Title
Essays in Labor Economics

Author
Khanna, Shantanu

Publication Date
2022

UNIVERSITY OF CALIFORNIA,
IRVINE

Essays in Labor Economics

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Shantanu Khanna

Dissertation Committee:
Distinguished Professor David Neumark, Chair
Professor Matthew Freedman
Professor John Duffy
Professor Michael McBride

2022

DEDICATION

To my mother, Savita Khanna.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank David Neumark (Chair) for the incredible guidance and support over the last few years. Working with David has been a singular privilege during my academic journey. I am also grateful to Matthew Freedman for his thoughtful mentorship and advice. I want to thank John Duffy and Michael McBride for their generous guidance and comments at critical moments during the Ph.D. Every interaction with my committee members made me a better economist and academic. I would also like to thank co-authors on papers that are not in this dissertation, especially Ashwini Deshpande, Deepti Goel, Udayan Rathore, Mark Granberg, and Ben Hyman.

I am indebted to my family, friends, and colleagues for their unwavering support throughout the process; Savita Khanna, Hillary Piccerillo, Kanishk Khanna, Jeevant Rampal, Sidra Haye, Zarak Sohail, Nishtha Sharma, Timothy Young, Brittany Bass, and Mark Hup. I would like to thank Melissa Valdez, Katie Holland, and Adam Cook for administrative support.

I would like to thank Elsevier for agreeing to publish the first and third chapter, and acknowledge their permission for me to reprint it here, as credited above. For Chapter 1, reprinted from *Labour Economics*, I would like to thank two anonymous referees and the editor, Arthur Van Soest. For Chapter 3, reprinted from the *Journal of Urban Economics: Insights*, I would like to thank two anonymous referees and the editor Stuart Rosenthal. I am also grateful to participants at the UCI Experimental Economics Workshop, the UCI Applied Microeconomics Workshop, and the UCI Labor Public Seminar.

I would like to thank the School of Social Sciences and the Department of Economics at the University of California, Irvine for financial support.

VITA

Shantanu Khanna

EDUCATION

Ph.D. in Economics (2022)	University of California, Irvine
M.Phil. in Economics	Delhi School of Economics, University of Delhi
Master of Arts in Economics	Delhi School of Economics, University of Delhi
Bachelor of Arts (Hons.) in Economics	University of Delhi

TEACHING

University of California, Irvine, Teaching Assistant

Intermediate Microeconomics, Probability and Statistics, Principles of Microeconomics, Principles of Macroeconomics, Managerial Economics

University of Delhi, Ad-hoc Assistant Professor of Economics

Intermediate Microeconomics, Mathematical Methods, Statistical Methods for Economics, Economic Development and the Indian Economy, Introductory Microeconomics, Principles of Macroeconomics

GRANTS & FELLOWSHIPS

Ford Foundation Grant, 2021-2023, "An Empirical Investigation of Child Labor in India" with U. Rathore

Werner Fellowship, Fall 2020, UC-Irvine School of Social Sciences

Economics Research Summer Fellowships, 2018, 2019, 2020, 2021 UC-Irvine Department of Economics

PROFESSIONAL SERVICE

Referee: *Review of Development Economics, Asian Development Review, Indian Economic Review*

CONSULTING

The World Bank, New Delhi (Social Observatory), Short Term Consultant, May - Sep 2015

The World Bank, New Delhi (Global Practice on Health, Nutrition and Population), Short Term Consultant, Nov 2014 - Apr 2015

PricewaterhouseCoopers Pvt. Ltd., Consultant, Government Reforms and Infrastructure Development, Aug 2009 - May 2010

SOFTWARE

Python/ STATA/ oTree

OTHER INFORMATION

Citizenship: India

Security Clearance: Special Sworn Status (U.S. Census Bureau)

ABSTRACT OF THE DISSERTATION

Essays in Labor Economics

By

Shantanu Khanna

Doctor of Philosophy in Economics

University of California, Irvine, 2022

Distinguished Professor David Neumark, Chair

This dissertation consists of three essays in labor economics. Broadly, they explore the labor market responses of firms and workers to various economic policies. The first chapter uses a lab experiment to study the impacts of salary history inquiry bans on wage bargaining. I find that these laws may not work as well as intended. The second chapter uses text analysis of reviews to explore the impacts of minimum wage changes on restaurant prices, hygiene, staff friendliness, wait times, and portion sizes. I find evidence of a rise in prices and improvements in staff friendliness with higher wages. While there is some weak evidence for deteriorating hygiene standards, there are no detectable impacts on portion sizes or wait times. Overall, these changes are associated with a small but significant drop in restaurant ratings. The third chapter explores the early impacts of the Opportunity Zone program on residents of targeted areas. The Opportunity Zone program, created by the Tax Cuts and Jobs Act in 2017, was designed to encourage investment in distressed communities across the U.S. We leverage restricted-access microdata from the American Community Survey and employ a matching approach to estimate causal reduced-form effects of the program. Our results point to little or no evidence of positive effects of the Opportunity Zone program on the employment, earnings, or poverty of zone residents.

Chapter 1

Salary History Bans and Wage Bargaining: Experimental Evidence

"Using salary history to determine compensation perpetuates a system that pays women less than their male counterparts"

-Susan Eggman (D), principal author of California AB 168.[2]

"The Ordinance violates employers' First Amendment rights to ask about, and rely on, wage history, and is not supported by any tangible evidence that these practices perpetuate wage discrimination."

-The Chamber of Commerce for Greater Philadelphia on the Salary History Ordinance.[3]

I thank David Neumark, John Duffy and Michael McBride for numerous discussions and valuable suggestions. I am also grateful to participants at the UCI Experimental Economics Workshop and the UCI Applied Microeconomics Workshop, especially Sidra Haye, Brittany Bass, Timothy Young and Patrick Julius. I would also like to thank Arthur Van Soest and two anonymous referees at *Labour Economics* for valuable comments.

[2]https://a13.asmdc.org/press-release/eggman-bill-would-ban-employers-seeking-applicants%E2%80%99-salary-history

[3]http://news.chamberphl.com/2017/04/the-chamber-challenges-the-philadelphia-salary-history-ordinance

1.1 Introduction

As part of the larger push for Equal Pay legislation in the United States, several states and municipalities have recently passed laws regarding revelation of salary histories. While the details of these laws differ across jurisdictions, they have several common features. These laws explicitly prohibit employers from seeking salary history information of potential employees during the hiring process.[4] While some laws are silent on employees volunteering their salary history information, others have explicit clauses that allow for this possibility (California and Connecticut, for instance). In fact, the California law states "If an applicant voluntarily and without prompting discloses salary history information to a prospective employer, nothing in this section shall prohibit that employer from considering or relying on that voluntarily disclosed salary history information in determining the salary for that applicant".[5] Table 1 lists the municipalities and states that have passed these laws, with the date of enforcement and restrictions, if any.[6]

The stated goal of these laws is to prevent the persistence of low salaries among groups that were underpaid due to discrimination to begin with, especially women. The aim is to ensure that low paying jobs held early in a woman's career (or poorly negotiated ones) do not have an impact on future salaries. Supporters claim that the potential impact of this law will be to raise the relative wages of women compared to similarly qualified men. Opponents claim that it would complicate the hiring process and increase litigation risk. They believe that open discussion about salary history as well as future salary expectations actually benefit employees in making decisions about jobs. Since this is a contentious issue, the laws have faced backlash from some members of the business community. For instance,

[4] Some laws also require employers to provide pay scale information for job vacancies to applicants upon reasonable request, but I do not examine this aspect of the laws in my experiment.

[5] `https://leginfo.legislature.ca.gov/faces/billCompareClient.xhtml?bill_id=201720180AB168`

[6] For an up to date list, see: `https://www.hrdive.com/news/salary-history-ban-states-list/516662/`

the Greater Philadelphia Chamber of Commerce filed a lawsuit questioning the validity of the law and filed a motion for preliminary injunction. More importantly, two states (Michigan and Wisconsin) have passed laws that explicitly ban the passage of such laws by local governments within the state with the intention of protecting employers' right to solicit salary information from prospective employees. Barring these two exceptions, the rapid expansion of these laws across various states suggests that several other states may follow suit, and a federal law may be in the cards as well. Thus far, 16 states and 10 cities have passed some form of ban on salary history inquiries.[7]

Because the laws are recent, empirical evidence of the impact of these laws on salaries for new jobs is scant. Aside from issues of data and timing, identifying the causal impact of such a law using observational data may be challenging since the laws are sometimes announced months or years before being enforced, and even when enforced, employers are only liable for penalties at a much later date, giving employers several months to address salary disparities. In some cases, there are other contemporaneous policy changes which may have confounding effects making it harder to disentangle the impact of the law under scrutiny.[8] Real world bargaining is complex and different features of these laws may have countervailing effects on wage outcomes. The main focus on this paper, therefore, is to examine the impact of certain aspects of such policies in a simple game-theoretic framework and to test those predictions in the lab. The controlled setting of the lab enables comparisons between treatment conditions which would not be possible with observational data. I examine the intended benefits of the policy by comparing the private information case with a perfect information benchmark. The efficacy of these laws may be affected by the fact that the laws either explicitly or implicitly allow prospective employees to voluntarily reveal their salary history to employers. Employers may treat non-revelation of salary history as a negative signal, since those with

[7]Civic Impulse. (2018). H.R. 2418 — 115^{th} Congress: Pay Equity for All Act of 2017. Retrieved from `https://www.govtrack.us/congress/bills/115/hr2418`

[8]One example of this is California's amendment AB 1676 to the Equal Pay Act, which prohibits employers from justifying a sex, race, or ethnicity-based pay differences solely on the grounds of prior salary. This amendment took effect on the 1^{st} of January, 2017

higher salaries would be much more likely to voluntarily disclose such information. To this end, I examine how the introduction of a voluntary information revelation choice can affect offers and final payoffs. Finally, I am also interested in understanding if individuals with an informational advantage behave strategically to misrepresent that information if they are free to do so, and how these results compare to other information regimes. Along with of perfect information, this regime is also an important policy counterfactual to consider, since it is unclear whether real world salary revelations are necessarily truthful.

I conduct an ultimatum bargaining experiment with different information treatments regarding disagreement payoffs (the outside option). While the outside option of the proposer (A) is always fixed and known, the outside option of the responder (B) can be one of two types, low or high. In the perfect information treatment (PER) B's outside option is known to A. In the second treatment (PRI), it is privately known, while in the third I introduce a voluntary disclosure choice for B to truthfully reveal their outside option before the proposer makes an offer (PTR). In the final cheap-talk treatment, I allow B the choice to misrepresent their outside option (PCT). In a way, the third and the fourth treatment represent important modifications to the private information case that lie at two extremes; the former representing an environment where all revelations are truthful while the latter allows for deception that has no explicit punishment or cost. Using this framework, I answer the following questions: Under which conditions are offers observed by those with the low outside option higher under private information as compared to perfect information? If offers are indeed higher, does adding a truthful information revelation choice before the offer is made affect the benefit of private information because of inferences the uninformed party can make based on revelation behavior? Do informed agents strategically misrepresent their outside option for perceived benefits? What implications do theory and the experimental results have for salary history bans?

I find that introducing a truthful revelation choice for the informed party (B) undermines some of the benefits of private information, and outcomes come closer to the perfect information case. Based on the revelation choices of B, therefore, A can make some inferences about their opponents outside option. Confirming theoretical predictions, cheap-talk outcomes resemble those under private information since revelation signals are uninformative. This is because the overwhelming majority of those with a low outside option either hide or misrepresent their outside option as high in the hope of eliciting higher offers. The superior outcomes for those with a low outside option under private information reflects the intended effect of a policy that gives an informational advantage to one party in bargaining. Inequality is lowest under private information and highest under perfect information. The experiment also demonstrates that the voluntary revelation feature of such policies may erode the informational advantage. Inequality is higher in this treatment as compared to the private information case. The similarity in results of the cheap-talk treatment and private information indicates that if we believe that potential employees can lie about their outside options and that employers know this, then their information is essentially private to begin with.

Admittedly, real world wage bargaining may be more complicated than the ultimatum bargaining framework chosen here. I consider only two outside options, whereas an employer may have a continuous probability distribution over a range of possible salaries in mind while negotiating with an employee. Salary histories may not be exactly known by employers (perfect information) and it isn't realistic to expect a situation in which they are private and we can compel all employees from revealing them (private information). To motivate these modeling choices, note that in a representative sample survey of U.S workers, Hall and Krueger (2012) find that only one-third did not consider their offer to be a take-it-or-leave-it offer, which suggests that ultimatum bargaining may not be uncommon. Half the respondents also said that their employers had learned their pay in their earlier jobs before making the offer that led to the current job. The private information case in this paper is a benchmark that indicates the policy intention, and is useful to study as such to reveal the full benefits of private information. It may also not be common knowledge that everyone who chooses to reveal their prior salary does so truthfully (truthful revelation treatment) or one in which lying about your previous salary is as easy as the click of a button without any chance of being found out or any real consequence (cheap talk). There is some evidence, however, that misrepresenting current salaries is a very real phenomenon. For instance, a labor market study of 60,000 respondents across 28 countries in Europe conducted by an international data company, Intelligence Group,[9] revealed that around 25% of Europeans lie about salaries during negotiations. This number is as high as 43% in countries like Germany and much lower in the United Kingdom and Greece. Given the sensitive nature of such a question, it is conceivable that these underestimate the actual degree of lying about previous salaries. A recent survey by Agan et al. (2020) also finds that only 5% of respondents believe with certainty that their current salary disclosures would be verified by potential employers. Thus, while understanding the limitations of modeling choices that yield clean and empirically testable predictions for this lab experiment is important when thinking about policy implications, there is evidence to suggest that even these simple cases may have wider implications for labor markets. A more general theoretical framework with less restrictive modeling choices which can also be adapted to study the salary history laws can be found in Cullen and Pakzad-Hurson (2019). My experiment tests the theoretical prediction on an extension of their model which also argues that these laws may be ineffective, for exactly the same reason as elaborated upon in this paper. This is discussed further in the next section.

The rest of this paper is organized as follows: Section 1.2 discusses selected literature most closely related to my experimental design and the policy question. The theoretical framework and hypotheses are discussed in section 1.3. Section 1.4 outlines the experimental design, section 1.5 presents the results of the experiment and section 1.6 concludes with a brief discussion of mechanisms and policy implications. The appendix contains supplemental figures, experiment instructions, and some screenshots from the experiment.

[9] `https://intelligence-group.nl/en/news/quarter-of-europeans-lie-about-their-salary`

1.2 Related Literature

The experimental literature on bargaining is vast (see Roth (1995); Dhami (2016); Rapoport et al. (2008); Murnighan (2008) for reviews). Here, I briefly mention results of experiments that directly relate to the framework and treatments examined in this paper. Güth et al. (1982) were the first to document the now well known finding of deviations from subgame perfection and a high number of equal split outcomes in ultimatum games. This has been modeled explicitly to take fairness considerations into account in several papers since. While I mention the game theoretic predictions in Section 1.3, my hypothesis will test the differences between treatment conditions as opposed to theoretical point predictions for each. Under private information about one of two pie sizes (large and small), Güth et al. (1996) found that proposers who knew about the exact size of the pie pretended fairness by offering a fair share of the smaller pie to the uninformed responders. In Kagel et al. (1996), players engaged in ultimatum bargaining over chips with different monetary payoffs assigned to chips, and different treatments about which party was informed about the exchange rates. Consistent with self-serving notions of fairness, the authors find that when only the proposer is informed about their higher payoff, they appear to go for more even splits in chips, but in the case where they have lower exchange rates, they propose unequal splits in chips, equalizing the monetary value in that case. Incomplete or private information in bargaining has most often been modeled this way. In my experiment, information does not relate to the size of the pie being divided or monetary value of chips, but rather the outside option of the responder.

Schmitt (2004) conducted ultimatum bargaining games with different outside options, information treatments and payoffs and found offers were rejected more often when responders had an outside option. While I choose one particular set of parameters across all treatment conditions, a recent paper Hennig-Schmidt et al. (2018) studied how systematically varying the size of the outside options in a perfect information setting can affect offers and rejections. They find that proposers and responders both increase their demands as outside options increase, with higher rejections in cases where there is a large difference between the outside options of the two players.

The final treatment in this paper is also related to another strand of experimental literature in bargaining on cheap talk, deception and retribution (Croson et al., 2003; Kriss et al., 2013; Boles et al., 2000). Croson et al. (2003) conduct an ultimatum game experiment with two-sided imperfect information focusing on the short-term and long-term consequences of cheap talk. Cheap talk messages took place before the proposer made an offer. They took the form of the responder misrepresenting their outside option, threatening rejection, making demands or the proposer misrepresenting the unknown pie size, making claims about the offers fairness or responding to threats. They find that responder lies about their outside option significantly increase offers. Threats of rejection had the same effect. They also find that when proposer lies were revealed (long-run effect) they made higher offers in subsequent play of the game but were also less likely to have them accepted. Besancenot et al. (2013) also conduct an ultimatum game where responders have imperfect information about the endowment of the proposers who can send messages about their endowment. They show that 88.5% of proposers understate their endowment by around 20%, and for 1 euro of gap between actual and stated amounts, reduce their proposals by 19 cents. Instead of relying on a continuous endowment distribution as in Besancenot et al. (2013), Kriss et al. (2013) rely on a binomial with similar results, finding relatively high dishonesty in comparison to other studies. Anbarcı et al. (2015) study the *taxicab* game, a modification of the standard ultimatum game, where messages are sent along with offers by the proposer and the responder has to accept or reject after observing the message. Treatments vary the probability that the actual offer will be seen as well (0 or 0.5). Results show that lying is very common but not universal even when it is completely undetectable. As the probability of observing actual offers increases, offers increase and the extent to which messages overstate them falls.

My focus is less on the extensive margin impact of this policy and whether employers choose to offer employment or who they decide to interview. Thus, instead of studying market interactions between employers and employees (analyzed using market experiments), I focus on the dyadic interaction between an already paired employer-employee at the beginning of the bargaining process. There is evidence that employers use salary history information to "weed out" some applicants, insofar as wage history may reflect productivity. My experiment abstracts away from productivity concerns. Even so, it is important to discuss results from a recent paper by Barach and Horton (2017) that mainly addresses these concerns. They conducted a field experiment in which treated employers in an online labor market could not observe the wage histories of their applicants, whereas control employers could. They find that without access to applicant salary histories, employers responded by enlarging the pool of applicants they considered by evaluating about 7% more applicants than control employers. They also evaluated those applicants more intensively, asking more detailed questions. Treated employers were also more likely to evaluate workers with lower previous wages. Applicants who were "called-back" by treated employers had about 7% lower average past wages. Workers hired by employers who could not observe past wages had about 16% lower past average wages. The explanation provided for these findings is one of "bargain hunting", wherein treatment made low experience/low wage bid[10] applicants appear to be a better bargain relative to high experience/high wage bid applicants, as firms with less information infer the workers have closer to the average level of productivity. Treatment did not alter the probability of bargaining, but when bargaining occurred, workers hired by treated employers struck a more favorable wage bargain; they were offered and accepted wages which were 9% more than their initial bid as compared to workers bargaining with control employers. Based on these results, the authors claim that the salary history bans would have their intended effect as they would help relatively less experienced workers get their foot in the door. Elsewhere, the authors have acknowledged that these results may not generalize to low skilled or entry level jobs.[11] Other key features of the laws that could

[10] Workers could compete in this market by submitting wage bids

[11] `https://www.nytimes.com/2018/02/16/business/economy/salary-history-laws.html`

not be studied through this field experiment is the impact of group identity as well as the voluntary revelation of salary history. The applicants in this experiment were unaware that employers could not see their compensation history.

A recent working paper by Cullen and Pakzad-Hurson (2019) studies the equilibrium impact of pay transparency laws in the context of online labor markets. These are distinct from the laws that motivate this paper in that they relate to workers' knowledge of the pay of peers, as opposed to firms knowledge of salary histories. Even though online labor markets employ only a small percentage of the U.S. workforce, they provide an interesting context because they are less regulated. As the authors point out, internet platforms such as eLance and UpWork have moved in the opposite direction of the laws I study, as they include accounts of past contract payments on worker profiles. While the main model of Cullen and Pakzad-Hurson (2019) is set up to study their empirical application of online markets and pay transparency across peers, their model is general enough to account for several extensions which includes the case I study here. This is explored in Appendix H of their paper (pp. 45) where they write, "we have reason to believe laws prohibiting employers from asking about pay history will be ineffective...workers with the highest wages will find it in their interests to disclose their work history to new potential employers. The unraveling logic of Theorem 4 holds, and all workers should voluntarily reveal their previous wages in equilibrium." The results of my experiment are consistent with this hypothesis. The authors also speculate that these bans may have larger effects in online labor markets where communication between workers and employers is limited, as in Barach and Horton (2017).

With regards to salary history bans in particular, a recent working paper by McNichols (2019) is the first to look at the early impacts of this policy in California. This paper uses a synthetic control approach and finds that statewide female-male earnings ratio increased from 0.77 to 0.81, with results being primarily driven by male-dominated industries. Finally, in a survey of 504 Americans conducted in November 2019, Agan et al. (2020) find that about a quarter of respondents were asked about their salary at some point during the application process. The survey results also suggests the possibility of unraveling, as workers with more attractive salaries are very likely to disclose, and since even though those with lower salaries are less likely to disclose at all, they are more likely to do so as others disclose salaries.

1.3 Theoretical framework and Hypotheses

1.3.1 Theory

Consider a two player ultimatum bargaining interaction with an outside option with risk-neutral agents. In the first stage, A makes an offer $w \in [0, \overline{W}]$ which player B can accept or reject in Stage 2. The monetary payoffs will be as follows.

If B accepts:

$$U_A = \overline{W} - w \tag{1.1}$$

$$U_B = w \tag{1.2}$$

If B rejects:

$$U_A = O^A \tag{1.3}$$

$$U_B = O^B \tag{1.4}$$

where w is the offer, $\overline{W}$ is the maximum offer possible, and O^A and O^B are referred to as the outside options. While O^A is assumed to be fixed and known by both A and B throughout, $O^B \in \{O^B_L, O^B_H\}$, that is, it can either take a low value or a high value with probabilities P_L and P_H, respectively. I will also assume $O^A + O^B < \overline{W}$ (gains from trade) and $O^A > 0$, $O^B_H > O^B_L > 0$. The following information treatments refer to whether O^B is privately known by B or common knowledge (known by A).

1.3.1.1 Perfect information (PER)

Under perfect information, O^A and O^B are known to both A and B. There is a unique subgame perfect Nash equilibrium (SPNE), in which the proposer (A) makes an offer $w = O^B$, and B accepts. I will assume throughout that B will accept in case of indifference so as to keep the notation simple.

1.3.1.2 One-sided private information (PRI)

In this case, only B knows the exact value (realization) of their outside option. A knows that $O^B \in \{O^B_L, O^B_H\}$.

Since there are only two possible outside options for B, in equilibrium, A will offer either O^B_L or O^B_H.

Expected payoff from offering O^B_L:

$$EU_A(w = O^B_L) = P_L[\overline{W} - O^B_L] + P_H[O^A] \tag{1.5}$$

This is because all the low B types will accept the offer, giving A the remaining surplus, whereas all the high B types will reject it, which will give A their outside option.

Expected payoff from offering O_H^B:

$$EU_A(w = O_H^B) = \overline{W} - O_H^B \tag{1.6}$$

In this case, all the B types will accept the offer.

Thus, if $EU_A(w = O_H^B) > EU_A(w = O_L^B)$, then A should offer O_H^B, otherwise A should offer O_L^B. The former is accepted by all B types, whereas only those with $O^B = O_L^B$ will accept the latter. Thus, the choice of offers depends on the parameters O_L^B, O_H^B, $\overline{W}$, O^A, P_L and P_H.

1.3.1.3 One-sided private information with truthful revelation (PTR)

In this modification of the ultimatum game, O^B can be low or high with certain probabilities as before, but there is an added decision for B to reveal their outside option (O^B) to A *before* an offer is made (Stage 0). A result that follows immediately is that there is no incentive for any player B type with O_H^B as their outside option not to reveal their outside option, in any equilibrium.

Knowing this, any non-revelation by player B signals to A that they are dealing with a low B type ($O^B = O_L^B$ is assigned a probability of 1). It follows that even low B types gain nothing by non-revelation and we end up with a separating Perfect Bayesian equilibrium (PBE). Thus, any advantage for the informed parties in the case of private information should unravel completely when revelation is made a strategic choice. Offers under this information treatment should therefore be much like the perfect information case.

The argument carries through to the case where $O^B \sim U[O_L^B, O_H^B]$ or any other continuous distribution by analogy. [12]. These results may have some implications for the efficacy of the

[12]Essentially if there is any interor threshold value for the oustide option above which B types reveal their

salary history bans. As demonstrated above, under certain conditions the case of one-sided private information without a chance for truthful revelation (Section 1.3.1.2), analogous to a "don't ask, don't tell (your previous salary)" policy may better for the informed group with a low outside option when compared to the case where truthful revelation is a choice.

1.3.1.4 One-sided private information with cheap-talk revelation (PCT)

Next, consider a modification of the previous case where instead of only allowing for a choice to not reveal or reveal their outside option truthfully, revelation involves a choice from the set of both outside options $\{O_L^B, O_H^B\}$. Thus, we allow for the possibility that B can misrepresent their outside option.

Observing a signal of O_H^B is uninformative, as there is an incentive for B with a low outside option to claim that their outside option is O_H^B as well. The offer in response to such a signal should therefore be the same as derived in the section 1.3.1.2 on private information, as the prior beliefs are relied on.

1.3.2 Hypotheses

The following parameterization was chosen: $\overline{W} = 30$ (the size of the pie to be divided), $O^A = 5$ (A's outside option), $O_L^B = 5$ and $O_H^B = 20$ (B's low and high outside option, respectively). Both outside options were equally likely for B ($P_L = P_H = \frac{1}{2}$).

The low and the high outside option are on either side of an equal split (15) and chosen to be far apart. Another consideration that was kept in mind was based on an extension of this current paper which will use a more complex bargaining protocol, namely, Rubinstein's

outside option, then a non-revelation signals to A that they are dealing with a B type with an outside option which can be at maximum equal to the threshold. Those B types at the threshold will have an incentive to deviate, lowering the threshold itself. This unraveling continues until full transparency.

offer-counteroffer bargaining instead of ultimatum bargaining. The SPNE prediction under that bargaining protocol lies in between the two outside options (approximately a 17-13 split in favor of the first mover when the continuation probability is set at 0.75). Thus, there exists a tension between the high outside option and these theoretical and behavioral predictions.

Most importantly, this parameter choice also ensures that under the assumption of log utility (risk aversion), the proposer (A) has an equal expected utility from choosing the low or the high outside option as their offer in the private information treatment. Under the assumption of expected payoff maximization (risk-neutrality), theory would predict A would always choose the low outside option.[13]

For most bargaining situations tested in the lab, point predictions based on sub-game perfection are rarely borne out. As mentioned before, other regarding preferences are a concern and equal splits are focal points. Therefore, hypotheses relate to differences between treatments (taken pair-wise). The central outcome of interest, will be the offers observed or accepted by B with a low outside option (w_l). I will consider the mean, median, and the distribution of these offers across treatments. Other outcomes of interest will be the differences between offers observed or accepted by those with a low outside option and those with a high outside option, and the inequality measures under different information regimes. For the main outcome of interest w_l, I have the following hypotheses:

Hypothesis 1 and Hypothesis 2 ($H1$ and $H2$): Signaling

The prediction under private information is that under the chosen parameterization with risk-neutrality, A will always make the offer of $w = O_L^B$. This is equivalent to the perfect information offer to those with a low outside option. Alternatively, we could anticipate that private information will yield higher offers to those with the low outside option as compared to perfect information. Under truthful revelation, the prediction is that any informational advantage to those with the low outside option will unravel completely. Thus, comparing

[13] Expression (1.5) would equal 15, whereas expression (1.6) would be 10.

private information and truthful revelation is analogous to comparing private and perfect information. This leads to:

$$H1 \text{ Null: } {w_l}^{PRI} = {w_l}^{PTR}$$
$$H1 \text{ Alternative: } {w_l}^{PRI} > {w_l}^{PTR}$$

The theory in section 1.3.1.3 predicted equivalence between truthful revelation and perfect information regimes. Yet, if enough of B types with a high outside option choose to hide their outside option, it is possible that the outside option information under the truthful revelation treatment stays private, leading to spillovers to those with the low outside option (as under private information). I test this behavioral hypotheses as the alternative in $H2$:

$$H2 \text{ Null: } {w_l}^{PTR} = {w_l}^{PER}$$
$$H2 \text{ Alternative: } {w_l}^{PTR} > {w_l}^{PER}$$

The separating equilibrium consistent with revelation choices having a strong signaling value should be evidenced by rejecting the null under $H1$ and not rejecting it under $H2$.

