0.125
	(0.027)***	(0.015)***	(0.028)***	(0.016)***	(0.028)***	(0.026)***
2009	-0.093	-0.113	-0.130	-0.137	-0.125	-0.123
	(0.027)***	(0.011)***	(0.026)***	(0.013)***	(0.027)***	(0.025)***
2014	-0.097	-0.121	-0.141	-0.149	-0.132	-0.134
	(0.026)***	(0.012)***	(0.025)***	(0.014)***	(0.026)***	(0.025)***
2018	-0.105	-0.132	-0.151	-0.159	-0.138	-0.141
	(0.030)***	(0.012)***	(0.027)***	(0.015)***	(0.028)***	(0.026)***
Fraction of immigrants			0.388	0.221	-0.517	-0.430
			(0.060)***	(0.068)***	(0.203)**	(0.196)**
English skill level (O*NET)					-0.023	0.037
					(0.009)***	(0.055)
Avg. English ability						0.133
						(0.070)*
constant	0.024	-0.074	-0.009	-0.065	0.063	-0.218
	(0.019)	(0.032)**	(0.021)	(0.032)**	(0.039)	(0.175)
R^2	0.03	0.71	0.15	0.72	0.22	0.23
N	1,607	1,607	1,607	1,607	1,607	1,607
Occ FE		X		X		X

Notes: Data are sourced from the 1980, 1990, and 2000 IPUMS U.S. Census and from the ACS five-year pooled samples: 2005-2009 (for 2009), 2010-2014 (for 2014), and 2014-2018 (for 2018). The sample consists of employed adult females aged 25-64. The dependent variable is the residual intra-occupation wage gap between natives and immigrants. Each column represents a different model specification and displays the weighted results from a pooled cross-sectional regression. The weights used correspond to the sample size of the dependent variable. Robust standard errors are reported in parentheses. *, ** and *** indicates significance at 10%, 5% and 1% significance level,respectively.

into occupations where they perform relatively well. Consequently, on average, the intra-occupation wage penalty tends to increase with the proportion of the national workforce composed of immigrants. Specifically, a one percentage point increase in the fraction of immigrants within a given occupation and year raises the intra-occupational wage penalty for immigrant men by 14.7 percentage points.

Turning to the results for women, the coefficient on the fraction of the immigrant workforce remains positive and significantly different from zero, despite the inclusion of occupation fixed effects, as illustrated in column 4. Immigrant women appear to be sorting into "better" occupations, but this may diminish the extent of their relative disadvantage within occupations over time. Specifically, the self-selection bias-adjusted effect of a one-percentage-point increase in the share of the national workforce comprised of immigrant women is associated with a reduction of 22.1 percentage points in the intra-occupational wage penalty. Using only these covariates understates the severity of the wage penalty and, therefore, may not be enough to explain the dramatic widening of the wage gap between native-born and immigrant women within occupations.

Finally, holding constant these measures of immigrant shares, a one percentage point increase in the English level requirement of an occupation is generally associated with positive intra-occupational earning effects for both men and women. However, this effect is primarily driven by occupation controls, suggesting that highly proficient immigrants select occupations requiring high English proficiency levels.

Having established the results at the national level, I now turn to the local labor market to continue the analysis across states, followed by metropolitan areas.

1.5.2 Intra-occupational wage gap across states

The national-level results reveal significant heterogeneity in the intra-occupational wage penalties experienced by immigrants across the U.S. labor market. However, it is important to emphasize that if the occupation fixed effects are insufficient to account for the level of occupational clustering, the estimates may: (i) be biased; and (ii) distort the impact of sorting, potentially exaggerating the intra-occupation wage penalty. In a related framework, Orrenius and Zavodny (2007) observed that endogeneity bias, driven by occupational sorting, decreases as additional area- and year-specific controls are incorporated into the model. Therefore, it may be valuable to investigate the results at a more granular level by exploiting the variation within occupations across states and over time. This approach would help identify the regional intra-occupational wage gaps between natives and immigrants.

Analysis at the state level may reveal interesting spatial variation in the wage penalty given the idiosyncrasies inherent to the labor markets across states. These idiosyncrasies are not only limited to local economic conditions, state policies, and attitudes towards immigrants but also include immigrants' proclivity to settle in particular states. For instance, over the sample period, Mexicans make up 10% of the sample population in California and less than one percent in Maryland. On the other hand, Indians and Chinese account for just over one percent of residents in California and less than one percent in Maryland. Although not the most dramatic, these summaries suggest that immigrants from specific birthplaces have developed ethnic enclaves in specific states. Furthermore, Patel and Vella (2013) documents that this phenomenon is also evident in specific occupations. Suppose the rate of immigrant concentration across states is growing on a larger scale. In that case, the crucial question is whether the likelihood of an immigrant working in an identical occupation as their compatriot affects the residual wage gap in that occupation. I explore this possibility by constructing a variable, `frac_bpl`, that measures the share of immigrants from a given birthplace in a given occupation, state, and year relative to total employment in that occupation, state, and year. The intuition here is to calculate the level of ethnic concentration in each occupation.[24]

I augment equations 1.1 and 1.2 as follows:

$$ln(wage_{i,o,s,t}) = \lambda + \beta \mathbf{X}_{i,o,s,t} + \tau occ_{o,s,t} + \omega_{o,t} occ * immig + \varepsilon_{i,o,s,t} \tag{1.3}$$

[24] In other words, I want to distinguish between occupations that are considered "immigrant-dense" and those that are "ethnically-dense".

where o, s, and t indexes occupations, states, and survey years, respectively. $wage_{i,o,s,t}$ measures the log weekly wages in occupation o in state s, survey year t. As before, the vector X contains the same conditioning variables from equation 1.1 to capture the worker's demographic characteristics, and likewise, occ is a vector of over 300 three-digit occupations.

$$\hat{\omega}_{o,s,t} = \kappa + \eta frac\_immig + \delta frac\_bpl + \nu_{o,s,t} \tag{1.4}$$

As before, the second regression explicitly controls for the occupation-specific features, namely $frac\_immig$ and $frac\_bpl$. The former accounts for the competition among immigrants that sort into identical occupations. The latter reflects the impact of ethnic networks, the tendency of immigrants from certain birthplaces to sort into identical occupations as their compatriots. Therefore, η is identified using the variation within occupations across states over time.

Equations 1.3 and 1.4 are estimated using pooled data over the full sample period 1980-2018 and separately by gender. The results are illustrated in table 1.4. I estimate various specifications of equation 1.4 that include state and period fixed effects as well as with and without occupation fixed effects. As a reminder, the occupation fixed effects help to account for self-selection bias. In contrast, the state fixed effects absorb the time-invariant characteristics of the state that might affect wages.

Columns 1 and 3 indicate that, in the absence of occupation controls, higher immigrant shares within an occupation appear to be a positive and significant determinant of the intra-occupation wage gap across states. However, when occupation controls are added to the model, as shown in columns 2 and 4, the estimated effect on the wage gap is now negative and statistically significant. A one-percentage-point increase in the share of male immigrants within a state's workforce is associated with a 21-percentage-point increase in the wage penalty. In contrast, for immigrant women, the effect is smaller; a one-percentage-point increase in the proportion of immigrant women in the state's workforce raises the intra-occupation wage penalty by 12 percentage points.

Regarding the impact of ethnic clustering on the intra-occupation wage gap across states, the results indicate a positive effect that is statistically significantly different from zero. Notably, in columns 1 and 3, when occupation controls are omitted from the model and the overall fraction of immigrants is accounted for, an increase in the share of the state's immigrant workforce from the same birthplace in a given occupation and year narrows the intra-occupation wage gap about 11-12 percentage points. Columns 2 and 4 report the associated change in the intra-occupation wage gap when all occupation-specific measures are considered. Interestingly, the ethnic network effect magnitudes are much stronger for men and women. For example, the effect of immigrants from the same birthplace clustering in identical occupations rose by 19.8 percentage points (from 12.2% to 32.0%) for men and by 5.7 percentage points (from 11.2% to 16.9%) for women, *ceteris paribus.*

The impact of ethnic clustering appears to be diametrically opposed to the effect of immigrant competition on the intra-occupation wage penalty. However, an increase in the immigrant share of a state's workforce within a given occupation and year mechanically raises the likelihood that immigrants will share the same birthplace. My results suggest that, in a local labor market setting, if immigrants from the same birthplace complement each other more effectively than other immigrants and natives in identical occupations, an increase in the size of the ethnic network will influence the relative price of labor within that occupation. This *complementary* effect may not necessarily stem from differences in skills among co-nationals but could instead reflect better job-worker match quality or enhanced bargaining power due to network connections. Consequently, wage regressions may underestimate the impact of ethnic clustering on the intra-occupation wage gap if the overall immigrant share in the workforce is not accounted for.

1.5.3 Intra-occupational wage gap at the metro-area level

Before proceeding to the estimation results, it bears noting that to avoid cells with small sample sizes, I follow the literature by making the following adjustments to the local labor market. Firstly, the metropolitan area is defined as the metro area of the individual's place of

Table 1.4: Determinants of variation in intra-occupational wage gap across states

	Males		Females	
	(1)	(2)	(3)	(4)
1990	-0.040	-0.043	-0.040	-0.031
	(0.009)***	(0.007)***	(0.013)***	(0.010)***
2000	-0.098	-0.082	-0.079	-0.048
	(0.016)***	(0.015)***	(0.020)***	(0.019)**
2009	-0.052	-0.052	-0.100	-0.090
	(0.008)***	(0.007)***	(0.011)***	(0.009)***
2014	-0.037	-0.045	-0.109	-0.099
	(0.008)***	(0.007)***	(0.011)***	(0.010)***
2018	-0.035	-0.032	-0.117	-0.114
	(0.008)***	(0.007)***	(0.011)***	(0.011)***
Fraction of:				
Immigrants	0.121	-0.209	0.277	-0.118
	(0.043)***	(0.032)***	(0.031)***	(0.034)***
Immig. by birthplace	0.310	0.736	0.710	0.783
	(0.138)**	(0.131)***	(0.230)***	(0.218)***
constant	-0.158	-0.079	-0.155	-0.205
	(0.009)***	(0.041)*	(0.009)***	(0.045)***
R^2	0.02	0.31	0.09	0.34
N	7,529	7,529	6,497	6,497
Occ FE		X		X

Notes: Data are from the IPUMS ACS 2014-2018 5-year (pooled) sample. The sample is employed adult males and females aged 25-64. The dependent variable is the residual state-level intra-occupation wage gap between natives and immigrants. Each column represents a different model specification and shows the weighted results from a pooled cross-section regression. The weights used are the sample size of the dependent variable. Robust standard errors are reported in parentheses. *, ** and *** indicates significance at 10%, 5% and 1% significance level,respectively.

work. Secondly, the sample is restricted to the top 121 immigrant-dense metropolitan areas.[25] Finally, the second stage of the analysis is done separately by year due to differences in Census boundaries.[26] The unit of observation is an individual working in the metropolitan area, m, in

[25]This number is not arbitrary, these metro areas have at least 1,000 immigrant workers with a minimum of 100 from the same birthplace.

[26]In particular, the Census Bureau redraws Public Use Microdata area (PUMA) boundaries every decade so that it is more difficult to pool the cross-section of surveys while maintaining consistent metro areas codes. An alternative approach is to create a crosswalk using the 2013 Census delineations for metropolitan

occupation, o in survey year t. The model specification is analogously defined as in equations 1.3 and 1.4. As mentioned previously, endogeneity bias poses a significant issue because immigrants' location choices are not random. This concern becomes even more pronounced at the metropolitan area level, where immigrants may be more likely to settle in high-wage regions. To address this concern, I include controls for the area, state, and occupation fixed effects to absorb unobserved factors that may influence immigrants' settlement choices so that identification rests solely on the wage differential within occupations across metro areas.

The estimated results are shown in Table 1.5. I focus on the critical occupation-specific attributes that capture immigrant clustering to determine their effects on the intra-occupation wage gap at the localized level. For expositional convenience, I discuss the results for the survey years 2000-2018.

Immigrants appear to compete more strongly with each other within occupations at the metro level. Specifically, the effect of an increase in the immigrant share of the workforce in a given metro area, occupation, and year is negatively correlated with the intra-occupation wage gap for men and women. Interestingly, this result holds even without occupation controls (for men) and is nearly twice as large as the estimates at the state level. For example, a one percentage point increase in the share of immigrant men raised the intra-occupation wage gap by 43.3 percentage points in 2000, before declining sharply by 25.8 percentage points in 2018. In contrast, the corresponding effects for immigrant women are less pronounced, showing a decline of 8.3 percentage points in 2000 and 25.7 percentage points in 2018. One possible explanation for these results is that the immigrant population in the U.S. has become more homogeneous over time. The Bureau of Labor Statistics reported that in 2018 and 2019, Hispanics made up nearly half of the immigrant labor force, while Asians accounted for about one-quarter (Bureau of Labor Statistics, 2023). This pattern, however, is not a recent phenomenon; it broadly characterizes the composition of the foreign-born workforce over the past two decades. Consequently, the elasticity of

statistical areas (MSAs) to match with Census metro codes for previous survey years. However, the number of identifiable metro areas declined significantly, affecting cell sizes. Aydemir and Borjas (2011) find significant attenuation bias due to measurement error from small sample sizes when estimating the wage impact of immigration. Considering this, I do not pool survey years.

substitution between immigrants and natives may have been greater in the past than it is now. In essence, immigrants in today's labor market are relatively more homogeneous and, as a result, are closer substitutes for each other than for natives. While immigrants are increasingly sorted into better occupations, they compete intensely within those occupations, which exacerbates the intra-occupation wage gap. It is worth noting, however, that the persistent and substantial increase in the wage penalty faced by immigrant women within occupations remains puzzling. This suggests the possibility that these estimates may be influenced by other unobserved factors, such as labor force participation decisions.

The estimated coefficients on the measure of ethnic clustering are consistent with the relationship uncovered at the state level. Notably, for immigrant men, controlling for occupation fixed effects and immigrant competition within the same occupations reveals that an increase in the share of immigrants from the same birthplace within the same occupation is associated with a reduction in the intra-occupation wage penalty of up to 52.1 percentage points in 2000, which moderated to 18.5 percentage points by 2018. Notice that again, even at the localized level, the overall fraction of immigrants in a given occupation and the share from the same birthplace have competing effects on the intra-occupation wage gaps. However, the positive wage benefits immigrant men enjoyed through ethnic membership within a given occupation have substantially weakened over time. Thus, the negative wage penalty for immigrants being more substitutable to each other has dominated the overall impact on the intra-occupation wage penalty for immigrant men. In other words, the wages of immigrant men in the same occupations as their native-born counterparts within a given metropolitan area have, on average, declined over time. This finding is compelling and suggests that occupations have become more ethnically diverse over time. As such, an increase in the proportion of immigrants from a specific ethnic group now has a less pronounced impact on wages than it did previously.

Regarding the relationship between ethnic clustering and the intra-occupation wage gap at the metro level, the estimated effects are inconclusive across years for immigrant women. Notably, in column 4, which accounts for immigrant share and occupation controls, the effect of a one percent increase in the share of co-national immigrant women in a given occupation, metro area, and year fluctuates between -0.26 and 0.16 percentage points and is mostly insignificant. On the other hand, for immigrant men, ethnic clustering narrows the intra-occupation wage gap by 0.19 to 0.52 percentage points over the sample period and is statistically different from zero, except in 2018. The results for immigrant women highlight important implications for understanding the intra-occupation wage gap in the local labor market. For example, exploring other occupation features that emphasize gender differences might prove crucial to further explain immigrants' wage disadvantage within occupations in a given metro area and year. Overall, the characteristics investigated in this paper—such as English language proficiency, the immigrant workforce share in an occupation, and ethnic clustering—seem to better explain the intra-occupation wage differential for immigrant men compared to immigrant women.

1.6 Conclusion

The native-immigrant wage gap is a long-standing issue that still deserves further attention. It is pertinent for immigrants to select the "right" occupations because this affects how well they assimilate into the U.S. labor market. This paper explores the native-immigrant wage gap within occupations by investigating whether the specific features of the occupation help explain this gap. I use data from the Census and O*NET to compose three attributes of the occupation that capture occupation sorting, ethnic clustering, and English language skill requirement of a given occupation. Finally, I exploit the variation in occupations over time to identify the effects of these characteristics on the intra-occupation wage gap at the national level and across local labor markets, namely the state and metro area.

The findings appear consistent across labor markets. However, the reliability of the estimated effects hinges on the assumption that the various fixed effects can sufficiently control for the endogeneity bias. The residual approach used adjusts for observable human capital characteristics that potentially affect wages within a given occupation. However, it is likely that immigrants' location choices, particularly at the localized level, depend on other uncontrolled factors that determine the intra-occupation wage. To the extent that this concern

Table 1.5: Determinants of variation in intra-occupational wage gap across metro areas, 2000-2018

	Men		Women	
	(1)	(2)	(3)	(4)
Year 2000:				
Share of immig.	-0.181	-0.433	0.072	-0.083
	(0.075)**	(0.054)***	(0.078)	(0.057)
Immig. share by birthplace	0.137	0.521	-0.275	-0.268
	(0.182)	(0.123)***	(0.237)	(0.155)*
constant	-0.186	-0.118	-0.122	-0.100
	(0.011)***	(0.036)***	(0.010)***	(0.122)
R^2	0.01	0.42	0.00	0.34
N	2,274	2,274	2,215	2,215
Occ FE		X		X
Year 2009:				
Share of immig.	-0.095	-0.333	0.007	-0.255
	(0.050)*	(0.053)***	(0.065)	(0.052)***
Immig. share by birthplace	0.038	0.328	-0.021	-0.019
	(0.176)	(0.139)**	(0.184)	(0.120)
constant	-0.237	-0.101	-0.139	-0.089
	(0.012)***	(0.046)**	(0.010)***	(0.119)
R^2	0.01	0.51	0.00	0.45
N	2,393	2,393	2,317	2,317
Occ FE		X		X
Year 2014:				
Share of immig.	-0.159	-0.342	0.009	-0.206
	(0.063)**	(0.043)***	(0.049)	(0.044)***
Immig. share by birthplace	0.183	0.297	0.072	0.004
	(0.134)	(0.113)***	(0.135)	(0.117)
constant	-0.206	-0.003	-0.168	-0.221
	(0.012)***	(0.058)	(0.010)***	(0.028)***
R^2	0.01	0.45	0.00	0.40
N	2,378	2,378	2,302	2,302
Occ FE		X		X
Year 2018:				
Share of immig.	-0.067	-0.258	-0.065	-0.257
	(0.061)	(0.062)***	(0.052)	(0.048)***
frac_bpl	0.084	0.185	0.368	0.162
	(0.116)	(0.143)	(0.120)***	(0.110)
constant	-0.228	-0.152	-0.172	-0.062
	(0.012)***	(0.057)***	(0.010)***	(0.110)
R^2	0.01	0.42	0.00	0.26
N	2,048	2,048	2,051	2,051
Occ FE		X		X

Notes: Data are from the 2000 U.S. Census and 2009, 2014, and 2018 pooled 5-year ACS. The sample is employed adult males and females aged 25-64. The dependent variable is the residual intra-occupation wage gap between natives and immigrants. Each column represents a different model specification and is weighted using the sample size of the dependent variable. Robust standard errors are reported in parentheses. *, ** and *** indicates significance at 10%, 5% and 1% significance level, respectively.

is legitimate, the estimated effects of the occupation-specific characteristics on the intra-occupation wage gap should be viewed as upper-bound estimates and the first step of a two-part analysis. The second step would be to implement an instrumental variable method like Patel and Vella (2013) or employ a shift-share instrument like Card (2001), Borjas and Cassidy (2019), and Sharpe and Bollinger (2020) to account for the endogeneity bias due to the non-random location choices of immigrants. Moreover, I focus on a narrow set of occupation-specific attributes, which does not account for other possible occupation features that may influence the wage gap. Nevertheless, it is important to note that my findings represent an initial step in understanding how immigrants perform in identical occupations compared to observationally equivalent natives. This research is ongoing, and future iterations of this paper will address the limitations discussed.

