Essays on Economic Behavior and Design

Essays on Economic Behavior and Design

A dissertation presented

by

Neil Thakral

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Business Economics

Harvard University

Cambridge, Massachusetts

April 2018

Dissertation Advisor: **Professor Lawrence F. Katz**

Author: **Neil Thakral**

Essays on Economic Behavior and Design

Abstract

This dissertation consists of essays on economic behavior and design, with an emphasis on understanding how people make decisions in labor markets and social programs.

The first two chapters, written jointly with Linh T. Tô, examine how reference points influence behavior. The first chapter studies the dynamics of reference dependence by asking how the timing of earnings within a day affects labor supply decisions. We find that money is not fungible over time, inconsistent with a standard neoclassical model of intertemporal optimization as well as alternative behavioral models invoking daily income targets. We reconcile these views by proposing a model of adaptive reference points. The second chapter empirically investigates the speed of adjustment of the reference point, a key degree of freedom in models of reference dependence. We show that reference points tend to adjust more readily in the direction of gains rather than losses. Our results are inconsistent with the idea of reference points based on rational expectations, which would imply no such asymmetry in reference-point adjustment.

The final chapter explores the design of policies for allocating public housing. I present a model of public-housing allocation and investigate the design of allocation mechanisms that are strategy-proof, or not subject to strategic manipulation. I characterize a new mechanism—the Multiple-Waitlist Procedure (MWP)—which allows applicants to optimally trade off their preferences for different units and waiting times by choosing among a set of waiting lists. Using estimates of preferences for public housing in Pittsburgh, I find that a counterfactual change from existing allocation mechanisms to the MWP would lead to substantial welfare gains.

Contents

List of Tables

List of Figures

Acknowledgments

I am grateful to Isaiah Andrews, Pedro Bordalo, David M. Cutler, Stefano DellaVigna, Kfir Eliaz, Itay P. Fainmesser, Henry S. Farber, Josh Feng, Peter Ganong, Edward L. Glaeser, Benjamin Golub, Jerry R. Green, Kevin He, Sonia Jaffe, Maximilian Kasy, Lawrence F. Katz, Marc Kaufmann, Scott Duke Kominers, Maciej H. Kotowski, David Laibson, Robin S. Lee, Xiaosheng Mu, Parag A. Pathak, Matthew Rabin, Joshua Schwartzstein, Jesse M. Shapiro, Andrei Shleifer, Ran I. Shorrer, Tayfun Sönmez, Elie Tamer, and Linh T. Tô for helpful comments.

Chapter 1

Daily Labor Supply and Adaptive Reference Points[1]

1.1 Introduction

Reference dependence plays a contentious role in labor supply. Ever since Camerer et al. [1997] showed a negative relationship between average daily wages and the number of hours worked each day for taxi drivers in New York City, studies on the daily labor-supply decisions of workers who can flexibly choose their work hours have fallen into a clear dichotomy. The suggestion that workers may have a daily income target led to studies that either support the reference-dependence model if there is evidence of such a target, or uphold the neoclassical model with the rejection of an income target. This paper bridges these perspectives through the notion of adaptive reference points and quantifies the speed of reference-point adjustment.

We find two main empirical results by analyzing a dataset consisting of all New York City (NYC) cab fares in 2013 using non-parametric methods. First, we document excess sensitivity of labor-supply decisions to daily earnings. A driver who has already worked for 8.5 hours, for example, is on average 3 percent more likely to end a shift when cumulative

[1]Co-authored with Linh T. Tô

daily earnings are 10 percent higher. This is contrary to the prediction of the neoclassical model that small windfalls should only trivially affect the marginal utility of lifetime wealth and therefore should leave labor-supply decisions unchanged. Second, we find stronger labor-supply reductions in response to earnings that accumulate more recently. One might naturally expect that the timing of income is economically irrelevant—an extra dollar earned at 12 PM is no different from an extra dollar earned at 11 AM from the perspective of a driver at 1 PM—or, if anything, recent earnings might be more informative about future earnings opportunities, which should make a driver less likely to end a shift in response to more recent earnings. However, we find that a driver who has already worked for 8.5 hours, for example, responds seven times more strongly to an additional dollar earned in the eighth hour of the shift compared to one earned in the fourth hour. The patterns persist for stopping decisions at different hours throughout the shift as well as for shifts that start at different hours of the day. We use an empirical Monte Carlo exercise to highlight the importance of the non-parametric methodology, as prior methods can yield spurious results due to functional-form assumptions.

