

How Difficult is Tipping?

Using Structural and Non-Structural Approaches to Estimate Decision Costs

Job Market Paper ([Latest Version](#))

Kwabena B. Donkor*

October, 2019

Abstract

Does a menu of recommended tips presented with a bill influence how much customers tip? Analyzing three quarters of a billion passenger tips in New York City Yellow taxis, we use changes in the menu presented to passengers to nonparametrically estimate that the decision cost of not following a menu is about \$1.89 (16% of the average taxi fare of \$12.17). To disentangle the mechanisms behind decision costs, we use a model in which customers' choices are based on their beliefs about the social norm tip. They incur a norm deviation cost for not conforming to the tipping norm and a cognitive cost from computing a non-menu tip. Our estimate of the distribution of passenger beliefs about the social norm tip averages at 19.8% of the taxi fare. Customers incur a norm deviation cost (shame) of \$0.42 when they tip five percentage points less. The cognitive cost of calculating a non-menu tip ranges from \$1.26 to \$1.41 on average. We also find that taxicabs currently present customers with a nearly tip-maximizing menu, and this menu increases tips by 12.4% relative to not presenting a menu. Taxicab companies appear to have learned over time to converge to the tip-maximizing menu. Our welfare calculations suggest that the current tip menu increases welfare by \$180.25 million per year relative to not presenting a menu.

Keywords: Menu Suggestions, Defaults, Norms, Decision Costs, Firm Learning.

JEL Codes: D91, D12, L80, L92

*Kwabena Donkor: Ph.D. candidate at UC Berkeley's Agricultural and Resources Economics Department (k.b.donkor@berkeley.edu). First and foremost, I thank God for helping me through this whole process. I thank Jeff Perloff, Stefano DellaVigna, Miguel Villas-Boas, Przemyslaw Jeziorski, Aprajit Mahajan, Giovanni Compiani, Ben Handel, Joseph Farrel, and Dmitry Taubinsky for their comments, advice, and guidance at the various stages of this project. Thanks to Kareem Haggag, Yousef Bensarti, Neil Thakral, Supreet Kaur, Larry Karp, Sofia Villas-Boas, Alain de Janvry, Ethan Ligon, Elisabeth Sadoulet, James Salle, and Catherine Wolfram for their invaluable comments. Thanks to all the students and faculty at the Stanford GSB Marketing seminar, UCB Haas Marketing seminar, UCB Psyc & Econ lunch seminar, UCB IO seminar, and the ARE ERE lunch seminar for comments and discussions. Last, I thank Hanan Wasse and all ARE students for their support and contributions. All errors are mine.

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.

Since tipping is not compulsory and requires costly effort, a passenger should either choose a menu tip to avoid the cost of computation or choose not to tip at all. However, a significant number of passengers (38% of 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. Hence, passengers who choose non-menu tips help to identify the cost (“shame”) of not obliging to the tipping norm (norm deviation cost).

We use both a nonparametric approach and a structural model to assess the effects from using menus. The advantage of the nonparametric approach is that we are able to identify a lower and upper bound of the unobservable decision cost (norm deviation cost + cognitive cost) of switching from a menu tip to one’s preferred tip through the following observation. 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. Suppose I observe a passenger picking the menu option 15% when presented, but then switches to their preferred tip (say 11%) when faced with the 20% menu. I conclude that the decision cost to switch is more than the difference between the 15% tip and the preferred tip and less than the difference between the 20% tip and the preferred tip. This effectively places observable monetary bounds on the decision cost of switching. With this reasoning, we assume that passengers who choose non-menu tips reveal their preferred tip. We then use how passenger tipping choices are affected by changes in the tip menu presented in Yellow taxicabs to identify bounds on the distribution of decision costs. This approach does not allow us to decompose decision costs into norm deviation and cognitive costs. To separately identify these two components, we place structure on our tipping behavior model and estimate the parameters of the model by method of moments. The structural model recovers the unobserved distribution of beliefs about the tipping norm within the population of taxi passengers, and separately identifies the norm deviation and cognitive cost.

Because social norms are important phenomena that deeply guide human behavior, an empirical analysis may help us gain a better understand of how norms influence consumer choices. However, in a field setting, norms are difficult to quantify and study in a scientific manner. Our model enables us to empirical analyze peoples unobserved beliefs about the social norm tip and the cost they face for not conforming to the norm. While an empirical estimate 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, policy counterfactuals are also informed by this separation of overall decision costs. For example, a menu that will maximize either

¹Sixty percent choose menu options and very few (at most 2%) choose not to tip. Cash tips are not observed in the data and therefore not part of our analysis sample.

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 (16% of the average taxi fare of \$12.17). Using the structural model, we estimate that the unobserved distribution of beliefs about the tipping norm across passengers averages at 19.8% of the taxi fare, which is around the average tip rate in the data (19%). The norm deviation cost associated with not conforming to the norm is large relative to the fare. For instance, a passenger who decides to tip five percentage points less than her perception of the norm incurs a norm deviation cost of about \$0.42 (3.5% of the average fare). We estimate that the average cognitive cost of computing a non-menu tip is between \$1.26 and \$1.41 (about 10% to 12% 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 16.87% to 18.96% (i.e., a 12.4% 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 \$180.25 million per year under the current tip menu relative to case where passengers are not offered 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

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](#)).

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.

A large 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 \(2001\)](#); [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](#); [O'donoghue and Rabin, 2001](#)). Consumers may perceive 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 effects, 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 ([Bernheim et al., 2015](#)).

Passenger tipping decisions for 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

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](#)).

tipping changed across menus and which menu extracted the most tips for taxi drivers.

A study close to ours is [Haggag and Paci \(2014\)](#). Using a clever regression discontinuity design, they explore whether menus with higher default tip amounts induce consumers to tip more. Using NYC Yellow taxi trips from 2009, they find that higher tip suggestions increase the amount tipped, but may cause some passengers to avoid tipping altogether. Another study close to ours is [Thakral and Tô \(2019\)](#). This study uses a change in the NYC Yellow taxi fare rate in 2012 in a structural model to evaluate the welfare consequences of adherence to the social norm of tipping in taxicabs and the dynamics of social norms. Our study differs from these two studies in several ways. First, neither study estimate the total decision costs involved when passengers tip, but we do. Second, neither study estimate passengers unobserved beliefs about the social norm tip, but we do. Third, although [Thakral and Tô \(2019\)](#) only estimates the cost of deviating from the norm, we separately estimate both the norm deviation cost and the cost of computing a non-menu tip when passengers do not follow the menu tip options. Third, we estimate the tip-maximizing menu and assess the implications of different tip menus for the social welfare from tipping.

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 study presents empirical evidence from a field setting on how social norms affect consumer choices. The literature in psychology refers to one’s expectation of what others do in a comparable situation as a descriptive norm. These descriptive norms influence one’s beliefs and preferences, and hence one’s choices (see discussion in [Bicchieri and Dimant \(2019\)](#)). In the context of tipping, consumers usually have a belief about a socially acceptable amount to tip and try to act accordingly. However, people’s perceptions about social norms are mostly unobservable or cannot be empirically quantified. This paper uses a novel approach to recover peoples unobserved beliefs about the social norm tip. In addition, we quantify the cost that consumers face for not conforming to the social norm.⁴

⁴In a related study about tipping in restaurants, [Azar \(2005\)](#) presents a theoretical discussion about the social norm of tipping and its implications for social welfare.

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. Both [Bernheim et al. \(2015\)](#) and [Goldin and Reck \(forthcoming\)](#) provide a theoretical framework that thoroughly discusses the welfare implications of default options. With the use of our structural 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 2 describes the tipping systems used in NYC Yellow taxis and gives a summary of our analysis data. Section 3 lays out a model for tipping in taxis. Section 4, uses a nonparametric approach to estimate decision costs. Section 5 presents a structural model used to estimate the social norm tip and disentangle decision costs into cognitive and norm deviation costs. Section 6, conducts counterfactual exercises to predict the tip-maximizing menu and evaluate its implications for social welfare. Section 7 concludes.

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 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 in 2011. Prior to 2012, VTS offered a menu of dollar amounts 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

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.

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 6.2.

Our 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 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). CMT trips from 2010 and 2011 are what we use in our nonparametric estimation of decision costs, and the data from 2014 are what we primarily use in our structural estimation. Between January 2010 and

⁶In the appendix section, Figure A.1 shows a typical screen displaying menu tip options, and figure A.2 shows the menu options 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.

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 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 from 2010 to 2011. 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 the menu tips decreased by about one-fifth after the menu 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 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%.

3 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). Azar (2010) tests for evidence of these two motivations by studying customers in restaurants. The author finds no evidence that customers tip strategically but, they do tip for social/psychological reasons. We therefore consider strategic tipping to be inconsequential and ignore it in our model of tipping.

We model a passenger’s decision to tip when provided with a menu of tips. Passenger i gives a tip of $t_i\%$ (a percentage of the taxi fare). She believes that the social norm tip is $T_i\%$ (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 , 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

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

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 $-c_i \times \mathbf{1}\{t_i \notin D\}$

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 $t_i \notin D$ if the benefit of tipping at that rate (denoted as B_i) is greater than choosing a menu tip rate d_j of $j = 1, 2, \dots, n$ in tip menu D . That is

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

$$B_i = (d_j - t_i)F > \Delta\nu + c_i \tag{2}$$

Equation (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.⁹

First, we use a nonparametric approach to place bounds on the total decision costs from tipping. This approach uses a change in the menu, which caused a change in the share of passengers who choose non-menu tips. This approach requires two weak assumptions (presented below), but does not allow us to separately identify the norm deviation and cognitive costs. In a second approach, we add stronger, structural assumptions, which allow us to separately identify the norm deviation and cognitive costs.

4 Nonparametric Estimation of Decision Costs

In 2011, the menu provider CMT changed its menu options from 15%, 20%, and 25% to 20%, 25%, and 30%. Figure 1 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. The key intuition is that, all else equal, changing the menu options (d_j) changes the values on both sides of the inequality in equation (2). Hence, the sign of the inequality might change for passengers who are on the

⁹If 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.

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 at a menu option.

This approach is nonparametric. Estimating bounds on these decision costs does not require us to make assumptions about how passengers decide how much to tip. However, we need to make two assumptions for this exercise:

A1 - Decision costs are constant across years.

A2 - One’s perception of the tipping norm T_i is jointly independent of the menu of tips and the taxi fare.

We corroborate assumption A1 by comparing the distribution of tips across different years where there were no changes in the tip menu.¹⁰ In order to test A2, we need information on the unobserved tipping norm T_i . However, we are able to verify A2 with the aid of a structural model presented in section 5 below.¹¹

4.1 Constructing Bounds

By inspecting non-menu tips in figure 1, we find significant increases in passenger tip rates below 20% after the CMT menu change; however other tips remain unchanged. Therefore, to compute the bounds for decision costs, we restrict attention to tips at or below 20%.¹² Thus, the relevant menu options were 15% and 20% before the change and only 20% after.

Assume that passenger i prefers to give a non-menu tip rate t_i on taxi fare F . For example, suppose t_i is 10% and F is \$10. Now suppose that before the removal of the 15% menu option, this passenger tipped 15%, but after, she chose a 10% tip instead. This implies her benefit from tipping at 10% when 15% was a menu option is less than the decision costs to compute 10%. That is, $B_i = (d_j - t_i) \times F = (0.15 - 0.10) \times \$10 = \$0.50 < \Delta\nu + c_i$ (equation (2)). Notice 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. Therefore, a lower bound for the decision cost of tipping $t_i = 10\%$ is \$0.50. Of course, her actual decision cost could be larger; hence this amount is a lower bound.

¹⁰See figure A.10 in appendix section A.5. Benzarti (2017) uses a similar approach.

¹¹Details can be found in appendix section A.6.

¹²After 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.

Similarly, $|0.20 - 0.10| \times \$10 = \1 is an upper bound of her decision cost of computing a 10% tip. If the cost of deciding to tip 10% is more than \$1, then she benefits by choosing the 20% menu option. Generally, for a given fare F , the lower and upper bounds for the decision costs of switching from 15% to some non-menu tip t_i are given by $[|0.15 - t_i|F, |0.20 - t_i|F]$.

4.2 Estimates of 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 $[|0.15 - t_i|F, |0.20 - t_i|F]$.

We use the same data as in figure 1, but focus on tip rates that are 20% and below, which reduces the number of trips by about 20%. 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. First, we group taxi fares into 29 non-overlapping bins of width \$2: $[\$3, \$5], (\$5, \$7], (\$7, \$9] \dots (\$59, \$61]$ 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 tip falls within $[0.5\%, 1.5\%]$, 2% is the share whose tips fall within $(1.5\%, 2.5\%]$, and so forth.

Second, we compute the share of tips in each tip bin for the subset of the data that fall in a particular fare bin. For example, figure 2 show the distribution of the subset of tips for fares between \$21 and \$23. The figure shows that higher shares of passengers choose a non-menu tip after 15% is removed from the menu.¹⁶

Third, for each tip bin, we use the midpoint of the fare bin to compute the lower and upper bounds for decision costs. For example, for all taxi fares that fall within fare bin $(\$9, \$11]$, \$10 is used to compute the relevant bounds. We then combine the increase in the share of passengers $\Delta S_{(t,F)}$ who tip $t\%$ to construct bounds for the CDF of decision costs. For example, figure 3 shows the computed bounds for the CDF of decision cost conditional on 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 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.4 in the appendix section A.2.1 shows similar figures that correspond to other fare bins.

tip rate of 10%.¹⁷

Suppose that the midpoint of the estimated bounds of decision costs is similar to the true decision cost. If so, 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. Figure 4 shows the estimated CDF. From this distribution, the average cost decision cost of making a calculated choice versus choosing a menu tips is \$1.89 (16% of the average taxi fare \$12.17).

As an alternative to the nonparametric approach above, we conduct a novel but complementary exercise that identifies the distribution of decision costs as well. This alternative approach uses a different source of variation, data and strategy to estimate decision costs. Specifically, we employ a semiparametric approach that uses changes in the share of passengers who choose non-menu tips as the taxi fare increases to estimate the distribution of decision costs. Estimates from this semiparametric approach are similar to the nonparametric approach. In fact, the average decision cost is estimated to be \$1.64 (14% of the average taxi fare \$12.17) compared to \$1.89, the average from the nonparametric approach. Details of this approach of estimating decision costs are in section A.4 of the appendix.

5 Parametric Estimation of Decision Costs

The nonparametric approach 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 deviating from the perceived tipping norm and the cognitive cost involved in 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, independent of social influences. 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 our tipping behavior model.

We continue to rely on assumption A2 and also specify a particular utility or loss function. This extra structure allows us to separately identify the norm deviation cost, the cognitive cost, and the social norm across taxi passengers.

¹⁷Figure A.5 in appendix section A.3 shows the computed bounds for other tip rates.

5.1 The Structural Model

Analogous to equation (1), passenger i chooses a tip to maximize her utility or minimize her loss represented by

$$U_i = \underbrace{-t_i F_i}_{\text{Tip paid}} - \underbrace{\theta (T_i - t_i)^2}_{\text{Norm deviation cost}} - \underbrace{c_i \times \mathbf{1}\{t_i \notin D\}}_{\text{Cognitive cost}} \quad (3)$$

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—which we assume is quadratic. 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.¹⁸ The third term, $-c_i \times \mathbf{1}\{t_i \notin D\}$, is passenger i 's cognitive cost of computing her preferred tip, where c_i is a fixed dollar cost of calculating $t_i F_i$, and $\mathbf{1}\{t_i \notin D\}$ is an 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 and zero otherwise.¹⁹

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 .

Let B_i be the benefit from choosing $t_i \notin D$ rather than a (higher) menu default, $B_i = (d_j - t_i) F_i$. The cost of doing that 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

¹⁸Because the norm deviation cost of deviating from 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 perception of the tipping norm T_i . So passenger i 's disutility from tipping can be written as

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

However, using this model does not affect any of our estimates from equation (3) in a significant way. This is because, in equation (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 (3) in our analysis.

¹⁹This model is similar to the one presented in Azar (2004). The main difference is that the current model takes into account that consumers are presented with a menu, and they incur a cognitive cost when they don't choose from the menu.

rate if the benefit B_i of tipping her preferred tip t_i exceeds the extra cost from not choosing a menu tip:

$$B_i = (d_j - t_i)F_i > \theta [(T_i - t_i)^2 - (T_i - d_j)^2] + c_i \quad (4)$$

Given that $t_i < d_j$, 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 5a shows a binned scatter plot of the share of passengers who choose a menu tip at different levels of the fare. It is clear from the figure that passengers are less likely to choose from the menu at higher fares. With equation (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). We ignore the cognitive cost c_i because of the indicator function $\mathbf{1}\{t_i \notin D\}$. From the first-order condition, we find that the optimal tip is

$$t_i^* = T_i - \frac{1}{2\theta}F_i \quad (5)$$

According to the first-order condition, passenger i 's preferred tip t_i^* is less than her perception of the social norm T_i . 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.

