

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays in Microeconomics

Permalink

<https://escholarship.org/uc/item/33w1h1nf>

Author

Donkor, Kwabena Baah

Publication Date

2020

Peer reviewed|Thesis/dissertation

Essays in Microeconomics

by

Kwabena B. Donkor

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jeffrey M. Perloff, Chair
Associate Professor Aprajit Mahajan
Professor Stefano DellaVigna

Spring 2020

Essays in Microeconomics

Copyright 2020
by
Kwabena B. Donkor

Abstract
Essays in Microeconomics

By

Kwabena B. Donkor

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Jeffrey M. Perloff, Chair

A person's behavior is influenced by tradeoffs between (1) personal preferences versus social expectations, and (2) choosing a menu or default option versus actively computing a preferred choice. To understand the mechanisms behind the decision-process, one must understand the empirical relationship between decision-makers' preferences and their sensitivity to social expectations and menu or default options. The first part of this thesis uses both a revealed preference approach and a model that empirically identifies the stated tradeoffs to examine tipping in New York City (NYC) taxi cabs. It uses the estimated model to analyze the welfare implications of presenting menus or default options.

The second part uses insights from the first part to examine why individuals in the less wealthy neighborhoods of NYC give cab drivers more generous gratuities than their counterparts in wealthier areas.

The final part examines how the Affordable Care Act (ACA) affected seasonal farmworkers' choice of health insurance coverage, medical services utilization, and jobs with employer-provided benefits.

The first chapter asks the question: Does a menu of recommended tips presented with a bill influence how much customers tip? The answer to this question depends on customers' ideal tips, how much customers are affected by their beliefs about a socially acceptable tip, and the effort costs of computing a tip versus choosing from a menu.

First, changes in the menu presented to passengers in NYC taxis are used to nonparametrically estimate that the cost of actively computing a tip and not following a menu is about \$1.89 (15.53% of the average taxi fare of \$12.17). Second, a model is used to explain the mechanisms behind the decision-process. In the model, passengers' tipping choices depend on their perception of a socially acceptable tip (social norm tip), the shame from given less (norm deviation cost), and the difficulty of calculating a tip (cognitive cost). An estimate of the distribution of beliefs about the social norm tip averages about 20% of the taxi fare. Customers incur a norm deviation cost of tipping five percentage points less than the norm of between \$0.30 and \$0.38. The cognitive cost of calculating a non-menu tip ranges from \$1.10 to \$1.32 on average.

The model predicts that taxicabs currently present customers with a nearly tip-maximizing menu, and this menu increases tips by 14.65% relative to not presenting a menu. Taxicab companies appear to have learned over time to converge to the tip-maximizing menu. Welfare calculations suggest that the current tip menu in NYC cabs increases overall welfare by \$1.08 per taxi trip relative to presenting no menu.

The second chapter documents that customers in low-income neighborhoods of NYC give more generous tips than their counterparts in wealthier areas. Several studies suggest similar counterintuitive findings. However, it is difficult to establish causality in these findings because wealth cannot be easily randomized. This chapter relies on insights and the model from the first chapter to explain the relationship between income and gratuity. The findings suggest that the distribution of beliefs about the social norm tip in the wealthiest areas averages about 20.57% of the taxi fare compared to 29.78% in the poorest areas. In contrast, the norm deviation cost is higher in wealthier neighborhoods. For example, a five-percentage point deviation from the norm in the wealthiest areas is almost twice the cost in the poorest areas (\$0.26 versus \$0.14). The cognitive cost of calculating a tip is lower in poorer areas compared to richer areas—a reflection of a higher opportunity cost of time in wealthier neighborhoods.

The third chapter (coauthored with Susan Gabbard and Jeffrey Perloff) studies how individual ACA policies affected seasonal farmworkers’ choice of health insurance coverage, medical services utilization, and jobs with employer-provided benefits. Seasonal agricultural workers are an important target population for the ACA. These workers have low incomes, have relatively little health insurance coverage, and face many job-related health risks. The chapter concentrates on four ACA policies that were likely to affect farmworkers. First, Medicaid expansion. Second, health insurance premium subsidies. Third, an individual mandate to maintain health insurance coverage. Fourth, the ACA prohibited insurance companies from setting insurance policy prices based on pre-existing health conditions.

The findings suggest that the ACA decreased the use of employer-provided coverage and substantially increased the share of workers with government insurance. Workers with pre-existing health conditions consumed more medical services, relative to those without pre-existing health conditions. The ACA did not reduce emergency room visits. However, it induced workers with pre-existing conditions to make greater use of private doctors and private clinics, hospitals, and community health centers.

Menus, default options, and social norms are important phenomena that deeply guide human behavior. However, in a field setting, the preferences of decision-makers and their perceptions of norms are difficult to observe, quantify, and study scientifically. The first two chapters of this thesis empirically identify unobserved consumer preferences to gain a better understanding of how menus, defaults, and norms influence consumer tipping choices in NYC taxi cabs. The findings are useful in preference identification, and considering more general “nudges,” such as those that are widely used by businesses and policy makers. The ACA law targeted people like farmworkers who are low-paid, traditionally had low rates of coverage, and suffered from more health challenges than most other workers. The third chapter is the only study to examine the effects of the ACA on farmworkers’ use of medical services, where they go for treatment, and whether their job provides health and non-health benefits.

Dedication

I dedicate this dissertation to my mother Mrs. Mary Attah Kofie Donkor, who inspired and taught me from an earlier age to put God first, work hard, pursue excellence, and provide myself no excuse not to pursue higher education.

Contents

1	How Difficult is Tipping? Nonparametric and Parametric Estimates of Decision Costs	1
1.1	Introduction	1
1.2	Taxi Tipping Systems and Data	5
1.3	The Structural Model	7
1.4	Nonparametric Estimation of Decision Costs	8
1.5	Semiparametric Estimation of Decision Costs	11
1.6	Parametric Estimation of Decision Costs	13
1.6.1	The Parametric Model	13
1.6.2	Estimation Procedure	15
1.6.2.1	Estimating Tipping Norms and Norm Deviation Costs . . .	15
1.6.2.2	An Upper Bound of Cognitive Costs for Non-Menu Tips . .	18
1.6.2.3	Minimum Distance Estimation of Cognitive Costs	19
1.6.3	Model Performance	21
1.7	Analysis of Menu Tips	21
1.7.1	Tip-Maximizing Menu	22
1.7.2	Evolution of Menu Tip Options	23
1.7.3	Welfare	23
1.8	Conclusions	25
2	Lower Income, Higher Gratuity: Evidence from Tipping in NYC Taxis	35
2.1	Introduction	35
2.2	Context and Data	37
2.2.1	Identifying Low Income and Wealthy Passengers	39
2.2.2	NYC Income Disparities	39
2.2.3	Trip Characteristics	39
2.3	Income and Tipping within NYC Boroughs	40
2.4	Possible Reasons the Less Wealthy Tip More	41
2.5	Model	42
2.5.1	Estimating Tipping Norms and Norm Deviation Costs	43
2.5.1.1	The Heckman Selection Model	43
2.5.1.2	The Effect of Income on Beliefs about Tipping Norms . . .	44
2.5.2	Estimating Cognitive Cost	45
2.5.3	Model Fit	46
2.6	Conclusions	46

3	The Effect of the Affordable Care Act on Agricultural Workers	56
3.1	Introduction	56
3.2	Data	58
3.3	Specification	59
3.4	ACA Effects on Health Insurance Coverage	61
3.5	ACA Effects on the Use of Medical Services	63
3.6	Employer-Provided Benefits	65
3.7	Conclusions	66
	Bibliography	75
A	Appendices for Chapters 1 and 2	82
A.1	New York City Yellow Taxi Tipping Systems	82
A.2	Empirical Support for Model Assumptions	83
A.3	Figures for Nonparametric Approach	88
	A.3.1 Tips Before and After CMT Tip Menu Change by Fare	88
	A.3.2 Bounds On the Conditional CDFs of Decision Costs	94
A.4	Figures for Semiparametric Approach	98
A.5	Empirical Support for Parametric Approach	103
A.6	Figures for Welfare Calculations	104

List of Figures

1.1	(A) Distribution of Tips (% of Taxi Fare) Before and After Menu Change, (B) Nonparametric Bounds of CDF of Decision Cost (Tip = 10%), (C) Nonparametric Unconditional CDF of Decision Costs	26
1.2	(A) Predicted Probability of Tipping 10% (B) Predicted Probability of Tipping 20% (C) Semiparametric Bounds of CDF of Decision Cost (Tip = 10%), (D) Semiparametric and Nonparametric Unconditional CDF of Decision Costs	27
1.3	Distribution of Beliefs about Tipping Norm	28
1.4	CDF of (1) Upper Bound of Cognitive Costs and (2) The Norm Deviation Cost for Non-Menu Tips	28
1.5	Model Fit	29
1.6	Grid Search for Tip-Maximizing Menu	30
2.1	Median Household Incomes in New York City	48
2.2	Tipping Rates in New York City	49
2.3	Distribution of Beliefs About Tipping Norms & Preferred Non-Menu Tips	50
2.4	Model Fit	51
3.1	Farm Worker's Health Insurance Coverage	67
3.2	Farm Worker's Medical Services Use	68
A.1	NYC Yellow Taxi Payment Screen with Menu Tip Options	82
A.2	Changes in Menu Tip Options Over Time by Vendor	83
A.3	Overlapping Distribution of Tip (%) in Years with No Change in CMT Tip Menu	85
A.4	Distribution of Tips < 20%	88
A.5	Distribution of Tip (%) less than 20% Before and After CMT Tip Menu Change	89
A.6	Conditional CDFs of Decision Costs	95
A.7	Predicted Probabilities by Level of Taxi Fare of Choosing Tips < 20%	99
A.8	Conditional CDFs of Decision Costs	100
A.9	Share of Menu Tips and Average Tip by Fare	103
A.10	Distribution of Tip (%) by Type of Tip Menu	105

List of Tables

1.1	Taxi Trip Characteristics	31
1.2	Structural Estimates	32
1.3	Evolution of Menu Tip Options	33
1.4	Welfare Estimates (Trip Level)	34
2.1	Taxi Trip Characteristics: Share or Mean (Standard Deviation)	52
2.2	OLS: Tipping in NYC Boroughs and by Income	53
2.3	Probit: First Step Heckman Selection Model	54
2.4	Structural Estimates	55
3.1	Eligibility and Tax Penalty	68
3.2	Summary Statistics by Farmworker's Legal Status and Period. Share or Mean (Standard Deviation)	69
3.3	Multinomial Logit Estimates, Health Insurance	70
3.4	Marginal Effects, Health Insurance	71
3.5	Multinomial Logit Estimates, Use of Medical Services	72
3.6	Marginal Effects, Use of Medical Services	73
3.7	Linear Probability Models of Employers' Provision of Fringe Benefits	74
A.1	Evidence for Assumption A3	87
A.2	Heckman Selection Correction Estimates: First-Step Probit Estimates	103
A.3	OLS Estimates of Tipping Norms and Norm Deviation Cost Parameter	104

Acknowledgements

First and foremost, I thank God for guiding and sustaining me throughout my education.

To my beloved wife Hanan Wasse, you have been my backbone through this journey. Thank you for your compassion, love and support. Thank you for reading through all my half-baked drafts, correcting typos, and giving me feedback. You are indeed a blessing I do not deserve. May God richly bless you.

Professor Jeffrey M. Perloff, you have been a great friend, father figure, mentor, and counselor throughout my time at UC Berkeley. I am deeply indebted to you for the grace you have extended towards me over the years: always looking out for me and for supporting me. I will miss the countless hours we spent discussing ideas, funny stories, and sharing our individual taxi experiences. I look forward to continuing this incredible relationship we have.

Professor Stefano DellaVigna, thank you for teaching me almost all that I know about behavioral economics. Thank you for the countless hours you spent in office hours discussing ideas, guiding me, and giving me feedback. I could not have come this far without your counsel.

Professor Aprajit Mahajan, thank you for all the countless hours we spent in your office working through mathematical models, discussing ideas, and econometrics. I am especially grateful for your advice, support, chairing my Orals committee, and all the feedback for my research and the job market process.

To Sofia Villas-Boas, Miguel Villas-Boas, Ben Handel, Dmitry Taubinsky, Elisabeth Sadoulet, Maximilian Auffhammer, Meredith Fowlie, and Catherine Wolfram thank you for your invaluable advice and guidance. Thank you to all the students and faculty at UC Berkeley's Agricultural and Resource Economics Department, Haas Marketing seminar, Industrial Organization seminar, Psychology & Economics lunch seminar, for your support and contributions.

I am deeply thankful and grateful to my coauthors Susan Gabbard, at JBS International, CA, and Jeffrey Perloff at UC Berkeley. I thank Daniel J. Carroll, at the Department of Labor, who helped to obtain the National Agricultural Workers Survey dataset and made many useful comments. I am also grateful to the Institute for Research on Labor and Employment for their generous dissertation fellowship and the Giannini Foundation for their research grants.

Most importantly, to all my family, I am grateful for you and appreciate all your support, love, and encouragement. I could not have come this far without learning from you all.

Chapter 1

How Difficult is Tipping? Nonparametric and Parametric Estimates of Decision Costs

1.1 Introduction

Not much is known about why defaults are so powerful in influencing consumer choices. As acknowledged by Bernheim et al. (2015), “costly decision making is notoriously difficult to model” due to the decision processes and potential mechanisms involved when defaults are presented to consumers. This study begins to address this gap in the literature using the large and newly available data from New York City (NYC) Yellow taxis combined with new econometric techniques.

Over the past decade, the introduction of new touch-screen payment technologies has influenced commercial transactions and tipping practices, increasing the revenue potential of several businesses. These touch-screen payment systems present consumers with a menu of default tips as well as options to leave a custom tip or no tip. This technology is used in all NYC Yellow taxicabs. When a customer pays for a NYC Yellow taxicab trip with a credit card, a screen shows the fare. Also, the screen suggests three possible tip rates, and provides the option of giving a non-menu tip (or no tip) instead. We use a new model to measure the effects of presenting customers with such a menu. In our model, customers have beliefs about the social norm of tipping, incur a norm deviation cost for not conforming to the norm, and incur a cognitive cost if they calculate a non-menu tip. Using the model, we estimate the distribution of beliefs about the unobserved social norm tip, the norm deviation cost, and the cognitive cost. In addition, we estimate the tip-maximizing menu and evaluate its implication for consumer utility and overall societal welfare.

We use data from about three quarters of a billion trips in NYC Yellow taxicabs over six years. There are two key sources of variation in our taxicab data. First, changes in the menu of tip options across years provide variation in both the share of passengers who opt for non-menu tips and the amount of tips received by taxi drivers. Second, Yellow taxis use two different credit card technologies with different menus in some years.

Although tipping is not obligatory, most consumers conform to this custom of paying

extra in addition to their bill. However, determining how much to give as a tip involves costly effort (cognitive cost). When a menu of tip suggestions is provided, consumers can avoid the cognitive cost of computing their preferred tip by choosing a suggested tip. Thus, passengers who choose menu options help to identify the cognitive cost of computing non-menu tips.

Reasoning from the fact that tipping is not compulsory and requires costly effort, it may be best to choose from a menu to avoid the cost of computation or choose not to tip at all. However, 38% of the observed tips in our analysis sample give a non-menu tip.¹ This group finds it beneficial to conform to the custom of tipping in addition to incurring the effort cost of computing their preferred non-menu tip. Therefore, passengers who choose non-menu tips help to identify the cost (“shame”) of not conforming to the tipping norm (norm deviation cost).

We use three approaches to assess the effect of using menus on consumer choices. The first is nonparametric. For this approach, we use how changes in menu tip options alter passengers’ tipping choices to identify bounds on the decision cost (norm deviation cost + cognitive cost) of switching from a menu option to choosing a preferred non-menu tip. The second approach is semi-parametric and akin to the nonparametric approach. The difference is that, the semi-parametric approach involves making some parametric assumptions to account for observable trip characteristics: which are not controlled for in the nonparametric approach. The first two approaches do not allow us to decompose decision costs into norm deviation and cognitive costs. Therefore, we result to a third approach, parametric, that allows us to separately identify the norm deviation and cognitive cost components. We do this by placing parametric assumption on our tipping behavior model. This allows us to estimate primitives of the model including passengers’ beliefs about the tipping norm. An added advantage of the parametric model is that it allows us to perform several counterfactual exercises.

Because social norms are important phenomena that deeply guide human behavior, an empirical analysis may help us gain a better understanding of how norms influence consumer choices. However, in a field setting, norms are difficult to quantify and study in a scientific manner. Our model remains agnostic about the underlying model that generates people’s beliefs about the norms of how much to give as a tip. Instead, the model enables us to empirically estimate these unobserved beliefs and analyze the cost individuals face for not conforming to them. Estimates of the cost of deviating from a norm separate from the cognitive cost of computing one’s final choice are interesting in and of themselves. Furthermore, these two components of decision costs individually inform policy. For example, a menu that will maximize either welfare or the profits of a firm depends on whether the menu options reflects the preferences of consumers and minimizes the cost of computation. Therefore, without separating decision costs into norm deviation and cognitive costs, such an exercise will be challenging.

From the nonparametric approach, we estimate that the decision cost associated with tipping averages \$1.89 (15.53% of the average taxi fare of \$12.17). After controlling for trip characteristics via the semi-parametric approach, the average decision cost decreases to

¹Sixty percent of passengers choose menu options and about two percent choose not to tip. Cash tips are not observed in the data and therefore not part of our analysis sample. See Hoover (2019) for more information on the differences between trips paid for with cash and credit cards.

\$1.64 (13.48% of the average taxi fare). From our parametric model, we estimate that the unobserved distribution of beliefs about the tipping norm averages at about 20% of the taxi fare, which is around the average tip rate in the data (19%). The norm deviation cost is large relative to the fare. For instance, a passenger who decides to tip five percentage points less than her belief about the norm incurs a norm deviation cost between \$0.30 and \$0.38 (2.5% - 3.1% of the average fare). We estimate that the average cognitive cost of computing a non-menu tip is between \$1.10 and \$1.32 (9% - 10.8% of the average taxi fare).

Finally, we use estimates from the model in counterfactual exercises to find the tip-maximizing menu and evaluate its implications for social welfare. We assess how the overall welfare from using the tip-maximizing menu compares to the case where consumers are not offered a menu, and for the case where consumers are presented with a menu that maximizes their utility. According to the counterfactual exercises, the tip-maximizing menu increases tips from 15.83% to 18.15% of the taxi fare (i.e., a 14.65% increase in the tip rate). We find that the current menu of tips in taxicabs nearly maximizes the tips received by drivers. However, that was not always the case. It took a few years of trying various menus before settling on using the current menu. The companies appear to have learned over time to converge to this menu. In our welfare calculations, the current tip menu increases the tips received by drivers and the utility of passengers. All else equal, overall welfare increases by \$1.08 per taxi trip relative to not presenting a menu.

Tipping is a major economic activity. According to Shierholz et al. (2017), annual tips from restaurants alone are \$37 billion (about 5% of the 2019 projected sales in restaurants). For most workers in the hospitality industry tips are about 20% of their income, and over 50% for those who earn a tipped wage. In 2007, the NYC Yellow taxis began the practice of presenting customers with a tip menu (Grynbaum, 2009). In 2009, the tech company Square started providing different establishments with electronic credit card readers that prompt customers to choose from a menu of tips. Square has since popularized this technology by making these electronic devices accessible to both small local businesses and large corporations around the United States.² Anecdotes suggest that tip menus compel consumers to tip and increase the amount tipped. Our findings are consistent with these claims.³ Although this study uses taxicab data, it has wide implication as similar menus are widely used in many other industries as well.

An extensive literature discusses how and why menu suggestions and default options affect consumer choice behavior. According to Thaler and Sunstein (2003), defaults and menus should have little to no effect on choices if consumers are fully rational. However, over the past two decades, a plethora of empirical evidence has shown that defaults affect consumers' behavior. For example, defaults affect (1) savings behavior: Madrian and Shea

²For example, the café chain Starbucks agreed in 2012 to invest \$25 million in Square and converted all its electronic cash registers to the ones offered by Square (Cohan, 2012). The grocery chain Whole Foods Market followed suit and announced in 2014 that it would roll out Square registers across some of its stores (Ravindranath, 2014).

³According to a New York Times article, the tips that taxi drivers receive doubled after the installation of electronic devices that present passengers with a menu (Grynbaum, 2009). Fast Company reported that some companies who changed to using Square registers saw about a 40% to 45% increase in customer tips, and that Square is on target to accrue about a quarter of a billion dollars annually for its clients from customer tips alone (Carr, 2013).

(2001); Choi et al. (2002, 2004); Carroll et al. (2009); DellaVigna (2009); Beshears et al. (2009); Blumenstock et al. (2018), (2) organ donations: Johnson and Goldstein (2003); Abadie and Gay (2006), (3) health insurance contracts: Handel (2013), (4) contract choice in health clubs: DellaVigna and Malmendier (2006), (5) tipping behavior: Haggag and Paci (2014), (6) marketing: Brown and Krishna (2004); Johnson et al. (2002), and (7) electricity consumption: Fowlie et al. (2017).

There are several explanations for the default effect. An important one is that, some consumers procrastinate on making important decisions or choices if the benefits of such actions are not immediate. Such consumers would rather opt for a menu or default option in the interim and defer active decision-making to some future date instead (O’Donoghue and Rabin, 1999, 2001). Consumers may perceive a default or menu options as a source of information indicating how to make choices, or which choices are the status quo (Beshears et al., 2009). Thus, they may find it unsettling to choose a different option. Other mechanisms involved when defaults are presented to consumers include the endorsement effect, social norms, calculation and switching costs, et cetera. Nevertheless, empirical estimates from the field on the economic importance of the mechanisms that drive the default and menu effect is limited (Jachimowicz et al., 2019). This is because, it is challenging to model and assess the decision processes and mechanisms involved when consumers are presented with default options.

Passenger tipping decisions from NYC Yellow taxicab trips provide several advantages for assessing these explanations. First, consumers cannot defer tipping to a later date; thus, self-control problems (e.g., naïveté, present bias, and procrastination) are ruled out as explanations for the default effect in this context. Another advantage is that different menus were offered to passengers over the period of this study, enabling us to assess how passenger tipping changed across menus and which menu extracted the most tips for taxi drivers.

There are two studies in the same context as ours. First, Haggag and Paci (2014) use a regression discontinuity design to explore whether menus with higher default tip amounts induce consumers to tip more. They use NYC Yellow taxi trips from 2009 for their analysis. They find that higher tip suggestions increase the amount tipped but may cause some passengers to avoid tipping altogether. Second, Thakral and Tô (2019) use a change in the NYC Yellow taxi fare rate in 2012 to evaluate the dynamics and welfare consequences of adherence to the social norm of tipping. In contrast to these studies, this paper provides the first comprehensive model that identifies and evaluate the mechanisms involved in the decision process when passengers are presented with a default tip menu. In addition, this study estimates passengers unobserved beliefs about the social norm tip, the cost of deviating from the norm, the tip-maximizing menu, and evaluates the implications of different tip menus for societal welfare.

This study contributes to the literature on behavioral industrial organization by measuring how consumer-switching costs affect firm profits. That is, the consumer faces costs when switching from one option to another. Beggs and Klemperer (1992), use a threshold model to show that competitive firms have an incentive to exploit switching costs in ways that can increase firm profits. DellaVigna and Malmendier (2004) show that some profit-maximizing firms design contracts that introduce switching costs and back-loaded fees to extract more profits—by taking advantage of consumers with time-inconsistent preferences and naive beliefs. Taxi drivers have an incentive to obtain a tip menu that will extract the highest tips

possible from passengers. In this paper, switching costs arise from switching from a menu to a non-menu tip.

Our paper contributes to the literature on preference identification in settings with framing effects (for example, Bernheim and Rangel (2009), Rubinstein and Salant (2011), Benkert and Netzer (2018), and Goldin and Reck (2019)). Specifically, we present empirical evidence from a field setting on how unobserved norms affect consumer choices in the presence of default suggestions. We take an approach similar to Goldin and Reck (2019) by relying on a revealed preference framework to recover population preferences (the distribution of beliefs about the social norm tip) from a subset of observed choices (non-menu tips). An added advantage of using our approach to recover peoples’ unobserved beliefs is that, we remain agnostic as to how these beliefs or norms are formulated.⁴ In addition, we quantify the cost that consumers face for not conforming to the social norm.

Finally, our study contributes to the literature on choice architecture. In particular, the implications of a profit-maximizing menu for social welfare. Carroll et al. (2009) presents a theoretical model to determine the optimal 401k-enrollment policy for different choice situations. With the use of our parametric model, we provide empirical evidence from a field setting on how menus affect profits and the utility of consumers. We also estimate the welfare implications of different counterfactual menus.

The rest of the paper proceeds as follows; section 1.2 describes the tipping systems used in NYC Yellow taxis and gives a summary of our analysis data, section 1.3 lays out a structural model for tipping in taxis, sections 1.4 and 1.5 uses a nonparametric and a semi-parametric approach respectively to estimate decision costs, section 1.6 presents a parameterization of our structural model to estimate the social norm tip and disentangle decision costs into cognitive and norm deviation costs, section 1.7 conducts counterfactual exercises to predict the tip-maximizing menu and evaluate its implications for social welfare, and section 1.8 concludes.

1.2 Taxi Tipping Systems and Data

Virtually all NYC Yellow taxicabs use electronic devices provided by two vendors to collect credit and debit card payments. The vendors are Creative Mobile Technologies (CMT) and VeriFone Incorporation (VTS), which roughly supply equal shares of the electronic devices.⁵ These devices record information such as the fare, tip, trip distance, geo codes of pickup and drop-off locations, date and time of trip, and other trip characteristics.

Because all Yellow taxicabs look similar, a passenger cannot tell which vendor operates the electronic transmission device within a particular cab. At the end of a ride, a digital screen in the back of the taxicab shows the trip expenses. A passenger opts to pay with cash or to use the screen to pay with a credit or debit card. For credit or debit card

⁴The literature in psychology has several refinements for what norms are and how they are formed. One example is a descriptive norm. That is, one’s expectation of what others do in a comparable situation. Another is, other’s expectation of what one is supposed to do in a comparable situation. These different formulations of norms may influence one’s beliefs and preferences differently, and hence one’s choices (see discussion in Bicchieri and Dimant (2019)).

⁵We ignore a third vendor, Digital Dispatch Systems, because it provided less than 5% of the electronic transmission devices in use between 2009 and August 2010.

payments, passengers are provided with a menu of suggested tips. The passenger may leave no tip, choose one of the suggested menu options, manually key in any amount, or provide a separate cash tip.