We interpret these facts as evidence of a daily income effect and a violation of fungibility of money over time, even within a single day. We consider a number of alternative explanations including learning about future earnings, option value, and liquidity constraints, and show that these explanations cannot account for the patterns we observe. Furthermore, the income effect does not decrease with experience. We also consider the possibility that daily earnings are correlated with unobserved determinants of labor supply such as effort or fatigue and show that the results hold when using instrumental-variables strategies that rely on variation in earnings due to distance or tips.

The findings are not only inconsistent with the neoclassical model but also with canonical behavioral explanations, necessitating an alternative formulation. Existing work, which invokes reference dependence and loss aversion to explain income-targeting behavior, does not account for the violation of fungibility for money earned at different times within a day. To the extent that a reference level influences decisions, our findings imply that the reference level must adjust within a day. A daily-level target for income does not permit stronger reactions to more recent earnings, and a reference level that adjusts instantaneously likewise does not make any distinction based on the timing of income, producing under some specifications behavior that resembles the neoclassical prediction. Our formulation consists of a slow-adjusting reference point, which incorporates experiences earlier in the day to a greater extent and results in stopping behavior that depends more strongly on recent earnings.

We explore several classes of models that can potentially explain the patterns of behavior that the data reveal. We develop structural models of daily labor supply and use the data on stopping decisions to estimate the parameters using maximum likelihood. As a benchmark, we consider models based on income targeting, expectations-based loss aversion [Kőszegi and Rabin, 2006], and salience [Bordalo et al., 2015], which can account for income effects but not the violation of fungibility within a day. We then propose two models that incorporate adaptive reference points, one based on loss aversion and one based on salience, and estimate the parameters of these extended models. We assess these explanations by comparing the estimated models' predictions about the magnitude and timing pattern of the income effect, as well as by deriving and testing an additional prediction about how the magnitude of the income effect changes around the reference point. Both models capture the main qualitative features of the data, namely that labor supply reduces in response to earnings with stronger effects for more recent earnings. However, our model of loss aversion tends to overstate some of the quantitative features, as the maximum likelihood estimate of the coefficient of loss aversion implies a magnitude of the income effect and a change in behavior around the reference point that exceed what we observe in the data. We discuss how the results suggest an important role for stochastic reference points and provide support for models that predict a lower degree of loss aversion over money such as the notion of news utility in Kőszegi and Rabin [2009].

We highlight the implications of our findings for three different areas of work: daily labor supply, structural behavioral economics, and reference dependence.

The main empirical contribution of this paper is to provide evidence of daily income effects and a violation of fungibility by studying how workers adjust labor supply in response to small changes in accumulated daily earnings. By contrast, most studies in settings where workers can choose their own daily hours focus on the effects of transitory *wage* changes on labor supply. Camerer et al. [1997] present evidence of negative wage elasticities in two out of three samples of NYC cabdrivers, but the findings in this literature are mixed.[2] Two recent papers revisit the setting of cabdrivers in NYC using a more comprehensive dataset: Morgul and Ozbay [2014] estimate elasticities for each of four months in 2013 using data from all trips taken in all taxi cabs in NYC and find a negative elasticity during only one of the months; Farber [2015] uses a sample of 13 percent of all cabdrivers between 2009 and 2013 and finds negative elasticities for only one-third of day-shift drivers and one-seventh of night-shift drivers. One potential concern with elasticity estimates in this setting is the implicit assumption that cabdrivers treat the average daily wage as parametric to their labor-supply decisions. Farber [2005] argues that the assumption of a parametric daily wage rate is unreasonable and instead proposes a model in which cabdrivers decide whether to stop working at the end of each trip. The stopping model serves as a starting point for our analysis, as we share the focus on daily income effects rather than elasticity estimates.

Despite the growing number of papers estimating high-frequency labor-supply elasticities, few authors investigate daily income effects. Using trip-level data for 21 cabdrivers, Farber [2005] concludes based on the stopping model that cumulative daily earnings do not significantly influence labor-supply decisions.[3] Due to the limited sample, however, Goette

[2] Papers finding negative elasticities include: Chou [2002] and Agarwal et al. [2015] on taxi drivers in Singapore; Ashenfelter et al. [2010] and Doran [2014] on taxi drivers in NYC; Dupas et al. [2016] on bicycle-taxi drivers in Kenya; Chang and Gross [2014] on pear packers in California; and Nguyen and Leung [2015] on swordfish fishermen in Hawaii. Papers finding positive elasticities include: Jonason and Wållgren [2013] on taxi drivers in Stockholm; and Stafford [2015] on lobster fishermen in Florida.