Another implication of the first-order condition is that the optimal tip rate falls as the fare increases ($\frac{dt_i^*}{dF} = -\frac{1}{2\theta} < 0$). This observation generally holds in the data. Figure 5b, a binned scatter plot of the average tip rate across different fare levels, shows that the average tip rate falls with the fare.²⁰

5.1.1 Assumptions

To estimate the parameters in the proposed model (equation (3)), we rely on assumption A2 and an additional assumption, A3.

Again, assumption A2 holds that passenger i 's perception of the social norm, T_i , is jointly independent of the taxi fare F_i and the tip menu D . The new assumption is:

²⁰Some passengers use other heuristics such as tipping a fixed dollar amount or rounding off the taxi 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 this behavior later in our analysis.

A3 - The cognitive cost c_i is 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. 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 such as 12% and 14%, which are relatively harder 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).²¹

5.2 Estimation Procedure

The parameters to be estimated are a passenger’s perception of 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 (5), to estimate the unobserved distribution of both T_i and θ . This leaves the distribution of c_i to be estimated, which we compute via a Minimum Distance Estimator.

Specifically, the first-order condition allows us to empirically estimate the distribution of T_i and θ . The advantage here is that, first, we need not make any distributional assumptions regarding T_i . Second, θ is directly estimated in the same equation used to recover T_i . In addition, because this approach allows us to empirically estimate the distribution of T_i , we can now test assumption A2. We do this by comparing estimates of T_i before the CMT tip menu change to estimates of T_i after the CMT menu change and the TLC taxi fare increase.²²

5.2.1 Estimation of T_i and θ

We estimate equation (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

$$t_i = \alpha_T + \beta F_i + \varepsilon_i, \tag{6}$$

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.²³ The challenge with estimating equation (6) using all

²¹These empirical observations are further discussed in section A.7 of the appendix.

²²Details of this exercise can be found in section A.6 of the appendix.

²³Note 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 (6). However, we expect this

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 (6) are likely biased in an OLS regression. However, we observe t_i^* for the subsample of passengers who choose non-menu tips. We focus on this subsample to estimate θ and the distribution of T_i by relying on assumptions A2 and A3, and correcting for the possible sample selection concern for this subset of tips.

The link between (5) and (6) is that, α_T as the population mean 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 perception of the tipping norm relative to the the population mean. To recover an estimate of T_i (denoted as \hat{T}_i), note that

$$\hat{\varepsilon}_i \equiv t_i^* - \hat{\alpha}_T + \frac{1}{2\hat{\theta}}F_i$$

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

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

Notice that, equation (6) can be rewritten in the regression form as

$$E[t_i^*|F_i, c_i] = \alpha_T + \beta F_i + D[\varepsilon_i|F_i, c_i]$$

According to A2, $T_i \perp (F_i, D)$, thus $\varepsilon_i \perp F_i$. So, the decision to choose a non-menu tip depends solely on one's cognitive cost c_i . This conclusion follows because the cognitive cost c_i is what influences a passenger to choose a menu option or otherwise. 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 (6) may be problematic due to sample selection. The concern here is the possibility that $E[\varepsilon_i|c_i] \neq 0$. For example, the subsample of passengers who give non-menu tips may systematically have lower levels of cognitive costs compared to passengers who choose menu tips. To address this concern, we employ a 2-step Heckman selection correction approach and compare it side-by-side with the naïve OLS estimates.

In the first step of the selection model, we use a probit equation to estimate the probability

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.

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 an added instrument, which is the taxi driver’s report of the number of passengers on each trip. We reason that a passenger is more likely to choose a menu tip if they have co-passengers, but their preferred tip is the same with or without co-passengers.²⁴

When passengers are shown the tip screen at the end of the trip, they may not want to delay other co-passengers by taking the time to deliberate on how much to tip. Lastly, it is important to note that the number of co-passengers does not enter the utility function defined in equation (3). Thus, the number of co-passengers is an excluded instrument with respect to the structure of our model. Table A.3 column (2) in section A.6 of the appendix shows the results from the first step probit estimation of the Heckman selection correction model. In the second step, we estimate equation (6) 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 possible selection bias.

We 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 and thus might affect the estimate of the distribution of the perceived tipping norm T_i . We control for this round-number bunching by including an indicator variable for round number tips in our regression equations to capture the rounding effects. Then, when we estimate T_i , we omit the contribution of the round-number indicator.²⁵

Estimates of T_i and θ : Generally, we find that the average perceived norm T_i across all passengers is to tip around 20% of the taxi fare. Tipping five percentage points less than the norm results in a norm deviation cost of about \$0.42 (3.5% of the average taxi fare of \$12.17). We also find that the selection concerns about using the subsample of non-menu tips do not have a significant impact on estimating θ and the distribution of T_i .

Table 2 Panel A compares estimates of equation (6) using OLS in column (1) to estimates from the Heckman selection correction model in column (2). Both columns show the same statistically significant coefficient estimate on the taxi fare: $\hat{\beta} = 0.003$. This estimate implies that $\hat{\theta} = \frac{1}{2\hat{\beta}} \approx 166.7$. The units of θ are dollars per percent squared ($\$/\%^2$). Therefore, the

²⁴A concern that might violate the exclusion restriction is that passengers who ride in groups may be inherently different from those who ride alone (e.g., business travelers versus tourists). However, the data at hand does not allow us to distinguish between different types of travelers.

²⁵This 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. However, our estimates do not change even if we do not control for round-dollar tip amounts.

dollar value of the norm deviation cost for tipping five percentage points²⁶ less than the norm (20%) is $\theta \times (d - t)^2 = 166.7 \times (0.20 - 0.15)^2 = \0.42 . Thus, for the average taxi fare of \$12.17, the passenger saves \$0.61 at a cost of \$0.42. The constant shows an estimate of the average perceived tipping norm across all passengers. The OLS estimate of the average perceived norm of the tip rate is 0.205, while the selection model estimate is 0.198. These two estimates are similar but statistically different. The coefficient on the Mills ratio term in column (2) is very small (0.005) but statistically different from zero. 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 θ and the distribution of T_i .

Using the results from Table 2 column (1), figure 7a compares the reduced-form estimate of the distribution of the perceived tipping norm T_i and the implied preferred tip t_i^* against the observed tip t_i in the data. The figure shows that the distribution of the perceived tipping norm (solid line) is approximately uniformly distributed between 12% and 21%. Very little to no mass of the distribution is observed beyond the stated range.

With \hat{T}_i and $\hat{\theta}$ in hand, we can compute the preferred tip \hat{t}_i^* using equation (5). We depict this as an un-shaded bar graph in figure 7a. The majority of passengers' preferred tip rates are below 20%. Figure 7b shows an analogous figure using the Heckman selection correction model. The two figures are more or less the same.

5.2.2 Upper Bound of Cognitive Costs for Passengers with Non-Menu Tips

Passengers who give non-menu tips find it more beneficial to compute their preferred tip instead of choosing a menu tip. Therefore, with knowledge of their perception of the norm tip T_i , θ , and the fare F_i , we can use equation (4) to compute the relevant cognitive cost above which such passengers would opt for a menu tip as follows

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

For each non-menu tip rate t_i , we use the nearest menu tip rate d_j from above as the reference menu option for computing \bar{c}_i . For example, for a non-menu tip rate of 15%, the relevant menu option is 20%, and for a non-menu tip rate of 23%, the relevant menu option is 25%, and so forth.²⁷ Notice that we can also compute the relevant norm deviation costs as $\hat{\theta} \left(\hat{T}_i - t_i \right)^2$.

²⁶We choose 5% because the average non-menu tip rate is around 15%.

²⁷Because there is no menu option above 30%, we do not use non-menu tips above 30% in this analysis. Thirty seven percent of non-menu tip rates are above 30%.

Figure 6 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. We find that the averages of these distributions are \$1.04 for cognitive cost, and \$0.30 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.

5.2.3 Estimating Cognitive Cost for All Passengers

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 for passengers who choose non-menu tips (figure 6).

The passenger’s objective is to give a tip that maximizes her utility. However, there is no analytical solution to equation (3) and hence no corresponding closed-form expression. This is because the derivative of the indicator function $\mathbf{1}\{t_i \notin D\}$ is not well defined. We circumvent this problem by using a Monte Carlo procedure of an algorithm chooses one of the menu options or a non-menu tip. The algorithm follows these steps:

- Step 1: For each observed taxi fare F_i , there is a random draw of \hat{T}_i from the distribution estimated in the section above and a draw of a corresponding c_i from an exponential distribution with rate parameter λ .
- Step 2: \hat{t}_i^* is then computed as defined in equation (5) using F_i , \hat{T}_i , and $\hat{\theta}$.
- Step 3: Using equation (3), the utility levels for leaving a non-menu tip $\hat{t}_i^*(U^{\hat{t}_i^*})$ and all of the menu tips: U^{d1} ($d_1 = 20\%$), U^{d2} ($d_2 = 25\%$), and U^{d3} ($d_3 = 30\%$) are computed.
- Step 4: The algorithm then chooses the tip that results in the highest utility by comparing the four levels from step 3.

To find a value of λ such that the model (equation (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.

As a primary set of moments used to identify λ , we construct sample statistics by dividing tip rates into 50 non-overlapping one percent bins, namely 1%, 2%, 3%...50%. 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 a simulated method of moments algorithm to estimate λ . We proceed as follows: 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 algorithm minimizes the criterion function $Q(\lambda|\hat{T}_i, \hat{\theta}) = g'\hat{W}g$, where \hat{W} is some positive-definite weight matrix that is a function of the realized data. Effectively, when minimizing the criterion function $Q(\lambda|\hat{T}_i, \hat{\theta})$, we match the sample statistics to their simulated analogues under the model.

We employ a two-step procedure to compute the model parameters. In the first step, an identity matrix is used as a preliminary weight matrix to estimate λ —denoted as $\hat{\lambda}$. In the second step, $\hat{\lambda}$ is used to predict a set of realized tips via equation (3). Next, the predicted tips are used to compute $m(\lambda|\hat{T}_i, \hat{\theta})$ —the model analogue of the empirical moments \hat{m} . We then calculate the vector of moment conditions as $g(\hat{\lambda}|\hat{T}_i, \hat{\theta}) = [\hat{m} - m(\hat{\lambda}|\hat{T}_i, \hat{\theta})]$. We assume independence across the moments so that the covariance between the moment conditions is set to zero. In the final stage of step 2, the diagonal of the inverted variance-covariance matrix of the moment conditions is used as a weight matrix (i.e., $\hat{W} = [\text{diag}\{gg'\}]^{-1}$) to compute the final parameter estimates.²⁸ 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.

Generally, λ is identified by the share of passengers who choose menu tips, that is, passengers who fall within tip bins 20%, 25%, and 30%. 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 tip rates. Thus, the shares of passengers in the 20%, 25%, and 30% tip bins (i.e., the share of passengers who choose menu tips) identifies λ and hence c_i .

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 estimates 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 (3). The standard error is defined as the standard deviation of the distribution of parameter estimates computed from all 1000

²⁸The theory suggests that the best choice of a weight matrix is the inverse of the covariance of the moment conditions.

bootstrap samples.

Estimates of Cognitive Cost c_i : Table 2, Panel C, shows the MDE estimates of c_i . We find that the average cognitive cost of deciding on one’s preferred tip rate and computing the corresponding dollar amount is \$1.26 (10% of the average taxi fare \$12.17) when we use the OLS estimates, and \$1.42 (12% of the average taxi fare) when we use the selection correction model estimates instead.²⁹

5.2.4 Model Performance

Figures 8a and 8b compare the observed data to the model predicted distribution of tips that correspond to using estimates from the OLS and the selection correction model respectively. Both models perform well by mimicking the point masses at all the three default options. Generally, the predicted non-menu tips are distributed similarly to the observed data as well. Visually both approaches work well in predicting the realized data. The χ^2 goodness of fit test supports what we observe in the figure. Between the data and model predictions shown in figure 8a, the statistics from the test are $\chi^2 = 0.17598$ and with a p-value of 0.99 and, for figure 8b, $\chi^2 = 0.22752$ with a p-value of 0.99. Thus, there is no significant difference between the observed tips and the tips predicted by the model.

The average decision cost for tipping obtained from adding the estimates of the norm deviation costs of tipping 5% less than the norm (\$0.48) and the cognitive cost (between \$1.26 and 1.41) is between \$1.68 and \$1.83 (14% to 15% of the average taxi fare \$12.17). This amount is similar to our estimate of \$1.89 from the non-parametric approach.

6 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.

²⁹We estimate that the average cognitive cost is \$1.10 for the model where passengers are not penalized for choosing a menu option greater than their perception of the tipping norm (see footnote 18).

Also, with all the parameters of the model estimated, we can compute the taxi fare thresholds (\bar{F}_i) where each passenger will be indifferent between choosing to compute their preferred tip and choosing a tip menu. That is,

$$\bar{F}_i = \frac{c_i - \hat{\theta} \left[\left(\hat{T}_i - t_i^* \right)^2 - \left(\hat{T}_i - d_j \right)^2 \right]}{d_j - t_i^*}$$

Figure A.12 in the appendix shows the distribution of \bar{F}_i which averages at about \$18 with a median of about \$15.

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 wage³⁰ 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 one person firm (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 tipping in taxicabs, the two reasons stated above are identical since drivers keep all the earnings (taxi fares + tips) from driving.

6.1 Tip-Maximizing Menu of Tips

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 full characterization of the tip-maximizing menus, which may include but not be limited to presenting some combination of dollar tip amounts and percentages.

We proceed by setting the model parameters to estimates from table 2, Panels A and B, and then comparing predictions from the OLS model in column (1) to those from the Heckman selection correction model in column (2). We present the results in Panel C of table (2).

Our procedure is to estimate equation (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. Given that our model parameters were estimated using data from 2014 only, we also conduct this exercise using data from 2010, when the menus of tips were different. Doing this helps us to gauge the sensitivity of our results to using different samples.

To fix ideas, we first consider the case where drivers are restricted to showing 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. We then proceed to search for the two menu tip options that will maximize the tips received by drivers.³¹

Following the procedure in the last paragraph, we continue to increase the number of menu tip options until the average tip stays practically the same upon adding more menu

³⁰This is a base wage below the minimum wage that is paid to employees who receive a substantial portion of their earnings from tips.

³¹There are detailed descriptions of these two cases in section A.8 of the appendix.

options. Figures 9a and 9b plot the average tip across all observed fares as the number of menu tip options increase—for both estimates that correspond to the OLS approach and the Heckman selection correction model respectively. The figure 9a shows that the average tip increases no further than about 19% after showing three or more tip-maximizing menu options (using the OLS approach). We therefore conclude that showing three menu options is tip maximizing. The corresponding predicted tip rates from the OLS approach are 21%, 27%, and 33% (table 2, Panel C, column (1)). Estimates from the Heckman selection model are similar to results from the OLS approach (Figure 9b, and table 2 Panel C column (2)).

It is important to note that the tip-maximizing menu proposed by our model (21%, 27%, and 33%) is very similar to the menu currently offered to passengers (20%, 25%, and 30%). We find similar results when we use data from 2010 to predict the tip-maximizing menu (table 2, Panel C). Another 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.³²

6.2 Evolution of Menu Tip Options

We examine differences in the average tip rate across the various tip menus presented by the two electronic credit card machine vendors in NYC Yellow taxicabs between 2010 and 2014. We then assess which of the menus induced passengers to tip more and how they compare to the model suggested tip-maximizing menu.

For each period where different sets of tip menus were presented to passengers, we use our model to predict an analogous set of tips where the tip menu is set to the model tip-maximizing menu (21%, 27%, and 33%). The model predictions are made using data from the relevant period.

We consider three main periods in this analysis (period 1: January 2010 - January 2011, period 2: February 2011-December 2011, and period 3: 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 figure A.2.