Chapter 2

SNAPing to the Beat: Revisiting the Impact of Immigrants' Food Stamp Eligibility

2.1 Introduction

Does access to food stamps disproportionately affect immigrants' economic outcomes?[1] This question has regained prominence among researchers and policymakers, given the extensive policy changes implemented since the COVID-19 pandemic.[2] Two changes warrant particular mention due to their relevance to this study: (i) the Department of Homeland Security's (DHS's) publication of the final rule on *Public Charge Ground of Inadmissibility* on September 9, 2022, and (ii) the decision by the State Legislature of Massachusetts and Governor Maura Healey in December 2023 to extend SNAP benefits to all legal immigrants meeting

[1]The Food Stamp Program, commonly known as food stamps, was renamed as the Supplemental Nutrition Assistance Program (SNAP) in 2008 under the 2008 Farm Bill. In this paper, I consistently use the term SNAP to denote the federal food assistance program previously known as food stamps.

[2]The Immigrant Legal Resources Center (ILRC) has documented over 100 policy changes impacting non-citizens and asylum seekers in the U.S. The report is summarized here: `https://www.ilrc.org/100-policy-changes-have-devastated-immigrants-and-asylum-seekers`. Similarly, reports from the American Immigration Council have concentrated on the ramifications of COVID-19 on the U.S. immigration system (Loweree and Walter, 2020).

the program's income criteria.

The DHS's rule became effective on December 23, 2022, and "*determines whether non-citizens are inadmissible to the United States because they are likely at any time to become a public charge.*"[3] It confirms that public charge determinations do not consider SNAP participation. Therefore, certain lawfully present non-citizens may apply for or receive SNAP benefits without fear of becoming a public charge. This new rule may potentially alter the existing narrative surrounding the fear of participating in government safety net programs among immigrants, which creates a "chilling effect" leading to adverse economic outcomes, particularly concerning labor supply. Similarly, the impact of Governor Healey's decision to extend state-funded SNAP benefits may reflect a shift in state policy aimed at proactively addressing the substantial influx of immigrants that has steadily increased across states since the pandemic.

For researchers, the question of whether safety net programs help or harm individual economic outcomes is an empirical one that warrants frequent evaluation and reassessment as policy changes are implemented (or repealed) at various government levels, and as new and improved data and econometric techniques become available. This paper uses these developments as motivation to reexamine the impacts of immigrants' access to SNAP.

The existing literature typically employs the 1996 federal welfare law, which introduced unprecedented restrictive eligibility criteria for non-citizen immigrants, who previously had the same eligibility as citizens, to examine whether SNAP disproportionately affects immigrants' economic outcomes (Fix and Passel, 2002).[4] More recent studies, such as East (2018b), exploit the variation in the timing of implementation and subsequent restoration of immigrants' eligibility to examine the effects of SNAP and find strong evidence of labor supply disincentives along the extensive and intensive margins across demographic groups.[5]

[3]Further information is provided at the Federal Register's website: `https://www.federalregister.gov/documents/2022/09/09/2022-18867/public-charge-ground-of-inadmissibility`.

[4]The Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (hereinafter "welfare reform").

[5]This finding is consistent with other studies examining the effect of immigrants' participation in federal safety net programs on their economic outcomes. See, for example, Lofstrom and Bean (2002), Borjas (2004), Haider et al. (2004), and Kaestner and Kaushal (2005).

However, this unique policy setting includes a staggered adoption of treatment within a two-way fixed effects (TWFE) model, which can lead to biased and oppositely signed estimates due to "negative weighting" problems if unadjusted (Goodman-Bacon, 2021). Unsurprisingly, in recent years, an overwhelming array of innovative papers has emerged in the literature on correcting these problems and the biases of the TWFE estimators with heterogeneous treatment effects. This new wave of research, commonly referred to by practitioners as the "Differences-in-Differences (DiD) credibility revolution," has sought to address the econometric issues arising from the differential timing of treatment effects. Consequently, previous studies leveraging the staggered rollout of welfare reform (treatment) between documented non-citizens (treated) and natives (controls) within a DiD or TWFE specification may potentially provide invalid inferences. In fact, Baker et al. (2022) illustrate that these biases can result in Type-I and Type-II errors, which have broader implications for the conclusions drawn from such assessments, ultimately impacting how policymakers make normative decisions.

In this paper, I revisit the impact of SNAP access on immigrants' labor market outcomes, focusing on the prior published results of East (2018b). This analysis is motivated by the new array of econometric studies that have provided robust alternative estimators and statistical packages to correct the shortcomings of the TWFE and staggered DiD models. This is not a frivolous exploration. Given the recent headlines from the Wall Street Journal and the Pew Research Center, reflecting public perceptions that immigration is one of the most crucial issues in the 2024 election, it becomes even more imperative for policymakers to accurately assess the impacts of one of the largest social safety net programs in the U.S. on a population that constitutes a significant and growing segment of the workforce.[6]

By reexamining the impact of SNAP access on labor market outcomes, this study aims to provide nuanced insights that can inform both federal and state-level policy decisions. Leveraging advanced econometric techniques, this research addresses the complexities inherent in evaluating the effects of welfare programs on vulnerable economic groups, offering a

[6]See Parti and Hackman (2024) for the Wall Street Journal article and Pew Research Center (2024) for the article from said organization.

timely contribution to the ongoing policy discourse. The potential labor supply disincentives identified in prior research underscore the need for effective policy interventions that safeguard the well-being of immigrants while fostering economic participation. The starting point for this balanced approach is accurate positive analysis.

I begin the analysis by concisely reviewing the recent literature that has catalyzed the "DiD credibility revolution." I then delve into the policy background of the welfare reform, which underscores the necessity for more refined empirical approaches to accurately recover causal effects. Following this, I replicate the main findings of East (2018b), which employs two reduced-form specifications of the DiD methodology to estimate the intent-to-treat effect of immigrants' access to SNAP. By juxtaposing these results with those obtained using the bias-corrected estimator proposed by Callaway and Sant'Anna (2021), I aim to highlight the differences and improvements in causal inference brought about by advanced econometric methods.

The rest of the paper is organized as follows. Section 2 presents an overview of the canonical DiD model and recent advancements, details the welfare reform, and explores the interplay between the two in obtaining more precise labor supply elasticities. Section 3 reviews the relevant literature and discusses the anticipated effects of SNAP on labor supply. Section 4 examines the data and outlines the empirical strategy employed. Section 5 presents the results, and Section 6 offers concluding remarks.

2.2 Background

2.2.1 An overview of the DiD Model

This section gives a succinct description of the basic setup of the DiD model, with the aim of emphasizing the empirical challenges that the recent innovations in the DiD literature have addressed.[7]

The DiD estimation is an established and commonly used method of quasi-experimental

[7]See Abadie (2005); Snow (1856) for classic references of the DiD model and Baker et al. (2022); Cunningham (2021); Roth et al. (2022) for more recent comprehensive narrations.

research design. Its prominence is perhaps due to the elegance of identifying the average treatment effect on the treated (ATT). The canonical setup of the DiD model involves two periods, $t = 1, 2$, representing the pre- and post-treatment periods, respectively. The units, indexed by i, are assigned to either a treated group, which receives a specific intervention in the second period, $D_i = 1$, or a control group,$D_i = 0$, which remains untreated throughout both periods (Roth et al., 2022). The key identifying assumption is the parallel trends assumption, which posits that, in the absence of treatment, the average outcome, Y_{it}, for both treated and control units would follow the same time trend. Additionally, the no anticipation assumption must hold, which stipulates that all units have untreated potential outcomes in the pre-treatment period. These assumptions enable the researcher to estimate the treatment effect by comparing the changes in average outcomes between treated and control units across both periods. Specifically, using classic notation, the DiD estimate is $(\bar{y}^{post}_{treat} - \bar{y}^{pre}_{treat}) - (\bar{y}^{post}_{control} - \bar{y}^{pre}_{control})$, as used by Goodman-Bacon (2021) or $\mathbb{E}[Y_{i2}(1) - Y_{i2}(0)|D_i = 1]$ following the potential outcomes framework of Rubin (2005).[8] This estimate is the causal estimand, which measures the ATT in the post-treatment period by using the outcome changes of the control units as a counterfactual for the treated units. It is essential to understand the structure of the DiD design, as recent extensions in the literature build upon this framework.[9] In what follows, I highlight two of the recent methodological advances that are most relevant to this paper.

The preceding outline of the DiD model implies a rudimentary estimation using Ordinary Least Squares (OLS) of the following regression:

$$y_{it} = \alpha + \beta_1 D_i + \beta_2 Post_t + \beta_3 (D_i * Post_t) + \varepsilon_{it} \tag{2.1}$$

where β_3 yields the ATT in a two-period, two-group (2 x 2) case; D_i is the binary variable for the treated units; and $Post_t$ represents the indicator variable for units in the treated period, $t = 2$.

[8] Goodman-Bacon (2021) notes that this DiD estimate can also be derived from the coefficient on the interaction of a treated group dummy and a post-treatment period dummy, i.e., β_3 in equation 2.1.

[9] For example, Goodman-Bacon (2021), Sant'Anna and Zhao (2020), Sun and Abraham (2021), Gardner (2022), and de Chaisemartin and D'Haultfœuille (2020).

The advantages of using regression analysis for this setup are well-documented. This approach not only allows researchers to obtain point estimates and their standard errors but also flexibly accommodates the inclusion of covariates, dynamic treatment effects, and multiple units and periods. When multiple units and time periods are involved, the regression DiD model in equation 2.1 is augmented with α_i and γ_t, to represent the unit and period fixed effects. This extension, commonly referred to by practitioners as the TWFE model, serves as the foundational model for estimating the average treatment effect on the treated (ATT) in contexts with staggered treatment timing across various fields.[10].

2.2.2 The Challenges of implementing DiD with multiple periods and heterogeneous treatment timing.

The pervasive use of staggered DiD designs can be attributed to their ability to yield robust causal estimates even in the presence of heterogeneous treatment timing. For instance, the welfare reform transferred the responsibility of determining eligibility rules for legal immigrants for many safety net programs to the states. Consequently, states could utilize their own funds to restore benefits to immigrants, and only nine states chose to reinstate immigrants' access to SNAP.[11] The "Fill-in" states enacted their state-funded replacement coverage between 1997 and 1998. Although the 2002 Farm Bill restored eligibility at the federal level for some groups of immigrants, the welfare legislation resulted in staggered treatment adoption in immigrant eligibility that varied across states and over time between 1997 and 2003. Consequently, the TWFE DiD would yield credible estimates for the dynamic effect of SNAP access for treated immigrants. However, Sun and Abraham (2021) demonstrate that these estimates become difficult to interpret when there are variations in treatment effects across treated groups. Notably, it becomes problematic if the average treatment effect differs for states that restored eligibility in 1997 compared to those that

[10] For example, Baker et al. (2022) show that out of 744 papers published in top-five finance and accounting journals, 407 papers (54.7%) employ a staggered DiD design.

[11] These nine states, henceforth referred to as the "Fill-in" states, are California, Connecticut, Maine, Massachusetts, Minnesota, Nebraska, Rhode Island, Washington, and Wisconsin.

did so in 1998. Goodman-Bacon (2021) shows that the TWFE estimations can make "bad" comparisons between units already treated, leading to coefficients with signs opposite to the individual-level treatment effects.

Goodman-Bacon (2021) derived the mathematical proofs demonstrating the significant biases that arise in these settings. Given the complexity of these derivations, which are beyond the scope of this paper, I will instead focus on the intuition behind the decomposition of the DiD estimand to highlight the challenges associated with using the TWFE DiD model when treatment effects vary over time. The core issue is that, when comparing the differences between treated and control units across pre- and post-treatment periods, the TWFE model uses "forbidden comparisons," particularly between later-treated and earlier-treated units. If the treatment effect is heterogeneous, earlier-treated units cannot serve as valid controls because their initial response to the treatment is not accurately known. This results in the problem of negative weighting, where earlier-treated units receive a negative weight as they are subtracted from later-treated units at time t. The challenge with TWFE models, when treatment timing varies, is that earlier-treated units should have a positive weight since they have already been exposed to the treatment. Robust DiD estimators that address this issue focus on identifying treatment effects by excluding already-treated units as controls. I discuss one of these estimators in section 2.6.1.

2.2.3 *Overview of SNAP*

SNAP is one of the largest means-tested government benefit programs aimed at reducing food insecurity in the United States (Hoynes and Feldstein, 2015).[12] In 2016, SNAP reached more than 44 million individuals at a cost of $70 billion, a significant increase from fewer than 10 million individuals who participated in the program in the early 1970s (United States Department of Agriculture, 2016). SNAP is the least restrictive of the food and nutrition programs: essentially, all families with total family income below 130% of the poverty line are eligible, regardless of their size and structure. In this sense, SNAP is considered a "need-

[12] Other food and nutrition programs include School Lunch, Women, Infants, Children (WIC), and School Breakfast.

based universal safety net," and as an entitlement, once eligible, low-income families can maintain a basic level of consumption from the benefit amount received.

Benefit amounts are modest, averaging only about \$1.40 per person per meal in the fiscal year 2017 (Center on Budget and Policy Priorities, 2018). This amounts to \$254 and \$126 per month on average for the typical SNAP household and recipient, respectively. A family's benefit amount is a function of adjusted family income, the number of eligible family members, and a maximum benefit amount, which is nationally determined (East, 2018a). The benefit formula assumes that families will spend 30% of their net income on food and is expressed as:

$$Benefit(\$) = \textit{Max Benefit(Number eligible in family)} - 0.30 * [\textit{Family income}] \qquad (2.2)$$

Consequently, this formula guarantees a benefit level even if the family has no income. Alternatively, since all members of the family can be eligible, as family income increases, benefits are reduced by the benefit reduction rate (30%). Because of this inverse relationship between labor income and benefit receipts, standard economic theory suggests that such cash or near-cash safety net programs produce strong disincentives for work. That is, the federally defined guaranteed amount produces an income effect, and the benefit reduction rate reduces the net wage, which leads to an income and substitution effect (Fayaz Farkhad and Meyerhoefer, 2018). In fact, as shown in the next section, several studies find significant labor supply disincentives for some groups of SNAP participants (East, 2018b; Eissa and Hoynes, 2006; Moffitt, 1983).

Despite the potential work disincentives associated with food stamp participation, changes to eligibility rules complicate labor supply decisions. A significant shift occurred with the 1996 PRWORA, which imposed work stipulations on Able-Bodied Adults Without Dependents (ABAWDs) receiving SNAP benefits.[13] Additionally, the legislation outlined conditions encouraging all SNAP participants to meet general work requirements to remain

[13]ABAWDs are defined as individuals between ages 18-50 who are not responsible for a child or incapacitated household member and are medically fit to work. For further details, see: https://www.fns.usda.gov/snap/able-bodied-adults-without-dependents-abawds

eligible (Fayaz Farkhad and Meyerhoefer, 2018). These work requirements include but are not limited to, participating in state training and employment programs, actively seeking and registering for work, and maintaining employment or a work effort of at least 30 hours per week.

2.2.4 *Immigrants' Eligibility and SNAP Policy Changes*

The PRWORA also had implications for non-citizens. Specifically, the welfare reform eliminated SNAP eligibility for legal immigrants. Prior to the welfare reform in 1996 (pre-PRWORA), legal immigrants' eligibility was identical to that of natives for most safety net programs, while unauthorized immigrants were not eligible. Temporary legal residents such as those on student visas or temporary work permits were also generally not qualified for these programs (Bitler and Hoynes, 2011). Following the passage of PRWORA and the Illegal Immigration Reform and Immigrant Responsibility Act (IIRIRA), new eligibility rules emerged for legal permanent residents -LPRs- (both those entering the U.S. before or after PRWORA) and naturalized citizens. Notably, for LPRs who arrived in the U.S. prior to PRWORA (known as pre-enactment immigrants), federal law extended eligibility for Temporary Assistance for Needy Families (TANF) and Medicaid, while those who arrived after PRWORA (post-enactment immigrants) were prevented from receiving these program benefits until they had been in the U.S. for five years. However, legislative changes for Food Stamps and Supplemental Security Income (SSI) were more restrictive: both pre-and post-enactment immigrants were ineligible to receive benefits until they accumulated 40 quarters of work history (Bitler and Hoynes, 2011). A relatively small group of immigrants faced different eligibility rules. Refugees, asylum seekers, those who served in the military, and naturalized citizens were unaffected by the welfare reform rules.

Against this background, I follow East (2018b) by defining the "treated immigrant" sample as those whose year of arrival predates 1997 and less than 15 years before being observed. These sample restrictions allow for the isolation of immigrants who were only affected by changes in Food Stamp eligibility. In particular, the 15-year cutoff eliminates immigrants who would have completed 40 quarters of work and thus were not affected by the reform. The 1997 cutoff removes possible confounding effects as immigrants were affected by eligibility rules of other safety net programs.

These various program eligibility changes illustrate significant heterogeneity in state immigration policies across the country. Considering this as well as the current polarizing immigration climate, it is important to understand the policy implications of changing immigrants' eligibility on various economic outcomes.

2.3 Related Literature and Expected Effects

In this paper, I study the effects of changes in Food Stamp eligibility among immigrants on their earnings, income, and occupational attainment. While these outcomes are primarily understudied in literature, many studies examine the link between Food Stamp participation and labor supply. Examining the empirical literature on the labor supply effects of SNAP might be helpful because labor supply theory predicts changes in wages in response to labor supply shifts. Therefore, I begin this section by highlighting the early literature on the labor supply effects of SNAP participation.

Earlier studies of the effect of Food Stamp participation on labor supply are largely based on structural estimation, possibly due to the lack of variation in program eligibility rules or benefit amounts across geographic locations or over time. Moffitt and Fraker (1988) is the first paper to use this method to estimate the effects of U.S. transfer programs on labor supply for a sample of female heads of households. In particular, they model the joint response of hours of work (zero, part-time, or full-time) to participation in AFDC and Food Stamps and find that Food Stamps have work disincentive effects of a one-hour reduction in work per week (9% reduction off the mean) among recipients. Using the same methodology, Keane and Moffitt (1998) extend this work by estimating the labor supply effects of female-headed households' simultaneous participation in multiple welfare programs (AFDC, Food Stamps, Medicaid, and housing benefits). They find larger wage elasticities of labor supply than Moffitt and Fraker (1988), although the simulations for the overall effect of Food Stamps on labor supply were not explicitly mentioned. The policy simulations of Chan (2013)'s dynamic structural model of labor supply, welfare, and Food Stamp participation reveal that the economy accounts for half of the increase in female labor supply. Using an alternative method, Hagstrom (1996) estimates a nested multinomial logit model to identify the effects of Food Stamps on married couples' labor supply. The author finds smaller labor supply effects on changes in food stamp benefits in comparison to previous studies on single parents.