In 2009–2010, CMT’s menu options were 15%, 20%, and 25%. It increased these amounts to 20%, 25%, and 30% starting February 9, 2011. Prior to 2012, VTS offered a menu of dollar amounts (\$2, \$3, and \$4) for fares under \$15, and choices of 20%, 25%, and 30% for larger fares. From 2012 on, it offered only the percentage choices. Therefore, the data set contains information on three sets of menus.⁶

To take advantage of the menu changes and differences across the two vendors, we use data from 2010 to 2015. The Taxi and Limousine Commission (TLC) compiles all the taxi trip data from the transmission devices in all active taxicabs. There were 725,441,461 taxi trips over the stated period. However, information on tipping is available only for credit and debit card transactions, which were used in roughly half of the trips, 427,142,274. We further limit the sample to rides that began and ended in New York City, had standard rate fares with no tolls, and had a positive tip recorded.⁷ The resulting sample covers 386,273,769 trips.

Some taxi screens display menu tip suggestions only as percentages, while others show both the percentages and the corresponding dollar amount. For example, between 2009 and 2012, VTS displayed corresponding dollar amounts for its percentage tip menu but CMT did not (Haggag and Paci, 2014). Moreover, since 2012, CMT and VTS have used menus with the same three tip options: 20%, 25%, and 30%. However, CMT calculates tips on the total fare: the sum of the base fare, the MTA tax, the tolls, and the surcharge. In contrast, VTS calculates tips on only the base fare and the surcharge. To avoid these complications, we use only CMT’s data (except in section 1.7.2). The data set reports only the dollar amount tipped by passengers. For example, if the tip percentage is 20% and the fare was \$10, the tip would be reported as \$2 in the data set. We convert that dollar amount to a percentage in our analyses.⁸

Table 1.1 shows summary statistics from the sample of trips in CMT taxicabs only. It shows trips from January 2010 through January 2011 (column 1), trips from February 2011 to December 2011 (column 2), and trips from 2014 (column 3). Between January 2010 and January 2011, CMT presented passengers with a menu that showed 15%, 20%, and 25%. Thereafter, the menu changed to 20%, 25%, and 30%. The major change between 2011 and

⁶Figure A.1 in the appendix shows a typical screen displaying menu tip options, and figure A.2 shows the menu options by vendor and when they changed.

⁷Because passengers often pay for the taxi fare using a credit card but give the driver a cash tip, we cannot infer that a lack of a credit card tip implies that no tip was given.

⁸We account for possible rounding errors by considering any tip that falls in the range between 19.99% and 20.01% as the lowest menu option (20%), tips in the range between 24.99% and 25.01% as the middle menu option (25%), and tips that fall in the range between 29.99% and 30.01% as the highest menu option (30%). For standard rate fares, passengers are charged \$2.50 upon entering the cab. Thereafter, every fifth of a mile or every minute when the cab travels less than 12 mph increases the fare by an additional \$0.40. After September 3, 2012, the Taxi and Limousine Commission (TLC) increased the travel rate from \$0.40 to \$0.50. A \$0.50 Metropolitan Transportation Authority (MTA) tax was added to all fares after September 2009. An additional \$0.50 night surcharge charge is added for trips between 8pm–6am, and a \$1 surcharge for trips picked up between 4pm–8pm on weekdays. Trips between Manhattan and JFK airport are charge at a flat rate. Trips outside NYC and other non-standard rate fares are listed at www.nyc.gov/html/tlc/html/passenger/taxicab_rate.shtml.

2014 occurred in 2012, when the TLC increased the taxi fare by about 17%. Before the menu change (column 1), the average tip amount was about \$1.77, which increased to \$1.95 after the 2011 change (column 2). However, the average taxi fare remained around \$10. After the CMT menu change, the average tip rate increased by 8%, from 17.82% before the menu change to 19.19% thereafter. The share of passengers who choose menu tips decreased by about one-fifth after the change (from 59.7% to 48.3%). In 2014 the share of passengers who choose menu tips returned to 60.6%. The fare increase in 2012 resulted in a higher average fare of \$12.17 by 2014. The average tip amount increased to \$2.27 in 2014 while the average tip rate remained at 19.06%.

1.3 The Structural Model

Why do people tip? The literature on tipping posits that people tip for strategic reasons intended to encourage better future services, and for psychological reasons such as societal pressures to conform to social norms (see Azar (2007) for a review). Repeated interactions between passengers and drivers are highly unlikely in NYC. Travelers hail the cabs nearest to them and there are approximately 13,500 Yellow taxicabs searching for passengers in the city. Therefore, passengers have no strong incentives to strategically tip drivers in order to receive better future taxi services. For that reason, we ignore strategic tipping in our model.⁹

We model a passenger's decision to tip when provided with a menu of tips. Passenger i gives a tip of $t_i\%$ of the taxi fare. She believes that the social norm tip is $T_i\%$ of the taxi fare (T_i may differ across passengers). Both t_i and T_i are tip rates. If t_i is less than the social norm tip T_i , she incurs a norm deviation cost $\nu(T_i, t_i)$ —a function that shows the degree to which she dislikes deviating from the norm. In addition, if t_i is not one of j tip rate options d_j in menu D_k , she incurs a cognitive cost c_i to compute the dollar tip amount on her taxi fare F_i . The norm deviation cost plus the cognitive cost (if any) is her total decision cost $[\nu(T_i, t_i) + c_i]$. At the end of the taxi ride, passenger i chooses a tip to maximize her utility or minimize her loss represented by

$$\text{Max}_{t_i} U = \underbrace{-t_i F_i}_{\text{Tip paid}} - \underbrace{\nu(T_i, t_i)}_{\text{Norm deviation cost}} - \underbrace{c_i \times \mathbf{1}\{t_i \notin D_k\}}_{\text{Cognitive cost}} \quad (1.1)$$

Decision Costs

The first term $-t_i F$ is passenger i 's expenditure from tipping at rate t_i . The second term $-\nu(T_i, t_i)$ reflects her disutility from deviating from T_i . The third term $-c_i \times \mathbf{1}\{t_i \notin D_k\}$ captures passenger i 's cost of computing her tip if she does not choose a menu tip option. Therefore, passenger i chooses her non-menu rate $t_i \notin D_k$ if the benefit of tipping at that rate (denoted as B_i) is greater than choosing a menu tip rate d_j (for $j = 1, 2, \dots, n$) in tip menu D_k . That is

$$B_i = (d_j - t_i)F > \underbrace{[\nu(T_i, t_i) - \nu(T_i, d_j)]}_{\Delta\nu} + c_i$$

⁹Azar (2010) finds no evidence that customers in restaurant tip strategically. However, the author finds evidence that customers tip for social/psychological reasons.

$$B_i = (d_j - t_i)F > \Delta\nu + c_i \quad (1.2)$$

Equation (1.2) implies that, all else equal, there is a fare threshold $\bar{F}_i$ above which passenger i computes her preferred non-menu tip. We reason that passengers have a rule of thumb for tipping from prior taxi ride experiences. That is, passenger i has a sense of the fare threshold $\bar{F}_i$ above which she computes her preferred tip, else she opts for a menu tip instead.¹⁰

We utilize both a nonparametric and semi-parametric approach to place bounds on the total decision costs from tipping. The nonparametric methodology utilizes variation in non-menu tips because of an adjustment in tip menu options. The semi-parametric approach utilizes variation in the share of passengers who pick non-menu tips as a function of the taxi fare. Both methodologies require weak assumptions (introduced beneath). However, these methods do not permit us to separately identify the norm deviation cost and cognitive costs. In a third approach (parametric), we include parametric assumptions that permit us to independently distinguish between norm deviation costs and cognitive costs.

1.4 Nonparametric Estimation of Decision Costs

Whether or not a passenger decides to choose a menu option depends on the “stakes” of the decision B_i . However, changing the tip menu option d_j in equation (1.2) causes the stakes B_i to change. This provides identification help to estimate bounds on the decision costs passengers face when tipping. In 2011, the tip menu provider CMT changed its menu options from 15%, 20%, and 25% to 20%, 25%, and 30%. Figure 1.1A shows the distribution of tips before and after CMT’s menu change. There is a clear increase in the share of passengers choosing non-menu tips below 20% after 15% is removed from the tip menu.

We use this menu change as a natural experiment to estimate bounds on the monetary cost of deciding to choose a tip different from the menu options. The key intuition is that, all else equal, changing the menu option(s) d_j changes the values on both sides of the inequality in equation (1.2). Hence, the sign of the inequality might change for passengers who are on the margin of choosing a non-menu tip. Thus, the identifying variation for this exercise is the change in the share of passengers who choose non-menu tips after the menu changes. If a passenger chooses a non-menu tip, then she finds it beneficial to incur the costs associated with deciding to tip at her preferred rate instead of choosing a menu option. This approach is nonparametric. That is, we remain agnostic about the underlying model that governs how passengers decide to tip. However, we make two assumptions for this exercise.

A1 - One’s belief about the tipping norm T_i is jointly independent of the menu of tips and the taxi fare.

¹⁰If a passenger takes the same ride each time, she might learn over time to compute her preferred tip, hence driving down her cognitive cost to zero. However, the taxi fare is calculated based on the time spent in the taxicab and the distance traveled. Thus, the taxi fare is not deterministic, but depends on traffic conditions and the route the driver takes. We think it is more difficult to learn to compute the relevant tip for different taxi ride lengths and durations than to use the stated fare threshold rule of thumb. This reasoning may not hold for passengers who are tourist or do not often take taxis. However, this concern should be mitigated given that we exclude airport rides to and from JFK airport in our analysis.

A2 - Decision costs are independent of the menu of tips and constant over time.

Assumption A1 suggests that the differences in the observed tipping choices under the two menus are due to the change in menu options and not differences in the passengers observed under the different menu. Assumption A2 suggests that decision costs are uncorrelated with the menu of tips, and they remain stable over time. The frame of the tip menus used in this analysis follow the same structure—both menus present percentage tip options. Therefore, the difficulty in computing one’s preferred non-menu tip should not change because a menu option is added or removed. Other empirical support for these two assumptions are presented in the appendix (section A.2). We later relax assumption A1 with the aid of a parametric model presented in section 1.6.

Constructing Bounds

By inspecting non-menu tips in figure 1.1A, we only find significant increases in tip rates below 20% after the CMT menu change.¹¹ We therefore restrict attention to tips at or below 20% to compute bounds for decision costs. The relevant menu options are 15% and 20% before the change, and only 20% after. The following instructing example presents the insight for the nonparametric estimation of bounds on decision costs.

Suppose there are two menus, each with a single menu tip 15% and 20% of the taxi fare respectively. Each passenger has a preferred tip different from the menu options, but in order to pick this tip, one must pay the decision cost to switch from the menu. For either menu, the passenger has a choice between either paying the difference between the available menu tip and their preferred tip or incur the decision cost of switching from the menu. For a fare of \$10, suppose we observe a passenger picking the menu option 15% when presented, but then switches to her preferred tip (say 10%) when faced with the 20% menu. We conclude that the decision cost to switch is more than \$0.50 (the difference between the 15% tip and the preferred tip), and less than \$1 (the difference between the 20% tip and the preferred tip).¹² This effectively places observable monetary bounds on the unobserved decision cost of switching from a menu option. From equation (1.2), that is $(0.15 - 0.10) \times \$10 = \$0.50 < \Delta\nu + c_i < \$1 = (0.20 - 0.10) \times \10 . With this reasoning, we assume that passengers who choose non-menu tips reveal their preferred tip. Therefore, for a given fare F_i , the lower and upper bound for the decision cost of switching from 15% to some non-menu tip t_i after the menu change is given by $[|0.15 - t_i|F_i, |0.20 - t_i|F_i]$.

¹¹After the menu change, there were significant increases in the share of passengers who tip at the menu options that remained unchanged (20% and 25%). A possible explanation for the increase at 20% is that some passengers who chose the 15% menu option now choose 20%. For the increase at 25%, the compromise effect may be a possible explanation. This is the hypothesis that, consumers are more likely to choose a middle option out of a selection rather than the extremes. Our method of computing bounds on decision costs cannot be applied to the changes in the share of passengers at the menu options. Thus, we do not include them in our calculations.

¹²Notice that, if her decision cost is less than \$0.50, then we should have observed her choosing 10% when the 15% menu option was available. Similarly, if the cost of deciding to tip 10% is more than \$1, then she benefits by choosing the 20% menu option.

Estimating Bounds

Our goal is to recover bounds on the distribution of decision costs for passengers who switch to choosing non-menu tips after the CMT menu change.¹³ Let $\Delta S_{(t,F)}$ represent the increase in the share of passengers who choose a non-menu tip t for a taxi fare F after 15% is removed from the menu. For each $\Delta S_{(t,F)}$, we compute the corresponding bounds as $[|0.15 - t|F, |0.20 - t|F]$.

We focus on passengers who tip 20% of the taxi fare or less for this analysis. This reduces our 2010-2011 analysis sample by one fifth. Given the reasoning behind how the bounds are estimated, we should not observe significant changes in the share of passengers who choose tips above 17.5% after the menu change. For example, after the 15% menu option is removed, passengers whose preferred tip is 19% should find it more beneficial to pick the 20% menu option rather than calculating 19%.¹⁴

To compute the bounds of decision costs, we proceed in three steps. We group taxi fares into 29 non-overlapping bins of width \$2: [\$3, \$5], (\$5, \$7], (\$7, \$9]...(\$59, \$61]¹⁵, and then categorize tips into 20 non-overlapping tip rate bins of width one percent: 1%, 2%, 3%...20%. Thus, the 1% bin is the share of all passengers whose tips falls within [0.5%, 1.5%] of the taxi fare, 2% is the share whose tips fall within (1.5%, 2.5%] of the taxi fare, and so forth. Then for each tip bin, we compute the shares of tips that fall within each fare bin before and after the menu change.¹⁶ We take the difference between the two shares as an estimate of $\Delta S_{(t,F)}$. The midpoint of each fare bin is then used to compute the lower and upper bounds of decision costs. For example, for all taxi fares that fall within fare bin (\$9, \$11], \$10 is used to compute the relevant bounds. So, for each tip rate t , we construct bounds for the corresponding CDF of decision costs by combine the shares $\Delta S_{(t,F)}$ with their subsequent estimated bounds. For example, figure 1.1B shows the computed bounds for the CDF of decision cost conditional on a tip rate of 10%.¹⁷

Suppose that the midpoint of the estimated bounds of decision costs is similar to the true decision cost. Then we can use the midpoints of all the conditional CDFs in conjunction with the relevant shares $\Delta S_{(t,F)}$ to estimate an unconditional CDF of decision costs. The solid line in Figure 1.1C shows the estimated CDF. From this distribution, the average decision cost of making a calculated choice versus choosing a menu tips is \$1.89 (15.53% of the average taxi fare of \$12.17).

The nonparametric estimate of bounds on decision costs does not control for trip characteristics such as day of week, time of day, distance travelled, duration of trip, weather conditions etc. Therefore, if any of these factors systematically influences the choices of passengers, then the estimated bounds on decision costs is biased. In addition, the data limits us from observing the choices of the same passenger under CMT's initial and later menu. Therefore, the assumption that the distribution of decision costs and fares remain the same before and after the change in the menu may not hold. In the following section, we utilize a

¹³It is important to note that passengers with the highest decision cost will almost always choose menu options. Thus, our estimated bounds on the distribution of decision costs are censored from above.

¹⁴We find this to be the case after excluding round dollar tip amounts. However, our analyses are unaffected even if we don't exclude such tips.

¹⁵We group fares into bins because the data is sparse for taxi fares above \$50.

¹⁶Figure A.5 in the appendix shows the distribution of tips for different levels of the taxi fare.

¹⁷Figure A.6 in the appendix shows the computed bounds for all other other tip rates.

semi-parametric approach to address some of the limitations of the nonparametric approach.

1.5 Semiparametric Estimation of Decision Costs

We use a semi-parametric strategy akin to the nonparametric approach to estimate new bounds on the distribution of decision costs. There are three main innovations with this approach. First, we rely on changes in the share of non-menu tips as a function of the taxi fare to identify bounds on decision costs. Second, we use a semi-parametric strategy that allows us to account from trip level characteristics and other observables. Third, we use data from trips in 2014 when all Yellow taxis presented passengers with the same menu tip options (20%, 25%, and 30%), and there were no major changes in the taxi industry. Therefore, our semi-parametric estimates avoid the potential impact of changes in the menu of tips on unobserved passenger behavior, potential time differences in decision costs, and account for observable characteristics. We rely on assumption A1 and an additional assumption A3. The new assumption is

A3 - Decision costs are jointly independent of the taxi fare F_i and one's preferred tip t_i^* .

Because we do not observe decision costs, there is no straightforward way to test A3. However, we find the data to be consistent with assumption A3. For example, we do not find a large share of passengers tipping at 10% relative to other non-menu tip rates that may be relatively more difficult to compute. We also find that passengers are no more likely to tip at non-menu tip rates for fares where tip rate computations (percent to dollar conversions) may be easier (e.g., fares that are multiples of \$10).¹⁸

In 2014, all taxicabs presented 20%, 25%, and 30% to passengers as menu tip options. We reason that, if a passenger chooses a tip different from the menu options, then she reveals a preference for her tip relative to the menu options. According to the model in section 1.3, such a passenger deems it economical to incur the decision costs associated with offering her preferred tip instead of choosing a menu option. For this analysis, we restrict attention to taxi trips in CMT cabs from 2014 where passengers tipped 20% of the taxi fare or less. For these passengers, we assume that 20% (the lowest menu tip option) is what they would have chosen had they decided to choose a menu tip option. Recall that t_i is the observed tip rate in the data, $d_j = 20\%$ of the taxi fare is the relevant menu option, and F_i is the taxi fare. It follows from equation (1.2) that the benefit (B_i) for a passenger choosing to tip $t_i < d_j$ is $B_i = (0.20 - t_i) \times F_i$. Hence, a passenger who tips $t_i < 20\%$ reveals that her benefit from tipping t_i is greater than her decision costs associated with computing her preferred tip.

Constructing Bounds

The key intuition to estimate bounds on decision cost is that, all else equal, a change in the fare F_i changes the values on both sides of the inequality from equation (1.2). Hence, the sign of the inequality will likely change for passengers who are on the margin of choosing a non-menu tip. For such a passenger, the benefit from tipping at their preferred tip rate is approximately equal to the decision cost involved in computing their preferred tip (i.e., $B_i \approx \Delta\nu + c_i$). We compute the bounds on decision costs as follows. Suppose that F_i

¹⁸These empirical observations are further discussed in section A.2 of the appendix.

increases by ΔF , then for a passenger on the margin of choosing a non-menu tip rate t_i , we bound her decision costs of deviating from the menu ($d_j = 20\%$) to choosing t_i as $(0.20 - t_i) \times F < \Delta\nu + c_i < (0.20 - t_i) \times (F_i + \Delta F)$.

In addition, we estimate the share of passengers with preferred tip t_i and decision cost $\Delta\nu + c_i$ as follows. Denote $p(t_i|F_i, d_j = 20\%, X_{it})$ as the probability of choosing a tip t_i conditional on the taxi fare F_i , the menu tip option $d_j = 20\%$, and a vector of observed trip characteristics X_{it} . Suppose that F_i increases by ΔF , then it follows from equation (1.2) that $\Delta p_{(t,F)} = p(t_i|F_i + \Delta F, d_j, X_{it}) - p(t_i|F_i, d_j, X_{it}) \geq 0$. $\Delta p_{(t,F)}$ is the change in the probability of choosing one's preferred tip t_i relative to the menu tip option when F_i increases by ΔF . When ΔF is small (a marginal increase in the fare), $\Delta p_{(t,F)}$ represents the share of passengers who reveal that their benefit from giving their preferred tip is approximately equal to their cost of deliberating and computing their preferred tip (i.e., $B_i = (0.20 - t_i) \times (F_i + \Delta F) \approx \Delta\nu + c_i$).

Estimating Bounds

To implement the procedure above, we first impose a parametric structure by estimating an ordered logistic regression. In this regression the outcome variable is the tip rate categorized into 20 non-overlapping bins of width one percent (namely 1%, 2%, 3%...20%) and the covariates are the taxi fare, month of year, day of week, hour of day, holidays, weather condition, and hourly temperature and precipitation.¹⁹ Setting all the covariates to their sample average, we estimate the predicted probabilities for choosing each tip rate in the outcome variable as functions of the taxi fare. Figure 1.2A shows the predicted probability estimates of choosing 10% as a function of the taxi fare. As expected, the probability of choosing to tip of 10% is increasing as the fare increases.²⁰ In contrast, figure 1.2B shows that the probability of choosing the menu tip option 20% is decreasing as the fare increases.

With the predicted probabilities for tip rates from 1% through 19%, we compute both the change in the probability of choosing each tip rate and the corresponding bounds on decision cost for small increments in the fare. We then combine the estimates ($\Delta p_{(t,F)}$ and the estimated bounds) to construct bounds on the distribution of decision costs. For example, figure 1.2C shows the empirical estimate of bounds on the CDF of decision cost for passengers who tip 10% of the taxi fare.²¹ We use the midpoints of the estimated bounds of decision costs from all the other non-menu tip rates to estimate an unconditional CDF of decision costs. The dashed line depicted in figure 1.2D shows the semi-parametric estimate of the CDF of decision costs. The distribution averages at \$1.64 (13.48% of the average taxi fare of \$12.17). Note that the average decision cost from the semi-parametric approach is \$0.25 lower than what we estimated using the nonparametric approach in section 1.4 (\$1.89)—shown as a solid line in figure 1.2D. Some of the difference in the two estimates may stem from the fact that the semi-parametric estimate is purged of potential biases from trip characteristic and unobserved impacts of changes in tip menus.

The semi-parametric approach has some limitations similar to the nonparametric approach. First, we are not able to recover the full distribution of decision costs, since the

¹⁹The outcome variable is categorized as follows; for example, 15% is defined as the share of passengers whose tip fall within the range of 14.5% and 15.5%.

²⁰Figure A.7 in the appendix shows the estimated predicted probabilities for all the other tip rates.

²¹Figure A.8 in the appendix shows the computed bounds for all other other tip rates.

sample is limited to passengers who tip less than 20%. The second limitation is that only fares within the range of \$3 - \$30 are used in this exercise. Thus, the support of the estimated decision costs is censored.²²

1.6 Parametric Estimation of Decision Costs

Both the nonparametric and semi-parametric approaches of estimating decision costs provides evidence that decision costs are large relative to the taxi fare. However, we are not able to distinguish between the norm deviation cost of not conforming to the perceived tipping norm and the cognitive cost involved when computing one's preferred tip.

It is necessary to account separately for the social pressures that regulate decision-making versus the effort required to make a decision. In fact, consumers may feel obligated to conform to social norms that go against their personal desires. Thus, in the tipping context, where social norms matter for decision making, it is important to distinguish between norm deviation and cognitive costs and quantify their economic significance. To do that, we place more structure on the tipping behavior model from section 1.3. Specifically, we rely on assumptions A1 - A3 and specify a parametrized utility function. This extra structure allows us to separately identify the norm deviation cost, the cognitive cost, and the distribution of beliefs about the social norm tip across taxi passengers.

1.6.1 The Parametric Model

Each passenger is characterized by four random variables (T_i, F_i, D_k, c_i) drawn from some underlying distribution. Analogous to equation (1.1), passenger i chooses a tip to maximize her utility or minimize her loss represented by

$$U_i = \underbrace{-t_i F_i}_{\text{Tip paid}} - \underbrace{\theta (T_i - t_i)^2}_{\text{Norm deviation cost}} - \underbrace{c_i \times \mathbf{1}\{t_i \notin D_k\}}_{\text{Cognitive cost}} \quad (1.3)$$

Decision Costs

The first term $-t_i F_i$ is her expenditure from tipping t_i (a percentage of the fare). The second term $-\theta (T_i - t_i)^2$ is her norm deviation cost—disutility for not conforming to what she believes is the social norm. The scalar θ is the norm deviation cost parameter. Of course, passenger i avoids the norm deviation cost if she tips T_i . However, if she deviates from tipping T_i , then her norm deviation cost increases with the size of the percentage point deviation.²³ The third term, $-c_i \times \mathbf{1}\{t_i \notin D_k\}$, is passenger i 's cognitive cost of computing her preferred tip, where c_i is a fixed dollar cost of calculating $t_i \times F_i$, and $\mathbf{1}\{t_i \notin D_k\}$ is an

²²The support of the fare is restricted to the \$3 - \$30 range because, \$3 is the lowest taxi fare, and \$3 - \$30 is the range within which the change in the predicted probabilities ($\Delta p_{(t,F)}$) is nonnegative for all non-menu tips.

²³Because the norm deviation cost of not conforming to the norm is symmetric, passenger i will nonetheless experience a utility loss if she chooses a tip larger than T_i . However, it may be intuitive that one would likely feel ashamed or experience disutility for choosing a tip that is less than T_i but not for a tip equal to or larger than T_i . We therefore conduct an exercise where passenger i is assumed to face no disutility from choosing a tip that is larger than her belief about the tipping norm T_i . So passenger i 's disutility from tipping can be written as

indicator function that equals one if t_i is not one of the options d_j of $j = 1, 2, 3$ in tip menu D_k and zero otherwise.

The dollar amount of the tip $t_i F_i$ enters linearly into the utility function. Hence the utility function is quasi-linear in money. This assumption is relatively innocuous given that tips are a small amount compared to the wealth of customers. We remain agnostic as to how passenger i determines T_i and assume that all the processes involved, including warm glow, are subsumed in passenger i 's formulation of T_i .

$B_i = (d_j - t_i)F_i$ is the benefit from choosing $t_i \notin D_k$ rather than a (higher) menu default. The cost of tipping t_i is that her norm deviation cost rises from $\theta (T_i - d_j)^2$ to $\theta (T_i - t_i)^2$. In addition, she incurs a cognitive cost of c_i . Thus, she tips at her preferred rate if the benefit B_i of tipping her preferred tip t_i exceeds the extra cost from not choosing a menu tip:

$$B_i = (d_j - t_i)F_i > \theta [(T_i - t_i)^2 - (T_i - d_j)^2] + c_i \quad (1.4)$$

For $t_i < d_j$, it follows that $\frac{dB_i}{dF} = d_j - t_i > 0$. That is, the benefit of computing one's ideal tip is larger at higher fares. Therefore, passengers will be more likely to choose non-menu tips at higher fares. Figure A.9A in the appendix confirms this observation from the model. With equation (1.4) in mind, we reason that, to decide on the optimal tip, passengers have a rule of thumb for tipping from prior taxi ride experiences. That is, passenger i has a sense of a fare threshold $\bar{F}_i$ above which she computes her preferred tip, else she opts for a menu tip instead.