For this analysis, Table 3 presents three panels (A, B, and C) that correspond to the three periods being considered. Each panel has three columns that report the average tip rate for CMT rides (column (1)), VTS rides (column (2)), and model prediction (column (3)). Panel A corresponds to the first period where CMT presented 15%, 20%, and 30%, while VTS presented a different set of tips for fares under \$15 than for higher fares. During

³² See section A.8 in the appendix for details.

the first period, in VTS cabs, passengers saw three options in dollars for fares under \$15 (\$2, \$3, and \$4), and three options in percentages for fares \$15 and above (20%, 25%, and 30%). Panel A shows that, on average, passengers tip at higher rates in VTS cabs (20.68%) relative to CMT cabs (17.81%). Our model predicts that the tip-maximizing menu would have resulted in an average tip rate of 19.56%. Compared to both VTS and the model, CMT used an inferior menu.

In period 2, 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 both VTS and the model prediction.

In period (3), 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 rate in VTS cabs dropped by two percentage points (from 20.66% to 18.55%). The average tip rate remained about the same as in period 2 for both CMT and the model prediction in this period.

With respect to presenting passengers with a menu that show percentages as options, the similarity across the data and model in the third period suggests that taxicabs are currently showing the tip-maximizing menu. The evolution over time of the menu of fares is consistent with taxi companies learning to converge to a menu that maximizes tips.

6.3 Welfare

When firms present customers with menus, it has implications for their profits and the utility of consumer who choose from the menu. Figure A.13 shows that the choice of menu options may either increase or decrease tips—above or below the tips that passengers would give without a menu of tips. 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 (21%, 27%, and 33%) to tippers. 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 (3)) is quasi-linear in money. Thus, the social welfare from tipping is the sum of the utility that consumers get from tipping and the tip revenues

that drivers receive. We therefore use the structural estimates to compute the dollar value of a passenger’s utility from tipping. From equation (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. However, 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 4, reports welfare calculations averaged 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 -\$3.47 (29% of the average taxi fare of \$12.17), and average tip received by drivers is \$2.05 (16.8% of the average taxi fare). Therefore, the average welfare (loss) from tipping on a taxi trip without a tip menu is $-\$3.47 + \$2.05 = -\$1.41$ (12% 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 other tip menus mentioned above. Table 4 Panel B shows the results. Relative to the no-menu case, the previous tip menu (15%, 20%, 25%) decreased the utility loss from tipping by \$1.06 but had no effect on tip revenue. Therefore, welfare increased by \$1.06 (a 75% increase) per taxi trip on average.³³ The current tip menu in taxis (20%, 25%, and 30%) increases both tip revenue and the utility from tipping relative to the no-menu case. In particular, the utility loss from tipping decreases by \$0.78 (a 23% decrease), and the revenue from tips rises by \$0.25 (a 12% increase). In sum, welfare increases by \$1.03 (a 73% increase) under the current menu relative to the no-menu case. The social welfare under both the current and previous menus are very similar. However, passengers transfer about \$0.25 of their gains from tipping under the previous menu over to drivers under the current menu. Showing the tip-maximizing menu yields similar results to 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

³³The estimated tipping behavior under the previous menu is an out-of-sample test of our model. Figure A.15a shows the model predicted distribution of tip rates under the previous menu. Our model over predicts the share of passengers who choose the lowest menu option (15%) and under predicts the share that choose the middle option (20%). However, our model predicts the same average tip as found in data from tips under the previous menu.

same procedure in section 6.1 and find that showing 10%, 16%, and 25% is utility maximizing. Showing this menu enhances the social welfare from tipping the most. Interestingly, under this menu, there is no change in the rate at which passengers tip compared to the no-menu case or the previous menu. However, the utility loss from tipping decreases by 33% (a gain of \$1.14 for passengers on average). Under this menu, welfare increases by \$1.15 (a 79% increase) relative to the no-menu case.³⁴

The 2014 Taxi fact book reports that there are about 175 million taxi rides annually. To put the trip level welfare estimates in perspective, we can rescale all the estimates in Table 4, by multiplying all the estimates by 175 million. The main takeaway is that, relative to the no-menu case, the current taxi menu increases the welfare from tipping by about \$180.25 million annually.³⁵

7 Conclusion

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 tipping 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.

We present both a nonparametric and a structural analysis of tipping behavior. These two analyses provide consistent results. Our nonparametric estimate of the average decision cost—the combination of the norm deviation and cognitive costs—is about \$1.89 (16% of the average taxi fare of \$12.17). To disentangle the norm deviation and cognitive costs, we add structure to our model of decision-making. We estimate the distribution of passenger beliefs about the social norm tip and it averages at 19.8% of the taxi fare, which almost exactly equals 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

³⁴In fact, forcing people is 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 = -166.7 \times (0.20 - 0)^2 = -\6.67 . This is five times worse than the no-menu case.

³⁵Figure A.15 in the appendix shows the distribution of tips under the different scenarios discussed above.

a norm deviation cost of \$0.42 (3.5% of the average taxi fare). The estimated cognitive cost of calculating a non-menu tip ranges from \$1.26 to \$1.41 (10% to 12% of the average taxi fare) on average.

We use the structural model to investigate a number of what-if questions. For example, our simulations suggest that, the current menu increases the amount of tips received by 12.4% compared to not showing a menu. The simulations also show that the current number and level of menu options in NYC Yellow taxicabs nearly maximizes tips. Thus, 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. In addition, our findings suggest that the current menu increases the social welfare from tipping by \$180.25 million each year relative to not presenting a tip 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 business and policy makers.

References

- 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. [4](#)
- Azar, O. H. (2004). What sustains social norms and how they evolve?: The case of tipping. *Journal of Economic Behavior & Organization*, 54(1):49–64. [13](#)
- Azar, O. H. (2005). The social norm of tipping: Does it improve social welfare? *Journal of Economics/ Zeitschrift fur Nationalokonomie*, 85(2):141–173. [5](#)
- Azar, O. H. (2007). The Social Norm of Tipping: A Review. *Journal of Applied Social Psychology*, 37(2):380–402. [8](#)
- Azar, O. H. (2010). Do people tip because of psychological or strategic motivations? An empirical analysis of restaurant tipping. [8](#)
- Beggs, A. and Klemperer, P. (1992). Multi-Period Competition with Switching Costs. *Econometrica*, 60(3):651–666. [5](#)
- Benzarti, Y. (2017). How Taxing Is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs. Technical report, NBER. [10](#)
- 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. [1](#), [4](#), [6](#)
- 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. [4](#)
- Bicchieri, C. and Dimant, E. (2019). Nudging with care: the risks and benefits of social information. *Public Choice*. [5](#)
- Blumenstock, J., Callen, M., and Ghani, T. (2018). Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan. *American Economic Review*, 108(10):2868–2901. [4](#)
- 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. [4](#)
- Carr, A. (2013). How Square Register’s UI Guilts You Into Leaving Tips. *Fast Company*. [4](#)
- 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. [4](#), [6](#)
- 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. [4](#)
- 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. [4](#)

- Cohan, P. (2012). Will Square’s Starbucks Deal Spark the End of Cash? *Forbes Magazine*. 3
- DellaVigna, S. (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature*, 47(2):315–372. 4
- DellaVigna, S. and Malmendier, U. (2004). Contract Design and Self-Control: Theory and Evidence. *The Quarterly Journal of Economics*, 119(2):353–402. 5
- DellaVigna, S. and Malmendier, U. (2006). Paying Not to Go to the Gym. *American Economic Review*, 96(3):694–719. 4
- 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. 4
- Goldin, J. and Reck, D. Revealed Preference Analysis with Framing Effects *. *Journal of Political Economy*. 6
- Grynbaum, M. (2009). In New York, Taxi Revenue and Tips From Credit Cards Rise. *The New York Times*. 3, 4
- Haggag, K. and Paci, G. (2014). Default Tips. *American Economic Journal: Applied Economics*, 6(3):1–19. 4, 5, 7
- Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–2682. 4
- 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. 4
- 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. 4
- Johnson, E. J. and Goldstein, D. (2003). Do Defaults Save Lives? *Source: Science, New Series*, 302(5649):1338–1339. 4
- 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. 17
- Madrian, B. C. and Shea, D. F. (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *Source: The Quarterly Journal of Economics*, 116(4):1149–1187. 4
- O’Donoghue, T. and Rabin, M. (1999). Doing It Now or Later. *American Economic Review*, 89(1):103–124. 4
- O’Donoghue, T. and Rabin, M. (2001). Choice and Procrastination. *Source: The Quarterly Journal of Economics*, 116(1):121–160. 4
- Ravindranath, M. (2014). Whole Foods to use Square for store checkout. *The Washington Post*. 3
- 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*. 3

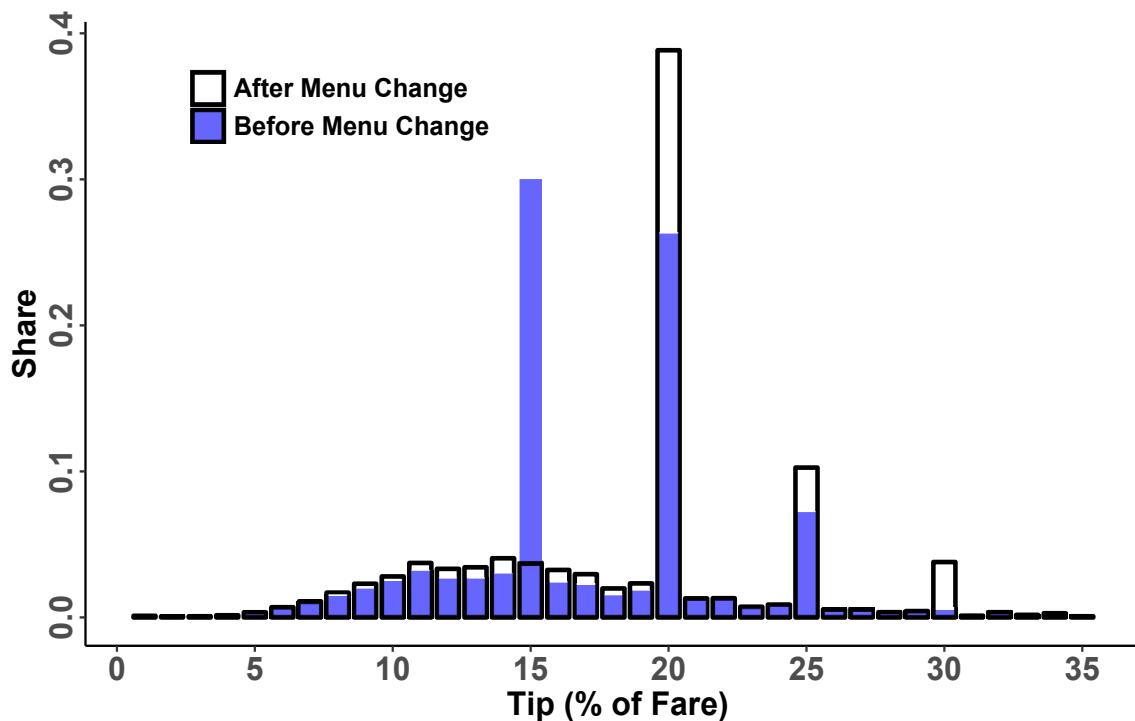
Stiles, R., Siegel, L., Garber, J., Neger, H., and Ullah, A. (2014). 2014 Taxicab Fact Book. Technical report, NYC Taxi & Limousine Commission.

Thakral, N. and Tô, L. T. (2019). Tipping and the Dynamics of Social Norms. *Mimeo.* 5

Thaler, R. H. and Sunstein, C. R. (2003). Libertarian Paternalism. *Source: The American Economic Review*, 93(2):175–179. 4

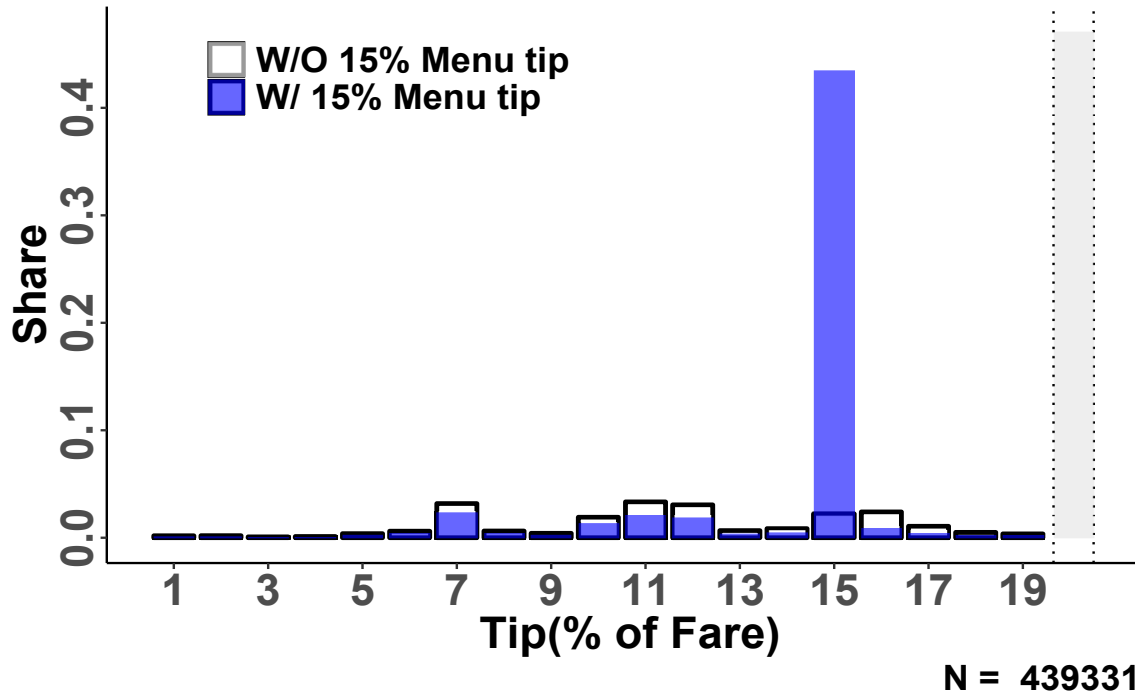
Figures

Figure 1: Distribution of Tip (%) Before and After CMT Tip Menu Change



Notes: This figure shows the distribution of tips before and after CMT—a New York City Yellow taxi electronic credit card machine vendor—changed the menu of tips that is presented to taxi passengers in 2011. CMT presented customers with three tip options in percentages (15%, 20%, and 25%) before the menu change. Afterward, 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 menu change. The bars in this figure 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 data used for this figure are from 2010 and 2011 standard rate taxi fares paid for via a CMT credit card machine along with a positive tip.

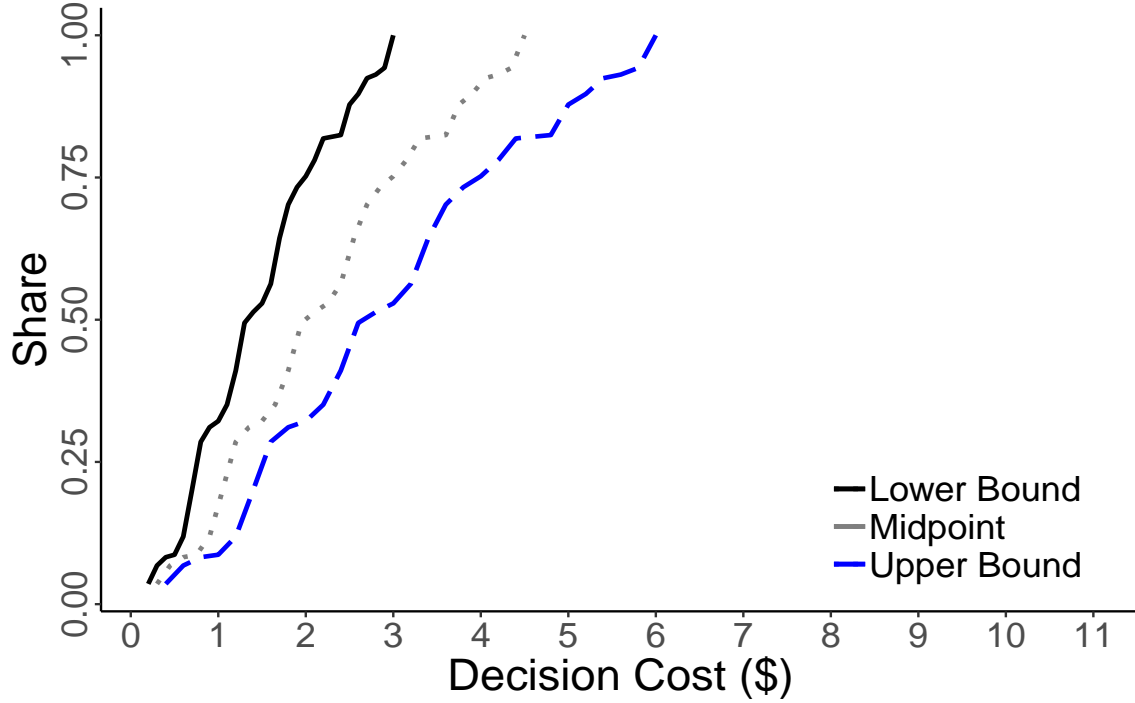
Figure 2: Distribution of Tip (%) for Fares Between \$21 and \$23, Before and After CMT Tip Menu Change



Notes: Figures 4a shows the distribution of positive tips truncated at the tip rate 19.5% for the subset of taxi trips whose fare falls within the range of (\$21, \$23]. The figures show the distribution of tips 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. CMT presented customers with three tip options in percentages (15%, 20%, and 25%) before the menu change. Afterward, 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 paid via a CMT credit card machine are used in this figure.