More recent studies implement quasi-experimental methods by utilizing the variation across counties in the timing of the program's roll-out in the 1960-70s (Hoynes and Schanzenbach, 2012).[14] In fact, Hoynes and Schanzenbach (2012) is the first to estimate the impact of the Food Stamp program on the extensive and intensive margins of labor supply, earnings, and family cash income using quasi-experimental methods. Overall, the authors find modest reductions in employment and hours worked when the program was introduced, with the larger effects for low-educated single mothers. East (2018b) builds upon these findings by utilizing modern state changes in the program's eligibility rules. As noted in the subsequent section, legal immigrants lost eligibility under the PRWORA in 1996. However, some states choose to restore eligibility to immigrants without additional eligibility restrictions. In this context, East (2018b) examines how the loss, and subsequent restoration of eligibility, affected immigrants' labor supply outcomes.

2.3.1 *Expected Effects of SNAP Eligibility*

As mentioned in the previous section, SNAP benefits are designed to supplement earnings in low-income households or help smooth consumption for families unable to work. Standard models of labor supply theory suggest that Food Stamps, like any other means-tested programs, alters the household's labor-leisure tradeoff given the benefit formula. In particular, the food stamp benefit is largest at zero hours of work, where individuals receive the food stamp "guaranteed amount", but decreases at the legislated benefit reduction rate for some positive hours of work. The initial benefit reduction rate in the food stamp program is

[14] Almond et al. (2011) and Hoynes et al. (2016) also employed this variation to examine the effects of the Food Stamp program on children's and families' outcomes.

30%, which is the lowest among other means-tested transfer programs in the United States (Hoynes and Schanzenbach, 2012).[15] Modern changes in SNAP include a deduction for 20 percent of earnings, which implies that the effective benefit reduction rate is even lower than 30%.

Considering this combination of a guaranteed income and benefit reduction rate, the simple static labor supply model unambiguously predicts a decline in employment (extensive margin), hours worked (intensive margin), and earnings (if wages are fixed). The rationale for these predicted declines is due to the income effect of the benefit, which is predicted to increase their leisure (reduce their hours worked if leisure is a normal good). Alternatively, it is possible that individuals can be pushed out of the labor market due to the combined negative income and substitution effects from the benefit reduction rate (Hoynes and Schanzenbach, 2012).

However, in the modern context of the program, the Center on Budget and Policy Priorities (2018) reported that the SNAP benefit formula contains an important work incentive: for every extra dollar a SNAP recipient earns, the SNAP benefits decline by only 24 to 36 cents. As such, the benefit formula favors earned income over unearned income given the 20% income deduction. Consequently, families that receive SNAP have a strong incentive to work more hours or search for better-paying jobs. In fact, SNAP increases a family's take-home income by roughly 15 to 21%, depending on the number of hours worked. For example, a mother with two children who works for 35 hours at $10 per week, increases her monthly income by 21% when her SNAP benefits are added. As a result, the modern effects of SNAP on earnings are somewhat unclear. On one hand, there could be work disincentives by tying SNAP benefit receipts to income levels, but on the other hand, SNAP's work requirements create work incentives. Ultimately, the effect of SNAP eligibility on work incentives is an empirical question that deserves further investigation.

[15] In comparison to benefit reduction rates of 67% or 100% in AFDC.

2.4 Data

To examine the impact of the change in SNAP eligibility on various economic outcomes, I utilize the Annual Social and Economic (ASEC) Supplement to the Current Population Survey (CPS) from 1995-2007. The ASEC, conducted annually in March, contains rich information about employment, union membership, health insurance, and taxes. While the Basic Monthly CPS, is the official source of labor force statistics of the United States, the ASEC gathers data on social and economic indicators, from which the official poverty rate of the United States is computed (Flood and Pacas, 2016a). It provides annual estimates based on a survey of more than 75,000 households.[16]

Beyond the social and economic variables that are central to the ASEC, and important for my analysis, the survey includes information on household Food Stamp receipts and the dollar amount of benefits received in the past year, which are not available in any other supplement of the CPS. The ASEC/CPS also reports information on citizenship status and year of immigration for individuals, but this only began in 1994. Due to the reported issues with the 1994 survey weights for immigrants, I am only able to observe one pre-welfare reform period.[17] It is worth mentioning that the sample end period is deliberately chosen to avoid the possible confounding effects of other immigration policies that were implemented in the latter part of the decade. Notably, the Secure Communities program, which is one of the largest federal immigration enforcement programs in the U.S., was rolled out gradually in each county between 2008 and 2014, until the entire country was covered (Alsan and Yang, 2018).

In addition, I use the "Merged Outgoing Rotation Groups" (MORG) supplement, which are extracts of the Basic Monthly CPS (conducted during the household's fourth and eighth month in the survey), that contain information on weekly labor hours/earnings as a continuous variable.[18] From this, I construct dummy variables to measure full-time and part-time

[16]The survey contains detailed questions covering the social and economic characteristics of each person who is a household member as of the interview date.

[17]The CPS weights were adjusted in 1996 to account for immigrants (Schmidley and Robinson, 1998).

[18]Income questions in the ASEC supplement are on an annual basis and refer to income received during the previous calendar year.

work if the person is working more than 35 and at most 20 hours per week, respectively. The CPS includes weights that are designed to be representative of the population and eliminate biases attributable to oversampling or attrition. As such, all results use the CPS-provided weights.

From the ASEC, I construct two groups: *natives* (which are defined as those born in the U.S., Puerto Rico, or other U.S. outlying areas, or those born abroad but of American parentage) and *immigrants* (which are defined as any foreign-born individual). It is worthwhile to mention that although I cannot observe unauthorized immigrants from the survey, they were never eligible for SNAP. As such, my immigrant sample consists of all other immigrants including those with permission to remain in the United States but without legal permanent resident status as well as legal permanent residents -LPRs- with less than 10 years of arrival). For this analysis, naturalized citizens and immigrant veterans are grouped as natives since they remained eligible after the welfare reform. I restrict my sample to individuals ages 25-59, whose head of household has at most a high school education since these households are more vulnerable to the effects of the welfare reform. The age cutoff at 59 ensures that I examine those individuals just before the early retirement age for Social Security (62). The rationale for starting the sample at age 25 is to ensure that individuals have an adequate level of labor market experience.

The unit of analysis is at the household level. However, household characteristics such as household immigration status are assigned using the values of the head of household. In this analysis, I do not consider mixed households. In other words, both spouses are categorized as either natives or immigrants. However, I do not restrict the children's birthplace among immigrant households. As such, there will be immigrant households with U.S.-born children, but they were not affected by the welfare reform since they remained eligible.[19]

[19] At worst, the household benefit amount would have declined since the parents lost eligibility.

2.5 Empirical Strategy

2.5.1 Replicating East's (2018b) results

The goal of this analysis is to first replicate East's main results, which estimate the effects of SNAP eligibility on both the extensive and intensive margins of labor supply across four demographic groups. I begin with East's double difference model, leveraging geographic (state of residence) and temporal (time of survey) variation among treated immigrants to identify the effects using the following equation:

$$Y_{ist} = \alpha + \beta Elig_{st} + \delta_1 X_{ist} + \delta_2 C_{st} + \gamma_s + \tau_t + \varepsilon_{ist} \tag{2.3}$$

where Y_{ist} is the outcome of interest for individual i in state s observed in survey year t. $Elig_{st}$ is an indicator variable equal to one if treated immigrants are eligible for SNAP based on their state of residence and year of survey. The coefficient β is the intent to treat effect of SNAP eligibility on the given outcome. One might be concerned that state-specific characteristics could affect the outcomes of interest. To address this concern, a number of state-specific controls are added. In particular, γ_s, which is a vector of state fixed effects to absorb time-invariant characteristics of the state as well as C_{st}, which is a vector of controls for the state unemployment rate in the year of the survey and the year before the survey. I account for time-varying national shocks that may affect the outcomes over time with time-fixed effects, τ_t. Finally, in X_{ist}, I control for demographic characteristics at the household and individual level including age, education, marital status, race, number of years since migration, number of children, and number of children under age 5. As discussed earlier, the static (and dynamic) TWFE DiD specifications are inherently biased if treatment effects are heterogeneous and treatment timing varies.

East (2018b) uses the following triple difference model:

$$\begin{split} Y_{istn} = \alpha + \beta_1 Elig_{st} + \beta_2 Elig_{st} * TI_n + \delta_1 X_{istn} + \delta_2 C_{st} + \delta_3 C_{st} * TI_n + \gamma_s + \tau_t \\ + TI_n + TI_n * \gamma_s + TI_n * \tau_t + \varepsilon_{istn} \end{split} \tag{2.4}$$

where n denotes whether the individual is in the treated immigrant group or the native control group. Since immigrants are inherently different from natives, a separate dummy variable is included to indicate treated immigrants, TI_n. For the same reason, state by year controls interacted with treated immigrants are added, $C_{st} * TI_n$, since economic conditions and state policies may have heterogeneous effects across citizenship status. As additional controls, state by treated immigrant, $TI_n * \gamma_s$, and year by treated immigrant, $TI_n * \tau_t$, fixed effects are added. In this specification, the variables of interest are β_1 and β_2. Notably, the coefficient β_1 captures the effect of treated immigrants' eligibility on native outcomes while the β_2 measures the differential effect of treated immigrants' eligibility on treated immigrants' outcomes in comparison to natives. This implies that the sum of these two coefficients, $\beta_1 + \beta_2$, gives the estimated impact of treated immigrants' eligibility on treated immigrant outcomes. It is important to reiterate that identification in equation (2.4) comes from the variation in treated immigrants' outcomes relative to natives' outcomes that occur across states and over time. Equations (2.3) and (2.4) are estimated using CPS weights and the standard errors are clustered at the state level.

2.6 Results

Table 2.1 displays the replication results for East (2018b). Each column represents a separate regression, incorporating various controls across four demographic groups: married men, married women, single women, and single men. For brevity, only the coefficient on *Elig* (in the double difference model) or *Elig*TI* (in the triple difference model) is reported. Panel A presents the results for the extensive margin of labor force supply, while Panels B to D illustrate outcomes across three different measures of hours worked: usual weekly hours worked, full-time work ($\geq$ 35 hours per week), and part-time work ($\leq$ 20 hours per week).

Consistent with East, my analysis reveals a statistically significant effect of SNAP eligibility on the extensive margin of labor supply for married women, with similar magnitudes but negligible effects for married and single men. Specifically, the coefficient from the model in column (6) indicates that access to SNAP reduces the likelihood of work for married

Table 2.1: East's Replication - Effect of Eligibility on Whether Working Last Week and Hours Usually Work

	Married Men			Married Women			Single Women			Single Men		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A: Work Last Week												
Elig	0.006	-0.008		0.016	-0.004		-0.013	-0.005		0.009	-0.009	
	(0.015)	(0.006)		(0.016)	(0.006)		(0.022)	(0.012)		(0.031)	(0.011)	
Elig ∗ TI		-0.020	-0.12		-0.028	-0.29		-0.004	-0.020		-0.018	0.021
		(0.015)	(0.015)		(0.014)**	(0.015)**		(0.027)	(0.032)		(0.031)	(0.035)
Mean Y	0.82	0.79	0.79	0.40	0.59	0.59	0.55	0.61	0.61	0.78	0.69	0.69
N	38,861	885,295	885,296	49,812	993,420	993,426	17,289	533,152	533,152	13,790	446,763	446,763
B: Weekly Hours												
Elig	-1.143	-0.011		-0.603	-0.234		-0.664	-0.380		-0.134	-0.239	
	(0.322)***	(0.29)		(0.545)	(0.172)		(1.128)	(0.247)		(0.582)	(0.177)	
Elig ∗ TI		-1.049	-1.051		-0.653	-0.661		-0.527	-0.622		0.372	0.182
		(0.317)***	(0.335)***		(0.014)	(0.532)		(1.189)	(1.267)		(0.649)	(0.745)
Mean Y	40.04	40.39	40.39	35.26	34.84	34.84	34.05	28.98	28.98	39.60	39.29	39.29
N	4,660	83,762	83,762	4,741	83,278	83,280	3,071	83,732	83,732	1,898	49,367	49,367
C: Full-time Work (≥35 Hours)												
Elig	-0.067	0.003		-0.043	-0.014		0.009	0.002		0.006	-0.013	
	(0.012)***	(0.004)		(0.030)	(0.011)		(0.062)	(0.010)		(0.024)	(0.009)	
Elig ∗ TI		-0.063	-0.066		-0.015	-0.029		0.050	0.027		0.025	0.029
		(0.013)***	(0.014)***		(0.029)	(0.027)		(0.067)	(0.065)		(0.021)	(0.023)
Mean Y	0.94	0.96	0.96	0.74	0.71	0.71	0.70	0.49	0.49	0.91	0.90	0.90
N	4,660	83,762	83,762	4,786	83,278	83,280	3,094	84,266	84,266	1,898	49,367	49,367
C: Part-time Work (≥20 Hours)												
Elig	0.051	0.001		0.049	0.018		-0.050	-0.026		-0.005	0.008	
	(0.012)***	(0.005)		(0.029)*	(0.013)		(0.054)	(0.013)**		(0.020)	(0.012)	
Elig ∗ TI		0.046	0.046		0.025	0.045		-0.058	-0.086		-0.011	-0.003
		(0.011)***	(0.012)***		(0.026)	(0.022)*		(0.063)	(0.058)		(0.024)	(0.023)
Mean Y	0.05	0.03	0.03	0.19	0.22	0.22	0.21	0.28	0.28	0.07	0.07	0.07
N	4,660	83,762	83,762	4,786	83,278	83,280	3,094	84,266	84,266	1,898	49,367	49,367

Notes: Data are taken from the 1995-2007 ASEC supplement of the CPS. The sample is adults aged 25-59 whose head of household has less than a High School education. Treated immigrants are foreign-born individuals who entered the U.S. before 1997 and less than 15 years prior to being surveyed. *Elig* is equal to one if treated immigrants are eligible for SNAP, based on the year of survey and state of residence. *TI* is equal to one if the individual is in the treated immigrant group. All models include controls for state and year-fixed effects as well as demographic characteristics. The results are weighted using the given CPS weights. Standard errors are clustered by state and shown in parentheses. *$p<0.10$, **$p<0.05$ and ***$p<0.01$.

women by 2.9 percentage points, closely aligning with East's estimate of 2.8 percentage points. For single women (heads of household), the effect is smaller, with a reduction in the likelihood of work by 2 percentage points. This finding contrasts with East's results in two key aspects: East reports a larger effect—a 4.8 percentage point decline in the likelihood of work—and this estimate is statistically significant. I have ruled out differences in sample size as a plausible explanation for the discrepancy and instead focus on the diagnostic model results provided by Callaway and Sant'Anna (2021) (referred to hereafter as CS) for further insights. In Panel A, this discrepancy is the only deviation from East's results. Notably, the null effects of SNAP eligibility on the extensive margin of labor supply for single men align with expectations, given their typically low rates of SNAP participation. Furthermore, the analysis reveals no significant impact of SNAP eligibility for treated immigrants on the labor supply of native individuals. This finding suggests that unobserved time-varying trends within states are unlikely to confound the observed effects, reinforcing the *validity* of the results. However, as discussed, while TWFE models generally facilitate "good" comparisons between treated and untreated units, they face challenges due to the presence of "bad" comparisons among units that have already received treatment. These problematic comparisons contribute to negative weighting issues, which can undermine the credibility of the estimates. Therefore, it cannot be entirely ruled out that the estimates may be compromised.

I now turn to the results in Panels B to D, where the Earner Study/ORG supplement is employed, resulting in a smaller sample size compared to Panel A. Focusing first on married men using the preferred specification in column (3), the effect of SNAP eligibility on usual weekly hours worked is statistically significant, indicating a reduction of 4.8% in weekly hours—mirroring East's estimate. Additionally, the likelihood of working full-time decreases by 7%, and the likelihood of working part-time decreases by 4%, both figures also aligning with East's findings. These results are particularly insightful given the insignificant findings for married men along the extensive margin.

The results for married women, single women, and single men along the intensive margin are less striking. While both married and single women appear to reduce their usual weekly hours and the likelihood of working full-time, these effects are not statistically different from zero. Furthermore, the estimated effects for single men are inconsistent in direction, negligible in magnitude, and statistically insignificant across various specifications. The limitation of the TWFE design may be more evident here, given the smaller sample size of the ORG supplement. The reduced sample size can amplify the influence of poor comparisons, leading to biased and inconsistent overall effects. To address these limitations, I now turn to the difference-in-differences estimator proposed by Callaway and Sant'Anna (2021), which is designed to handle multiple time periods and heterogeneous treatment effects.

2.6.1 The diagnostic approach of Callaway and Sant'Anna (2021)

Several diagnostic tools are available to address the limitations of TWFE regressions in settings with staggered treatment timing. Among these tools, the estimator proposed by CS is widely regarded as comprehensive, addressing several key issues through the statistical packages developed by the authors.[20] These packages, `csdid` and `did`, are available for implementation in `Stata` and `R`, respectively, and offer the following capabilities:

i. Accommodating multiple periods and variations in treatment timing; and

ii. Providing disaggregated group-time average treatment effects, event-study type estimates, as well as overall treatment effect estimates.

Before presenting the results from the CS estimator, it is essential to address how the authors overcome the negative weighting issue that is inherent in Two-Way Fixed Effects (TWFE) models. The CS estimator mitigates this problem by utilizing all possible "clean" 2x2 comparisons, where never-treated or not-yet-treated groups are used as valid control groups for estimating the ATT. These comparisons are aggregated to estimate the overall treatment effect. A straightforward illustration, using notation from Roth et al. (2022), is the group-time average treatment effect on the treated, $ATT(g, t)$, which represents the average treatment effect at time t for units first treated at time g. For instance, in the SNAP context, $ATT(1997, 1999)$ denotes the average treatment effect in 1999 for states that

[20]The authors have also provided several accompanying resources, including vignettes, recordings of discussions on the implementation of their estimators, and examples with real data.

initially restored SNAP eligibility to immigrants in 1997. Roth et al. (2022) emphasize that the identification of $ATT(g,t)$ relies on "clean" comparisons, which involve comparing the expected change in outcomes for the treated group g between periods $g-1$ and t with that of a control group that has not yet been or was never treated at time t.[21]

I utilize the `did` package to estimate the group-time average treatment effects of SNAP eligibility on the extensive margin of labor supply across four demographic groups of treated immigrants. This approach involves reestimating Panel A in Table 2.1 using the CS estimator. To illustrate the issues associated with TWFE models, I present both estimators in Table 2.2. Specifically, the overall ATT from the CS estimator is displayed in the even-numbered columns (2, 4, 6, and 8), while the estimates from the TWFE DiD are shown in the odd-numbered columns (1, 3, 5, and 7) for the four demographic groups.

Table 2.2: Effect of eligibility on whether working last week using the TWFE DiD and CS estimators

	Married Men		Married Women		Single Women		Single Men	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A: Work Last Week								
Elig	0.006	-0.020**	0.016	-0.010**	-0.013	-0.005**	0.009	-0.007**
	(0.015)	(0.003)	(0.016)	(0.001)	(0.022)	(0.012)	(0.031)	(0.001)

Notes: Data are taken from the 1995-2007 ASEC supplement of the CPS. The sample is adults aged 25-59 whose head of household has less than a High School education. Treated immigrants are foreign-born individuals who entered the U.S. before 1997 and less than 15 years prior to being surveyed. *Elig* is equal to one if treated immigrants are eligible for SNAP, based on the year of survey and state of residence. Point estimates in the odd-numbered columns are derived from the TWFE DiD model, whereas those in the even-numbered columns are obtained using the CS estimator. The TWFE models include controls for both state and time fixed effects, as well as demographic characteristics, and are weighted using the provided CPS weights. Standard errors for all models are clustered by state and shown in parentheses. **p<0.10*, ***p<0.05* and ****p<0.01*.