[3] Also see recent field experiments by Andersen et al. [2014] in India, which concludes that the labor supply of vendors in their Betel Nut Experiment remains unchanged in response to unexpected cash windfalls in the morning, and Dupas et al. [2016] in Kenya, which examines the effects of unexpected cash windfalls on a group of workers who have daily cash needs.

et al. [2004] express concerns about the ability to identify income effects, which our paper addresses by using comprehensive administrative data. In addition, we document biases that arise due to functional-form assumptions in the stopping model (see Section A.3.2). We circumvent these biases using a non-parametric methodology and, contrary to the result in Farber [2005], find evidence of substantial daily income effects. Moreover, we find that income effects depend on recency: a worker is disproportionately more likely to end a shift in response to a dollar received later in the shift than if the same dollar were received earlier in the shift.

As a second contribution of the paper, our model-comparison exercises provide a detailed comparison between different behavioral mechanisms that can potentially explain our findings. An emerging body of work in structural behavioral economics generally focuses on estimating the parameters of a single behavioral model to test the null hypothesis of the neoclassical model.[4] Our structural models of daily labor supply serve several purposes: testing against the neoclassical model, testing whether reference points adjust, comparing the magnitudes of the income effects that the estimates predict, and testing an additional implication about moments of the data that the models were not designed to fit. Our results highlight that rejecting the neoclassical prediction does not imply a validation of a particular alternative. The few existing papers that analyze multiple behavioral theories generally reject the neoclassical model and broadly conclude support for the behavioral alternatives. Hastings and Shapiro [2013] find that consumers substitute to lower octane gasoline to an extent that cannot be explained by income effects when prices rise, consistent with their implementations of both loss aversion [Kőszegi and Rabin, 2006] and salience [Bordalo et al., 2013]. Busse et al. [2014] analyze the effect of weather on car purchases and obtain results predicted by projection bias [Loewenstein et al., 2003] and also consistent with salience theory [Bordalo et al., 2013]. Barseghyan et al. [2013] demonstrate the importance of probability weighting for explaining observed levels of risk aversion in

[4]See, for example, Conlin et al. [2007], Crawford and Meng [2011], DellaVigna et al. [2012], Grubb [2012], and Laibson et al. [2015].

insurance deductible choices, but cannot conclude whether loss aversion [Kőszegi and Rabin, 2006] or disappointment aversion [Gul, 1991] also plays a role. The present paper highlights the link between features of the data and assumptions of various behavioral models.

Prior work on daily labor supply tends to equate the question of whether workers behave according to the predictions of the neoclassical model with the sign of a single parameter, the elasticity of labor supply with respect to average daily wages, interpreting a positive elasticity as evidence in favor of the neoclassical model and a negative elasticity as evidence in favor of reference dependence and loss aversion. We depart from this paradigm for three reasons: first, within-day variation in the wage rate can bias elasticity estimates in either direction (see Section A.2); second, reference dependence does not imply negative labor-supply elasticities if changes in wages are anticipated [Kőszegi and Rabin, 2006]; third, simply rejecting the neoclassical prediction does not imply a validation of a particular alternative theory.

A third contribution of the paper is to provide field evidence on reference-point formation and adjustment. A growing body of work shows that reference dependence influences behavior in a variety of settings.[5] However, much of the existing evidence on how reference points adjust comes from decisions in game shows and lab experiments [Post et al., 2008, Gill and Prowse, 2012, Song, 2016]. More recently, DellaVigna et al. [2017] find evidence of reference-point adjustment in the context of job search over a horizon of several months. By contrast, the use of high-frequency labor-supply data allows us to detect the speed of reference-point adjustment at the daily level. Moreover, while the model in DellaVigna et al. [2017] involves a reference-dependent utility function with a backward-looking reference point, a number of recent papers provide evidence for forward-looking expectations-based reference points.[6] We take expectations-based reference points as a benchmark and express

[5]See, for example, Barberis et al. [2001], Fehr and Goette [2007], Card and Dahl [2011], and Allen et al. [2016].

[6]For example, see lab evidence from Abeler et al. [2011], Gill and Prowse [2012], Karle et al. [2015], and Sprenger [2015] as well as field evidence from Post et al. [2008], Card and Dahl [2011], Ericson and Fuster [2011], and Pope and Schweitzer [2011].

updated reference points as a function of lagged expectations. Our formulation remains consistent with the notion of reference points as recent expectations in Kőszegi and Rabin [2006] but generates some backward-looking features, which can also account for empirical results that find an influence of past prices on behavior in various settings [Odean, 1998, Genesove and Mayer, 2001, Hastings and Shapiro, 2013].