Hypothesis 3 ($H3$): Cheap Talk

Under the cheap talk treatment, revelation signals will be uninformative, which leads to a null that there is equivalence between this treatment and private information. The alternative is simply that these two treatments are different from each other.

$$H3 \text{ Null: } {w_l}^{PCT} = {w_l}^{PRI}$$
$$H3 \text{ Alternative: } {w_l}^{PCT} \neq {w_l}^{PRI}$$

1.4 Experimental Design

1.4.1 Treatment design

The treatment design corresponds to the four cases discussed in the previous section, namely, perfect information, private information, truthful revelation and cheap talk. To avoid order effects, I use a 4 by 1 between-subject design (see Table 1.2). I ran two sessions for each treatment condition. The sessions were conducted with undergraduate students at a large public university in the United States between August and November, 2018. The experiment was coded using the python-based software OTree (Chen et al., 2016). There were a total of 174 participants across all treatments and 1,740 distinct bargaining interactions. 70.3% of participants were female.

1.4.2 Sessions

Participants were first randomly assigned to the role of either A or B and retained this role throughout the session to aid in learning. A session consisted of participants engaging in 20 distinct repetitions of ultimatum bargaining with random re-matching for each. Hereafter, each of these 20 interactions will be referred to as trials. In each trial, half the B's were assigned a low outside option of 5, while the other half were assigned a high outside option for 20. Typically, a trial began with A making an offer to B anywhere between 0 and 30, which B could then accept or reject. In the truthful revelation treatment, the only modification to this was that B first had an added decision to either click "Do Not Reveal" or "Reveal" regarding their assigned outside option for that trial *before* A made an offer. In the cheap talk treatment, they could either reveal 5, or reveal 20, or not reveal anything. As examples, screen-shots of the revelation decision page, the offer page and the accept/reject decision page for the truthful revelation treatment are provided in the Appendix.[14]

Apart from the show-up payment, two trials were selected at random for payment. Participants earned points in each trial and 3 points were equivalent to 1 US dollar. On-screen instructions were supplemented by a hard copy of instructions provided at the beginning of the session (See the Appendix for instructions for all four treatments). Five short comprehension questions were also asked to ensure participants understood the instructions well. At the end of the session, participants were asked their major and gender. They also answered three standard Cognitive Reflection Test (CRT) questions[15] and were asked to briefly describe the general strategies they followed during the session.

1.5 Results

To get a broad picture of what the data look like across treatments, I first pool all the observations from each treatment. Recall that there were a total of 1,740 distinct bargaining interactions across the four treatments. The values in the parentheses of column 2 in Table 1.2 provide the number of observations for each treatment. To begin with, therefore, each distinct bargaining outcome of a one-shot ultimatum game is treated as an observation. In section 1.5.6, I explicitly account for the dependence of observations induced by each participant engaging in 20 repeated one-shot trials by using panel regressions to estimate average treatment effects.

[14]Screen-shots for the other treatments are analogous and available with the author upon request.
[15]CRT questions are provided in the Appendix

1.5.1 Aggregate Analysis of Offers (Pooled Sample)

Consider the average offers made by A to both the low and high B types across the four information treatments shown in Figure 1.1. Comparisons both across and within the two panels of the figure are informative. Under private information (PRI), without any information on the outside option, we observe very little difference across the two outside options. While large differences across the two panels in case of perfect information (PER) are expected since the exact outside option is known in each trial, the same result holds under the truthful revelation treatment (PTR). This is important from a policy perspective since part of the goal of salary history bans is to reduce the differences in offers made to those with high outside options and those with low outside options. Adding a truthful revelation choice to private information does resemble perfect information in this respect. We will explore distributional implications more thoroughly in the next section.

Examining the offers to those with low outside options and comparing w_l across treatments (left panel of Figure 1.1) lends support to the alternative hypothesis for $H2$, and the null hypothesis for $H1$ and $H3$. In particular, the mean offers observed by B with a low outside option are highest under cheap talk (14.7) and private information (14.4), and these are not statistically different from each other. While the average offer under truthful revelation (13.6) in the left panel of Figure 1.1 is lower than those under private information or cheap talk, these differences are not significant. Thus, even though adding the revelation choice may undermine the benefits of private information to some extent, the average under truthful revelation are still higher than that under perfect information (12.4). In later sections, I explore how robust this result is accounting explicitly for individual fixed effects and examine how the revelation behavior of B under the truthful revelation and the cheap talk treatment can explain some of these findings. Figure 1.2 depicts only the accepted offers and tells a similar story, when we compare across treatments for accepted offers by low B types.

1.5.1.1 Non-parametric tests

Next, I turn to non-parametric tests of equality of medians. This is used to determine if the offers faced by those with the low outside option arise from the same distributions (or distributions with the same median) across treatments, without relying on normality assumptions. The results for the test of equality of medians is presented in Table 1.3.[16] The test of equality of medians indicates that the Null of $H1$ and $H2$ are rejected while the null of $H3$ can not be rejected. Taken together, these results show that offers made to those with a low outside option are highest under cheap-talk and private information (PCT and PRI), are lowest under perfect information (PER), with the truthful revelation treatment (PTR) lying somewhere in the middle (that is: ${w_l}^{PCT} \approx {w_l}^{PRI} > {w_l}^{PTR} > {w_l}^{PER}$). This result is robust even if I drop all the observations that are exactly at the median. This result also holds if we consider final payoffs for low B types instead of offers. It is important to recall that the pooled sample looks only at the averages over the twenty trials that comprised each treatment. For more evidence of learning behavior of participants while making offers, see Figure A.1 in the Appendix. The learning behavior indicates that increasing the number of trials would reinforce this conclusion, if offers continued as per trends.

1.5.2 Distributions of Offers and Inequality

1.5.2.1 Comparing offer distributions across treatments

In this section, I consider the distribution of offers under different treatments. Figure 1.3 shows the histogram of offers across the four treatments with a chosen bin size of 1 point. In the perfect information treatment, when A was faced with B with a low outside option of 5, the most common offer was a proposal of an equal split, and the second most common

[16] The results of the Mann-Whitney test are consistent with the result for the test of equality of medians.

offer was 10. When faced with a B with 20 as their outside option, the most common offer was 20, followed by the equal split.

Consistent with the null hypothesis of Hypothesis 2, I find that the pattern of the two most common offers are identical in the truthful revelation and perfect information treatments (15 and 10 to those who have an outside option of 5, and 20 and 15 to those with an outside option of 20). Thus, it appears that the revelation behavior by B that preceded A's offer was informative for A in inferring which type they were faced with. However, there is also some suggestive evidence that this was imperfect. For instance, 40% of offers to those with a low outside option under truthful revelation were at 15 exactly, whereas under perfect information the corresponding proportion at 15 was 32%. Another difference was that about 10% of those with a low outside option got an offer of 20 under truthful revelation while only 3% did under perfect information. Both these factors led to a higher average offer under truthful revelation compared to perfect information, as was documented in Figures 1.1 and 1.2.

Under private information, the most common offers were 15, followed by 20 and 10, and the distributions were similar across the two outside options, which is expected since the outside option was unknown, by definition. There is also some mass between 15 and 20 under this treatment. Furthermore, the three most common offers to those with an outside option of 5 are the same under private information and cheap talk (these are 15, 20 and 10), lending support to the null under Hypothesis 3.

For more formal tests of the differences in distributions, consider the results of the Kolmogorov-Smirnov (K-S) tests presented in Table 1.4. In line with the hypotheses, we compare differences in the offer distributions observed by low-type B for all treatment condition pairs among the first three treatments as well as an additional comparison between private information and cheap talk. The first row tests the hypothesis that offers for that treatment contain smaller values than that for the comparison treatment group. The second row tests the same for the other treatment. The value under the difference column represents $D^+ = Max_x\{F(x) - G(x)\}$ for row 1 and $D^- = Min_x\{F(x) - G(x)\}$ for row 2. $F(x)$ and $G(x)$ are the CDFs of offers to those with the low outside option for the pair of treatments under comparison.

Comparing perfect and private, we see that offers under perfect information are significantly smaller, and the reverse is not true. The value 0.30 indicates that the proportion of offers *below* the difference maximizing offer value was 30% higher under perfect than under private. Offers under truthful revelation are smaller than under private (panel 2), but the coefficient (0.099) is not statistically significant at the 10% level, though only marginally so (p=0.104). The third panel indicates that offers under perfect information are smaller than under truthful revelation, so low B types retained some of their information advantage.

1.5.2.2 Comparing inequality across treatments

Reducing inequality (in general, or between groups with low and high outside options) is a key goal of policies that motivated this experiment. In this section, therefore, I examine this aspect using the Gini coefficient.[17] I focus on examining inequality in the observed offers received and subsequent payoffs earned by B. Using the Pyatt (1976) decomposition, the Gini is broken down into three components: within-group inequality, between-group inequality, and an overlap term. Here, the two groups correspond to the two outside options of B (low and high). This will enable us to compare overall inequality across treatments, as well as the part of inequality that stems from differences in payoffs between the groups with

[17]The Gini coefficient is a summary measure of inequality that lies between 0 (perfect equality) and 1 (maximum inequality). For a discrete distrubution, the Gini is defined as:

$$G = \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} |y_i - y_j|}{2n^2\overline{y}}$$

Thus, the Gini is half the average of the absolute difference of all pairs of offers or payoffs (y) normalized by the mean.

different outside options.

The between-group term reflects the part of the total Gini coefficient that is due to differences in average offers or payoffs between the two groups with outside options. This can be interpreted as the Gini that would result if every offer for the two subgroups would be replaced by their subgroup means. The within group term reflects the inequality in offers within the two groups.[18] The residual term depends on the extent to which the income distributions of the two groups overlap, and would be zero if there are no overlaps in offers between the two groups.[19]

1.5.2.2.1 Inequality in offers made by A: Results are presented in Table 1.5. The last row of Table 1.5 shows that overall inequality as measured by the Gini is highest under the perfect information treatment, but not too dissimilar across the other three treatments. Looking at the decomposition, we see a strong similarity between private information and cheap talk, with a very small contribution of the between component. This is expected as average offers made are very similar across the two outside options, either because they were unknown (private information) or they were uninformative if revealed (cheap talk). Comparing across treatments, the between group component is largest under perfect information since outside options were known and offers reflected that. The truthful revelation treatment seems to mitigate some of the between group inequality (share of 25.6%) but falls short of the private information benchmark (4.9%).

1.5.2.2.2 Inequality in final payoffs for B: Next, consider inequality in the final payoffs earned by player B. The results are presented in Table 1.6. A notable difference between Table 1.5 and Table 1.6 is that the between group component is much larger and

[18] This term will be the weighted sum of the Gini coefficients for each outside option where the weights reflect income shares and group sample sizes.

[19] See Lambert and Aronson (1993) for a graphical analysis of these three terms using Lorenz curves.

the overlap component much smaller when payoffs are considered as opposed to offers. This is to be expected as those with a high outside option often rejected offers to walk away with 20 points (rejection behavior is discussed further in Section 1.5.5).

Once again, I find that overall inequality is highest under perfect information (0.184), and very similar under cheap talk (0.129) and private information (0.125). The Gini for the truthful revelation treatment (0.143) lies somewhere in the middle: inequality is lower compared to perfect information but higher than under private information or cheap-talk. About 72% of this higher total inequality in truthful revelation as compared to private information stems from having a higher between-group component.[20] Thus, if a policy-makers' objective function is to reduce between-group inequality, keeping information private (PRI) would be more effective than allowing a revelation choice (PTR), whether we consider offers or final payoffs.

1.5.2.3 A short note on fairness

The distributional analysis above also reveals that a large number of offers were at an equal split of 15. In some rare cases, offers that are lower than the outside option are also accepted by B. This could be an indication of other regarding preferences. The labor market analogy for such behavior would be a worker accepting a lower wage than the current pay that they are getting, and is therefore extremely unlikely to be anything other than an artifact of a lab experiment. There are 68 cases out a total of 1740 distinct interactions where agent B accepts an offer lower than their outside option, and 46 of these were offers exactly at 15. Dropping these cases from the analysis does not alter any of the conclusions of the paper. Results from Tables 1.3 through Table 1.6 are unchanged, and estimates are nearly identical. [21]

[20]Calculated from columns 2 and 3 of Table 1.6 as $\frac{0.096-0.083}{0.143-0.125} = \frac{0.013}{0.018} = 0.72$

[21]All Figures and Tables based on dropping these 68 observations are available from the author upon request.

1.5.3 Truthful revelation choices and subsequent offers (PTR)

The revelation treatment is a key treatment of interest because this treatment involved a revelation choice by B before A could make an offer. We examine these choices and the associated offers in further detail in this section. Looking at all the bargaining trials, only 15% of low type B revealed their outside option, as compared to 62% for those with a high outside option. This 47% percentage point difference is statistically significant. Figure 1.4 shows the percentage of participants that revealed their outside option for each trial. The upward trend in the second panel of Figure 1.4 indicates learning to reveal for those with the high outside option of 20. The coefficient on the slope of the best fit line has a p-value of 0.06. This is also consistent with the statistically significant trend in offers to those with the high outside option for this treatment seen in Figure A.1 in the Appendix.

The average offer to those who had an outside option of 5 and chose not to reveal was 13.7 and it was 13.2 for those who did disclose it, but this difference was not significant. For those with the high outside option of 20, revelation yielded offers on average of 17.7 whereas non-revelation yielded offers of 13.7. This difference is statistically significant. This could explain why there were patterns of learning to reveal for those with a high-outside option. Since those assigned a high outside option chose not to reveal their outside option sizable proportion of trials (38%), the informational advantage for B was partially maintained. This could benefit those with a low outside option, though not as much as if information was completely private and revelation was not a choice at all. Among the trials where the high outside options were hidden, over 77% of offers were rejected. This could therefore induce higher subsequent offers from A that could spillover to trials where outside options were low (and hidden). The distribution of offers made for these revelation decisions are given in Figure 1.5. We find that in 11% of trials, those with 5 as their outside option who chose to hide it got offers of 20, a finding consistent with a partial information advantage.

1.5.4 Cheap talk revelation choices and subsequent offers (PCT)

Recall that under cheap talk, B could decide to not reveal their outside option, or choose between claiming their outside option as either 5 or 20. The 380 distinct trials in this treatment were split equally into those where B's outside option was low or high (190 each). Table 1.7 provides the distribution of revelation decisions across the two outside options. In 57% of trials where B had an outside option of 5 it was instead misrepresented as 20. In 21% of trials, low B types revealed the truth, and in another 22% cases they chose not to reveal at all. High B-types revealed their outside option truthfully in 75% of trials. As theory would suggest, observing a revealed outside option of 20 is uninformative for A since all of these claims were unverifiable. Table 1.7 also provides the average offers made by A for different revelation decisions, across the two outside options. Note that these offers are very similar across outside options for each revelation decision (they are not statistically different from each other). Comparing across rows, note that claiming an outside option of 20 does yield a higher average offer than claiming it as 5 or hiding it. Figure A.2 in the Appendix shows the distribution of offers made by A based on the revelation decisions of B across the outside options.

1.5.5 Accept/ Reject Decisions

Table 1.8 shows the acceptance rate across treatments and outside options. This is defined as the percentage of trials where the offer was accepted. Naturally, this decision depends on the offer received, but it is still worth exploring unconditional decisions to get a sense of the efficiency of a particular information regime. Acceptance rates are much higher for those with a low outside option since the effective surplus to be divided is larger. Comparing across treatments for B's with the low outside option, acceptance rates are highest under cheap talk and lowest under perfect information. Recall that cheap talk had the highest offers on average. For B's with the high outside option, acceptance rates are lowest under private and highest under perfect information, which is to be expected as in the former there is no way to communicate that the outside option is high whereas in the latter it is known to A. The truthful revelation treatment lies between the perfect and private information treatments.

1.5.6 Panel Regressions

While the analysis so far assumed that observations were independent, since each participant engaged in 20 trials, I now turn to panel data methods to make accurate inferences about differences across treatments. As is typical in between-group analyses with repeated trials in the experimental literature, I use random effects panel regressions to explicitly account for dependence (Moffatt, 2015). I test the hypotheses separately by regressing the offers made by A to low B types on binary treatment dummies. Results are presented in Table 1.9. Participant controls for the CRT score[22] and gender are included in the regressions.

Offers observed by low type B under private information are significantly higher than perfect information. Thus, even under fairly weak conditions we can say that there is an informational advantage from private information. We cannot reject the null for Hypothesis 1 or Hypothesis 2, leading to some mixed evidence for signaling and the separating equilibrium. Based on earlier results, it is clear that there is evidence for signals being informative and some of the benefits of private information eroding, but also that the unraveling is not complete as the offers under truthful revelation are higher that perfect information, even if not significantly so. Finally, as per theoretical predictions, average offers under the cheap-talk treatment are not different from those under private information (cannot reject the null of $H3$). Adding an additional control for session fixed effects does not alter any of these results

[22] Since there are three standard CRT questions, participants could score 0, 1, 2 or 3. We control for CRT scores by including dummies, with 0 as the omitted group.

(See Table 1.10).

1.6 Discussion and Policy Implications

1.6.1 Discussion

Using the framework of ultimatum bargaining with an outside option, this paper demonstrates that adding a revelation choice to private information makes an impact on the offer distribution and consequently the mean, median, and the Gini coefficient. Some of the benefits that accrue to the informed party with the low outside option under private information are lost in the truthful revelation treatment as proposers can make meaningful inferences about the informed party based on their revelation choices. A parallel result is that this information setup lies somewhere in between the private information and perfect information case. The strategy-box comments participants entered at the end of the experiment displays some of this thinking at work. For example, one particular participant in the role of B wrote:

"I revealed when I had the outside option of receiving 20 because that would pressure the A participant into providing an offer higher either close or it or even above it as a sort of bribe to assure that I accepted. If I received an outside option of only 5 rather than 20, I would not reveal it because the mystery would leave A unsure of what to offer."

The following demonstrates that such a strategy was recognized by an A:

"I was participant A throughout the entire experiment. The method I used would be putting my bet on the fact that if B's outside option was 5, humans tend to NOT want to reveal it."

Finally, a B participant recognizing this line of thinking, went for complete transparency: *"As player B, always reveal... If my outside option is 5, it's a bit more complicated. Logically* *whether I reveal it or not, A will figure out that my outside option is 5. However, if I do reveal it, I am showing vulnerability to A and could potentially prompt him or her to have pity on me and make a higher offer. It's pretty rude, but the goal in this scenario is to play on A's emotions to have him/her offer me more points."*

It is also interesting to note the behavioral line of thinking demonstrated at the end of the comment above, where the participant banks on a reward for being transparent and honest about their low outside option. However, I find no significant differences in the offers made to those who chose to reveal their low outside option and those who chose to hide it.

Under the cheap talk treatment, a typical line of reasoning demonstrated by B was *"In the experiment, I am B. I will always bluff and say I have 20 as my outside option."* and consequently by an A *"At first I felt like the B person was telling the truth about their outside option, but toward the end I felt like they were lying."*. Some of these insightful comments may reveal some unexplored behavioral channels behind the treatment effects observed both in this paper and elsewhere in the literature which may be worth exploring in future work.

1.6.2 Policy Implications

The results of this experiment indicate asymmetric information in bargaining may be advantageous for the informed party under certain conditions, but also that signals based on a choice to reveal that information may be informative. In many ways, the full benefits of private information are hard to achieve realistically purely because in most real-world scenarios, revealing one's outside option is always a choice.

The case I demonstrate under which the salary history bans may be most effective for those with a low wage, is one in which all potential employees are barred from revealing their outside option. That is the standard private information case studied in the bargaining literature and perhaps best reflects the intent of the policy. Between-group inequality in offers across those with low and high outside options is also very low under this treatment.

In the case where revelations are truthful, the choice to reveal one's private outside option in bargaining can be informative in and of itself.[23] This can undermine the benefits of private information. This may be one reason salary history laws may not be as effective as conventional wisdom would suggest. Results also suggest that if enough of those with high-outside options hide their private information, then some of the positive spillovers may benefit those with low outside options and outcomes would not be as poor as under perfect information. But, I also find some suggestive evidence that there may be some learning involved in the experiment in revealing high outside options, and that eventually this benefit may disappear altogether, as predicted by theory. Finally, the cheap talk treatment demonstrates an extreme case in which potential employees could easily misrepresent represent their outside option (previous salary). Results here suggest that if prospective employees could lie about their current salaries in bargaining for new jobs, that scenario would be akin to private information to begin with, and thus, once again the laws may not lead to significant gains for disadvantaged groups.

1.6.3 Extensions

A natural extension is to use a more complex bargaining protocol with several (finite) stages of bargaining with offers and counteroffers or infinitely repeated Rubinstein's bargaining. In these experiments, B can use richer signals for communicating the outside option through counteroffers. This would also provide more opportunities to learn within a given interaction. Such protocols would also generate data on the stage of bargaining at which revelations or agreements occur. Since the results suggested that there is some learning to reveal under

[23]Naturally, salary history revelations may not always be truthful in reality as they are, by definition, in the truthful revelation treatment in this experiment. Thus, we can consider this case as one in which the risks of misrepresenting salary history are very high.

the truthful revelation treatment, another extension would be to increase the total number of trials to 40 instead of 20.

While the focus in this paper was just on the revelation behavior of the responder, and comparing outcomes from that treatment the case where outside options are known (perfect information), another extension would be to add the choice for the proposer to *ask* for their opponents outside option in an additional stage before revelation. This richer setting may better reflect the status quo, though we would not expect this to change results since the proposer should always ask about the outside option of the responder. Additionally, we could add a third outside option instead of two, to see the unraveling logic better. In the current experiment there is no incentive for the second highest type to reveal since they are also the lowest type. [24]

Finally, given the large literature on behavioral differences in bargaining between men and women, another extension could be to make gender identity explicit in the experiment with A's being males and B's being females in one set of sessions and the roles reversed in another.

[24] I thank a reviewer for proposing both of these extensions.

Tables

Table 1.1: Salary History Inquiry Bans in the U.S.

State	Date of Enforcement	Restrictions
New York	9-Jan-2017	State Agencies, Statewide + All Employers in NYC, Albany, Suffolk, Westchester
Louisiana	25-Jan-2017	City Agencies in New Orleans
Puerto Rico	8-Mar-2017	All employers, Commonwealth-wide
Oregon	6-Oct-2017	All employers, Statewide
Delaware	14-Dec-2017	All employers, Statewide
California	1-Jan-2018	All employers, Statewide
New Jersey	1-Feb-2018	All employers, Statewide
Utah	1-Mar-2018	City Agencies in Salt Lake City
Kentucky	17-May-2018	City Agencies in Louisville/ Jefferson County Metro Government
Massachusetts	1-Jul-2018	All employers, Statewide
Vermont	1-Jul-2018	All employers, Statewide
Missouri	26-Jul-2018	All employers in Kansas City
Pennsylvania	4-Sep-2018	State Agencies, Statewide + Philadelphia (on hold) + Pittsburgh city agencies
Connecticut	1-Jan-19	All employers, Statewide
Hawaii	1-Jan-2019	All employers, Statewide
Illinois	15-Jan-2019	State Agencies, Statewide + Chicago City Agencies
Georgia	18-Feb-2019	City Agencies in Atlanta
North Carolina	2-Apr-2019	State Agencies, Statewide
South Carolina	23-May-2019	All employers in Richland County
Mississippi	13-Jun-2019	All employers in Jackson
Washington	28-Jul-2019	All employers, Statewide
Maryland	14-Aug-2019	County agencies in Montgomery
Alabama	1-Sep-2019	All employers, Statewide
Maine	17-Sep-2019	All employers, Statewide
Ohio	1-Mar-2020	Employers with 15 or more employees in Cincinnati and Toledo
Colorado	1-Jan-21	All employers, Statewide

Table 1.2: 4 by 1 Between-Subject Design

Information Treatment	Num. of Subjects (Trials)	Avg. $
Perfect Information (PER)	44 (440)	$15.85
Private Information (PRI)	48 (480)	$16.40
Truthful Revelation (PTR)	44 (440)	$16.44
Cheap Talk (PCT)	38 (380)	$16.13
Total	174 (1740)	$16.21

Table 1.3: Non-parametric test for equality of medians

PER and PRI	PTR and PRI	PER and PTR	PRI and PCT
χ^2	χ^2	χ^2	χ^2
36.5***	4.5**	16.5***	0.665

*** indicates significance at 1%, ** at 5%, * at 10%

Table 1.4: Kolmogorov-Smirnov Test for offers faced by low B types

Perfect and Private		Truthful-revelation and Private		Perfect and Truthful revelation		Private and Cheap talk	
Smaller Group	Diff.	Smaller Group	Diff.	Smaller Group	Diff.	Smaller Group	Diff.
PER	0.308***	PTR	0.099	PER	0.214***	PRI	0.132**
PRI	-0.001	PRI	-0.001	PTR	-0.027	PCT	-0.050

*** indicates significance at 1%, ** at 5%, * at 10%

Table 1.5: Gini coefficient decomposition for offers made by A, across treatments

	Perfect		Private		Truthful Revelation		Cheap Talk	
	Gini	Share (%)	Gini	Share (%)	Gini	Share (%)	Gini	Share (%)
Between	0.084	47.2	0.008	4.9	0.042	25.6	0.006	3.9
Overlap	0.023	13.1	0.073	45.2	0.045	27.7	0.076	46.2
Within	0.071	39.7	0.081	49.9	0.077	46.7	0.083	49.9
Total Gini	0.179 (0.006)	100	0.162 (0.007)	100	0.164 (0.007)	100	0.165 (0.007)	100

*Standard Errors in parentheses for the total Gini are based on 1000 bootstrap replications

Table 1.6: Gini coefficient decomposition for payoffs for B, across treatments

	Perfect		Private		Truthful Revelation		Cheap Talk	
	Gini	Share (%)	Gini	Share (%)	Gini	Share (%)	Gini	Share (%)
Between	0.131	71.0	0.083	66.5	0.096	67.2	0.075	58.1
Overlap	0.004	2.0	0.003	2.1	0.004	2.5	0.007	5.7
Within	0.050	27.0	0.039	31.4	0.043	30.3	0.047	36.3
Total Gini	0.184 (0.008)	100	0.125 (0.007)	100	0.143 (0.008)	100	0.129 (0.008)	100

*Standard Errors in parentheses for the total Gini are based on 1000 bootstrap replications

Table 1.7: Cheap talk revelation decisions by B and associated average offers by A

	Outside Option: 5		Outside Option: 20	
	Decision Shares	Avg. Offer	Decision Shares	Avg. Offer
Did not reveal	22%	12.6	12%	12.6
Revealed 5	21%	12	13%	12.2
Revealed 20	57%	16.4	75%	15.9
Total	100% (N=190)	14.7	100% (N=190)	15

Table 1.8: Percentage Accepted (All Trials)

	Perfect	Private	Truthful Revelation	Cheap Talk
Outside Option: 5	80.5	92.5	89.1	94.0
Outside Option: 20	59.1	19.2	41.3	31.6
Number of Trials	440	480	440	380

Table 1.9: Panel Regressions: Offer observed by low B types on Treatment Dummies

	PER to PRI	H1:PTR to PRI	H2:PER to PTR	H3:PRI to PCT
D1 (0 if Perfect, 1 if Private)	2.32*** (0.64)			
D2 (0 if Truthful Revelation, 1 if Private)		0.94 (0.84)		
D3 (0 if Perfect, 1 if Truthful Revelation)			1.27 (0.83)	
D4 (0 if Private, 1 if Cheap Talk)				0.4 (0.69)
Gender and CRT Controls	Yes	Yes	Yes	Yes
Constant	12.75*** (0.55)	14.08*** (0.67)	12.95*** (0.72)	14.3*** (0.64)
N	442	442	440	412

* indicates significance at 10%, ** at 5%, and *** at 1%

Table 1.10: Panel Regressions: Offer observed by low B types on Treatment Dummies with Session Fixed Effects

	PER to PRI	H1:PTR to PRI	H2:PER to PTR	H3:PRI to PCT
D1 (0 if Perfect, 1 if Private)	2.63*** (0.64)			
D2 (0 if Truthful Revelation, 1 if Private)		0.93 (0.87)		
D3 (0 if Perfect, 1 if Truthful Revelation))			1.28 (0.84)	
D4 (0 if Private, 1 if Cheap Talk				0.35 (0.66)
Gender and CRT Controls	Yes	Yes	Yes	Yes
Session Fixed Effects	Yes	Yes	Yes	Yes
Constant	13.21*** (0.58)	14.06*** (0.79)	12.98*** (0.80)	15.2*** (0.80)
N	442	442	440	412