Overall, the findings reveal that the pattern of labor supply disincentives is consistently observed across demographic groups when using the CS estimator, contrasting sharply with the TWFE DiD estimates. Notably, the TWFE estimates exhibit oppositely signed coefficients and larger magnitudes across all demographic groups except married men. This

[21] The CS estimator generalizes the parallel trends and no anticipation assumptions, including scenarios where the parallel trends assumption holds only conditionally on covariates.

discrepancy suggests that the TWFE DiD estimates may be upwardly biased in assessing the impact of SNAP eligibility on the likelihood of work among treated immigrants. This bias is particularly concerning if policy decisions regarding immigrants' access to SNAP rely on models that use TWFE DiD, especially when the treatment effects vary.

The CS estimator indicates a modest reduction in the likelihood of employment, where the largest margin of response is observed among married men, showing a 2 percentage point decrease in their likelihood of working. Single women and single men exhibit the smallest margin of response, with a 0.5 and 0.7 percentage point decline in the likelihood of employment, respectively. The result for low-educated single men is particularly notable because East (2018b) reports that this demographic group has very low SNAP participation rates—just 6% in her sample—and did not anticipate any effect of eligibility on their labor supply. In contrast, East finds that single women have the highest SNAP participation rates, at 31%, and anticipated the largest effect on this group. These economic patterns observed in the summary statistics were not revealed in the CS estimates, even after controlling for biases in the TWFE DiD model estimated in equation 2.3. This may provide weakly suggestive evidence that the choice of diagnostic estimator is important for addressing bias. As noted, the CS estimator aggregates all possible 2x2 comparison groups. Although clean control groups are used, I employ repeated cross-sectional data, leading to many of these 2x2 groups having relatively small sample sizes, given that there are a total of 6,623 early and late treated units. Consequently, additional estimators are necessary to validate and better understand the true effects of SNAP eligibility on immigrants' labor supply.

Other estimators, known as imputation estimators, such as those proposed by Borusyak et al. (2024) and Liu et al. (2024), utilize only untreated observations as controls to estimate group and period fixed effects. In particular, Borusyak et al. (2024) derive the weighted sums of estimated treatment effects by imputing the untreated potential outcomes from the fixed effects. This approach results in a robust and unbiased estimator that remains efficient even under unrestricted treatment effect heterogeneity and when some restrictions are imposed on treatment effects. It addresses a limitation of the CS estimator and similar estimators, for which the efficiency properties are less well understood.[22]

The estimator proposed by Borusyak et al. (2024) is notably versatile and applicable to settings with time-varying controls, triple-difference designs, and certain non-binary treatments. It would be particularly suitable for comparing the TWFE triple-difference model in equation 2.4 and for providing additional insights into the differences between CS and TWFE DiD estimates of the labor supply response among treated immigrant single men and women. I plan to explore this comparison in future research.

2.7 Conclusion and Future Work

This paper evaluates the impact of changes in immigrants' eligibility for the Supplemental Nutrition Assistance Program (SNAP) on both the extensive and intensive margins of their labor supply. SNAP is one of the largest safety net programs in the U.S.; however, causal evidence on its impact on labor costs is limited, likely due to the lack of variation in federal program design and welfare policies. Consequently, understanding both the costs and benefits of this program is crucial for designing optimal policy.

The shift in immigrant eligibility under the Welfare Reform Act of 1996 provides a well-suited quasi-experimental setting for estimating the effects of SNAP. This setting is characterized by staggered treatment adoption, as certain states altered their eligibility rules for immigrants at different times. Specifically, after the 1996 welfare reform restricted SNAP access for foreign-born individuals, nine states restored eligibility in 1997 and 1998, prior to the federal restoration of eligibility in 2003. Recent advancements in econometric methods have demonstrated that the TWFE DiD design becomes invalid when treatment timing varies.

Employing newer econometric techniques and leveraging the staggered changes in immigrants' SNAP eligibility rules, I find that estimates from the TWFE DiD model exhibit opposite signs and are upwardly biased. By utilizing the unbiased estimator proposed by

[22] An additional advantage of imputation estimators is the efficiency gained from leveraging more information to estimate unit and period effects.

Callaway and Sant'Anna (2021), I reveal modest reductions at the extensive margin of labor supply when immigrants have access to SNAP.

These findings are particularly timely given the ongoing discussions about further extending SNAP benefits to immigrant households. The evidence from this analysis provides valuable insights into the potential impacts of broadening SNAP eligibility for immigrants on a wider scale. I argue that practitioners and policymakers should adopt not only a data-driven approach to policymaking but also a robust methodological framework that incorporates techniques at the frontier of the field. One pertinent question to explore is whether eligibility leads immigrants to change jobs or alter their labor market behavior in other ways.

Chapter 3

Effects of Large Oil Shocks in the Presence of Downward Nominal Wage Rigidity: Evidence using Micro-data

3.1 Introduction

Hamilton (1996, 2003, 2011) finds that oil price shocks have asymmetric effects on key macroeconomic variables. Specifically, Hamilton's "net oil price increase" model predicts that a rise in oil prices generates a larger drop in economic activity when oil prices advance above their highest level in the previous three years. Bachmeier and Keen (2023) argue that macroeconomic variables respond asymmetrically to oil price shocks because some workers are subject to a downward nominal wage rigidity (DNWR). When wages are fully flexible, a large rise in oil prices reduces labor demand by so much that the nominal wage declines. A DNWR constraint, however, prevents the nominal wage from falling for some workers. The firms that employ those workers respond by further reducing their output and labor demand, which leads to oil price shocks having asymmetric effects.

The macroeconomic literature on DNWR is extensive and includes models that examine its impact on various topics such as the inflation-unemployment trade-offs (Daly and Hobijn, 2014), asymmetric responses (Abbritti and Fahr, 2013), the output-inflation trade-offs (Benigno and Antonio Ricci, 2011), optimal monetary policy (Abo-Zaid, 2013), and the optimal steady-state inflation rate (Fagan and Messina, 2009). These macroeconomic models include a DNWR constraint because many individual-level studies find nominal wages decline infrequently, indicating some degree of stickiness in those wages. Despite variations in measures of wage rigidity based on data sources, sample periods, and wage periodicity, early individual-level studies have confirmed the presence and extent of DNWR in the U.S. (Akerlof et al., 1996; Altonji and Devereux, 2000; Card and Hyslop, 1997; Gottschalk, 2005; Kahn, 1997; Lebow et al., 1995).[1] More recent literature investigates the degree of nominal wage stickiness by examining the cyclicality and persistence of nominal wage change distributions over time (Fallick et al., 2016; Jardim et al., 2019; Kurmann and McEntarfer, 2019).[2]

Our paper adopts a different perspective, shifting the focus from identifying the extent of DNWR to assessing its broader implications. Specifically, we examine how profit-maximizing firms adjust their demand for different types of labor—blue-collar versus white-collar—under two conditions: (i) the presence of a binding DNWR constraint and (ii) the impact of an exogenous oil shock. The response of these firms to a downwardly rigid nominal wage constraint, in the context of rising marginal costs of production due to an oil price shock, remains unclear. Whether firms' cost-saving strategies involve reducing the labor services of specific types of workers—and the underlying reasons for such decisions—remains an open question. Understanding the implications of the oil shock and its interaction with DNWR for both workers and firms is essential for formulating effective policy responses. We argue that explanations for the macroeconomic response to oil price shocks may benefit from emphasizing the central role of overtime compensation as a key margin where firms adjust

[1]Kuroda and Yamamoto (2007) provide a comprehensive review of studies from industrialized countries confirming DNWR using micro-data since the 1970s.

[2]Two particularly relevant papers in this domain are Jo (2024) and Grigsby et al. (2021). Jo (2024) utilizes household survey data to show that the share of workers experiencing no wage changes is more countercyclical compared to the share facing wage cuts. In contrast, Grigsby et al. (2021) analyze administrative payroll data and find that nominal wage declines for job stayers are significantly less frequent than previously suggested by the literature, with over one-third of workers experiencing no change in their annual base wage.

production costs.

We start by showing a simple approach commonly employed in the literature to examine the existence of DNWR. Utilizing data from the Current Population Survey (CPS) from 1982 to 2018, we analyze year-over-year changes in nominal wages for hourly-paid workers. Consistent with previous studies, we find that, on average, wages are more likely to remain at zero than to decrease over the sample period. We then use simulations to illustrate how a DNWR constraint affects the response of the wage distribution before discussing the expected effects on firms' demand for labor services.

Using the demand for and supply of labor model, we show that an exogenous increase in oil prices leads firms to reduce their demand for labor services. A binding DNWR constraint implies that real wages partially decline, resulting in a reduction in employment, as measured by labor hours in our paper. Based on the model's predictions, we analyze the response of nominal wage growth, hours worked, and overtime pay using microdata from the Current Population Survey (CPS) covering the years 1982 to 2018. To our knowledge, previous research has not addressed the impact of large oil price increases on the margin of adjustment for firms at the individual level.

Our key finding, which contributes new insights to the literature, is that the primary adjustment margin for firms in response to large oil price increases is overtime pay. Additionally, our analysis reveals that hours worked decline in most industries following significant oil price increases, consistent with firms' reluctance to reduce wages. Finally, we did not find evidence suggesting that large oil price increases affect the distribution of nominal wage growth or hours worked. These null results may reflect limitations within our dataset rather than a true lack of effect.

The remainder of the paper is organized as follows: Section 2 provides an overview (DNWR) and its effect on our main outcomes of interest. Section 3 details the methodological framework employed in the analysis, while Section 4 presents and discusses the results. Finally, Section 5 offers concluding remarks and implications of the findings.

3.2 Background

The notion that wages are more rigid downward than upward has been established in several studies using microdata (Card and Hyslop, 1997; Kahn, 1997; Lebow et al., 2003; McLaughlin, 1994 and Daly and Hobijn, 2014).[3] DNWR implies that nominal wages do not easily adjust downwards due to various frictions, such as contractual agreements, worker morale, and institutional settings. Indeed, the archetypal evidence of sticky wages is a large spike at zero in a wage change distribution for a given period, accompanied by less frequent wage declines than wage increases.

This pattern is evident in Figure 3.1, which plots the distribution of the annualized change in a worker's nominal hourly wages from 1983 to 2018 using the Current Population Survey (CPS) data. The spike at zero represents the fraction of workers whose year-over-year growth rate in nominal hourly wages was zero, accounting for nearly 30 percent of the wage change distribution. The observations to the right of zero denote the share of workers with positive nominal growth rates (hourly wage raises), while those to the left of zero depict the fraction of workers facing negative nominal growth rates (nominal hourly wage cuts). Over the sample period, 48 percent of workers experienced wage increases, while 23 percent faced wage cuts on average (see Table A.1).[4] Figure 3.1 shows that wage increases were relatively more common than wage cuts, which is unsurprising given the period's technological growth and average inflation rate of 2.7 percent.

At the same time, another strand of literature has moved beyond validating the inertia in nominal wage rates to assessing the frequency of changes in observed wages, which is a crucial factor in determining labor adjustment costs. Notably, recent work by Barattieri et al. (2014) utilizes hourly wage data from the Survey of Income and Program Participation

[3]The empirical macroeconomics literature has extensively documented the slow cyclical adjustment of wages using aggregate wage data (see Calvo, 1983; Taylor, 2016 and Erceg et al., 2000 for alternative wage-setting schemes). Christiano et al. (2005), Smets and Wouters (2005), and Beraja et al. (2019) offer a theoretical overview of the dynamic stochastic general equilibrium (DSGE) models incorporating nominal wage rigidities.

[4]Jo (2024) provides a survey of other studies that report the variations in the proportion of workers experiencing wage increases, cuts, or no wage changes. Our estimates in Table A.1 are broadly consistent with these studies.

Figure 3.1: Distribution of annual nominal hourly wage changes, 1983 - 2018

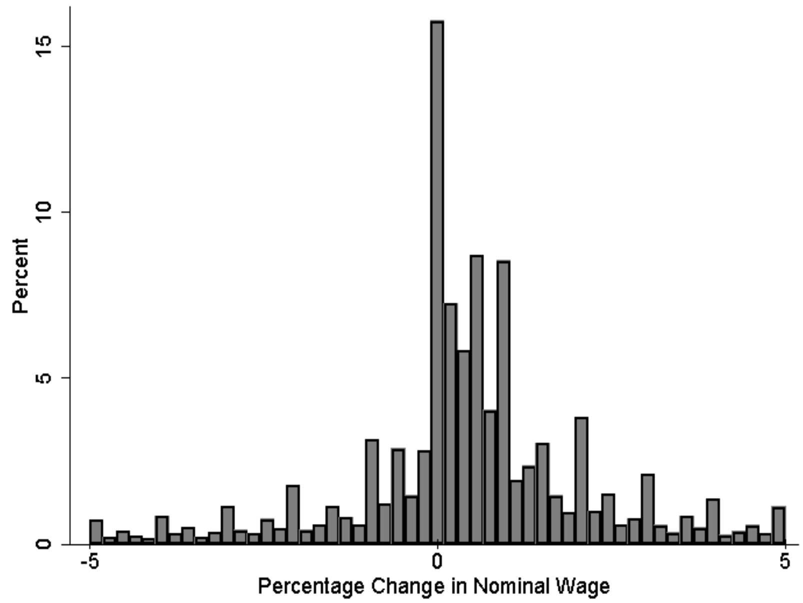

Source: CPS and authors' calculations. The histogram illustrates the year-over-year growth in nominal wages over the sample period. The nominal wage changes are expressed in percentage points and calculated for the full sample of hourly workers (233,713 observations), where each observation represents a worker-year. The tallest bar, located at zero, represents the fraction of workers who experienced no change in their nominal hourly wage over the period. Observations to the right(left) of zero indicate workers who experienced wage increases(decreases).

(SIPP), which tracks individuals for up to 48 months with interviews every four months, adjusting for substantial measurement error. Their main results indicate that, on average, the quarterly probability of a change in nominal wages for hourly-paid workers is 21.1%, with changes in wages being much more likely when workers change jobs.

Taken together, the pattern displayed in Figure 3.1 and the findings of Barattieri et al. (2014) enhance our understanding of wage-setting behavior by firms. In this context of nominal wage stickiness, a pertinent question arises: How do firms adjust their labor demand in response to a real aggregate shock, such as an exogenous oil price increase?

In this section, using simulated data, we first demonstrate how an oil shock affects the wage distribution under the assumptions of flexible and downwardly rigid nominal wages. We then discuss the expected effect of an oil shock on firms' demand for labor with DNWR.

3.2.1 *Distribution of Wages*

The simulations demonstrating how a DNWR constraint affects the response of the wage distribution to an oil price shock are intended to be expository. They aim to provide motivation for the empirical work and guide its interpretation rather than serve as an accurate empirical depiction of the U.S. economy. By illustrating the theoretical implications of DNWR, these simulations help to highlight the potential mechanisms through which wage rigidity can amplify the effects of an oil shock on the labor market.

Specifically, the change in the fully flexible nominal wage, $\Delta w\tilde{a}ge_t^-$, is always less than or equal to the change in actual nominal wage, $\Delta wage_t^-$, when some workers face a DNWR:

$$\Delta wage_t^- \geq \Delta w\tilde{a}ge_t^- \tag{3.1}$$

DNWR is present if the nominal wage of some workers falls less in response to a negative economic shock than it would in an economy with fully flexible nominal wages. In the extreme case where all workers are constrained by the DNWR, the change in the nominal wage, $\Delta wage_t$, must be non-negative:

$$\Delta wage_t \geq 0. \tag{3.2}$$

3.2.2 *Effect of an oil shock with flexible nominal wages*

The standard Keynesian model predicts that a negative supply shock—in our case, an increase in oil prices—raises production costs, compelling firms to increase prices and reduce output. This scenario occurs because higher oil prices, a critical input for many industries, directly increase the cost of production. Firms, facing higher costs, pass these onto consumers in the form of increased prices to maintain their profit margins. Consequently, the higher prices reduce aggregate demand, as consumers' purchasing power diminishes. This reduction in demand leads firms to cut back on production, further exacerbating the economic slowdown, similar to the economic fluctuations experienced in the mid and late 1970s.

Moreover, this decline in output decreases labor demand, which drives down the real wage. If the fall in the real wage outweighs the rise in the price level, then the nominal wage will decrease in the economy, absent a DNWR constraint (Blanchard and Galí, 2009 and Bachmeier and Keen, 2023). How will the full distribution of wage growth respond to an oil price shock, Δoil_t, in the absence of DNWR?

Assume nominal wage growth, Δw_{it}, takes the form:

$$\Delta w_{it} = \beta \Delta oil_t + X_{it} + \xi_{it}, \tag{3.3}$$

where β is the sensitivity (elasticity) of wage growth for worker i in period t to an oil price shock, X_{it} is the set of factors other than the oil shock that influence wage growth in period t, and $\xi_{it} \sim N(0, 1)$. For example, $\beta = -0.1$, implies that wage growth decreases by one percentage point for each ten percentage point increase in the price of oil.

Figure 3.2 shows the simulated quantiles of wage growth when $\beta = -0.1$ and $\Delta oil_t = 0.10$, 0.00, or -0.10 (a 10% positive oil price shock, no oil shock, and a 10% negative oil price shock).[5] The only effect of a change in the price of oil is shifting the mean. A higher oil price shifts wage growth distribution to the left by the same amount at all quantiles. Figure 3.3 shows that a 10% increase in the oil price causes the nominal wage growth rate to fall by 1% at all quantiles. This constant decline is not surprising given that the linear model estimated in equation 3.3 assumes that all workers experience the same percentage decline in the wage rate after an oil price increase.

3.2.3 *Effect of an oil shock with DNWR*

We now introduce two modifications to the model from the previous section. First, the nominal wage cannot fall for a fraction of workers, ϕ, due to a binding DNWR constraint. Second, the response of wage growth to an oil shock varies across different workers. In other words, the presence of a DNWR constraint for some, but not all, workers implies that the nominal wage growth rate will respond heterogeneously to an oil price shock.

[5] The figure is based on simulating wage growth for 1,000,000 workers.

Figure 3.2: Quantiles of wage growth conditional on an increase in oil prices

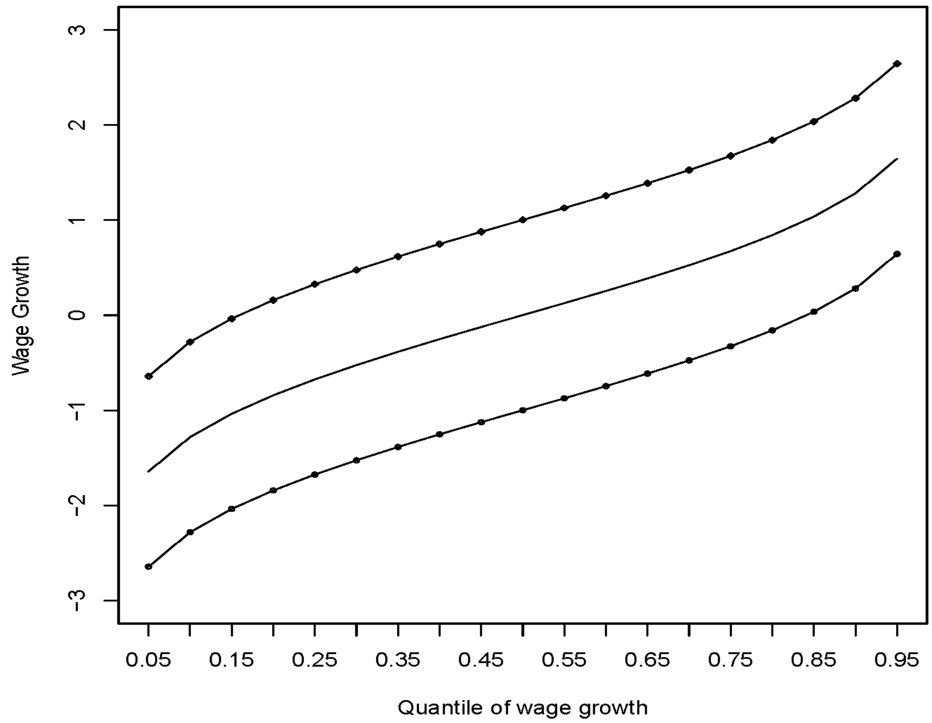

Notes: The figure illustrates the simulated wage growth by quantiles, conditional on an oil price increase, based on a sample of 1,000,000 workers for $\beta = -0.1$ from equation 3.3. The predictions are derived from a linear model.