We now solve for the preferred tip by maximizing equation (1.3). We ignore the cognitive cost c_i because of the indicator function $\mathbf{1}\{t_i \notin D_k\}$. From the first-order condition, we find that the optimal tip is

$$t_i^* = T_i - \frac{1}{2\theta} F_i. \quad (1.5)$$

According to the first-order condition, passenger i 's preferred tip t_i^* is less than her belief about the social norm tip T_i . The preferred tip rate falls as the fare increases $\left(\frac{dt_i^*}{dF} = -\frac{1}{2\theta} < 0\right)$.²⁴ Therefore, when deciding on how much to tip, a passenger tries to save a little bit by trading off the dollars lost to tipping at the social norm against the shame from being a cheapskate. Some passengers may use other heuristics such as tipping a fixed dollar amount or rounding off the taxi fare to a specific dollar amount. For example, a passenger facing a fare of \$9 many decide to tip \$1 to round off her total trip expense to \$10. We will account for such behavior in our analysis.

Our model permits us to relax assumption A1 by allowing the tip menu to affect passengers' belief about the tipping norm. Suppose that $T_i = \tilde{T}_i + \gamma_k D_k$, where $\tilde{T}_i$ is passenger i 's

$$U_i = \begin{cases} -t_i F - \theta (T_i - t_i)^2 - c_i \mathbf{1}\{t_i \notin D_k\} & , \text{ if } t_i < T_i \\ -t_i F & , \text{ if } t_i \geq T_i \end{cases}$$

However, using this parameterization does not affect any of our estimates from equation (1.3) in a significant way. This is because, in equation (1.3), the only case where a passenger may tip above T_i is if she chooses a menu tip larger than T_i which rarely occurs in the model setup. We therefore proceed with equation (1.3) in our analysis. However, using this new setup does not impact our results.

²⁴This observation generally holds in the data. Figure A.9B in the appendix shows that the average tip rate falls as the fare increases.

belief about the tipping norm, D_k is a vector of different menus enumerated as $k = 1, \dots, n$, and γ_k is a vector of coefficients that denote the differential impact of each menu on a passenger's belief about the tipping norm. With this reasoning, the first order condition (equation (1.5)) can be rewritten as

$$t_i^* = \tilde{T}_i + \gamma_k D_k - \frac{1}{2\theta} F_i. \quad (1.6)$$

1.6.2 Estimation Procedure

The parameters to be estimated from our model are a passenger's belief about the social norm T_i , the norm deviation cost parameter θ and the cognitive cost c_i of computing one's preferred tip t_i^* . Given the structure of the utility function, we are able to rely on the first-order condition, equation (1.5), to estimate the unobserved distribution of T_i and the norm deviation cost parameter θ . The advantage here is that, we need not make any distributional assumptions regarding T_i . In addition, θ is directly estimated in the same equation used to recover T_i . This leaves the distribution of c_i to be estimated, which we compute via a Minimum Distance Estimator.

1.6.2.1 Estimating Tipping Norms and Norm Deviation Costs

We estimate equation (1.5) using an ordinary least squares regression (OLS), where all components of the regression equation have structural interpretations linked to the proposed model. Specifically, the equation to be estimated is

$$t_i = \alpha_T + \beta F_i + \varepsilon_i, \quad (1.7)$$

where t_i is the observed tip rate in the data, α_T is the constant term, F_i is the observed taxi fare, and ε_i is the residual.²⁵ The challenge with estimating equation (1.7) using all observed tips is that, for passengers who choose tips from the menu, we do not know what tips they would have given otherwise. As a result, the coefficient estimates from equation (1.7) are likely biased in an OLS regression. However, we observe t_i^* for the subsample of passengers who choose non-menu tips. Our approach is to estimate θ and the distribution of T_i using the subsample of non-menu tips. We will then extrapolate the estimates to all passengers by adjusting for possible sample selection.

The link between equation (1.5) and (1.7) is that, α_T is the population average tipping norm $E[T_i]$, $\beta = \frac{1}{2\theta}$, and the residual term $\varepsilon_i = T_i - \alpha_T$ represents the difference in passenger i 's belief about the tipping norm relative to the the population average. To recover an

²⁵Note that, the outcome variable in the equation is the tip rate (i.e., $\frac{\text{Tip}}{\text{Taxi Fare}}$) and the main covariate is the taxi fare. Thus, division bias might be a concern for estimating equation (1.7). However, we expect this bias to be insignificant in our setting for two reasons. First, there is little to no measurement error in the data on tips and fares. Second, the lowest taxi fare is \$3—hence the outcome variable (tip rate) does not have a case where the numerator (tip) is divided by \$0 or a very small fare.

estimate of T_i (denoted as $\hat{T}_i$), note that

$$\begin{aligned}\hat{\varepsilon}_i &\equiv t_i^* - \hat{\alpha}_T + \frac{1}{2\hat{\theta}}F_i \\ \hat{\alpha}_T + \hat{\varepsilon}_i &\equiv t_i^* + \frac{1}{2\hat{\theta}}F_i = \hat{T}_i\end{aligned}$$

Therefore, the constant term plus the residuals is an estimate of the unobserved realization of peoples beliefs about the tipping norm.

Notice that, equation (1.7) can be rewritten in the regression form as

$$\begin{aligned}E[t_i^*|F_i, c_i] &= \alpha_T + \beta F_i + E[\varepsilon_i|F_i, c_i] \\ &= \alpha_T + \beta F_i + E[\varepsilon_i|c_i]\end{aligned}$$

The above result follows because $\mathbb{A}1 \implies T_i \perp (F_i, D) \implies (\alpha_T + \varepsilon_i) \perp (F_i, D)$, thus $\varepsilon_i \perp F_i$.²⁶ Therefore, the decision to choose a non-menu tip depends solely on one's cognitive cost c_i . Unfortunately, we do not know the relationship between ε_i and c_i . Therefore, using the subsample of passengers who choose non-menu tips to estimate equation (1.7) may be problematic due to sample selection. The concern here is the possibility that $E[\varepsilon_i|c_i] \neq 0$. That is, selecting to give a non-menu tip may be systematically related to the cognitive costs of a decision-maker. Thus, our estimates will be biased if we use the subsample of non-menu tips. To address this concern, we employ a 2-step Heckman selection correction approach. The idea here is to use a *decision quality instrument* to correct for potential sample selection among passengers who choose non-menu tips (Goldin and Reck, 2019). This instrument must be an element of the decision-making environment that affects a passenger's decision to choose a menu tip, but unrelated to her belief about the tipping norm or the cost of computing her preferred non-menu tip.

In the first step of the selection model, we use a probit equation to estimate the probability of choosing a non-menu tip using the entire sample. The outcome variable is a dummy variable that equals one if the passenger chooses a non-menu tip and zero otherwise. The independent variables are the taxi fare, and the added decision quality instrument is the taxi driver's report of the number of passengers on each trip. We reason that a passenger faces a greater time pressure to choose if they are with other co-passengers, but the time pressure is unlikely to affect their preferred tip or the difficulty of computing it. A concern that might violate the exclusion restriction is that the preferences of an individual who travels in a group may be impacted by the preferences of his or her co-riders.²⁷ Lastly, it is important to note that the number of co-passengers does not enter the utility function defined in equation (1.3). Thus, the number of co-passengers is an excluded instrument with respect to the structure of our model.

In the second step, we estimate equation (1.7) using the subsample of passengers who choose non-menu tips while including the estimated Inverse Mills Ratio from the first-step probit regression to correct for sample selection bias.

²⁶The maintained assumption is that, assumption $\mathbb{A}1$ holds independent of the cognitive cost c_i .

²⁷For example, passengers who ride in groups may decide to share the bill. This may impact the preferred tip because, the group's preferred tip may differ from each traveler's privately preferred tip.

Estimates of Tipping Norms and Norm Deviation Costs: We find that the average perceived norm among passengers is to tip around 20% of the taxi fare. Tipping five percentage points less than the norm results in a norm deviation cost between \$0.30 and \$0.38 (2.5% - 3.1% of the average taxi fare of \$12.17). We find limited evidence that the CMT menu change significantly impacted passengers’ beliefs about the tipping norm.

We first use 2014 CMT taxi trips for our analysis. 2014 was a period when the NYC Yellow taxi industry had essentially reached a steady-state. This is because, regardless of vendor, all taxis presented passengers with the same menu tip options (20%, 25%, and 30%), and there were no changes in the fare or major developments in the taxi industry. We then use CMT taxi trips from 2010 and 2011 to analyze how different tip menus affect passengers’ beliefs about the tipping norm by relying on the CMT menu changed over this period.

Table 1.2 presents the parameter estimates from our structural model. In column (1), we use CMT taxi trips from 2014, and in column (2), we use CMT taxi trips from 2010-2011. Panel A shows the Heckman selection correction model estimates of the norm deviation cost parameter and the mean of the distribution of passengers’ beliefs about the tipping norm.²⁸ Estimates from Panel A column (1) corresponds to equation (1.7) and are statistically significant and precisely estimated. We recover the distribution of passengers’ beliefs about the tipping norm by augmenting the residual from the estimated regression with the estimated constant term. Figure 1.3A shows this distribution. Passengers’ beliefs about the tipping norm are mostly within the range of 12% and 25% of the taxi fare. The coefficient on the taxi fare $\hat{\beta} = -0.00329$. This estimate implies that the norm deviation cost parameter $\hat{\theta} = \frac{1}{2\hat{\beta}} \approx 152.24$. The norm deviation cost increases with the size of the percentage point deviation from one’s belief of the social norm tip. For example, the dollar value of the norm deviation cost for tipping five percentage points less than one’s perceived norm is $\theta \times (T_i - t_i^*)^2 = 152.24 \times (0.05)^2 = \0.38 . Thus, for the average taxi fare of \$12.17, the passenger saves \$0.61 at a cost of \$0.38. The constant term α_T is estimated as 0.198 and this implies that the average perceived tipping norm across all passengers is to tip 19.8% of the taxi fare.

Following equation (1.6), we turn to analyzing the impact of different menus on passengers’ beliefs about the tipping norm. We use CMT trips from 2010 to 2011—the period where CMT changed its menu choices from 15%, 20%, and 25% to 20%, 25%, and 30%. We estimate equation (1.7) with an added indicator variable for the period after the menu change. Panel A column (2) presented the results. Our estimate of the constant term suggests that when passengers are presented with the tip menu showing 15%, 20%, and 25%, the average of passengers’ beliefs about the tipping norm is 20.22% of the taxi fare. However, the coefficient on the indicator variable suggests that, when the menu changes to show 20%, 25%, and 30% the average tipping norm decreases to 19.56% of the taxi fare. While the difference between these two averages is statistically significant, it is not economically significant. In particular, the decrease in the average tipping norm after the menu change is only three hundredths of the average norm before the change. Therefore, the CMT tip menu change had very little impact on passengers’ beliefs about the tipping norm. The similarities between the distribution of passengers’ beliefs about the social norm tip is more apparent

²⁸Column (1) of table A.2 in the appendix presents the results from the first step probit estimation of the Heckman selection correction model.

in figure 1.3B, which shows the two distributions (before and after menu change) overlaid. It is also remarkable that passengers average perceived tipping norm remained very similar over the period of 2010 to 2014 (20.22% in 2010-before menu change, 19.56% in 2011-after menu change, and 19.8% in 2014). On the other hand, our estimate of the norm deviation cost parameter is 117.92 ($\hat{\beta} = -0.00423$) for trips in 2010-2011. Therefore, tipping about five percentage points less than one's perceived tipping norm comes at a norm deviation cost of \$0.30. This is \$0.08 less than our estimate from 2014 (column (1)), suggesting that the norm deviation cost increased by 21% between 2010 and 2014.

Table A.3 in the appendix present OLS estimates analogous table 1.2 Panel A. The OLS estimates do not account for sample selection, however, the estimates are very similar to those from table A.3 Panel A. Therefore, while the coefficients on the Inverse Mills Ratio term in table 1.2 Panel A are statistically significant, the similarity in estimates across the two models is comforting. We take this to mean that sample selection concerns are inconsequential when using non-menu tips to estimate the norm deviation cost and the population distribution of beliefs about the tipping norm.

We also find that some passengers provide tips that are round-dollar amounts and this creates mass points in the empirical distribution of the dollar value of tips. These passengers may possibly be using some heuristic that may not be captured in our model. We control for this round-number bunching by including an indicator variable for round number tips in our regression equations to capture these rounding effects.²⁹

1.6.2.2 An Upper Bound of Cognitive Costs for Non-Menu Tips

With estimates of the distribution of passengers' beliefs of the social norm tip and the norm deviation cost parameter, cognitive cost is the final primitive of the model left to be estimated. With no further parameterization, knowledge of passengers' beliefs about the social norm tip and the cost of deviating from these beliefs is sufficient to estimate an upper bound on the cognitive cost for the subset of passengers who give non-menu tips. The idea is that, passengers who give non-menu tips find it more beneficial to compute their preferred tip instead of choosing a menu tip. Therefore, if we know a passenger's belief about the norm T_i , the norm deviation cost parameter θ , and the fare F_i , we can use equation (1.4) to compute the level of cognitive cost above which such passengers would opt for a menu tip. That is,

$$\bar{c}_i = (d_j - t_i)F_i + \theta [(T_i - t_i)^2 - (T_i - d_j)^2].$$

For each observed non-menu tip rate t_i and the corresponding fare F_i , we compute T_i as defined in equation (1.5). We then choose the nearest menu tip rate d_j above t_i as the reference menu option for computing $\bar{c}_i$. For example, for a non-menu tip rate of 17%, the relevant menu option is 20%, and for a non-menu tip rate of 23%, the relevant menu option is 25%, and so forth.³⁰ We go further to estimate the distribution of norm deviation costs among this subset of passengers as well. For each passenger, we compute their relevant norm deviation costs as $\theta (T_i - t_i)^2$.

²⁹This approach is similar to what Kleven and Waseem (2013) used to capture the effect of self-employed workers who report round-number income amounts for tax purposes.

³⁰Because there is no menu option above 30%, we do not use non-menu tips above 30% in this analysis.

Using trips from 2014 and parameter estimates from table 1.2 (Panel A, column (1)), figure 1.4 shows the distribution of the upper bound of the cognitive cost and the distribution of the norm deviation cost for passengers who give non-menu tips. The averages of these distributions are \$0.95 for cognitive cost, and \$0.33 for the norm deviation cost.

Passengers who choose menu options likely have higher decision costs relative to those that we observe giving non-menu tips. Therefore, to estimate the distribution of cognitive costs for all passengers (both non-menu tips and menu tips) we make an assumption about the distribution of cognitive costs within the population of passengers and use a Minimum Distance Estimator to estimate the relevant parameter(s) of the assumed distribution.

1.6.2.3 Minimum Distance Estimation of Cognitive Costs

We assume cognitive cost c_i to be exponentially distributed with rate parameter λ . The choice of an exponential distribution for cognitive cost is inspired by the estimated upper bound of cognitive costs for passengers who choose non-menu tips (figure 1.4). In addition, the estimated distribution of decision costs from both the nonparametric and semi-parametric approaches resample exponential distributions (figure 1.2D).

The passenger's objective is to give a tip that maximizes her utility. However, there is no analytical solution to equation (1.3) and hence, no corresponding closed-form expression. This is because the derivative of the indicator function $\mathbf{1}\{t_i \notin D_k\}$ is not well defined. We circumvent this problem by using a Monte Carlo procedure of an algorithm that chooses one of the menu options or a non-menu tip. The algorithm follows these steps:

1. For each observed taxi fare F_i , there is a random draw of T_i from the estimated distribution of passenger beliefs about the social norm tip, and a random draw of a cognitive cost value c_i from an exponential distribution with rate parameter λ .
2. The preferred tip t_i^* is then computed as defined in equation (1.5) using F_i , T_i , and θ .
3. Using equation (1.3), the utility levels for leaving a non-menu tip $t_i^*(U^{t_i^*})$ and all three menu tips: U^{d1} , U^{d2} , and U^{d3} are computed.
4. The algorithm then chooses the tip that results in the highest utility by comparing the four levels from step 3.

To identify a value of λ such that the model (equation (1.3)) predicts a realization of tips that matches the observed data as closely as possible, we match a vector of model predicted moments to those computed from the observed data. We use a simulated method of moments (SMM) algorithm that proceeds as follows.

Henceforth, quantities with the carets denotes estimates of population statistics. Let $g(\lambda|\hat{T}_i, \hat{\theta}) = [\hat{m} - m(\lambda|\hat{T}_i, \hat{\theta})]$ be a vector of moment conditions, where $\hat{m}$ is the vector of sample statistics (empirical moments from the data) and $m(\lambda|\hat{T}_i, \hat{\theta})$ is the model analogue of $\hat{m}$. Therefore, the SMM algorithm minimizes the criterion function $Q(\lambda|\hat{T}_i, \hat{\theta}) = g' \widehat{W} g$, where $\widehat{W}$ is some positive-definite weight matrix that is a function of the realized data. When minimizing the criterion function $Q(\lambda|\hat{T}_i, \hat{\theta})$, we match the sample statistics to their simulated analogues under the model. In particular, we employ a two-step procedure to compute the model parameters. In the first step, an identity matrix ($\widehat{W}_1 = I$) is used as

a preliminary weight matrix to estimate λ . Then, the estimated λ (denoted as $\hat{\lambda}_1$) is used to predict a set of realized tips via equation (1.3). Next, the predicted tips are used to compute $m(\hat{\lambda}_1|\hat{T}_i, \hat{\theta})$ —the model analogue of the empirical moments $\hat{m}$. We then calculate the vector of moment conditions as $g(\hat{\lambda}_1|\hat{T}_i, \hat{\theta}) = [\hat{m} - m(\hat{\lambda}_1|\hat{T}_i, \hat{\theta})]$. In step two, we assume independence across the moments so that the covariance between the moment conditions is set to zero. We then take the diagonal of the inverted variance-covariance matrix of the moment conditions from step one and use it as a weight matrix (i.e., $\widehat{W}_2 = [\text{diag}\{gg'\}]^{-1}$) to compute the final parameter estimates—via the same SMM algorithm.³¹ Therefore, the algorithm in the second step minimizes the squared distance between the empirical and the model predicted moments using a metric that is determined by the estimated weight matrix $\widehat{W}_2$.

Generally, λ is identified by the share of passengers who choose menu tips. Note that, if there is no cognitive cost for computing one’s preferred tip, then we should not find a significant share of passengers choosing from the menu relative to other non-menu tip rates. Thus, the shares of passengers who choose menu tips identifies λ and hence c_i .

As a primary set of moments used to identify λ , we construct sample statistics by dividing tip rates into 35 non-overlapping one percent bins, namely 1%, 2%, 3%...35%. Each statistic is defined as the share of passengers whose tip falls within a particular bin. For example, the estimated moment for passengers who tip 10% of the taxi fare is defined as the share of passengers who give a tip that is between 9.5% and 10.5% of their taxi fare.

We use the “optim” package that is implemented in the *R* statistical software as the numerical optimization algorithm to compute λ . This algorithm finds the parameter estimates that minimize the criterion function $Q(\lambda|\hat{T}_i, \hat{\theta})$. To avoid selecting a local minimum, we search for the model parameter estimate over 500 iterations of the algorithm and choose the estimate that results in the smallest minimized value of $Q(\lambda|\hat{T}_i, \hat{\theta})$. We compute standard errors using a bootstrapped procedure where 1000 independent draws of tips are constructed by a random resampling of tips generated via equation (1.3). The standard error is defined as the standard deviation of the distribution of parameter estimates computed from all 1000 bootstrap samples.

Minimum Distance Estimates of Cognitive Cost: Table 1.2, Panel B, presents the SMM estimates of the average cognitive cost. Because cognitive costs is assumed to be exponentially distributed with parameter λ , Panel B shows an estimate of the average cognitive cost $\frac{1}{\lambda}$ separately for the data from CMT trips in 2014 (column (1)) and for the CMT trips from 2010-2011 (column (2)). It is important to note that the change in the CMT tip menu in 2011 provides an extra source of variation that helps to identify λ . In particular, the menu change presents variation in menu options which serve as extra moments that help to identify λ .³² To reduce the computational burden, we randomly sample five million observations each from 2014 and 2010-2011 respectively for this analysis. Panel B also reports both the first and second step estimates of λ from the SMM algorithm.

³¹The theory suggests that the best choice of a weight matrix is the inverse of the covariance of the moment conditions.

³²We use 70 moments in the column (2) of panel B. Thirty-five moments from the period before the menu change and 35 after the menu change.

Using trips from 2014 (column (1)), the estimate of the average cognitive cost of computing one’s preferred non-menu tip is \$1.34 (11% of the average taxi fare of \$12.17). The estimate of the average cognitive cost is reduced to \$1.14 (9.4% of the average taxi fare) when using the data from 2010-2011 (column (2)). The first step estimates from the SSM algorithm are very similar to the second step estimates. This suggests that our estimates are not driven by the choice of weighting matrix.

1.6.3 Model Performance

The model performs the best for the period when passengers were presented with 20%, 25%, and 30% as tip menu options. Specifically, In 2014 and 2011, figures 1.5A and 1.5B show that our model mimics the point masses at all the three menu options and more or less at all other non-menu tip rates. On the other hand, the model does not perform well when predicting the 2010 distribution of tips in CMT taxis when passengers were presented 15%, 20%, and 25% as menu tip options (figure 1.5C). While the model performs well in predicting tipping behavior under the menu with 20%, 25%, and 30%, a χ^2 goodness of fit test suggests significant differences between the observed tips and the model predicted tips (test results are presented in the notes of figure 1.5).

As another model performance check, we combine the estimated components of decision costs estimates from the parametric model and compare them to the nonparametric and semi-parametric estimates. Adding the norm deviation costs of tipping five percentage points less than the norm (\$0.30 - \$0.38) and the cognitive cost of computing a non-menu tip (\$1.16 - \$1.34), we get that the average decision cost lies between \$1.46 and \$1.72 (i.e., between 12% and 14% of the average taxi fare \$12.17).³³ These parametric estimates are in similar magnitude as the nonparametric and semi-parametric estimates of the average decision costs of \$1.89 and \$1.64 respectively.

1.7 Analysis of Menu Tips

Given the proposed model and estimated parameters, we conduct counterfactual exercises to find the menu of tips that will maximize the tips that drivers receive from passengers. This exercise is of interest for two separate reasons that go beyond the context of tipping in taxicabs. First, for workers who receive a tipped wage³⁴ or depend on tips to supplement their income, we may want to construct a menu that will extract high enough tips in order to raise their earnings. Second, if we consider a cab driver as a sole proprietor, then implementing a set of menu options that maximizes tips directly impacts firm profits. Thus, this exercise is relevant for firms where tips are a direct source of revenue. However, in the contexts of

³³We choose a five-percentage point deviation from the tipping norm because, the estimated average norm deviation cost among passengers who give non-menu tips (\$0.33, see figure 1.4). Per our estimate of the norm deviation cost parameter, \$0.33 is roughly the cost of tipping five percentage points less than one’s belief of the tipping norm. Second, the average non-menu tip in 2014 is 4.25 percentage points less than the estimated average tipping norm. We round 4.25 up to 5.

³⁴This is a base wage below the minimum wage that is paid to employees who receive a substantial portion of their earnings from tips.

tipping in taxicabs, the reasons stated above are identical since drivers keep all the earnings (taxi fares + tips) from driving.

1.7.1 Tip-Maximizing Menu

To find the tip-maximizing menu, we need to know two things: (1) the number of menu options to show passengers, and (2) the corresponding tip rate for each option. It is important to note that this exercise is a computation of the tip-maximizing menu given a menu that presents customers with percentage tip options—as is currently the case. Thus, this is not a full characterization of the tip-maximizing menu, which may include but not be limited to presenting some combination of dollar tip amounts and percentages.

For this exercise, we proceed by setting the model parameters to estimates from column (1) of table 1.2. The procedure is to estimate equation (1.3) by fixing the model parameters and then setting menu tip options as the free parameters to be evaluated for values that maximize the average tip. To fix ideas, we first consider the case where drivers are restricted to show passengers a single menu tip option. We then search over a grid of tip rates between 0% and 100% to find the tip rate that our model predicts as increasing the average tip the most. figures (1.6)A show the results from the grid search for a single menu option that maximizes the average tip. The model predicts that the average tip rate is maximized when passengers are shown 22% as the menu option. The average tip when passengers are not shown a menu is 15.83%. However, with the 22% menu tip option, the average tip rate increases to 17.85%. This is a 12.76% increase in the average tip. Figure 1.6A also shows that presenting a menu option below 13% depresses the average tip rate. In sum, presenting a menu and the choice of option(s) may have either positive or negative consequences on profits.

We then proceed to search for menu tip options that will maximize tips from a menu restricted to showing only two options. Figure 1.6B shows a three-dimensional surface that characterizes the average tip from the different combinations of two menu tips options between 0% and 50%. Our model predicts that showing 20% and 27% as menu tips maximizes tips. This menu increases the average tip to 18.02%—a 13.83 % increase compared to the average tip without a menu. An inspection of figure 1.6B also suggests that certain choice combinations of the two options can either increase or decrease tips.

We continue to increase the number of menu tip options until the average tip stays practically the same upon adding more menu options. Figure 1.6C plots the average tip across all observed fares as the number of menu tip options increase. Figure 1.6C shows that the average tip increases no further after showing three or more tip-maximizing menu options. We therefore conclude that showing three menu options is tip maximizing. The corresponding predicted tip rates are 20%, 26%, and 32%. Figure 1.6D shows the model predicted distribution of tips for the tip maximizing menu. With this menu, the average tip rate increases to 18.15%. This is a 14.65% increase in the average tip relative to not presenting a menu. It is important to note that the tip-maximizing menu proposed by our model (20%, 26%, and 32%) is very similar to the menu currently offered to passengers (20%, 25%, and 30%). A main insight from this exercise is that, certain choice combinations of menu tip options are revenue decreasing. That is, there are some menu tip options that drive tips below what passengers would have given absent the menu.

1.7.2 Evolution of Menu Tip Options

We examine passenger tips across the different tip menus in NYC Yellow taxis from 2010 to 2014, and assess which menu induced passengers to tip the most. We consider three main periods in this analysis. The first period is from January 2010 to January 2011. The second period is from February 2011 to December 2011. The third period is from 2013 to 2014. In the first two periods, the two taxi vendors CMT and VTS provided passengers with different menus, and in the last period, both vendors provided the same set of menus. Details of the menus are presented in section 1.2. Table 1.3 presents three panels (A, B, and C) that correspond to the three periods being considered. Each panel has two columns that report the average tip rate for CMT rides (column (1)) and VTS rides (column (2)).