* indicates significance at 10%, ** at 5%, and *** at 1%

Figures

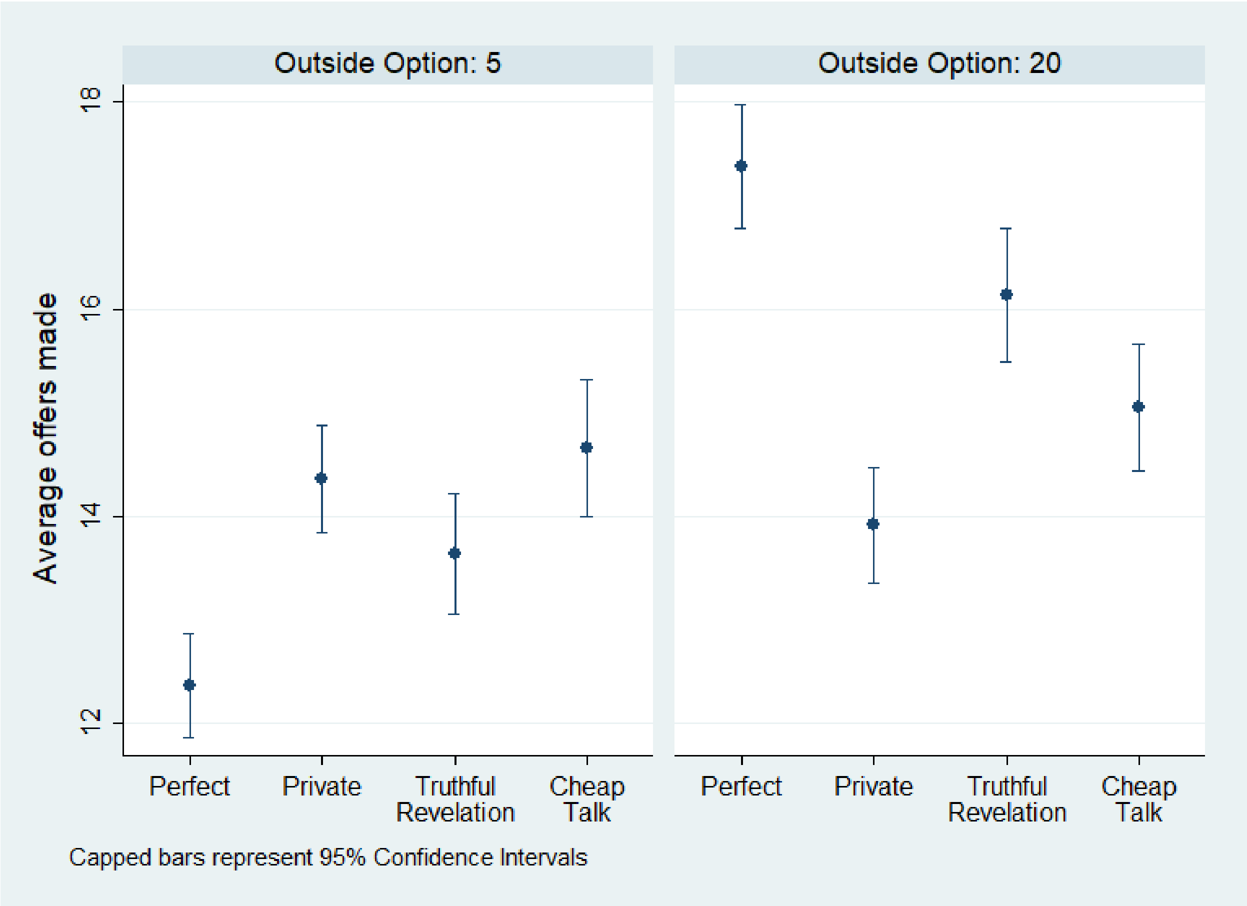

Figure 1.1: Average offers made across treatments, by B's Outside Options

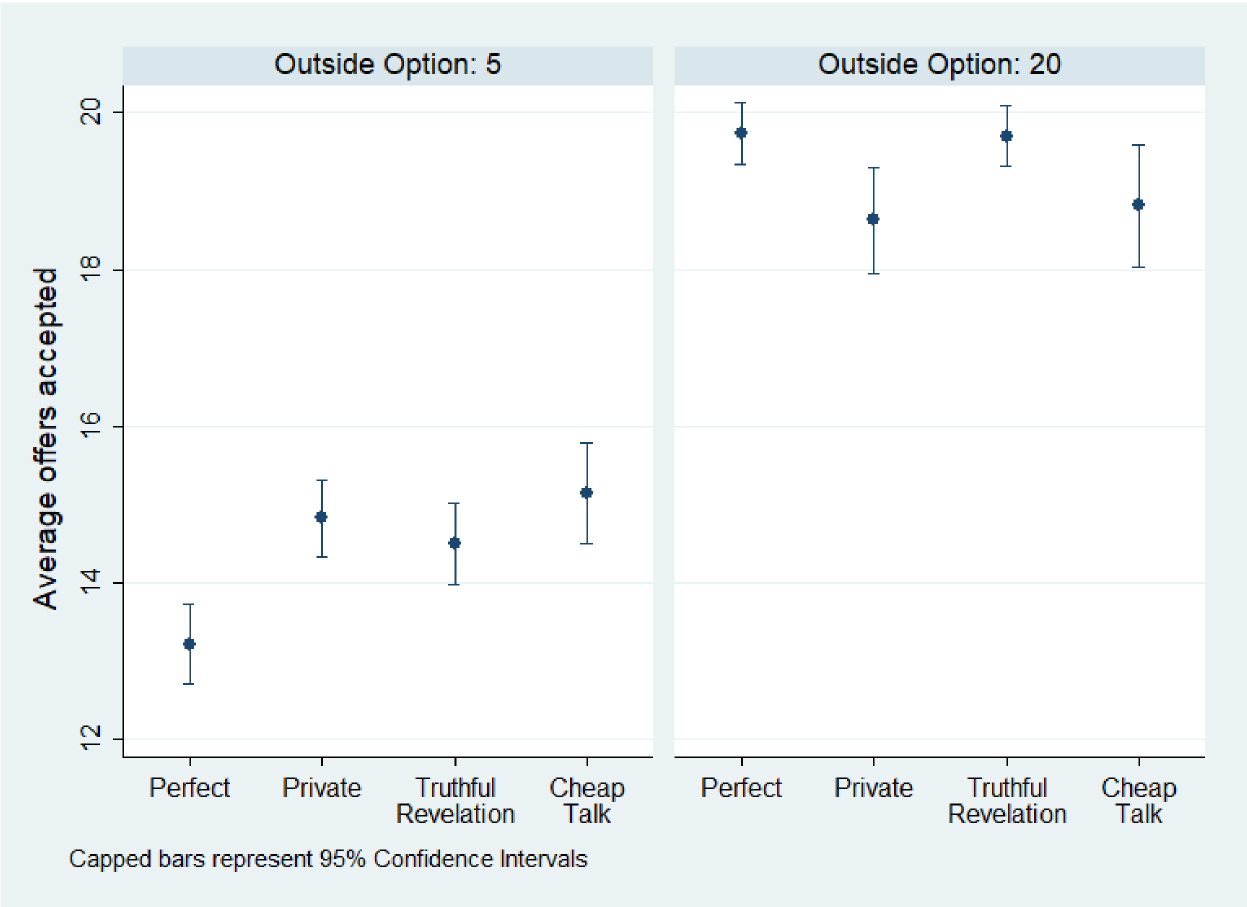

Figure 1.2: Average accepted offers across treatments, by B's Outside Options

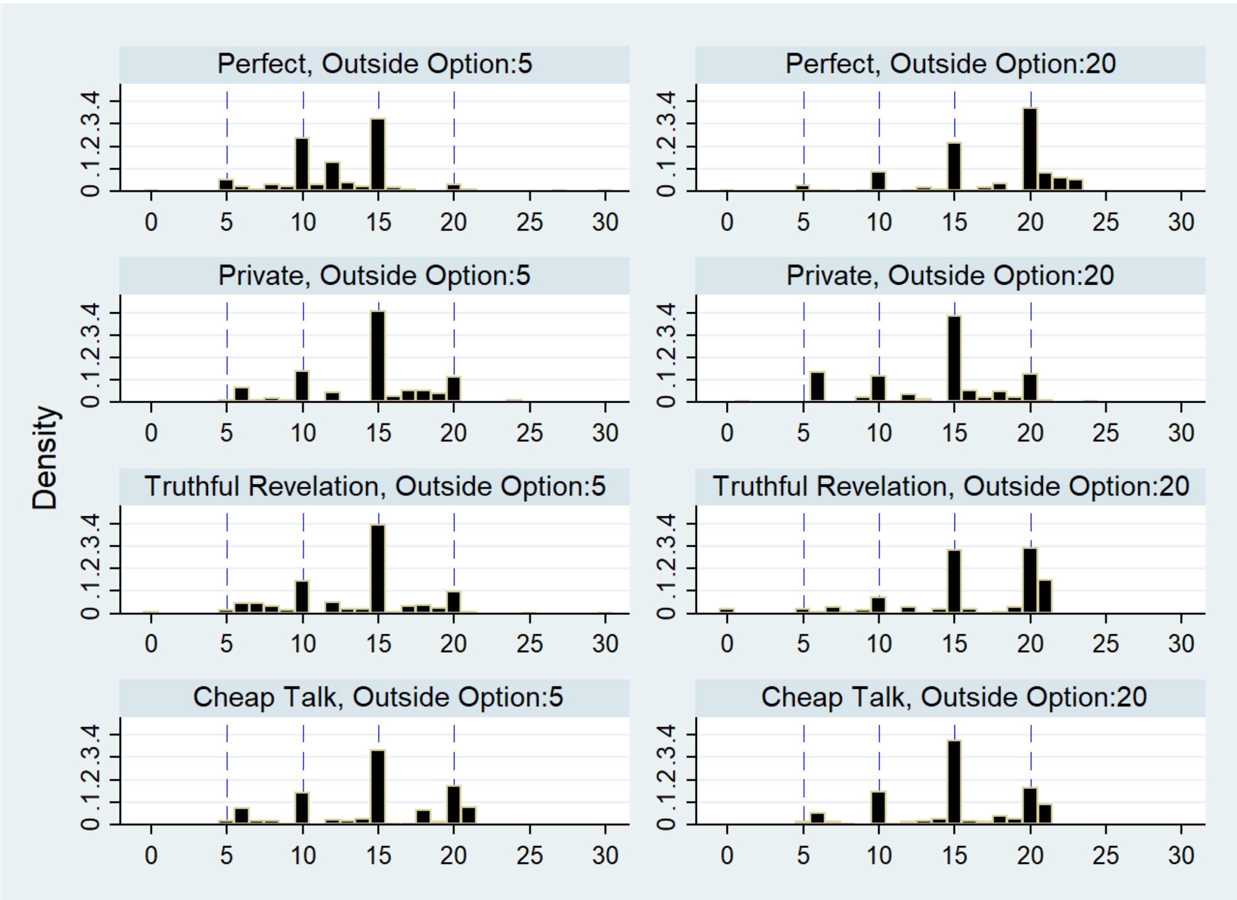

Figure 1.3: Offers Histogram across treatments, by B's Outside Options

Figure 1.4: Percent of Subjects Revealing Outside Option: Truthful-Revelation (PTR) Treatment

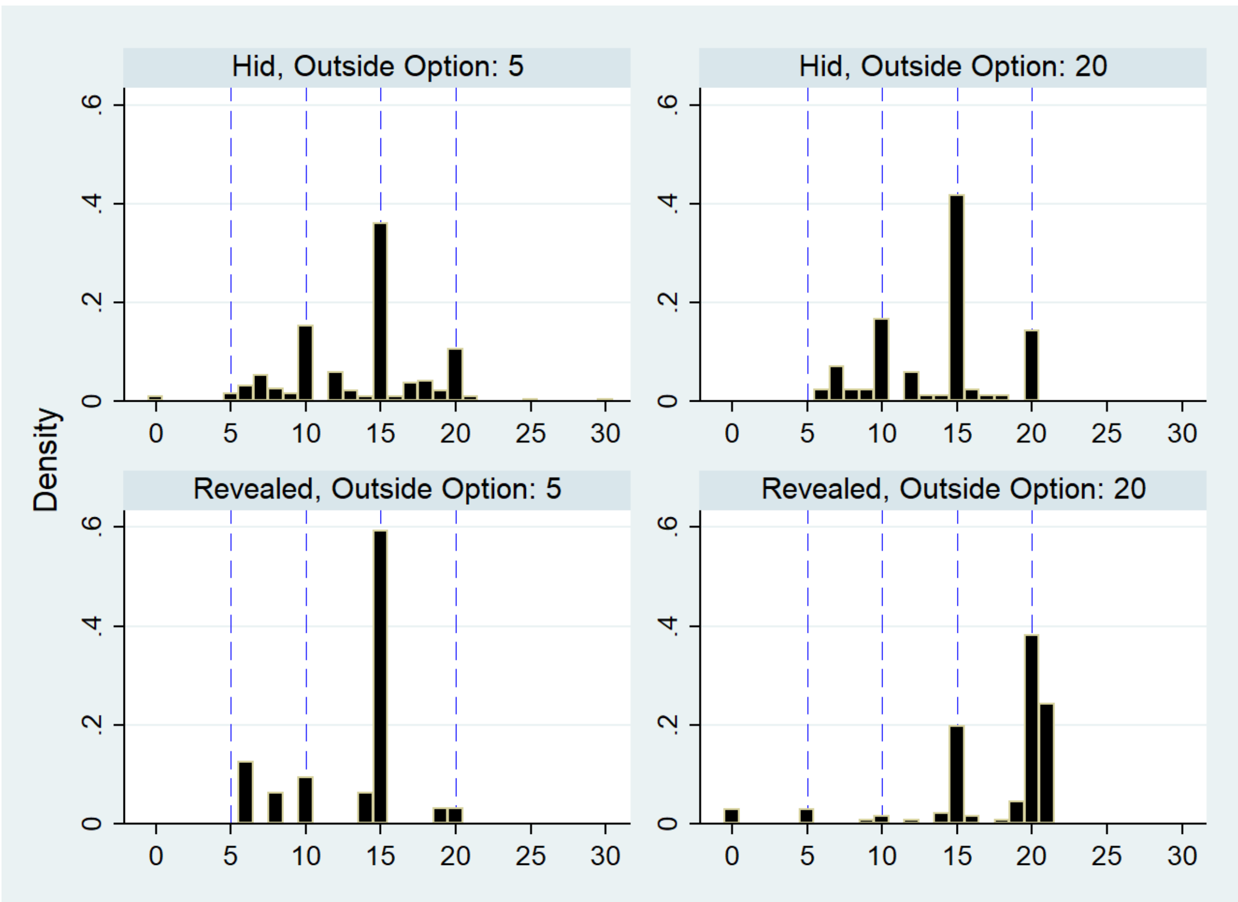

Figure 1.5: Truthful Revelation: Offer Histograms by outside option and revelation decision

Chapter 2

Many Channels of Adjustment to a Higher Minimum Wage: Evidence from Restaurant Reviews

2.1 Introduction

Although nearly all research on firms' responses to higher minimum wages focuses on changes in employment, firms employing low wage workers can make adjustments along several margins. In this paper, I use review data for restaurants and employ natural language processing methods to simultaneously examine many margins empirically, including some that are difficult to detect with conventional data and methods. I also examine the net impact of

I thank David Neumark, Matthew Freedman, John Duffy, Michael McBride, Marion Aouad, Arthi Vellore, Emily Owens, Damon Clark, Amihai Glazer, participants of the UC-Irvine Labor-Public seminar, and the UC-Irvine Applied Microeconomics Workshop for valuable feedback on this paper.

minimum wage changes on restaurant quality as proxied by the ratings associated with these reviews to understand the overall effect of the adjustments restaurants make on the "value" of their product.

The debate on the employment effects of minimum wage changes has led to a vast literature (see Dube, 2019; Neumark and Shirley, forthcoming, for recent discussions). Schmitt (2015) reviews the employment debate and focuses on other adjustment channels that may reconcile some of the differences in the literature (also see Manning (2021)). He considers alternative channels that could arise from assuming a competitive model, an institutional model, or a dynamic monsopsony framework, reviewing other research work that studies each of these in greater detail. Apart from a reduction in employment, competitive models suggest that higher minimum wages can lead to higher prices, reductions in fringe benefits, or reductions in worker training, though the empirical evidence on this is mixed (Belman and Wolfson, 2014). Institutional models do not assume that firms are operating at peak efficiency and allow for other adjustments, such as greater managerial efforts, reliance on productivity enhancing activities, the efficiency wage effect, and the mitigating effects that result from the improved spending power of low-wage workers. Another suggested channel is an upgrade in skill level of workforce in response to a higher minimum wage (Clemens et al., 2021). The dynamic monopsony model suggests an increase in employment and reduced turnover as possible consequences of higher minimum wages.

Most literature on the effects of higher minimum wages on employment focuses exclusively on this margin, and indirectly appeals to some of the above alternative adjustment channels as potential explanations for small or null effects. Among the recent papers that focus on other channels, many only examine a single adjustment mechanism. Thus, while there is now considerable collective evidence on alternative channels like prices, there is very little on others, since these may be difficult to study using typical labor market or firm level data, especially simultaneously.

The minimum wage literature often focuses on restaurants or food services because they rely heavily on minimum wage workers and employ a significant share of these workers. Restaurants also have fairly complex production technology with several potential margins to adjust, making this an interesting sector to study multiple margins of adjustment. Advances in text analysis methods mean that reviews of restaurants may be an informative source to detect responses to higher minimum wages. In this paper, I use text analysis and natural language processing methods on reviews to provide further evidence on the adjustments that can be reliably detected through reviews. Specifically, I examine impacts on prices, friendliness, hygiene, portion sizes, and wait times.

Another contribution of this paper is to study the impact of higher minimum wages on the overall customer experience as proxied by ratings. [2] Studying various adjustments together with ratings helps us better understand the broad implications of minimum wage changes, and the mechanisms underlying the net impacts on this useful summary measure. Since the different adjustments restaurants could make in response to minimum wage increases may have countervailing effects on ratings for service-oriented establishments, the overall effect is theoretically ambiguous, but empirically testable. For instance, if employment falls, the retained employees may have to handle larger workloads and deal with additional tasks where they are less productive. Alternatively, a higher wage for workers, *ceteris paribus*, may motivate them to work harder and exert more effort which would reflect in the actions of customer facing personnel in service-oriented establishments. This efficiency wage effect may lead to higher ratings.[3] Additionally, employers could cut costs which may be reflected in poorer quality of food or service, or lower hygiene standards. Higher menu prices could

[2]Ratings may also be an important outcome to examine because recent research suggests that they can impact revenues and reservations of restaurants. For instance, Luca (2011) finds that a one star increase in restaurant ratings increase revenues by 5 to 9 percent, an effect driven mainly by independent restaurants. Anderson and Magruder (2012) find that a half-star increase in Yelp ratings causes restaurants to fill reservations more frequently.

[3]Reich et al. (2005) find that in response to living wage policies implemented at the San Francisco Airport in 1999, employers surveyed were more likely to report an improvement in overall work performance, morale, absenteeism, and disciplinary issues.

lower ratings holding everything else constant.

I use data from over 16 million reviews of about 370,000 establishments across 18 metropolitan areas (MSAs) in the United States and Canada. Combining natural language processing methods with difference-in-difference techniques for identification, I find evidence of a rise in prices and improvements in staff friendliness with higher wages. I find weaker evidence for deteriorating hygiene, and no significant effects for portion sizes or wait times. An estimate of price elasticity recovered from review texts of 0.05 is in line with recent literature, and lends support to the competitive model. However, improvements in staff friendliness may indicate efficiency wage effects or changes by managers to improve operational efficiency that is more in line with institutional models. The findings on improved friendliness and null impacts on wait times are corroborated by interviews with managers surveyed in Hirsch et al. (2015), where they indicated the importance of maintaining speedy customer service and improving performance standards in response to higher wages. However, managers also indicated cross-training of employees as a cost saving measure, which may be consistent with the evidence of deteriorating standards of hygiene found here, and elsewhere in the literature (Chakrabarti et al., 2020). Finally, I find evidence that higher minimum wages are associated with a drop in ratings of restaurants and other food-service establishments.[4]

Aside from these specific conclusions, the methods I develop and implement demonstrate that unstructured natural language found in review data analyzed with computational linguistics methods may be informative in examining the response of firms to minimum wage changes as well as other policies. Gentzkow et al. (2019) provide an excellent introduction on the potential of text as data for use in economic research. In their review of applications and methods they conclude that "virtually all of the methods applied to date, including those we would label as sophisticated or on the frontier, are based on fitting predictive models to simple counts of text features. Richer representations, such as word embeddings, and lin-

[4] As a useful placebo check, I find no impact on ratings of other (non-restaurant) establishments that are much less reliant on minimum wage workers.

guistic models that draw on natural language processing tools have seen tremendous success elsewhere, and we see great potential for their application in economics." (pp. 569). The text analysis in my paper relies on one such word embedding technique - Word2Vec (Mikolov et al., 2013), to better capture the meaning of words and the relationships between them. Words in the text of reviews and the adjustment mechanisms are represented as vectors in a space where similar words (and their combinations) are co-located. [5] The outcomes that capture the adjustments studied in this paper are based on "similarity scores" that rely on models which account for the syntactical richness of natural language used in reviews. To verify that these outcomes capture the concepts meaningfully, I use validation techniques detailed in Section 4.2, making use of the ratings associated with the reviews, or the "$" sign (a proxy for price range) associated with the business.

The rest of this paper is organized as follows. Section 2 provides a brief overview of recent work on the impact of minimum wages on non-employment margins, especially those that use digital data. Section 3 describes the data and provides some descriptive statistics. Section 4 outlines the identification strategy as well as the construction and validation of the text-based outcomes. Section 5 presents the results, and Section 6 concludes by placing the findings in the context of the channels of adjustments literature.

2.2 Brief Review of Related Literature

While there is limited evidence on some of the adjustments studied here, this paper also supplements findings from recent work on others. [6] The one margin aside from employment on which there is fairly extensive evidence is prices, where some degree of pass-through is documented. Aaronson et al. (2008) use panel of store-level restaurant prices as well as

[5] A recent application of these methods can be found in Burn et al. (2019) who study the relationship between ageist language in job ads and hiring discrimination.

[6] Clemens (2021) focuses on the relevance of non-employment margins, and is an excellent review of the recent work that explores these margins.

aggregated Consumer Price Index data to document price increases in response to federal minimum wage hikes. Allegretto and Reich (2018) find that almost all of the minimum wage increase was passed on through higher prices when San Jose increased its minimum wage in 2013. Using firm level price index data from Hungary, Harasztosi and Lindner (2019) find that three-fourths of a large minimum wage increase was passed on to consumers.

Turning to other margins that are harder to measure, work by Chakrabarti et al. (2020) suggests that minimum wage hikes are associated with lower hygiene levels in restaurants due to a possible increase in the task demands on the retained staff. Using a difference-in-difference framework with Seattle restaurants as treated and those in Bellevue (in the same county) as controls, they find that a \$0.25 increase in real minimum wage is associated with a 8% increase in total health violation scores.

Recent work has started to rely on digital data in creative ways to further our understanding of the impacts of minimum wage increases. Luca and Luca (2019) use Yelp data for the San Francisco Bay area and conclude that restaurants with lower ratings are closer to the margin of exit, and are disproportionately driven out of business by increases in the minimum wage. Direct evidence on the impact of minimum wages on ratings is rare. A working paper by Crain (2018) uses a scraped dataset of online menu items and restaurant quality for restaurants across three states (New York, Massachusetts and New Jersey) to estimate the price and quality responses to minimum wage increases in 2017. The author finds higher pass-through in prices among smaller firms, and significant changes in restaurant quality, with quality declining among low quality restaurants and increasing for high quality ones. The author relies mainly on Grubhub, with a secondary dataset from Yelp. Using the former dataset to examine the quality results further, the author finds that changes in quality are driven by changes in order accuracy and timeliness of deliveries as opposed to food quality. These ratings could reflect different adjustments than those captured in this paper using text reviews largely based on customer experience in person. The other margins captured in the text analysis in my paper also help in the interpretation of the ratings responses found here and elsewhere in the literature. The data used in this paper allow me to exploit several minimum wage changes since the establishment of Yelp.com (2004) across 18 metropolitan areas in the United States and Canada. Other non-restaurant establishments less affected by minimum wage changes serve as a useful placebo check. Moreover, since reviews are linked to businesses and users using identifiers, I use both business fixed effects and business-user pair fixed effects to assess impact on ratings, which can help address concerns of compositional changes driving the results. The large sample of review texts also enables a text analysis with greater precision, and is critical for studying the multiple margins of response.

2.3 Data and Descriptive Statistics

I use datasets released by Yelp for use by researchers. [7] The data released in 2020 and 2021 together contain information on nearly 370,000 businesses across eighteen major metropolitan areas in two countries (United States and Canada) with information from nearly 17 million reviews written by over 3.9 million users. The data include businesses that have had at least 3 reviews older than 14 days and only includes the Yelp recommended reviews. [8] The eighteen metropolitan areas include fourteen in the United States: Las Vegas (NV), Pittsburgh (PA), Phoenix (AZ), Charlotte (NC and SC), Cleveland (OH), Madison (WI), Champaign (IL), Boulder (CO), Orlando (FL), Atlanta (GA), Boston (MA), Columbus (OH), Portland (OR), Austin (TX) and four in Canada, namely, Toronto (ON), Calgary (AB), Montreal (QC), and Vancouver (BC).

[7] This data is made available under the annual Yelp data challenge and has been used in several research papers. The data are updated for each iteration of the challenge.

[8] Yelp recommended reviews are about 75% of the total reviews and excludes reviews that their algorithm detects as biased, fake or unreliable based on user characteristics (`https://www.yelp-support.com/article/Why-would-a-review-not-be-recommended?l=en_US`). Luca and Zervas (2016) provide evidence that this alogirthm does a good job in identifying fake reviews by validating these with sting operations conducted by Yelp starting in 2012.

Aside from average ratings, establishment level information includes hours of operation, street address, geographic coordinates, number of reviews and various attributes like type of business, cuisine for restaurants, parking availability and others. The review level data has individual review texts, associated ratings from 1 to 5, the date of the review, and individual dummies for whether the review was useful, funny or cool. Finally, the user-level data has the user id, first name, number of reviews, the average rating they give, date of first review, their friend lists, number of fans, number of useful, funny or cool votes they received and so on.

I link this review data with both the users and businesses using identifiers. Integrated information provided in this format by Yelp on businesses, reviews and users has advantages over scraped data which is typically unable to link all this information. For instance, this allows me to add not just business fixed effects, but also user fixed effects, or business-user pair fixed effects for more credible identification that accounts for compositional changes. Essentially, the identification here relies on the same user reviewing the same restaurant at different minimum wage levels.

The critical component of the data from the point of view of this study is the exact time and date on which the review was written. This means each review can be associated with with a particular minimum wage level in that state at that time. Since the data contain the entire set of reviews for the included businesses, the reviews begin as early as October 2004 (when Yelp was founded) until December 2019. [9] This means that several minimum wage changes across these eighteen metro areas can be exploited. The primary sources for the minimum wage data were the Department of Labor [10] for the fourteen U.S. metro areas and the Government of Canada [11] for the four metro areas in Canada. Tables 2.1 and 2.2 provides details of the minimum wage levels over the relevant time period (2004-2019) for

[9] The data for 2020 is only available for 8 metropolitan areas, and is dropped from the analysis to restrict attention to the period unaffected by the COVID-19 pandemic.

[10] `https://www.dol.gov/whd/state/stateMinWageHis.htm`

[11] `http://srv116.services.gc.ca/dimt-wid/sm-mw/rpt2.aspx`

these cities. While the tables list a single minimum wage for a given year for simplicity, in the analysis I exploit the exact month that minimum wage changes came into effect.

Table 2.3 provides some descriptive statistics for the data. The final data used in this analysis drops less than 0.03 percent of observations from the raw data, excluding only those that did not belong to the 18 metropolitan areas, because of data entry errors and observations where the business, user and review files did not merge perfectly. The average number of reviews for a business is 43, and 79% of businesses in the data are classified as open. Using the attributes of businesses, I classified 149,244 businesses as restaurants operating in the food-service industry more generally. While restaurants make up about 40% of the total businesses, they account for over 71% of all reviews written. [12] The second panel of the table provides information user characteristics. On average, each user has written about 18 reviews in total.