For workers not subject to the DNWR constraint, wage growth is determined by:

$$\Delta w_{it} = \beta oil_t + X_{it} + \xi_{it}, \tag{3.4}$$

and for workers subject to the DNWR constraint, it is:

$$\Delta w_{it} = max\left(0, \Delta\widetilde{w}_{it}\right), \tag{3.5}$$

where

$$\Delta\widetilde{w}_{it} = \beta_i oil_t + X_{it} + \xi_{it}, \tag{3.6}$$

Figure 3.3: Effect of a 10% increase in oil prices on wage growth

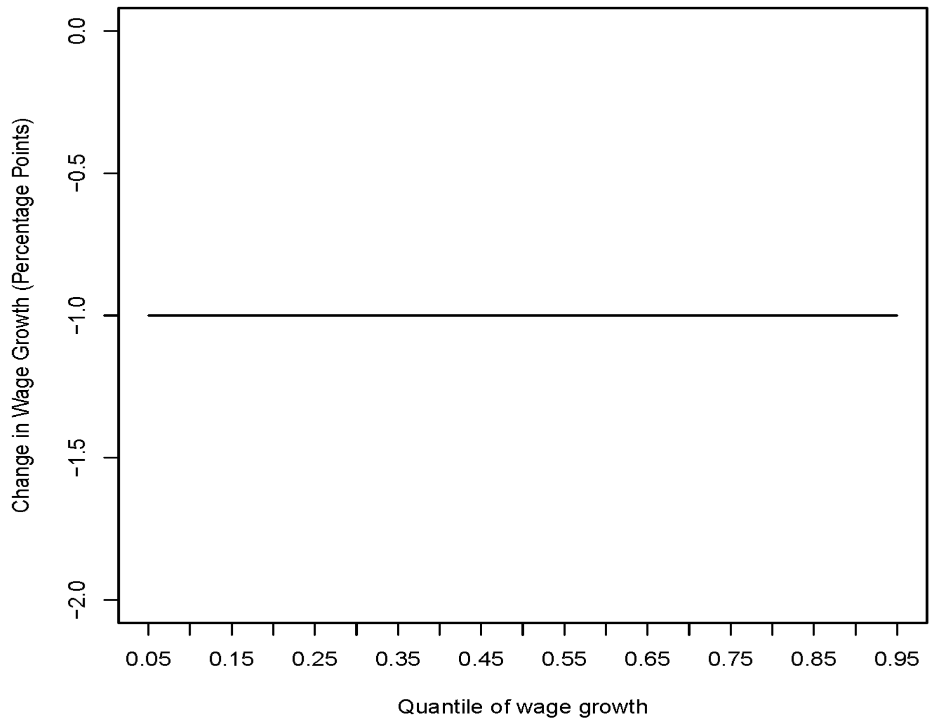

Notes: The figure shows the simulated changes in wage growth in response to a 10% increase in oil prices, as predicted by a linear model.

and

$$\beta_i \sim U\left(\beta^{min}, 0\right); \xi_{it} \sim N\left(0, 1\right). \tag{3.7}$$

The effect of an oil price shock on the distribution of wage growth is a nonlinear function of the distribution of $\widetilde{w}_{it}$, the percentage of workers subject to the DNWR constraint, and the distribution of β_i.

Figure 3.4 presents the simulated quantiles of wage growth for $\phi = 0.6$, $\beta^{min} = -0.2$, and $\Delta oil_t = 0.10$, 0.00, or -0.10. The flat part of each distribution is the set of quantiles for which the DNWR constraint binds for the workers experiencing a sufficiently low ξ_{it}. As the price of oil increases, the constraint binds for a larger portion of workers, with a smaller share seeing positive wage growth. Figure 3.5 shows the effect of the positive oil shock ($\Delta oil_t = +0.10$ vs. $\Delta oil_t = 0.00$) at different quantiles of wage growth. The far-right captures the workers experiencing large positive ξ_{it} and β_i close to zero. The far-left captures workers sensitive to oil shocks, who experience large negative ξ_{it} shocks and do not face the DNWR constraint. The oil price increase does not affect wage growth in the middle quantiles, from about 0.35 to 0.52, because those workers face a binding DNWR constraint with or without the oil price shock. The nominal wage growth rate for the other quantiles either remains the same or experiences a more moderate change. The specifics of this depend on whether the DNWR constraint holds for a combination of β_i and ξ_{it}.

Figure 3.4: Wage growth by quantile conditional on the change in oil prices

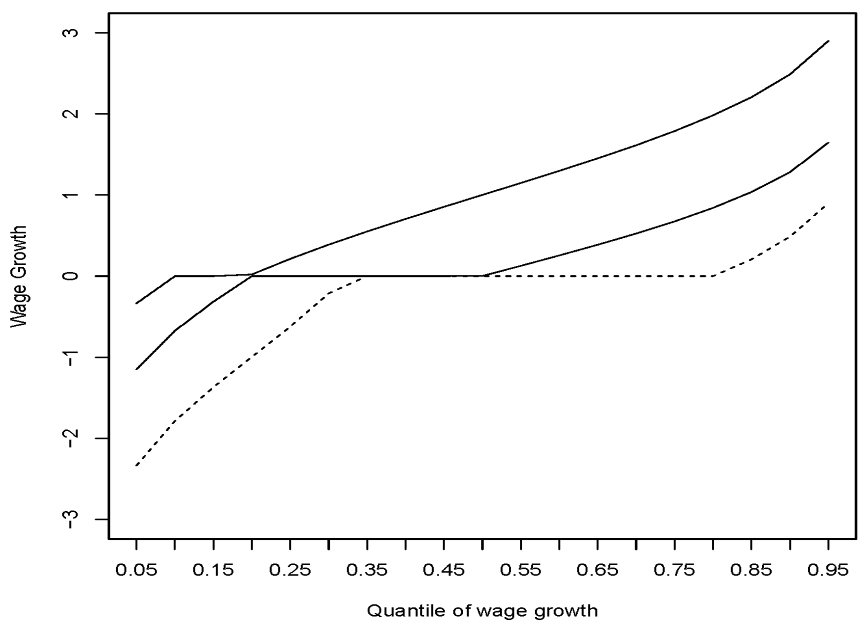

Notes: The figure illustrates the heterogeneous response of wage growth across various quantiles to changes in oil prices. The simulations are based on a fraction of workers constrained by DNWR, $\phi = 0.6$, and the changes in oil prices of $\Delta oil_t = 0.10$, 0.00, or -0.10.

3.2.4 *Expected effect of an oil shock on employment with DNWR*

Following our simulation analysis from the previous section, the decline in real wages due to an oil shock can have several repercussions for employment. Consider the case when wages

Figure 3.5: Effect of a positive oil price shock on wage growth by quantile

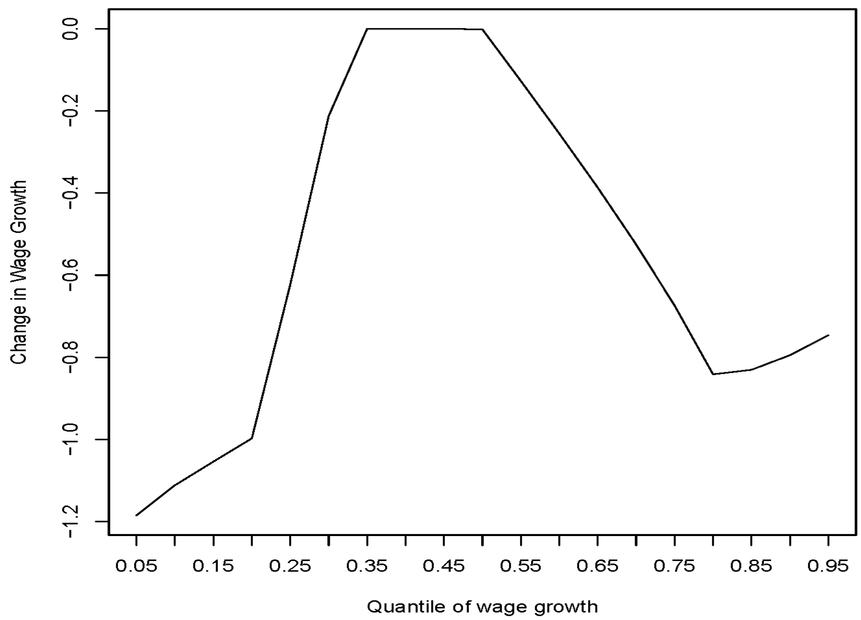

Note: The figure shows the effect of the positive oil shock ($\Delta oil_t = +0.10$ vs. $\Delta oil_t = 0.00$) at different quantiles of wage growth.

are flexible. As real wages fall, the substitution effect may lead workers to supply less labor services since the opportunity cost of leisure decreases, potentially causing individuals to choose more leisure over labor. On the other hand, the income effect might compel workers to increase their labor services to compensate for the loss in real income. Workers may need to work more hours or take on additional jobs to maintain their standard of living. As such, for workers faced with a non-binding DNWR constraint, the overall response of employment to a decline in real wages is ambiguous and depends on the relative strength of the substitution and income effects.

However, the presence of a binding DNWR constraint complicates this scenario. As shown in the previous section, the oil price shock does not affect the wage growth for the fraction of workers who are constrained by DNWR. With DNWR, workers cannot easily negotiate lower nominal wages even if they are willing to accept them to retain their jobs as firms adjust to the oil shock. Daly and Hobijn (2014) emphasize that these workers are trapped with a nominal wage higher than what would prevail under a flexible wage regime.

At this wage, these workers will provide fewer labor services than intended due to the oil shock. In other words, we anticipate that workers will unambiguously reduce their labor services along the intensive margins when constrained by DNWR.[6]

Moreover, we hypothesize that DNWR may introduce heterogeneity across occupations in terms of employees' responses at the intensive margin. While we provide a detailed discussion of worker classification by occupation in Section 3.3.1, it is useful to mention it here to enhance the coherence of our argument. We categorize workers as "blue-collar" and "white-collar" based on the occupational classification used by (Autor and Dorn, 2013). Faced with downwardly sticky nominal wages, we expect that adjustments along the intensive margin will be more pronounced for one type of worker.

The nature of blue-collar occupations makes them more susceptible to automation and substitution by capital services compared to white-collar occupations. In response to a positive oil price shock, profit-maximizing firms are compelled to economize by reducing factor inputs, including labor. Consequently, the optimal mix of cost-saving inputs chosen by the firm remains uncertain *ex-ante*. Since our analysis does not address firms' demand for capital services, it would be speculative to assert that firms will substitute capital with blue-collar workers. Instead, we hypothesize that the aggregate reduction in labor hours may disproportionately impact one type of worker.

Figure 3.6 depicts the impact of an exogenous oil shock on firms' demand for labor services. Firms reduce their demand for labor services from L^D to $L^{D'}$ in response to a positive oil price shock, which also increases the overall price level. A binding DNWR constraint implies that the real wage rate cannot adjust to the new competitive equilibrium at ω_1^*. Instead, the rising price level and the binding DNWR imply the real wage can only decline to ω^i, which is equivalent to the size of the increase in the price level. Consequently, hours worked decreases from $L^*_{N_0}$ to L^i_N, reflecting a larger reduction in hours worked than would occur if wages were allowed to adjust flexibly.

In the next section, we discuss our data and introduce the econometric model we use to explore the effect of an oil shock under the condition of DNWR on wages and hours worked.

[6] Intensive margin adjustments relate to the number of hours worked by those employed.

Figure 3.6: Labor market response of a positive oil shock

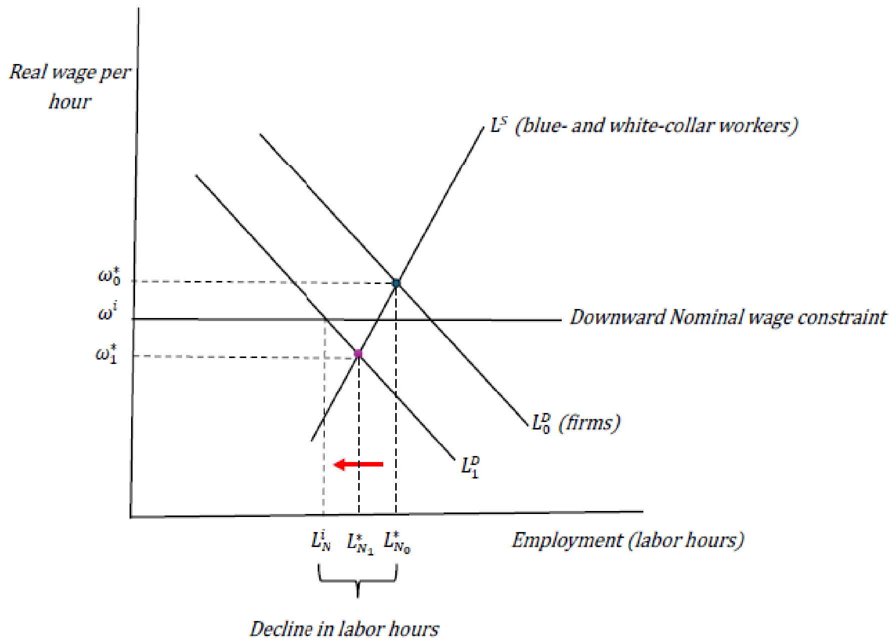

Notes: On the vertical axis is the real wage rate—the nominal wage divided by the price level—while the horizontal axis represents labor employed, expressed as aggregate hours worked per week. The figure illustrates the labor market response to a positive oil price shock, which causes firms to reduce their demand for labor services from L^D to $L^{D'}$. A binding DNWR constraint implies that wages cannot adjust to the new competitive equilibrium wage at ω_1^*. Instead, wages decline partially to ω^i. Consequently, hours worked decreases from $L_{N_0}^*$ to L_N^i, reflecting a larger reduction in hours worked than would occur if wages were allowed to adjust flexibly.

3.3 Methodology

3.3.1 Data

We use data from the Current Population Survey (CPS) for our analysis (Flood et al., 2023). This nationally representative and publicly available survey of over 65,000 households provides individual-level data on demographics, education, and a wide selection of labor force statistics such as employment status and wages.[7] One advantage of the CPS disaggregated wage data is that it bypasses the inherent composition bias found in aggregate wage statistics.

[7]The CPS is commonly used in studies that examine aggregate wage growth in the U.S. (see for example Card and Hyslop, 1997; Elsby, 2009; Solon et al., 1994 and Daly and Hobijn, 2022).

The CPS provides several self-reported statistics on wage and salary earnings by periodicity, worker class, and occupation.

We utilize the basic monthly samples along with the specialized supplement known as the Outgoing Rotation Group (ORG)/Earner Study. This supplement includes information on hourly earnings and work schedules that are not included in the basic monthly surveys. Notably, the ORG asks additional questions to households in months four and eight of the sample rotation. These questions are more explicit in the periodicity for which respondents find it easiest to report their earnings. In our analysis, the respondents are characterized as hourly workers according to the ORG variable `hourwage`, where the respondents declare how much they earned per hour in their current job.[8] An additional question is asked if individuals report a different periodicity of earnings other than hourly. The question determines whether they are indeed paid by the hour and, if so, the hourly rate.[9] One can also use other ORG variables to impute hourly wages for non-hourly workers by taking the ratio of weekly earnings (`earnweek`) to usual hours worked (`uhrsworkorg`). Borjas (1980), however, showed that this approach leads to division bias, which arises from potential reporting errors in the weekly hours worked. As a result, our empirical analysis is restricted to hourly-rated workers to circumvent the division bias.

The distinctive structure of the CPS facilitates both cross-sectional and panel analysis. For the latter, we leverage the CPS rotation pattern to construct linked one-year panels of individuals. Specifically, the basic monthly CPS employs a 4-8-4 rotation pattern, where U.S. households are surveyed intermittently over a total of eight months. Respondents are first interviewed consecutively for four months, then excluded from the sample for the following eight months, after which they are surveyed again for another four months. The months during which households are interviewed are referred to as "month-in-sample" (`MISH`) and are numbered from one to eight, i.e., `MISH 1` through `MISH 8`. We utilize this sampling design and follow the innovative method of linking participants across months as outlined by Rivera Drew et al. (2014) to exploit the longitudinal features of the data. Notably,

[8] The variable `hourwage` excludes other types of compensation such as overtime pay, tips, or commissions. Therefore, only the base pay rate is used in our analysis.

[9] This question is coded as a dichotomous variable, "paidhour", in the Earner Study questionnaire.

the microdata from IPUMS-CPS provides unique identifiers for each household member throughout the rotation pattern.[10] We use these identifiers to match individuals across monthly samples based on validation characteristics such as age, race, and sex.[11]

We use monthly samples from January 1982 to January 2018 and consider workers between 16 and 64 years of age. Our linking criteria ensure that we observe individuals over the 8-month interview period. Specifically, we link individuals from month-in-sample 1 to 4 and again when they are interviewed during month-in-sample 5 to 8. This structure is illustrated in Table A.2 and implies that we observe an individual for a year. The importance of this structure for our analysis is twofold: it allows us to ascertain any changes in an individual's labor force indicators from their first interview to their last, and it helps us circumvent a known limitation of the CPS, which does not track employer information, making it challenging to clearly define who remained with the same employer or who switched to a new employer. For individuals in `MISH 4` and `MISH 8`, the CPS classifies them as "outgoing rotation groups" since they will not be interviewed during the following month. Since we are interested in assessing the response of wages and hours worked to an oil shock, we rely on the more accurately classified labor force information provided in MISH 4 and 8, known as ORG variables.

Following the Federal Reserve Bank of Atlanta's methodology for constructing their job-stayer series, we validate job stayers by comparing the respondent's reported industry and occupation when they are first observed in the sample (MISH 1) to the final month of the survey rotation (MISH 8). If the respondent has a different occupation code and industry classification, we conclude that the worker has changed jobs and exclude them from our main analysis. Therefore, we focus on the ORG variables for employed job stayers in a given occupation.[12] Our final linked sample averages nearly 7,400 observations annually over the 35-year sample period.

[10]The key identifiers are `CPSID` and `CPSIDP`. `CPSID` is a 14-digit number that uniquely identifies households across months, and `CPSIDP` identifies individuals for every instance they appear in the CPS.

[11]Participants who cannot be linked across months are discarded.

[12]Kambourov and Manovskii (2009) demonstrate that the costs of switching occupations are significant, highlighting a 12 percentage point difference between the weekly earnings losses of displaced workers who switched occupations compared to those whose subsequent jobs were in the same occupation.