Panel A corresponds to the first period where CMT presented 15%, 20%, and 30%. In VTS cabs, passengers saw three options in dollars (\$2, \$3, and \$4) for fares under \$15, and three options in percentages (20%, 25%, and 30%) for fares above. Panel A shows that, on average, passengers tip at higher rates in VTS cabs (20.68%) relative to CMT cabs (17.81%). In the first period, CMT used an inferior menu compared to VTS. In the second period, CMT taxicabs changed their tip menu to show 20%, 25%, and 30%, while VTS cabs maintained the same menu as in the first period. Panel 2 shows that the CMT menu change increased the average CMT tip rate by about 7.5% (from 17.81% to 19.16%). The average tip rate remained the same for tips in VTS taxis. In the third period, both CMT and VTS cabs presented passengers with the same menu of tip options (20%, 25%, and 30%). Panel C shows that the average tip in VTS cabs dropped by two percentage points (from 20.66% to 18.55%). There was on change in the average tip in CMT taxis.

Both vendors currently present 20%, 25%, and 30% as options to passengers. With respect to presenting passengers with a menu that shows percentages as options, the results from table 1.3 suggests that the current menu induces passengers the most. Interestingly, these menu options are very similar to our model predicted tip-maximizing menu. Thus, the convergence in the tip menu across the two vendors over time is consistent with taxi companies learning overtime to us a menu that maximizes tips.

1.7.3 Welfare

When firms present customers with menus, it has implications for their profits and the utility of consumer who choose from the menu. Figures 1.6A and 1.6B show that the choice of menu options may drive tips above or below what passengers would give without a menu. Secondly, passengers will be more likely to forgo computing their preferred tip when presented with menu options that are close enough to their preferences. Hence, avoiding some of the decision costs involved to actively compute their preferred tip.

We evaluate how the revenues from tips and the utility from tipping are affected under different tip menus. First, we look at the case where consumers are not provided with a menu of tip options. Second, we consider the case of the previous CMT tip menu (15%, 20%, and 25%). Third, we consider the case of the current tip menu (20%, 25%, and 30%). Fourth, we consider the case of presenting the tip-maximizing menu (20%, 26%, and 32%) to passengers. Fifth, we estimate the menu that maximizes the utility of tippers and evaluate how it impacts the revenue from tips.

The utility from tipping (equation (1.3)) is quasi-linear in money. Thus, the social welfare from tipping is the sum of the dollar value of utility that consumers get from tipping and the tip revenue that drivers receive. We therefore use the parameter estimates from the model to compute the dollar value of a passenger’s utility from tipping. From equation (1.3), the utility from tipping is always less than zero, even for the case where the passenger decides not to leave a tip. This is because, in addition to the payment of the tip to the driver, the consumer incurs decision costs (norm deviation and/or cognitive costs) as well. Hence, the welfare from tipping is negative or at most zero. The welfare estimates do not account for a passenger’s utility from the whole taxi ride experience. We assume that all unobserved aspects of the taxi ride are similar on average. Therefore, our estimates only pertain to the aspect of the taxi trip that involves tipping.

Table 1.4, reports welfare calculations at the taxi trip level for the five scenarios stated above. Columns (1) and (2) report estimates of the utility loss from tipping, and the tip revenue received by drivers respectively. Column (3) is the welfare from tipping (i.e., the sum of columns (1) and (2)). Panel A shows that when passengers are not presented with a tip menu during a taxi trip, the average utility loss from tipping is -\$3.429 (26.68% of the average taxi fare of \$12.17), and the average tip received by drivers is \$1.924 (15.8% of the average taxi fare). Therefore, the welfare loss from tipping on a taxi trip without a tip menu is $-\$3.429 + \$1.924 = -\$1.504$ (12.35% of the average taxi fare).

We consider the estimates from the no-menu case as the baseline and compute changes in the welfare components from tipping under the four remaining tip menus mentioned above. Table 1.4 Panel B shows the results. Relative to the no-menu case, the previous tip menu (15%, 20%, 25%) increases the welfare from tipping by \$1.265 (or an 84% increase). The components of the welfare increment originate from an increase in both the utility from tipping by \$1.097 and tip revenue for drivers by \$0.167. The current tip menu (20%, 25%, and 30%) also increases both tip revenue and the utility from tipping relative to the no-menu case. In particular, the current menu increases the utility from tipping by \$0.80, and the revenue from tips rises by \$0.28. In sum, welfare increases by \$1.08 (a 72% increase) under the current menu relative to the no-menu case. The overall welfare increase under the previous menu is higher than the current menu (\$1.265 versus \$1.081). Relative to the previous menu, consumers lose \$0.297 under the current menu. This loss is composed of a transfer of \$0.114 to drivers, and a deadweight loss of \$0.183. The tip-maximizing menu yields similar results as the current tip-menu.

For the last scenario, we use our model to estimate a three-option tip menu that maximizes the utility from tipping (or minimizes the utility loss from tipping). We follow the same procedure in section 1.7.1 and find that showing 9%, 15%, and 25% is utility maximizing. Unsurprisingly, consumer surplus is highest under this menu (\$1.212) compared to the surplus under all the other menus. However, there is no change in the rate at which passengers tip compared to the no-menu case. The overall welfare under the utility maximizing menu increases by \$1.217 relative to the no-menu case. It is important to note that, forcing passengers not to tip at all decrease welfare the most. For example, with the estimated tipping norm of 20%, the welfare from not tipping at all is $-\theta \times (d - t)^2 = -152.24 \times (0.20 - 0)^2 = -\6.01 . This is four times worse than the no-menu case.³⁵

³⁵The 2014 Taxi fact book reports that there are about 175 million taxi rides annually. To put the trip level

1.8 Conclusions

Firms find that menu suggestions and default options are powerful tools that influence consumers’ behavior. Many influential studies have also examined their use in setting policy. However, few studies have examined the mechanisms at work and the welfare implications of such tools in a field setting. This study focuses on how default tip suggestions in NYC taxis affect consumers’ behavior. The advantage of restricting our study to tipping is that we avoid a number of complications that vexed previous researchers. For example, because customers cannot delay choosing a tip, we do not have to consider behavioral biases due to naiveté, present bias, and procrastination.

We develop a model that allows us to empirically estimate the unobserved beliefs about the social norm tip, the norm deviation cost of not conforming to the social norm, and the cognitive cost of calculating a non-menu tip. Relying on the model, we conduct a nonparametric, semi-parametric, and a parametric analysis of tipping behavior. All these methods provide consistent results. The nonparametric estimate of the average decision cost (a combination of norm deviation cost and cognitive costs) is about \$1.89 (15.53% of the average taxi fare of \$12.17). After accounting for observable factors and trip characteristics using a semi-parametric approach, our estimate of the average decision cost falls to \$1.64 (13.5% of the average taxi fare). We then proceed to separately identify the norm deviation and cognitive costs by adding parametric assumptions to our model of tipping decision-making. We estimate the distribution of passenger beliefs about the social norm tip and it averages around 20% of the taxi fare, which is close the average observed tip (19%). The estimated norm deviation cost varies with the size of the deviation. For example, tipping five percentage points less than the norm imposes a norm deviation cost between \$0.30 and \$0.38 (2.5% to 3.1% of the average taxi fare). The estimated cognitive cost of calculating a non-menu tip ranges from \$1.10 to \$1.32 (9% to 10.4% of the average taxi fare) on average.


We use the parametric model to investigate a number of what-if questions. For example, compared with not presenting a tip menu, the current menu increases the tip revenue by 14.65%, and the overall welfare from tipping by \$1.08 on average per taxi trip. Our simulations also show that the current menu in NYC Yellow taxicabs nearly maximizes tips. In fact, the two Yellow taxi credit card machine vendors (CMT and VTS) appear to have converged over time to present passengers with the tip-maximizing menu.

We believe that our findings are not limited to tipping in taxicabs. Obviously, the tip analysis applies to other service industries such as restaurants, delivery services, bars, and hotels. Our results that the size of norm deviation and cognitive costs are relatively large may be useful in considering more general “nudges,” such as those that are widely used by businesses and policy makers.

welfare estimates in perspective, we can rescale all the estimates in table 1.4, by multiplying by 175 million. For example, the current taxi menu increases the welfare from tipping by about \$190 million annually relative to the no-menu case.


Figures

Figure 1.1: (A) Distribution of Tips (% of Taxi Fare) Before and After Menu Change, (B) Nonparametric Bounds of CDF of Decision Cost (Tip = 10%), (C) Nonparametric Unconditional CDF of Decision Costs




Notes: Panel A shows the distribution of tips in CMT Yellow taxis before and after the menu of tips presented to passengers changed from showing 15%, 20%, and 25% to show 20%, 25%, and 30%. The bars in Panel A are non-overlapping bins of width 1% for tips between 0.5% and 35.5% of the taxi fare. The tips rates are truncated at 35.5% because the share becomes essentially zero. Panel B shows the lower and upper bounds for the CDF of decision costs computed nonparametrically for passengers who tip 10% of the taxi fare (tips rates between 9.5% and 10.5%). Panel C shows the nonparametric estimate of the unconditional CDF of decision costs.

Figure 1.2: (A) Predicted Probability of Tipping 10% (B) Predicted Probability of Tipping 20% (C) Semiparametric Bounds of CDF of Decision Cost (Tip = 10%), (D) Semiparametric and Nonparametric Unconditional CDF of Decision Costs




Notes: Panels A and B show the predicted probabilities for tipping 10% and choosing the menu tip rate of 20% respectively as functions of the taxi fare. The range of fares are between \$3 and \$30. Panel C shows the lower and upper bounds for the CDF of decision costs computed semiparametrically for passengers who tip 10% of the taxi fare (tips rates between 9.5% and 10.5%). Panel D shows the both the semi-parametric and nonparametric estimate of the unconditional CDF of decision costs. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip.

Figure 1.3: Distribution of Beliefs about Tipping Norm



Notes: Panel A shows the estimated distribution of passengers' beliefs about the tipping norm in CMT taxis in 2014. Panel B shows the estimated distribution of passengers' beliefs about the tipping norm before and after the menu of tips presented to passengers changed from showing 15%, 20%, and 25% to show 20%, 25%, and 30% in CMT taxis over the period of 2010-2011. For the left panel, a chi-square goodness of fit test between the distribution of beliefs about tipping norms before and after the menu change yields a χ^2 statistic of = 362900 and a P-value $< 2.2e-16$. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive non-menu tip.

Figure 1.4: CDF of (1) Upper Bound of Cognitive Costs and (2) The Norm Deviation Cost for Non-Menu Tips



Notes: This figure shows the distribution of upper bounds of cognitive costs for passengers who choose non-menu options and the distribution of their norm deviation cost. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive non-menu tip.





Figure 1.5: Model Fit

Notes: These figures illustrate how the parametric model fits the observed data by showing the observed distribution of tips against the model predicted distribution of tips. The left panel shows the fit of the model for passenger tips in CMT taxis in 2014 when the tip menu showed 20%, 25%, 30%. A chi-square goodness of fit test between the model and the observed data yields a χ^2 statistic of = 1726300 and a P-value < 2.2e-16. The middle panel shows the fit of the model for passenger tips in CMT taxis in 2011 right after the tip menu changed from showing 15%, 20%, and 25% to show 20%, 25%, 30%. A chi-square goodness of fit test between the model and the observed data yields a χ^2 statistic of = 1860100 and a P-value < 2.2e-16. The right panel shows the fit of the model for passenger tips in CMT taxis in 2010 when the tip menu showed 15%, 20%, and 25%. A chi-square goodness of fit test between the model and the observed data yields a χ^2 statistic of = 1692300 and a P-value < 2.2e-16. The bars in this all the panels are non-overlapping bins of width 1% for tips between 0.5% and 35.5% of the taxi fare. The tips rates are truncated at 35.5% where the share becomes essentially zero. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip.

Figure 1.6: Grid Search for Tip-Maximizing Menu



Notes: Panels A and B plot the results from a grid search for a one-option menu and a two-option menu respectively that will maximize the average tip received from passengers. Panel C plots the average tip from a grid search of the tip-maximizing menu as a function of the number of menu options. Panel D shows the model prediction of the distribution of tips for a three-option menu that maximizes tips.

Tables

Table 1.1: Taxi Trip Characteristics

	Before Menu Change	After Menu Change	
	<u>1/2010 – 1/2011</u>	<u>2/2011 – 12/2011</u>	<u>2014</u>
	(1)	(2)	(3)
Menu of Tips	[15%, 20%, 25%]	[20%, 25%, 30%]	
Tip (\$)	1.77 (1.88)	1.95 (1.23)	2.27 (1.51)
Taxi Fare (\$)	10.22 (5.33)	10.42 (5.24)	12.17 (6.69)
Tip Rate (% of Taxi Fare)	17.82% (7.77%)	19.19% (8.59%)	19.06% (7.01%)
Menu Tip (%)	18.22% (3.38%)	21.64% (2.96%)	21.40% ((2.93%))
Non-Menu Tip (%)	17.22% (11.40%)	16.70% (11.12%)	15.75% (17.46%)
Share of Menu Tips	59.7%	48.3%	60.6%
Observations	28,305,969	31,227,439	41,620,454

Notes: This table presents means (standard deviation) across different trip characteristics. Column (1) presents trip characteristics one year before the CMT menu change, and column (2) presents trip characteristics about a year after the change. Column (3) presents trip characteristics four years after CMT’s menu change. The data used are standard rate NYC Yellow taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip.

Table 1.2: Structural Estimates

	<u>2014</u> (1)	<u>2010 – 2011</u> (2)
<i>Panel A: Heckman Selection Estimates of α_T and θ</i>		
	Dependent Variable: Tip Rate	
Taxi Fare	–0.00328*** (0.00001)	–0.00423*** (0.00001)
Inverse Mills Ratio	0.00481*** (0.00034)	0.00320*** (0.00020)
1(Post Menu Change)		–0.00640*** (0.00005)
Constant	0.19801*** (0.00051)	0.20220*** (0.00027)
Norm Deviation Cost Parameter $\hat{\theta} \left(= -\frac{1}{2\beta} \right)$	152.2359	118.298
1 st -Stage Instrument	Number of Passengers	Number of Passengers
1(Round Number Tip)	Yes	Yes
Observations with Non-Menu Tips	16,394,917	25,206,358
R ²	0.01641	0.04198
<i>Panel B: Simulated Method of Moments Estimates of c_i</i>		
	Mean Cognitive Cost (\$): $c_i = 1/\lambda$	
1 st -Step Estimate (weight matrix: $\widehat{W} = I$)	1.38021*** (0.00522)	1.1073*** (0.00262)
2 nd -Step Estimate (weight matrix: $\widehat{W} = [\text{diag}\{gg'\}]^{-1}$)	1.33586*** (0.00531)	1.13705*** (0.00873)
Observations	5,000,000	5,000,000

Notes: This table reports estimates of the primitives in the structural model. In column (1), we use CMT taxi trips from 2014, and in column (2), we use CMT taxi trips from 2010-2011. Panel A reports estimates of the tipping norm T_i and the norm deviation cost parameter θ from the second step of the 2-step Heckman selection correction model. Using a Simulated Method of Moments algorithm, Panel B reports estimates of the two-step procedure of estimating the cognitive costs incurred by passengers when they opt to compute their preferred non-menu tip. We use the whole analysis sample in Panel A. However, in Panel B, we randomly sample five million observations each from 2014 and 2010-2011 respectively to reduce the computational burden of the SMM algorithm. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip. The standard errors in Panel A are robust white standard errors and in Panel B, the standard errors are computed as the standard deviation of the distribution of parameter estimates computed from 1000 bootstrap samples. *p<0.1, **p<0.05, ***p<0.001.

Table 1.3: Evolution of Menu Tip Options

	CMT (1)	VTs (2)
<i>Panel A: Trips from Jan 2010 - Jan 2011</i>		
Tip menu for taxi fare < \$15	[15%, 20%, 25%]	[\$2, \$3, \$4]
Tip menu for taxi fare $\geq$ \$15		[20%, 25%, 30%]
Average tip for all fares	17.81%	20.68%
Average tip for fare < \$15	18.00%	21.19%
Average tip for fare $\geq$ \$15	17.51%	16.61%
Observations	27,574,410	28,658,477
<i>Panel B: Trips from Feb 2011 - Dec 2011</i>		
Tip menu for taxi fare < \$15	[20%, 25%, 30%]	[\$2, \$3, \$4]
Tip menu for taxi fare $\geq$ \$15		[20%, 25%, 30%]
Average tip for all fares	19.16%	20.66%
Average tip for fare < \$15	19.37%	21.21%
Average tip for fare $\geq$ \$15	18.00%	17.66%
Observations	31,960,044	30,339,659
<i>Panel C: Trips from 2013 - 2014</i>		
Tip menu for all taxi fares	[20%, 25%, 30%]	[20%, 25%, 30%]
Average tip for all fares	19.07%	18.55%
Average tip for fare < \$15	19.42%	18.74%
Average tip for fare $\geq$ \$15	17.96%	17.93%
Observations	83,107,354	84,332,924

Notes: This table reports the average tip rate across the different menus of tips presented to passengers in NYC Yellow taxis over time. Panels A through C correspond to one of three periods where at least one of the two Yellow tax credit card machine providers (CMT and VTS) changed the menu of tips presented to passengers. Column (1) shows the average tip rate offered by passengers in CMT cabs. Column (2) is analogous to column (1) but for passengers in VTS cabs. Each panel also reports the average tip rate separately for trips where the taxi fare is less than \$15. Only standard rate taxi trips, with no tolls, paid for via a credit card machine where passengers leave a positive tip are used in this table. The sample restriction are standard rate taxi trips paid via a CMT or VTS credit card machine along with a positive tip.

Table 1.4: Welfare Estimates (Trip Level)

	Utility (Loss) from Tipping Consumer Surplus (CS) (1)	Tip Revenue Producer Surplus (PS) (2)	Welfare CS + PS (3)
<i>Panel A: Baseline</i>			
No tip menu	-\$3.429	\$1.924	-\$1.504
<i>Panel B: Change relative to no tip menu</i>			
Previous tip menu [15%, 20%, 25%]	\$1.097	\$0.167	\$1.265
Current tip menu [20%, 25%, 30%]	\$0.800	\$0.281	\$1.081
Tip-maximizing menu [20%, 26%, 32%]	\$0.802	\$0.282	\$1.084
Consumer utility- maximizing menu [9%, 15%, 25%]	\$1.212	\$0.005	\$1.217

Notes: This table reports estimates of the effect of different tip menus on social welfare at the taxi trip level. In column (1), we use the parametric estimates to compute the dollar value of a passenger's utility from tipping. In column (2), we compute the tip revenue that drivers receive. Social welfare is calculated in column (3) as the sum of the utility that consumers get from tipping and the tip revenues that drivers receive. The utility from tipping is always less than zero, even for the case where the passenger decides not to leave a tip. This is because, in addition to the payment of the tip to the driver, the consumer incurs decision costs (norm deviation or cognitive costs).

Chapter 2

Lower Income, Higher Gratuity: Evidence from Tipping in NYC Taxis

2.1 Introduction

It is challenging to establish a causal effect of wealth on generosity or on the concern for others (pro-sociality). This is because, it is difficult to randomize wealth. Analyzing about five million New York City (NYC) taxi trips with actual tipping choices, we document a negative relationship between passenger tips and wealthier neighborhoods. We use a model of tipping behavior to empirically identify economically relevant parameters that explain how wealth relates to tipping.

Related studies of tipping posit that racial minorities are poor tippers compared to whites (Lynn and Thomas-Haysbert, 2003). Income and tipping perceptions have been suggested as explanations for the racial gap in tipping (Lynn, 2011; Lynn et al., 2012; Lynn and Brewster, 2015b,a; Lynn, 2004; Thomas-Haysbert, 2002). These studies mostly rely on surveys of servers or online survey respondents who report their preferred tipping choices when presented with hypothetical dinning situations. This study differs from most previous studies in that we study actual tipping choices from taxi passengers in NYC.¹ We use a structural model of tipping to recover passenger tipping preferences and how they relate to observed tips. We do not study race because it is unobserved in our data.

More broadly, most past studies find a negative relationship between wealth and pro-sociality while others find the opposite. Using lab experiments, Piff et al. (2010) document that participants who report higher socioeconomic statuses are less likely to engage in pro-social behaviors such as: generosity, charity, trustworthiness, and helpfulness.²

Other studies document mixed findings or that the rich are more pro-social than the poor. Factors correlated with high social classes such as income and educations are found to be positively related to donations (Gittell and Tebaldi, 2006). A few studies also find a U-shaped relationship between income and cash and non-cash charitable giving (James and

¹A related study in our context is Haggag and Paci (2014). Analyzing 2009 NYC Yellow taxi data, the authors find that providing passengers with higher tip suggestions increases the amount tipped by passengers.

²Several other studies present similar findings (e.g., Wang and Murnighan (2014); Wang et al. (2014); Grossmann and Varnum (2011); Dubois et al. (2015)).

Sharpe, 2007; Liddell and Wilson, 2016; Schervish and Havens, 1995b).³

We study the relationship between income and tipping. Tipping is a significant economic activity. Each year, there is about \$37 billion in tips from restaurants alone. This is about five percent of the projected sales from restaurants in 2019. Tipping is an interesting economic activity to study because it presents consumers with an opportunity to engage in prosocial behavior that is not obligatory. How much one tips also depends on his or her beliefs about the social norm tip and the guilt or shame associated with giving less than the norm. These factors may differ across passengers with different income levels and hence affect their tipping choices differently.⁴

We use a newly available dataset from NYC Green taxis. Green taxis were put in place to address the lack of access to legal street hail taxi services (Yellow taxis) from outside the core of Manhattan. Compared to Yellow cabs, Green cabs presents an opportunity to study how passenger tipping behavior varies across the different boroughs of NYC and by income level.

NYC has stark income inequality. For example, the median household income in Manhattan is more than twice that of the Bronx. There is also substantial variation in income across different neighborhoods within the boroughs of NYC. From an econometric standpoint, it is more desirable to assess a homogeneous group. So, while we are able to look at the differences in passenger tipping behaviors across boroughs, we are also able to look at the differences within boroughs.

We find that passengers in boroughs with higher incomes give less generous tips, and the tips are also lower in the wealthier neighborhoods within all the boroughs of NYC. Endogeneity problems such as social pressures, social signaling, and contextual factors distort choices (Dellavigna et al., 2012). These distortions may have differential impacts on the rich versus the poor. So, while we control for borough specific effects in our analysis, we are unable to control for unobserved differences between the rich and the poor. Thus, the negative relationship that we document between wealth and gratuity is not purely causal.

To explain the negative relationship between income and tipping, we model how passengers tip in NYC taxis. A machine in the cab presents passengers with a set of menu tip options (20%, 25%, and 30%) at the end of the ride. Passengers can choose one of the menu options, key in their preferred custom tip, or decide not to tip at all. In our model, passengers' choices are based on their beliefs about the social norm tip. They incur a norm deviation cost for not conforming to the norm and a cognitive cost from computing a non-menu tip. We allow both the norms and the cost parameters to vary by income level.

The model allows us to recover the unobserved tipping preferences of passengers. First, we estimate a distribution of the unobserved beliefs about the tipping norm. Second, we estimate a norm deviation cost, which is a measure of how binding tipping norms are. Last, we estimate the cognitive cost of actively computing one's preferred tip. These parameters together provide some insights into why we observe a negative relationship between wealth and gratuity.

We find that the average belief about the social norm tips in the wealthiest neighborhoods

³Including donors and non-donors in the analysis has been found to transform the U-shaped relationship into a positive relationship (Schervish and Havens, 1995a).

⁴Other factors that may affect tipping include service quality, the gender, race, posture etc. See (Azar, 2010; Lynn, 2007) for reviews.

is 20.57% of the taxi fare. In contrast, the average belief about the tipping norm among the lowest income group is 29.78% of the taxi fare. Thus, the average tipping norm in the lowest income group is 45% higher than that of the wealthiest group.

In contrast, we find that the cost of deviating from one’s belief of the tipping norm is highest among the wealthy. For example, tipping five-percentage points less than one’s belief of the norm is associated with a norm deviation cost is \$0.26 (2.1% of the average taxi fare of \$12.29) among the wealthiest passengers versus \$0.14 (1.1% of the average taxi fare) among the lowest income group.

The cost of actively computing one’s tip is also higher in wealthier neighborhoods, which may reflect a higher opportunity cost of time for the rich versus the poor.

Unlike Andreoni et al. (2017), our findings suggest that financial stressors do not crowd out pro-social behavior—at least in the context of tipping.⁵ While our model does not establish a causal link between wealth and gratuity, it suggests that the less wealthy have preferences of a more generous gratuity.

Our study also contributes to the study of preference identification (for example, Bernheim and Rangel (2009); Rubinstein and Salant (2011); Benkert and Netzer (2018); Goldin and Reck (2019)). This study recovers unobserved behavioral primitives that are important in explaining some of the differences in prosocial behaviors across different wealth groups (Keltner et al., 2014).

Contrary to studies that suggest that poverty inhibits decision making (e.g., Mullainathan and Shafir 2013; Mani et al. 2013; Carvalho et al. 2016), we find no evidence of a negative effect in our setting. In fact, we find that the cost of calculating one’s preferred tip versus choosing from a menu is higher among the wealthy. However, our finding may reflect the fact that the wealthy may have a higher opportunity cost of time compared to the less wealthy.

We describe the context, data and how we identify rich and poor passengers in the next section. The following section estimates and discusses the relationship between income and tipping. We then present and empirically estimate a model of passenger tipping that we use to explain the relationship between income and tipping. The final section concludes.

2.2 Context and Data

Prior to the summer of 2013, Yellow cabs were the only taxis which could legally pick up passengers in response to street hails in NYC. Using GPS data, the Taxi and Limousine Commission (TLC) found that 95% of taxi pickups only occur in Manhattan below 96th street, and at John F Kennedy and LaGuardia airports. Thus, other NYC boroughs were under served and lacked access to safe and legal taxi rides. In response, the State Livery Law was passed and signed into law in December 2011. The new law allowed the TLC to issue taxi permits that allowed new taxis called the Boro (Green) taxis to provided services to the underserved areas in NYC.

⁵In a field experiment where semi-transparent envelopes containing money are intentionally misdelivered to different households, Andreoni et al. (2017) observe if the envelopes are returned. They find that the wealthier households are twice as likely to return envelopes relative to the less wealthy. Because this study uses an experimental treatment, they can control for some of the selection concerns that plague most studies on wealth and pro-sociality.

The service provided to passengers by Yellow and Green taxicabs are essentially identical with two main exceptions. First, Yellow taxicabs are painted yellow, and Green taxicabs are painted green. Second, both Yellow and Green taxicabs are permitted to pick up passengers in the Bronx, Brooklyn, Queens, Staten Island and Northern Manhattan. However, Green taxis cannot pick up passengers in the Hail Exclusionary Zone (HEZ). That is, Manhattan south of West 110th street and East 96th street, John F Kennedy Airport, and LaGuardia Airport.