2.4 Methods

2.4.1 Identification Strategy

I use a difference-in-difference (DID) strategy to estimate the causal effects of minimum wage hikes. To keep the identification strategy transparent, the main analysis focuses on the time period 2011-2019 and a single large minimum wage hike in Phoenix, Arizona in 2017. The Arizona minimum wage was increased by nearly 25%, from \$8.05 to \$10.00 per hour. This analysis also relies on using the seven metro areas that serve as "pure" controls or a "never-treated" group. This mitigates some of the issues generated by staggered roll out

[12] The most popular non-food related categories of businesses in the data are Shopping, Beauty & Spas, Home Services, Health & Medical, Automotive, Local Services, Active Life, Hair Salons, Auto Repair, Nail Salons, Event Planning & Services, Hotels & Travel, Fashion, Real Estate, Doctors, Professional Services, Pets, Arts & Entertainment, Home & Garden, Financial Services, Fitness & Instruction, Dentists and so on.

designs (Callaway and Sant'Anna, 2020; Sun and Abraham, 2020). These are Las Vegas, Champaign, Charlotte, Pittsburgh, Madison, Austin and Atlanta. Figure 2.1 shows the (annual) average weekly wage in the food service industry for these metropolitan statistical areas over this period using the Quarterly Census of Employment and Wages (QCEW) data. The figure clearly demonstrates the impact of this minimum wage change on average wages of workers in this industry in Phoenix. [13]

The baseline specification for the DID analysis is the following:

$$Y_{rjst} = \beta \times Treat_s \times Post_t + \mu_j + \eta_t + \epsilon_{rjst} \tag{2.1}$$

In equation (2.1), Y_{rjst} represents an outcome associated with review r for business j in state s at time t. Outcomes can be "star" ratings between 1 to 5, or can represent the closeness of the review text to an adjustment mechanism defined in the section below. μ_j and η_t represent business and time fixed effects. Business fixed effects control for time invariant business characteristics, and year fixed effects control for factors changing each year that are common to all businesses. I extend this difference-in-differences model to an event study design in which I estimate treatment effects by year:

$$Y_{rjst} = \sum_{\substack{t=2011 \\ t\neq 2016}}^{2019} \beta_t \times (Treat_s \times D_t) + \mu_j + \eta_t + \epsilon_{rjst} \tag{2.2}$$

The event study design not only allows us to trace out the time pattern of the minimum wage impacts on Phoenix establishments after 2017, but also helps us assess the validity of the parallel trends assumption by estimating differences between outcomes in each year

[13] Arizona's minimum wage for tipped workers also increased commensurately from \$5.05 to \$7 an hour in 2017. Even and Macpherson (2014) study the impacts of changes in the tipped minimum on full service restaurants and find reductions in employment in the full-service restaurant industry and tipped workers. In Appendix Figure B.1, I provide average weekly wages from the QCEW data for workers for workers in Arizona and the control states for both limited service and full service restaurants. Wages appeared to increase for workers in both categories (QCEW data includes tips). In this paper, I focus on all restaurants in the review data, and do not make a distinction between the two.

prior to minimum wage hike (relative to 2016, the omitted year). I also supplement this DID analysis (including event study designs) with a synthetic control approach.

In supplementary analysis, I use reviews from all 18 metropolitan areas from 2004-2019 using two-way fixed effects panel regressions with a continuous minimum wage variable to estimate equations of the form:

$$Y_{rjst} = \beta \times MW_{st} + \mu_j + \eta_t + \epsilon_{rjst} \tag{2.3}$$

2.4.2 Constructing and validating the text-based outcomes

Here, I describe the procedure I use to detect restaurant's adjustment mechanisms from the language of review texts. Figure 2.2 shows the empirical distribution of the number of words in each of the 11.3 million reviews for restaurants in the entire data. The mean number of words is 105, while the median is 75. The first step in text processing was to remove the frequently used stop words (such as "of", "a", "the", "and", "to", "in" ,"for"). I then used a word2Vec model where each word is represented as a vector in a vector space with a dimensionality of 300. This model was first trained on the Google News corpus. Using this model, we can calculate the semantic similarity scores between words, such that words that appear more often together in similar contexts have higher similarity scores. Cosine similarity scores between two vectors representations can range between values from -1 to 1, where scores closer to 1 indicate a high similarity and scores closer to -1 implies that the words are unrelated (unlikely to appear in similar contexts). Based on this model, I calculate the cosine similarity score between each word in a review and a word (or phrase) that represents a particular adjustment mechanism of interest. As an example, the notion of how expensive a restaurant is can be expressed in various ways. Words with high cosine similarity scores with the word "expensive" include impractical, fancier, onerous, uneconomical, pricy, pricey and so on. Using these scores helps account for the various ways in which language can be used to express a particular concept, as opposed to searching for particular words. The choice of phrases that represent the adjustment mechanisms of interest is informed by looking at token or tri-gram (three-word phrase) frequencies of how these are typically referred to in review language. The list of words and phrases that represent each mechanism is as follows:

- Prices: "Expensive", "Cheap"
- Efficiency Wage: "Friendly", "Rude"
- Hygiene: "Clean", "Dirty"
- Portion Sizes: "Large" + "Portion", "Small" + "Portion" [14]
- Wait Times: "Wait" + "Long" + "Time", "Wait" + "Short" + "Time"

First, I calculate the cosine similarity scores for each token (or word) in a given review and each of the vector representations in the list above. Therefore, for each review, I obtain a distribution of scores. For example, if a review consists of 100 words (after removing stop words) I calculate the cosine similarity score between each of these words and the word (vector) "expensive". The next step is to recover the 95^{th} percentile of cosine similarity score as the building block of the outcome we will examine. I look at higher percentiles within a review in order to capture the words (and associated scores) in the review that are more likely to be related to the adjustment mechanism, on average.

To validate that these cosine similarity scores are informative, as a first example, I provide a figure that relates the 95^{th} percentile of the cosine similarity score from the distribution of words within a review and the word "friendly". Since each review is also associated with a

[14] Note that the vector space in which words are represented allows for an internally consistent arithmetic. This means that mathematical operations can be done in a way that preserves the meanings of words. For instance, the vector that results from "king" - "man" +"woman" will be close to "queen" in this vector space.

customer rating from one-star up to five-stars, we can plot the distribution *across* reviews (of the 95^{th} percentile score *within* each review) for each of these rating levels. This is shown in the top panel of Figure 2.3. It is evident that reviews with five stars had the highest scores (right-most), and those with progressively lower ratings are shifted to the left, in the expected order. This is also reflected in the vertical lines, that show the mean values for these distributions. In the bottom panel of the figure, I show the analogous figure with similarities with the word "rude". As expected, the order of the distributions is now reversed, with the one-star review distribution with the greatest semantic similarity scores, and the five-star distribution with the least similarity.

It is important to note that lower values indicate of cosine similarity implies that the terms are unrelated but not opposites. To create a single, continuous summary measure that captures the closeness of a review text to an adjustment mechanism such that low values are related to one end of the mechanism (lower prices, rudeness, good hygiene) and high values indicate the other (higher prices, friendliness, poor hygiene), I use the (standardized) distance between these scores for each review. Thus, the outcome for each review r is defined for the six adjustment mechanisms as follows:

$$Y_r^{Price} = \Theta(CS(\text{``}Expensive\text{''}, Token_r)) - \Theta(CS(\text{``}Cheap\text{''}, Token_r))$$

$$Y_r^{Efficiency} = \Theta(CS(\text{``}Friendly\text{''}, Token_r)) - \Theta(CS(\text{``}Rude\text{''}, Token_r))$$

$$Y_r^{Hygiene} = \Theta(CS(\text{``}Dirty\text{''}, Token_r)) - \Theta(CS(\text{``}Clean\text{''}, Token_r))$$

$$Y_r^{Portion} = \Theta(CS(\text{``}Small\ portion\text{''}, Token_r)) - \Theta(CS(\text{``}Large\ portion\text{''}, Token_r))$$

$$Y_r^{Waiting} = \Theta(CS(\text{``}Wait\ long\ time\text{''}, Token_r)) - \Theta(CS(\text{``}wait\ short\ time\text{''}, Token_r))$$

, where Θ represents the 95^{th} percentile of the distribution of CS scores *within* a review. In Figure 2.4, I present the final (standardized) outcome that captures efficiency wage effects (proxied by friendliness or rudeness of staff) and its relationship to review ratings. Figure 2.5 shows hygiene and wait time outcome histograms (validated with review stars), and price adjustments (validated with Yelp $ signs). As further validation of the information contained in these outcomes, I provide example review texts for a random sample of 100 reviews drawn from within the first percentile and above the ninety-ninth percentile for each of the outcome variables in an online appendix. [15]

Finally, even though we have picked the 95^{th} percentile value of the cosine similarity scores *within* a review to construct an outcome, most reviews will not be informative about many adjustment mechanisms especially considering average review length. Therefore, instead of only looking at mean impacts using ordinary least squares, I also consider impacts on higher percentiles *across* reviews by employing quantile regressions. I calculate the difference-in-difference estimates for 95^{th}, 90^{th}, and 85^{th} percentiles of the outcome distribution by estimating a modified version of equation (2.1) where $Post_t = 1$ for 2017 and $Post_t = 0$ for 2016.

2.5 Results

2.5.1 Ratings

Before turning to the text of reviews to test for evidence of various adjustment mechanisms, I first examine the impact of minimum wages on ratings. Ratings here serve as a proxy or summary measure for overall quality of service. While some potential adjustments such as higher prices or lower hygiene could impact ratings negatively others such as improvements in operational efficiency or efficiency wage effects could lead to higher ratings.

[15] Link to Online Appendix.

I focus on the impact of the minimum wage hike in Arizona that came into effect in 2017 on ratings. The 2017 Arizona minimum wage rise was one of the largest in our data for the United States from $8.05 to $10.00. Since this was a recent minimum wage hike in the Phoenix metropolitan area, there is a considerable amount of data in terms of number of establishments and associated reviews. The comparison states of Nevada, Illinois, North and South Carolina, Pennsylvania, Wisconsin, Georgia and Texas had constant minimum wages throughout this period. Figure 2.6 shows the annual average ratings using reviews from restaurants in the top panel. First, it is evident from the top panel that since 2012 average restaurant ratings were on the rise. Second, while Phoenix restaurants had higher ratings throughout, trends looked quite similar to the comparison areas until the minimum wage hike in 2017. Ratings in Phoenix declined in 2017, and fell below the average ratings of restaurants in the control group. The latter continued to rise in line with previous trends.[16]

To estimate the impact on ratings formally, I use a difference-in-difference (DID) framework in Equation (2.1), where $Treat = 1$ for Phoenix, Arizona and 0 for control areas. $Post$ takes a value of 0 for 2011 to 2016 and 1 for 2017 to 2019. The DID coefficient of interest is β. All specifications include year and month fixed effects to account for overall time trends affecting all establishments (including seasonal variation). Table 2.4 shows that the DID coefficient is negative and significant. To account for possible compositional changes in the restaurants being reviewed, columns 2 includes business fixed effects to effectively exploit the within-establishment variation. Results are similar with this restriction. Column 3 further adds business-user pair fixed effects.[17] The nearly two dollar increase in minimum wage in Arizona was associated with a 0.03 to 0.04 “stars” reduction in average ratings for restaurants, relative to the control states. [18] The results are robust to only including a single

[16]In contrast to this, Appendix Figure B.2 shows that trends continued to be similar through 2017 for treated and control states for all other (non-food) establishments. This serves as a useful placebo check.

[17]This restriction means identification relies on a very small subset of reviews which makes the estimate somewhat imprecise (p=0.11)

[18]The results for placebo establishments corresponding to the specifications in Table 2.4 are presented in Appendix Table B.1. For these (non-food service) establishments, the coefficients are not significant and close to zero.

pre (2016) and post (2017) year, as well as using only a single neighboring state (Nevada) as a control.

Event study estimates for restaurants are presented in the bottom panel of Figure 2.6, where we see that the parallel trends appear to be satisfied, even accounting for alternative levels of clustering standard errors. As a robustness check, I also employ a synthetic control approach (Abadie et al., 2010) that matches on pre-trends of average annual ratings for restaurants since 2011. Figure 2.7 presents the results for restaurants. This figure is very similar to the top panel of Figure 2.6, and encouragingly I find that the method assigns the largest weights (0.488 and 0.123) to the two most geographically proximate and demographically similar control regions of Austin, Texas and Las Vegas, Nevada, respectively. Wisconsin, Pennsylvania, North Carolina, Georgia, South Carolina and Illinois are all assigned weights less than 0.1 (in descending order).

Finally, in Table 2.5, I present results from estimating Equation (2.3) where the coefficient of interest represents the rating change associated with a one dollar change in minimum wage. This specification utilizes all the minimum wage changes for 18 metro areas since 2004. The table confirms that a higher minimum wage is associated with a significant drop in ratings for restaurants. Column 2 introduces Business-User pair fixed effects and reinforces this result, even though identification relies on a much smaller subset of reviews.

2.5.2 Restaurant Prices

I use two approaches to examine price adjustments to minimum wages using restaurant review texts. I do this using two approaches. The first searches the text of reviews for an explicit mention of numeric dollar prices and recovers numbers following the "$" symbol. The advantage of this numeric approach is that I can estimate the elasticity of prices faced by the consumers with respect to minimum wages which is readily interpretable and can be placed in the context of the existing literature. An important caveat here is that identification relies on the reviews that explicitly mentioned a dollar amount. This restricts the sample to about one million (about 10 percent) of reviews. The second approach uses the outcome based on the semantic similarity of the review text to the word "expensive" relative to the word "cheap". This approach was outlined in detail in section 2.4.2 above for all the adjustment mechanisms I explore in this paper. This has the benefit of using all reviews in the data.

To check whether the extracted prices are sensible and informative, Figure 2.8 shows how the average extracted dollar prices from reviews correspond to Yelp's "$" signs for the restaurant that was reviewed. These average prices are about $9,$17, $31, and $41 for the four categories in increasing order of expensiveness. The event study figure for extracted (log) prices is shown in Figure 2.9. The Arizona minimum wage hike appeared to raise prices for Phoenix restaurants relative to the controls.

Table 2.6 presents the results from a panel regression of log prices on the log of minimum wages. This relies on the specification of equation (2.3) (Table 2.5), where we use all the data since 2004. [19] For the specification with business fixed effects (column 1), the minimum wage price elasticity is about 0.05 and statistically significant. This implies that a 10 % increase in minimum wage leads to a 0.5 % rise in the prices, as experienced and mentioned by consumers. This is in line with other estimates of price elasticity in the literature (Allegreto and Reich , 2018). Adding business-user pair fixed effects (column 3) limits the variation in the sample further to repeat reviewers for restaurants that mentioned prices explicitly. I find that the estimates are still positive though no longer significant. [20]

Next, we can explore price adjustments using the language based outcome. As mentioned at the end of Section 2.4.2 above, I present the DID estimates for the language based outcomes

[19] Table 2.6 is limited to the 14 U.S. metro areas. I exclude Canada from the analysis since the extracted prices rely on the US $ symbol, which is often not used in Canadian reviews.

[20] Appendix Table B.3 shows that for other (placebo) establishments the elasticity is not significant in either specification.

at the higher quantiles of the distribution in column 1 of Table 2.7. Looking at higher percentiles of the outcome focuses on the reviews most likely to mention language highly correlated with price adjustments, I find some evidence of a higher propensity for language associated with "expensive" (relative to "cheap") for Phoenix, though these are imprecisely estimated. Overall, looking at the evidence from both of these approaches, it does appear that there is an increase in prices in Phoenix restaurants after the minimum wage hike relative to the control MSAs.

2.5.3 Friendliness Outcome

In this section, I examine whether higher minimum wages are associated with improvements in staff friendliness. Once again, I present the DID estimates for these language based outcomes at the higher quantiles of the distribution. For staff friendliness (relative to rudeness), these results are presented in column 2 of Table 2.7. I find evidence of a significant improvement in this outcome at higher percentiles in Phoenix after the minimum wage rise relative to the control regions. This is especially true of reviews at the 90^{th} and 95^{th} percentile of this outcome, which are likely the most informative in capturing the adjustment mechanism.

2.5.4 Hygiene

To the study the association between minimum wages and overall levels of hygiene, we consider the text based outcome that takes on higher values when reviews are associated more closely with worse hygiene, based on higher semantic similarities with the word "dirty" relative to "clean". Column 3 of Table 2.7 shows the difference-in-difference estimate for this outcome at the 95^{th}, 90^{th} and 85^{th} percentile. The 95^{th} percentile has the largest estimate, which is significant at the 10% level. This would imply that the most informative reviews about this adjustment mechanism experienced a larger increase in Phoenix after the minimum wage rise relative to the control areas, indicating worsening hygiene levels.

2.5.5 Wait Times and Portion Sizes

Analogous to the case of price adjustments, I approach wait times in two different ways. The first is to extract wait times from text using any numeric mention of wait times preceding the word “minutes”. [21] The event-study presented in Figure 2.10 does not provide strong evidence of lack of clear pre-trends, or a significant change in wait times after the minimum wage change. The language based outcome on wait times constructed as a standardized measure based on mentions of long relative to short wait times also does not appear to show clear effects. This is true even for the most informative reviews at higher percentiles of the outcome (column 4 of Table 2.7). The estimates for these standardized outcomes are close to zero. From the point of view of customers, therefore, there appears to be no detectable changes in wait times.

The last column of Table 2.7 presents the results for the change in the outcome on portion sizes at the top three vigintiles of the outcome distribution. Recall that higher values of this outcome are associated with smaller portion sizes. I find no clear evidence that restaurants in Arizona changed their portion sizes significantly in response to the minimum wage change, relative to restaurants in the control areas. Once again, these estimates are close to zero for all the top three viginitiles (95^{th}, 90^{th} and 85^{th} percentiles).

[21] This extraction can be validated by considering the average wait times across Yelp $ signs. These results are presented in appendix Figure B.3. I find that wait times are significantly shorter for “$” sign restaurants relative to “$$” restaurants, which in turn have shorter wait times than the most expensive “$$$” and “$$$$” restaurants.

2.6 Conclusion

Restaurants and other food service establishments can adjust to changes in minimum wages in many different ways. Combining natural language processing with differences-in-differences methods for identification, this paper examines the impact of minimum wage changes on prices, staff friendliness, hygiene, wait times, and portion sizes. I find evidence of improvements in staff friendliness, higher prices, and deteriorating hygiene. Extracting prices from text of reviews indicates that a 10 percent increase in minimum wage increases prices by 0.5 percent. I find no discernible impacts on portion sizes or wait times.

It is important to place these findings in the context of the small but growing literature on alternative channels of adjustment, which typically relies on very different methods and data. To further explore the small effects of a federal minimum wage hikes between 2007-2009 on employment and hours worked, Hirsch et al. (2015) surveyed of 66 store managers to examine alternative channels of adjustments. In response to how they would adjust to minimum wage changes, managers cited the importance of improving performance standards, hiring more experienced workers, and maintaining speedy customer service. Since enhanced operational efficiency requires more training and not less, very few managers actually reported decrease in training as a cost saving response. My findings on an improvement in staff friendliness are in line with the "institutional" model of the labor market. It is conceivable that the improved friendliness reflects increased performance standards. This finding is also consistent with the efficiency wage theory that predicts greater work effort by workers either by increasing the cost of losing a job, or through the reciprocity effect of the "gift" of higher wages. Aside from discouraging over-time work, interviews with managers in Hirsch et al. (2015) also revealed that cross training workers is an important tool to improve operating efficiency. My finding of a decline in cleanliness may reflect this, and is also consistent with Chakrabarti et al. (2020), who cite increased task demands on retained staff as a possible explanation. The fact that I do not find any impacts on wait-times is also consistent with managerial priorities on maintaining the speed of customer service. Therefore, while I do find some evidence in support of institutional models overall, the increase in prices observed is also consistent with the competitive model of labor markets.

Finally, the net impact of the changes on ratings is negative, though small in magnitude. A dollar increase in minimum wage is associated with a fall in ratings of 0.01 to 0.02 off a base of approximately 3.8 out of 5 (about 0.5 %). These negative impact on ratings are robust to the inclusion of business fixed effects and business-user fixed effects, as well to clustering at the zipcode, state or business level. While these effects may seem small, they can have still have more serious implications for a large number of restaurants, especially because of the way summary ratings are displayed on the Yelp page of businesses. Recall that even though ratings by users can take values from 1 to 5, the main page for the restaurant summarizes the average ratings in half-star intervals. Since a 3.24 average would be rounded to 3, and a 3.26 would be rounded to 3.5, even a small shift in response can have implications for firms. Both Luca (2011) and Anderson and Magruder (2012) have used these rounding thresholds in RD designs, and found significant impacts of ratings on revenues and reservations. Thus, these costs may be important as ratings are being seen as increasingly valuable signals of restaurant quality by consumers. Along with rising prices, this may be another channel through which consumers "pay" for higher minimum wages. Because of the various responses available to firms, the incidence of higher minimum wages likely extends beyond employment effects. In addition to the impacts on consumers, examining these other margins simultaneously in the context of restaurants can offer additional insights into firm and worker behavior more generally.

Tables

Table 2.1: Minimum Wages for the U.S States in Data, 2004-2019 in USD

	AZ Phoenix	IL Champaign	NC Charlotte	SC	NV Las Vegas	OH Cleveland-Columbus	PA Pittsburgh	WI Madison	CO Boulder	FL Orlando	GA Atlanta	MA Boston	OR Portland	TX Austin
2004	5.15	5.50	5.15	5.15	5.15	5.15	5.15	5.15	5.15	5.15	5.15	6.75	7.05	5.15
2005	5.15	6.50	5.15	5.15	5.15	5.15	5.15	5.15	5.15	6.15	5.15	6.75	7.25	5.15
2006	5.15	6.50	5.15	5.15	5.15	5.15	5.15	5.70	5.15	6.40	5.15	6.75	7.50	5.15
2007	6.75	7.50	6.15	5.85	6.15	6.85	6.25	6.50	6.85	6.67	5.15	7.50	7.80	5.15
2008	6.90	7.75	6.15	6.55	6.33	7.00	7.15	6.50	1.02	6.79	5.85	8.00	7.95	5.85
2009	7.25	8.00	6.55	7.25	6.85	7.30	7.15	6.50	7.28	7.21	6.55	8.00	8.40	6.55
2010	7.25	8.25	7.25	7.25	7.55	7.30	7.25	7.25	7.24	7.25	7.25	8.00	8.50	7.25
2011	7.35	8.25	7.25	7.25	8.25	7.40	7.25	7.25	7.36	7.31	7.25	8.00	8.80	7.25
2012	7.65	8.25	7.25	7.25	8.25	7.70	7.25	7.25	7.64	7.67	7.25	8.00	8.95	7.25
2013	7.80	8.25	7.25	7.25	8.25	7.85	7.25	7.25	7.78	7.79	7.25	8.00	9.10	7.25
2014	7.90	8.25	7.25	7.25	8.25	7.95	7.25	7.25	8.00	7.93	7.25	8.00	9.25	7.25
2015	8.05	8.25	7.25	7.25	8.25	8.10	7.25	7.25	8.23	8.05	7.25	9.00	9.25	7.25
2016	8.05	8.25	7.25	7.25	8.25	8.10	7.25	7.25	8.31	8.05	7.25	10.00	9.75	7.25
2017	10.00	8.25	7.25	7.25	8.25	8.15	7.25	7.25	9.30	8.10	7.25	11.00	10.25	7.25
2018	10.50	8.25	7.25	7.25	8.25	8.30	7.25	7.25	10.20	8.25	7.25	11.00	10.75	7.25
2019	11.00	8.25	7.25	7.25	8.25	8.55	7.25	7.25	11.10	8.46	7.25	12.00	11.25	7.25

Table 2.2: Minimum Wages for the relevant Canadaian Provinces in the Data, 2004-2019 in CAD

	AB Calgary	ON Toronto	QC Montreal	BC Vancouver
2004	7.50	7.50	7.50	8.00
2005	7.50	7.50	7.60	8.00
2006	7.50	7.75	7.75	8.00
2007	8.00	8.00	8.00	8.00
2008	8.40	8.75	8.50	8.00
2009	8.80	9.50	8.50	8.00
2010	8.80	10.25	9.50	8.00
2011	9.40	10.25	9.65	8.00
2012	9.75	10.25	9.90	8.75
2013	9.95	10.25	10.15	9.50
2014	10.20	11.00	10.35	10.25
2015	11.20	11.25	10.55	10.45
2016	12.20	11.40	10.75	10.85
2017	13.60	11.60	11.25	11.35
2018	15.00	14.00	12.00	12.65
2019	15.00	14.00	12.50	13.85

Table 2.3: Descriptive Statistics for Businesses, Users and Reviews

Businesses		
	Mean	S.D
Business Rating	3.59	(0.99)
Number of Reviews	43.47	(126.52)
Open	0.79	(0.41)
Restaurant Dummy	0.40	(0.49)
N	369,844	
Users		
	Mean	S.D
User Reviews	17.82	(61.59)
Average rating given	3.64	(1.19)
Elite Dummy	0.03	(0.17)
N	3,901,337	
Reviews		
	Mean	S.D
Rating	3.72	(1.47)
N	*16,653,518*	

Table 2.4: Ratings Impact for Restaurants: DID Specification

	(1)	(2)	(3)
Post X Treatment	-0.037**	-0.026**	-0.035
	(0.015)	(0.012)	(0.024)
Constant	3.764***	3.768***	3.696***
	(0.017)	(0.002)	(0.003)
Month FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Business FE	No	Yes	No
Bus-User FE	No	No	Yes
N	5567805	5567130	337599

AZ treated, NV IL NC SC PA WI TX GA Controls. SE clustered at zipcode.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.5: Ratings Impact for Restaurants: Panel Regressions

	(1)	(2)
Min Wage	-0.007**	-0.022***
	(0.003)	(0.006)
Constant	3.802***	3.875***
	(0.028)	(0.053)
Month FE	Yes	Yes
Year FE	Yes	Yes
Business FE	Yes	No
Bus-User FE	No	Yes
N	11371172	691204
R-Square	0.19	0.75
Adj. R-Square	0.18	0.54

Restaurant reviews for all 18 metropolitan areas. SE clustered at zipcode.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.6: Panel Regressions of Log Prices on Log of MW for Restaurants

	(1)	(2)
Ln(MW)	0.047*	0.190
	(0.028)	(0.134)
Constant	2.262***	1.902***
	(0.059)	(0.281)
Month FE	Yes	Yes
Year FE	Yes	Yes
Business FE	Yes	No
Bus-User FE	No	Yes
N	995152	31384
R-Square	0.34	0.84
Adj. R-Square	0.29	0.70

Sample restricted to restaurant reviews where dollar prices mentioned explicitly.