Considering the differential impacts revealed by our simulated results of an oil shock on workers constrained by DNWR compared to those unaffected by this wage stickiness, we further categorize workers into "blue-collar" and "white-collar" occupations. We utilize the occupational classification of Autor and Dorn (2013) where blue-collar occupations encompass roles involving routine-manual tasks. This includes mining and oil extraction, construction, and manufacturing occupations. Conversely, white-collar occupations are defined by nonroutine cognitive tasks and include executive, administrative, and managerial occupations.[13] This dichotomy allows us to explore how an oil price shock differentially impacts these occupational categories.

Our main sample comprises hourly-paid employees in both blue-collar and white-collar occupations who have not changed jobs over the past year. Summary statistics, as presented in Table 3.1, reveal that the average hourly wage and labor force participation rate are higher for white-collar employees compared to their blue-collar counterparts. Additionally, white-collar workers have lower average weekly hours worked and are less likely to have experienced wage cuts during the sample period, suggesting greater downward stickiness in nominal wages among these workers. Conversely, blue-collar workers are more likely to experience zero wage changes over the same period.

3.3.2 The Oil Shock

A critical step in the analysis is selecting an appropriate measure of oil price shocks. Bachmeier and Keen (2023) advocates for employing a modified version of the net oil price increase measure, as initially proposed by Hamilton (1996, 2003). The goal is to identify periods in which the price of oil rises above its recent high, defined as the highest oil price of the previous three years. The intuition behind this measure is that these periods identify oil prices that are well above the normal range in which oil prices are expected to fluctuate. As a result, the existing level of wages reflects an environment of higher oil prices, so constraints on downward nominal wage movements will be binding.

[13] See Table A.3 in the Appendix for a concise list of occupations classified as blue-collar and white-collar.

Table 3.1: Summary statistics of main sample

	Blue-collar workers	*White-collar workers*
Hourly wage	\$12.75	\$15.54
Weekly hours worked	37.00	35.02
Labor force participation	0.79	0.83
Age	40	40
Education		
High School Diploma	0.68	0.35
College Degree	0.32	0.64
Wage change distribution		
$\Delta\omega < 0$	0.20	0.17
$\Delta\omega = 0$	0.23	0.20
$\Delta\omega > 0$	0.57	0.63
Observations	130,353	128,623
Share of female employees	0.35	0.71

Notes: Data are based on authors' calculations using the 1982-2018 Current Population Survey. The sample comprises employed adults aged 25-64 who have not changed occupations in the past year. *High School Diploma* refers to the fraction of workers with at most a high school degree, whereas *College Degree* denotes workers with at least one year of college education. The wage distribution shows the proportion of workers experiencing wage cuts ($\Delta\omega < 0$), no change in wages ($\Delta\omega = 0$), or wage increases ($\Delta\omega > 0$) over the sample period. The results are weighted using the weights provided by the CPS.

One concern with this measure of oil price shock is that it may capture the effects of increased world aggregate demand as opposed to lower oil supply, which may cause output, wages, and hours worked to move together. To mitigate this concern, we use the world aggregate demand shock of Kilian (2009) as a control variable in our regressions. That shock is a measure of the increase in the price of oil due to higher world aggregate demand. Although the published version of his paper used data through 2007, Kilian publishes updated shock data on his website, allowing us to include it in our regressions for all of the years in our sample.

3.3.3 Model setup

The simulations in the previous section demonstrated the asymmetric relationship between oil price shocks and nominal wage growth across different quantiles of the wage distribution. Large positive oil price shocks negatively impact wages at both the highest and lowest quantiles, while having no discernible effect at the middle quantiles. Consequently, we expect a similar impact of the oil shock on the distribution of nominal wage growth when the DNWR constraint is binding. To test this hypothesis, we estimate quantile regressions using the methodology of Lipsitz et al. (2016).

To identify the effects of the positive oil shock, we estimate the following wage growth model as a function of the oil shock and other relevant macroeconomic covariates:

$$\Delta w_{it} = \beta_0 + \beta_1 \pi_{t-1} + \beta_2 u_t + \beta_3 D_t + \beta_4 \eta_t^{oil} + \beta_5 (\eta_t^{oil} * \zeta_j) + \varepsilon_{it}. \tag{3.8}$$

where Δw_{it} denotes the year-over-year percentage change in nominal hourly wages for worker i in period t. π_{t-1} and u_t represent the one-period lag of inflation and unemployment rate, respectively. D_t represents the measure of changes in aggregate demand driving oil prices, while η_t^{oil} is the positive oil shock from the high oil regime in Bachmeier and Keen (2023). Our quantile regression model allows β_4 to vary depending on the specific quantile of the wage growth distribution. Additionally, the industry-level response to the oil shock is incorporated by interacting positive changes in oil prices with a vector of 34 industries, ζ_j.[14] The industries are based on the time-consistent Census Bureau industrial classification system and chosen based on their energy dependence.[15] Therefore, β_5 represents the differential effect of the industry-level response to the oil shock on nominal hourly wages. We estimate equation 3.8 across 19 quantiles of wage growth, ranging from 0.05 to 0.95 in increments of 0.05. The estimated coefficients of β_5 in equation 3.8 are presented as our baseline model. We then proceed to analyze the impact on weekly hours worked.

[14] A list of industries along with the number of observations for each is presented in Table A.4.

[15] We use the `ind1990` variable from the CPS. Industries for which oil price changes have no significant impact on production costs were excluded, as were those with insufficient data.

3.4 Empirical Results

3.4.1 Baseline results

Effects on the Distribution of Wages

We begin by reporting the estimated coefficients from equation 3.8 for blue-collar workers in selected industries. These estimates are presented in Figure 3.7, where each circle represents the point estimate of the oil shock, and the horizontal bars display the bootstrapped standard errors at a given quantile for the average worker in each industry. Across all industries, we find that the oil shock has a slight positive effect on nominal hourly wage growth (not exceeding 3 percentage points) at the lower end of the wage distribution, specifically from quantiles 0.05 to 0.45. For example, in the Mining industry, an exogenous oil price shock positively affects the percent change in nominal hourly wages for blue-collar workers up to the 0.25 quantile, beyond which the impact starts to have a decreasing effect at the 0.65 quantile. Although the positive effect of the oil shock at lower quantiles might initially appear counterintuitive, it is crucial to consider this effect in the context of nominal hourly wage growth at a given quantile. Notably, the average nominal wage growth for workers in the Mining industry at the 0.25 quantile is -4% over the sample period.[16] This suggests that the oil shock exacerbates the negative wage growth for these workers at lower quantiles. We interpret this pattern as indicating that the DNWR constraint is non-binding at quantiles below 0.30 for Mining workers.

A similar interpretation applies to the response of nominal wage growth at upper quantiles, where the oil shock tends to compress the magnitude of wage growth towards zero. Just as with lower quantiles, workers in Mining at quantiles above 0.60 encounter a non-binding DNWR constraint. For the middle quantiles, specifically between 0.30 and 0.50 in Mining, the oil shock appears to have no effect, consistent with our earlier analysis suggesting that this is where the DNWR constraint is binding. Overall, this general pattern is observed

[16]The average nominal wage for workers at the 0.25 quantile is \$12 per hour. A rough calculation reveals that the oil shock results in a decline of less than 0.05% in nominal hourly wage growth at this quantile.

across all industries (see Figure 3.7).[17] In other words, the industry-specific response of nominal wage growth to an exogenous oil shock is heterogeneous: the oil shock widens(narrows) the negative(positive) wage growth at lower(upper) quantiles, with no effect at the middle quantiles. However, these results are not statistically significant across any of the industries, making it challenging to draw definitive conclusions regarding the impact of oil price shocks on the distribution of nominal wage growth.

The null effects observed on nominal wage growth may be attributed to the low statistical power, stemming from the small number of observations in each industry-quantile cell.[18] Although we use the Frisch–Newton interior point method to estimate the fit of the quantile regressions, which is appropriate for our dataset, a more precise approach might have been the Frisch–Newton algorithm for sparse quantile regression (Koenker and Ng, 2005). This algorithm is computationally more robust for handling panel data where the model dimension is large due to interaction terms, yet the number of non-zero elements in the matrix structure is small-—such as in the case of our oil shock variable. As a robustness check, we plan to employ the Frisch–Newton algorithm for sparse quantile regression in a future version of this paper.

Effects on annual percent change in hours worked

To capture the impact on the quantity of labor supplied along the intensive margin, we re-estimate equation (3.8) with the annualized change in weekly hours worked ($\Delta hours_{it}$) as the outcome variable.[19] To avoid the limited number of observations within industry-quantile cells, we use 10 quantiles (q), spanning from 0.05 to 0.95 in increments of 0.10. We then

[17] This pattern is also evident in industries not depicted here, as well as in the analysis of white-collar workers.

[18] Our one-year panel data may be insufficient to capture adequate within-industry variation at each quantile.

[19] We use the ORG variable `uhrsworkorg`, which "reports the total number of hours the respondent, who is paid hourly, usually works per week at their main job," to calculate the annual percent change. Using the annual percent change in hours worked as a metric is common practice. This measure is standard for assessing labor productivity, as reported by the Bureau of Labor Statistics, alongside the annual percent change in output. However, for the sake of easier interpretation, we utilize the level form of this variable in equation 3.10.

Figure 3.7: Effect of exogenous oil price on nominal hourly wage growth on selected industries

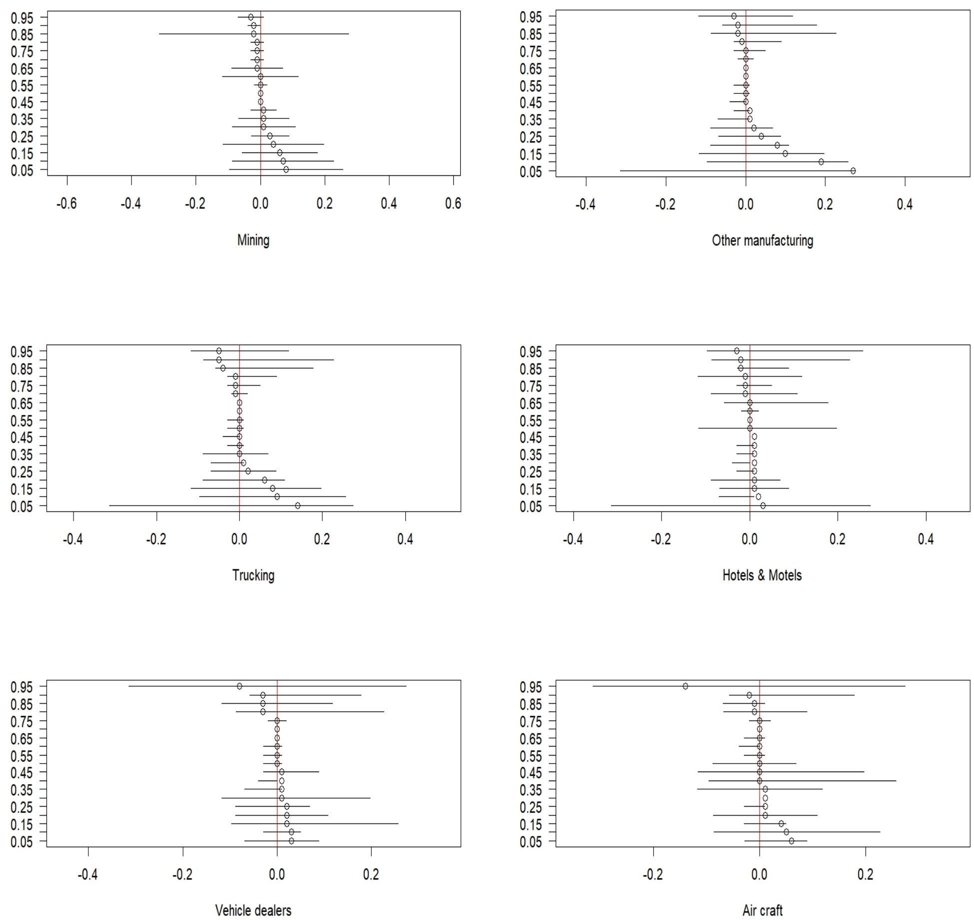

Notes: The figure displays the estimated coefficients of β_5 in equation (3.8) using quantile regressions for the following industries, listed in clockwise order: Mining, Manufacturing, Hotels and Motels, Aircraft, Vehicle Dealers, and Trucking. Each circle represents the point estimate of the oil shock's effect on nominal hourly wages at a specific quantile, while the horizontal bars indicate the bootstrapped standard errors at that quantile for the average worker in each industry.

include the nominal wage growth at each quantile (ω_q) interacted with industry response (ζ_j) and the oil shock, as a worker's industry has a substantial impact on their wage, which may differentially affect hours worked (Krueger and Summers, 1988). Specifically, we estimate the following quantile regression:

$$\Delta hours_{it} = \beta_0 + \beta_1 \pi_{t-1} + \beta_2 u_t + \beta_3 D_t + \beta_4 \eta_t^{oil} + \beta_5 (\eta_t^{oil} * \Lambda) + \varepsilon_{it}. \tag{3.9}$$

where $\Lambda = (\zeta_j * \omega_q)$ and all variables are as previously defined in equation 3.8. Here, β_5 captures the differential effect of the industry-level response to the oil shock at a given wage quantile on hours worked. In other words, it measures how variations in the industry-level response to the oil shock across different wage quantiles influence hours worked. The estimated coefficients of β_5 for selected industries are displayed in Figure 3.8.

Across all industries, we observe that average annual hours worked decreased at nearly all wage quantiles in response to a positive oil shock. Moreover, the magnitude of the effects increases at higher wage quantiles. For example, in the construction industry, the oil price shock reduces average annual hours worked by approximately 3 percentage points at wage quantiles below 0.35. However, at wage quantiles above 0.45, the decrease in average annual hours worked approaches 10 percentage points. This pattern is generally consistent across the selected industries, as illustrated in Figure 3.8. Although these findings align with our *a priori* expectations discussed in Section 3.2.4, the estimated effects are not statistically different from zero. As a robustness check, we plan to use four-digit NAICS codes, which classify industries with a higher level of aggregation. This higher aggregation would increase the sample size, which may provide more power to identify statistical effects at any given quantile.

In light of this, we use a fixed effects (FE) model, where the oil shock is invariant at all quantiles of our labor market outcomes. We augment equations 3.8 and 3.9 as follows:

$$Y_{i_o t} = \beta_0 + \beta_1 \pi_{t-1} + \beta_2 u_t + \beta_3 D_t + \beta_4 \eta_t^{oil} + \beta_5 X_{i_o t} + \lambda_t + \gamma_o + \varepsilon_{i_o t}. \tag{3.10}$$

where $Y_{i_o t}$ represents the outcomes of interest, specifically the annual percent change in nominal hourly wages and usual weekly hours worked for worker i in occupation type o in period t. o indicates whether a worker is classified in a blue-collar or white-collar occupation. $X_{i_o t}$ controls for worker characteristics such as age, sex, and education, while λ_t represents the time fixed effects that account for time-varying national shocks. Additionally, γ_o captures the occupation-specific features that might affect $Y_{i_o t}$. All other variables remain as defined in equation 3.8.

In contrast to our previous analysis, which examined the industry-level response to an exogenous oil shock on the outcomes of interest, this section shifts the focus to the impact across different occupation types. Specifically, we investigate whether the oil shock disproportionately affects nominal wage growth and weekly hours worked among workers in blue-collar versus white-collar occupations.

3.4.2 Nominal Wage Growth

Table 3.2, Panel A presents the result of estimating equation (3.10) with nominal wage growth as the response variable. Columns (1) and (2), show the estimates for workers in blue-collar and white-collar occupations, respectively. For blue-collar workers, we find that a positive oil price shock depresses nominal hourly wage growth by roughly 4 percentage points over the sample period. When scaled by the mean, it represents a 0.45% decline in nominal hourly wage growth. In the context of our measure of an exogenous oil shock—-defined as instances when oil prices exceed its recent highs-—this estimated effect should be considered an upper bound on the impact of oil price shocks on nominal hourly wage growth for blue-collar workers. A similar analysis can be made for workers in white-collar occupations where the magnitude of the oil shock decreases nominal wage growth by 0.44%. Similar to the estimates from our quantile regression model, the response of nominal hourly wage growth to an exogenous oil shock in our fixed effects (FE) model is statistically indistinguishable from zero. At this point, we cannot draw any further conclusions except that firms appear less likely to adjust nominal wages in response to significant increases in oil prices. Consequently, we next examine whether firms are adjusting their marginal costs via the intensive margin.

3.4.3 Hours Worked

Turning to hours worked, Panel B of Table 3.2, we use the level of usual hours worked per week for easier interpretation of the results.[20] The most striking pattern observed is that

[20]The estimates remain qualitatively similar when using the annual percent change in hours worked.

Figure 3.8: Effect of exogenous oil price on usual hours worked per week on selected industries

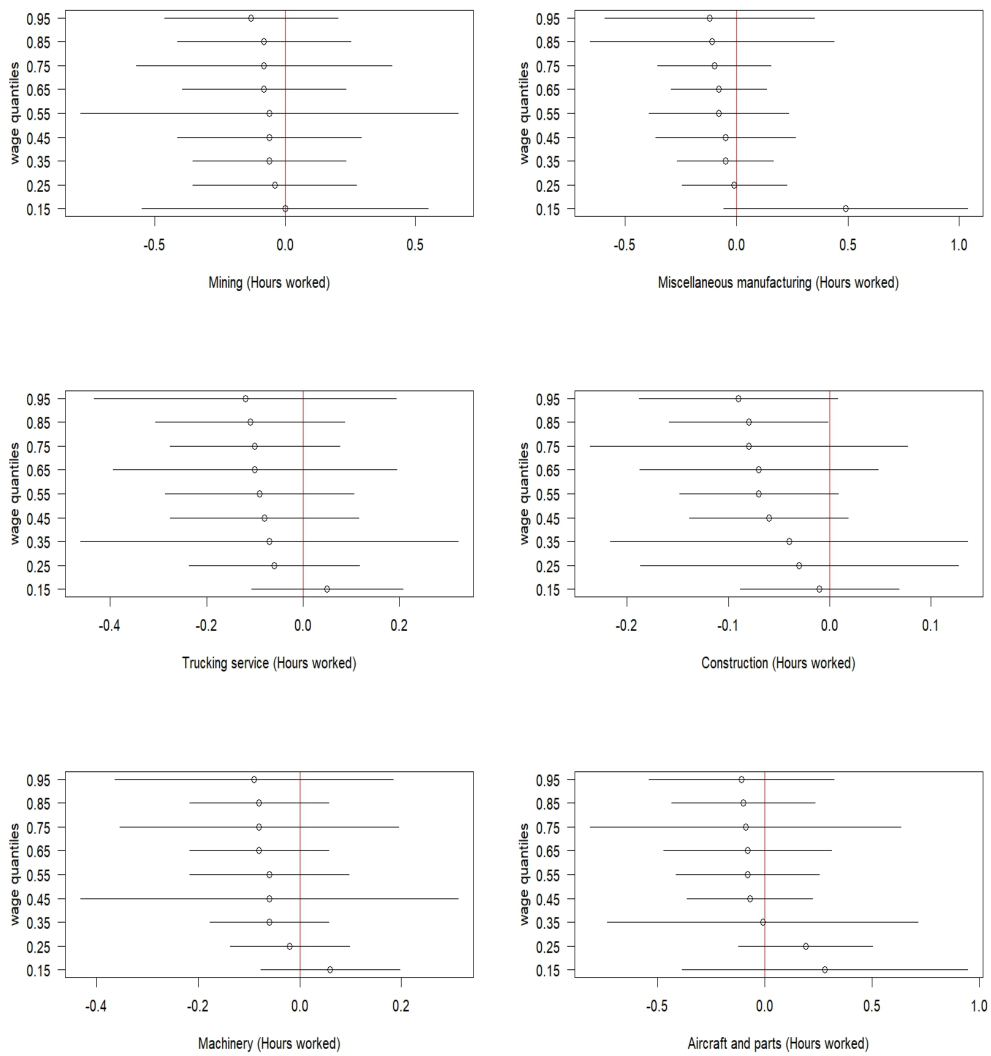

Notes: The figure displays the estimated coefficients of β_5 in equation (3.9) using quantile regressions for the following industries, listed in clockwise order: Mining, Manufacturing, Construction, Aircraft, Machinery, and Trucking. Each circle represents the point estimate of the oil shock's effect on annualized change in usual hours worked per week at a specific wage quantile, while the horizontal bars indicate the bootstrapped standard errors at that quantile for the average worker in each industry.