Unlike the Yellow taxis, Green cabs pick up a significant number of passengers across the NYC boroughs. For pickups in Green (Yellow) taxis, 38% (2%) are in Brooklyn, 28% (92%) are in Manhattan, 28% (1%) are in Queens, 6% ($< 1\%$) are in the Bronx, and less than 1% ($< 1\%$) are in Staten Island (Gonzalez et al. 2016). Green cabs provide a unique opportunity to study passenger behavior across all the five NYC boroughs.

We use taxi trip information from Green taxis as the main data source. The TLC provides a dataset that has information on all active taxis since August 2013. We use data from August 2013-2014. We do not use data beyond 2015 because a 30-cent surcharge was introduced. The surcharge may affect the behaviors of passengers. We also exclude Staten Island from the analysis because the data are sparse.

There were about 12,873 drivers and 7,159 green taxis by the end of 2014. The Green taxi data has records for 17,047,812 trips over the study period. We restrict the sample to standard rate fares with no tolls where passengers tip a positive amount. Because passenger tipping information is available for only credit card transactions, we drop trips with cash transactions. It is not uncommon that passengers pay for the taxi fare using a credit card and then tip using cash. Thus, we leave out zero tips from the analysis.

Because Green cabs are not allowed to pick up passengers in the core of Manhattan, our analysis of tipping will be incomplete for passengers in Manhattan. We supplement our data by adding a random sample of 500,000 Yellow taxi trips from Manhattan. The Yellow taxi data is also collected by the TLC and has the same information as in the Green taxi data.

The data reports the taxi fare, tip, trip distance, geocodes for pickup and drop-off locations, date and time of trip, and other trip characteristics. The trip expenses in a NYC taxi are as follows. For standard rate fares, passengers are charged \$2.50 and a 50-cent MTA tax upon entering the cab. Then for every fifth of a mile or for every minute where the cab travels less than 12mph the fare increases by an additional 50 cents. There is also an additional 50 cent night surcharge charge for trips between 8pm and 6am and a \$1 surcharge for trips picked up between 4pm and 8pm on weekdays.

When a customer pays for her taxi fare with a credit card, a screen shows the fare details and suggests three menu tip options. The screen also provides the option for the passenger to input her preferred tip or no tip instead. Creative Mobile Technologies (CMT) and VeriFone Inc. (VTS) are the two vendors that provide the electronic screen and devices that collect credit and debit card payments in both Yellow and Green taxicabs. Both vendors presented passengers with the same tip menu suggestions: 20%, 25%, and 30%. The only difference across the two vendors is how the menu tips are computed. CMT percentage tips are calculated on the total fare (i.e., fare + MTA tax + Surcharge) and VTS calculates tips on only the fare plus surcharge. The trip data only provide the dollar amount tipped by passengers. We classify a tip as a menu tip if the dollar amount of the tip as a percentage of the fare is equal to a menu option.

2.2.1 Identifying Low Income and Wealthy Passengers

There are no data on passenger characteristics. Thus, we do not know where passengers reside or their income. For each trip, we use the census tracts at the pickup and drop-off points to assign location and income. To increase the likelihood that passengers reside in a specific borough, we limit our analysis to trips that start and end within the same NYC borough. As a proxy for a passenger’s income, we use the average of the median household income at the pickup and drop-off census tract locations. The income data we use is from the 2010 American Community Survey. For customers who don’t start or end in areas where they reside, and for tourists, this measure would be inaccurate. Although these approximations are imperfect, this is the best we can do given the data. Our final analysis sample has a total of 4,734,022 taxi trips.

2.2.2 NYC Income Disparities

NYC has stark income inequality and diversity across and within its boroughs. For example, in 2014, Manhattan had the highest median household income (\$78,863), then Staten Island (\$73,714), followed by Queens (\$59,328), then Brooklyn (\$49,715), and the lowest is the Bronx (\$34,941). The median household income in Manhattan is 2.26 times that of the Bronx. Figure 2.1 shows the distribution of median household incomes across census tracts within the boroughs of NYC. Unsurprisingly, incomes in Manhattan are generally higher than all other boroughs. In contrast, incomes in the Bronx are lower. Although incomes differ across boroughs, there is substantial variation within each borough. The inequality and disparities across and within the boroughs of NYC present a good opportunity to examine differences in the tipping behavior of taxi passengers across location and income.

2.2.3 Trip Characteristics

We compare taxi trip characteristics across the boroughs of NYC. Table 2.1 presents a summary. The average tip rate is lower in boroughs with higher average median household incomes. The average tip rate is 18.69% in Manhattan, 20.82% in Queens, 20.91% in Brooklyn, and 24.28% in the Bronx. However, passengers across all boroughs tip around \$2.30 – \$2.40 on average. The taxi fares across boroughs are in a similar range as well, between \$11.29 and \$12.77. Passengers in Brooklyn are 78% more likely to choose menu tips compared to Queens (68%), the Bronx (65%), and Manhattan (62%).

The main take away from table 2.1 is that passengers in boroughs with lower incomes are more generous in tipping. For example, while the average median household incomes in Manhattan is more than double (2.26 times) that of the Bronx, the tip rate in the Bronx is 30.9% higher than tips in Manhattan.

Although passengers in boroughs with lower incomes give relatively higher tips, it is not evident that the less wealthy areas are more generous. Passengers may be very different across boroughs. Therefore, we analyze the tips of passengers within the same borough. Passengers in the same borough likely share common characteristics compared to those in other boroughs.

2.3 Income and Tipping within NYC Boroughs

There is income inequity across NYC boroughs (table 2.1). However, figure 2.1 shows evidence of significant disparities in median household incomes within boroughs as well. We take advantage of the variation within boroughs to assess whether passengers who pick taxis in poorer neighborhoods within a borough are more generous when tipping.

Figure 2.2 shows the distribution of tipping rates across census tracts within each NYC borough. A close inspection of the figure shows the general patterns of lower tips in boroughs with higher incomes. For example, passengers picked up in most census tracts within Manhattan tip 20% of the taxi fare or less. In the Bronx and Brooklyn, passengers mostly tip above 20% of the taxi fare. However, there is substantial variation in tipping rates within most boroughs.

We estimate the relationship between income and tipping within the boroughs of NYC using a linear regression. We take the average median household income at a passenger's pickup and drop-off location as a proxy for income.⁶ We regress the tip rate on the fare, income, a dummy variable for each borough, and the interactions of income and borough dummies. We use the interactions to determine the relationship between income and tipping within each borough. A negative coefficient on the interaction terms will suggest a negative relationship between income and the generosity of passenger tips.

We control for hour of day, day of week, month, year, vendor, and type of taxi. Dummies for holidays and hourly measures of temperature, temperature squared, precipitation and precipitation squared are included in our specification.⁷

Results

Table 2.2 reports the coefficients from our OLS specification. Passengers reduce their tip rate by 0.22 percentage points for every dollar increase in the taxi fare. All the coefficients on the borough dummies are in line with what we would expect. Relative to passengers in Manhattan who on average tip 18.69%, passengers in Queens tip 2.42 percentage points more, passengers in Brooklyn tip 2.72 percentage points more, and those in the Bronx tip 6.65 percentage points more.

Income is negatively related to tipping across all boroughs. The interaction of income and borough dummies suggest that a ten thousand dollar increase in income decreases the tip rate by 0.085 of a percentage point in Manhattan, 0.201 of a percentage point in Queens, 0.211 of a percentage point in Brooklyn, and 0.6 of a percentage point in the Bronx.

In a robustness exercise, we check for nonlinearities in the effect of income on tipping. Including a squared term to the equation suggests a U-shape relationship which bottoms out at an income of \$94,231. However, adding a cubic term suggests that the relationship between income and tips is generally downward sloping. Because of the high variability in tip rates, we truncate the top and bottom five percent of tips within each borough to check if the results still hold. For this exercise, we do not observe a negative relationship between

⁶Our estimates remain largely the same when we use the income from either the pickup or drop off location for our analysis.

⁷Our weather data is from Central Park and LaGuardia Airport weather stations: <https://www.wunderground.com>.

income and tips within all boroughs. Thus, the top and bottom five percent of tips may drive some of our results.

To summarize, passengers who take trips in less wealthy neighborhoods are more generous in tipping than their counterparts in richer areas. Interestingly, boroughs with lower incomes have higher average tips and a stronger negative relationship between income and tips. In the following section, we explore some of the reasons why the less wealthy may give more generous tips.

2.4 Possible Reasons the Less Wealthy Tip More

The literature in psychology and sociology posits that one’s identity is a primary determinant of behavior. Individuals identify with different social categories. They associate with groups who have shared interests or characteristics. However, affiliation to a group may require that one’s behavior or commitment aligns with set rules or guidelines. One therefore risks being shunned by group members, punishment, or guilt for not conforming to group standards. Thus, a sense of self and belonging is bound in conforming to group prescriptions which in turn affects choices and economic outcomes (Akerlof and Kranton, 2000).

With identity being a driver of behavior, we briefly discuss three potential reasons why we observe that passengers in less wealthy areas of NYC are more generous in tipping taxi drivers than those in wealthier areas.

First, passengers in less wealthy neighborhoods likely belong to similar income brackets or work in similar occupations where tips are a significant part of one’s income. As a result, they are more aware of the needs of drivers and therefore more generous in tipping relative to passengers in wealthier neighborhoods. Second, passengers in less wealthy areas may tip more to bolster their sense of self or salve a diminishing self-image (Akerlof and Kranton, 2000). Individuals in less wealthy areas have fewer resources and therefore more likely to put their self-interest first. So, they may believe that they are perceived as less generous. To compensate or counter such perceptions, passengers in poorer neighborhoods may choose to tip more. They may also think of tipping as a status good that is relatively less expensive compared to other status goods and therefore consume more of it (Landis and Gladstone, 2017). Third, passengers in wealthier neighborhoods may view giving lower tips as a reflection of their bargaining prowess. Thus, giving lower tips bolsters their sense of self.⁸

The reasons provided above are not exhaustive but are potential explanations for why we observe a negative relationship between income and tipping. It is clear that one’s level of wealth may change his or her incentives which may in turn affect his or her behavior. Thus, the question of whether wealthier individuals are inherently more selfish or less pro-social cannot be parsed as a *ceteris paribus* comparison. Endogeneity problems such as social pressures, social signaling, and contextual factors distort choices.

In sum, we are unable to cleanly estimate the causal effect of wealth on tipping. In the following sections, we use a model of tipping behavior to empirically identify economically relevant parameters that shed light on the adverse relationship between income and tipping.

⁸Thanks to Zakaria, a veteran cab and limousine driver in the Bay Area who in conversation, brought up this point as being his experience over his decades of driving.

2.5 Model

NYC taxis present passengers with a set of menu tip options (20%, 25%, and 30%) at the end of the ride. Passengers can choose one of the menu options, key in their preferred custom tip, or choose not to tip at all. We reason that passengers' choices are based on their beliefs about the social norm tip. They incur a norm deviation cost for not conforming to the norm and a cognitive cost from computing a non-menu tip. We make the following assumptions for this exercise.

First, one's belief about the tipping norm is jointly independent of the menu of tips and the taxi fare. Second, norm deviation and cognitive costs are independent of the menu of tips, the taxi fare and one's preferred tip. Empirical evidence in support of these assumptions are presented in the appendix, section A.2.

At the end of the taxi ride that costs F_i , passenger i gives a tip of $t_i\%$, which may differ from what she believes to be the social norm tip is $T_i\%$. Both t_i and T_i may differ across passengers. Instead of computing her tip, passenger i can pick $d_j \in D$, where d_j is one of $j = 1, 2, 3$ alternatives in tip menu D . The tradeoff faced by the passenger is to compute her preferred tip and incur a cost c_i or choose a menu option and avoid c_i . Passenger i chooses t_i to maximize her utility represented by

$$\text{Max}_{t_i} U = \underbrace{-t_i F_i}_{\text{Tip paid}} - \underbrace{\theta (T_i - t_i)^2}_{\text{Norm deviation cost}} - \underbrace{c_i \times \mathbf{1}\{t_i \notin D\}}_{\text{Cognitive cost}} \quad (2.1)$$

$t_i F_i$ is the dollar amount of the tip paid. $\theta(T_i - t_i)^2$ is the norm deviation cost—disutility for not conforming to the norm—which we assume to be quadratic. θ is the norm deviation cost parameter. The norm deviation cost is zero if the passenger i tips T_i . However, the norm deviation cost increases with the size of the percentage point deviation from T_i . $c_i \mathbf{1}\{t_i \notin D\}$ captures passenger i 's cognitive cost for computing a non-menu tip. c_i is a fixed cost that differs across passengers and $\mathbf{1}\{t_i \notin D\}$ equals one if t_i is not an option in D . The maintained assumption is that passenger i always computes her tip if she does not choose from the menu.⁹

So far, the current model is identical to the one in chapter 1, section 1.6.1. To analyze the tipping decision process for passengers with different income levels, we extend the model by allowing all parameters to depend on one's income level.

We assume that passenger i 's belief of the social norm is $T_i = T_i^* + \alpha Y_i + \varepsilon_i$. Thus, passenger i 's belief of the social norm tip has a private component T_i^* and an income component Y_i . α is the marginal effect of income on one's belief about the social norm tip. Our earlier analysis suggests that α is negative. ε_i is a residual term with zero mean. In addition, we assume that one's norm deviation cost and cognitive cost depends on income (i.e., is $\theta(Y_i)$ and $c_i(Y_i)$). Ignoring the cognitive cost, the first-order condition is

$$t_i^* = T_i^* + \alpha Y_i - \frac{0.5}{\theta(Y_i)} F_i + \varepsilon_i \quad (2.2)$$

⁹Passengers may also use other heuristics such as tipping a fixed dollar amount or round off the fare to a specific dollar amount e.g., a passenger facing a fare of \$9 many decide to tip \$1 to round off her total trip expense to \$10. We account for such behavior when estimating the model.

Passenger i 's preferred tip t_i^* is less than her belief of the social norm T_i . Therefore, when deciding on how much to tip, a passenger trades off the dollars lost to tipping at the social norm against the shame from being a cheapskate.

We estimate the parameters of the model in two steps. First, we use a heckman selection model to estimate passengers' beliefs about the tipping norm and the norm deviation cost. Second, we use a simulated method of moments (SMM) algorithm to estimate the cognitive cost.

2.5.1 Estimating Tipping Norms and Norm Deviation Costs

All the parameters in equation (2.2) can be estimated by regressing the observed tip rate on a constant term, income and the taxi fare. The constant term is the population mean of the passengers' private component of the tipping norm $E[T_i^*]$, the coefficient on income is the empirical analogue of α , and the coefficient on fare $\frac{0.5}{\theta}$ gives an estimate of the reciprocal of the norm deviation cost parameter. Another advantage here is that the residual term $\varepsilon_i = t_i - T_i^* - \alpha Y_i + \frac{0.5}{\theta(Y_i)} F_i$ can be rewritten as $\varepsilon_i + T_i^* + \alpha Y_i = t_i + \frac{0.5}{\theta(Y_i)} F_i = T_i$. Therefore, the constant term plus an estimate of α times income, plus the residual is an estimate of the unobserved realization of passengers' beliefs about the tipping norm.

There are two main challenges. First, because the norm deviation cost is a function of income, we cannot recover an estimate of $\theta[Y_i]$ without making functional form assumptions. Second, there is a challenge with using all observed tips in the data to estimate equation (2.2). This is because, for passengers who choose tips from the tip menu, we do not know what they would have chosen without the menu.

To address the first concern, we categorize income into four groups. The first group are passengers with median household incomes less than \$35K, the second group have incomes between \$35K and \$70K, the third group has incomes between \$70K and \$100K, and the fourth group are passengers with incomes above \$100K. We then regress the tip rate on the income group dummies and their interactions with the fare. The coefficient on the income group dummies allow us to identify level differences in average beliefs about the norm tip across income groups. The coefficients on the interaction represent differences in the norm deviation costs across the four income groups.

To address the second concern, we use the Heckman selection correction technique. We use the subsample of passengers who choose non-menu tips to estimate equation (2.2) while correcting for sample selection.

2.5.1.1 The Heckman Selection Model

In the first step of the selection model, we use a probit model to estimate the probability of choosing a non-menu tip using the entire sample. We use the number of co-passengers on a taxi trip as a *decision quality instrument* to correct for sample selection (Goldin and Reck, 2019).

The assumption here is that the instrument affects a passenger's decision to choose a menu option, but it does not affect a passenger's preferred tip and her cognitive cost of calculating a non-menu tip. The reasoning here is that passengers are likely engaged in conversations and under more time pressure when riding with co-passengers. As a result,

they are more likely to choose a menu tip when exiting the cab as taxi meter stops at the end of the ride and the transaction screen is then presented.

Our assumptions are likely violated if passengers split the bill. One’s private preference may differ from that of the group. However, the number of passengers on the ride does not enter into one’s utility function for tipping. Thus, per the model, the number of co-passengers is an excluded instrument. As added instruments, we also include variables such as an indicator for tips of round dollar amounts, borough fixed effects, hour of day, day of week, month, and year fixed effects.

In the second step, we use the subset of non-menu tips to estimate equation (2.2) using the income group dummies and their interactions with the fare. We estimate the Inverse Mills Ratio from the first stage probit regression, which is included in the equation to correct for sample selection. We also include borough fixed effects to parse out borough specific effects on the social norm tip.

2.5.1.2 The Effect of Income on Beliefs about Tipping Norms

Passengers in less wealthy neighborhoods have higher tipping norms relative to their counterparts in wealthier areas. However, tipping norms are more binding in wealthier areas than in less wealthy neighborhoods.

Table 2.3 reports estimates from the first stage Heckman selection probit equation. Our instrument is relevant; the number of co-passengers increases the likelihood of choosing menu tips. Table 2.4 panel A reports the Heckman selection estimate of beliefs about tipping norms and the norm deviation cost parameters.

The average belief about the tipping norm is to tip 20.57% of the taxi fare in areas with median household incomes over \$100K and in areas with incomes between \$70K and \$100K. The average belief of the tipping norm is 23.61% in neighborhoods with incomes between \$35K and \$70K, and 29.78% in areas with incomes below \$35K. Thus, the average tipping norm in the lowest income group is 45% higher than that of the wealthiest income group. Figure 2.3 panel A shows the distribution of beliefs about the tipping norm for each income group.

We turn to estimates of how binding passengers’ tipping behaviors are to their beliefs about the tipping norm. We measure this using their “shame” or cost of deviating from their belief of the norm.

The norm deviation cost parameter θ is derived from the coefficient on the fare. Suppose the coefficient on the fare is β , then the norm deviation cost parameter $\theta = -\frac{0.5}{\beta}$. The norm deviation cost is calculated as θ times the squared percentage point deviation of one’s tip from one’s belief of the norm. For example, tipping five percentage points less than one’s belief of the tipping norm is associated with a dollar cost of $\theta \times 0.05^2$. The norm deviation cost increases with the size of the percentage point deviation from one’s belief of the norm. For example, the dollar value of the norm deviation cost for tipping ten percentage points less than one’s belief of the norm ($\theta \times 0.1^2$) is four times higher than tipping five percentage points less ($\theta \times 0.05^2$).

The cost of a five-percentage point deviation from one’s belief of the social norm tip is \$0.26 ($= 105.26 \times 0.05^2$) for passengers with incomes above \$100K, \$0.28 ($= 111.61 \times 0.05^2$) for passengers in areas with incomes between \$70K and \$100K, \$0.20 ($= 80.39 \times 0.05^2$) for

passengers in areas with incomes between \$35K and \$70K, and \$0.14 ($= 50.87 \times 0.05^2$) for passengers in areas with incomes less than \$35K. The norm deviation cost ranges from 2.1% of the average taxi fare of \$12.29 among passengers in wealthier neighborhoods to 1.1% of the average taxi fare for passengers in low income areas. All estimates are statistically significantly different from zero at the one percent level. In addition, the norm deviation cost in neighborhoods with incomes between \$70K and \$100K are not statistically different from that of those in neighborhoods with incomes above \$100K.

With an estimate of the distribution of beliefs about the tipping norm and the norm deviation cost parameter, we can use equation (2.2) to compute the tips that passenger will give in the absence of menu tip options. Panel B in figure 2.3, shows the distribution of preferred non-menu tips by income level. Unsurprisingly, the preferred non-menu tips for passengers in wealthier neighborhoods are generally lower than their counterparts in less wealthy neighborhoods.

To summarize, on average, beliefs about the social norm tip is higher in less wealthy neighborhoods than in richer areas. In contrast, the cost of deviating from the norm is higher for passengers in wealthier neighborhoods than for their counterparts in less wealthy areas.

2.5.2 Estimating Cognitive Cost

With the income level estimates of the distributions of beliefs about the social norm tip and the corresponding norm deviation cost parameters in hand, we proceed to estimate the cognitive cost c_i of computing a non-menu tip. We use a SMM approach for this exercise.

We proceed with a simulation for the following reasons. The objective of a passenger (equation (2.1)) is to give a tip that maximizes her utility. However, because of the indicator function $1\{t_i \notin D\}$ in the assumed utility function, there is no analytical solution for c_i . Thus, we use a Monte Carlo algorithm to find the parameters of an assumed distribution of c_i such that the model predicts a realization of tips that matches the observed data as closely as possible.

We estimate the cognitive cost separately for each of the four income groups. We follow the same estimation procedure as in chapter 1. We assume that cognitive cost c_i is exponentially distributed with rate parameter λ . As moments to identify λ , we divide tips into non-overlapping one percent bins, namely 1%, 2%, 3%...35%. Each bin is a moment: defined as the share of passengers whose tip falls within that respective bin.

The share of passengers who choose menu options identifies λ . The intuition here is that, in the absence of cognitive costs, we should not find a large share of passengers choosing the menu tip options.

We take a two-step approach to estimate λ . In the first step, we use the identity matrix as the waiting matrix to estimate λ . In the second step, we use the first step estimates to compute a new weight matrix. The new weight matrix is the diagonal of the inverted variance-covariance matrix of the moment conditions. We then use the estimated weight matrix in our final estimation of λ . We estimate standard errors as the standard deviation of the distribution of parameter estimates from 1000 bootstrap samples.

SMM Estimates of Cognitive Costs

Table 2.4, panel B presents estimates of the cognitive cost of computing non-menu tips by income group. Cognitive costs are much higher among passengers in the top two income groups (above \$70k) compared to those in less wealthy neighborhoods. Specifically, the average cognitive cost is \$1.49 (12.12% of the average taxi fare of \$12.29) for passengers in neighborhoods with incomes above \$100K, \$1.60 (13% of the average taxi fare) for those in neighborhoods with incomes between \$70K and \$100K, \$1.12 (9% of the average taxi fare) for those in neighborhoods with incomes between \$35K and \$70K, and \$1.21 (9.8% of the average taxi fare) for those in neighborhoods with incomes below \$35K.

2.5.3 Model Fit

Using the estimated parameters, the model predicts tips that are consistent with the observed data. Figure 2.4 shows the model predicted tips against the observed tips by income level. The model performs well predicting both the share of passengers who choose menu options and non-menu options as well. However, a chi-squared goodness of fit test rejects that the observed data and the model predictions have no significant differences.

In sum, the model suggests that the less wealthy have beliefs aligned with being more generous towards tipping cab drivers compared to the wealthy. This behavior is in line with the observation of an adverse relationship between income and gratuity. On the other hand, it is more costly to deviate from norms among the wealthy than the poor. This explains why there is more variability in tips among the poor versus the rich. The relatively higher cognitive cost of computing non-menu tips may reflect a higher opportunity cost of time for the wealthy passengers compared to passengers in low-income areas.

2.6 Conclusions

This study finds that taxi drivers receive more generous tips in the low-income neighborhoods compared to the wealthier parts of NYC. This finding holds within each borough as well. There are several reasons why we may observe this relationship. Less wealthy passengers identify with drivers, are more aware of their needs, thus give more generous tips. The less wealthy may view tipping as a status good and thus use it to signal that they are not poor or salve perceptions that they may be cheap. Another potential reason may be that the rich view given lower tips as a way to maintain control or a sign of their ability to bargain. These reasons make it difficult to assign causality to the observed adverse relationship between income and gratuity.


We make progress by using a model of tipping behavior to empirically assess parameters that may explain why we observe such behavior. In the model, passengers have a belief about the social norm tip. This belief is influenced by a passenger's income level, which we determine empirically from the data. Passengers face a norm deviation cost if they decide not to adhere to the norm, and this cost varies by income. Last, passengers face a cognitive cost if they decide to actively compute their preferred tip versus choosing a suggested tip. Cognitive costs are also allowed to vary by income level.

Empirical estimates from our model suggests that passengers in less wealthy areas believe that one should tip taxi drivers at higher rates compared to their wealthier counterparts. However, the beliefs of the less wealthy are less binding because, they have norm deviation costs that are only half that of their wealthier counterparts. We also find that cognitive costs are lower among the less wealthy than wealthier passengers.

Our use of a structural model sheds light on the potential mechanisms that drive the observed negative relationship between income and gratuity. The differences we find in preferences, adherence to norms and computation costs are potentially relevant for social welfare as it relates to policies concerning income and giving.


Figures

Figure 2.1: Median Household Incomes in New York City



Source: 2010 American Community Survey.


Figure 2.2: Tipping Rates in New York City




Notes: Tips are computed from Green Taxis in NYC. The sample restriction are standard rate taxi trips, with no tolls, paid for via a credit card machine along with a positive tip. For Manhattan, Yellow taxi data is used as well.

Figure 2.3: Distribution of Beliefs About Tipping Norms & Preferred Non-Menu Tips

A) Beliefs About Tipping Norms by Income




B) Preferred Non-Menu Tips by Income



Notes: Panel A and B show the estimated distribution of passengers' beliefs about the social norm tip and their preferred non menu tips respectively. The data used are standard rate NYC Green and Yellow taxi trips from 08/2013-2014, with no tolls, paid for via a credit card machine along with a positive tip.

Figure 2.4: Model Fit



Notes: These figures illustrate how the parametric model fits the observed data across different income groups by showing the observed distribution of tips against the model predicted distribution of tips.

Tables

Table 2.1: Taxi Trip Characteristics: Share or Mean (Standard Deviation)

	Manhattan	Queens	Brooklyn	Bronx	All Trips
Tip Rate $\times 100$	18.69% (22.47)	20.82% (28.75)	20.91% (36.93)	24.28% (117.05)	19.97% (32.81)
Fare	\$12.77 (6.79)	\$12.07 (7.39)	\$11.87 (5.72)	\$11.29 (5.97)	\$12.29 (6.53)
Tip	\$2.3 (1.8)	\$2.39 (2.27)	\$2.4 (2.34)	\$2.36 (6.26)	\$2.36 (2.22)
Menu Tips	62%	68%	74%	65%	68%
Median Household Income	\$78,863	\$59,328	\$49,715	\$34,941	\$61,024
Observations	2,076,904	808,221	1,779,822	69,075	4,734,022

Notes: The data used are standard rate NYC Green and Yellow taxi trips from August 2013-2014, with no tolls, paid for via a credit card machine along with a positive tip. Median Household Income are estimates from the 2010 American Community Survey.