1.2 Data

1.2.1 Background

Our study uses trip-level data provided by the New York City Taxi and Limousine Commission (TLC) for every fare served by NYC medallion taxicabs in 2013. The "trip sheets" consist of detailed information about each fare, including identification numbers for the driver and car, start and end times for each trip, pick-up and drop-off locations, tips paid by credit card, and the fare charged. These data are collected and transmitted electronically in accordance with the Taxicab Passenger Enhancements Project (TPEP).[7]

Prior to TPEP, cabdrivers were required to fill out trip sheets by hand to record and store information on paper about each fare. By 2008, all medallion taxicabs had implemented a series of technology-based service improvements—including credit/debit card payment systems, passenger information monitors, text messaging between TLC and drivers, and automated trip sheet data collection—due to a March 2004 mandate by the Board of Commissioners of the TLC. Along with these service improvements, the automated trip sheet data include Global Positioning System (GPS) coordinates for pick-up and drop-off locations, which are available for over 98 percent of the data.

In each trip, the fare is determined by the meter. The standard city rate is calculated by combining a base rate of $2.50, any surcharges, and an additional amount that depends on

[7]Haggag and Paci [2014] use the TPEP data in a study on default tip suggestions, which analyzes fares that were paid with a credit card and not subject to any surcharges. Farber [2015] also uses these data in a replication and extension of Camerer et al. [1997].

Figure 1.1: *Supply of cabs throughout the day*

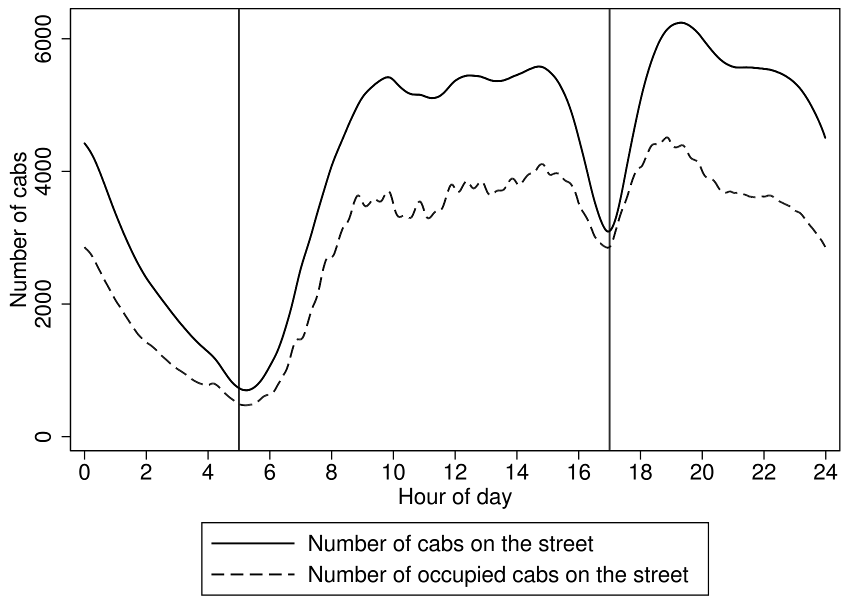

Note: The figure depicts the average number of cabs that are on the road at any given minute of the day in our cleaned data. The solid line depicts the supply pattern of cabs searching or carrying passengers. The dashed line depicts the supply pattern of cabs with passengers.

the distance/time driven.[8] Prior to September 4, 2012, the fare increases by \$0.40 per unit (approximately 0.2 miles); afterwards the incremental charge is \$0.50 per unit.[9]

Figure 1.1 depicts the average number of cabs that are on the road working during each minute of the day. The systematic drops in the number of cabs available in the early morning and early evening reflect the common institutional arrangement whereby two drivers share the same cab (typically switching at 5 AM and 5 PM). The TLC regulates the maximum amount that can be charged to lease a cab for a twelve-hour shift, with a "lease

[8]The TLC reports three categories of surcharges: New York State Tax Surcharge of \$0.50 for every ride on or after November 1, 2009; Night Surcharge of \$0.50 between 8 PM and 6 AM; and Peak hour Weekday Surcharge of \$1.00 between 4 PM and 8 PM. Source: `http://www.nyc.gov/html/tlc/html/passenger/taxicab_rate.shtml`.

[9]More precisely, the TLC defines a *unit* as follows: driving one-fifth of one mile at a speed of at least 6 miles per hour; or 60 seconds when the cab is not in motion or is traveling at less than 12 miles per hour. Source: `http://www.nyc.gov/html/tlc/html/passenger/taxicab_rate.shtml`.

cap" of roughly $130 depending on the day of the week and the time of the shift.[10]

In addition to institutional constraints, weather can potentially affect labor-supply decisions. Our study uses minute-level weather data (temperature, precipitation, and wind speeds) from the National Centers for Environmental Information collected at five locations around NYC. We match each trip from the TPEP data with the weather conditions at the closest station during the minute when the trip ends.