Figure 3: Bounds on the CDF of Decision Costs

Tip Rate = 10%



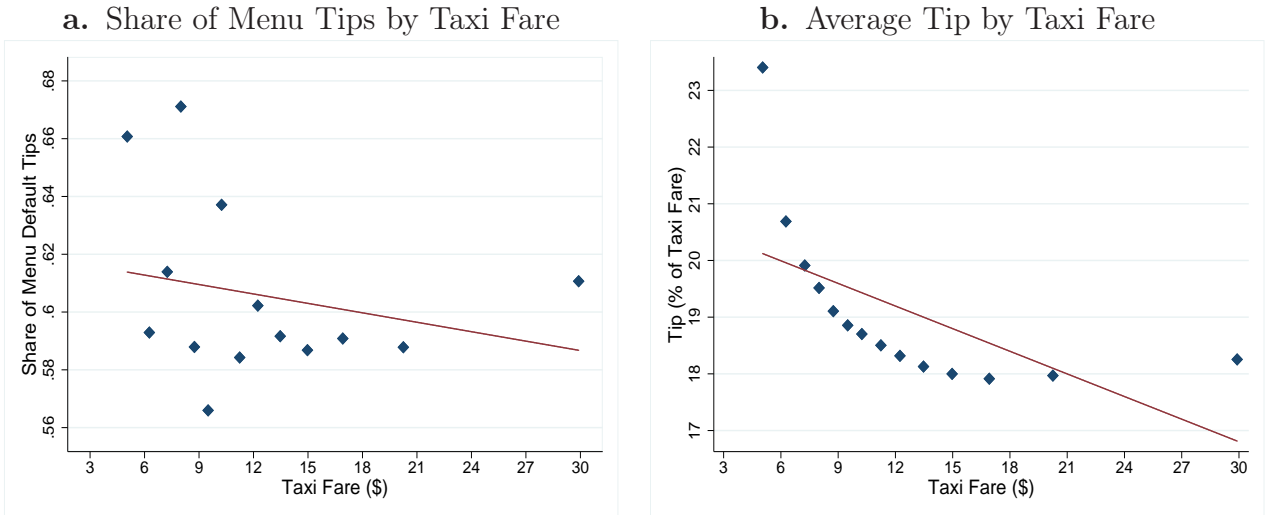
Notes: Figure 5 shows the lower and upper bounds for the CDF of decision costs computed for passengers whose tip rates fall within the range of (9.5%, 10.5%]. The computation of these bounds relies on CMT's (a New York City taxi credit card machine vendor) change in the menu of tips that is presented to taxi passengers in 2011. CMT presented customers with three tip options in percentages (15%, 20%, and 25%) before the menu change. Afterward, CMT removed the lowest tip option (15%) and added a higher percentage option (30%), so that it offered 20%, 25%, and 30%. These bounds are computed using the increase in the share of passengers who tip at the 2% and 10% respectively 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_i is given by $[|0.15 - t|F, |0.20 - t|F]$. Data from 2010 and 2011 standard rate taxi trips paid via a CMT credit card machine along with a positive non-menu tip that is not round-number dollar amount.

Figure 4: Unconditional CDF of Decision Costs



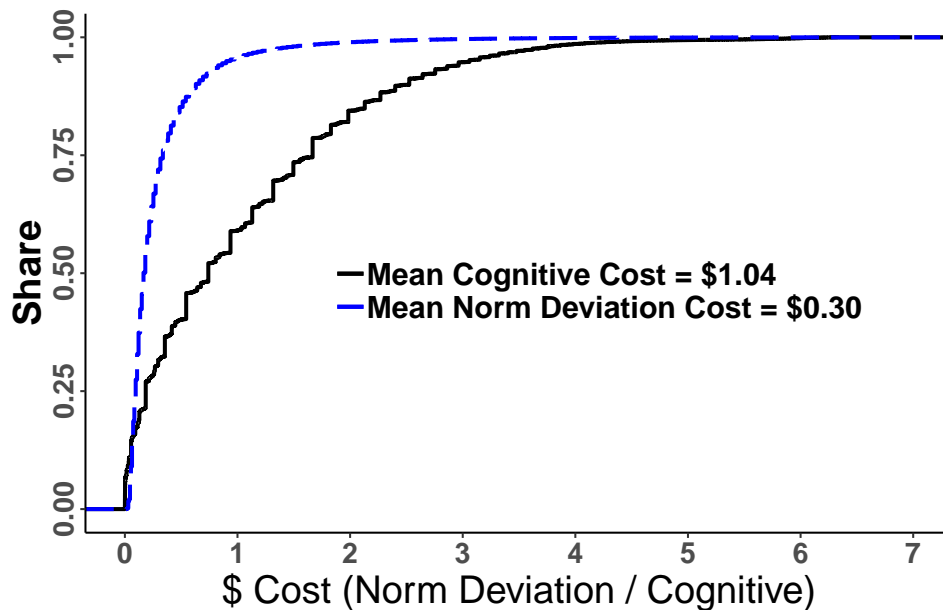
Notes: This figure shows the distribution of decision costs. In this figure, we assume that the midpoints of the estimated bounds of decision costs across all tip rates are similar to the true decision costs. We construct the CDF using information from the estimated conditional CDFs of decision costs (section 4). Data from 2010 and 2011 standard rate taxi trips paid for via a CMT credit card machine along with a positive non-menu tip that is not round-number dollar amount.

Figure 5



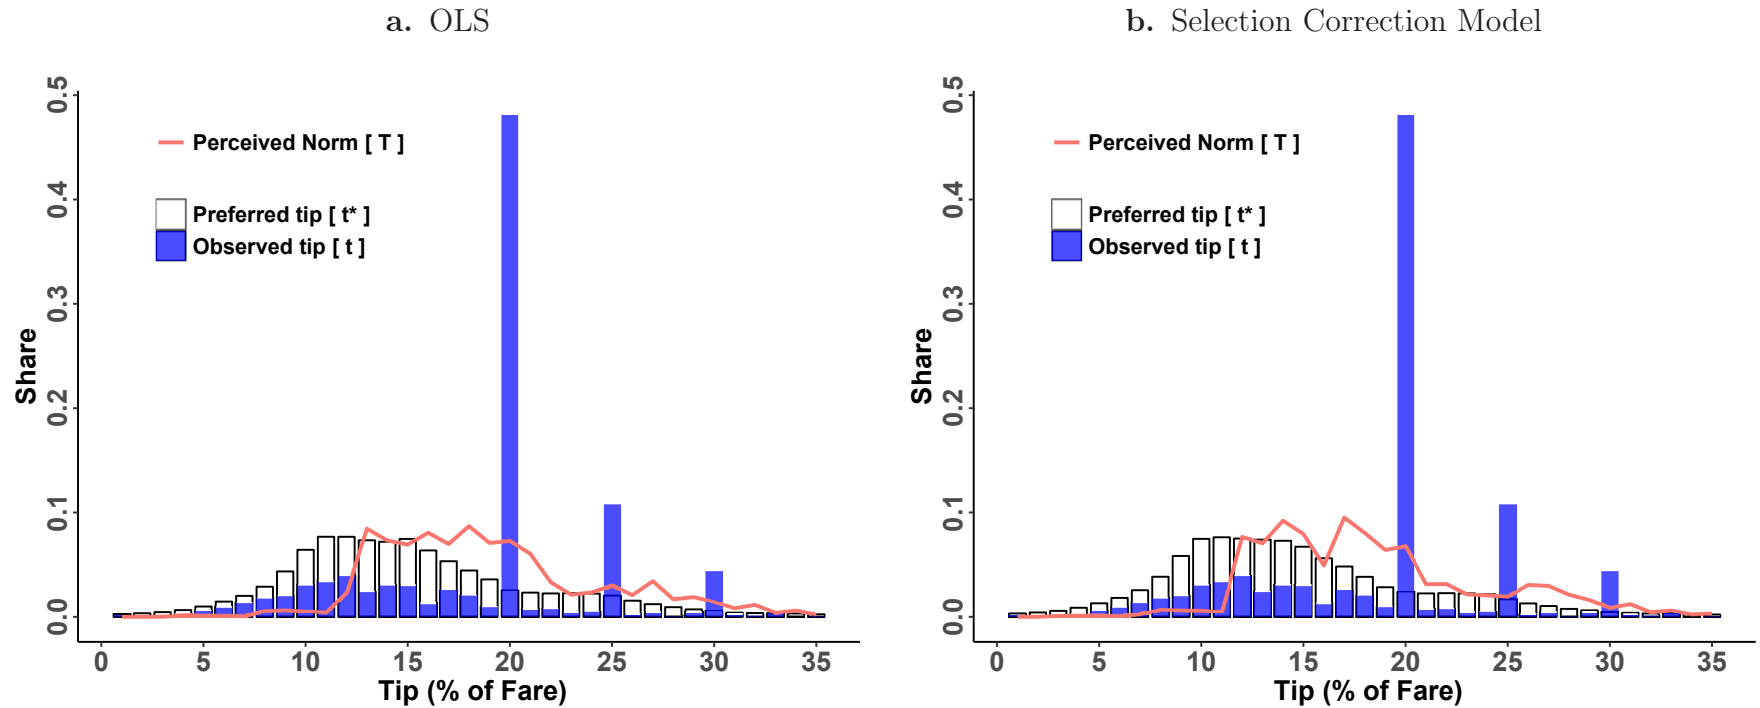
Notes: Figures 5a is a binned scatter plot that illustrates the relationship between the share of passengers who choose any one of the suggested menu tips presented at the end of a taxi ride at different levels of the taxi fare. Figures 5b is a binned scatter plot that illustrates the average tip rate at different levels of the taxi fare. The data used in both figures are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive tip amount.

Figure 6: CDF of The Norm Deviation Cost and The Upper Bound of Cognitive Costs for Non-Menu Tips



Note: 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. Following equation (5) and using the estimate of θ (166.7), we compute the upper bound of cognitive cost for non-menu tips as $c_i = (d_j - t_i)F_i - \hat{\theta} \left[\left(\hat{T}_i - t_i \right)^2 - \left(\hat{T}_i - d_j \right)^2 \right]$. We compute the norm deviation costs as $\hat{\theta} \left(\hat{T}_i - t_i \right)^2$. The data are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive non-menu tip that is not round-number dollar amount.

Figure 7: Estimated Distribution of the Tipping Norm T_i

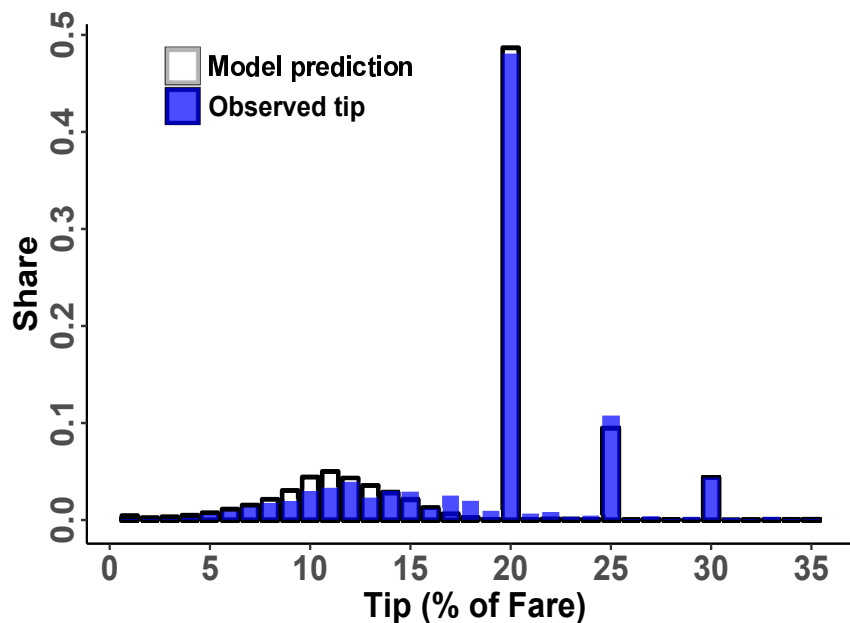


35

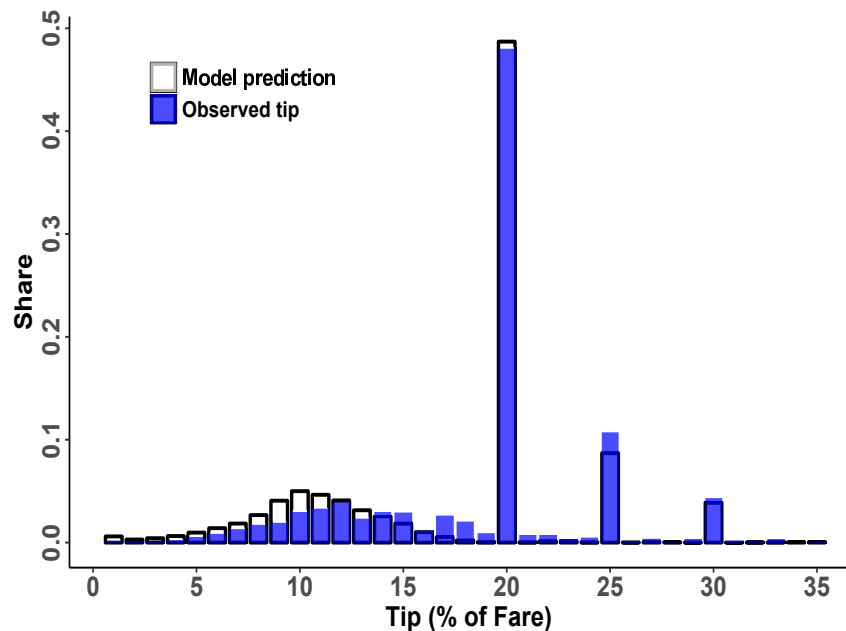
Notes: Figures 7a and 7b depict the empirical estimates—via our structural model—of the distribution of the perceived tipping norm T_i and the implied distribution of the preferred tip t_i^* within the population of the taxi passengers. These distributions are compared to the observed tips (shaded bars). Figure 7a shows estimates of T_i recovered from an OLS regression of the observed non-menu tip rates regressed on a constant term and the taxi fare. Figure 7b accounts for possible sample selection and shows estimates of T_i recovered from a 2-step Heckman selection correction model. The data used are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive non-menu tip. The bars in this figure 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.