SE clustered at zipcode.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.7: DID Results for Text-Based Outcomes: Coefficients at 95th, 90th, and 85th percentiles

	Price	Friendliness	Hygiene	Wait Times	Portion Sizes
DID at 95th	0.009	0.017***	0.015*	0.005	0.001
	(0.008)	(0.006)	(0.009)	(0.008)	(0.009)
DID at 90th	0.006	0.008*	0.008	0.001	0.001
	(0.006)	(0.004)	(0.005)	(0.003)	(0.006)
DID at 85th	0.007	0.006	0.003	0.002	-0.002
	(0.005)	(0.005)	(0.005)	(0.003)	(0.006)
N			1599733		

AZ is treated, NV, IL, NC, SC, PA, WI, TX, GA are controls

Figures

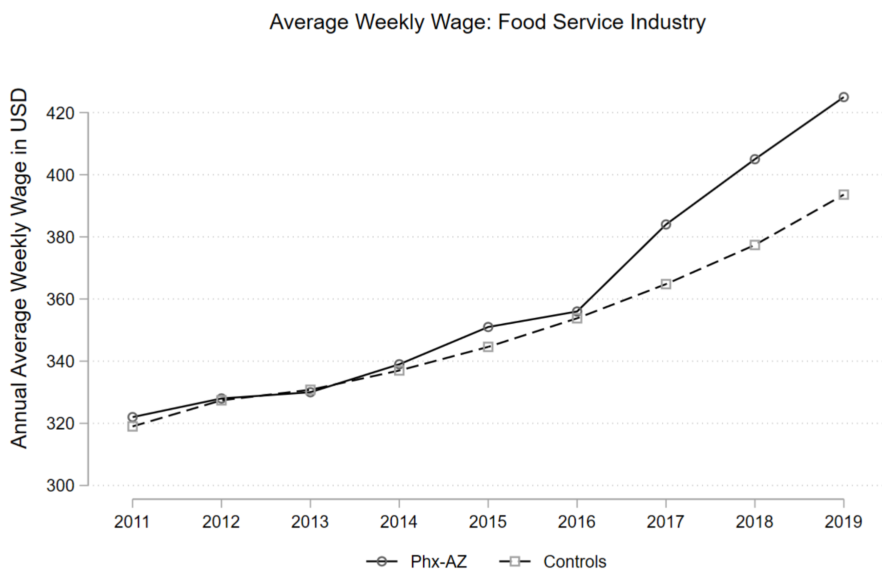

Figure 2.1: Average Annual Weekly Wages in Food Service (QCEW Data)

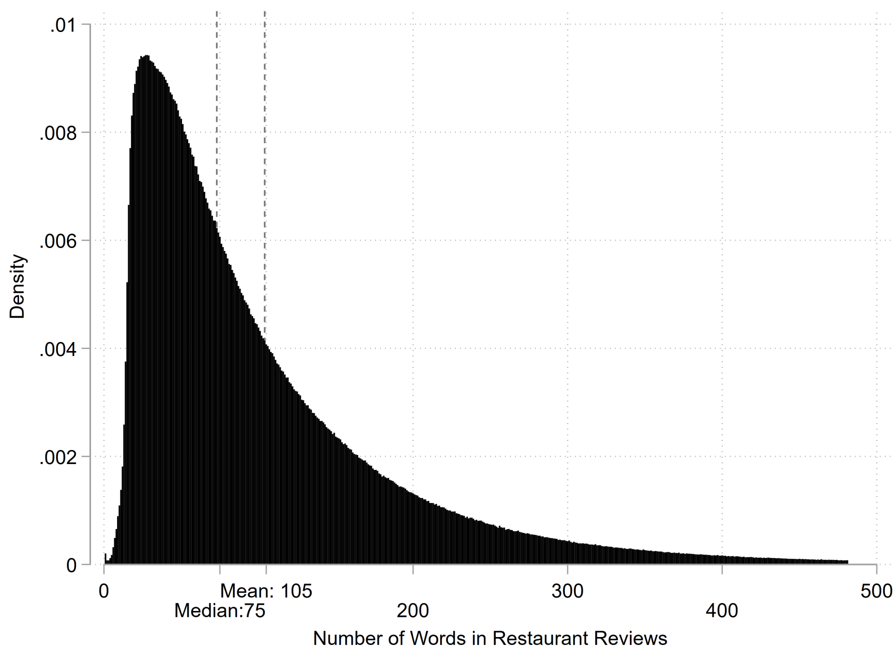

Figure 2.2: Histogram of Word Counts in Review Texts for Restaurants

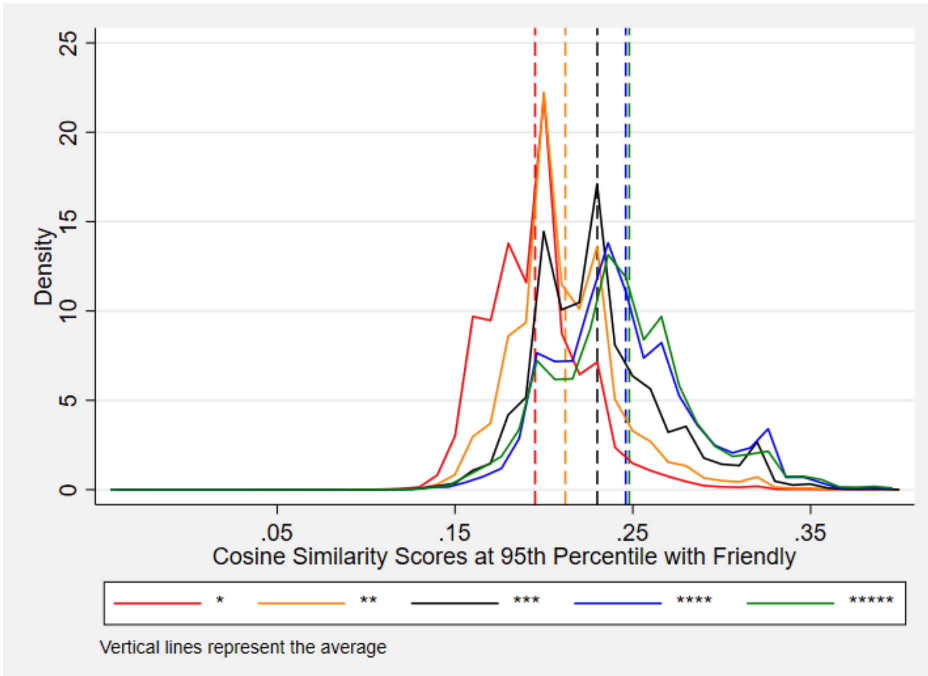

(a) Similarity with "Friendly"

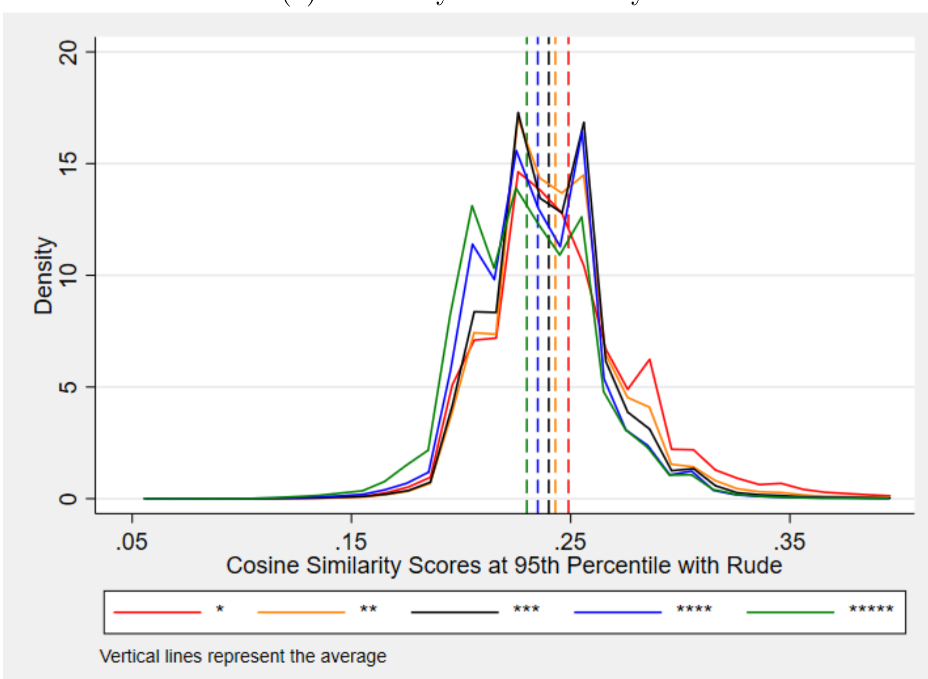

(b) Similarity with "Rude"

Figure 2.3: Distribution of Cosine Similarity Scores, by Review Ratings (Stars)

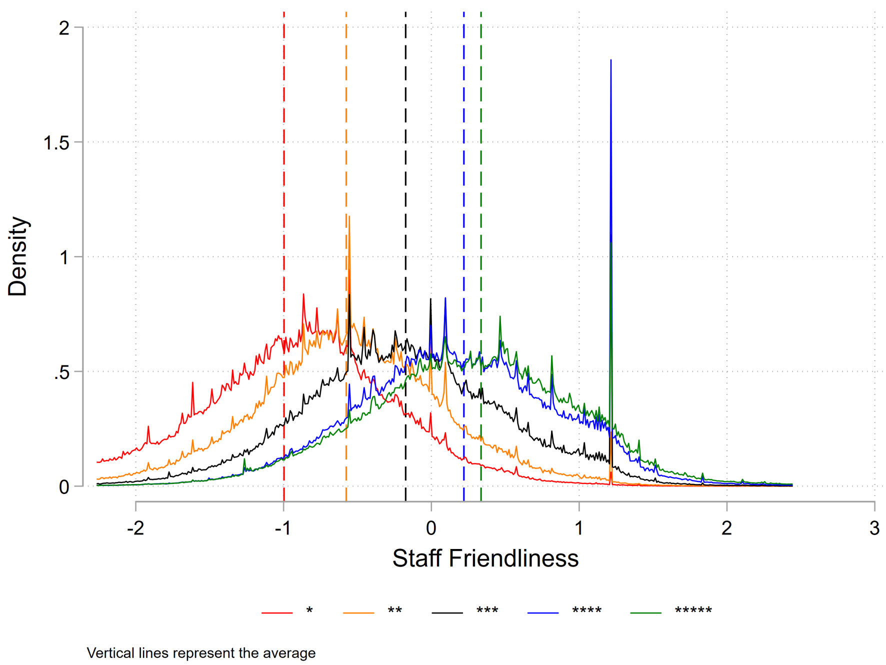

Figure 2.4: Distribution of “Friendliness” Outcome, by Review Ratings (Stars)

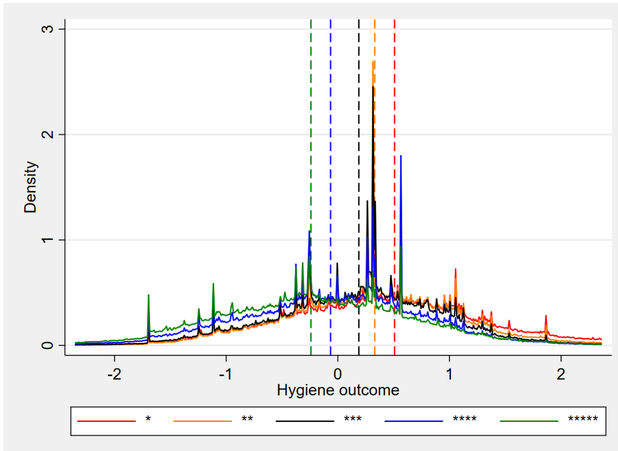

(a) Distribution of Hygiene Outcome, by stars

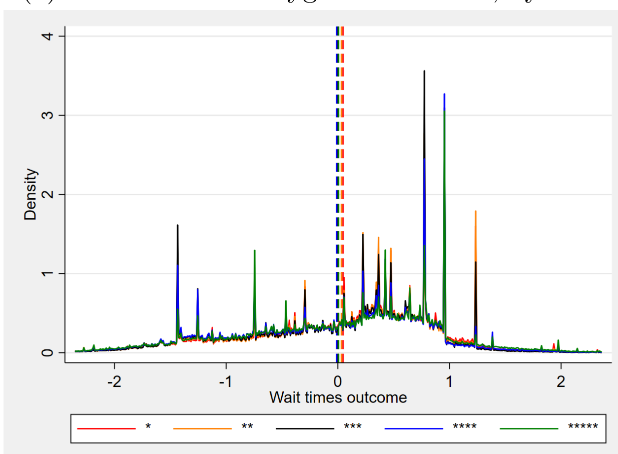

(b) Distribution of wait-time Outcome, by stars

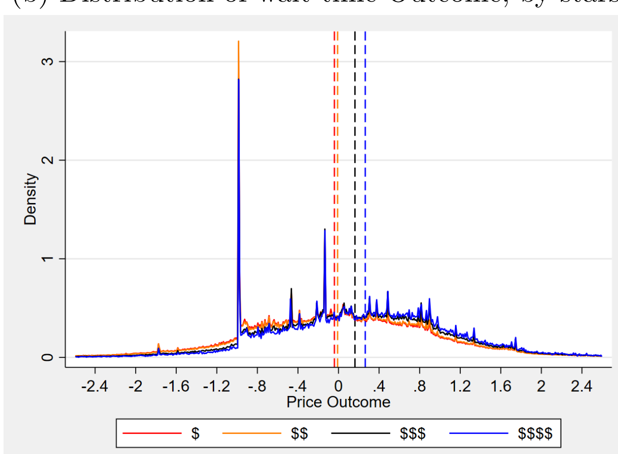

(c) Distribution of Price Outcome, by Yelp $ signs

Figure 2.5: Validating Text-Based Outcomes

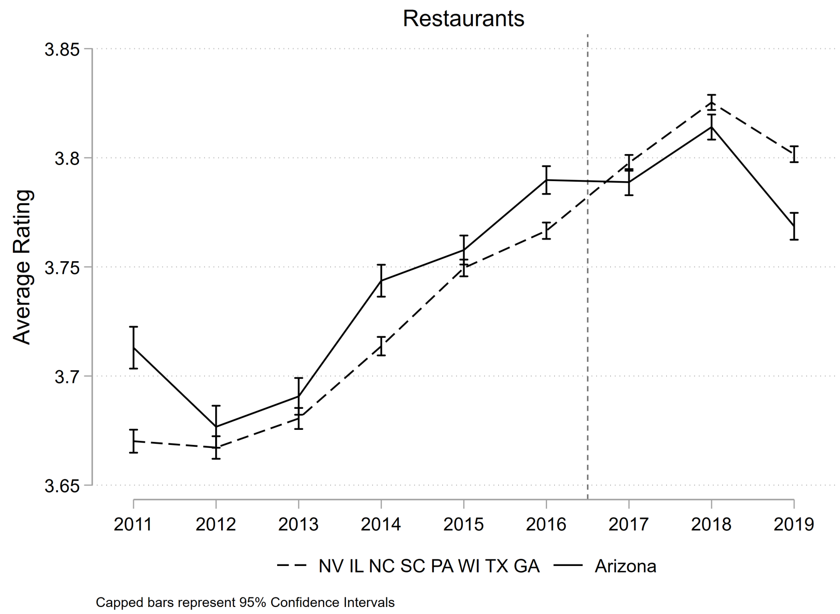

(a) Annual average ratings for restaurant reviews in Phoenix, AZ and control, 2011-2019

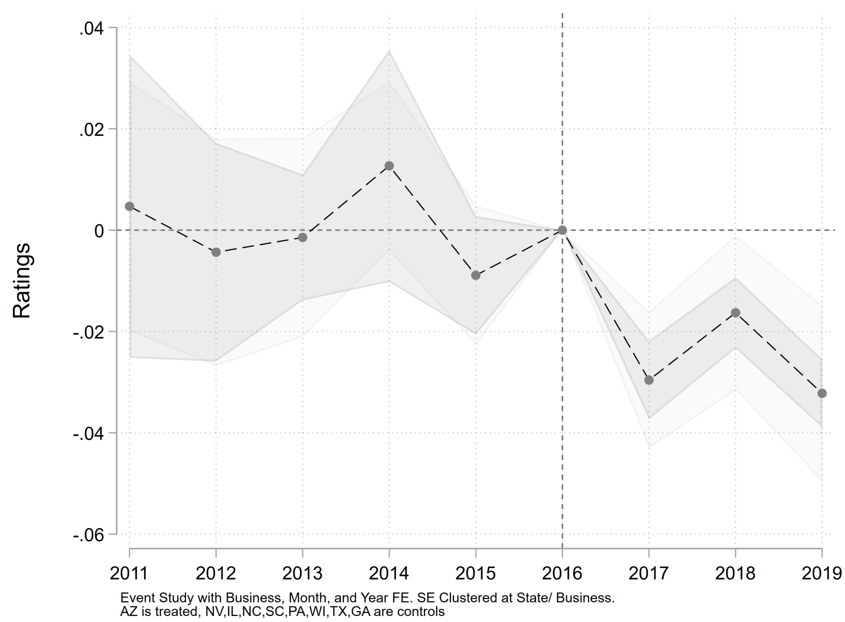

(b) Event-Study Estimates for Impact on Restaurant Ratings

Figure 2.6: Ratings associated with Restaurant Reviews

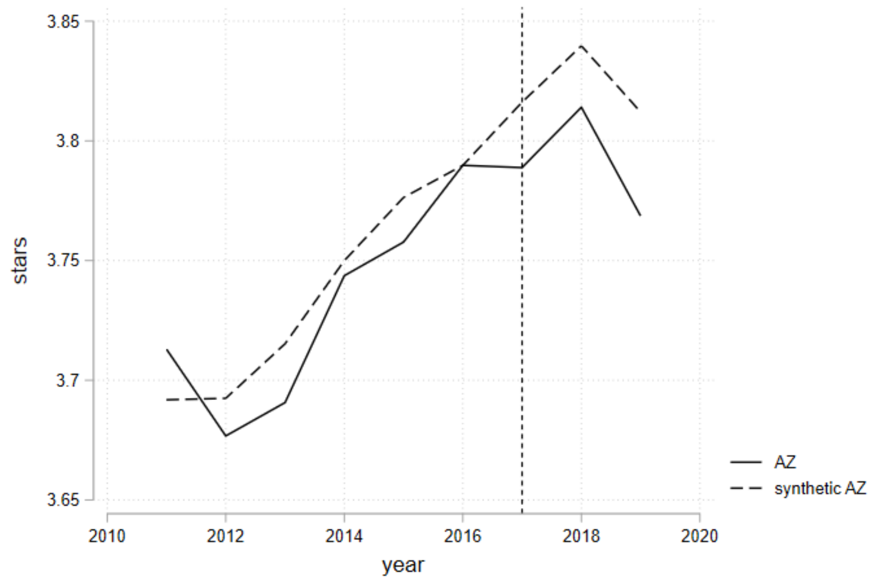

Figure 2.7: Synthetic Control for Impact on Restaurant Ratings

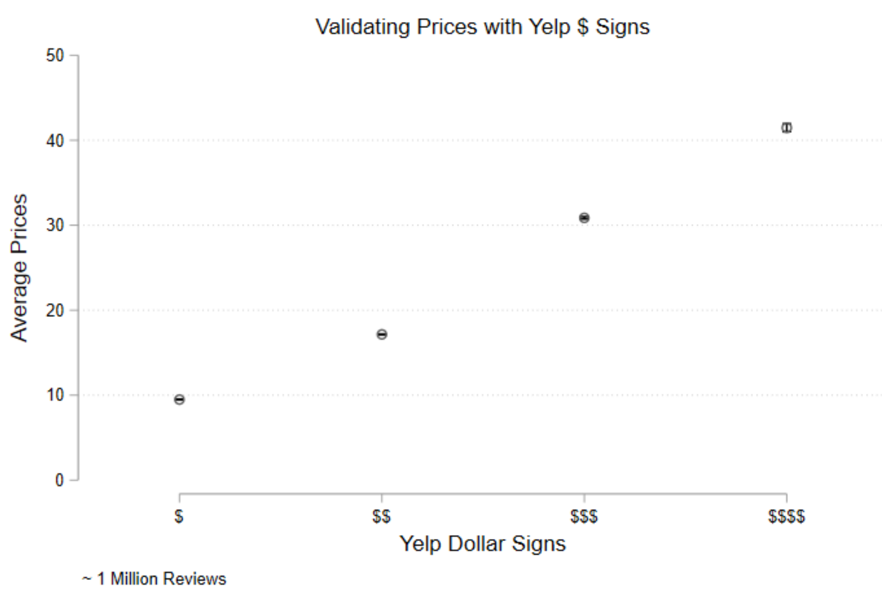

Figure 2.8: Average Extracted Dollar Prices by Yelp “$” Signs

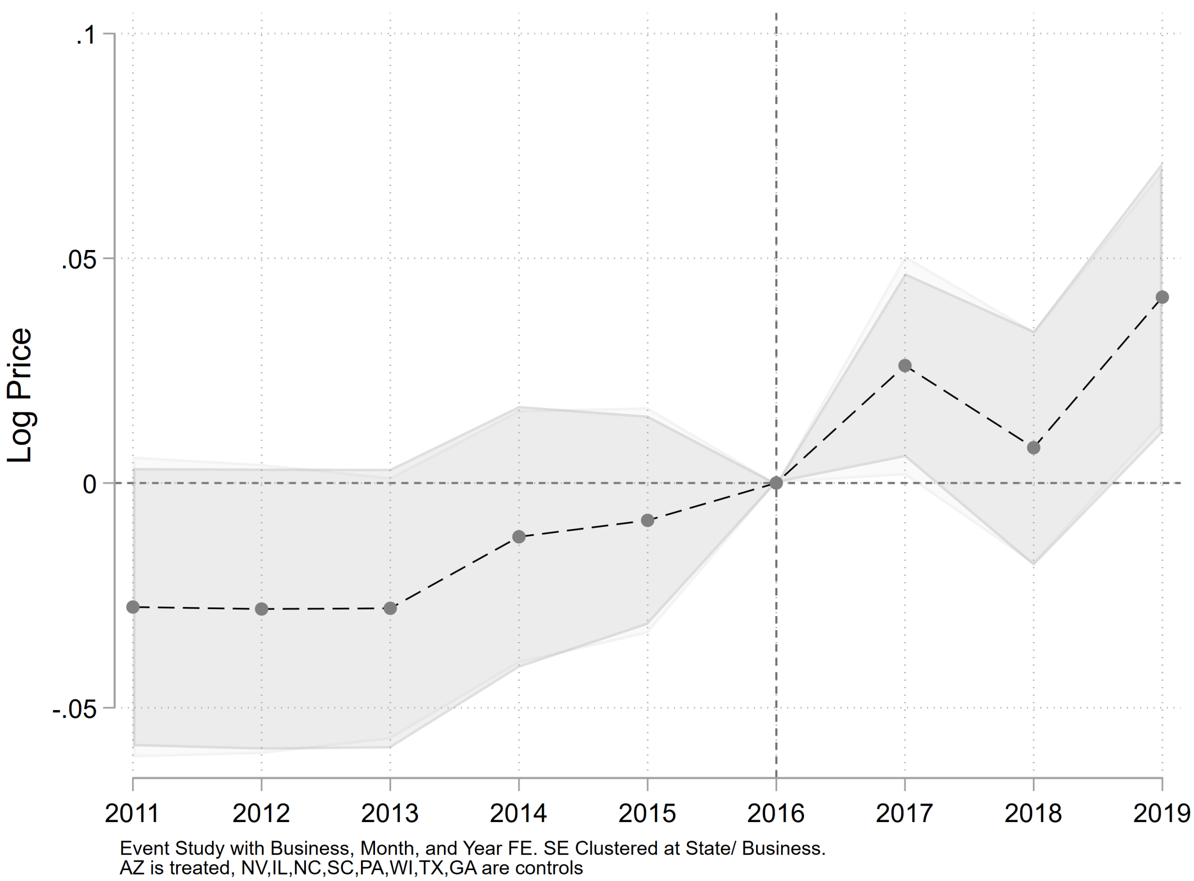

Figure 2.9: Event Study Estimates for Log of Prices

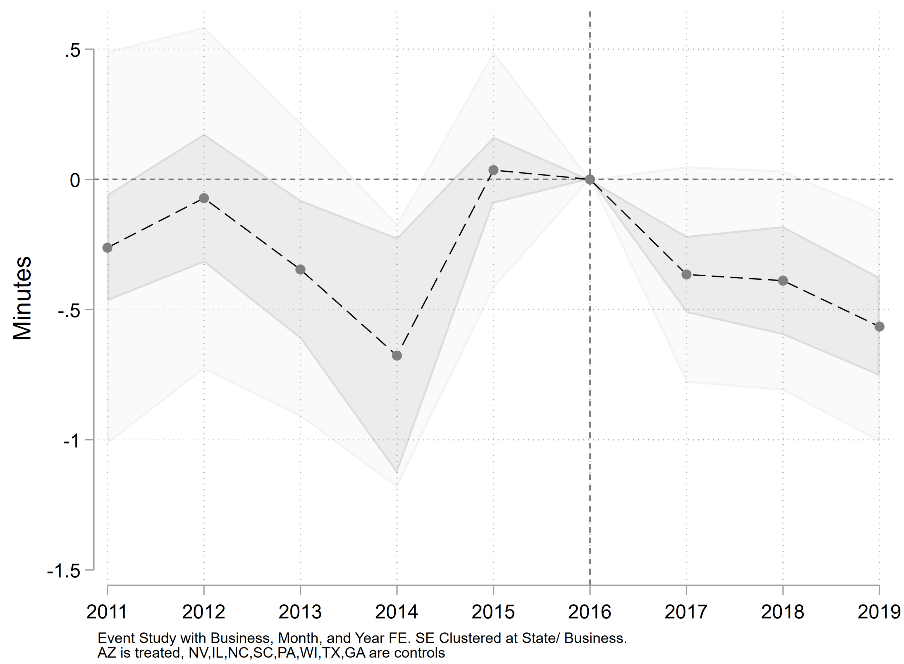

Figure 2.10: Event Study Estimates for Wait Times (in minutes)

Chapter 3

The Impacts of Opportunity Zones on Zone Residents

(with Matthew Freedman and David Neumark)

We are grateful for helpful comments from Stuart Rosenthal and two anonymous referees. We also thank Timothy Bartik, Aaron Hedlund, Rebecca Lester, Shawn Rohlin, Brett Theodos, and participants at the Brookings Institution conference "Opportunity Zones: The Early Evidence" for comments on earlier versions on this paper. This paper uses restricted-access data from the U.S. Census Bureau. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at the UC Irvine Federal Statistical Research Data Center under FSRDC Project Number 2146 (CBDRB-FY21-P2146-R8858, CBDRB-FY21-P2146-R9150, CBDRB-FY21-P2146-R9197).

3.1 Introduction

There is a lack of clear evidence that the most prominent place-based policy - enterprise zones – have created jobs and raised incomes for the least-advantaged people in neighborhoods with high concentrations of low-income residents (see the review in Neumark and Simpson (2015)). Nonetheless, with strong encouragement from the Trump Administration, the Tax Cuts and Jobs Act of 2017 created a new place-based policy: "Opportunity Zones." Opportunity Zones are targeted at disadvantaged census tracts and are intended to spur economic development. Opportunity Zone incentives are directed at investors in property, allowing deferral or avoidance of federal taxes on capital gains from investments in these zones. In this paper, we provide early evidence on the impacts of Opportunity Zone designation on residents of zones, focusing in particular on employment, earnings, and poverty. We take advantage of restricted-access microdata from the American Community Survey (ACS) for 2013-2019 to explore the program's impacts at a geographically granular level. We estimate effects for tracts designated as Opportunity Zones using a control group of eligible, but not designated tracts matched on the basis of trends in outcomes prior to the program's introduction. With our data, we can study effects up to about one-and-a-half years after enactment of the zones. Overall, we find limited evidence that Opportunity Zone designation has positive effects on the economic conditions of local residents. Our preferred estimates based on an inverse probability weighting (IPW) approach point to effects of Opportunity Zone designation that are economically small and generally statistically indistinguishable from zero. Specifically, we estimate that following Opportunity Zone designation, employment rates of residents do not change, with statistically insignificant yet fairly precise estimates that are very near zero. We can rule out increases in employment rates larger than 0.2 percentage point with 95% confidence. Estimated effects on average earnings of employed residents of designated tracts are positive, but are economically small and not consistently statistically significant. We additionally find that, if anything, zone designation is associated with a slight increase in local poverty rates, although the evidence is largely consistent with no effect. Notably, a difference-in-differences approach that ignores differential pre-designation trends implies positive effects on zone resident employment rates and reductions in poverty rates, with effects that are both statistically significant and economically meaningful. The problem of differential pre-designation trends is also apparent from an event-study analysis. Hence, an approach that assumes that zone selection was orthogonal to tracts' economic trajectories gives the misleading impression of substantial positive effects of zone designation on residents. Given that Opportunity Zone designations were first announced in 2018, we are at the beginning of research on the impacts of the program. [2] In recent work, Arefeva et al. (2021) leverage establishment-level data (the Your-economy Time Series) and find that Opportunity Zone designation increased employment growth relative to comparable (eligible, but not chosen) tracts significantly (by 3.0 to 4.5 percentage points), with the growth spread across industries. Atkins et al. (2020) study Opportunity Zone effects on job postings (from Burning Glass) by zip code, distinguished by whether the zip code contains at least one Opportunity Zone tract or not. They find only limited evidence of any effects of zone designation on job postings or posted salaries. Other recent work has studied the effects of Opportunity Zone designations on investment and real estate markets. While Corinth and Feldman (2021) find no impacts of zone designation on commercial investments, Sage et al. (2019) document significant positive effects on the prices of some types of commercial properties. Frank et al. (2020) also find positive effects of Opportunity Zone designation on commercial real estate transactions, building permits, and construction employment. However, Chen et al. (2019) find little effect of Opportunity Zone designation on residential property prices. The main contribution of our paper is that we identify the impacts of the Opportunity Zone program on zone residents as opposed to businesses, workers, or property values. To the extent that a major motivation for the Opportunity Zone program was improving outcomes for residents

[2] Earlier work on the federal New Markets Tax Credit, which is the most similar prior place-based policy, found a positive impact on investment, mainly via real-estate investment, coupled with a modest and costly poverty reduction effect (Freedman, 2012). Lester et al. (2018) discuss the similarities and differences between the New Markets Tax Credit and Opportunity Zones.

of distressed communities - as evidenced by the criteria for designating Opportunity Zones being based largely on the economic circumstances of residents - the impacts of the program on residents is of paramount importance. Past work on place-based policies suggests that even those programs that are effective at creating jobs may not deliver benefits to residents of targeted places (Busso et al., 2013; Freedman, 2015; Reynolds and Rohlin, 2015). The institutional structure of the Opportunity Zone program raises concerns that any job creation or investment spurred by the program may have limited benefits for local residents (Gelfond and Looney, 2018; Eastman and Kaeding, 2019). Our data on the economic conditions of those living in Opportunity Zones allow us to speak directly to the program's benefits for residents. By examining impacts on residents, we also provide evidence comparable to that for enterprise zones and other place-based policies, (e.g, Freedman (2012); Busso et al. (2013); Neumark and Young (2019)). An additional contribution of our evaluation is the use of rich, granular demographic and economic information available in the confidential ACS together with alternative empirical approaches based on selecting suitable sets of comparison groups. We first leverage our detailed data to establish that tracts designated as Opportunity Zones were on different trajectories, particularly with respect to employment and poverty rates, than tracts eligible but not designated. These differential trends undermine a simple difference-in-differences approach to estimating the causal effects of Opportunity Zone designation using eligible but not designated tracts as a control group. We therefore use a refined comparison group, assigning weights to control tracts to better match the evolution of outcomes for treated tracts prior to Opportunity Zone designation. By addressing differences in underlying trends in outcomes across designated and non-designated (but eligible) tracts, our matching approaches deliver more credible estimates of the program's effects on residents of targeted areas.