Table 3.2: Regression estimates of positive oil price change on labor supply indicators

	Blue-collar workers	*White-collar workers*
Panel A: Nominal wage growth		
Positive oil price change	-0.035	-0.023
	(0.024)	(0.020)
Intercept	0.109*	0.029
	(0.059)	(0.064)
Mean Y (%)	7.749	5.222
Percent Effect (%)	-0.446	-0.444
Time & Occupation FE	X	X
Macro & worker control variables	X	X
Observations	104,104	91,532
Panel B: Weekly hours worked		
Positive oil price change	-2.746	0.372
	(2.185)	(1.997)
Intercept	13.245	35.574***
	(11.101)	(10.794)
Mean Y (*hours per week*)	37	35
Percent Effect (%)	-7.427	1.075
Time & Occupation FE	X	X
Macro & worker control variables	X	X
Observations	70,774	61,897
Panel C: Overtime pay		
Positive oil price change	-0.072***	-0.069***
	(0.026)	(0.021)
Intercept	-0.420***	-0.245**
	(0.120)	(0.122)
Mean Y (*share of workers*)	0.226	0.155
Percent Effect (%)	-32.006	-44.685
Time & Occupation FE	X	X
Macro & worker control variables	X	X
Observations	78,783	72,982

Notes: Data are obtained from the 1982-2018 Current Population Survey. The sample comprises employed adults aged 25-64 who have not changed occupations in the past year. The dependent variable *Nominal wage growth* represents the annual percentage change in hourly wages, while *Weekly hours worked* refers to the ORG variable that reports the number of hours an hourly worker typically works per week at their primary job. The dependent variable in Panel C denotes the likelihood of receiving overtime pay. Occupation fixed effects are based on the time-consistent Census Bureau industrial classification system (i.e., the `occ1990` variable). Macro control variables include contemporaneous unemployment, lagged inflation, and the Kilian measure of changes in aggregate demand driving the oil price. Controls for worker characteristics include age and education. Standard errors, shown in parentheses, are clustered at the 3-digit industry level. The results are weighted using the weights provided by the CPS. *$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$ indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

the impact of the oil shock disproportionately affects blue-collar workers more negatively compared to their white-collar counterparts. Specifically, weekly hours worked by blue-collar workers decrease by 7%, while white-collar workers experience a 1% increase in their weekly hours. Although the effects are not statistically different from zero, they are practically important and worth discussing for two reasons: (1) the average weekly hours worked may mask the true response of hours worked to the oil shock, and (2) the heterogeneity in the response of hours worked across occupation types could signal a form of labor hoarding by firms to manage the rising production costs associated with the oil shock.

A possible concern regarding our specification in equation 3.10 is that there is little variation in our measure of hours worked. Although our analysis focuses on employees who have remained in the same job for at least one year, we did not explicitly account for the probability that workers might shift between full-time and part-time status due to various factors, including macroeconomic fluctuations driven by oil shocks. Consequently, using average weekly hours worked may obscure variations in employment flows—specifically, the differences in hours worked across blue-collar and white-collar occupations and over time.[21] Daly and Hobijn (2014) proposed a decomposition technique for aggregate wage growth that could be applicable in our context.[22] This method would enable us to decompose average weekly hours worked into components reflecting compositional changes due to transitions out of full-time employment (where weekly hours worked are < 35) and transitions into full-time employment (where weekly hours worked are ≥ 35). Implementing this technique would require transforming our measure of hours worked to capture changes in weekly hours worked (or using the growth rate) and potentially adjusting our unit of analysis to the local labor market level. In addition, we cannot completely rule out the possibility of self-reported measurement error in our outcomes of interest—hourly wage and weekly hours worked—despite using the Earner Study Supplement, where these variables are more precisely defined. Barattieri et al. (2014) document significant measurement error in household surveys and propose

[21] Figure A.1 illustrates the little variation in average weekly hours worked across blue-collar and white-collar occupations over time, with the trend line for white-collar workers remaining essentially flat.

[22] The authors' technique is analogous to standard shift-share designs, as outlined by Adao et al. (2019), which provides an overview of typical shift-share methodologies.

a method for correcting it. We consider these important areas for future exploration.

Overall, our findings regarding annual nominal wage growth and hours worked are economically significant. Specifically, our estimates indicate that nearly one in four hourly paid workers may experience a nominal wage cut of approximately 1% due to firms' adjustments to rising marginal costs resulting from the exogenous oil shock.[23] This suggests that exogenous oil price increases adversely workers, particularly blue-collar employees. The decline in nominal wage growth, combined with a 7% reduction in the usual weekly hours worked by blue-collar workers, necessarily leads to a decrease in labor income for these individuals. Furthermore, the results regarding hours worked suggest that businesses partly adjust to the oil shock by modifying the weekly hours of existing employees along the intensive margin.[24] This pattern may reflect a broader phenomenon of labor hoarding, where firms are hesitant to alter their existing workforce during periods of weaker growth or reduced demand.[25] This explanation aligns with the observed robustness of the U.S. labor market, even as growth moderated in recent years.

3.4.4 Overtime Pay

We also investigate the effect of the oil shock on the likelihood of receiving overtime pay as a potential channel through which firms adjust production costs. We utilize the ORG variable `otpay`, a binary indicator that denotes whether the respondent receives overtime pay, tips, or commissions. We discuss our findings and their implications in the subsequent section.

Table 3.2 Panel C presents the estimates from (3.9) regarding overtime pay. These results tell a clear story of a reduction in the likelihood of employees receiving overtime pay in the aftermath of a large oil price increase. Specifically, the estimates are statistically significant and suggest that, in response to higher marginal costs imposed by the oil shock, firms reduce the probability of paying overtime compensation to blue-collar workers by 7

[23] Table A.1 indicates that, on average, 23% of workers experience wage cuts over the sample period.

[24] Estimates from equation 3.10 for our full sample indicate that the oil shock reduces usual weekly hours worked by 3%.

[25] Biddle (2014) explains that labor hoarding is a profit-maximizing strategy firms employ to manage the costs associated with hiring, firing, and training workers.

percentage points, which represents a 32% decrease from the mean. For white-collar workers, the estimates indicate a 45% reduction from the mean. To further contextualize these results, we also examine the likelihood of working overtime hours, defined as a full-time employee working more than 40 hours per week. The idea is that overtime pay impacts the marginal cost of an employee only when they transition from having no overtime hours to accruing positive overtime hours (Grigsby et al., 2021). The findings, presented in Table 3.3, indicate that the rate of overtime hours worked decreases by 20 percentage points for blue-collar workers and by 15 percentage points for white-collar workers.

Taken together, these results are economically meaningful. Firms, faced with higher marginal costs due to an exogenous oil shock and constrained by downwardly rigid nominal wages, may be willing to sacrifice laying off workers for reducing overtime hours. The cost savings from this trade-off can be significant, particularly in light of federal regulations on overtime work. Under the Fair Labor Standards Act (FLSA), non-exempt employees are entitled to an overtime premium of 1.5 times their regular pay rates for hours worked over 40 in a workweek.

It is important to note that our data is limited to specific measures of compensation and does not include a separate variable for reported overtime pay. Although this amount could be imputed based on the Fair Labor Standards Act (FLSA), doing so would introduce additional noise into the data. Nonetheless, we believe these results are promising and plan to extend this research in several ways. Given that overtime hours are highly responsive to macroeconomic fluctuations and may vary considerably across industries, we aim to assess firms' adjustments to an exogenous oil shock at the industry level.

Table 3.3: Regression estimates of positive oil price shock on the rate of working overtime hours

	Blue-collar workers	*White-collar workers*
	(1)	(2)
Overtime hours		
Positive oil price change	-0.202***	-0.154***
	(0.061)	(0.052)
Intercept	-1.061***	-0.770***
	(0.295)	(0.273)
Mean Y (*share of workers*)	0.073	0.044
Time & Occupation FE	X	X
Macro & worker control variables	X	X
Observations	70,774	61,897

Notes: Data are obtained from the 1982-2018 Current Population Survey. The sample comprises employed adults aged 25-64 who have not changed occupations in the past year. The dependent variable, *Overtime Hours*, is a binary indicator variable that signifies whether a full-time employee's weekly hours worked exceed 40 hours. Occupation fixed effects are based on the time-consistent Census Bureau industrial classification system (i.e., the `occ1990` variable). Macro control variables include contemporaneous unemployment, lagged inflation, and the Kilian measure of changes in aggregate demand driving the oil price. Controls for worker characteristics include age and education. Standard errors, shown in parentheses, are clustered at the 3-digit industry level. The results are weighted using the weights provided by the CPS. *$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$ indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

3.5 Conclusion

The literature on DNWR in macro models is extensive, yet remains relatively underdeveloped when it comes to microdata. This paper enhances our understanding of the relationship between oil price shocks and DNWR on firms' demand for labor services by utilizing microdata from the CPS spanning the past three decades. Through a series of simulations, we demonstrate that firms reduce their demand for labor services in response to increased production costs from an exogenous oil price increase. This reduction in labor demand necessarily implies lower employment, measured as weekly hours worked in our paper. Our results suggest heterogeneity in how firms adjust weekly hours across different occupations: labor hours for employees in blue-collar occupations decrease by 7%, whereas those for white-collar counterparts increase by 1%. However, we could not reject the null hypothesis that these estimates are statistically indistinguishable from zero. This result may reflect limitations in our dataset rather than the lack of an effect.

In addition to nominal wage growth, we examined overtime pay as a possible channel of adjustment. Our results reveal that overtime compensation is a key margin for firms to adjust to large oil price increases. Specifically, the likelihood of receiving overtime pay declines by nearly 45% across both blue-collar and white-collar workers due to a reduction in overtime hours. Future versions of this paper will explore additional margins of adjustment, including changes in employment, to provide a unified perspective on how firms respond to exogenous oil price increases. Nevertheless, our findings suggest that explanations of the macroeconomic response to oil price shocks could benefit from emphasizing the pivotal role of overtime pay as a key adjustment margin for firms.

Bibliography

Abadie, A. (2005). Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19.

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

Abo-Zaid, S. (2013). Optimal monetary policy and downward nominal wage rigidity in frictional labor markets. *Journal of Economic Dynamics and Control*, 37(1):345–364.

Adao, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.

Akerlof, G. A., Dickens, W. T., Perry, G. L., Gordon, R. J., and Mankiw, N. G. (1996). The macroeconomics of low inflation. *Brookings papers on economic activity*, 1996(1):1–76.

Almond, D., Hoynes, H., and Schanzenbach, D. (2011). Inside the war on poverty: The impact of food stamps on birth outcomes. *The review of economics and statistics*, 93(2):387–403.

Alsan, M. and Yang, C. (2018). Fear and the safety net: Evidence from secure communities. Working Paper 24731, National Bureau of Economic Research.

Altonji, J. G. and Devereux, P. J. (2000). The extent and consequences of downward nominal wage rigidity. *Research in Labor Economics*, 19:383–431.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–97.

Aydemir, A. and Borjas, G. J. (2011). Attenuation Bias in Measuring the Wage Impact of Immigration. *Journal of Labor Economics*, 29(1):69–113.

Bachmeier, L. J. and Keen, B. D. (2023). Modeling the asymmetric effects of an oil price shock. *International Journal of Central Banking*, 19(3):1–47.

Baker, A. C., Larcker, D. F., and Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2):370–395.

Barattieri, A., Basu, S., and Gottschalk, P. (2014). Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics*, 6(1):70–101.

Bartolucci, C. (2014). Understanding the native-immigrant wage gap using matched matched employer-employee data: Evidence from germany. *ILR Review*, 67(4):1166–1202.

Benigno, P. and Antonio Ricci, L. (2011). The inflation-output trade-off with downward wage rigidities. *American Economic Review*, 101(4):1436–1466.

Bentolila, S., Michelacci, C., and Suarez, J. (2010). Social contacts and occupational choice. *Economica*, 77(305):20–45.

Beraja, M., Hurst, E., and Ospina, J. (2019). The aggregate implications of regional business cycles. *Econometrica*, 87(6):1789–1833.

Bitler, M. and Hoynes, H. W. (2011). Immigrants, welfare reform, and the U.S. safety net. Working Paper 17667, National Bureau of Economic Research.

Bizopoulou, A. (2019). Job Tasks and Gender Wage Gaps within Occupations. Working Papers 124, VATT Institute for Economic Research.

Blanchard, O. and Galí, J. (2009). The macroeconomic effects of oil price shocks. *International Dimensions of Monetary Policy*, 373.

Blau, F. D., Brummund, P., and Liu, A. Y.-H. (2013). Trends in occupational segregation by gender 1970-2009: Adjusting for the impact of changes in the occupational coding system. *Demography*, 50(2):471–494.

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865.

Bleakley, H. and Chin, A. (2004). Language skills and earnings: Evidence from childhood immigrants. *The Review of Economics and Statistics*, 86(2):481–496.

Bollinger, C. R. and Hirsch, B. T. (2006). Match bias from earnings imputation in the current population survey: The case of imperfect matching. *Journal of Labor Economics*, 24(3):483–519.

Borjas, G. (1994). The economics of immigration. *Journal of Economic Literature*, XXXII:1667–1717.

Borjas, G. (1995). Assimilation and changes in cohort quality revisited: What happened to immigrant earnings in the 1980s? *Journal of Labor Economics*, 13(2):201–45.

Borjas, G. (2003). The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *The Quarterly Journal of Economics*, 118(4):1335–1374.

Borjas, G. (2015). The slowdown in the economic assimilation of immigrants: Aging and cohort effects revisited again. *Journal of Human Capital*, 9(4):483–517.

Borjas, G., Friedman, R. B., and Katz, L. (1997). How much do immigration and trade affect labor market outcomes? *Brookings Papers on Economic Activity*, 28(1):1–90.

Borjas, G. J. (1985). Assimilation, changes in cohort quality, and the earnings of immigrants. *Journal of Labor Economics*, 3(4):463–489.

Borjas, G. J. (2004). Food insecurity and public assistance. *Journal of Public Economics*, 88(7):1421–1443.

Borjas, G. J. and Cassidy, H. (2019). The wage penalty to undocumented immigration. *Labour Economics*, 61:101757.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, page rdae007.

Bratsberg, B. and Terrell, D. (2002). School quality and returns to education of us immigrants. *Economic inquiry*, 40(2):177–198.

Brown, R. S., Moon, M., and Zoloth, B. S. (1980). Incorporating occupational attainment in studies of male-female earnings differentials. *Journal of Human Resources*, 15(1):3–28.

Bureau of Labor Statistics (2023). Foreign-born workers: Labor force characteristics- 2023.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398.

Camarota, S. A. (1997). The effect of immigrants on the earnings of low-skilled native workers: Evidence from the june 1991 current population survey. *Social Science Quarterly*, 78(2):417–431.

Card, D. (2001). Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64.

Card, D. and Hyslop, D. (1997). Does inflation" grease the wheels of the labor market"? In *Reducing inflation: Motivation and strategy*, pages 71–122. University of Chicago Press.

Cassidy, H. (2019). The decline in earnings of childhood immigrants in the united states. *International Migration Review*.

Cassidy, H. (2020). The decline in earnings of childhood immigrants in the united states. *The International Migration Review*, 54(1):58–82.

Cassidy, H. and Dacass, T. (2019). Occupational licensing and immigrants. Working paper 2019-9, Center for Growth and Opportunity.

Center on Budget and Policy Priorities (2018). Chart book: Snap helps struggling families put food on the table. Technical report, Center on Budget and Policy Priorities.

Chan, M. K. (2013). A dynamic model of welfare reform. *Econometrica*, 81(3):941–1001.

Chiswick, B. and Miller, P. (1995). The endogeneity between language and earnings: International analyses. *Journal of Labor Economics*, 13(2):246–288.

Chiswick, B. and Miller, P. (2009a). Earnings and occupational attainment among immigrants. *Industrial Relations: A Journal of Economy and Society*, 48(3):454–465.

Chiswick, B. R. (1978). The effect of americanization on the earnings of foreign-born men. *Journal of Political Economy*, 86(5):897–921.

Chiswick, B. R. and Miller, P. W. (2009b). Educational mismatch: are high-skilled immigrants really working at high-skilled jobs and the price they pay if they aren't?

Chiswick, B. R. and Miller, P. W. (2009c). The international transferability of immigrants' human capital. *Economics of Education Review*, 28(2):162–169.

Chiswick, B. R. and Miller, P. W. (2012). Negative and positive assimilation, skill transferability, and linguistic distance. *Journal of Human Capital*, 6(1):35–55.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy*, 113(1):1–45.

Cunningham, S. (2021). In *Causal Inference: The Mixtape*. Yale University Press.

Daly, M. C. and Hobijn, B. (2014). Downward nominal wage rigidities bend the phillips curve. *Journal of Money, Credit, and Banking*, 46(S2):51–93.

Daly, M. C. and Hobijn, B. (2022). The importance of the part-time and participation margins for real wage adjustment. *Journal of Money, Credit and Banking*, 54(S1):89–111.

Damm, A. P. (2009). Ethnic enclaves and immigrant labor market outcomes: Quasiexperimental evidence. *Journal of Labor Economics*, 27(2):281–314.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.

Dechief, D. and Oreopoulos, P. (2012). Why do some employers prefer to interview matthew, but not samir? new evidence from toronto, montreal, and vancouver. *SSRN Electronic Journal*.

Demoussis, M., Giannakopoulos, N., and Zografakis, S. (2010). Native–immigrant wage differentials and occupational segregation in the greek labor market. *Applied Economics*, 42(8):1015–1027.

Dolton, P. J. and Kidd, M. P. (1994). Occupational access and wage discrimination. *Oxford Bulletin of Economics and Statistics*, 56(4):457–474.

Dustmann, C., Schönberg, U., and Stuhler, J. (2016). The impact of immigration: Why do studies reach such different results? *Journal of Economic Perspectives*, 30(4):31–56.

East, C. (2018a). The effect of food stamps on children's health: Evidence from immigrants changing eligibility. *Labour Economics*, 51:202–226.

East, C. (2018b). Immigrants' labor supply response to food stamp access. *Journal of Human Resources*.

Eissa, N. and Hoynes, H. W. (2006). Behavioral responses to taxes: Lessons from the eitc and labor supply. *Tax policy and the economy*, 20:73–110.

Elsby, M. W. (2009). Evaluating the economic significance of downward nominal wage rigidity. *Journal of Monetary Economics*, 56(2):154–169.

Erceg, C. J., Henderson, D. W., and Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of monetary Economics*, 46(2):281–313.

Fagan, G. and Messina, J. (2009). Downward wage rigidity and optimal steady-state inflation. Technical report, ECB Working paper.

Fallick, B., Lettau, M. K., and Wascher, W. L. (2016). Downward nominal wage rigidity in the united states during and after the great recession. Technical report, Federal Reserve Bank of Cleveland.