Figure 8: Model Fit

a. OLS (Distribution of Tips)

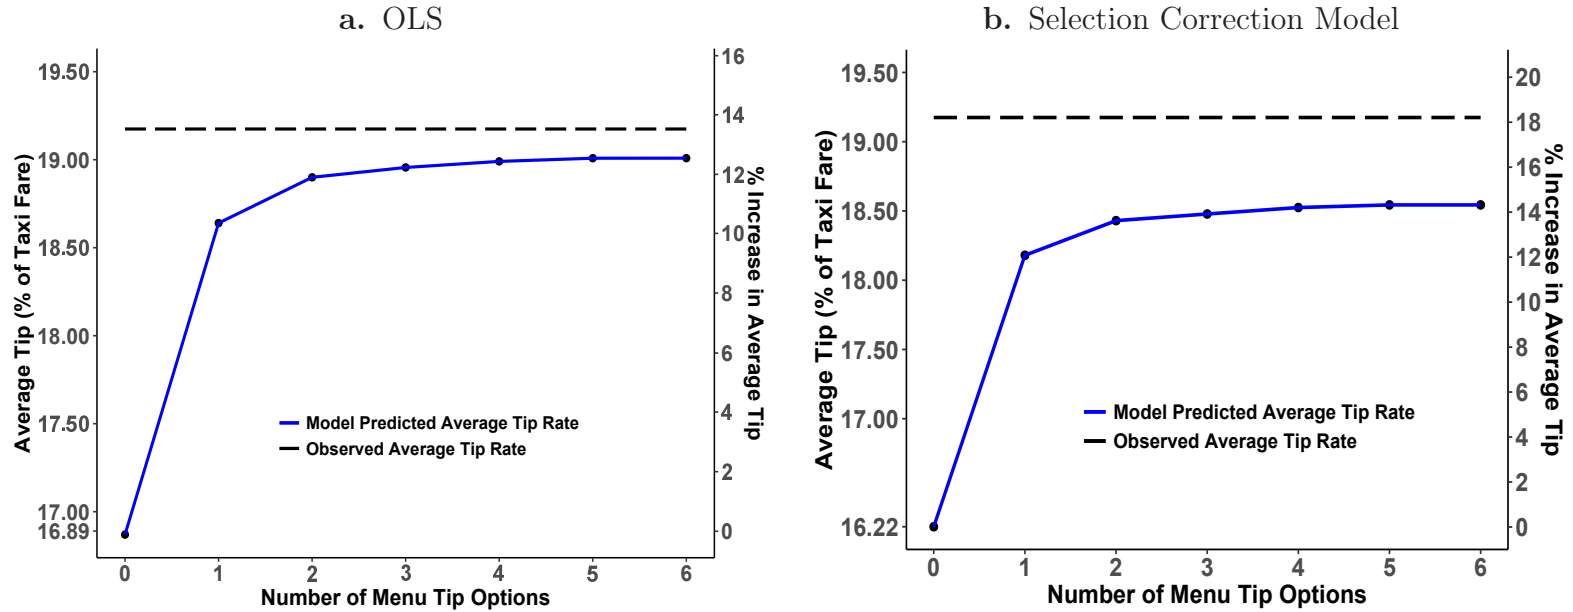


b. Selection Correction Model (Distribution of Tips)



Notes: These figures illustrate how our structural model fits the observed data by depicting the distribution of the observed tips versus the distribution of tips predicted by our model. Figure 8a shows the fit of the model when model estimates are computed using the perceived tipping norm T_i and θ (the marginal disutility for deviation from the norm) recovered from an OLS regression. The statistics from a chi-square goodness of fit test for figure 8a are $\chi^2 = 0.17598$ and a P-value of 0.99. Figure 8b is analogous to figure 8a but shows the model fit when we use the estimate of T_i and θ from the 2-step Heckman selection correction model. The statistics from a chi-square goodness of fit test for figure 8b are $\chi^2 = 0.22752$ and a P-value of 0.99. The data used are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive non-menu tip. The bars in this figure 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.

Figure 9: Average Tip by Number of Tip-Maximizing Menu Tip Options



Notes: Given a menu that presents customers with percentage tip options, Figures 9a and 9b plot the average tip rate predicted by the model for showing the tip-maximizing menu as the number of menu options increases. 9a corresponds to the case where tipping norm T_i and θ (the marginal disutility for deviation from the norm) are recovered from an OLS regression, and 9b corresponds to the case where tipping norm T_i and θ are recovered from a 2-stage Heckman selection correction model. The data used are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive tip.

Tables

Table 1: CMT Taxi Trip Characteristics, Mean (Standard Deviation)

	Before Tip Menu Change		After Tip Menu Change		
	Jan 2010 - Jan 2011		Feb 2011 - Dec 2011		2014
	(1)		(2)		(3)
Menu of Tips	[15%, 20%, 25%]		[20%, 25%, 30%]		
All 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 (%)	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 reports the summary statistics of standard rate NYC Yellow taxi trips paid via a CMT credit card machine along with a positive non-menu tip. It shows the differences in trip characteristics before and after CMT changed the menu of tips that is presented to taxi passengers. Column (1) presents trip characteristics one year before the 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.

Table 2: Structural Model Estimates (Panels A & B) and Counterfactual Exercises (Panel C)

	OLS	Selection Correction
	(1)	(2)
Panel A:		
<i>Estimates of α_T and θ</i>		
	Dependent Variable: Tip Rate	
Taxi Fare ($\hat{\beta} = \frac{1}{2\hat{\theta}}$)	−0.0033***(0.00001)	−0.0033***(0.00001)
Constant ($\hat{\alpha}_T$)	0.2046***(0.0001)	0.1980***(0.0001)
$\hat{\theta}$ ($= \frac{1}{2\hat{\beta}}$)	166.667	166.667
Inverse Mills Ratio		0.0048***(0.0003)
ρ		0.028
1 st -Stage Added Instrument		# of Passengers
1(Round Number Tip)	YES	YES
N. obs.	16,394,858	16,394,858
Panel B:		
<i>Minimum Distance Estimates of c_i</i>		
Ave. Cognitive Cost (\$) $c_i = 1/\lambda$	1.26***(0.000)	1.41***(0.000)
N. obs.	41,620,454	41,620,454
Panel C:		
<i>Counterfactual Exercise</i>		
2014 Data	Tip-Maximizing Menu	[21%, 27%, 33%]
	N.obs.	41,620,454
2010 Data	Tip-Maximizing Menu	[21%, 27%, 32%]
	N. obs.	25,091,514

Notes: In this table, panels A and B report estimates from the structural model of passenger tipping behavior—discussed in section 5. Panel C reports counterfactual estimates of the tip-maximizing menu using trips from both 2010 and 2014.—discussed in Section 6.1. Only standard rate taxi trips paid via a CMT credit card machine along with a positive tip are used in this table. Column (1) reports estimates that correspond to recovering the tip norm T_i and θ (the marginal disutility for deviation from the norm) using the subsample of trips with non-menu tips in an OLS regression that does not account for sample selection bias. Column (2) is analogous to column (1) but reports estimates of the T_i and θ from the second step of the 2-step Heckman selection correction model. Panel A uses the subsample of trips with non-menu tips from trips in 2014 and reports reduced-form estimates of the average tip norm T_i and θ for the population of taxi passengers. Panel B reports estimates using the sample of all taxi tips in a minimum distance estimation of the distribution of cognitive costs incurred by passengers who opt to compute their preferred non-menu tip. *p<0.1, **p<0.05, ***p<0.001.

Table 3: Evolution of the Menu Tip Options

	CMT	VTS	Model Prediction
	(1)	(2)	(3)
<i>Panel A: Trips from Jan 2010 - Jan 2011</i>			
Tip menu for taxi fare < \$15	[15%, 20%, 25%]	[\$2, \$3, \$4]	[21%, 27%, 33%]
Tip menu for taxi fare \geq \$15		[20%, 25%, 30%]	
Average tip for all fares	17.81%	20.68%	19.56%
Average tip for fare < \$15	18.00%	21.19%	20.09
Average tip for fare \geq \$15	17.51%	16.61%	16.25
N. obs.	27,574,410	28,658,477	56,232,887
<i>Panel B: Trips from Feb 2011 - Dec 2011</i>			
Tip menu for taxi fare < \$15	[20%, 25%, 30%]	[\$2, \$3, \$4]	[21%, 27%, 33%]
Tip menu for taxi fare \geq \$15		[20%, 25%, 30%]	
Average tip for all fares	19.16%	20.66%	19.50%
Average tip for fare < \$15	19.37%	21.21%	20.05%
Average tip for fare \geq \$15	18.00%	17.66%	16.39%
N. obs.	31,960,044	30,339,659	62,299,703
<i>Panel C: Trips from 2013 - 2014</i>			
Tip menu for all taxi fares	[20%, 25%, 30%]	[20%, 25%, 30%]	[21%, 27%, 33%]
Average tip for all fares	19.07%	18.55%	19.50%
Average tip for fare < \$15	19.42%	18.74%	19.91%
Average tip for fare \geq \$15	17.96%	17.93%	16.01%
N. obs.	83,107,354	84,332,924	167,440,278

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. Column (3) shows the average tip rate predicted by our model if the tip-maximizing menu were presented to passengers. Each panel also reports the average tip rate separately for trips where the taxi fare is less than \$15. Only standard rate taxi trips paid for via a credit card machine where passengers leave a positive tip are used in this table.

Table 4: Trip Level Welfare Estimates by Type of Tip Menu

	Utility (Loss) from Tipping	Tip Revenue	Welfare
	Consumer Surplus (CS)	Producer Surplus (PS)	CS + PS
	(1)	(2)	(3)
<i>Panel A: Baseline</i>			
No tip menu	-\$3.47	\$2.05	-\$1.41
<i>Panel B: Change Relative to No Tip Menu</i>			
Previous tip menu [15%, 20%, 25%]	\$1.06	\$0	\$1.06
Current tip menu [20%, 25%, 30%]	\$0.78	\$0.25	\$1.03
Tip-maximizing menu [21%, 27%, 33%]	\$0.73	\$0.26	\$0.99
Utility-maximizing menu [10%, 16%, 25%]	\$1.14	\$0.01	\$1.15

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 structural 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. From equation (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). For welfare estimates at the level of yearly-trips rescale the all the estimates by multiplying by 175 million (total number of taxi trips in 2014).

Appendix


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 in the menu of tips presented to passengers between 2009 and the menus presented by the two main NYC Yellow taxi credit card machine providers (CMT and VTS). The figure also shows whether the provider presented different tip menus for different levels of the fare.

A.2 Tips Before and After CMT Tip Menu Change

Figure A.3: Trend of Tips from VTS and CMT Taxicabs Before and After CMT Tip Menu Change



Notes: This figure shows the evolution of average tip rate from 2010 to 2011 for the two vendors. The tip rate is always higher in VTS cabs across time. While the tip rate in VTS cabs is higher than the CMT trend, both trends progress in a very similar manner in 2010. There is an apparent jump of about 1.5 percentage point in the CMT trend right at the beginning of 2011 when CMT makes the changes in their menu options. After the jump, both the VTS and the CMT trends are relatively flat and continue to trend similarly. There is no apparent change in how the trends for both vendors are evolving.

Table A.1: Trip Characteristics

Tip Bin	Share Before (%)	Share After (%)	Difference (%)
[0%, 0.5%]	2.01	2.02	-0.01
(0.5%, 10%]	6.74	7.67	-0.93
(10%, 14.5%]	12.21	15.31	-3.10
(14.5%, 15.5%]	28.08	3.62	24.46
(15.5%, 19.5%]	8.83	10.23	-1.39
(19.5%, 20.5%]	24.58	36.31	-11.73
(20.5%, 24.5%]	5.23	6.03	-0.80
(24.5%, 25.5%]	6.78	9.78	-2.99
(25.5%, 29.5%]	2.04	2.25	-0.21
(29.5%, 30.5%]	0.60	3.77	-3.17
>30.5%	2.89	3.01	-0.12

Note: This table shows changes in distribution of tips before and after one of the NYC Yellow taxi electronic credit card machine vendors (CMT) changed its menu of tip options that is presented to taxi passengers. Until 2010, CMT presented customers with 15%, 20%, and 25% as tip suggestions and later changed the menu to show 20%, 25%, and 30%. The data in this table is from 2010 and 2011 standard rate taxi trips paid for via a CMT credit card machine along with a positive tip.

A.2.1 Tips Before and After CMT Tip Menu Change by Fare

Figure A.4: Distribution of Tip (%) less than 20% Before and After CMT Tip Menu Change

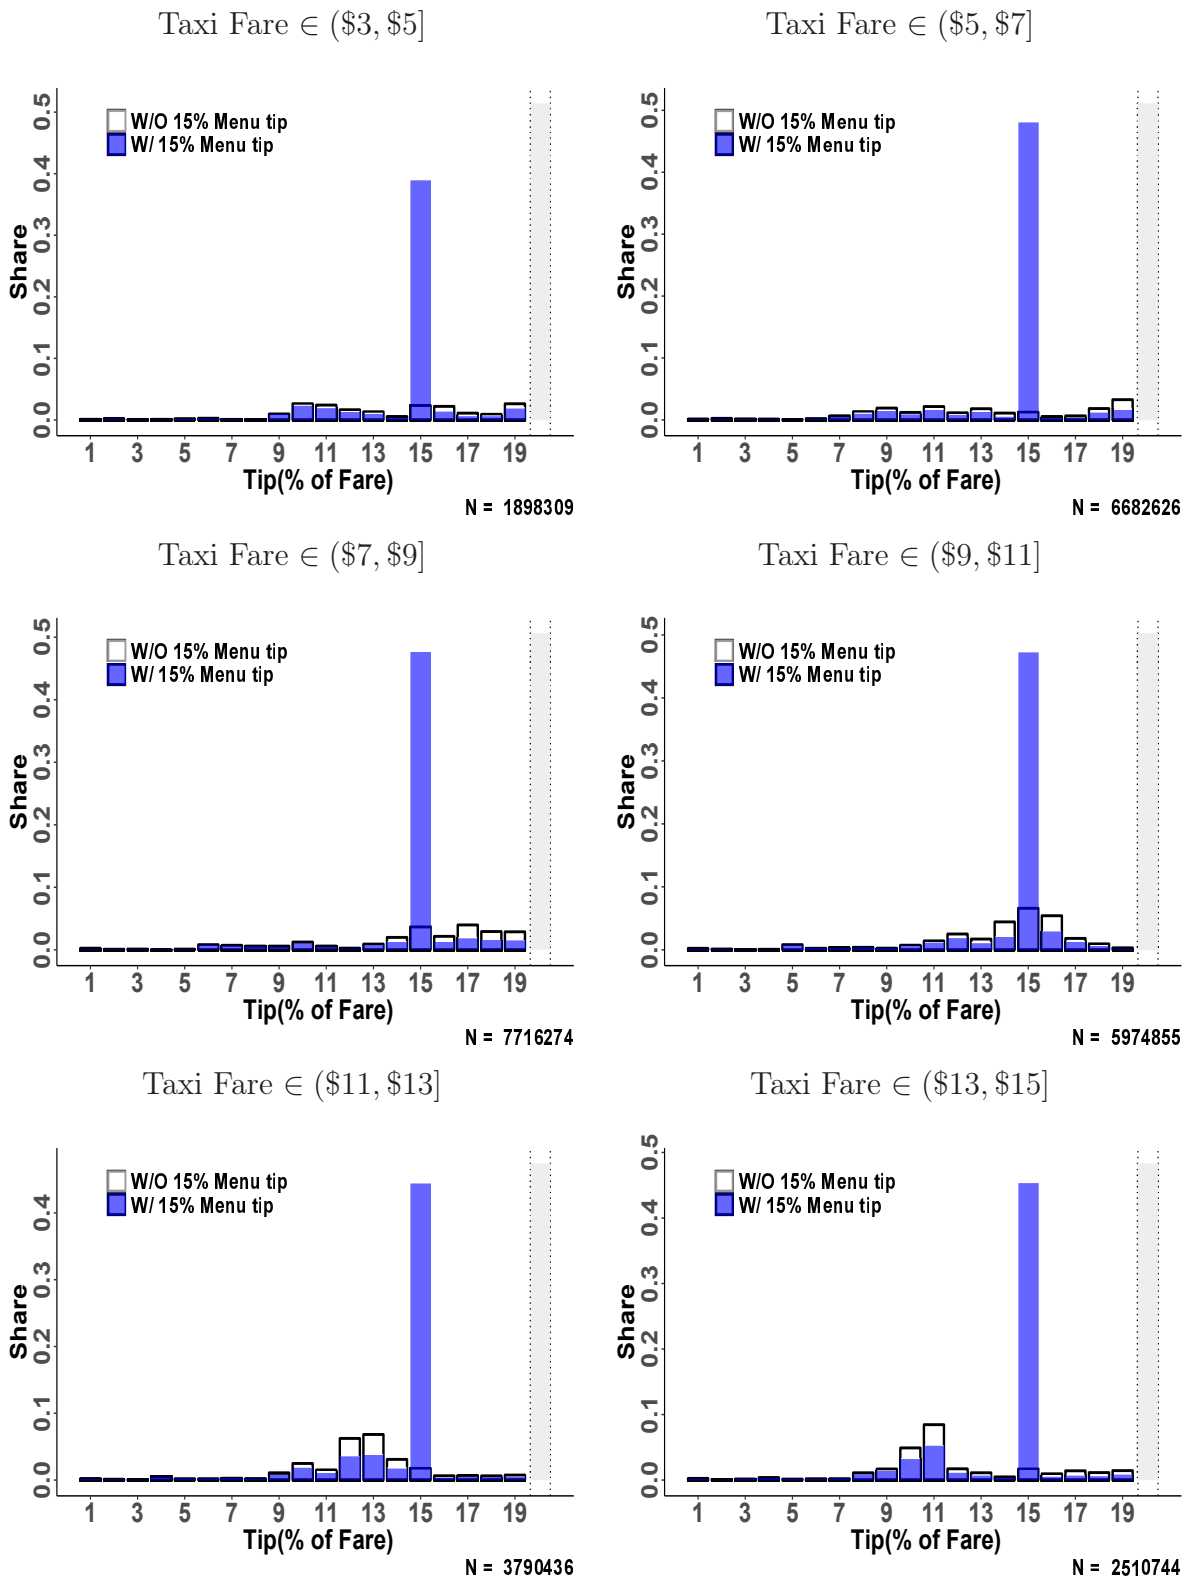




Figure A.4 continued


Taxi Fare \in (\$15, \$17]