3.2 The Opportunity Zone Program

The Opportunity Zone program was introduced as part of the 2017 Tax Cut and Jobs Act (TCJA). The Opportunity Zone program provides preferential tax treatment for capital gains from investments in certain designated census tracts. There are a number of tax benefits associated with investing in Opportunity Zones: temporary deferment of taxes owed on realized capital gains from liquidating an asset if those gains are invested in businesses or real estate in Opportunity Zones, a basis step-up for realized capital gains that are reinvested in Opportunity Zones, and non-taxation of capital gains on Opportunity Zone investments if those investments are held for ten years or more (Theodos et al., 2018; Department of Treasury, 2018). The TCJA legislation provided for designation of Opportunity Zones in the 2018 tax year. It allowed governors to designate as Opportunity Zones up to 25% of census tracts in their state that qualified as "low-income communities" (LICs), as well as some tracts contiguous with LICs. An LIC is a tract with a poverty rate of at least 20% or median family income less than or equal to 80% of the greater of metropolitan area or statewide median family income (just statewide for rural tracts). Tracts within a federal Empowerment Zone, with population below 2,000, and contiguous with one or more LICs are also LICs. By law, 95% of Opportunity Zone tracts had to be LICs; governors were permitted to choose some additional tracts to designate as Opportunity Zones if the tracts were contiguous with an LIC and had median income less than 125% of the median income of the LIC with which it was contiguous. In total, 42,176 tracts were eligible to be Opportunity Zones - 31,864 LICs and 10,312 non-LIC contiguous tracts. Nationwide, governors selected 8,762 tracts as Opportunity Zones; 97% (8,532) were LICs while only 3% (230) were non-LIC contiguous tracts. States made their designations by June 2018 (Theodos et al., 2018; Department of Treasury, 2018). Figure 3.1 maps Opportunity Zones across the United States. The map suggests that Opportunity Zones are diffusely distributed geographically, in contrast to the more focused approach taken in, for example, the federal Empowerment

Zone program (Busso et al., 2013). Prior work suggests that place-based policies can be more effective if they are carefully targeted at distressed areas where there is the potential to generate agglomeration externalities (Glaeser and Gottlieb, 2008; Moretti, 2010). However, evidence on the selection process for Opportunity Zones points to only limited attention to such strategic considerations. Theodos et al. (2018) find that tracts selected as Opportunity Zones were more economically distressed than other eligible tracts, but that there was only a limited amount of targeting toward more disadvantaged neighborhoods with lower access to capital. Similarly, Alm et al. (2021) find that selection processes for Opportunity Zones were generally not sophisticated, and may have in fact favored areas that were already better positioned to grow in the future. Some evidence points to political favoritism in governors' zone selections (Alm et al., 2021; Frank et al., 2020; Eldar and Garber, 2020), but there is also evidence indicating that governors largely rubber-stamped recommendations for zone designations that came from mayors (Duarte et al., 2021). Several papers have also highlighted that, at least along some dimensions, tracts that were designated as Opportunity Zones were on different trajectories than tracts eligible but not designated (Frank et al., 2020; Atkins et al., 2020) - consistent with our evidence. According to the Internal Revenue Service, Opportunity Zones represent a tool "to spur economic development and job creation in distressed communities" (Internal Revenue Service, 2020). In principle, the program could affect employment in designated tracts, even in the short run. The funds flowing into qualified areas under the program can finance a wide variety of projects, including infrastructure, commercial or industrial real estate, and new or existing businesses. Some of these projects might be associated with new, albeit transitory, construction employment. Others may be associated with more durable jobs in different industries. Still others may be labor neutral or even labor displacing (Patrick, 2016; Brachert et al., 2019; Criscuolo et al., 2019). Given that our data cover a relatively short period after the creation of Opportunity Zones, it is important to consider the possible dynamics of the effects of the policy. It is conceivable that there will be more meaningful changes in zone economic conditions as more Opportunity

Zone capital is deployed in the future. Conceptually, however, we would expect the short-run effects on employment and wages to be larger than the long-run effects. Opportunity Zones could generate some immediate job growth from luring investment to an area. In the longer run we would expect the tax benefits to be capitalized into land values, increasing property prices and driving employment rates and real wages back toward their equilibrium levels, although this can be mediated by agglomeration and multiple equilibria (Glaeser and Gottlieb, 2008; Moretti, 2010; Bartik, 2020). Evidence indicates that long-run effects of one-time increases in local job opportunities do in fact tend to be smaller than short-run effects (Freedman, 2017; Garin, 2019). Recent research on U.S. as well as French enterprise zones also finds no evidence of growing effects many years after zone incentives are created (Neumark and Kolko, 2010; Gobillon et al., 2012; Givord et al., 2018; Neumark and Young, 2019).

3.3 Data

Our data on tracts eligible and designated as Opportunity Zones come from the Community Development Financial Institutions (CDFI) Fund at the U.S. Department of Treasury. [3] To construct outcomes, we take advantage of restricted-access American Community Survey (ACS) data for 2013-2019, which we accessed in a Federal Statistical Research Data Center (FSRDC). The advantage of the restricted-access ACS data is that we can measure outcomes at the tract-level on an annual basis; the public-use data only provide tract-level information averaged over five years. However, due to sample sizes and confidentiality restrictions, we are limited in the extent to which we can drill down to look at outcomes measured for sub-geographies (e.g., individual states) or examine heterogeneity in effects across areas with different initial conditions or other characteristics. We focus on the effects of Opportunity Zones on residents of designated areas. We construct three main outcome

[3] See `https://www.cdfifund.gov/opportunity-zones`

measures: the employment-to-population ratio for residents, average earnings of employed residents, and the poverty rate for residents. We aggregate the individual microdata to the tract-by-year level, using the person weights in the ACS. We only keep tracts that have complete information for all our outcomes of interest. [4] We additionally restrict attention to designated and eligible tracts that are LICs; while non-LIC contiguous tracts represent over one-fifth of tracts technically eligible, limits on how many of these tracts could be chosen as Opportunity Zones as well as a seeming preference for designating more distressed tracts led to only 230 non-LIC contiguous tracts being designated nationwide (3% of the total). Including the complete set of non-LIC contiguous tracts in the sample would, at least for the event-study and difference-in-differences analyses, necessitate using a disproportionate number of higher-income tracts as controls - controls that are less comparable to the final set of designated tracts. Taken together, these restrictions reduce our sample of designated Opportunity Zones to (a rounded) 7,600 tracts, and our sample of eligible but not designated tracts to (a rounded) 23,000. [5] We conduct our main analyses using this sample of tracts for the 2013-2019 period. Descriptive statistics for our sample, broken out by year and for designated Opportunity Zone tracts as well as eligible but not designated tracts, appear in Table 3.1. On average, the adult population of tracts in the sample is around 3,200. Consistent with earlier findings, tracts designated as Opportunity Zones have lower employment rates, lower average earnings, and higher poverty rates than tracts eligible but not designated as Opportunity Zones (e.g., Theodos et al. (2018); Frank et al. (2020)).

[4]We also exclude from the analysis Puerto Rico, where all eligible LICs were designated as Opportunity Zones.

[5]These counts of tracts are rounded for confidentiality reasons. While in principle we could estimate effects for LIC and non-LIC designated tracts separately, doing so would pose potential disclosure problems in light of the small number of non-LIC tracts that were selected as zones.

3.4 Empirical Approach

We begin by using an event-study framework to assess the comparability of designated Opportunity Zones and eligible but not designated tracts in terms of pre-existing trends in employment, earnings, and poverty. Our basic model is

$$y_{it} = \sum_{j=2013}^{2016} \left\{ \beta_j^{pre} \times OZ_i \times 1\left[j = t\right] \right\} + \sum_{k=2018}^{2019} \left\{ \beta_k^{post} \times OZ_i \times 1\left[k = t\right] \right\} + \gamma_i + \eta_t + \varepsilon_{it} \quad (3.1)$$

In equation 3.1, y_{it} is the outcome of interest for tract i in year t. OZ_i is a dummy that takes a value of 1 if tract i is designated as an Opportunity Zone and 0 if it is eligible but not designated; recall that the sample is restricted to designated Opportunity Zones and eligible but not designated LICs. The tract fixed effects in the model (γ_i) control for time invariant tract characteristics that could be correlated with Opportunity Zone designation and also independently affect outcomes. [6] The year fixed effects in the model (η_t) control for factors changing each year that are common to all tracts in the sample. Finally, β_j^{pre} and β_k^{post} are the estimated "effects" of Opportunity Zones for each pre- and post-treatment year. [7] These are measured relative to 2017. In some of the event-study specifications, we additionally include county-by-year fixed effects in an effort to account for potentially differential trends in outcomes across geographies at a higher level of aggregation than census tracts. This more saturated model effectively narrows the set of control tracts for any given treatment tracts to those more geographically proximate. While this limits the scope for potential unobservable time-varying factors to bias our estimates, it may amplify bias attributable to spillovers of Opportunity Zone effects across nearby tracts. As a point of comparison with our preferred estimates based on matching methods described below, we also run the simple difference-in-differences version of equation 3.1, in which β_j^{pre} is constrained to be zero and we estimate a single β_k^{post} parameter (where the post-period is defined as 2018-2019 or, in a robustness

[6] The tract fixed effect also subsumes the main effect for OZ_i.

[7] We put effects in quotes because in the pre-treatment period, at a minimum, we do not think of the estimates as capturing causal effects.

check, as only 2019). For the event-study and difference-in-differences analyses, we cluster standard errors at the tract level, which allows for arbitrary patterns of heteroskedasticity across tracts and serial correlation within tracts. Motivated by the results of our event-study analyses, which point to violations of the parallel trends assumption necessary to interpret the difference-in-differences estimates as causal, we proceed to construct a control group using a data-driven approach to weight potential comparison tracts. Our preferred results use inverse probability weighting (IPW) as well as the doubly-robust inverse probability weighted regression adjustment method. For these approaches, we construct the dependent variable as simply the average of the outcome for each tract over the post-designation period (defined as 2018-2019 or, in a robustness check, as only 2019) minus the average over the pre-designation period (defined as 2013-2017 or, for the robustness check, as 2013-2018). [8]. We want to compare this outcome for treated (Opportunity Zone) and control (eligible but not designated) tracts. With IPW, an estimate of the unobserved counterfactual of the average outcome for the treated tracts, if Opportunity Zone designation had not occurred, is constructed as a weighted average across non-treated tracts, where the weights are the inverse of the probability that the tract was not treated, adjusted for the probability of treatment.[9] These weights are estimated from a logit model, for which the underlying linear model for the latent variable is:

$$OZ_i^* = \alpha + \sum_{\tau=2013}^{2017} \{\rho_\tau emp_{i\tau} + \varphi_\tau earn_{i\tau} + \omega_\tau poverty_{i\tau}\} + v_i \tag{3.2}$$

That is, we predict Opportunity Zone designation for all tracts in our sample of LICs based on each tract's employment rate (emp), average earnings (earn), and poverty rate (poverty) in each year between 2013 and 2017 (i.e., over the entire pre-treatment period). Thus, the

[8]By construction, regressing this differenced dependent variable on a dummy for Opportunity Zone designation for our tract-level sample of LICs delivers identical point estimates as a standard difference-in-differences regression run on a tract-by-year sample with tract and year fixed effects (see columns (1), (3), and (5) of Appendix Table C.3).

[9]The expression for the weights for the non-treated tracts is $\frac{\hat{p}}{1-\hat{p}}$, where $\hat{p}$ are the predicted probabilities from Equation 3.2 - described below.

most weight is put on the non-treated tracts with the highest estimated probability of being treated based on the observables. The assumption is that the expected value of the weighted average of the outcome for the non-treated tracts equals the expected value of the outcome for the treated tracts if they were not treated, which is more plausible because we are using as controls tracts on trajectories more comparable to those of the treated tracts. While the IPW method above models the treatment, regression adjustment methods allow us to model the outcome to account for non-random treatment assignment. Regression adjustment methods construct counterfactuals by fitting linear regression models separately for the treated group and the control group and using the predicted values of the outcome for a given covariate vector as estimates of the potential outcomes. Averaging the covariate-specific treatment effect across treated tracts using these fitted values yields the ATT estimate. The regression-adjusted IPW method uses the IPW weights in the estimation of the regression adjustment models to estimate corrected regression coefficients, and therefore combines these two approaches. This estimator has the virtue of being "doubly robust." That is, it provides a consistent estimate as long as either the inverse probability weighting or the regression adjustment removes the confoundedness of treatment with unobservables - although both approaches, as this implies, use selection on observables. [10] In our application of regression-adjusted IPW, we model both the outcome and the treatment using the same set of covariates. [11] As we illustrate below, using the IPW and regression-adjusted IPW approaches, we can more credibly attribute differential changes in outcomes after Opportunity Zone designation to the program itself as opposed to continuations of pre-existing trends.

[10] Tan (2010) provides a detailed explanation of these estimators.

[11] With our matching approaches, the unit of observation is the census tract. We calculate heteroskedasticity robust standard errors following Wooldridge (2010)

3.5 Results

3.5.1 Event-Study Estimates

In Figure 3.2, we show event-study estimates for the three main outcomes discussed above. We show results based on models with just tract and year fixed effects as well as models including tract and county-by-year fixed effects. The graphs show point estimates and 95% confidence intervals. In each case, 2017 is the reference year; the figures report the interactions of OZ_i and the year dummy variables, with the interaction with the dummy variable for 2017 omitted, as in Equation 3.1. Event-study results for the employment rate of residents appear in Panel A of Figure 3.2. There is clear evidence of a differential pre-treatment trend in employment rates for those areas designated as Opportunity Zones relative to those areas eligible but not designated. In particular, relative to employment rates in other eligible tracts, employment rates in designated areas were trending upward prior to 2017, and the higher relative employment rates after 2017 appear to reflect merely the continuation of that trend. In Panel B of Figure 3.2, we see limited evidence of any Opportunity Zone effects on resident average earnings, although there is also less indication of a strong differential pre-trend in resident average earnings. However, Panel C of Figure 3.2 indicates a strong relative pre-treatment trend in resident poverty rates. The results suggest that, compared to poverty rates in other eligible areas, poverty rates in tracts that were designated as Opportunity Zones were already declining prior to designation, and that the post-treatment changes may reflect the continuation of the prior trend rather than the causal effect of Opportunity Zone designation. The clear violation of parallel trends, particularly for employment and poverty rates, would call into question any causal interpretation of results from a standard difference-in-differences approach that compares changes in outcomes for tracts designated Opportunity Zones to those eligible but not designated Opportunity Zones. [12] In light of

[12]For this reason, we present unadjusted difference-in-differences estimates only in the Appendix (see Appendix Table C.3).

the differential pre-treatment trends in outcomes and the resulting bias that would arise in a difference-in-differences framework that ignored these trends, we implement a matching strategy that balances treatment and control tracts on the pre-designation evolution of outcomes. We turn to our matching-based results in the next section.

3.5.2 Estimates using IPW

In this section, we present results from matching Opportunity Zone tracts to eligible but not designated tracts based on pre-treatment trends in outcomes. We begin by illustrating in Figure 3.3 that our IPW approach (without further regression adjustment) eliminates the differential pre-2017 trends in outcomes that existed for designated vs. all other eligible tracts. By design, our IPW approach ensures parallel trajectories in outcomes for designated Opportunity Zones and our (weighted) group of non-designated but eligible LIC tracts prior to 2017. [13] The contrast with the event study estimates from Figure 3.2 (reproduced in Figure 3.3 for comparison) is clear. Estimates of the effects of treatment on the treated using our IPW and regression-adjusted IPW approaches appear in Panels A and B of Table 3.2. For these results, the Opportunity Zone indicator takes a value of one for designated tracts in 2018 and 2019, and zero otherwise. We report estimates in which we define the post-treatment period as only 2019 in Appendix Table C.2; the results in that case are very similar, with some exceptions regarding statistical significance noted below. Our IPW estimates in Table 3.2 indicate that Opportunity Zone designation has no meaningful effect on the employment rates of residents on average (column (1)). The point estimates from both our IPW and regression-adjusted IPW approaches imply miniscule (negative) effects on employment rates. Based on these estimates, we can rule out with 95% confidence an effect size for employment rates larger than 0.2 percentage point. Our estimates of the effects of Opportunity Zone designation on average annual earnings, reported in column (2) of Table

[13] Descriptive statistics on the weights are reported in Appendix Table C.1. The large positive skewness indicates a long right tail, indicating that the weight is concentrated in a smaller number of control tracts.

3.2, are also economically small, but are positive and statistically significant. The IPW estimates of around \$350 represent 1% of the 2017 mean of average earnings in designated tracts (see Table 3.1). When we define the post-treatment period as 2019 only, the point estimates are smaller and not statistically significant (see Appendix Table C.2). Finally, we find that Opportunity Zone designation does not reduce poverty rates (and may even increase them slightly). In Table 3.2, the IPW and regression-adjusted IPW estimates of the treatment on the treated imply that designation is associated with a 0.4-0.6 percentage point (approximately 2%) increase in resident poverty rates. While the effect sizes are small, their significance implies that we can statistically rule out with 95% confidence any reduction in poverty due to Opportunity Zone designation. When we define the post-treatment period as only 2019, the IPW estimates for poverty rates remain positive but are less precise; in that case, the regression-adjusted IPW estimate is not statistically significant. Overall, then, a true effect of on poverty of zero or very close to zero is fully consistent with the estimates. By way of comparison, Appendix Table C.3 reports standard difference-in-differences estimates for the effects of zone designation on employment rates, average earnings, and poverty rates. For employment, the naïve difference-in-differences results suggest substantial gains from Opportunity Zone designation, with estimates indicating that zone designation is associated with a statistically significant 0.5 percentage point increase in resident employment rates. But as we saw from Figure 3.2, these estimates are misleading, as they reflect differential pre-existing trends between designated zones and eligible but not designated tracts; the IPW results in Table 3.2 indicate a much different conclusion, with a quite precise estimate that rules out material employment gains from Opportunity Zones. The estimated earnings effects in Appendix Table C.3 are more similar to those we obtain using IPW, which is not surprising given the lack of a pronounced pre-treatment trend in this outcome apparent in Figure 3.2. For poverty rates, like employment rates, the difference-in-differences estimates generate spurious evidence of beneficial effects from Opportunity Zone designation; specifically, they indicate statistically significant reductions in poverty rates of about 1 percentage point (4%).

The pre-trends in Figure 3.2, however, imply that the IPW estimates showing no effect or, if anything, a slight increase in poverty rates following zone designation are more credible.

3.6 Conclusion

We provide early evidence on the impacts of Opportunity Zone designation on residents of zones, estimating effects on employment, earnings, and poverty. We use restricted-access microdata from the American Community Survey (ACS) for 2013-2019 to explore the program's impacts at a geographically granular level, comparing outcomes in tracts designated as Opportunity Zones to those eligible but not designated. We show that credible evidence on the effects of zone designation requires using a control group of tracts matched on the basis of trends in outcomes prior to the program's introduction. Using an inverse probability weighting approach, we find little if any evidence of positive effects of Opportunity Zone designation on the economic conditions of residents of targeted neighborhoods. Our estimates of the effects of Opportunity Zone designation on employment rates are statistically indistinguishable from zero and sufficiently precise to rule out material increases. We find at best modest positive effects on the average earnings of zone residents. And we find no evidence that Opportunity Zones reduce local poverty rates. Other recent studies have found mixed evidence on the effects of Opportunity Zones on labor markets. While Atkins et al. (2020) document limited effects of zone designation on job postings, Arefeva et al. (2021) find that zone designation is associated with significant job growth. Evidence of job growth in Opportunity Zones does not necessarily contradict our results. As previously shown in the context of Empowerment Zones (Busso et al., 2013) and the New Markets Tax Credit (Freedman, 2015), the effects of place-based policies could be different for jobs in the zones vs. employment of zone residents. Institutional features of the Opportunity Zone program may be militating against meaningful positive impacts of the program on residents of targeted communities (Gelfond and Looney, 2018). The program lacks several features of place-based policies that have been found to be associated with stronger job creation effects, including directly incentivizing hiring, focusing the incentives on new hires rather than windfalls, and facilitating the recapture of tax benefits if job creation or investment goals are not met. [14] In addition, the Opportunity Zone selection process did not ensure that communities most in need were the ones ultimately eligible to receive tax-advantaged investments. Indeed, as our results suggest, many of the communities designated as Opportunity Zones likely would have experienced improvements in economic conditions even in the absence of zone designation. The Joint Committee on Taxation (2019) estimated that through 2023, the Opportunity Zone Program would cost approximately $3.5 billion each year in foregone tax revenues. The White House Council of Economic Advisors (2020) estimated that Qualified Opportunity Funds had raised as much as $75 billion in private capital by the end of 2019, although how much has been invested in Opportunity Zones is currently unclear (Government Accountability Office, 2020). We find limited evidence of any impacts of zone investment to date on zone residents. An important limitation of our study is that our estimates are "early," in the sense of extending only one-and-a-half years since Opportunity Zones were officially designated. It is possible that there will be more meaningful changes in zone economic conditions as more Opportunity Zone capital is deployed in the future. However, prior evidence from other place-based programs suggests that the long-run effects of the program are unlikely to be larger than the short-run effects (Neumark and Kolko, 2010; Gobillon et al., 2012; Givord et al., 2018; Neumark and Young, 2019). Additionally, given that the 2020 data will include a year strongly affected by the COVID-19 pandemic, with effects also extending into 2021, the data through 2019 may provide the most definitive evidence we can obtain for many years, barring future policy changes such as creating new Opportunity Zones or eliminating existing ones.

[14] See Neumark and Grijalva (2017) for evidence on features of hiring credits that more likely lead to net job creation, and Freedman et al. (2021) for evidence of positive job creation effects from an economic development policy that includes these features.

Table 3.1: Descriptive Statistics for Opportunity Zones and Other Low-Income Communities, 2013-2019

A. Treated Tracts (Opportunity Zone Tracts)								
	2013	2014	2015	2016	2017	2018	2019	All Years
Adult Population	3115	3151	3165	3195	3193	3207	3219	3178
	(1587)	(1583)	(1636)	(1652)	(1680)	(1680)	(1747)	(1653)
Resident Employment Rate	0.5010	0.5103	0.5185	0.5267	0.5312	0.5391	0.5471	0.5248
	(0.1379)	(0.1353)	(0.1362)	(0.1397)	(0.1424)	(0.1436)	(0.1475)	(0.1412)
Resident Poverty Rate	0.3144	0.3066	0.2919	0.2752	0.2614	0.2544	0.2415	0.2779
	(0.1735)	(0.1744)	(0.1692)	(0.1718)	(0.1738)	(0.172)	(0.1701)	(0.174)
Resident Average Earnings	28340	28860	30230	31470	32700	34270	35770	31660
	(15730)	(10880)	(11570)	(12750)	(13140)	(14120)	(14670)	(13600)
Tracts	7600	7600	7600	7600	7600	7600	7600	
B. Potential Control Tracts (Other Low-Income Communities Eligible but Not Designated)								
	2013	2014	2015	2016	2017	2018	2019	All Years
Adult Population	3173	3203	3232	3249	3269	3281	3283	3241
	(1542)	(1555)	(1597)	(1617)	(1649)	(1663)	(1699)	(1619)
Resident Employment Rate	0.5357	0.5430	0.5494	0.5568	0.5613	0.5663	0.5730	0.5551
	(0.1330)	(0.1325)	(0.1330)	(0.1362)	(0.1385)	(0.1396)	(0.1413)	(0.1369)
Resident Poverty Rate	0.2551	0.2525	0.2410	0.2226	0.2152	0.2093	0.1993	0.2279
	(0.1599)	(0.1564)	(0.1554)	(0.1560)	(0.1560)	(0.1554)	(0.1567)	(0.1578)
Resident Average Earnings	30480	31000	32240	33630	34820	36120	37900	33740
	(10930)	(10890)	(11850)	(13060)	(13050)	(13750)	(14850)	(12950)
Tracts	23000	23000	23000	23000	23000	23000	23000	

Notes: Data derived from the 2013-2019 American Community Surveys. Standard deviations in parentheses.

Table 3.2: Estimates of the Effects of Opportunity Zones on Residents

	(1)	(2)	(3)
	Employment Rate	Average Earnings	Poverty Rate
	A. IPW Treatment on the Treated Estimates		
Opportunity Zone	-0.0005	350.4***	0.0058***
	(0.0012)	(124.9)	(0.0017)
Tracts	30600	30600	30600
	B. Regression-Adj. IPW Treatment on the Treated Estimates		
Opportunity Zone	-0.0002	337.9***	0.0043***
	(0.0012)	(124.0)	(0.0016)
Tracts	30600	30600	30600

Notes: Data derived from the 2013-2019 American Community Surveys. The Opportunity Zone variable takes a value of one for designated tracts in 2018 and 2019. Panel A reports inverse probability weighted (IPW) estimates of the treatment on the treated. Panel B reports doubly robust regression-adjusted IPW estimates of the treatment on the treated. Standard errors in parentheses are heteroskedasticity robust. Statistically significant at * 10%, ** 5%, *** 1%.

Figures

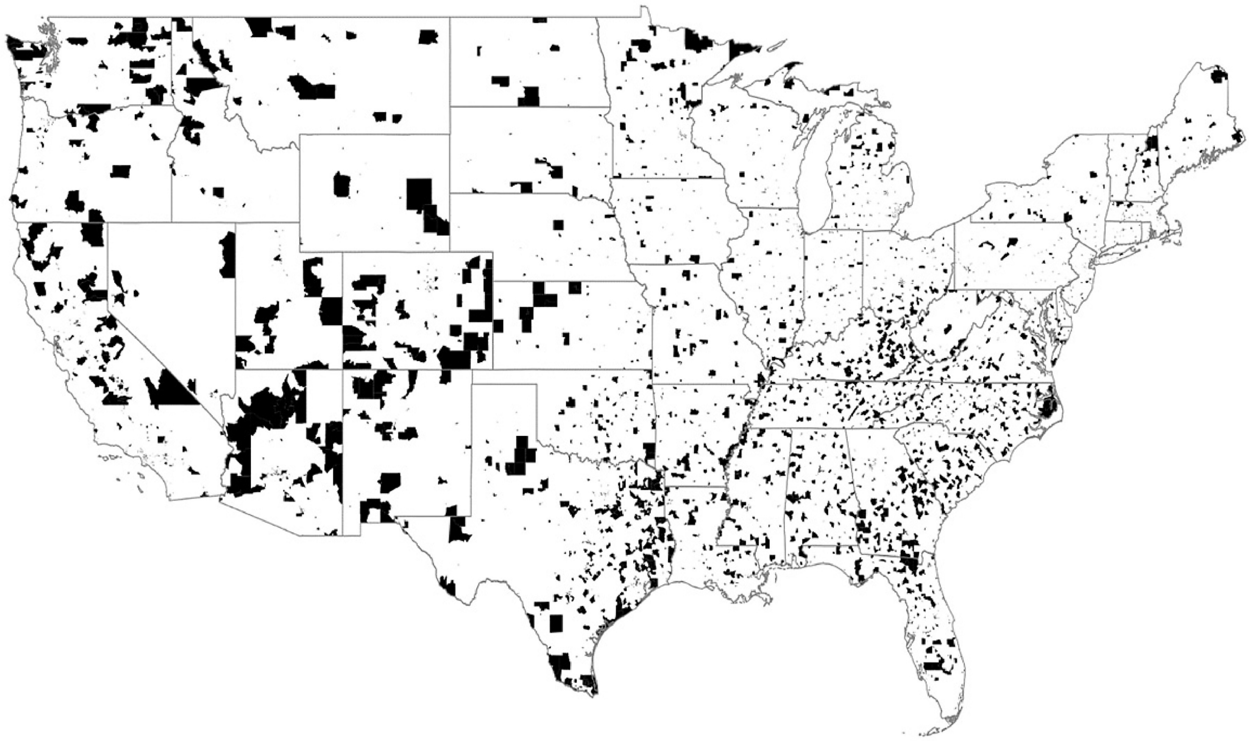

Figure 3.1: Opportunity Zones

Notes: Shaded areas are census tracts designated as Opportunity Zones. Information on Opportunity Zones is from the Community Development Financial Institutions (CDFI) Fund at the U.S. Department of the Treasury.