Table 2.2: OLS: Tipping in NYC Boroughs and by Income

	Dependent Variable: Tip Rate $\times$ 100
Fare	-0.222*** (0.002)
Bronx	6.647*** (0.368)
Brooklyn	2.722*** (0.093)
Queens	2.418*** (0.150)
Income/\$10K	-0.085*** (0.009)
Bronx*Income/\$10K	-0.515*** (0.104)
Brooklyn*Income/\$10K	-0.126*** (0.014)
Queens*Income/\$10K	-0.116*** (0.028)
Manhattan (Average Tip Rate)	18.69
Observations	4,734,022
$\bar{R}^2$	0.004

Notes: The data used are standard rate NYC Green and Yellow taxi trips from August 2013-2014, with no tolls, paid for via a credit card machine along with a positive tip. The excluded groups are trips in the borough of Manhattan and incomes over \$100K. We also include hourly temperature and precipitation and their quadratic transformations and fixed effects such as hour of day, day of week, month, and year.

Table 2.3: Probit: First Step Heckman Selection Model

	Dependent Variable: 1(Non-Menu Tip)
Fare	0.00666*** (0.00047)
Income $\leq$ \$35K	0.04797*** (0.00773)
Income (\$35K, \$70K]	0.03281*** (0.00659)
Income (\$70, \$100K]	0.04346*** (0.00722)
Fare*Income $\leq$ \$35K	-0.00642*** (0.00059)
Fare*Income (\$35K, \$70K]	-0.00400*** (0.00050)
Fare*Income (\$70, \$100K]	-0.00308*** (0.00054)
Bronx	-0.02388*** (0.00711)
Brooklyn	-0.22827*** (0.00205)
Queens	-0.14448*** (0.00259)
Number of Passengers	-0.01713*** (0.00073)
Constant	-0.60406*** (0.00899)
Observations	4,734,022
Log Likelihood	-1,374,375
Akaike Inf. Crit.	2,748,857

Notes: This table reports the first stage probit estimates of the Heckman selection correction model. In addition to Number of passengers in the taxi as an instrument, we include an indicator for round number tips, and fixed effects such as hour of day, day of week, month, and year. The data used are standard rate NYC Green and Yellow taxi trips from August 2013-2014, with no tolls, paid for via a credit card machine along with a positive tip. We report robust white standard errors in parenthesis. *p<0.1, **p<0.05, ***p<0.001.

Table 2.4: Structural Estimates

<i>Panel A: Second Step Heckman Selection Estimates</i>	
	Dependent Variable: Tip Rate
Fare	-0.00475*** (0.00025)
Income $\leq$ \$35K	0.09207*** (0.00416)
Income (\$35K, \$70K]	0.03042*** (0.0036)
Income (\$70, \$100K]	0.00270 (0.00398)
Fare*Income $\leq$ \$35K	-0.00420*** (0.0003)
Fare*Income (\$35K, \$70K]	-0.00147*** (0.00026)
Fare*Income (\$70, \$100K]	0.00027 (0.00029)
Bronx	0.12376*** (0.00379)
Brooklyn	0.04799*** (0.00108)
Queens	0.04230*** (0.00137)
Constant	0.20570*** (0.00336)
Norm Deviation Cost Parameter $\hat{\theta}$	
Income $\leq$ \$35K	55.86592*** (0.06274)
Income (\$35K, \$70K]	80.38585*** (0.13049)
Income (\$70, \$100K]	111.60714*** (0.76039)
Income $>$ \$100K	105.26316*** (0.03722)
R^2	0.00925
<i>Panel B: Simulated Method of Moments Estimates of Average Cognitive Cost ($1/\hat{\lambda}$)</i>	
Income $\leq$ \$35K	1.21262*** (0.03528)
Income (\$35K, \$70K]	1.11517*** (0.00309)
Income (\$70, \$100K]	1.60151*** (0.03071)
Income $>$ \$100K	1.49265*** (0.00868)

Notes: This table reports estimates of the parameters in the structural model. Panel A reports estimates of the beliefs about tipping norm T_i and the norm deviation cost parameter θ from the second step of the 2-step Heckman selection correction model. The excluded groups are trips in the borough of Manhattan and incomes over \$100K. We use the delta method to compute the standard errors for the estimates of θ . Panel B reports estimates from a two-step Simulated Method of Moments procedure of estimating the cognitive costs incurred by passengers when they opt to compute their preferred non-menu tip. In panel B, the standard errors are computed as the standard deviation of the distribution of parameter estimates computed from 1000 bootstrap samples. The data used are standard rate NYC Green and Yellow taxi trips from August 2013-2014, with no tolls, paid for via a credit card machine along with a positive tip. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.001$.

Chapter 3

The Effect of the Affordable Care Act on Agricultural Workers

3.1 Introduction

This study is the first to examine how individual Affordable Care Act (ACA) policies affected seasonal farmworkers' choice of health insurance coverage, medical services utilization, and jobs with employer-provided benefits.¹ Given that the objective of the Affordable Care Act (ACA) was to extend healthcare coverage to previously uncovered people, particularly those with low incomes and pre-existing medical conditions, seasonal agricultural workers are an important target population. These workers have low incomes, have relatively little health insurance coverage, and face many job-related health risks:² which results in high rates of health problems (Hansen et al., 2003; Guild et al., 2016; Mobed et al., 1992; Arcury and Quandt, 2007).

We concentrate on four ACA policies that were likely to affect farmworkers. First, the law required that states expand Medicaid to households with incomes less than 139% of the Federal Poverty Level (FPL). However, a 2012 Supreme Court ruling allowed states to opt-out of this expansion (ACA, P.L. 111-148, as amended).³ Furthermore states expanded Medicaid at different times. As a result, Medicaid eligibility varies by state over time and by the incomes and household sizes of workers.

Second, the law provided health insurance premium subsidies for households with incomes between 100% and 400% of the FPL who were not eligible for Medicaid.⁴ Third, an individual mandate required most citizens and legal residents to maintain health insurance coverage for at least nine months of each year or face a tax penalty unless they are exempt from

¹This chapter is coauthored with Susan Gabbard and Jeffrey Perloff.

²Farming ranks as one of the most hazardous industries by the Center for Disease Control and Prevention. Almost two-thirds (64%) of farmworkers and ranch workers report having pre-existing health conditions.

³Nine states expanded Medicaid coverage between 2010 and 2014. In 2014, 22 states expanded eligibility for Medicaid to the ACA recommended level. Nine more states expanded Medicaid eligibility between 2015 and 2016. See Simon et al. (2017).

⁴The Patient Protection and Affordable Care Act rules in 2020 cut the provision of premium subsidies in the health insurance market place (<https://s3.amazonaws.com/public-inspection.federalregister.gov/2019-08017.pdf>).

filing taxes.⁵ Fourth, the ACA prohibited insurance companies from setting insurance policy prices based on pre-existing health conditions. Insurers must set their prices based only on factors such as age, location, and smoking status.

In addition, the ACA required that insurers offer dependent child coverage until a child reaches the age of 26. However, that policy is unlikely to affect many farmworkers.

The ACA also mandated that firms with 50 or more full-time employees provide affordable health insurance packages for their workers or risk a penalty for failing to comply. These factors may have induced firms to redesign benefit plans made available to workers. We do not observe compensation packages directly in the data. However, we use self-reported information on benefits to analyze whether employers provided fewer benefits due to the ACA.

Several earlier studies found that the ACA raised coverage rates of previously uninsured low-income people and young adults; however, few examined the effect on farmworkers.⁶ Various studies have reported significant decreases in the uninsured rate after the implementation of the ACA. Some of these studies focused on evaluating the impact of Medicaid across states that expanded Medicaid coverage (e.g., Courtemanche et al. (2016); Kandilov and Kandilov (2019); Simon et al. (2017); Sommers et al. (2015); Frean et al. (2017); Kaestner et al. (2017), etc). Relatively few studies examined other ACA policies, such as the individual mandate and the premium subsidy (e.g., Frean et al. (2017); Fiedler (2018); Goldin et al. (2019)).

Various studies examined the effects of Medicaid expansion and most find an increase in coverage and use of medical services. Sommers et al. (2015) found a significant reduction in self-reports of being uninsured. Simon et al. (2017) reported an increase in the use of preventive care in states that expand Medicaid. Akosa Antwi et al. (2015b) found a decrease of 1.6 per 1000 young adults in the use of emergency room visits due to the Patient Protection and ACA dependent coverage expansions. For a more comprehensive literature review see Antonisse et al. (2018).

Similarly, several studies investigated the effects of the individual mandate. Fiedler (2018) argued that the individual mandate reduced the uninsurance rate from 24% to 39% among young adults with incomes above 400% of the FPL.

Frean et al. (2017) took a comprehensive approach, analyzing the effects of the ACA premium subsidies, the individual mandate, and Medicaid expansion on the uninsured rate. They found that Medicaid eligibility significantly increased coverage rates. However, their results show that both the premium subsidies and the individual mandates had little effect on raising coverage rates.

Our study evaluates all the policies that Frean et al. (2017) investigated, but for the population of farmworkers. Our study provides the first estimates of the effects of the ACA policies on the choice of coverage, choice of medical services utilized, and employer-provided benefits. Our analysis also allows us to evaluate the impact on individuals with pre-existing conditions and undocumented workers.

One of the few other studies to examine the ACA effect on farmworkers is Kandilov and

⁵The Tax Cut and Jobs Act (P.L. 115-97) repealed the individual mandate in 2019.

⁶For example, Akosa Antwi et al. (2015b); Gonzalez et al. (2016); Frean et al. (2017); Kaestner et al. (2017); Simon et al. (2017); Antonisse et al. (2018); Courtemanche et al. (2016).

Kandilov (2019). They used the same dataset we do to analyze differences in the coverage rates for states that expanded Medicaid versus those who have not. They found a significant increase in the use of government-provided insurance in states with Medicaid expansion. Our study differs from theirs by investigating the effects of more ACA policies, including premium subsidies, tax penalties, and protection for people with pre-existing medical conditions.

Unlike Kandilov and Kandilov (2019) (and all other previous studies), ours is the only one to examine the effects of the ACA on farmworkers’ use of medical services, where they go for treatment, and whether their job provides health and non-health benefits.⁷ Because this study is the first to examine how the ACA affected farmworkers’ use of various types of health providers, it is the first that addresses whether the ACA reduced the possibly inefficient use of emergency room services among farm farmworkers.⁸ Studies have document use of the emergency room by the insured and uninsured (Zhou et al., 2017). Others have focused on how medicaid or being newly insured affects the use of the emergency room (e.g., Finkelstein et al. (2016a,b); Taubman et al. (2014); Akosa Antwi et al. (2015a)).

The next section describes our data sources and variables. The following section discusses our specification. The next section explains our identification strategy. The next three sections discuss our analysis of the effects of the ACA policies on insurance coverage, health care use, and job benefits in turn. The final section presents a summary of our results and our conclusions.

3.2 Data

The National Agricultural Worker Survey (NAWS) is our primary data source. The U.S. Department of Labor started collecting the NAWS, as required by the Immigration Reform and Control Act of 1986. The NAWS randomly samples U.S. hired, seasonal crop workers over three cycles (seasons) each year. Using face-to-face interviews, it collects demographic, legal status, health insurance coverage and funding, medical services use, and job characteristics. The confidential version of the NAWS, which we use, provides detailed information on the county and state of residence for each surveyed farmworker.

Unlike most previous agricultural worker surveys that selected interviewees by where they lived, the NAWS selects people at their jobs. This approach is more likely to find undocumented workers, who are not citizens and green-card holders. Undocumented workers are over one-third (37%) of our sample.

We use data from 2010 through 2016. We do not examine 2007–2009 to avoid the Great Recession, and 2016 is the latest survey available.

The NAWS data set for 2010–2016 has 13,532 observations. We restrict our sample to farmworkers from 19 through 64 years of age for whom we have complete information on all relevant variables. This restriction reduces our sample by about 17%, leaving us with 8,557

⁷Unlike this study, Kandilov and Kandilov (2019) also analyzed the effects of Medicaid expansion on the labor supply of farmworkers. Other studies that examine the ACA effect or Medicaid on non-agricultural labor supply and found some evidence (Kaestner et al., 2017; Baicker et al., 2014; Antwi et al., 2013).

⁸Access to affordable healthcare can either increase or decrease the demand of emergency room services. The insight here is that, access to private and outpatient care may decrease the use of emergency room services. On the other hand, lowering out-of-pocket health costs may increase the use of emergency rooms.

farmworker observations.

We supplement our data set with information about Medicaid eligibility rules across states and national ACA subsidy and tax penalty rules, using information from state governments, the Kaiser Family Foundation, Healthinsurance.org, and the federal government.⁹

Generally, only documented workers qualify for Medicaid. States implemented Medicaid expansions at various times before and after the ACA took effect in 2014.¹⁰ The Kaiser Family Foundation lists the annual Medicaid eligibility household income thresholds for each state and the District of Columbia. We examined the Medicaid expansions that occurred between 2010 and 2016.

We combined these rules with household characteristics such as income and family composition to obtain an indicator variable for each farmworker’s eligibility for Medicaid. The first three columns of table 3.1 show the relationship between legal status, income, and Medicaid eligibility.

We construct an indicator variable for the ACA premium subsidy using information about the worker’s legal status from the NAWS and the ACA rules on the household composition and income threshold criteria for eligibility. Farmworkers who do not qualify for Medicaid are eligible for an ACA subsidy if their household incomes fall between 100% and 400% of the FPL.

From 2014, following the establishment of the ACA, through 2016, the last year of our sample, workers faced a tax penalty if they were covered by health insurance for fewer than nine months within the relevant tax year and met other criteria. Workers were exempt from the penalty if their household income was below the federal tax-filing threshold, their income was less than 139% of the FPL in Medicaid non-expansion states, they lived where no affordable coverage was available, or they were Native Americans or undocumented workers. We used the federal guidelines on household composition and income thresholds to calculate the prospective penalty for each farmworker based on their NAWS household characteristics. We label a farmworker as having a pre-existing health condition if he or she reports ever receiving a diagnosis of asthma, diabetes, high blood pressure, or heart disease.

3.3 Specification

We use a triple difference-in-differences strategy to identify the effects of the ACA policies. We estimate the effects of the ACA policies on five outcomes: (1) whether farmworkers have health insurance coverage, (2) their choice of the type of coverage, (3) whether they use of medical services, (4) the type of services they use, and (5) whether they work for an employer who provides various benefits.

We use the same explanatory variables in each of our analyses. In addition to the policy variables, we include demographic variables and state fixed effects (which we do not report in our tables to save space). The demographic variables are legal status, age, age squared,

⁹Medicaid eligibility by state across years is available at <https://www.kff.org/state-category/medicaid-chip/medicaidchip-eligibility-limits/>. For information on subsidy eligibility 2014–2016, see <https://www.healthinsurance.org/obamacare/will-you-receive-an-obamacare-premium-subsidy/#immigration>.

¹⁰See Simon et al. (2017) for a list of the Medicaid expansion and non-expansion states by year.

work experience, work experience squared, years of education, family income, household size, whether the worker reads and speaks English well, whether the individual works for a farm labor contractor (FLC), sex, white, Hispanic, season, and marital status. We include the FLC dummy because FLCs are less likely than growers to provide health insurance.

Our three main policy variables are a Medicaid-eligibility dummy, an ACA-subsidy-eligibility dummy, and a potential tax penalty variable. The tax penalty was due if a worker did not have health insurance coverage for at least nine months of the tax year.¹¹ The ACA-subsidy-eligibility dummy and the potential tax penalty depend on a farmworker's family income, which we assume is pre-determined.

These three policy measures differ across individuals. Our equations also include a post-ACA dummy, which captures the market impact of the ACA program and other time-related events since 2014.

A pre-existing-medical-condition dummy equals one for farmworkers who report having had a relevant health condition in the previous two years. We interact this variable with the policy dummies. We use the interaction to determine whether the ACA differentially affects the subset of farmworkers with pre-existing health conditions. By doing so, we measure the impact of the ACA requirement that prevents health insurance providers from price discriminating based on pre-existing health conditions.

Treatment and Control Groups

To identify the policy effects, we compare people who are eligible for Medicaid or an ACA subsidy to those who are not. After the enactment of the ACA, coverage of documented workers differed by state. Undocumented workers in our samples were ineligible for the ACA programs. We dropped Californian undocumented workers because California provided the same healthcare coverage for undocumented immigrants that it provided to citizens (Hayes et al., 2015).¹²

We identify the ACA policy impacts by comparing the behavior of documented workers who were affected by these policies after the law went into effect in 2014 to documented workers in the earlier period and undocumented workers throughout the entire period.

Although undocumented workers differ from documented workers in several ways, the shares of both groups who had health insurance and who used medical services were virtually constant before the ACA went into effect in 2014. Figure 1 shows that before the ACA, 2010–2013, the rate of insurance coverage was stable at around 41% for documented workers and 14% for undocumented workers. After the ACA went into effect, documented

¹¹The ACA tax penalty was \$0 for all workers before the ACA went into effect in 2014. Thereafter, the penalty was \$0 for all undocumented workers, who are not covered by the ACA, and for documented workers who were not required to file income taxes. For all other documented workers, the penalty in 2014 was the greater of \$95 per adult in the family plus \$47.50 per child (up to a \$285 maximum) or one percent of household income. In 2015, the penalty was the greater of \$325 per uninsured person plus \$162.50 per child (up to a \$975 maximum) or two percent of household income. In 2016, the penalty was the greater of \$695 per uninsured person plus \$347.50 per child (up to a \$975 maximum) or two percent of household income.

¹²During part of the ACA period, six states and the District of Columbia used their funds to provide Medicaid coverage to income-eligible children regardless of immigration status. However, we do not have children in our sample.

workers' insurance coverage rate rose steadily, reaching about 70% in 2016, while the rate for undocumented workers was unchanged.

The only information we have about the use of medical services is whether a worker reported using any medical service within the two years before the interview date. Figure 2 shows that the usage rate of medical services remained virtually constant throughout the entire period at about 70% for documented workers and 53% for undocumented workers.

Table 3.2 contains summary statistics for documented and undocumented farmworkers before and the period after the enactment of the ACA. Most demographics, employment, language, and other farm work-related variables were little changed after the ACA went into effect.

Our Medicaid eligibility indicator for documented workers nearly doubles from 12% before to 23% after. We calculated that two-thirds of documented farmers were eligible for the premium subsidy post-ACA. The average potential tax penalty for documented workers who lack coverage is \$593.11. However, the average penalty trended upward over time: \$288.60 in 2014, \$704.61 in 2015, and \$879.46 in 2016.

We record a farmworker as having a pre-existing medical condition if that worker reported having suffered from asthma, diabetes, high blood pressure, or heart disease. The share of workers with pre-existing conditions was essentially constant before and after the ACA at 22%–23% for documented workers and 10%–11% for undocumented workers.

One reason why documented workers are more likely to have pre-existing conditions than undocumented workers is that undocumented workers were eight years younger than documented workers pre-ACA and five years younger in the later period. The other main differences between them are that documented workers were wealthier, more likely to be white, read and speak English well, have more work experience, and are less likely to be Hispanic.

3.4 ACA Effects on Health Insurance Coverage

The Affordable Care Act's primary objective was to provide affordable healthcare for people with low incomes or pre-existing health conditions. We use a multinomial logit model to examine how the ACA policies affected the probability that a farmworker lacked coverage or had one of four types of health insurance coverage: employer, government, private, or other. "Employer" insurance is provided by the worker's employer or by the spouse's employer. "Government" insurance is provided by a government agency, directly or indirectly (as with ACA insurance). The worker or a family member paid for "Private" insurance. An unnamed source funded "Other" insurance.

Table 3.3 and 3.4 report the multinomial logit estimates and the marginal effects respectively of various demographic and policy variables evaluated at the mean of the sample. For a continuous variable, we calculate the marginal effect for a small change. However, for a discrete variable, we report the difference between the probability if the dummy equals one minus the probability if the dummy is zero.

Demographic Effects

The demographic variables have plausible effects. Being proficient in English raised a worker's probability of having insurance (lowered the probability of not having insurance) by 5.3%, increasing the probability of having government insurance by 3.8%.

A larger household increases the probability of having government insurance by 2.2%, but lowers the probability of having employer insurance by 1.3%. Employees of farm labor contractors (FLCs) were 5.5% less likely to have insurance than people who were directly hired by growers. Females were 5.7% more likely to have insurance, and 4% more likely to have government insurance than men. After 35 years old, being older raises the likelihood of being insured.

Whites were 4.4% more likely to have insurance, while the probability for Hispanics was 12.5% lower. Hispanics' insurance coverage was 7.1% lower for employer insurance, 3% less for private insurance, and 4% lower for other insurance.

Based on the "Income/\$1,000" variable, the marginal effect is a 0.5% increase in insurance coverage.¹³ Thus, a ten thousand dollars increase above the average income increases the probability that the worker has health insurance by about 7%.

Policy Effects

Some of the ACA policies substantially increased the share of documented workers with government insurance, while reducing the share that had employer-provided insurance or were uninsured. The ACA and state laws expanded Medicaid eligibility. For documented workers with pre-existing medical conditions, Medicaid eligibility increased the probability of having government insurance by 11.4%, increased other insurance by 4.4%, and lowered the probability of employer insurance by 5.6%. In net, the likelihood they had medical coverage rose by 9.2%.

For documented workers without pre-existing conditions, employer coverage fell by about the same amount (5.3%) as for those with conditions. Government insurance rose less (7.6%), and other insurance rose by 3.4%. The net effect was a 3.8% increase in medical coverage. This 5.5% ($= 9.2\% - 3.8\%$) difference in medical coverage between those with pre-existing conditions and those without conditions may reflect adverse selection.¹⁴

Being eligible for an ACA subsidy statistically significantly decreased the probability of having employer insurance by 7.5% for documented workers with pre-existing medical conditions, and by 5.3% for those without such conditions. It raises the probability of having government insurance by 4.3% for those with pre-existing conditions (but we can reject the null hypothesis of no effect at only the 10% level), and by 4.1% for those without the conditions. As a consequence of these offsetting effects, the subsidy did not have a statistically significant effect on the overall probability of having medical coverage for either group.

The marginal effect of a (\$100) increase in the potential ACA tax penalty was a 1.3% rise in overall medical coverage. However, we can reject the null-hypothesis of no effect at

¹³However, income also affects the probability of being eligible for Medicaid and a subsidy, which we do not include in this calculation.

¹⁴Panhans (2019) finds evidence of adverse selection due to the ACA in Colorado.

only the 10% level. If we rely on this point estimate, the effect grew from 2014–2016. The coverage rose by about 3.8% to 9.2% to 11.4% as the potential tax penalty rose from \$288.60 to \$704.61 to \$879.46 from 2014 through 2016.

These three policy variables do not capture the full effect of the ACA. For example, the ACA provided coverage of offspring under the age of 26 (though presumably such extended coverage was relatively unlikely to be used by farmworkers). It affected market prices. Also, the debate and advertising about how to sign up for ACA coverage made health insurance more salient to the public. To capture these “residual” policy effects, we included in our equation a post-ACA dummy for documented workers, which is one for the years when the ACA was in effect, 2014–2016.¹⁵ All else the same, the probability of having health insurance rose by 11% for documented workers with and without pre-existing medical conditions. Surprisingly, the biggest effect is an increase in private insurance rather than government insurance.¹⁶

The ACA policies should not affect undocumented workers, unless it indirectly affects them through changes in price or public information. As a robustness check, we re-estimated the model with the post-ACA dummy interacted with the undocumented worker dummy (and with the undocumented worker dummy interacted with the pre-existing medical conditions dummy) and did not find statistically significant effects. This (lack of a) result is consistent with Figure 1.

Being eligible for an ACA premium subsidy did not have a statistically significant effect at the 0.05 level for documented workers with or without pre-existing medical conditions. That may be because the subsidy caused people to shift coverage from private to government insurance.

Because we use an imperfect estimate of income, we are calculating Medicaid eligibility and subsidy eligibility with noise. Such noise may bias our estimates downward.

To summarize, the ACA policies affected only documented workers. The effects were larger for those with pre-existing medical conditions than for those without these conditions. Several of the ACA policies caused the use of government insurance to increase and employer-insurance to decrease. Consequently, although the overall share of workers with medical insurance rose (the uninsured rate fell), it increased much less than did government coverage. The net effect of subsidy eligibility on the total coverage rate was a wash. Moreover, the potential ACA tax penalty for not maintaining coverage was negligible. However, the effective policies can explain the substantial increase in coverage of documented workers in the post-ACA period.

3.5 ACA Effects on the Use of Medical Services

Some of the most important, unanswered questions during the ACA debate concerned how it would affect the consumption of medical services. Would people use substantially more

¹⁵We checked whether these factors affected undocumented workers, but the coefficient on the post-ACA dummy for undocumented workers was virtually zero.

¹⁶We estimated an alternative equation with year dummies rather than with a single post-ACA dummy. However, that greater flexibility was irrelevant because the coefficients on the year dummies were nearly identical.

medical services? Would expanding medical insurance coverage reduce inefficient uses of emergency room services?

We use the NAWS data to address these questions for farmworkers. Unfortunately, the NAWS does not report a measure of how intensively farmworkers used healthcare services. However, it does record whether a farmworker used any medical services in the 24 months before the interview.

The 24-month timeframe poses potential challenges, especially for estimating the impact of the ACA specific mandates. For example, for farmworkers surveyed in 2014 and 2015, we cannot distinguish whether their most recent healthcare service visit occurred after 2013—the ACA period.¹⁷

We estimate a multinomial logit to examine whether a farmworker used a medical service, and if so, from what type of provider. The NAWS divides workers' healthcare usage into six categories. "None" refers to farmworkers who report they have not used any medical services within the past 24 months. If they reported using a service, the NAWS asked them where they received their most recent treatment: an "Emergency Room" (ER), a "Hospital," a "Private Doctor or Clinic," a "Community Health Center or Migrant Health Center," or another type of provider. This "Other" category includes chiropractors, healers, dentists, and other unspecified providers. Table 3.5 and 3.6 report the multinomial logit estimates and the marginal effects respectively.

The effects of most of the demographic variables seem plausible. For example, the larger the household size, the more likely they are to get treatment. The probability of having received a treatment was higher for workers who were proficient in English, were female, were white, had more years of education, or did not work for an FLC. After 40, being older increased the probability of using medical services. Most of these same factors raised the probability that workers visited a private doctor or clinic for medical services. None of the demographic characteristics had a substantial effect on the likelihood of using the ER. Females, workers with large households, and very experienced workers were more likely to use hospitals.