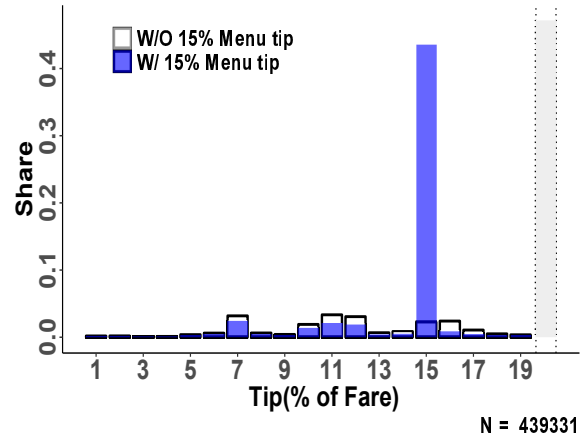
Taxi Fare \in (\$17, \$19]




Taxi Fare \in (\$19, \$21]



Taxi Fare \in (\$21, \$23]



Taxi Fare \in (\$23, \$25]



Taxi Fare \in (\$25, \$27]






Figure A.4 continued


Taxi Fare \in (\$27, \$29]




Taxi Fare \in (\$29, \$31]




Taxi Fare \in (\$31, \$33]



Taxi Fare \in (\$33, \$35]



Taxi Fare \in (\$35, \$37]



Taxi Fare \in (\$37, \$39]

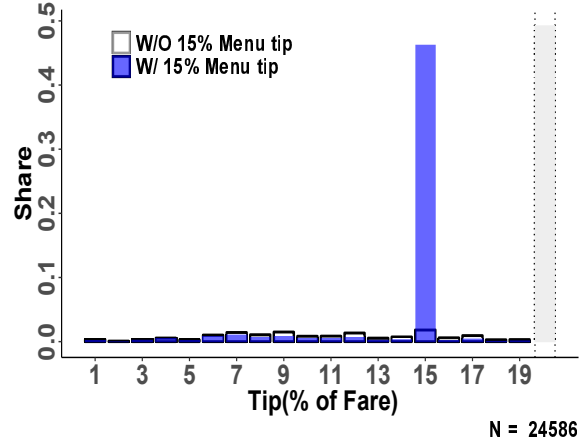




Figure A.4 continued


Taxi Fare \in (\$39, \$41]




Taxi Fare \in (\$41, \$43]




Taxi Fare \in (\$43, \$45]



Taxi Fare \in (\$45, \$47]



Taxi Fare \in (\$47, \$49]



Taxi Fare \in (\$49, \$51]


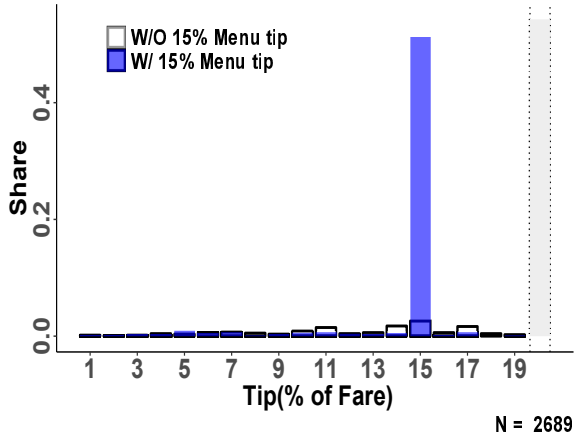



Figure A.4 continued


Taxi Fare \in (\$51, \$53]




Taxi Fare \in (\$53, \$55]



Taxi Fare \in (\$55, \$57]



Taxi Fare \in (\$57, \$59]



Taxi Fare \in (\$59, \$61]




Figure A.4 above 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 paid via a CMT credit card machine are used in these figures.

A.3 Bounds On the Conditional CDFs of Decision Costs

Figure A.5 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 paid 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.5: Conditional CDFs of Decision Costs