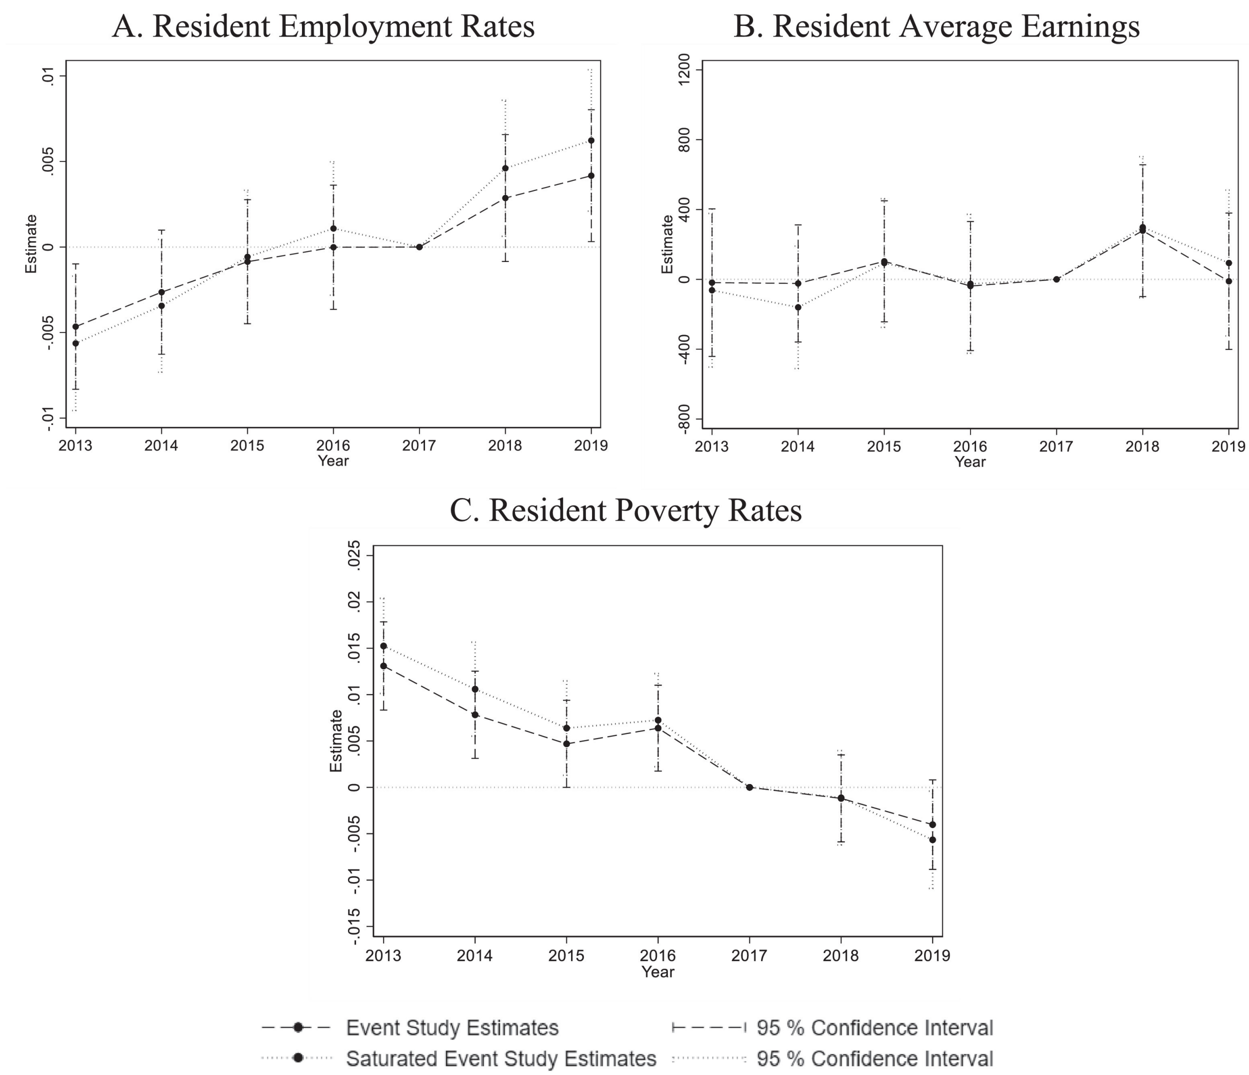

Figure 3.2: Event Study Estimates for Resident Outcomes

Notes: Data derived from the 2013-2019 American Community Surveys. The panels show point estimates and 95% confidence intervals from event studies comparing outcomes in Opportunity Zones to outcomes in eligible but not designated LICs. The dashed lines show results from an event study with only tract and year fixed effects. The dotted lines show results from an event study that additionally includes county-by-year fixed effects. 2017 is the omitted year. The confidence intervals are based on standard errors clustered at the tract level.

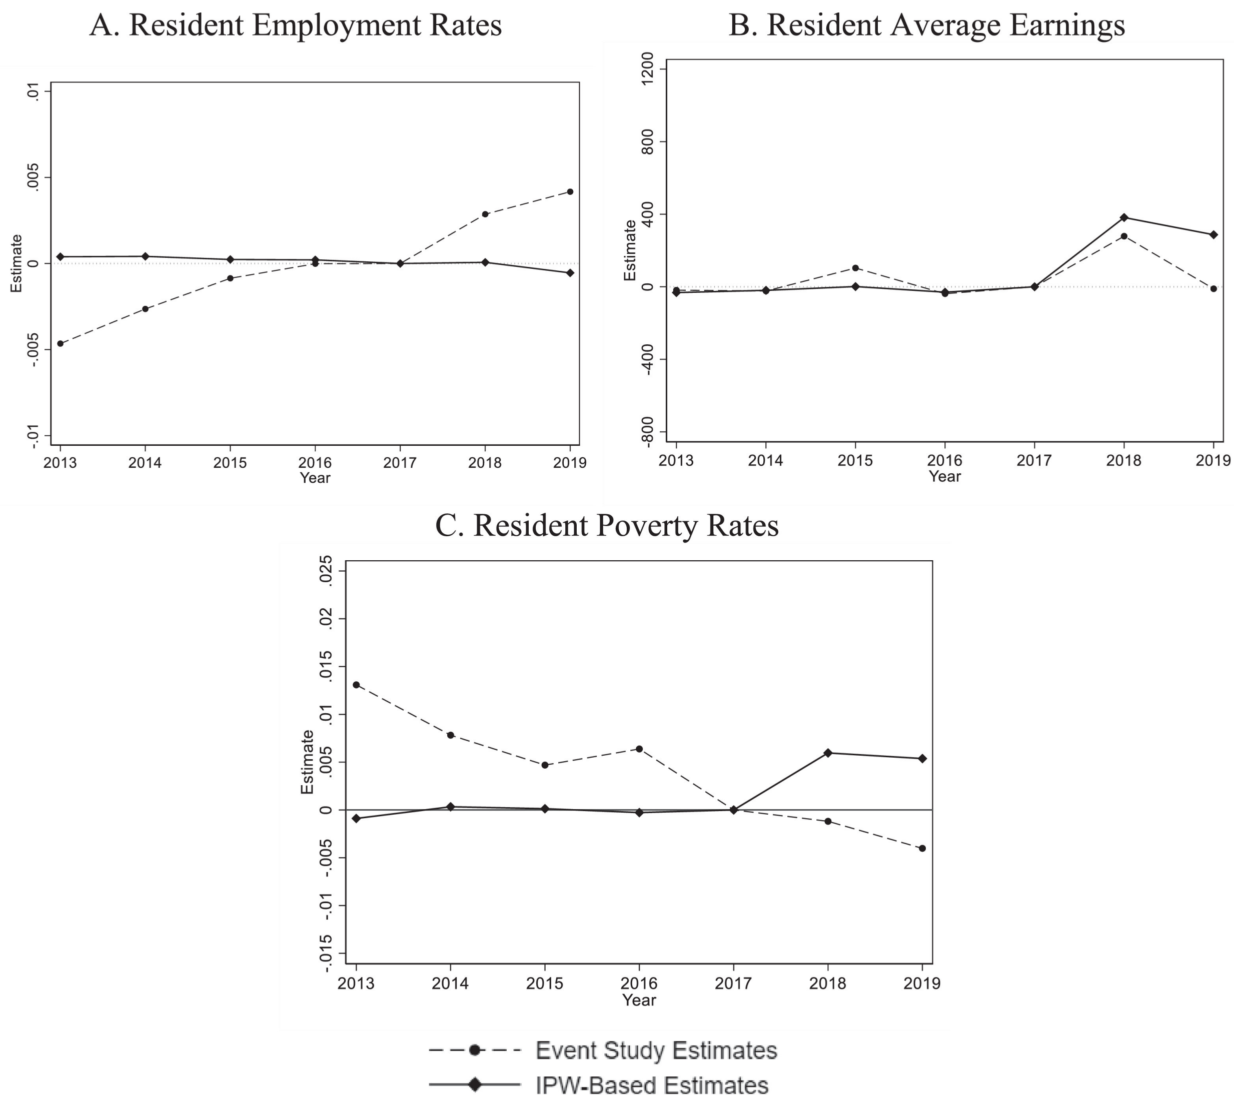

Figure 3.3: Event Study Estimates with Alternative Weighting Schemes

Notes: Data derived from the 2013-2019 American Community Surveys. The panels show point estimates from event studies using as controls all eligible but not designated LICs (reproducing the estimates with tract and year fixed effects in Figure 3.2 as well as using as controls eligible tracts weighted based on the estimated propensity to be treated (the IPW approach).

Bibliography

Aaronson, D., French, E., MacDonald, J., 2008. The minimum wage, restaurant prices, and labor market structure. The Journal of Human Resources 43, 688–720. URL: http://www.jstor.org/stable/40057364.

Abadie, A., Diamond, A., Hainmueller, J., 2010. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. Journal of the American Statistical Association 105, 493–505. URL: https://doi.org/10.1198/jasa.2009.ap08746, doi:10.1198/jasa.2009.ap08746, arXiv:https://doi.org/10.1198/jasa.2009.ap08746.

Agan, A.Y., Cowgill, B., Gee, L., 2020. Do workers comply with salary history bans? a survey on voluntary disclosure, adverse selection, and unraveling. AEA Papers & Proceedings 110. URL: https://ssrn.com/abstract=3522170.

Allegretto, S., Reich, M., 2018. Are local minimum wages absorbed by price increases? estimates from internet-based restaurant menus. ILR Review 71, 35–63. URL: https://doi.org/10.1177/0019793917713735, doi:10.1177/0019793917713735, arXiv:https://doi.org/10.1177/0019793917713735.

Alm, J., Dronyk-Trosper, T., Larkin, S., 2021. In the land of oz: designating opportunity zones. Public Choice 188, 503–523.

Anbarcı, N., Feltovich, N., Gürdal, M.Y., 2015. Lying about the price? ultimatum bargaining with messages and imperfectly observed offers. Journal of Economic Behavior & Organization 116, 346 – 360. URL: http://www.sciencedirect.com/science/article/pii/S0167268115001444, doi:https://doi.org/10.1016/j.jebo.2015.05.009.

Anderson, M., Magruder, J., 2012. Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. The Economic Journal 122, 957–989. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0297.2012.02512.x, doi:10.1111/j.1468-0297.2012.02512.x, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0297.2012.02512.x.

Arefeva, A., Davis, M.A., Ghent, A.C., Park, M., 2021. Job growth from opportunity zones. Available at SSRN 3645507.

Atkins, R., Hernandez-Lagos, P., Jara-Figueroa, C., Seamans, R., 2020. What is the impact of opportunity zones on employment outcomes? NYU Stern School of Business .

Barach, M.A., Horton, J., 2017. How Do Employers Use Compensation History?: Evidence from a Field Experiment. CESifo Working Paper Series 6559. CESifo Group Munich. URL: https://EconPapers.repec.org/RePEc:ces:ceswps:_6559.

Bartik, T.J., 2020. Using place-based jobs policies to help distressed communities. Journal of Economic Perspectives 34, 99–127. URL: https://www.aeaweb.org/articles?id=10.1257/jep.34.3.99, doi:10.1257/jep.34.3.99.

Belman, D., Wolfson, P.J., 2014. What Does the Minimum Wage Do? Kalamazoo, MI: W.E. Upjohn Institute for Employment. URL: https://research.upjohn.org/up_press/227/.

Besancenot, D., Dubart, D., Vranceanu, R., 2013. The value of lies in an ultimatum game with imperfect information. Journal of Economic Behavior & Organization 93, 239 – 247. URL: http://www.sciencedirect.com/science/article/pii/S0167268113000747, doi:https://doi.org/10.1016/j.jebo.2013.03.029.

Boles, T.L., Croson, R.T., Murnighan, J., 2000. Deception and retribution in repeated ultimatum bargaining. Organizational Behavior and Human Decision Processes 83, 235 – 259. URL: http://www.sciencedirect.com/science/article/pii/S074959780092908X, doi:https://doi.org/10.1006/obhd.2000.2908.

Brachert, M., Dettmann, E., Titze, M., 2019. The regional effects of a place-based policy–causal evidence from germany. Regional Science and Urban Economics 79, 103483.

Burn, I., Button, P., Corella, L.F.M., Neumark, D., 2019. Older Workers Need Not Apply? Ageist Language in Job Ads and Age Discrimination in Hiring. Working Paper 26552. National Bureau of Economic Research. URL: http://www.nber.org/papers/w26552, doi:10.3386/w26552.

Busso, M., Gregory, J., Kline, P., 2013. Assessing the incidence and efficiency of a prominent place based policy. American Economic Review 103, 897–947.

Callaway, B., Sant'Anna, P.H., 2020. Difference-in-differences with multiple time periods. Journal of Econometrics URL: https://www.sciencedirect.com/science/article/pii/S0304407620303948, doi:https://doi.org/10.1016/j.jeconom.2020.12.001.

Chakrabarti, S., Devaraj, S., Patel, P., 2020. Minimum wage and restaurant hygiene violations: Evidence from seattle. Managerial and Decision Economics 42. doi:10.1002/mde.3215.

Chen, D.L., Schonger, M., Wickens, C., 2016. otree—an open-source platform for laboratory, online, and field experiments. Journal of Behavioral and Experimental Finance 9, 88 – 97. URL: http://www.sciencedirect.com/science/article/pii/S2214635016000101, doi:https://doi.org/10.1016/j.jbef.2015.12.001.

Chen, J., Glaeser, E.L., Wessel, D., 2019. The (non-) effect of opportunity zones on housing prices. Technical Report. National Bureau of Economic Research.

Clemens, J., 2021. How do firms respond to minimum wage increases? understanding the relevance of non-employment margins. Journal of Economic Perspectives 35, 51–72. URL: https://www.aeaweb.org/articles?id=10.1257/jep.35.1.51, doi:10.1257/jep.35.1.51.

Clemens, J., Kahn, L.B., Meer, J., 2021. Dropouts need not apply? the minimum wage and skill upgrading. Journal of Labor Economics 39, S107–S149. URL: https://doi.org/10.1086/711490, doi:10.1086/711490, arXiv:https://doi.org/10.1086/711490.

Corinth, K., Feldman, N.E., 2021. A first look at the impact of opportunity zones on commercial investment and economic activity. Available at SSRN 3794396 .

Crain, C., 2018. Price and quality responses of the restaurant industry to increases in the minimum wage. Job Market Paper, Department of Economics, Univeristy of Iowa .

Criscuolo, C., Martin, R., Overman, H.G., Van Reenen, J., 2019. Some causal effects of an industrial policy. American Economic Review 109, 48–85.

Croson, R., Boles, T., Murnighan, J., 2003. Cheap talk in bargaining experiments: lying and threats in ultimatum games. Journal of Economic Behavior & Organization 51, 143 – 159. URL: http://www.sciencedirect.com/science/article/pii/S0167268102000926, doi:https://doi.org/10.1016/S0167-2681(02)00092-6.

Cullen, Z.B., Pakzad-Hurson, B., 2019. Equilibrium effects of pay transparency in a simple labor market. Harvard Business School Working Paper .

Dhami, S., 2016. The Foundations of Behavioral Economic Analysis. Oxford University Press. URL: https://books.google.com/books?id=DFNDjwEACAAJ.

Duarte, J., Umar, T., Yimfor, E., 2021. Rubber stamping opportunity zones. Unpublished Paper .

Dube, A., 2019. Impacts of Minimum Wages: Review of the International Evidence. Technical Report. HM Treasury and Department for Business, Energy & Industrial Strategy.

Eastman, S., Kaeding, N., 2019. Opportunity zones: What we know and what we don't. Tax Foundation Fiscal Fact 630.

White House Council of Economic Advisors, U.S., 2020. The Impact of Opportunity Zones: An Initial Assessment. Technical Report. Prepared for the House Committee on Ways and Means and the Senate Committee on Finance.

Eldar, O., Garber, C., 2020. Does government play favorites? evidence from opportunity zones. Evidence from Opportunity Zones (September 1, 2020). Duke Law School Public Law & Legal Theory Series .

Even, W.E., Macpherson, D.A., 2014. The effect of the tipped minimum wage on employees in the u.s. restaurant industry. Southern Economic Journal 80, 633–655. URL: https://onlinelibrary.wiley.com/doi/abs/10.4284/

0038-4038-2012.283, doi:https://doi.org/10.4284/0038-4038-2012.283, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.4284/0038-4038-2012.283.

Frank, M.M., Hoopes, J.L., Lester, R., 2020. What determines where opportunity knocks? political affiliation in the selection and early effects of opportunity zones, in: 113th Annual Conference on Taxation. NTA.

Freedman, M., 2012. Teaching new markets old tricks: The effects of subsidized investment on low-income neighborhoods. Journal of Public Economics 96, 1000–1014.

Freedman, M., 2015. Place-based programs and the geographic dispersion of employment. Regional Science and Urban Economics 53, 1–19.

Freedman, M., 2017. Persistence in industrial policy impacts: Evidence from depression-era mississippi. Journal of Urban Economics 102, 34–51.

Freedman, M., Neumark, D., Khanna, S., 2021. Combining Rules and Discretion in Economic Development Policy: Evidence on the Impacts of the California Competes Tax Credit. Technical Report. National Bureau of Economic Research.

Garin, A., 2019. Public investment and the spread of 'good-paying'manufacturing jobs: Evidence from World War II's big plants. Technical Report. Working Paper.

Gelfond, H., Looney, A., 2018. Learning from opportunity zones: How to improve place-based policies. Washington, DC: Brookings Institution .

Gentzkow, M., Kelly, B., Taddy, M., 2019. Text as data. Journal of Economic Literature 57, 535–74. URL: https://www.aeaweb.org/articles?id=10.1257/jel.20181020, doi:10.1257/jel.20181020.

Givord, P., Quantin, S., Trevien, C., 2018. A long-term evaluation of the first generation of french urban enterprise zones. Journal of Urban Economics 105, 149–161.

Glaeser, E.L., Gottlieb, J.D., 2008. The economics of place-making policies. Technical Report. National Bureau of Economic Research.

Gobillon, L., Magnac, T., Selod, H., 2012. Do unemployed workers benefit from enterprise zones? the french experience. Journal of Public Economics 96, 881–892.

Government Accountability Office, U.S., 2020. Opportunity Zones: Improved Oversight Needed to Evaluate Tax Expenditure Performance. Technical Report. Report GAO-12-30.

Güth, W., Huck, S., Ockenfels, P., 1996. Two-level ultimatum bargaining with incomplete information: an experimental study. The Economic Journal 106, 593–604. URL: http://www.jstor.org/stable/2235565.

Güth, W., Schmittberger, R., Schwarze, B., 1982. An experimental analysis of ultimatum bargaining. Journal of Economic Behavior & Organization 3, 367 – 388. URL: http://www.sciencedirect.com/science/article/pii/0167268182900117, doi:https://doi.org/10.1016/0167-2681(82)90011-7.

Hall, R.E., Krueger, A.B., 2012. Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. American Economic Journal: Macroeconomics 4, 56–67. URL: http://www.aeaweb.org/articles?id=10.1257/mac.4.4.56, doi:10.1257/mac.4.4.56.

Harasztosi, P., Lindner, A., 2019. Who pays for the minimum wage? American Economic Review 109, 2693–2727. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20171445, doi:10.1257/aer.20171445.

Hennig-Schmidt, H., Irlenbusch, B., Rilke, R.M., Walkowitz, G., 2018. Asymmetric outside options in ultimatum bargaining: a systematic analysis. International Journal of Game Theory 47, 301–329. URL: https://doi.org/10.1007/s00182-017-0588-4, doi:10.1007/s00182-017-0588-4.

Hirsch, B.T., Kaufman, B.E., Zelenska, T., 2015. Minimum wage channels of adjustment. Industrial Relations: A Journal of Economy and Society 54, 199–239. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/irel.12091, doi:10.1111/irel.12091, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/irel.12091.

Internal Revenue Service, U.S., 2020. Opportunity Zone Frequently Asked Questions. Technical Report.

Kagel, J.H., Kim, C., Moser, D., 1996. Fairness in ultimatum games with asymmetric information and asymmetric payoffs. Games and Economic Behavior 13, 100 – 110. URL: http://www.sciencedirect.com/science/article/pii/S0899825696900263, doi:https://doi.org/10.1006/game.1996.0026.

Kriss, P.H., Nagel, R., Weber, R.A., 2013. Implicit vs. explicit deception in ultimatum games with incomplete information. Journal of Economic Behavior & Organization 93, 337 – 346. URL: http://www.sciencedirect.com/science/article/pii/S0167268113000693, doi:https://doi.org/10.1016/j.jebo.2013.03.024.

Lambert, P.J., Aronson, J.R., 1993. Inequality decomposition analysis and the gini coefficient revisited. The Economic Journal 103, 1221–1227. URL: http://www.jstor.org/stable/2234247.

Lester, R., Evans, C., Tian, H., 2018. Opportunity zones: An analysis of the policy's implications. State Tax Notes 90, 19–18.

Luca, D.L., Luca, M., 2019. Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit. Working Paper 25806. National Bureau of Economic Research. URL: http://www.nber.org/papers/w25806, doi:10.3386/w25806.

Luca, M., 2011. Reviews, reputation, and revenue: The case of yelp.com. Harvard Business School Working Paper .

Luca, M., Zervas, G., 2016. Fake it till you make it: Reputation, competition, and yelp review fraud. Management Science 62, 3412–3427. URL: https://doi.org/10.1287/mnsc.2015.2304, doi:10.1287/mnsc.2015.2304, arXiv:https://doi.org/10.1287/mnsc.2015.2304.

Manning, A., 2021. The elusive employment effect of the minimum wage. Journal of Economic Perspectives 35, 3–26. URL: https://www.aeaweb.org/articles?id=10.1257/jep.35.1.3, doi:10.1257/jep.35.1.3.

McNichols, D., 2019. Information and the persistence of the gender wage gap; early evidence from california's salary history ban. SSRN, http://dx.doi.org/10.2139/ssrn.3277664 URL: http://dx.doi.org/10.2139/ssrn.3277664.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J., 2013. Distributed representations of words and phrases and their compositionality, in: Advances in neural information processing systems, pp. 3111–3119.

Moffatt, P., 2015. Experimetrics: Econometrics for Experimental Economics. Palgrave Macmillan. URL: https://books.google.com/books?id=HlkrCwAAQBAJ.

Moretti, E., 2010. Local labor markets. Technical Report. National Bureau of Economic Research.

Murnighan, J.K., 2008. Chapter 50 fairness in ultimatum bargaining, Elsevier. volume 1 of *Handbook of Experimental Economics Results*, pp. 436 – 453. URL: http://www.sciencedirect.com/science/article/pii/S1574072207000509, doi:https://doi.org/10.1016/S1574-0722(07)00050-9.

Neumark, D., Grijalva, D., 2017. The employment effects of state hiring credits. ILR Review 70, 1111–1145.

Neumark, D., Kolko, J., 2010. Do enterprise zones create jobs? evidence from california's enterprise zone program. Journal of Urban Economics 68, 1–19.

Neumark, D., Shirley, P., forthcoming. Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the united states? Industrial Relations: A Journal of Economy and Society n/a. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/irel.12306, doi:https://doi.org/10.1111/irel.12306, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/irel.12306.

Neumark, D., Simpson, H., 2015. Place-based policies, in: Handbook of regional and urban economics. Elsevier. volume 5, pp. 1197–1287.

Neumark, D., Young, T., 2019. Enterprise zones, poverty, and labor market outcomes: Resolving conflicting evidence. Regional Science and Urban Economics 78, 103462.

Patrick, C., 2016. Jobless capital? the role of capital subsidies. Regional Science and Urban Economics 60, 169–179.

Pyatt, G., 1976. On the interpretation and disaggregation of gini coefficients. The Economic Journal 86, 243–255. URL: http://www.jstor.org/stable/2230745.

Rapoport, A., Daniel, T.E., Seale, D.A., 2008. Chapter 62 asymmetric two-person bargaining under incomplete information: Strategic play and adaptive learning, Elsevier. volume 1 of *Handbook of Experimental Economics Results*, pp. 560 – 571. URL: http://www.sciencedirect.com/science/article/pii/S1574072207000625, doi:https://doi.org/10.1016/S1574-0722(07)00062-5.

Reich, M., Hall, P., Jacobs, K., 2005. Living wage policies at the san francisco airport: Impacts on workers and businesses. Industrial Relations: A Journal of Economy and Society 44, 106–138. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.0019-8676.2004.00375.x, doi:https://doi.org/10.1111/j.0019-8676.2004.00375.x, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.0019-8676.2004.00375.x.

Reynolds, C.L., Rohlin, S.M., 2015. The effects of location-based tax policies on the distribution of household income: evidence from the federal empowerment zone program. Journal of Urban Economics 88, 1–15.

Roth, A.E., 1995. Chapter 4 bargaining experiments, Princeton University Press. Handbook of Experimental Economics, pp. 253 – 348.

Sage, A., Langen, M., Van de Minne, A., 2019. Where is the opportunity in opportunity zones? early indicators of the opportunity zone program's impact on commercial property prices. Early Indicators of the Opportunity Zone Program's Impact on Commercial Property Prices (May 1, 2019) .

Schmitt, J., 2015. Explaining the small employment effects of the minimum wage in the united states. Industrial Relations: A Journal of Economy and Society 54, 547–581. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/irel.12106, doi:10.1111/irel.12106, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/irel.12106.

Schmitt, P.M., 2004. On perceptions of fairness: The role of valuations, outside options, and information in ultimatum bargaining games. Experimental Economics 7, 49–73. URL: https://doi.org/10.1023/A:1026210021955, doi:10.1023/A:1026210021955.

Sun, L., Abraham, S., 2020. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics URL: https://www.sciencedirect.com/science/article/pii/S030440762030378X, doi:https://doi.org/10.1016/j.jeconom.2020.09.006.

Tan, Z., 2010. Bounded, efficient and doubly robust estimation with inverse weighting. Biometrika 97, 661–682.

Joint Committee on Taxation, U.S., 2019. Estimates of Federal Tax Expenditures for Fiscal Years 2019-2023. Technical Report. Prepared for the House Committee on Ways and Means and the Senate Committee on Finance.

Theodos, B., Meixell, B., Hedman, C., 2018. Did states maximize their opportunity zone selections? Washington, DC: Urban Institute .

Department of Treasury, U.S., 2018. Irs announce final round of opportunity zone designations. URL: `https://home.treasury.gov/news/press-releases/sm0414`.

Wooldridge, J.M., 2010. Econometric analysis of cross section and panel data. MIT press.