Fayaz Farkhad, B. and Meyerhoefer, C. (2018). The impact of participation in snap on labor force decisions. Technical report, Agricultural and Applied Economics Association.

Ferrer, A., Green, D. A., and Riddell, W. C. (2006). The effect of literacy on immigrant earnings. *The Journal of Human Resources*, 41(2):380–410.

Fix, M. and Passel, J. (2002). Scope and impact of welfare reform's immigrant provisions.

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., Backman, D., Chen, A., Cooper, G., Richards, S., Schouweiler, M., and Westberry, M. (2023). IPUMS CPS: Version 11.0 [dataset]. Minneapolis, MN: IPUMS.

Flood, S. and Pacas, J. (2016a). Using the annual social and economic supplement with current population survey panels. Working Paper 2016-4, Minnesota Population Center, University of Minnesota.

Flood, S. and Pacas, J. (2016b). Using the annual social and economic supplement with current population survey panels. Technical report, Minnesota Population Center WP.

Fortin, N. M. and Huberman, M. (2002). Occupational gender segregation and women's wages in canada: An historical perspective. *Canadian Public Policy / Analyse de Politiques*, 28:S11–S39.

Friedberg, R. M. (1992). The labor market assimilation of immigrants in the united states: the role of age at arrival. *Brown University*.

Gabriel, P. E. and Schmitz, S. (1987). The relative earnings of native and immigrant males in the united states. *Quarterly Review of Economics and Business*, 27:91–101.

Gardner, J. (2022). Two-stage differences in differences. *arXiv preprint arXiv:2207.05943*.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Gottschalk, P. (2005). Downward nominal-wage flexibility: real or measurement error? *Review of Economics and Statistics*, 87(3):556–568.

Grigsby, J., Hurst, E., and Yildirmaz, A. (2021). Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review*, 111(2):428–471.

Hagstrom, P. A. (1996). The food stamp participation and labor supply of married couples: An empirical analysis of joint decisions. *Journal of Human Resources*, pages 383–403.

Haider, S., Schoeni, R., Bao, Y., and Danielson, C. (2004). Immigrants, welfare reform, and the economy. *Journal of Policy Analysis and Management*, 23(4):745–764.

Hamilton, J. D. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of monetary economics*, 38(2):215–220.

Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2):363–398.

Hamilton, J. D. (2011). Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic dynamics*, 15(S3):364–378.

Hauser, R. M. and Warren, J. R. (1997). Socioeconomic indexes for occupations: A review, update, and critique. *Sociological methodology*, 27(1):177–298.

Hersch, J. and Shinall, J. B. (2018). Imputation match bias in immigrant wage convergence. *Demography*, 55:1475–1485.

Hirsch, B. T. and Schumacher, E. J. (2004). Match bias in wage gap estimates due to earnings imputation. *Journal of Labor Economics*, 22(3):689–722.

Hoynes, H. and Feldstein, S. (2015). The future of SNAP: Improving nutrition policy to ensure health and food equity. *Berkeley Food Institute Policy Brief*.

Hoynes, H. and Schanzenbach, D. W. (2012). Work incentives and the food stamp program. *Journal of Public Economics*, 96(1-2):151–162.

Hoynes, H., Schanzenbach, D. W., and Almond, D. (2016). Long-run impacts of childhood access to the safety net. *American Economic Review*, 106(4):903–34.

Imai, S., Stacey, D., and Warman, C. (2019). From engineer to taxi driver? language proficiency and the occupational skills of immigrants. *Canadian Journal of Economics*, 52(3):914–953.

Jardim, E. S., Solon, G., and Vigdor, J. L. (2019). How prevalent is downward rigidity in nominal wages? evidence from payroll records in washington state. Technical report, National Bureau of Economic Research.

Jo, Y. (2024). Beyond the spike at zero: Understanding nominal wage rigidity through empirical and model-based approaches. Available at SSRN 4747444.

Kaestner, R. and Kaushal, N. (2005). Immigrant and native responses to welfare reform. *Journal of Population Economics*, 18(1):69–92.

Kahn, S. (1997). Evidence of nominal wage stickiness from microdata. *American Economic Review*, 87(5):993 – 1008.

Keane, M. and Moffitt, R. (1998). A structural model of multiple welfare program participation and labor supply. *International Economic Review*, 39(3):553–589.

Kerr, W. R. and Lincoln, W. F. (2010). The supply side of innovation: H-1b visa reforms and us ethnic invention. *Journal of Labor Economics*, 28(3):473–508.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American economic review*, 99(3):1053–1069.

Koenker, R. and Ng, P. (2005). A frisch-newton algorithm for sparse quantile regression. *Acta Mathematicae Applicatae Sinica*, 21:225–236.

Krueger, A. B. and Summers, L. H. (1988). Efficiency wages and the inter-industry wage structure. *Econometrica*, 56(2):259–293.

Kurmann, A. and McEntarfer, E. (2019). Downward nominal wage rigidity in the united states: new evidence from worker-firm linked data. Technical report, Drexel University School of Economics Working Paper Series WP.

Kuroda, S. and Yamamoto, I. (2007). Why are nominal wages downwardly rigid, but less so in japan? an explanation based on behavioral economics and labor market/macroeconomic differences. *Monetary and Economic Studies*, 25(2):45–88.

LaLonde, R. J. and Topel, R. H. (1992). *The Assimilation of Immigrants in the U. S. Labor Market*, pages 67–92. University of Chicago Press.

Lazear, E. P. (1999). Culture and language. *Journal of Political Economy*, 107(S6):S95–S126.

Lebow, D. E., Saks, R. E., and Wilson, B. A. (2003). Downward nominal wage rigidity: Evidence from the employment cost index. *Advances in Macroeconomics*, 3(1).

Lebow, D. E., Stockton, D. J., and Wascher, W. (1995). Inflation, nominal wage rigidity, and the efficiency of labor markets. Available at SSRN 3032.

Lipsitz, M., Belloni, A., Chernozhukov, V., and Fernández-Val, I. (2016). quantreg. nonpar: An r package for performing nonparametric series quantile regression. *The R Journal*, 8(2):370–381.

Liu, L., Wang, Y., and Xu, Y. (2024). A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. *American Journal of Political Science*, 68(1):160–176.

Liu, P., Zhang, J., and Chong, S. (2004). Occupational segregation and wage differentials between natives and immigrants: evidence from Hong Kong. *Journal of Development Economics*, 73(1):395–413.

Llull, J. (2017). Immigration, Wages, and Education: A Labour Market Equilibrium Structural Model. *The Review of Economic Studies*, 85(3):1852–1896.

Lofstrom, M. and Bean, F. D. (2002). Assessing immigrant policy options: Labor market conditions and postreform declines in immigrants' receipt of welfare. *Demography*, 39(4):617–637.

Loweree, J, R.-M. A. and Walter, E. (2020). The impact of covid-19 on noncitizens and across the u.s. immigration system. Special report, American Immigration Council.

Lubotsky, D. (2007). Chutes or ladders? a longitudinal analysis of immigrant earnings. *Journal of Political Economy*, 115(5):820–867.

McLaughlin, K. J. (1994). Rigid wages? *Journal of Monetary Economics*, 34(3):383–414.

Mincer, J. (1974). *Schooling, Experience, and Earnings*. National Bureau of Economic Research, Inc.

Moffitt, R. (1983). An economic model of welfare stigma. *The American Economic Review*, 73(5):1023–1035.

Moffitt, R. and Fraker, T. (1988). The effect of food stamps on labor supply: A bivariate selection model. *Journal of Public Economics*, 35(1):25–56.

Oreopoulos, P. (2011). Why do skilled immigrants struggle in the labor market? a field experiment with thirteen thousand resumes. *American Economic Journal: Economic Policy*, 3(4):148–71.

Orrenius, P. M. and Zavodny, M. (2007). Does immigration affect wages? a look at occupation-level evidence. *Labour Economics*, 14(5):757 – 773.

Parti, T. and Hackman, M. (2024). Why immigration is now the no. 1 issue for voters. Article, The Wall Street Journal.

Patel, K. and Vella, F. (2013). Immigrant networks and their implications for occupational choice and wages. *The Review of Economics and Statistics*, 95(4):1249–1277.

Peri, G. and Sparber, C. (2009). Task specialization, immigration, and wages. *American Economic Journal: Applied Economics*, 1(3):135–69.

Pew Research Center (2024). Immigration attitudes and the 2024 election.

Rivera Drew, J., Flood, S., and Warren, J. R. (2014). Making full use of the longitudinal design of the current population survey: Methods for linking records across 16 months1. *Journal of Economic and Social Measurement*, (3):121–144.

Roth, J., Sant'Anna, P. H., Bilinski, A., and Poe, J. (2022). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *arXiv preprint arXiv:2201.01194*.

Rubin, D. B. (2005). Causal inference using potential outcomes. *Journal of the American Statistical Association*, 100(469):322–331.

Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J., and Sobek, M. (2020). Ipums usa: Version 10.0 [dataset].

Sant'Anna, P. H. and Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122.

Schmidley, D. and Robinson, J. (1998). How well does the current population survey measure the foreign born population in the united states? Technical report, Population Division Working Paper.

Schoeni, R. F. (1998). Labor market outcomes of immigrant women in the united states: 1970 to 1990. *International Migration Review*, 32(1):57–77.

Sharpe, J. and Bollinger, C. R. (2020). Who competes with whom? using occupation characteristics to estimate the impact of immigration on native wages. *Labour Economics*, 66:101902.

Smets, F. and Wouters, R. (2005). Comparing shocks and frictions in us and euro area business cycles: a bayesian dsge approach. *Journal of Applied Econometrics*, 20(2):161–183.

Snow, J. (1856). On the mode of communication of cholera. *Edinburgh medical journal*, 1(7):668.

Solon, G., Barsky, R., and Parker, J. A. (1994). Measuring the cyclicality of real wages: how important is composition bias? *The quarterly journal of economics*, 109(1):1–25.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Taylor, J. B. (2016). The staying power of staggered wage and price setting models in macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 2009–2042. Elsevier.

United States Department of Agriculture (2016). Supplemental nutrition assistance program (SNAP).

Xie, Y., K.-A. . N. C. (2016). Between- and within-occupation inequality: The case of high-status professions. *The Annals of the American Academy of Political and Social Science*, 663(1):53–79.

Appendix A

Chapter 3

Figure A.1: Average weekly hours worked across workers in blue-collar and white-collar occupations, 1989-2018

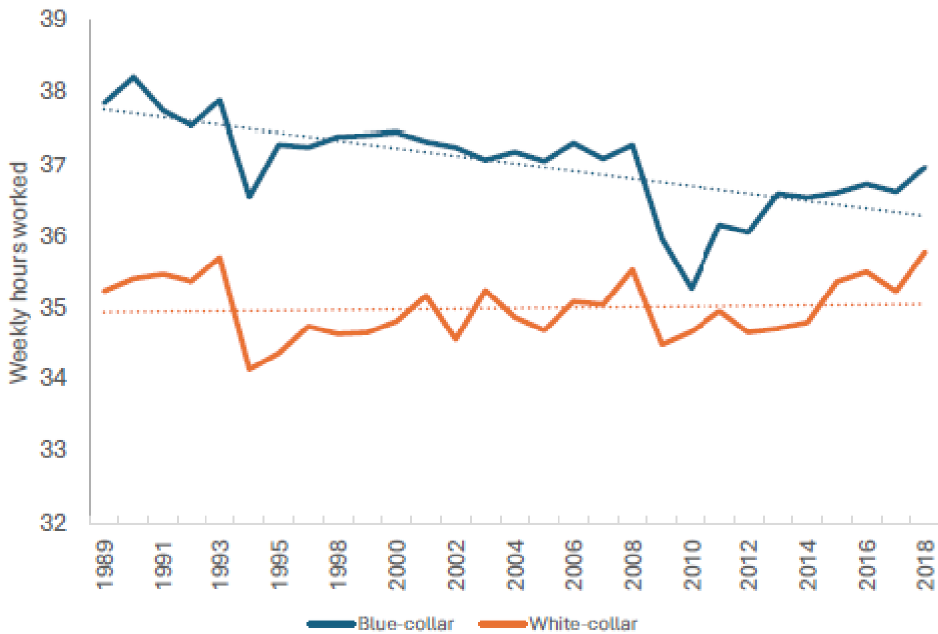

Notes: Data are based on authors' calculations from the 1982-2018 Current Population Survey. The sample includes employed adults aged 25-64 who have not changed occupations over the past year. The plot illustrates weekly hours worked by employees in both blue-collar and white-collar occupations. The variable is derived from the Earner Study Supplement of the CPS, specifically `uhrsworkorg`, which reports the typical number of hours an hourly worker works per week at their primary job. Over the sample period, the average weekly hours worked are 37 for blue-collar workers and 35 for white-collar workers.

Table A.1: Annual summary of inflation and nominal wage changes, 1983-2018

Year	N	PCE	CPI	$\%(\Delta\omega = 0)$	$\%(\Delta\omega < 0)$	$\%(\Delta\omega > 0)$
1983	19,525	10.16	3.16	39.40	16.14	44.46
1984	19,165	9.24	4.37	42.37	16.07	41.56
1985	19,533	8.85	3.53	40.25	16.61	43.14
1987	18,817	6.58	3.58	39.26	16.61	44.13
1988	19,451	8.25	4.10	38.86	16.29	44.85
1989	17,818	7.41	4.79	36.51	15.73	47.76
1990	19,081	6.49	5.42	36.26	16.32	47.42
1991	19,953	3.53	4.22	34.65	17.92	47.44
1992	19,761	6.44	3.04	36.94	18.15	44.91
1993	19,440	6.06	2.97	37.65	18.81	43.54
1994	18,636	6.04	2.60	36.15	20.33	43.53
1997	17,248	5.57	2.34	30.89	20.85	48.26
1998	17,653	6.15	1.55	30.03	20.43	49.54
1999	18,383	6.92	2.19	27.28	22.31	50.40
2000	18,843	7.69	3.37	26.97	22.54	50.49
2001	19,143	4.53	2.82	24.61	23.16	52.22
2002	17,610	3.89	1.60	25.65	24.31	50.04
2003	20,223	5.33	2.30	25.94	24.67	49.39
2004	21,492	6.35	2.67	28.16	24.69	47.15
2005	18,940	6.52	3.37	24.90	25.95	49.15
2006	20,177	5.80	3.22	24.01	24.72	51.27
2007	19,916	5.06	2.87	23.37	25.11	51.52
2008	20,656	3.11	3.81	24.51	24.66	50.83
2009	19,986	-1.58	-0.32	25.07	26.14	48.79
2010	19,970	3.73	1.64	28.74	27.13	44.13
2011	18,177	4.27	3.14	27.27	27.87	44.85
2012	18,175	3.26	2.07	25.69	28.33	45.98
2013	17,944	3.09	1.47	25.92	28.49	45.60
2014	17,917	4.27	1.62	24.88	28.06	47.05
2015	14,838	3.56	0.12	22.43	28.82	48.76
2016	17,226	3.49	1.27	21.47	28.50	50.03
2017	17,046	4.43	2.13	21.70	27.40	50.89
2018	16,639	4.84	2.44	20.91	28.18	50.91
Average	18,744	5.46	2.70	29.73	22.74	47.54
Median	18,892	5.63	2.81	27.28	23.73	47.60

Notes: Data are based on calculations by the authors using the Current Population Survey from 1983 to 2018. The table displays the unweighted number of hourly-paid workers, N, and the proportion of those experiencing no nominal wage change, $(\Delta\omega = 0)$, wage increases $(\Delta\omega > 0)$, and wage decreases $\%(\Delta\omega < 0)$ per year. Data for 1995 and 1996 are not shown because household identifiers were scrambled in 1995 to protect respondent confidentiality due to changes in geographic variables (Flood and Pacas, 2016b). This scrambling significantly impacted the number of observations for that year and the following one.

Table A.2: CPS Survey Design

Months in sample (MISH)	1	2	3	4	{8-month break}	5	6	7	8
CPS Basic Monthly	✓	✓	✓	✓	....	✓	✓	✓	✓
Earner Study/ORG Supplement				✓	....				✓
Linked sample				✓	....				✓

Notes: The CPS interviews household members for four consecutive months, then excludes them for eight months, and finally interviews them again for another four consecutive months. During the first interview, a household is assigned a value of 1 for the MISH variable. Households returning to the sample after an 8-month break are assigned a value of 5, and those that have completed their final interview receive a value of 8. Individuals with MISH codes of 4 or 8 are referred to as being in "outgoing rotation groups," as they will not be surveyed in the subsequent month. Our panel data links the same households interviewed in month 4 to their responses in month 8.

Table A.3: Examples of occupations classified as blue-collar or white-collar

Blue-collar	*White-collar*
Miscellaneous food prep workers	Legislators
Textile-cutting machine operators	Secondary school teachers
Recreation facility attendants	Librarians
Nursery farming workers	Retail sales clerks
Construction laborers	Managers of food-serving and lodging establishments
Farm managers, except for horticultural farms	Education instructors
Plumbers, pipefitters, and steamfitters	Insurance sales occupations
Tailors	Cashiers
Garage and service station-related occupations	Hotel clerks
Private household cleaners and servants	Kindergarten and earlier school teachers
Child care workers	Recreation workers
Janitors	Postmasters and mail superintendents
Farm workers	Sales demonstrators, promoters, models
Garbage and recyclable material collectors	Mail carriers for postal service
Barbers	Dietitians and nutritionists
Baggage porters	Editors and reporters
Kitchen workers	File clerks
Heavy equipment and farm equipment mechanics	Payroll and timekeeping clerks
Water and sewage treatment plant operators	Mail clerks, outside of post office
Supervisors of construction work	Data entry keyers
Power plant operators	Records clerks
Assemblers of electrical equipment	Announcers
Aircraft mechanics	Clergy and religious workers
Inspectors of agricultural products	Social workers
Timber, logging, and forestry workers	Speech therapists
Electricians	Door-to-door sales, street sales, and news vendors
Painters, construction, and maintenance	Statistical clerks
Carpenters	Photographers
Bus, truck, and stationary engine mechanics	Writers and authors
Repairers of mechanical controls and valves	Dental hygienists
Plant and system operators, stationary engineers	Registered nurses

Notes: Data are obtained from the Current Population Survey for the years 1982-2018. Occupations are categorized as 'blue-collar" and "white-collar" based on the classification used in Autor and Dorn (2013) and the `occ1990` variable, which provides a consistent long-term classification of occupations.

Table A.4: List of 34 industries and their corresponding observations

Industries	N
Mining	933
Oil and gas extraction	580
Construction	10632
Paper	1296
Plastic	797
Drugs and chemicals	377
Paints	79
Refining	212
Petro products	151
Rubber	442
Iron and steel	837
Engines	117
Machinery	4161
Vehicles	2631
Aircraft	642
Cycles	94
Other manufacturing	1016
Railroads	593
Urban transit	1064
Taxicab	62
Trucking	2463
Water transportation	108
Air transportation	651
Pipelines	49
Incidental to transportation	376
Electric light	1101
Gas supply systems	232
Electric and gas	125
Vehicle dealers	870
Gas stations	327
Fuel dealers	191
Auto rental	95
Auto repair	1094
Hotels and motels	2187

Notes: Data are obtained from the Current Population Survey for the years 1982-2018. The industries are classified according to the time-consistent Census Bureau industrial classification system, specifically the `ind1990` variable, and are selected based on their energy dependence. Industries for which oil price changes have no significant impact on production costs were excluded, as were those with insufficient data.