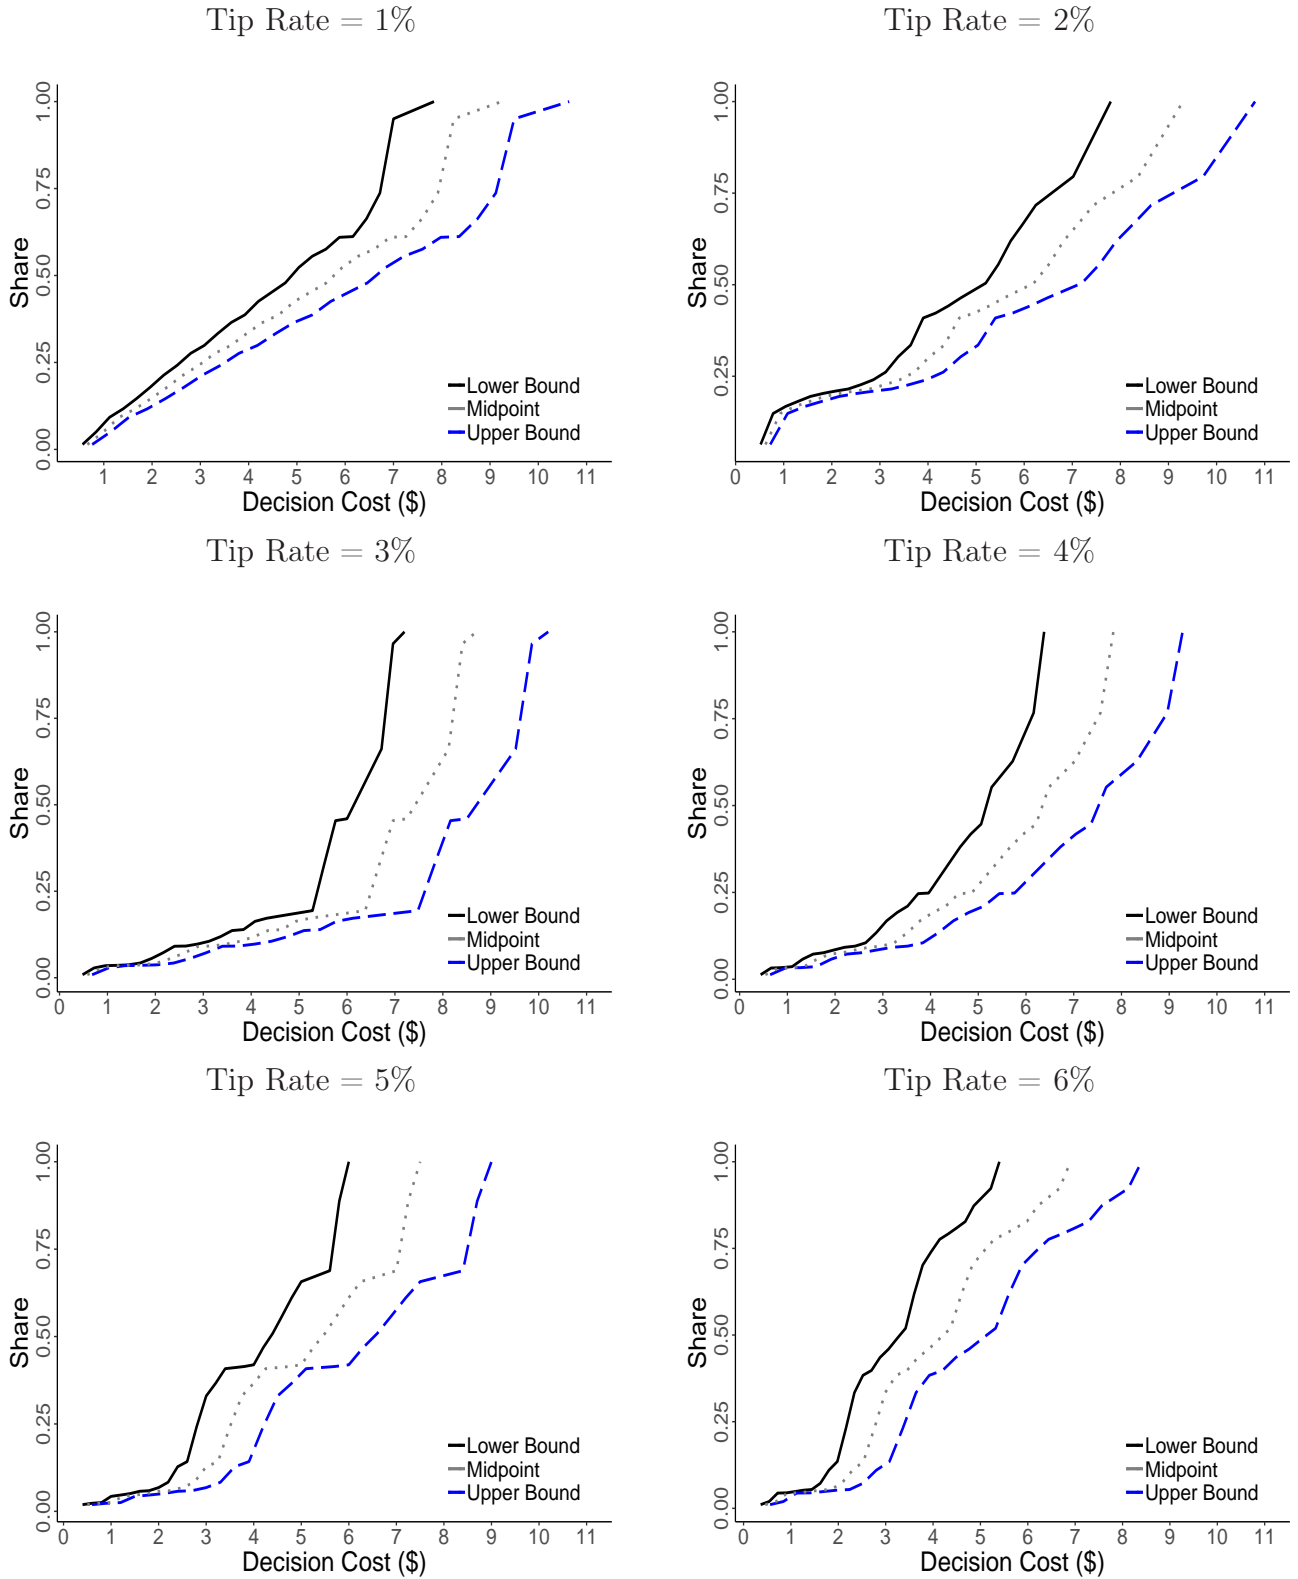


Figure A.5 continued

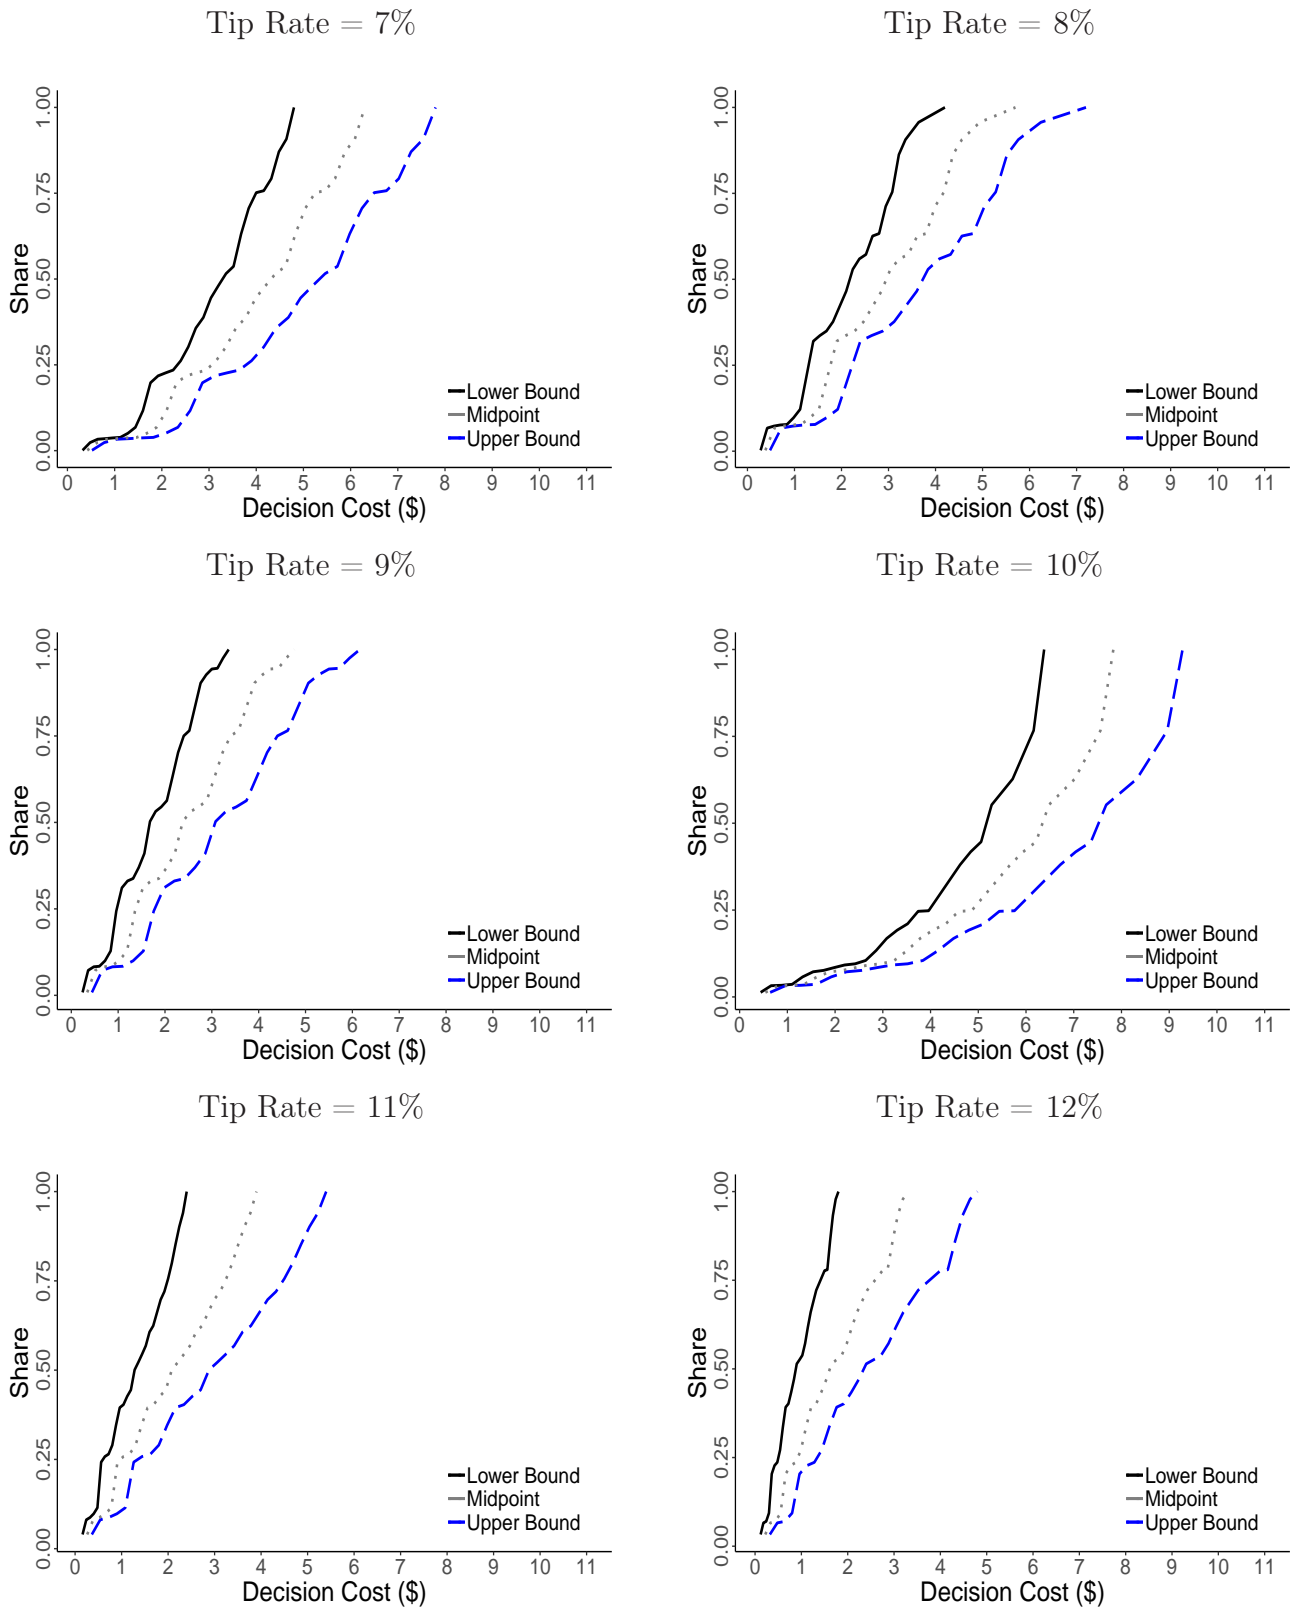
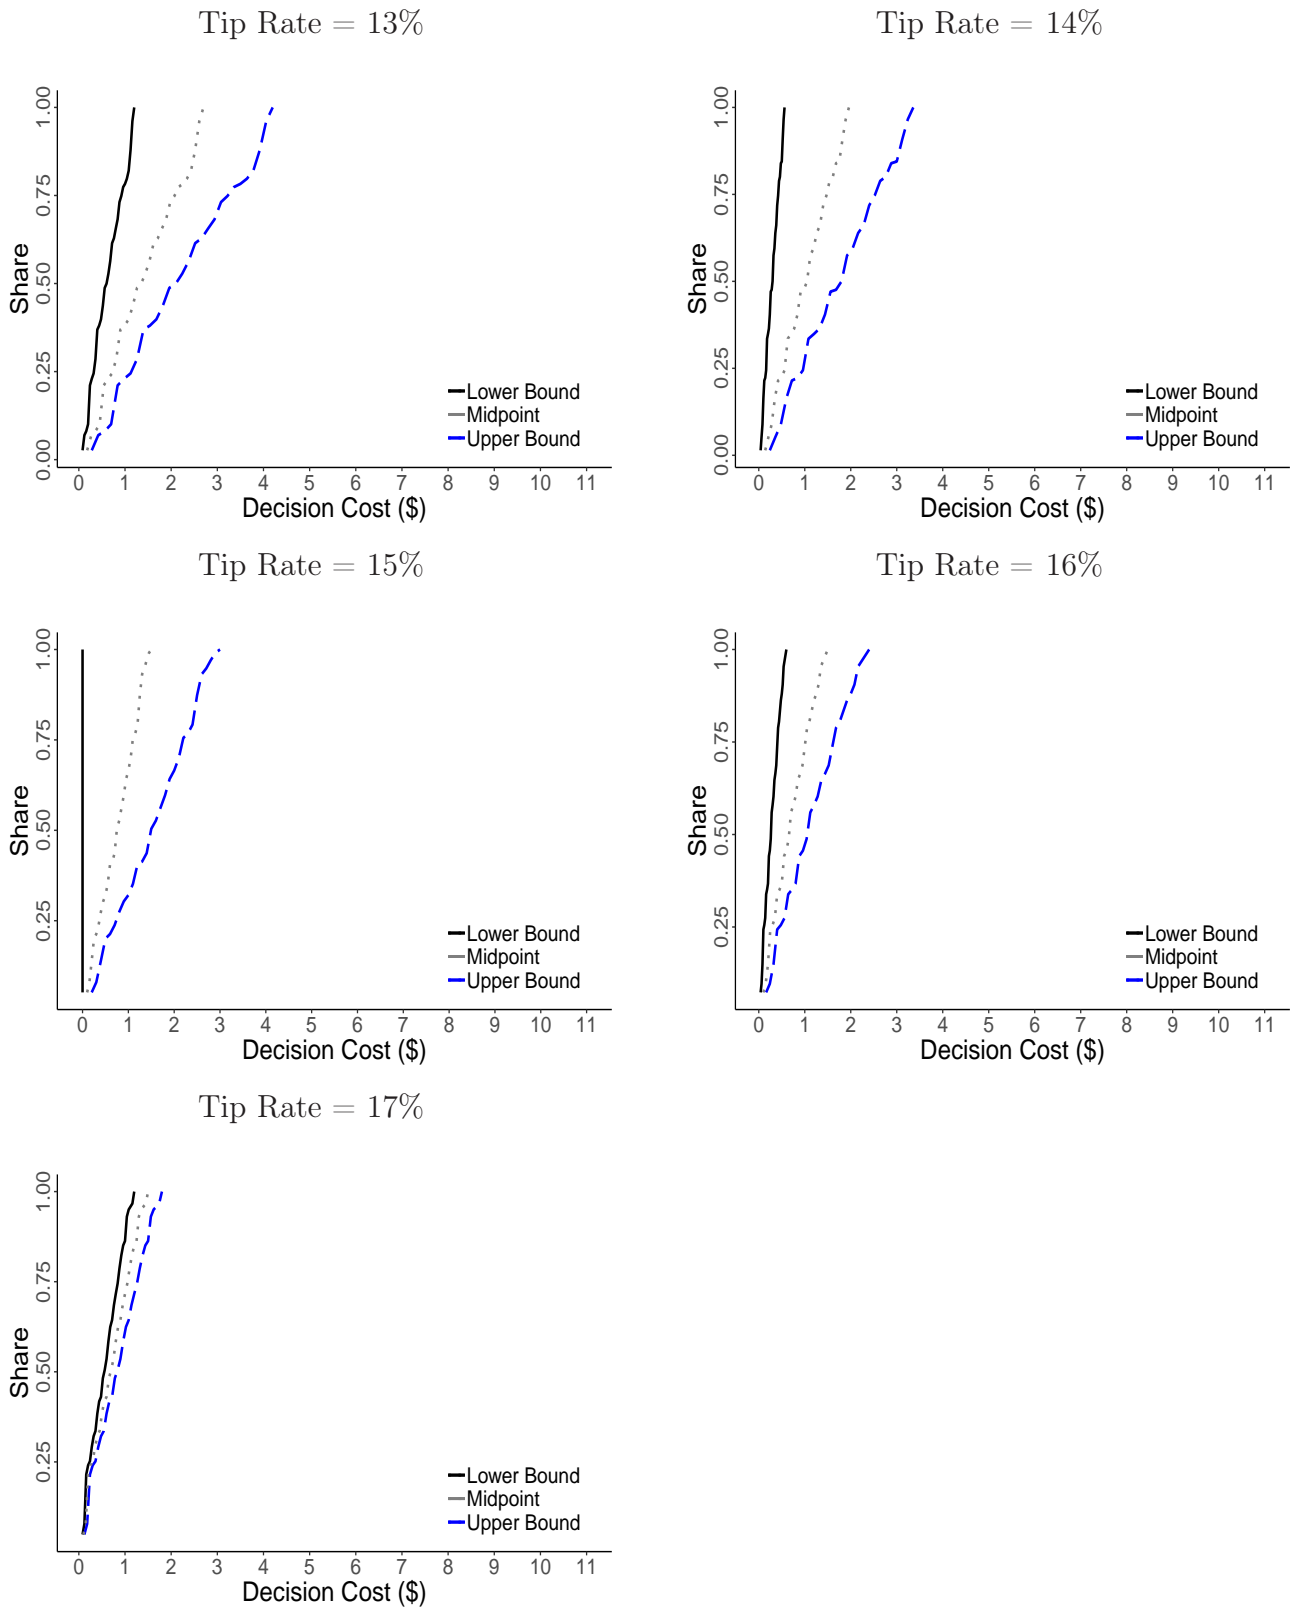


Figure A.5 continued



A.4 A Semiparametric Approach of Estimating Decision Costs

In this section, we use a novel approach complementary to the nonparametric approach from section 4 to estimate decision costs. There are three main innovations: (1) we use changes in the share of passengers who choose non-menu tips as the taxi fare increases to identify decision costs, (2) to compute decision costs, we use a semiparametric strategy, and (3) we use taxi trip data from a different time period (2014) for the estimation. This semiparametric approach of estimating decision costs provides an alternative estimate that can be compared to the nonparametric approach in section 4.

We reason that, if a passenger chooses a tip different from the menu tip options, then she reveals a preference for her tip relative to the menu options. According to the model in section 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. In 2014, all taxicabs presented 20%, 25%, and 30% to passengers as the menu of tip options. We maintain assumptions $\mathbb{A}2$ and $\mathbb{A}3$ (introduced in section 5 in the main text). For trips where passengers tip less than 20% of the taxi fare, we assume that 20% (the lowest menu tip option) is what a passenger would have chosen had she decided to choose a menu tip option. It therefore follows from equation (2) that the benefit (B_i) for a passenger choosing to tip $t_i < d_j = 20\%$ is $B_i = (0.20 - t_i) \times F$. The key intuition is that, all else equal, changing the fare (F_i) changes the values on both sides of the inequality in equation (2). Hence, the sign of the inequality might change for passengers who are on the margin of choosing a non-menu tip. Recall that $t_i\%$ is the observed tip rate in the data, $d_j = 20\%$ is the relevant menu option, and F is the taxi fare. Hence, a passenger who tips $t_i < 20\%$ reveals that her benefit B_i from tipping t_i is greater than her decision costs (denoted by K_i) associated with computing t_i . Notice that, for a passenger who is on the margin of choosing to tip t_i or $d_j = 20\%$, $B_i \approx K_i$.

From figure 5a, we know that passengers are more likely to choose non-menu tips as the taxi fare increases. With this observation in mind, denote $\Pr(t_i|F, d_j = 20\%)$ as the probability of choosing a tip t_i conditional on both the taxi fare F and the menu tip option $d_j = 20\%$. Suppose that F increases by ΔF ; then, it follows from equation (4) that

$$\Delta P = \Pr(t_i|F + \Delta F, d_j = 20\%) - \Pr(t_i|F, d_j = 20\%) \geq 0$$

So ΔP is the change in the probability of choosing one's preferred tip relative to the menu tip option when F increases by ΔF . When ΔF is small (a marginal increase in the fare),

then ΔP 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 tip (i.e., $B_i = (0.20 - t_i) \times (F + \Delta F) \approx K_i$). Therefore, for each $t_i < 20\%$, we can use the changes in the share of passengers ΔP who choose to tip t as the fare increases to estimate the corresponding K . With all the estimated “ K s” and “ ΔP s” in hand, we have all the necessary pieces to construct the distribution of decision cost for passengers who tip less than 20% of the taxi fare.


To implement the procedure above, we estimate an ordered logistic regression where the outcome variable is the tip rate categorized into 20 non-overlapping bins of width one percent (namely 1%, 2%, 3%... 20%) regressed on the fare and other trip characteristics and covariates. Each category of the outcome variable (for example 15% is defined as the share of passengers whose tip fall within the range of 14.5% and 15.5%). Figure A.6 illustrates the estimated predicted probabilities for choosing a non-menu tip for each tip rate as functions of the taxi fare. The figure shows that the probability of choosing any of the non-menu tips is increasing as the fare increases. In contrast, figure A.7 shows that the probability of choosing the menu tip option 20% is decreasing as the fare increases. The support of the fare is restricted to the \$3 - \$30 range. This is because \$3 is the lowest taxi fare, and \$3 - \$30 is the range within which the change in the predicted probabilities (ΔP) is nonnegative for all non-menu tips. The restricted support of the fare results in a censored support of the estimated decision cost K for each tip rate t .

From the estimated predicted probabilities for each of the non-menu tips, we compute K and ΔP using the tip and fare combinations. We use this information to construct the unconditional CDF of decision costs.

The dashed line in figure A.8 shows the estimated CDF of decision costs K using the approach detailed above. The average decision cost is estimated to be \$1.64 (14% of the average taxi fare of \$12.17). Note that this average decision cost is similar to what we estimated using the nonparametric approach in section 4 (\$1.89)—shown as a solid line in figure A.8.


The semiparametric approach has four main limitations. First, we are not able to recover the full distribution of decision costs, since the sample is limited to passengers who tip less than 20%. Second, we assume that the relevant menu option is 20%. If this is not the case, then the estimated distribution of decision costs is a lower bound for passengers who tip less than 20%. Third, the support of the estimated decision costs is censored because only fares within the range of \$3 - \$30 are used in this exercise. Fourth, this approach lumps the two sources of decision costs (norm deviation and cognitive costs) together.

Figure A.6: 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.

Figure A.7: Predicted Probability by Level of Taxi Fare for Choosing a 20% Tip



Note: This figure shows the estimated predicted probabilities for the menu tip rate 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.

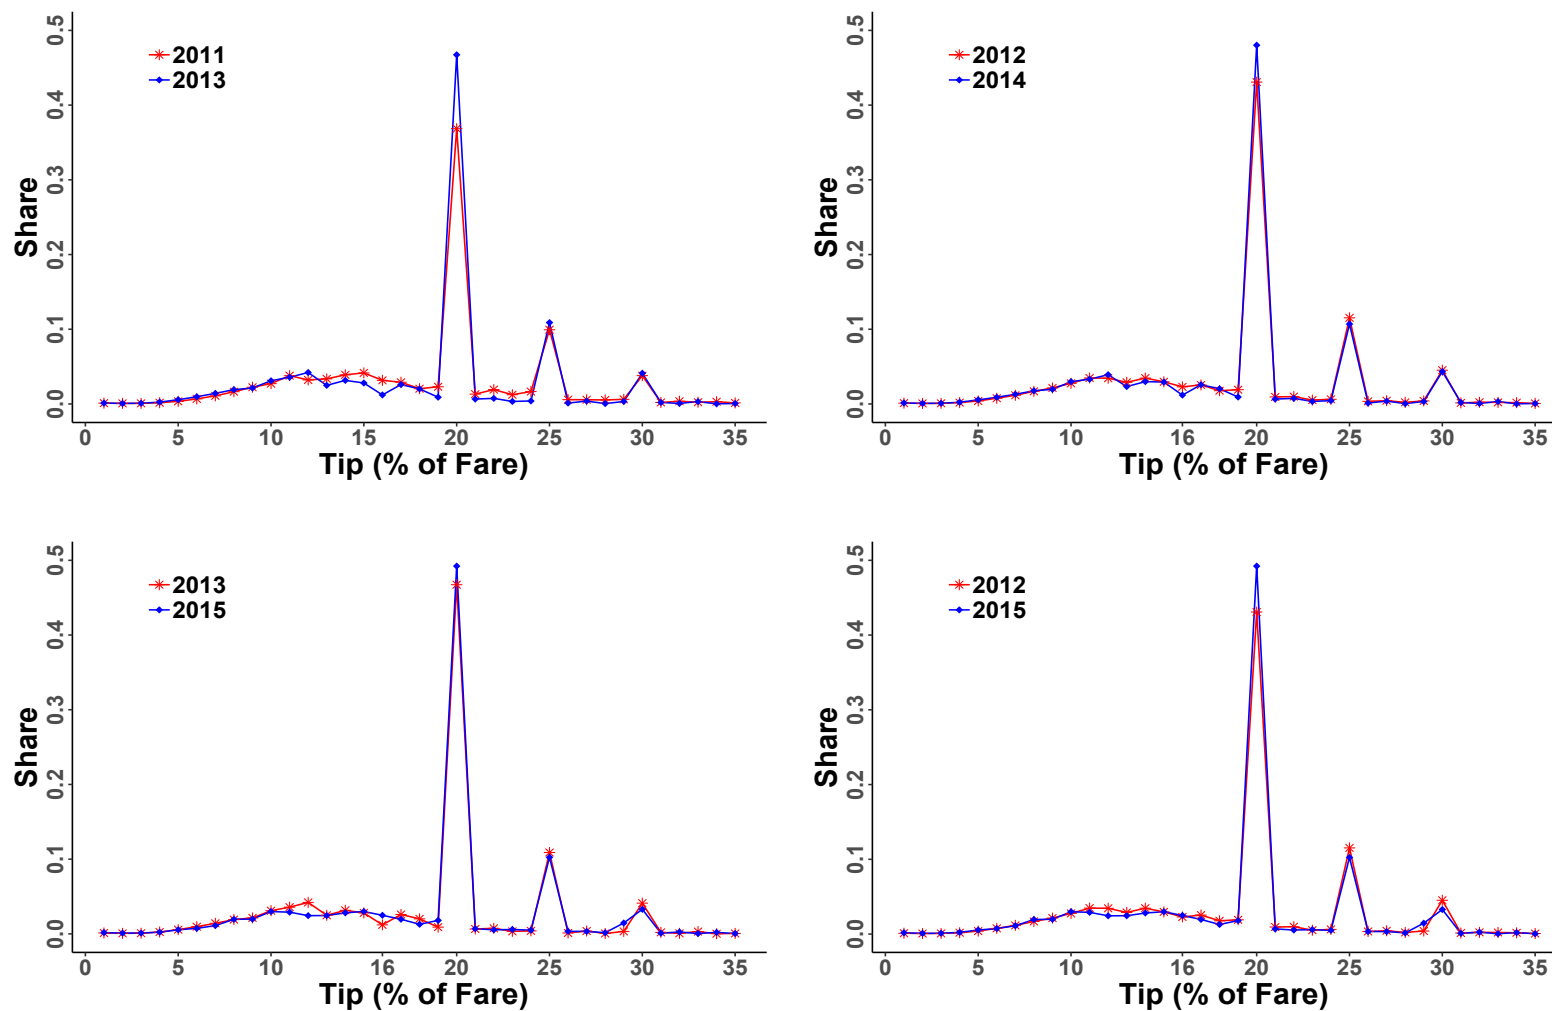
Figure A.8: Estimate of the Unconditional CDF of Decision Costs: Semiparametric Vs. Nonparametric Approach



Note: This figure shows the CDF of decision costs estimated using the semiparametric and nonparametric approach respectively. The dashed line is the estimated CDF from the semiparametric approach and the solid line is the CDF from the nonparametric approach.