Appendix A

Appendix For Chapter 1

A.1 Appendix: Supplemental Figures, Instructions, and Screenshots

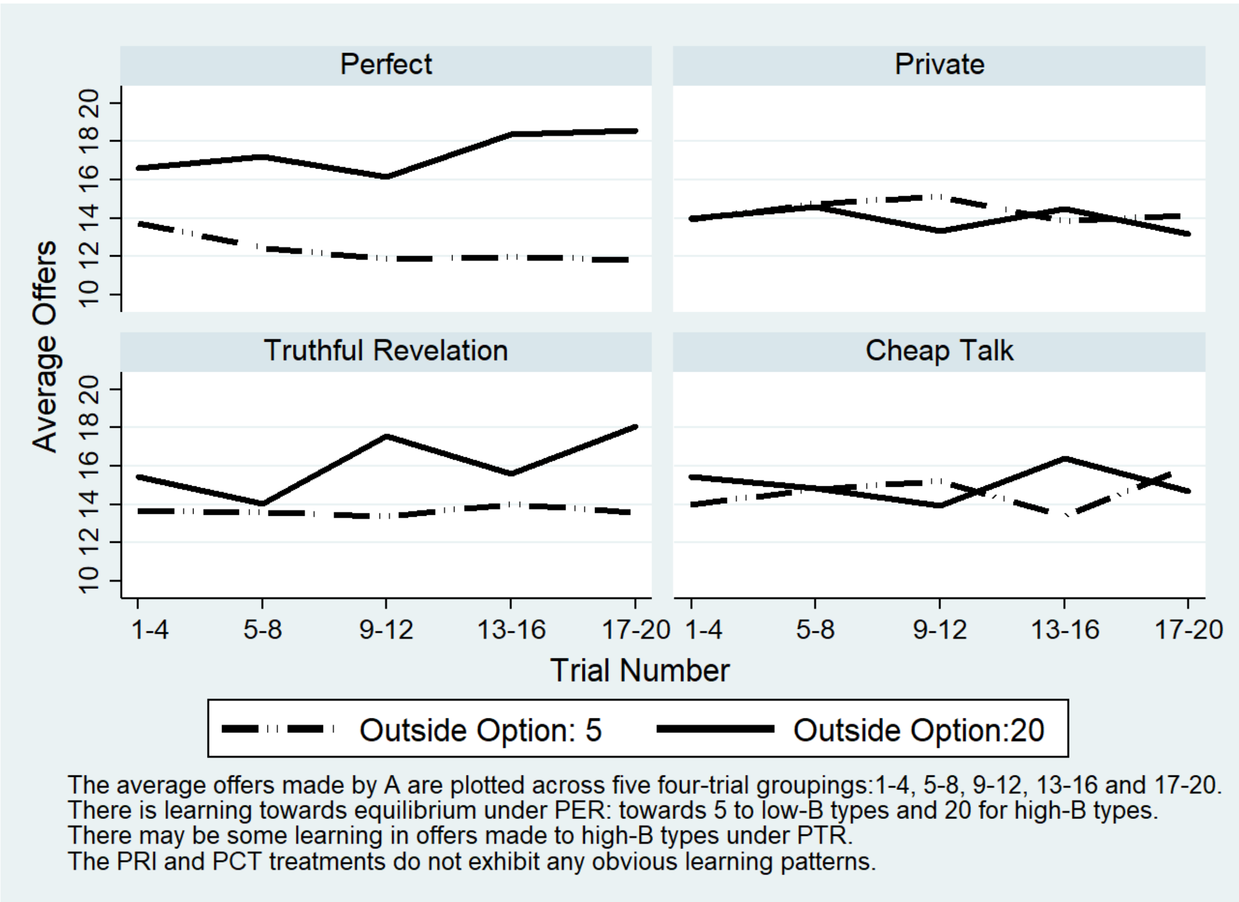

Figure A.1: Average offers across trials, by B's Outside Options

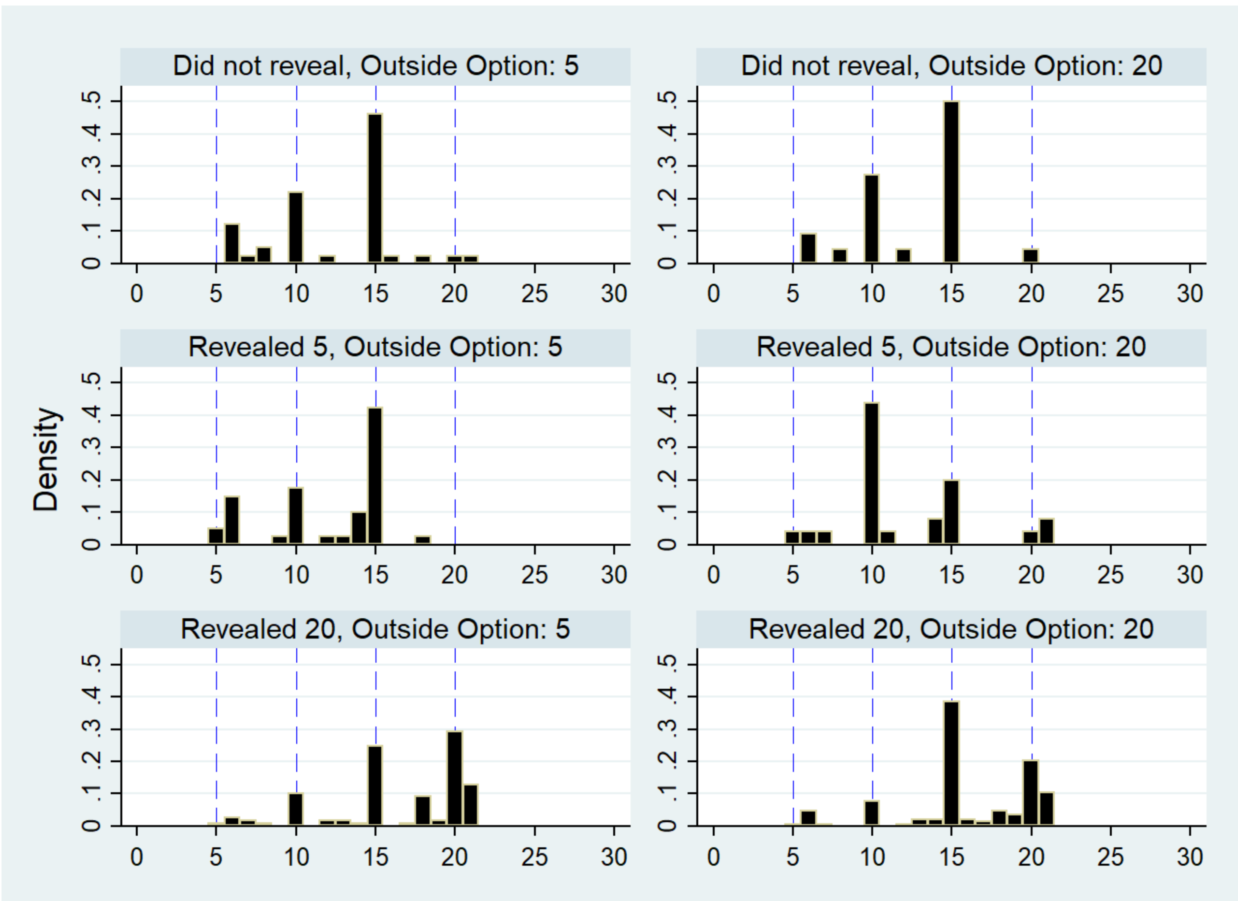

Figure A.2: Cheap Talk (PTR Treatment): Offer Histograms by outside option and revelation decision

Instructions

You will first be randomly assigned to the role of either participant *A* or participant *B* throughout this experiment.

In each trial, an *A* will be matched randomly and anonymously with a *B*.

After matching, Participant *A* makes an offer to *B* between 0 and 30.0 points.

Then, *B* makes a choice to accept or reject this offer.

If *B* **accepts**, then *B* gets the offered amount and *A* keeps the rest **(30 - Offer)**.

If *B* **rejects**, each party gets their **outside option**. *A*'s outside option is always 5. *B*'s outside option is randomly set in each trial as either 5 or 20.

In each trial, both *A* and *B* will be told *B*'s outside option value.

At the start of every trial, all participants will be randomly re-matched each time for a total of **20** identical trials.

Figure A.3: Instructions: Perfect Information (PER)

Instructions

You will first be randomly assigned to the role of either participant *A* or participant *B* throughout this experiment.

In each trial, an *A* will be matched randomly and anonymously with a *B*.

After matching, Participant *A* makes an offer to *B* between 0 and 30.0 points.

Then, *B* makes a choice to accept or reject this offer.

If *B* **accepts**, then *B* gets the offered amount and *A* keeps the rest **(30 - Offer)**.

If B **rejects**, each party gets their **outside option**. *A*'s outside option is always **5**.

B's outside option is randomly set in each trial as either **5** or **20**.

***B*'s exact outside option is known privately only to *B*.**

A only knows that B's outside option is either 5 or 20 with an equal chance.

At the start of every trial, all participants will be randomly re-matched each time for a total of **20** identical trials.

At the end of the experiment, the computer will randomly pick **two** trials for payment. **3 points = $1.**

Figure A.4: Instructions: Private Information (PRI)

Instructions

You will first be randomly assigned to the role of either participant *A* or participant *B* throughout this experiment.

In each trial, an *A* will be matched randomly and anonymously with a *B*.

After matching, Participant *A* makes an offer to *B* between 0 and 30.0 points.

Then, *B* makes a choice to accept or reject this offer.

If *B* **accepts**, then *B* gets the offered amount and *A* keeps the rest **(30 - Offer)**.

If B **rejects**, each party gets their **outside option**. *A*'s outside option is always **5**.

B's outside option is randomly set in each trial as either **5** or **20** (equally likely).

***B*'s exact outside option is known privately only to *B*.**

Before each trial begins, B will decide whether or not to reveal their outside option to A.

For every trial, all participants will be randomly re-matched each time for a total of 20 identical trials.

At the end of the experiment, the computer will randomly pick **two** trials for payment. **3 points = $1.**

Figure A.5: Instructions: Truthful revelation (PTR)

Instructions

You will first be randomly assigned to the role of either participant *A* or participant *B* throughout this experiment.

In each trial, an *A* will be matched randomly and anonymously with a *B*.

After matching, Participant *A* makes an offer to *B* between 0 and 30.0 points.

Then, *B* makes a choice to accept or reject this offer.

If *B* **accepts**, then *B* gets the offered amount and *A* keeps the rest **(30 - Offer)**.

If B **rejects**, each party gets their **outside option**. *A*'s outside option is always **5**.

B's outside option is randomly set in each trial as either **5** or **20** (equally likely).

***B*'s exact outside option is known privately only to *B*.**

Before each trial begins, B will decide to declare some information about their outside option to A.

They can choose ***(1) not to reveal anything, (2) claim it is 5, or (3) claim it is 20.*** *Thus, it is not necessary that B has to reveal their actual outside option.*

For every trial, all participants will be randomly re-matched each time for a total of 20 identical trials.

At the end of the experiment, the computer will randomly pick **two** trials for payment. **3 points = $1.**

Figure A.6: Instructions: Cheap-Talk (PCT)

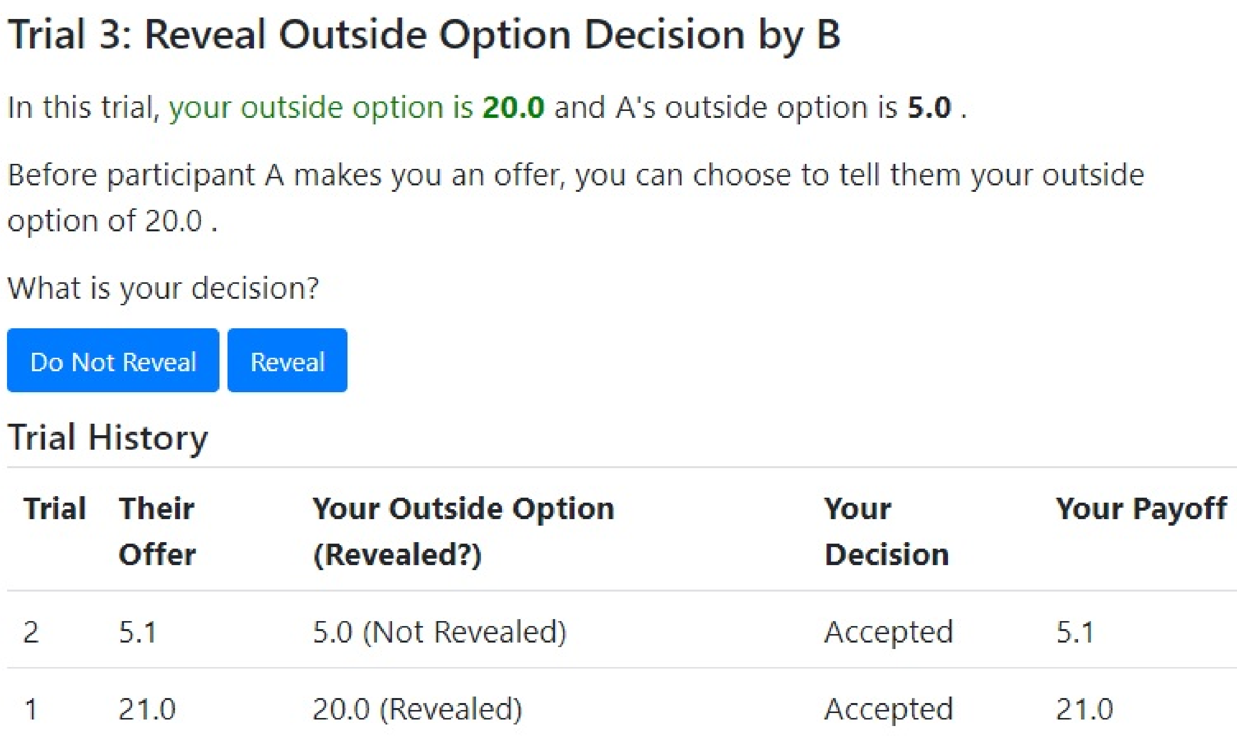

Trial 3: Reveal Outside Option Decision by B

In this trial, your outside option is **20.0** and A's outside option is **5.0** .

Before participant A makes you an offer, you can choose to tell them your outside option of 20.0 .

What is your decision?

Do Not Reveal | Reveal

Trial History

Trial	Their Offer	Your Outside Option (Revealed?)	Your Decision	Your Payoff
2	5.1	5.0 (Not Revealed)	Accepted	5.1
1	21.0	20.0 (Revealed)	Accepted	21.0

Figure A.7: Truthful Revelation treatment Reveal Decision Screenshot

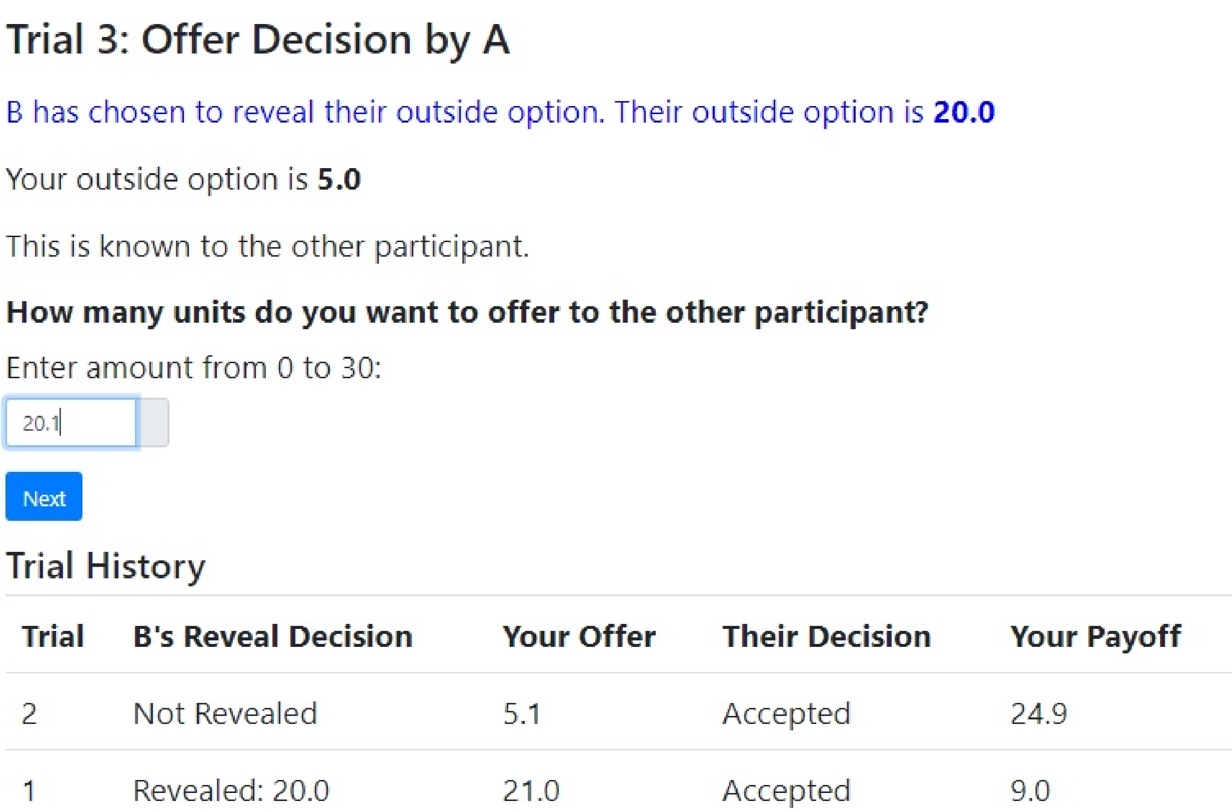

Trial 3: Offer Decision by A

B has chosen to reveal their outside option. Their outside option is **20.0**

Your outside option is **5.0**

This is known to the other participant.

How many units do you want to offer to the other participant?

Enter amount from 0 to 30:

20.1

Next

Trial History

Trial	B's Reveal Decision	Your Offer	Their Decision	Your Payoff
2	Not Revealed	5.1	Accepted	24.9
1	Revealed: 20.0	21.0	Accepted	9.0

Figure A.8: Truthful Revelation treatment Offer Decision Screenshot

Trial 3: Accept or Reject Decision by B

The other participant has offered you 20.1 .

Recall that in this trial your outside option is **20.0** and A's outside option is **5.0**

You have already revealed your outside option to participant A.

What is your decision?

ACCEPT REJECT

Trial History

Trial	Their Offer	Your Outside Option (Revealed?)	Your Decision	Your Payoff
2	5.1	5.0 (Not Revealed)	Accepted	5.1
1	21.0	20.0 (Revealed)	Accepted	21.0

Figure A.9: Truthful Revelation treatment Accept Decision Screenshot

Cognitive Reflection Test (CRT) Questions

Q1) A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost?

Q2) If it takes 5 machines 5 minutes to make 5 widgets, how many minutes would it take 100 machines to make 100 widgets?

Q3) In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how many days would it take for the patch to cover half of the lake?

Figure A.10: CRT Questions

Appendix B

Appendix For Chapter 2

Table B.1: Ratings Impact for Other (Placebo) Establishments: DID Specification

	(1)	(2)	(3)
Post X Treatment	0.008	-0.002	-0.027
	(0.015)	(0.009)	(0.025)
Constant	3.697***	3.685***	3.285***
	(0.024)	(0.002)	(0.005)
Month FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Business FE	No	Yes	No
Bus-User FE	No	No	Yes
N	2732913	2732310	206730

AZ treated, NV IL NC SC PA WI TX GA Controls. SE clustered at zipcode.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.2: Ratings Impact for Other (Placebo) Establishments: Panel Regressions

	(1)	(2)
Min Wage	-0.003	0.007
	(0.003)	(0.011)
Constant	3.686***	3.238***
	(0.025)	(0.097)
N	4628958	336064
R-Square	0.33	0.84
Adj. R-Square	0.30	0.70

Other (non-food) reviews for all 18 metropolitan areas. SE clustered at zipcode.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.3: Panel Regressions of Log Prices on Log of MW for Other (Placebo) Establishments

	(1)	(2)
Ln(MW)	-0.023	-0.086
	(0.037)	(0.287)
Constant	3.029***	3.222***
	(0.077)	(0.610)
N	320502	11764
R-Square	0.40	0.88
Adj. R-Square	0.29	0.76

Sample restricted to other (non-food establishment) reviews where dollar prices mentioned explicity.
SE clustered at zipcode.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

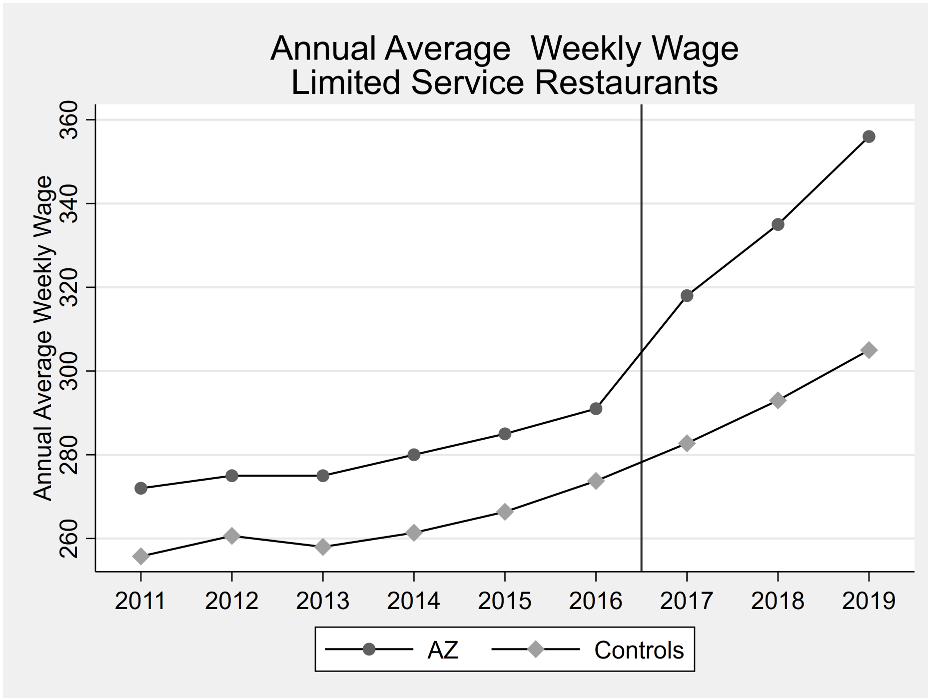

(a) Limited Service

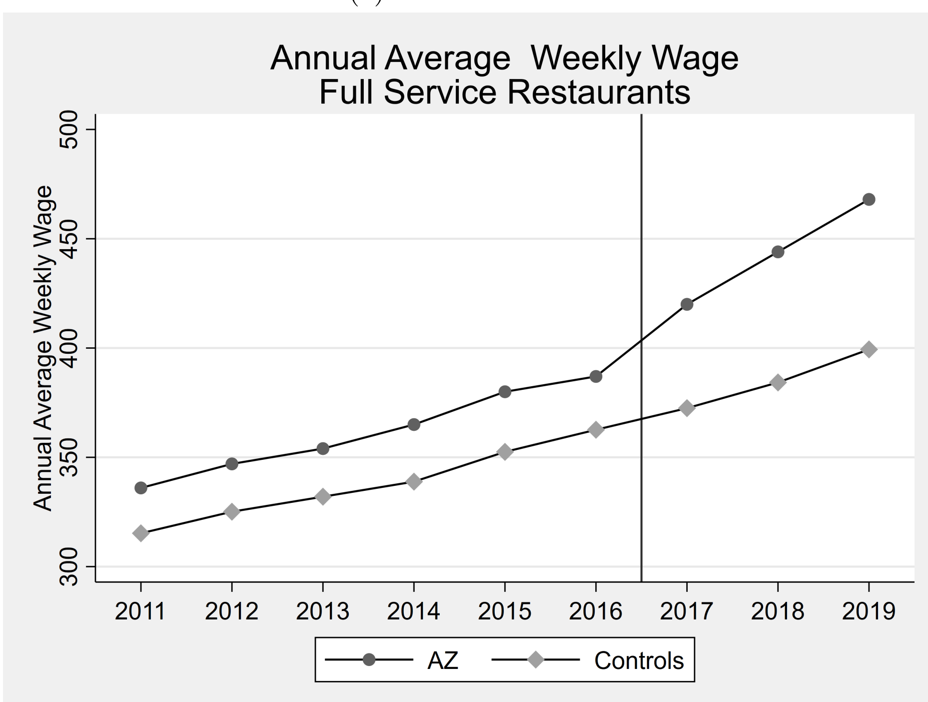

(b) Full Service

Figure B.1: Annual Average Wages in Restaurants (QCEW), Treated (AZ) and Control States

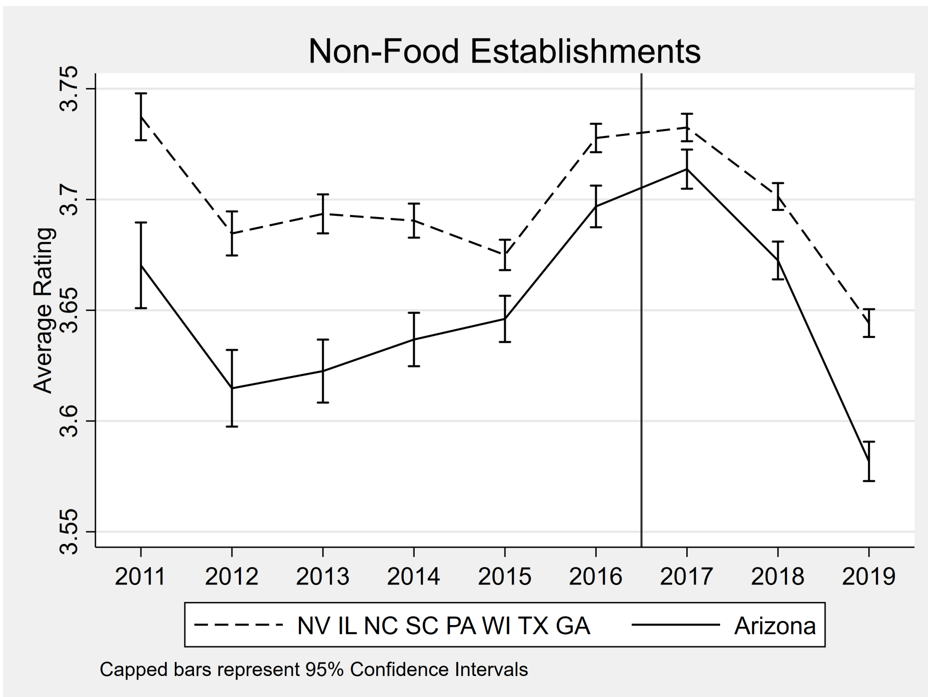

Figure B.2: Annual average ratings for placebo establishment reviews in Phoenix, AZ and control, 2011-2019

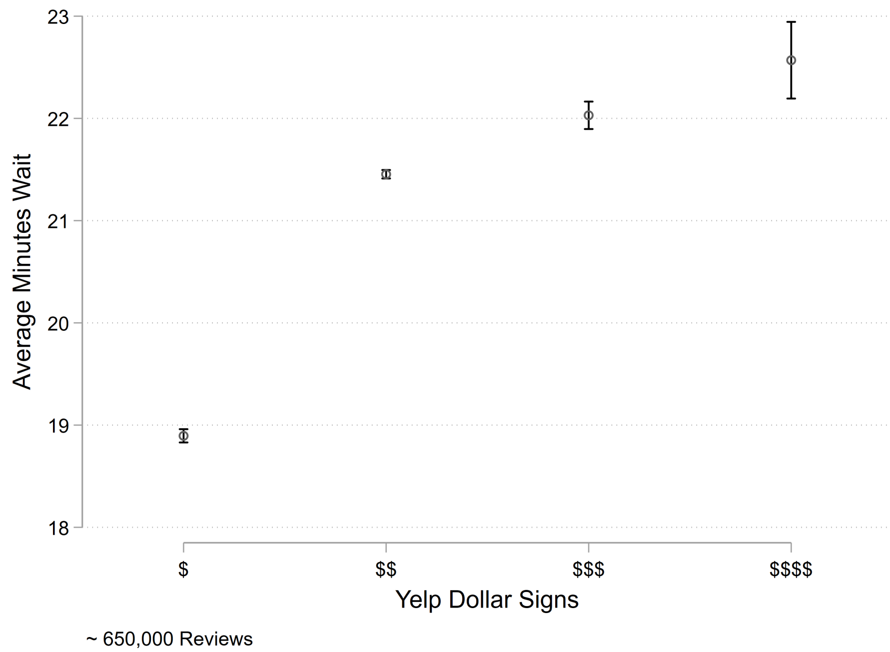

Figure B.3: Average Extracted minute wait times by Yelp “$” Signs

Appendix C

Appendix For Chapter 3

Table C.1: Summary Statistics for the Inverse Probability Weights Assigned to the Control Tracts

	IPW Weights
Mean	0.3294
Std. Dev.	0.1991
Skewness	5.42
Kurtosis	68.12
Control Tracts	23000

Table C.2: IPW and Regression-Adjusted IPW Estimates using Opportunity Zone Post-Treatment Period Defined as 2019

	(1)	(2)	(3)
	Employment Rate	Average Earnings	Poverty Rate
	A. IPW Treatment on the Treated Estimates		
Opportunity Zone	-0.0008	236.6	0.0045**
	(0.0016)	(162.9)	(0.0022)
Tracts	30600	30600	30600
	B. Regression-Adj. IPW Treatment on the Treated Estimates		
Opportunity Zone	-0.0004	213.9	0.0030
	(0.0016)	(162.1)	(0.0021)
Tracts	30600	30600	30600

Notes: Data derived from the 2013-2019 American Community Surveys. The Opportunity Zone variable takes a value of one for designated tracts in 2019. The pre-treatment period is defined as 2013-2018. Panel A reports inverse probability weighted (IPW) estimates of the treatment on the treated. Panel B reports doubly robust regression-adjusted IPW estimates of the treatment on the treated. Standard errors in parentheses are heteroskedasticity robust. Statistically significant at * 10%, ** 5%, *** 1%.

Table C.3: Unadjusted Difference-in-Differences Estimates of the Effects of Opportunity Zones on Residents

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment Rate		Average Earnings		Poverty Rate	
	A. Post-Treatment Period Defined as 2018-2019					
Opportunity Zone	0.0052***	0.0071***	129.3	227*	-0.0090***	-0.0113***
	(0.0012)	(0.0012)	(121.8)	(129.7)	(0.0015)	(0.0016)
	B. Post-Treatment Period Defined as 2019					
Opportunity Zone	0.0051***	0.0069***	-61.32	69.58	-0.0092***	-0.0121***
	(0.0015)	(0.0016)	(159)	(170.8)	(0.0019)	(0.0021)
Tract Fixed Effects	Y	Y	Y	Y	Y	Y
Year Fixed Effects	Y		Y		Y	
County-by-Year Fixed Effects		Y		Y		Y
Tracts	30600	30600	30600	30600	30600	30600
Observations	214200	214200	214200	214200	214200	214200

Notes: Data derived from the 2013-2019 American Community Surveys. Panel A reports difference-in-differences estimates with the post-treatment period defined as 2018-2019. Panel B reports difference-in-differences estimates with the post-treatment period defined as only 2019. The odd-numbered columns show estimates from models with tract and year fixed effects. The even-numbered columns show results from models with tract and county-by-year fixed effects. Standard errors in parentheses are heteroskedasticity robust and clustered at the tract level. Statistically significant at * 10%, ** 5%, *** 1%.