1.2.2 Descriptive Statistics

The raw data consist of information on about 41,000 unique drivers and 14,000 taxicabs taking around 173 million trips in 2013. To study cabdrivers' labor-supply decisions, we group trips into *shifts*. As in Haggag and Paci [2014], a shift is a sequence of consecutive trips that are not more than six hours apart from each other. In other words, a given trip is the last one in its shift if and only if it is followed by a period of at least six hours during which the driver does not pick up any more passengers. As in Farber [2005], we define a break as a long waiting time between fares: at least 30 minutes between a fare that ends in Manhattan and a fare that starts in Manhattan (over 85 percent of trips); at least 60 minutes between fares that start or end outside Manhattan but do not end at an airport; or at least 90 minutes between a fare that ends at an airport and the next fare. After eliminating shifts with missing or inconsistent information (described in Section A.1), about 75 percent of the observations (over 5.8 million shifts by over 37,000 drivers) remain, comprising over $1.5 billion in transactions for cab fares.

Table 1.1 provides summary statistics at the trip level. Over 85 percent of all trips start and end in Manhattan, and the median ride is 10 minutes long. The median fare is about $9.5, with 90 percent of fares falling below $22. We observe tips for the 54 percent of fares that are paid using a credit card. A driver collects an average of $2.48 in tips per trip, but

[10]During our sample period, the lease caps for standard vehicles were $115 for all AM shifts, $125 for Sunday-Tuesday PM shifts, $130 for Wednesday PM shifts, and $139 for Thursday-Saturday PM shifts. The lease caps for hybrid vehicles are $3 higher. Cabs can also be leased on a weekly basis, with a lease cap that is about six-sevenths of the sum of the daily lease caps. Source: `http://www.nyc.gov/html/tlc/downloads/pdf/lease_cap_rules_passed.pdf`.

Table 1.1: *Trip-level summary statistics*

	Median	Mean	Standard deviation
Ride duration (minutes)	10	12.6	9.2
Wait duration (minutes)	5	11.3	19.6
Fare (dollars)	9.5	12.2	9.4
Tip ratio (percent)	20	19.3	12.3

Note: This table reports summary statistics at the trip level for all 127 million NYC taxi trips in 2013 in the cleaned data (see Section A.1). Ride duration is the number of minutes between pick-up time and drop-off time. Wait duration is the number of minutes between dropping off a passenger and picking up a new passenger. Fare is the amount earned not including tips. Tip ratio is the tip divided by the fare, which is available for the 54% of trips with credit card as the payment type.

Figure 1.2: *Tip distribution by fare*

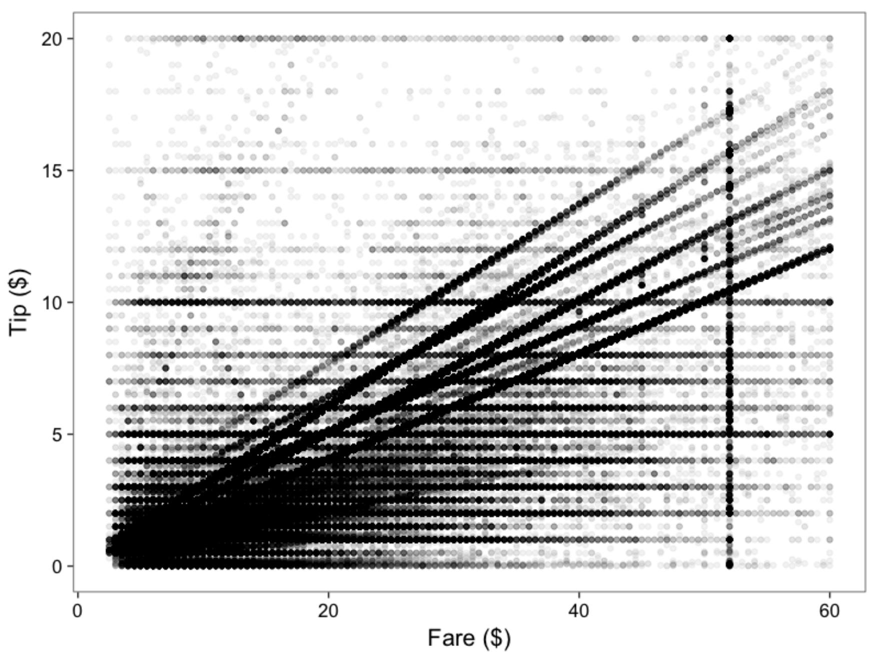

Note: The figure depicts the distribution of tips by fare in the sample for which we observe a tip when tips are between 0 and 20 dollars and fares are between 0 and 60 dollars. The level of darkness represents the density of the points.