A.5 Verifying Assumption A1

Figure A.9: Overlapping Distribution of Tip (%) in Years with No Change in CMT Tip Menu



Notes: These figures compare the distribution of tips five years after CMT changed its menu in 2011—from 15%, 20%, and 30%, to showing 20%, 25%, and 30%. These figures test assumption A1 that states that decision costs are constant across years.

A.6 Verifying Assumption A2

Assumption A2 - One's perception of the tipping norm T_i is jointly independent of the menu and the taxi fare.

A.6.1 Using A Structural Approach to Verify Assumption A2

To assess the validity of assumption A2, we compare estimates of the distribution of tipping norms before CMT (a Yellow taxi electronic credit card machine vendor) changed its menu of tips to estimates of the tipping norms after the menu change and after the Taxi and Limousine Commission increased the taxi fare by about 17%.

Specifically, with our structural model outlined in section 5, we compare the estimated distribution of tipping norms using taxi trips in CMT taxicabs from 2010 to the estimated distribution of tipping norms using taxi trips in CMT taxicabs in 2014. In 2010, passengers riding in CMT taxicabs were presented with a tip menu that showed 15%, 20%, and 30%. However, in 2014, CMT taxicabs presented passengers with a tip menu that showed passengers 20%, 25%, and 30%. In addition, the TLC increased the taxi fare by about 17% in 2012. Therefore, passengers in CMT taxicabs faced both a new menu of tip suggestions and higher fares in 2014. Thus, comparing the distribution of tipping norms in 2010 to that in 2014 serves as a test for assumption A2.

The data we use are standard rate NYC Yellow taxi trips in CMT taxis from 2010 and from 2014. The trips are limited to fares paid for using a credit/debit card, with no tolls, and with passengers leaving a positive tip. We use the same approach as outlined in section 5.2.1 for this exercise. That is, we estimate equation (6) using OLS, and with a Heckman selection correction model to account for sample selection concerns.

Table A.2 columns (1) and (2) present the OLS estimates of equation (6) for trips from 2010 and 2014 respectively. The coefficient estimates differ across the two different years. However, these estimates do not account for the selection concerns mentioned above.³⁶ After accounting for selection, we find that the estimated distribution of the tipping norm is not influenced much by either a change in the presented menu of tips or the change in the taxi fare. Table A.3 presents results from the first step probit estimation of the Heckman selection model and Table A.4 shows the second step of the selection model. All estimates corresponding to the data from 2010 are similar to that of 2014. Figure A.10b depicts the comparison of the estimated distribution of tipping norms from 2010 and 2014.

³⁶Figure A.10a compares the estimated distribution of tipping norms from 2010 to 2014 using the estimates from Table A.2.

Table A.2: OLS Estimates of α_T and θ from Taxi Trips in 2010 and 2014

	Dependent Variable: Tip Rate	
	2010	2014
	(1)	(2)
Fare	-0.001*** (0.00000)	-0.003*** (0.00001)
Constant ($\hat{\alpha}_T$)	0.187*** (0.0001)	0.205*** (0.0001)
$\hat{\theta}$	500	166.667
1(Round Number Tip)	Yes	Yes
Observations	11,389,524	16,394,858
R ²	0.004	0.016

Note: The data are standard rate NYC Yellow taxi trips in CMT taxi cabs from 2010 and 2014. The trips are limited to fares paid for using a credit/debit card, with no tolls, and with passengers leaving a non-menu tip greater than zero. *p<0.1; **p<0.05; ***p<0.01

Table A.3: Probit Estimation of First Step Heckman Selection Model

	Dependent Variable: 1(Tip = Non-Menu Tip)	
	2010	2014
	(1)	(2)
Taxi Fare	0.013*** (0.00004)	-0.029*** (0.00004)
Number of Passengers	-0.688*** (0.0005)	-0.442*** (0.0004)
1(Round Number Tip)	Yes	Yes
Outcomes correctly predicted	78%	79%
Observations	26,498,695	41,620,454
Log Likelihood	-10,306,715.000	-14,436,474.000
Akaike Inf. Crit.	20,613,435.000	28,872,954.000

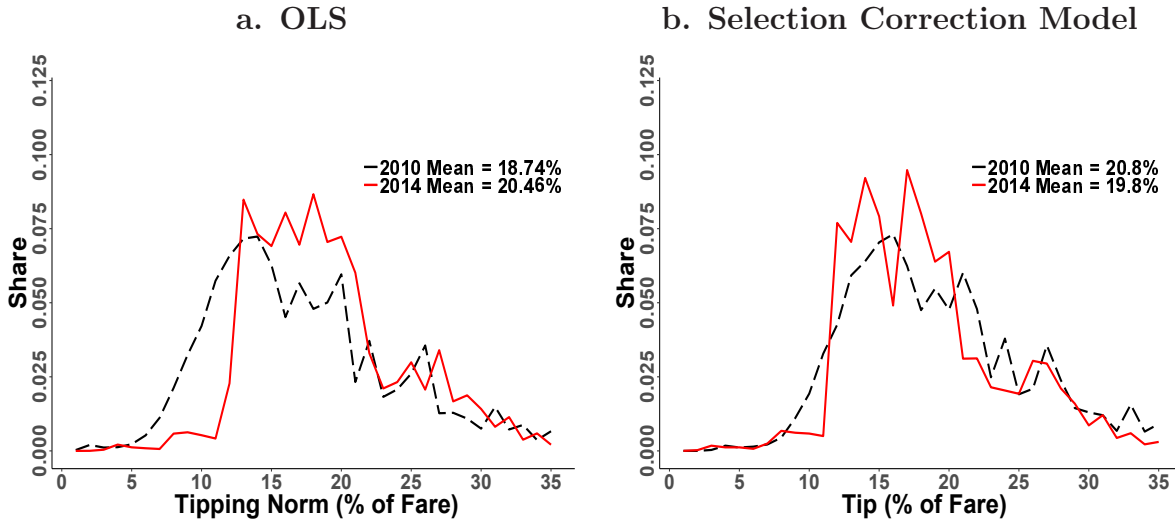
Note: The data are standard rate NYC Yellow taxi trips in CMT taxicabs from 2010 and 2014. The trips are limited to fares paid for using a credit/debit card, with no tolls, and passengers leave a positive tip. The added instrument this estimation is the number passengers in the taxicab ride as reported by the taxi driver. *p<0.1; **p<0.05; ***p<0.01

Table A.4: Heckman Selection Estimates of α_T and θ

	Dependent Variable: Tip Rate	
	2010	2014
	(1)	(2)
Fare	-0.001*** (0.00000)	-0.003*** (0.00001)
Constant ($\hat{\alpha}_T$)	0.208*** (0.0002)	0.198*** (0.0005)
$\hat{\theta}$	500	166.667
1(Round Number Tip)	Yes	Yes
Observations	26,498,695	41,620,454
ρ	-0.140	0.028
Inverse Mills Ratio	-0.015*** (0.0002)	0.005*** (0.0003)

Note: The data are standard rate NYC Yellow taxi trips in CMT taxicabs from 2010 and 2014. The trips are limited to fares paid for using a credit/debit card, with no tolls, and with passengers leaving a non-menu tip greater than zero. *p<0.1; **p<0.05; ***p<0.01

Figure A.10: Estimated Distribution of Tipping Norms T_i (2010 vs. 2014)



Notes: This figure illustrates the estimates of the distribution of tipping norms using our structural model for trips in CMT taxicabs from 2010 versus trips from 2014. Figures A.10a and A.10b show the distributions of tipping norms as estimated via OLS and via the Heckman selection model respectively. The test statistics from a chi-squared goodness of fit test for figure A.10a are $\chi^2 = 3.6472e^{-30}$ with p-value 0.99, and for figure A.10b, $\chi^2 = 3.6472e^{-30}$ with p-value 0.99. The data used are from 2010 and 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive non-menu tip. The tip rate is truncated at 35.5%, beyond which the shares become essentially zero.

Table A.5: Estimates of α_T and θ from Taxi Trips in 2014

	Dependent Variable: Tip Rate		
	OLS: Non-Menu Tips	Selection: Non-Menu Tips	OLS: All Tips
	(1)	(2)	(3)
Fare	-0.0033*** (0.00001)	-0.0033*** (0.00001)	-0.0012*** (0.000003)
Constant ($\hat{\alpha}_T$)	0.2046*** (0.0001)	0.1980*** (0.0005)	0.2206*** (0.00004)
$\hat{\theta}$	166.667	166.667	500
1(Round Number Tip)	Yes	Yes	Yes
Observations	16,394,858	16,394,858	41,620,454
R ²	0.0164	0.0164	0.0453

Note: The data are standard rate NYC Yellow taxi trips in CMT taxicabs from 2014. The trips are limited to fares paid for using a credit/debit card, with no tolls, and with passengers leaving a non-menu tip greater than zero. *p<0.1; **p<0.05; ***p<0.01

A.6.2 Empirical Evidence for Assumption A2

While we cannot formally test assumption A2 without the structural 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 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 A2.

A.7 Verifying Assumption $\mathbb{A}\mathfrak{3}$

Assumption $\mathbb{A}\mathfrak{3}$ - The cognitive cost c_i is 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 $\mathbb{A}\mathfrak{3}$. 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 .

A.7.1 c_i is Independent of One's Preferred Tip t_i^*

Because we do not observe c_i , we cannot directly 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.11](#) 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 $\mathbb{A}\mathfrak{3}$, the data are consistent with this assumption.

Figure A.11: 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.7.2 c_i is Independent of the Taxi Fare F

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.5 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 $\mathbb{A}3$, it does suggest that the data seems consistent with it.

Table A.6: Evidence for Assumption A3

<i>Dependent variable:</i>	
	1(Tip = 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 (0.076)
Constant	0.601*** (0.0001)
Observations	41,620,454
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, with no tolls, and with passengers leaving a positive tip. *p<0.1; **p<0.05; ***p<0.01

Figure A.12: Distribution of Predicted Taxi Fare Thresholds



Note: This figure plots the distribution of predicted taxi fare thresholds where passengers are indifferent between choosing to compute their preferred tip and choosing a tip menu. Using the estimated parameters from the model we compute the fare threshold as

$$\bar{F}_i = \frac{c_i - \hat{\theta} \left[\left(\hat{T}_i - t_i^* \right)^2 - \left(\hat{T}_i - d_j \right)^2 \right]}{d_j - t_i^*}.$$

The plot is limited to fares thresholds between \$0 and \$100.

A.8 Grid Search for Tip-Maximizing Menu Tip Options

Figures A.8 and A.9 below depict the process involved in estimating the tip-maximizing menu of tips. Specifically, the figures detail how we estimate the tip-maximizing menu for (1) the case where drivers are restricted to show passengers a single menu tip option, and (2) the case where drivers are restricted to show passengers two menu tip options.

For the case where drivers are restricted to showing passengers a single menu tip option, we 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 rate the most. Using the parameter estimates from table 2 column (1) (OLS results from Panels A and B), figures A.13.a1 and A.13.a2 show the results from the exercise above.

The average tip rate when passengers are not shown a menu is about 16.87%—shown


as a horizontal dashed line in figure [A.13.a1](#). Our model predicts that the average tip rate is maximized when passengers are shown 22% as the sole menu option (shown as a dotted vertical line in figure [A.13.a1](#)). Figure [A.13.a2](#) is a bar graph that shows the model predicted distribution of tips when 22% is presented to passengers as a menu tip (displayed as shaded bars) versus the distribution of preferred tips (displayed as overlaid un-shaded bars). With the 22% menu tip option, the average tip rate is 18.64%. This is an 11% increase in the average tip. Figure [A.13.a1](#) shows that presenting a menu tip below 13% depresses the average tip rate. In sum, the choice of a menu tip presented to passengers can have either positive or negative consequences.

Figures [A.13.b1](#) and [A.13.b2](#) show a similar exercise as in figures [A.13.a1](#) and [A.13.a2](#). Here, we search for the two menu tips that will maximize tips received. Figure [A.13.b1](#) shows a three-dimensional surface that characterizes the average tip from the different combinations of two menu tips between 0% and 50%. Our model predicts that showing 21% and 28% as menu tips maximizes tips. These two menu tips increase the average tip to 18.9%—a 12.4% increase compared to the average preferred tip rate. An inspection of figure [A.13.b1](#) also suggests that certain choice combinations of the two options can either increase or decrease tips.


We find very similar results when we conduct this exercise using estimates from the Heckman selection correction model estimates from table 2 column (2) (Panels A and B). Figure [A.14](#) shows a panel that is analogous to figure [A.13](#).

Figure A.13: OLS Results: Grid Search for Tip-Maximizing Menu Tip Options and Model Predicted Tips


a1. Grid Search for One Menu Tip Option




a2. Predicted Tips (One Menu Option)



b1. Grid Search for Two Menu Tip Options




b2. Predicted Tips (Two Menu Options)




Notes: Given a menu that presents customers with percentage tips, figures A.13a and A.13b show results from computing the tip-maximizing menu that show one and two menu options respectively. Panel A.13.a1 plots the results from the grid search for a single menu tip option that will maximize the average tip. Panel A.13.a2 shows the model prediction of the distribution of tips when the tip-maximizing menu option from panel A.13.a1 is presented as a menu tip versus the model predicted preferred tips when passengers are not presented with a menu. Panels A.13.b1 and A.13.b2 are analogous to panels A.13.a1 and A.13.a2 but for the case of choosing a menu that presents two menu options. The data used are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive tip amount.

Figure A.14: Selection Model Results: Grid Search for Tip-Maximizing Menu Tip Options and Model Predicted Tips


a1. Grid Search for One Menu Tip Option




a2. Predicted Tips (One Menu Option)



b1. Grid Search for Two Menu Tip Options




b2. Predicted Tips (Two Menu Options)




Notes: Given a menu that presents customers with percentage tips, figures A.14a and A.14b show results from computing the tip-maximizing menu that show one and two menu options respectively. Panel A.14.a1 plots the results from the grid search for a single menu tip option that will maximize the average tip. Panel A.14.a2 shows the model prediction of the distribution of tips when the tip-maximizing menu option from panel A.14.a1 is presented as a menu tip versus the model predicted preferred tips when passengers are not presented with a menu. Panels A.14.b1 and A.14.b2 are analogous to panels A.14.a1 and A.14.a2 but for the case of choosing a menu that presents two menu options. The data used are from 2014 standard rate taxi trips paid for via a CMT credit card machine along with a positive tip amount.

Figure A.15: Distribution of Tip (%) by Type of Tip Menu Versus No Menu Tips


a) Previous Tip Menu




b) Current Tip Menu



c) Tip-Maximizing Menu



d) Utility-Maximizing Menu



Note: This figure shows the model predicted distribution of tips for four different tip menus namely (15%, 20%, 25%), (20%, 25%, 30%), (21%, 27%, 33%), and (10%, 16%, 25%). The predictions are made using the estimated parameters of the model.