None of the policy variables had a statistically significant effect on ER use at the 5% level. For documented workers with pre-existing medical conditions, Medicaid eligibility increased the use of hospitals by 4.8%, private doctors or clinics by 13.9% and community, and migrant health centers by 18.7%. It reduced visits to other medical providers by 17.4%. For these workers, the main effect of being eligible for a subsidy was an 11.1% greater probability of going to a community or migrant health centers. For the documented workers without pre-existing conditions, none of the policy variables had a substantial effect on where workers went for medical care.

Documented workers with a pre-existing medical condition were nearly 21% more likely to have received treatment if they were Medicaid eligible. In contrast, for documented without pre-existing medical conditions, the marginal effect point estimate (which is statistically significantly different from zero at only the 10% level) is that Medicaid eligibility lowered the probability of medical use by 4.4%.

¹⁷We examined the importance of this problem by including post-ACA year specific effects of the various policies interacted with dummies for year 2014, 2015, and 2016. The year effects are nearly identical, which suggests that this problem may not be major.

For documented workers with pre-existing medical conditions, subsidy eligibility and the post-ACA dummy have a statistically significant effect on being treated that is different from zero at only the 10% level. However, the marginal effects are large. Subsidy eligibility raised the probability of using medical care by 11.4%, and the post-ACA dummy was associated with a 7.5% increase.

The subsidy eligibility variable for documented workers without pre-existing conditions is statistically significant at the 1% level. Surprisingly, these subsidy-eligible workers are 7.5% less likely to receive treatment, but where the service provider was relatively unchanged. The post-ACA dummy does not have a statistically significant effect.

The potential ACA tax penalty does not affect the probability of getting treatment. However, it is associated with a small but statistically significant reduction in the use of other services.

In summary, the ACA policies did not significantly affect ER use, but some of them increased the probability of going to a hospital, a private doctor or clinic, or community or migrant health center. These policies substantially increased the probability that documented farmworkers with pre-existing medical conditions sought treatment, but they did not have that effect for documented workers without pre-existing conditions.

3.6 Employer-Provided Benefits

A simple economic model would postulate that in a competitive market, firms that offer more benefits than rival firms are likely to pay a lower wage. Thus, a worker who took a relatively low-paying job to obtain health benefits before the ACA might choose a job that paid more but, lacked health insurance after the ACA.

However, the ACA could affect other benefits as well. Many agricultural labor markets are dual labor markets, where some firms pay high wages and provide good benefits, while others pay low wages and provide few if any benefits. Moretti and Perloff (2002) describe how such an outcome could arise due to firms' use of efficiency wages to increase productivity. Also, such an outcome could result from some employers taking advantage of relatively uninformed workers or undocumented workers, particularly those employed by FLCs.¹⁸ Thus, because some firms provide many benefits and others relatively few, the ACA might affect all benefits.

We examine whether workers who were eligible for Medicaid and ACA health insurance were less likely to work for an employer who provided various benefits than in the pre-ACA period.

We consider four benefits: health, incentive, season-end, and holiday benefits. The NAWS asks, "If you are injured at work or get sick as a result of your work, does your employer provide health insurance or pay for your health care?" Our incentive measure reflects whether the individual received an incentive based on the worker's effort or the employer's profit.

Whether workers receive benefits reflects market outcomes from both workers and employers' decisions. Table 3.7 reports linear probability estimates for each benefit individually.

¹⁸Similarly, undocumented workers may avoid using certain types of benefits (particularly health care) for fear that using them could leave them vulnerable to being deported. Moretti and Perloff (2000) found that undocumented agricultural workers were unlikely to make use of publicly funded welfare and health service programs, which is consistent with our results here.

The results provide little support for the hypothesis that the ACA induced changes in benefits. The ACA policy variables do not have a statistically significant effect at the 5% level for any benefit, including whether an employer provides medical coverage. Though not shown in the table, these policy variables also did not affect a composite incentive variable, which equals one if the worker received any benefit.

At first glance, we have a mystery when comparing these results to those in table 3.4. In Table 3.4, the ACA policies had moderately large and statistically significant effects on whether a worker has employer-provided medical coverage. In contrast, table 3.7 does not show a statistically significant effect on an employer health benefit. We believe that this difference is mainly due to wording differences in how the NAWS asked its questions. Presumably, the question as to whether the worker has employer-provided health insurance is unambiguous. In contrast, the benefit question as whether the employer would provide health insurance conditional on whether the worker was injured or became ill as a result of the job. Thus, we think the table 3.4 results are more reliable than those in table 3.7.

3.7 Conclusions

The purpose of the Affordable Care Act was to expand healthcare coverage, particularly to young and low-income people. Most previous studies of ACA effects have ignored farmworkers. However, the law targeted helping people like farmworkers who are low-paid, traditionally had low rates of coverage, and suffered from more health challenges than most other workers. Using the National Agricultural Workers Survey of seasonal agricultural workers, we examine how the ACA affected farmworkers' health insurance coverage, their use of medical care, and the types of benefits provided by their employers.

We compare the effects of the ACA on documented workers, who are directly affected by this law, to undocumented workers, who are not. Moreover, we show how the law's effects depend on whether workers had pre-existing medical conditions. We use a triple difference-in-difference framework to evaluate the effects of the ACA. Using variations in state laws, we can examine how differences in Medicaid eligibility (due to the ACA's expansion), eligibility for ACA subsidies, and the potential tax penalty for not maintaining coverage affect outcomes. We also use a catch-all dummy to capture other possible effects of the law.

Expanded Medicaid eligibility and the residual ACA policy variable substantially increased the healthcare coverage of documented workers, particularly those with pre-existing medical conditions. Although the share of documented workers with medical insurance rose significantly, it increased by much less than the increase in the share with government coverage, because many workers switch to government insurance from employer-provided insurance. The potential ACA tax penalty increased coverage, but was only statistically significant at the 10% level. The ACA subsidy may have encouraged workers to switch from employer-provided insurance to government insurance, so it had a negligible effect on overall medical coverage.

Proponents of the ACA hoped that expanded medical coverage would improve efficiency by reducing the reliance on emergency rooms for basic health care. However, these policies did not reduce ER use by this group. On the other hand, covered documented workers with pre-existing medical conditions were more likely to use hospitals, private doctors or clinics,


and community or migrant health centers.

The ACA policies substantially increased the share of documented workers with pre-existing medical conditions who sought medical care. However, they did not have this effect on those without such conditions.

Finally, we did not find that the share of workers receiving a variety of benefits changed significantly after the ACA started.


Figures

Figure 3.1: Farm Worker's Health Insurance Coverage



Source: National Agricultural Worker Survey, 2010–2016.

Figure 3.2: Farm Worker's Medical Services Use



Source: National Agricultural Worker Survey, 2010–2016.

Tables

Table 3.1: Eligibility and Tax Penalty

Legal Status	<u>Household Income</u> Federal Poverty Level (%)	Medicaid	ACA Premium Subsidy	Tax Penalty
Citizen & Green Card (> 5 years of residence)	< 139 %	1	1	No
	≤ 400%	1	1	Yes
	> 400%	No	No	Yes
Green Card (< 5 years of residence)	< 100%	No	Yes	No
	100%–400%	No	Yes	Yes
	> 400%	No	No	Yes
Undocumented ²	All	No	No	No

¹ Depends on the state's Medicaid eligibility laws.

² California provides Medicaid (Medi-Cal) using state funds for undocumented workers who would otherwise not qualify.

Table 3.2: Summary Statistics by Farmworker's Legal Status and Period. Share or Mean (Standard Deviation)

	<u>Pre-ACA:2010-2013</u>		<u>Post-ACA: 2014 - 2016</u>	
	Documented	Undocumented	Documented	Undocumented
<i>Insurance and Medical Characteristics</i>				
Health Insurance Coverage	41%	9%	63%	12%
Use Medical Services	68%	56%	68%	56%
Medicaid Eligible	12%	0%	23%	0%
Subsidy Eligible	0%	0%	67%	0%
Pre-existing Medical Condition	22%	10%	23%	11%
Potential ACA Tax Penalty (\$)	0	0	593.11 (300.93)	0
<i>Demographic and Work Characteristics</i>				
Married	66%	0.63	64%	0.63
Female	18%	0.23	20%	0.27
White	47%	0.29	38%	0.13
Hispanic	69%	1%	70%	99%
Citizen	53%	–	55%	–
Green Card	47%	–	45%	–
Read and Speak English Well	44%	3%	45%	3%
Farm Labor Contractor	8%	2%	8%	3%
Received Bonus Pay	46%	38%	47%	42%
Age	42.58 (12.55)	34.59 (9.38)	42.92 (13.07)	37.12 (9.68)
Family Income/1,000 (\$)	26.33 (13.17)	20.92 (9.69)	28.65 (13.77)	25.03 (11.18)
Family Income/FPL (\$)	1.77 (1)	1.32 (0.68)	1.95 (1.03)	1.51 (0.74)
Household Size	2.47 (1.79)	2.56 (1.82)	2.36 (1.71)	2.74 (1.84)
Education (years)	8.71 (4.17)	6.93 (3.28)	9.02 (4.13)	6.80 (3.51)
Work Experience (years)	23.05 (13.93)	11.94 (7.35)	23.07 (14.90)	13.94 (8.17)
Number of Observations	2,407	1,526	2,951	1,673

Source: National Agricultural Worker Survey, 2010–2016.

Table 3.3: Multinomial Logit Estimates, Health Insurance

	Employer		Government		Private		Other	
	Coef	SE	Coef	SE	Coef	SE	Coef	SE
<i>Documented Workers with</i>								
<i>Pre-Existing Medical Conditions</i>								
Medicaid Eligible	-0.171	0.234	1.476***	0.293	0.056	0.399	0.979***	0.353
Subsidy Eligible	-0.652**	0.326	0.405	0.378	-0.105	0.410	0.166	0.385
Post ACA	0.662**	0.337	0.820**	0.386	1.740***	0.432	-0.067	0.285
<i>Documented Workers without</i>								
<i>Pre-Existing Medical Conditions</i>								
Medicaid Eligible	-0.318	0.198	0.906***	0.179	-0.328	0.262	0.643**	0.285
Subsidy Eligible	-0.520***	0.180	0.351**	0.144	-0.485*	0.294	-0.084	0.241
Post ACA	0.644	0.456	0.840**	0.345	1.999***	0.406	-0.080	0.417
ACA Tax Penalty /\$100	0.073*	0.041	0.080	0.052	0.102**	0.044	0.101**	0.047
Income /\$1,000	0.059***	0.004	-0.008	0.006	0.027***	0.009	0.036***	0.005
FLC	-0.378**	0.188	0.074	0.082	-0.909***	0.333	-0.697***	0.193
Age	-0.025	0.019	-0.077***	0.028	-0.225***	0.042	-0.249***	0.045
Age Squared/100	0.050**	0.022	0.122***	0.033	0.260***	0.054	0.308***	0.057
Experience	0.032***	0.008	0.015	0.011	0.014	0.015	0.054***	0.016
Experience Squared/100	-0.048**	0.020	-0.028	0.019	0.011	0.025	-0.066**	0.033
Education	0.080***	0.014	0.044***	0.014	0.100***	0.027	0.070***	0.013
Married	0.215*	0.117	-0.108	0.102	-0.252	0.167	0.099	0.137
Female	0.137	0.159	0.588***	0.112	0.432***	0.142	0.395***	0.109
White	0.314**	0.122	-0.024	0.111	0.511***	0.115	0.586***	0.104
Hispanic	-0.996***	0.205	-0.214	0.154	-1.212***	0.239	-1.119***	0.237
English Proficient	0.276	0.171	0.556***	0.131	0.239	0.182	0.123	0.092
Household Size	-0.102***	0.036	0.247***	0.036	-0.038	0.072	-0.061	0.064
Pre-ACA shares,								
Documented Workers	0.186		0.100		0.027		0.098	

Notes: National Agricultural Workers Survey, 2010–2016, 8,557 observations. State fixed effects are not reported to save space. The base outcome is “Uninsured”. Standard errors are clustered at the level of the state. *p < 0.10, **p < 0.05, ***p < 0.01.

Table 3.4: Marginal Effects, Health Insurance

	Uninsured		Employer		Government		Private		Other	
	ME	ASE	ME	ASE	ME	ASE	ME	ASE	ME	ASE
<i>Documented Workers with</i>										
<i>Pre-Existing Medical Conditions</i>										
Medicaid Eligible	-0.092***	0.032	-0.056**	0.023	0.114***	0.022	-0.010	0.015	0.044***	0.016
Subsidy Eligible	0.017	0.050	-0.075***	0.023	0.043*	0.025	-0.001	0.012	0.015	0.016
Post ACA	-0.110***	0.040	0.037	0.031	0.045	0.027	0.056***	0.019	-0.028**	0.014
<i>Documented Workers without</i>										
<i>Pre-Existing Medical Conditions</i>										
Medicaid Eligible	-0.038*	0.022	-0.053***	0.018	0.076***	0.013	-0.018*	0.011	0.034**	0.015
Subsidy Eligible	0.025	0.023	-0.053***	0.016	0.041***	0.010	-0.015	0.010	0.002	0.011
Post ACA	-0.114**	0.057	0.032	0.035	0.045**	0.019	0.067***	0.013	-0.030*	0.016
ACA Tax Penalty /\$100	-0.013*	0.007	0.004	0.003	0.004	0.003	0.002*	0.001	0.003*	0.002
Income /\$1,000	-0.005***	0.001	0.005***	0.000	-0.002***	0.000	0.000	0.000	0.001***	0.000
FLC	0.055***	0.020	-0.021	0.016	0.024**	0.010	-0.029***	0.011	-0.029***	0.008
Age	0.016***	0.002	0.004*	0.002	-0.002	0.002	-0.007***	0.001	-0.012***	0.002
Age Squared/100	-0.023***	0.003	-0.004	0.003	0.005*	0.003	0.007***	0.002	0.014***	0.003
Experience	-0.005***	0.001	0.002***	0.001	0.000	0.001	-0.000	0.001	0.002***	0.001
Experience Squared/100	0.006***	0.002	-0.004*	0.002	-0.001	0.001	0.001	0.001	-0.003*	0.002
Education	-0.011***	0.002	0.005***	0.001	0.001	0.001	0.002***	0.001	0.002***	0.001
Married	-0.006	0.012	0.025**	0.012	-0.011	0.008	-0.012*	0.006	0.005	0.007
Female	-0.057***	0.017	-0.006	0.014	0.040***	0.008	0.010*	0.005	0.013**	0.005
White	-0.044***	0.013	0.019*	0.011	-0.015**	0.008	0.014***	0.004	0.026***	0.006
Hispanic	0.125***	0.018	-0.071***	0.022	0.016	0.011	-0.030***	0.008	-0.040***	0.013
English Proficient	-0.053***	0.016	0.015	0.016	0.038***	0.009	0.002	0.008	-0.003	0.005
Household Size	-0.004	0.005	-0.013***	0.004	0.022***	0.003	-0.002	0.002	-0.004	0.003
Pre-ACA shares,										
Documented Workers	0.59		0.186		0.100		0.027		0.098	

Notes: National Agricultural Workers Survey, 2010–2016, 8,557 observations. State fixed effects are not reported to save space. Standard errors are clustered at the level of the state. *p < 0.10, **p < 0.05, ***p < 0.01.

Table 3.5: Multinomial Logit Estimates, Use of Medical Services

	ER		Hospital		Private		Comm/Migrant		Other	
	Coef	SE	Coef	SE	Coef	SE	Coef	SE	Coef	SE
<i>Documented Workers with</i>										
<i>Pre-Existing Medical Conditions</i>										
Medicaid Eligible	1.599	1.269	1.411***	0.422	1.282***	0.248	1.631***	0.226	-1.220***	0.450
Subsidy Eligible	1.545	1.136	1.149*	0.614	0.381	0.361	0.929***	0.266	-0.372	0.578
Post ACA	-0.178	1.308	-0.618	0.509	0.661**	0.259	0.364*	0.193	0.488	0.383
<i>Documented Workers without</i>										
<i>Pre-Existing Medical Conditions</i>										
Medicaid Eligible	-0.468	0.564	-0.134	0.186	-0.131	0.138	-0.278*	0.168	-0.292*	0.163
Subsidy Eligible	-0.342	0.430	-0.280	0.217	-0.315**	0.141	-0.383**	0.184	-0.508**	0.208
Post ACA	-0.230	0.544	-0.241	0.172	-0.170	0.240	-0.058	0.156	0.342	0.212
ACA Tax Penalty /\$100	0.108*	0.060	0.031	0.020	0.039	0.030	0.003	0.032	-0.041	0.032
Income /\$1,000	0.006	0.011	0.008*	0.004	0.032***	0.003	0.008*	0.005	0.023***	0.005
FLC	-0.823	0.523	-0.271	0.209	-0.240**	0.096	-0.379***	0.082	-0.205	0.202
Age	0.019	0.085	-0.094***	0.030	-0.094***	0.022	-0.061***	0.020	-0.058**	0.028
Age Squared/100	-0.092	0.118	0.104***	0.033	0.126***	0.027	0.082***	0.025	0.063*	0.034
Experience	0.008	0.030	-0.001	0.015	0.024**	0.012	0.035***	0.013	0.027**	0.014
Experience Squared/100	0.077	0.062	0.041	0.030	-0.015	0.022	-0.055*	0.029	-0.044	0.027
Education	0.090	0.057	0.040**	0.019	0.064***	0.011	0.047***	0.013	0.043***	0.015
Married	0.029	0.352	-0.219	0.147	0.059	0.074	-0.198**	0.082	-0.098	0.076
Female	0.455*	0.252	0.968***	0.116	1.037***	0.077	1.239***	0.067	0.559***	0.092
White	0.196	0.310	-0.081	0.112	0.497***	0.072	-0.349***	0.078	0.411***	0.078
Hispanic	-0.755**	0.360	-0.295**	0.134	-0.529***	0.121	0.357***	0.127	-0.341	0.211
English Proficient	0.582	0.508	0.282**	0.141	0.466***	0.105	0.053	0.136	0.573***	0.144
Household Size	0.071	0.098	0.148***	0.023	0.063*	0.032	0.143***	0.026	0.147***	0.032
Pre-ACA shares,										
Documented Workers	0.010		0.068		0.318		0.162		0.118	

Notes: National Agricultural Workers Survey, 2010–2016, 8,557 observations. State fixed effects are not reported to save space. The base outcome is “None”. Standard errors are clustered at the level of the state. *p < 0.10, **p < 0.05, ***p < 0.01.

Table 3.6: Marginal Effects, Use of Medical Services

	None			ER			Hospital			Private			Comm/Migrant			Health Center			Other		
	ME	ASE	ME	ASE	ME	ASE	ME	ASE	ME	ME	ASE	ME	ME	ASE	ME	ME	ASE	ME	ME	ASE	ME
<i>Documented Workers with</i>																					
<i>Pre-Existing Medical Conditions</i>																					
Medicaid Eligible	-0.208***	0.037	0.009	0.012	0.048**	0.023	0.139***	0.046	0.187***	0.025	-0.174***	0.039									
Subsidy Eligible	-0.114*	0.067	0.011	0.011	0.051*	0.030	0.007	0.039	0.111***	0.024	-0.066*	0.039									
Post ACA	-0.075*	0.041	-0.004	0.012	-0.055*	0.030	0.088**	0.037	0.025	0.023	0.021	0.028									
<i>Documented Workers without</i>																					
<i>Pre-Existing Medical Conditions</i>																					
Medicaid Eligible	0.044*	0.024	-0.003	0.005	0.000	0.012	0.002	0.019	-0.028	0.023	-0.015	0.012									
Subsidy Eligible	0.075***	0.027	-0.001	0.004	-0.003	0.012	-0.015	0.020	-0.030	0.024	-0.026*	0.016									
Post ACA	0.011	0.025	-0.002	0.005	-0.013	0.012	-0.029	0.037	-0.003	0.029	0.036**	0.015									
ACA Tax Penalty /\$100	-0.003	0.005	0.001*	0.000	0.001	0.002	0.006*	0.003	-0.001	0.003	-0.005***	0.002									
Income /\$1,000	-0.004***	0.001	-0.000	0.000	-0.000	0.000	0.004***	0.000	-0.001	0.001	0.001***	0.000									
FLC	0.060***	0.019	-0.006	0.005	-0.005	0.011	-0.010	0.011	-0.038***	0.012	-0.002	0.015									
Age	0.015***	0.003	0.001	0.001	-0.003	0.002	-0.010***	0.003	-0.003	0.003	-0.001	0.002									
Age Squared/100	-0.019***	0.004	-0.001	0.001	0.003	0.002	0.014***	0.004	0.004	0.004	0.000	0.003									
Experience	-0.005***	0.002	-0.000	0.000	-0.001	0.001	0.002	0.002	0.004**	0.002	0.001	0.001									
Experience Squared/100	0.005	0.004	0.001	0.001	0.004**	0.002	0.000	0.003	-0.008**	0.004	-0.003	0.002									
Education	-0.011***	0.002	0.000	0.001	0.000	0.001	0.006***	0.001	0.003*	0.002	0.001	0.001									
Married	0.018	0.012	0.001	0.003	-0.011	0.008	0.025**	0.011	-0.028**	0.011	-0.005	0.005									
Female	-0.207***	0.009	-0.002	0.002	0.021***	0.007	0.081***	0.010	0.116***	0.009	-0.008	0.006									
White	-0.023**	0.011	0.001	0.003	-0.010	0.007	0.087***	0.010	-0.082***	0.010	0.028***	0.005									
Hispanic	0.033*	0.018	-0.005*	0.003	-0.012	0.009	-0.087***	0.018	0.090***	0.019	-0.018	0.015									
English Proficient	-0.065***	0.019	0.003	0.005	0.005	0.008	0.052***	0.014	-0.028	0.018	0.033***	0.009									
Household Size	-0.023***	0.003	-0.000	0.001	0.005***	0.002	-0.003	0.005	0.014***	0.004	0.007***	0.002									
Pre-ACA shares,																					
Documented Workers	0.32		0.010		0.068		0.318		0.162		0.118										

Notes: National Agricultural Workers Survey, 2010–2016, 8,557 observations. State fixed effects are not reported to save space. The base outcome is “None”. Standard errors are clustered at the level of the state. *p < 0.10, **p < 0.05, ***p < 0.01.

Table 3.7: Linear Probability Models of Employers' Provision of Fringe Benefits

	Health		Incentive		Season-End		Holiday	
	Coef	SE	Coef	SE	Coef	SE	Coef	SE
<i>Documented Workers with</i>								
<i>Pre-Existing Medical Conditions</i>								
Medicaid Eligible	0.041	0.025	-0.030	0.027	-0.037	0.037	0.028	0.034
Subsidy Eligible	0.020	0.039	-0.032	0.031	0.005	0.026	0.060	0.054
Post ACA	0.035	0.052	0.036	0.035	-0.029	0.028	-0.044	0.045
<i>Documented Workers without</i>								
<i>Pre-Existing Medical Conditions</i>								
Medicaid Eligible	-0.018	0.016	0.006	0.013	-0.026	0.020	-0.042	0.025
Subsidy Eligible	-0.036	0.024	-0.021	0.015	-0.009	0.016	-0.024	0.026
Post ACA	0.081*	0.042	0.012	0.017	-0.029	0.031	-0.003	0.035
ACA Tax Penalty /\$100	-0.002	0.003	0.000	0.002	0.005*	0.003	-0.000	0.003
Income /\$1,000	0.004***	0.001	0.002***	0.000	0.002***	0.001	0.005***	0.000
FLC	-0.010	0.025	-0.039***	0.007	-0.070***	0.009	-0.162***	0.025
Age	0.006*	0.003	0.001	0.002	-0.004	0.002	-0.003	0.003
Age Squared/100	-0.007*	0.004	-0.002	0.002	0.005	0.003	0.003	0.004
Experience	0.005**	0.002	0.001	0.001	0.005***	0.001	0.007***	0.002
Experience Squared/100	-0.006	0.004	0.000	0.002	-0.008***	0.002	-0.012***	0.004
Education	0.006***	0.001	0.000	0.001	0.002***	0.001	0.001	0.002
Married	-0.007	0.019	-0.003	0.008	0.009	0.012	-0.010	0.013
Female	-0.037**	0.016	-0.016**	0.007	-0.063***	0.013	-0.012	0.013
White	0.023	0.017	0.014*	0.007	0.006	0.009	-0.038***	0.011
Hispanic	-0.011	0.023	0.033*	0.019	0.067***	0.021	-0.056***	0.019
English Proficient	-0.008	0.015	0.007	0.010	0.025	0.021	0.009	0.014
Household Size	-0.008**	0.003	-0.004**	0.002	-0.002	0.003	-0.007**	0.003
$\bar{R}^2$	0.065		0.027		0.041		0.076	

Notes: National Agricultural Workers Survey, 2010–2016, 8,557 observations. State fixed effects are not reported to save space. Standard errors are clustered at the level of the state. *p < 0.10, **p < 0.05, ***p < 0.01.

Bibliography

- Abadie, A. and Gay, S. (2006). The Impact of Presumed Consent Legislation on Cadaveric Organ Donation: A Cross-Country Study. *Journal of Health Economics*, 25:599–620.
- Akerlof, G. A. and Kranton, R. E. (2000). Economics and Identity. *The Quarterly Journal of Economics*, (3):715–753.
- Akosa Antwi, Y., Moriya, A. S., Simon, K., and Sommers, B. D. (2015a). Changes in Emergency Department Use Among Young Adults After the Patient Protection and Affordable Care Act’s Dependent Coverage Provision. *Annals of emergency medicine*, 65(6):664–672.e2.
- Akosa Antwi, Y., Moriya, A. S., and Simon, K. I. (2015b). Access to Health Insurance and the Use of Inpatient Medical Care: Evidence from the Affordable Care Act Young Adult Mandate. *Journal of Health Economics*.
- Andreoni, J., Nikiforakis, N., and Stoop, J. (2017). Are the Rich More Selfish than the Poor, or Do They Just Have More Money? A Natural Field Experiment. *NBER Working Paper Series*, (23229).
- Antonisse, L., Garfield, R., Rudowitz, R., and Artiga, S. (2018). The Effects of Medicaid Expansion under the ACA: Updated Findings from a Literature Review. Technical report.
- Antwi, Y. A., Moriya, A. S., and Simon, K. (2013). Effects of Federal Policy to Insure Young Adults: Evidence from the 2010 Affordable Care Act’s Dependent-Coverage Mandate. *American Economic Journal: Economic Policy*, 5(4):1–28.
- Arcury, T. A. and Quandt, S. A. (2007). Delivery of Health Services to Migrant and Seasonal Farmworkers. *The Annual Review of Public Health*, 28:345–363.
- Azar, O. H. (2007). The Social Norm of Tipping: A Review. *Journal of Applied Social Psychology*, 37(2):380–402.
- Azar, O. H. (2010). Do People Tip Because of Psychological or Strategic Motivations? An Empirical Analysis of Restaurant Tipping.
- Baicker, K., Finkelstein, A., Song, J., and Taubman, S. (2014). The Impact of Medicaid on Labor Market Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment. In *American Economic Review*, volume 104, pages 322–328. American Economic Association.