Figure 1.3: *Shift-level summary statistics*

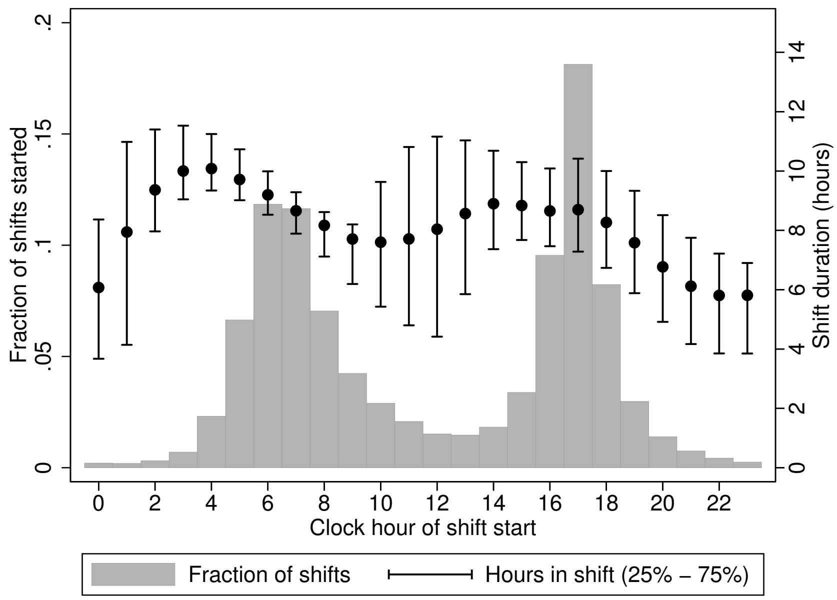

Note: The histogram depicts the distribution of shifts by the clock hour of when the shift starts between hour 0 and hour 23. For each clock hour, the distribution of duration of shifts starting at that hour is depicted by the bar graph, with the mean and interquartile range.

there is substantial variation in the rate of tipping, as Figure 1.2 shows. Given any fare between the minimum fare of $2.5 and $60, the associated tip can take any value between $0.50 and $20, with higher concentrations of rounded tips or fixed fraction of the fare, anywhere between 10 percent and 35 percent. Around 65 percent of shifts contain a tip of at least $5, and 20 percent of shifts contain a tip of between $10 and $20. Haggag and Paci [2014] provide further evidence on variation in the rate of tipping.[11]

Figure 1.3 displays the fraction of shifts starting at each hour of the day as well as the distribution of work hours. A typical shift consists of 22 trips, with 75 percent of shifts exceeding 7.2 hours. About 64 percent of the time in an average shift is spent with a

[11] For cabs that are equipped with credit-card machines from the largest vendor (accounting for 50 percent of cabs in NYC), there is a discontinuity in suggested tips when the fare reaches $15; Haggag and Paci [2014] exploit this discontinuity to show that default suggestions influence passenger tipping behavior.

Figure 1.4: *Pattern of wages throughout the day*

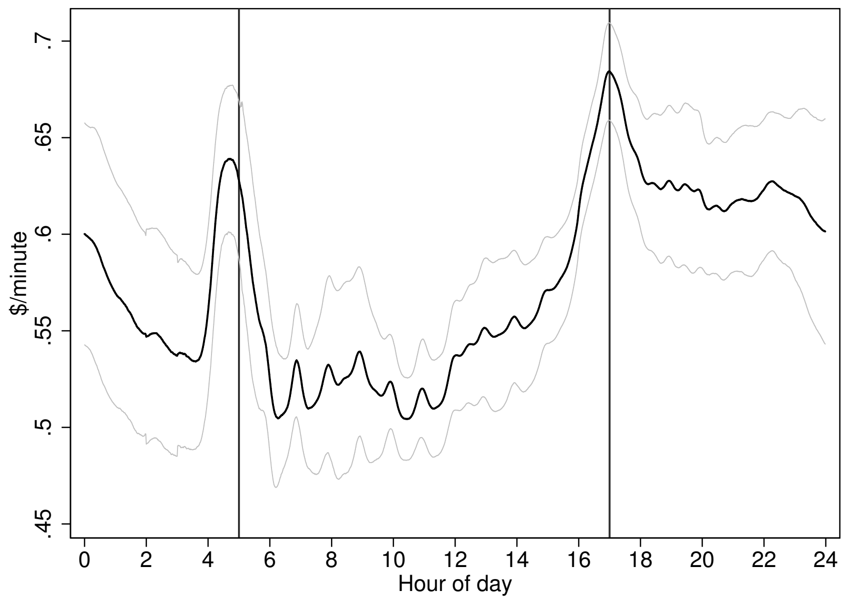

Note: The figure depicts the average market wage every minute throughout the day from hour 0 to hour 23. The market wage in each minute is the average of the per-minute wages of all drivers working during that minute, where a driver's per-minute wage is the ratio of the fare (not including tips) to the number of minutes spent searching for or riding with passengers for their current trip. Gray lines are one-standard-deviation bounds over the course of the year 2013.

passenger in the cab, 26 percent of the time is spent searching for the next passenger, and 13 percent of the time is spent on break.

The market wage varies considerably throughout the day. For each minute that a driver spends searching for or riding with passengers, we define the driver's per-minute wage as the ratio of the fare they earn for that trip to the number of minutes spent working (i.e., searching and riding). We define the market wage in each minute as the average of the per-minute wages of all drivers working during that minute. Figure 1.4 depicts the average wage during each minute of the day, and Figure A.1 shows how the wage pattern differs between weekdays and weekends.[12] The highest wages occur during the two hours with

[12]Section A.2 documents that the pattern of wages throughout the day can be a source of bias for elasticity

Figure 1.5: *Autocorrelation of residualized hourly market wage*

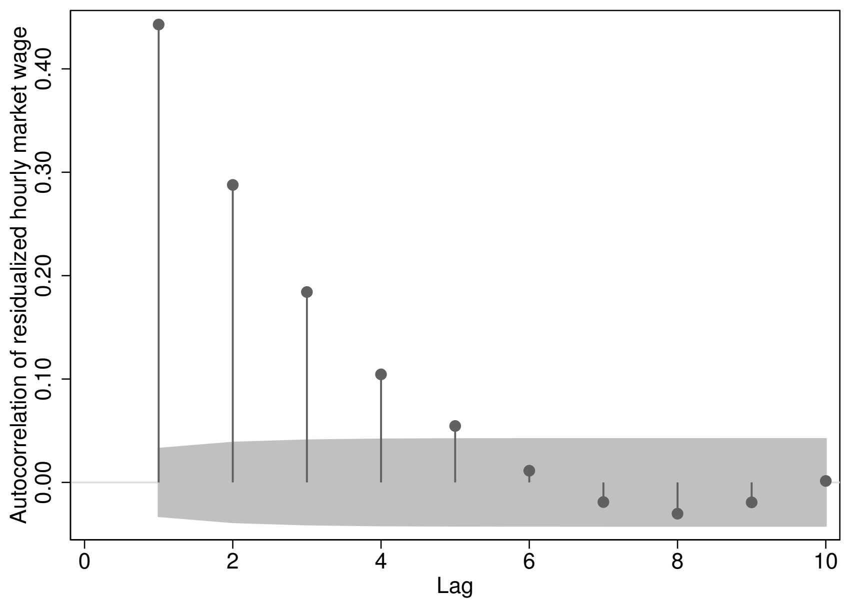

Note: The figure depicts the autocorrelation of hourly market wages indexed by hour of the calendar year 2013. The hourly market wage is the sum of the minute market wage in each hour, with the minute market wage computed as in Figure 1.4. The hourly market wage is residualized from a regression on a set of time and weather effects: an interaction between the hour of day and day of week, the week of the year, an indicator for federal holidays, an indicator for whether it rains during that hour, and indicators for high (over 80 degrees Fahrenheit) and low (under 30 degrees Fahrenheit) average hourly temperature. The shaded region denotes a 95-percent confidence band.

the lowest number of drivers, i.e., the transitions between AM and PM shifts each day (see Figure 1.1).[13] A cabdriver earns an average wage of about $31 per hour, which amounts to a gross income (excluding tips) of about $280 per shift, from which drivers may pay leasing fees and gasoline costs.[14] Figure 1.5 investigates the predictability of hourly wages. Residualizing hourly wages on a set of time effects (an interaction between the hour of day and day of week, the week of the year, and an indicator for federal holidays) and weather effects, which represent a substantial component of the variation in hourly wages but do not reflect transitory daily variation, the figure shows a positive autocorrelation.

1.3 Tests of Income Effects

A neoclassical model of intertemporal utility maximization predicts that daily hours of work do not respond to small changes in accumulated daily earnings. As Section A.3.1 discusses in more detail, we model the decision of a driver at the end of each trip to stop working or to continue working. After completing t trips in $h_{i,n,t}$ hours, driver i decides to end shift n when the disutility of effort for completing an additional trip outweighs the expected fare.[15] A key prediction of the model is that there are no daily income effects: cumulative daily earnings $y_{i,n,t}$ do not affect the decision to end a shift.

Letting $d_{i,n,t}$ indicate the decision to stop working, we test the prediction that daily income effects are inconsequential by expressing the probability that driver i ends shift t at

---

estimates.