- Beggs, A. and Klemperer, P. (1992). Multi-Period Competition with Switching Costs. *Econometrica*, 60(3):651–666.
- Benkert, J.-M. and Netzer, N. (2018). Informational Requirements of Nudging. *Journal of Political Economy*, 126(6):2323–2355.
- Bernheim, B. D., Fradkin, A., and Popov, I. (2015). The Welfare Economics of Default Options in 401(k) Plans. *American Economic Review*, 105(9):2798–2837.
- Bernheim, B. D. and Rangel, A. (2009). Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics. *The Quarterly Journal of Economics*, 124(1):51–104.
- Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2009). The Importance of Default Options for Retirement Savings Outcomes: Evidence from the United States. *Social security policy in a changing environment*, pages 167–195.
- Bicchieri, C. and Dimant, E. (2019). Nudging with care: the risks and benefits of social information. *Public Choice*.
- Blumenstock, J., Callen, M., and Ghani, T. (2018). Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan. *American Economic Review*, 108(10):2868–2901.
- Brown, C. L. and Krishna, A. (2004). The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice. *Journal of Consumer Research*, 31(3):529–539.
- Carr, A. (2013). How Square Register’s UI Guilts You Into Leaving Tips. *Fast Company*.
- Carroll, G. D., Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Optimal Defaults and Active Decisions. *The Quarterly Journal of Economics*, 124(4):1639–1674.
- Carvalho, L. S., Meier, S., and Wang, S. (2016). Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday. *American Economic Review*, 106(2):260–284.
- Choi, J., Laibson, D., Madrian, B., and Metrick, A. (2004). For Better or For Worse: Default Effects and 401(k) Savings Behavior. In David A. Wise, editor, *Perspectives in the Economics of Aging*, pages 81–121. University of Chicago Press, Chicago.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2002). Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance. *Tax Policy and the Economy*, 16.
- Cohan, P. (2012). Will Square’s Starbucks Deal Spark the End of Cash? *Forbes Magazine*.
- Courtemanche, C., Marton, J., Ukert, B., Yelowitz, A., and Zapata, D. (2016). Impacts of the Affordable Care Act on Health Insurance. *NBER Working Paper Series*, (22182).

- DellaVigna, S. (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature*, 47(2):315–372.
- Dellavigna, S., List, J. A., and Malmendier, U. (2012). Testing for Altruism and Social Pressure in Charitable Giving. *Source: The Quarterly Journal of Economics*, 127(1):1–56.
- DellaVigna, S. and Malmendier, U. (2004). Contract Design and Self-Control: Theory and Evidence. *The Quarterly Journal of Economics*, 119(2):353–402.
- DellaVigna, S. and Malmendier, U. (2006). Paying Not to Go to the Gym. *American Economic Review*, 96(3):694–719.
- Dubois, D., Rucker, D. D., Galinsky, A. D., Dubois, C., and Galinsky, D. D. (2015). When and Why Upper and Lower Class Individuals Behave Unethically. *Journal of Personality and Social Psychology*.
- Fiedler, M. (2018). How Did the ACA’s Individual Mandate Affect Insurance Coverage? Evidence from Coverage Decisions by Higher-Income People. Technical report.
- Finkelstein, A. N., Taubman, S. L., Allen, H. L., Wright, B. J., and Baicker, K. (2016a). Effect of Medicaid Coverage on ED Use—Further Evidence from Oregon’s Experiment. *New England Journal of Medicine*, 375(16):1505–1507.
- Finkelstein, A. N., Taubman, S. L., Allen, H. L., Wright, B. J., and Baicker, K. (2016b). Effect of Medicaid Coverage on ED Use—Further Evidence from Oregon’s Experiment. *New England Journal of Medicine*, 375(16):1505–1507.
- Fowlie, M., Wolfram, C., Spurlock, A., Annika, T., Baylis, P., and Cappers, P. (2017). Default Effects and Follow-On Behavior: Evidence from an Electricity Pricing Program.
- Frean, M., Gruber, J., and Sommers, B. D. (2017). Premium Subsidies, the Mandate, and Medicaid Expansion: Coverage Effects of the Affordable Care Act. *Journal of Health Economics*, 53:72–86.
- Gittell, R. and Tebaldi, E. (2006). Charitable Giving: Factors Influencing Giving in U.S. States. *Nonprofit and Voluntary Sector Quarterly*, 35(4):721–736.
- Goldin, J., Lurie, I., and McCubbin, J. (2019). Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach. *NBER Working Paper Series*, (Working Paper 26533).
- Goldin, J. and Reck, D. (2019). Revealed Preference Analysis with Framing Effects. *Journal of Political Economy*.
- Gonzalez, J., Kurland, B., Garber, J., Layva, E., Quintero, O., Sheriff, S., and Zakaryan, C. (2016). 2016 Taxi Factbook. Technical report, Taxi and Limousine Commission.
- Grossmann, I. and Varnum, M. E. W. (2011). Social Class, Culture, and Cognition. *Social Psychological and Personality Science*, 2(1):81–89.

- Grynbaum, M. (2009). In New York, Taxi Revenue and Tips From Credit Cards Rise. *The New York Times*.
- Guild, A., Richards, C., and Ruiz, V. (2016). Out of Sight, Out of Mind: The Implementation and Impact of the Affordable Care Act in U.S. Farmworker Communities.
- Haggag, K. and Paci, G. (2014). Default Tips. *American Economic Journal: Applied Economics*, 6(3):1–19.
- Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–2682.
- Hansen, E., Donohoe, M., and Mary’, S. (2003). Health Issues of Migrant and Seasonal Farmworkers. *Journal of Health Care for the Poor and Underserved*, 14(2):153–164.
- Hoover, H. (2019). Default Tip Suggestions in NYC Taxi Cabs. *Working Paper*.
- Jachimowicz, J. M., Duncan, S., Weber, E. U., and Johnson, E. J. (2019). When and why defaults influence decisions: a meta-analysis of default effects. pages 1–28.
- James, R. N. and Sharpe, D. L. (2007). The Nature and Causes of the U-Shaped Charitable Giving Profile. *Nonprofit and Voluntary Sector Quarterly*, 36(2):218–238.
- Johnson, E. J., Bellman, S., and Lohse, G. L. (2002). Defaults, Framing and Privacy: Why Opting In-Opting Out 1. *Marketing Letters*, 13:5–15.
- Johnson, E. J. and Goldstein, D. (2003). Do Defaults Save Lives? *Science*, 302(5649):1338–1339.
- Kaestner, R., Garrett, B., Chen, J., Gangopadhyaya, A., and Fleming, C. (2017). Effects of ACA Medicaid Expansions on Health Insurance Coverage and Labor Supply. *Journal of Policy Analysis and Management*, 36(3):608–642.
- Kandilov, A. M. G. and Kandilov, I. T. (2019). The Impact of ACA Medicaid Expansions on Agricultural Workers’ Health Insurance Coverage, Medical Care Utilization, and Labor Supply *. *Wor*.
- Keltner, D., Kogan, A., Piff, P. K., and Saturn, S. R. (2014). The Sociocultural Appraisals, Values, and Emotions (SAVE) Framework of Prosociality: Core Processes from Gene to Meme. *Annual Review of Psychology*.
- Kleven, H. J. and Waseem, M. (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *The Quarterly Journal of Economics*, 128(2):669–723.
- Landis, B. and Gladstone, J. J. (2017). Personality, Income, and Compensatory Consumption: Low-Income Extraverts Spend More on Status. *Psychological Science*, 28(10):1518–1520.

- Liddell, P. and Wilson, J. (2016). Individual Noncash Contributions, 2013. *Internal Revenue Service*.
- Lynn, M. (2004). Ethnic Differences in Tipping: A Matter of Familiarity with Tipping Norms. *Cornell Hotel and Restaurant Administration Quarterly*, 45(1):12–22.
- Lynn, M. (2007). Race differences in restaurant tipping: A literature review and discussion of practical implications. *Journal of Foodservice Business Research*, 9(4):99–113.
- Lynn, M. (2011). Race Differences in Tipping: Testing the Role of Norm Familiarity. *Cornell Hospitality Quarterly*, 52(1):73–80.
- Lynn, M. and Brewster, Z. W. (2015a). Racial and Ethnic Differences in Tipping. *Cornell Hospitality Quarterly*, 56(1):68–79.
- Lynn, M. and Brewster, Z. W. (2015b). Racial and Ethnic Differences in Tipping: The Role of Perceived Descriptive and Injunctive Tipping Norms. *Cornell Hospitality Quarterly*, 56(1):68–79.
- Lynn, M., Pugh, C. C., and Williams, J. (2012). Black-White Differences in Tipping: The Moderating Effects of Socioeconomic Status. *Cornell Hospitality Quarterly*, 53(4):286–294.
- Lynn, M. and Thomas-Haysbert, C. (2003). Ethnic Differences in Tipping: Evidence, Explanations, and Implications. *Journal of Applied Social Psychology*, 33(8):1747–1772.
- Madrian, B. C. and Shea, D. F. (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187.
- Mani, A., Mullainathan, S., Shafir, E., and Zhao, J. (2013). Poverty Impedes Cognitive Function. *Science*, 341(6149):976–980.
- Mobed, K., Gold, E. B., and Schenker, M. B. (1992). Occupational Health Problems Among Migrant and Seasonal Farm Workers. *Western Journal of Medicine*, 157(3):367–373.
- Moretti, E. and Perloff, J. M. (2000). Use of Public Transfer Programs and Privat Aid by Farm Workers. *Industrial Relations*, 39(1):26–47.
- Moretti, E. and Perloff, J. M. (2002). Efficiency Wages, Deferred Payments, and Direct Incentives In Agriculture. Technical report.
- Mullainathan, S. and Shafir, E. (2013). *Scarcity: Why Having Too Little Means So Much*.
- O’Donoghue, T. and Rabin, M. (1999). Doing It Now or Later. *American Economic Review*, 89(1):103–124.
- O’Donoghue, T. and Rabin, M. (2001). Choice and Procrastination. *The Quarterly Journal of Economics*, 116(1):121–160.
- Panhans, M. (2019). Adverse Selection in ACA Exchange Markets: Evidence from Colorado. *American Economic Journal: Applied Economics*, 11(2):1–36.

- Piff, P. K., Kraus, M. W., Côté, S., Cheng, B. H., and Keltner, D. (2010). Having Less, Giving More: The Influence of Social Class on Prosocial Behavior. *Journal of Personality and Social Psychology*, 99:771–784.
- Ravindranath, M. (2014). Whole Foods to use Square for store checkout. *The Washington Post*.
- Rubinstein, A. and Salant, Y. (2011). Eliciting Welfare Preferences from Behavioural Data Sets. *The Review of Economic Studies*, 79(1):375–387.
- Schervish, P. G. and Havens, J. J. (1995a). Do the Poor Pay More: Is the U-Shaped Curve Correct? *Nonprofit and Voluntary Sector Quarterly*, 24(1):79–90.
- Schervish, P. G. and Havens, J. J. (1995b). Explaining The Curve in the U-shaped Curve. *Voluntas*, 6(2):202–225.
- Shierholz, H., Cooper, D., Wolfe, J., and Zipperer, B. (2017). Employers would pocket \$5.8 billion of workers’ tips under Trump administration’s proposed ‘tip stealing’ rule Report. *Economic Policy Institute*.
- Simon, K., Soni, A., and Cawley, J. (2017). The Impact of Health Insurance on Preventive Care and Health Behaviors: Evidence from the First Two Years of the ACA Medicaid Expansions. *Journal of Policy Analysis and Management*, 36(2):390–417.
- Sommers, B. D., Gunja, M. Z., Finegold, K., and Musco, T. (2015). Changes in Self-Reported Insurance Coverage, Access to Care, and Health Under the Affordable Care Act. *JAMA - Journal of the American Medical Association*, 314(4):366–374.
- Stiles, R., Siegel, L., Garber, J., Neger, H., and Ullah, A. (2014). 2014 Taxicab Fact Book. Technical report, NYC Taxi & Limousine Commission.
- Taubman, S. L., Allen, H. L., Wright, B. J., Baicker, K., and Finkelstein, A. N. (2014). Medicaid Increases Emergency-Department Use: Evidence from Oregon’s Health Insurance Experiment. *Science*, 343(6168):263–268.
- Thakral, N. and Tô, L. T. (2019). Tipping and the Dynamics of Social Norms. *Mimeo*.
- Thaler, R. H. and Sunstein, C. R. (2003). Libertarian Paternalism. *Source: The American Economic Review*, 93(2):175–179.
- Thomas-Haysbert, C. D. (2002). The Effects of Race, Education, and Income on Tipping Behavior. *Journal of Foodservice Business Research*, 5(2):47–60.
- Wang, L. and Murnighan, J. K. (2014). Money, Emotions, and Ethics Across Individuals and Countries. *Journal of Business Ethics*, 125(1):163–176.
- Wang, L., Zhong, C. B., and Murnighan, J. K. (2014). The Social and Ethical Consequences of a Calculative Mindset. *Organizational Behavior and Human Decision Processes*, 125(1):39–49.

Zhou, R. A., Baicker, K., Taubman, S., and Finkelstein, A. N. (2017). The Uninsured Do Not Use the Emergency Department More—they Use Other Care Less. *Health Affairs*, 36(12):2115–2122.

Appendix A

Appendices for Chapters 1 and 2


A.1 New York City Yellow Taxi Tipping Systems

Figure A.1: NYC Yellow Taxi Payment Screen with Menu Tip Options



Notes: This is an example of a taxi screen displaying a menu of tip options and the taxi fare at the end of a taxi trip.

Figure A.2: Changes in Menu Tip Options Over Time by Vendor



Notes: This figure illustrates the changes and differences in the menus presented by the two main NYC Yellow taxi credit card machine providers (CMT and VTS).

A.2 Empirical Support for Model Assumptions

Assumption A1

A1 - One's belief about the tipping norm T_i is jointly independent of the menu of tips and the taxi fare.

While we cannot formally test assumption A1 without the parametric model, we examine whether the observed tip rate t_i is independent of the menu of tip options D . Specifically, we compare tipping decisions under two different tip menus D_1 and D_2 , where some of the options in D_2 are higher than the options in D_1 .

Suppose a passenger's preferred tip is t_i , which is not in either of the menus. She tips t_i if her decision cost is low enough not to benefit from choosing a menu tip option. Let $H(t_i|D_1)$ and $H(t_i|D_2)$ be the distribution functions of tips when passengers are shown D_1 and D_2 respectively. If t_i depends on the menu, then $H(t_i|D_1)$ and $H(t_i|D_2)$ will differ across the entire support of t_i . We expect that $H(t_i|D_2)$ will be shifted to the right of the distribution of tips under the menu with lower tip options $H(t_i|D_1)$. However, if $t_i \perp D$, then $H(t_i|D_1)$ and $H(t_i|D_2)$ will differ only around the neighborhood of the tip options across the two menus.


We use the CMT's tip menu change to assess whether t_i is independent of D by comparing the distribution of tips before and after the change. Figure 1.1 shows the distribution of tips before and after CMT's menu change. The figure shows stark differences in the share of passengers who choose menu options and in the share for tips within the neighborhood of the menu options. However, the two distributions remain relatively similar for non-menu tips. We take this as indirect evidence in support of A1.

Assumption A2

A2 - Decision costs are independent of the menu of tips and constant over time.

It is conceivable that traveler may learn after some time and become skilled at computing their preferred tip, subsequently alleviate their cognitive cost. If so, we ought to expect the share of travelers tipping at non-menu choices to increase after some time. To check for such a pattern, we contrast the distribution of tips across years where the tip menu stayed same in CMT taxis. That is the period between 2011 and 2015. We find no significant changes in the distribution of tips across the different years depicted in the figure below. We take this as partial evidence in support of A2.

Figure A.3: Overlapping Distribution of Tip (%) in Years with No Change in CMT Tip Menu



Notes: These figures compare the distribution of tips between 2011 and 2011. This period is five years after CMT—a Yellow taxi credit card machine vendor—changed its tip menu in 2011 from 15%, 20%, and 30%, to showing 20%, 25%, and 30%. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip. The points in the figure are estimates from non-overlapping bins of width 1% for tips between 0.5% and 35.5% of the taxi fare. The tips rates are truncated at 35.5% where the share becomes essentially zero.

Verifying Assumption A3

A3 - Decision costs are jointly independent of the taxi fare F_i and one's preferred tip t_i^* .

Because we do not observe c_i , there is no straightforward way to test A3. Therefore, we check first for evidence that c_i is independent of one's preferred tip t_i^* . Then, we check for evidence that c_i is independent of the taxi fare F .

$c_i \perp F_i$. The cognitive cost c_i associated with computing a tip is independent of the taxi fare. There is no straightforward way to test this assumption, because we do not observe c_i . However, we find it reasonable to assume that passengers find it easy to compute the dollar amount of their tip rate if the taxi fare is a multiple of \$10. Thus, if percent to dollar conversions are relatively easier for fares that are multiples of \$10, then passengers should be less likely to choose a menu tip option for these fares.

To test the previous statement, we regress a dummy variable that equals one if the tip is a menu tip option and zero otherwise on a set of dummies that indicate fares that are multiples of \$10. If it is significantly easier to calculate tips when the fare is a multiple of \$10, then c_i will be notably lower, and passengers will be less likely to choose menu tips. Hence, we should observe a negative coefficient on the dummy variable for fares that are multiples of \$10.

Table A.1 shows estimates from this regression. The coefficients on the dummy variables for fares that are a multiple of \$10 are all positive or not statistically significantly distinguishable from zero. This suggests that passengers are just as likely if not more likely to choose a menu tip option when the fare is a multiple of \$10 than otherwise. This is in direct opposition to what we predicted. Although, this observation is not sufficient evidence to establish assumption A3, it does suggest that the data seems consistent with it.

$c_i \perp t_i^*$. Since c_i is unobservable, we cannot measure c_i for all possible tip rates t_i . However, if we postulate that tips are smooth across all fares (that is, the distribution of t_i does not have point masses or holes), then we can take advantage of the fact that some percentage tips (such as 10%) are easy for passengers to compute. Then we can see whether these cases create point masses. More formally, suppose that the distribution of tips is smooth across all fares and a 10% tip rate (and possibly a 15% tip rate) is fairly easy to compute. Then there should be a point mass at 10% (and possibly at 15%) in the distribution of tip rates.


We use data from 2014 and restrict attention to tips less than 20% of taxi fare. We check for point masses at 10% and 15% in the distribution of tips. If 10% and 15% are fairly easy to compute, then a notably large share of passengers should be concentrated at these two rates relatively to other tips. Figure A.4 shows a bar graph of the shares of passengers whose tips fall in non-overlapping bins of width 1%. Most tips are concentrated between 8% and 18%. However, the shares of tips in bins that include 10% and 15% are not any higher than the majority of the other tip bins. Rather, the highest concentration of tips is at 12%. While this evidence is not a formal test of assumption A3, the data are consistent with this assumption.

Table A.1: Evidence for Assumption A3

	Dependent Variable 1(Menu Tip)
1(Taxi Fare = \$10)	0.104*** (0.0004)
1(Taxi Fare = \$20)	0.050*** (0.001)
1(Taxi Fare = \$30)	0.037*** (0.002)
1(Taxi Fare = \$40)	0.056*** (0.004)
1(Taxi Fare = \$50)	0.038*** (0.007)
1(Taxi Fare = \$60)	0.005*** (0.013)
1(Taxi Fare = \$70)	-0.003 (0.026)
1(Taxi Fare = \$80)	-0.065 (0.065)
Constant	0.601*** (0.0001)
Observations	41,620,580
R^2	0.002

Note: The data are standard rate NYC Yellow taxi trips in CMT taxi cabs from 2014. The trips are limited to fares paid for using a credit/debit card, has no tolls, and passengers leave a positive tip. $p < 0.1$; $p < 0.05$; *** $p < 0.01$

Figure A.4: Distribution of Tips $< 20\%$



Note: This plot shows the distribution of tips in non-overlapping bins of width 1% between 0.5% and 19.5% tip rate using data from CMT taxi trips in 2014.

A.3 Figures for Nonparametric Approach

A.3.1 Tips Before and After CMT Tip Menu Change by Fare

Figure A.5 shows the distribution of positive tips (truncated at the tip rate of 19.5%) before and after CMT—a New York City taxi credit card machine vendor—changed the menu of tips that is presented to taxi passengers in 2011. The figures correspond to the subset of taxi trips whose fare falls within different ranges of the taxi fares. CMT presented customers with three tip options in percentages (15%, 20%, and 25%) before the menu change. After CMT removed the lowest tip option (15%) and added a higher percentage option (30%), so that it offered 20%, 25%, and 30%. The shaded bars present the distribution of tips one year before the menu change, and the un-shaded bars show the distribution of tips about a year after the change. Data from 2010 and 2011 standard rate taxi trips, with no tolls, paid for via a CMT credit card machine are used in these figures.

Figure A.5: Distribution of Tip (%) less than 20% Before and After CMT Tip Menu Change




Figure A.5 continued




Figure A.5 continued




Figure A.5 continued





Figure A.5 continued




A.3.2 Bounds On the Conditional CDFs of Decision Costs


Figure A.6 shows the lower and upper bounds for the CDF of decision costs computed for passengers who tip at rates less than 18%. The computation of these bounds relies on CMT's change in the menu of tips that is presented to taxi passengers in 2011. These bounds are computed using the increase in the share of passengers who tip at a particular rate at different levels of the taxi fare. Generally, for a given fare F and tip rate $t < 20\%$, the lower and upper bounds for the decision cost of switching from 15% to some non-menu tip t is given by $[|0.15 - t|F, |0.20 - t|F]$. Data from 2010 and 2011 standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with positive non-menu tips (that are not round-number dollar amounts) are used in this figure.

Figure A.6: Conditional CDFs of Decision Costs


Tip Rate = 1%




Tip Rate = 2%




Tip Rate = 3%



Tip Rate = 4%



Tip Rate = 5%



Tip Rate = 6%




Figure A.6 continued






Figure A.6 continued



A.4 Figures for Semiparametric Approach

Figure A.6 shows the lower and upper bounds for the CDF of decision costs computed for passengers who tip at rates less than 20%. The computation of these bounds relies on CMT's change in the menu of tips that is presented to taxi passengers in 2011. These bounds are computed using the increase in the share of passengers who tip at a particular rate at different levels of the taxi fare. Generally, for a given fare F and tip rate $t < 20\%$, the lower and upper bounds for the decision cost of switching from 15% to some non-menu tip t is given by $[|0.15 - t|F, |0.20 - t|F]$. Data from 2010 and 2011 standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with positive non-menu tips (that are not round-number dollar amounts) are used in this figure.


Figure A.7: Predicted Probabilities by Level of Taxi Fare of Choosing Tips $< 20\%$




Note: This figure shows the estimated predicted probabilities for non-menu tip rates below 20% as functions of the fare. The probabilities are computed from an ordered logistic regression using data limited to trips with tip rates 20% or less and selected from CMT taxi trips in 2014. The range of fares used in this analysis is between \$3 and \$30. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip.

Figure A.8: Conditional CDFs of Decision Costs


Tip Rate = 1%




Tip Rate = 2%




Tip Rate = 3%



Tip Rate = 4%



Tip Rate = 5%



Tip Rate = 6%




Figure A.8 continued






Figure A.8 continued



A.5 Empirical Support for Parametric Approach

Figure A.9: Share of Menu Tips and Average Tip by Fare



Notes: Figures A.9a is a binned scatter plot that illustrates the relationship between the share of passengers who choose any one of the suggested menu tips presented at the end of a taxi ride at different levels of the taxi fare. Figures A.9b is a binned scatter plot that illustrates the average tip rate at different levels of the taxi fare. The data used in both figures are from 2014 standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip amount.

Table A.2: Heckman Selection Correction Estimates: First-Step Probit Estimates

	<u>2014</u>	<u>2010 – 2011</u>
Dependent Variable: 1(Tip=Non-Menu Tip)		
Taxi Fare	−0.02917*** (0.00004)	−0.04046*** (0.00004)
Number of Passengers	−0.44228*** (0.00041)	−0.40071*** (0.00035)
1(Post Menu Change)		0.11246*** (0.00043)
1(Round Number Tip)	Yes	Yes
Observations	41,620,591	59,861,675
Log Likelihood	−14,436,606	−20,633,368
Akaike Inf. Crit.	28,873,217	41,266,745

Notes: This table reports the first stage probit estimates of the Heckman selection correction model presented in Table II, Panel A. In column (1), we use CMT taxi trips from 2014, and in column (2), we use CMT taxi trips from 2010-2011. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive tip. Both columns are estimated without a constant term. We report robust white standard errors in parenthesis. *p<0.1, **p<0.05, ***p<0.001.

Table A.3: OLS Estimates of Tipping Norms and Norm Deviation Cost Parameter


	<u>2014</u>	<u>2010 – 2011</u>
	(1)	(2)
Dependent Variable: Tip rate		
Taxi Fare	−0.00326*** (0.00001)	−0.00420*** (0.00001)
1(Post Menu Change)		−0.00649*** (0.00005)
Constant	0.20458*** (0.00013)	0.20640*** (0.00007)
Norm Deviation Cost Parameter $\hat{\theta} \left(= -\frac{1}{2\beta} \right)$	151.9757	119.0844
1(Round Number Tip)	Yes	Yes
Observations with Non-Menu Tips	16,394,917	25,206,358
R ²	0.01639	0.04196

Notes: This table reports estimates of the tipping norm T_i and the norm deviation cost parameter θ from an OLS regression that does not account for sample selection bias. In column (1), we use CMT taxi trips from 2014, and in column (2), we use CMT taxi trips from 2010-2011. The sample restriction are standard rate taxi trips, with no tolls, paid for via a CMT credit card machine along with a positive non-menu tip. We report robust white standard errors in parenthesis. *p<0.1, **p<0.05, ***p<0.001.


A.6 Figures for Welfare Calculations

Figure A.10: Distribution of Tip (%) by Type of Tip Menu


a) Previous Tip Menu




c) Tip-Maximizing Menu



b) Current Tip Menu



d) Utility-Maximizing Menu



Note: This figure shows the model predicted distribution of tips for four different tip menus namely used for the welfare analysis presented in Table IV. The predictions are made using the estimated parameters of the model. The data shown are tips between 0.5% and 35.5% of the taxi fare. The tips rates are truncated at 35.5% where the share becomes essentially zero.