[13]This corroborates the criticism due to Goette et al. [2004] that controlling for clock-hour effects removes much of the variation in earnings and quitting in the data from Farber [2005] because wages are highest precisely during the hours when many cabdrivers are required to end their shifts.

[14]A medallion-cab lessor may agree to provide gasoline to drivers at no more than $21 per shift (or $126 per week). Source: `http://www.nyc.gov/html/tlc/downloads/pdf/lease_cap_rules_passed.pdf`.

[15]This assumption follows Farber [2005, 2008] and Crawford and Meng [2011], who suggest that not explicitly modeling option value is behaviorally reasonable. Moreover, for convex disutility of effort, this trip-by-trip stopping rule is consistent with maximizing the static objective function as long as the wage rate $y_n'(h_n)$ does not increase too rapidly. The pattern in Figure A.1 suggests that the stopping model may not apply to PM shifts on weekends, a point that we revisit when discussing empirical results in Section 1.3.1.

trip n non-parametrically as

$$\Pr(d_{i,n,t} = 1) = \sum_j \Big[\big(f_j(h_{i,n,t}) + \gamma_j y_{i,n,t} + X_{i,n,t}\beta_j + \mu_{i,j} \big) \mathbf{1}_{\{h_{i,n,t} \in H_j\}} \Big] + \epsilon_{i,n,t}, \tag{1.1}$$

where X consists of controls for location, time, and weather which can potentially be related to variation in earnings opportunities from continuing to work; $f(h)$ is a function of work hours; μ absorbs differences in drivers' baseline stopping tendencies; and H_j partitions the minutes of the shift into intervals to allow a time-varying relationship between each of the covariates and the probability of stopping. The parameters γ_j capture the effect of an additional dollar on the probability of stopping during the j^{th} time interval under the assumption that cumulative daily earnings are uncorrelated with unobserved determinants of the value of stopping (such as effort or fatigue) or the value of continuing (such as future earnings opportunities) conditional on the full set of time-varying covariates, which Section 1.3.4 discusses in more detail. The model predicts that $\gamma_j = 0$ for all j, i.e., that the decision to end a shift is unrelated to cumulative daily earnings. Incorporating positively autocorrelated hourly wages as in Figure 1.5 would yield the prediction that $\gamma_j \leq 0$, as greater recent earnings may indicate higher continuation values.

By interacting all covariates with a fine partition of the minutes in the shift, the regression does not impose functional-form assumptions that constrain the relationship between stopping and its potential determinants. For instance, this enables the relationship between hours and the probability of stopping to be driver specific, whereas a parametric model with driver fixed effects would force that for any pair of drivers one of them has a uniformly higher or lower predicted probability of stopping at the end of any given trip conditional on the other covariates. As another example, a standard fixed-effects model might suggest that drivers are more likely to stop at 4 PM, when it rains, or when a trip ends near the taxi garage regardless of how many hours they have worked, whereas the non-parametric formulation allows the marginal effect of each variable on the probability of stopping to vary flexibly throughout the shift. Section A.3.2 conducts an empirical Monte Carlo exercise that validates this approach and demonstrates that stopping models estimated in prior work can yield spurious results.

1.3.1 Estimation of the Stopping Model

This section evaluates the prediction of the neoclassical model that cumulative daily earnings do not affect labor supply decisions. We partition the shift into 10-minute intervals and estimate the stopping model using Equation (1.1).[16] Table 1.2 presents the results, with each row corresponding to a more comprehensive set of controls than the previous one. To interpret the magnitude of the income effect, the table reports the marginal effect of a 10 percent increase in cumulative earnings on the probability of ending a shift at 8.5 hours, which is approximately the median stopping time. The estimates in column (1) use variation in earnings conditional on an extensive set of covariates that capture the value of stopping (hours worked so far on the shift) and the value of continuing (expectations about future earnings possibilities). Columns (2) and (3) contain estimates from an alternative estimation strategy that uses distance between pick-up and drop-off locations to instrument for earnings, which we discuss further in Section 1.3.4.

All specifications consist of controls for minutes spent working, including indicators for the number of minutes with passengers in each hour. The specification in row 1 with no additional controls shows that higher cumulative daily earnings are associated with greater stopping probabilities, contrary to the prediction of the neoclassical model. Row 2 shows that using within-driver variation in earnings only strengthens the estimated effect. This contrasts with the results from Farber [2005], in which the positive effect of cumulative daily earnings on the probability of ending a shift becomes insignificant after accounting for interdriver differences in stopping probabilities. Section 1.3.3 discusses and compares the magnitudes in more detail.

In the remaining specifications, we include precisely measured controls to account for additional factors that can potentially affect labor supply and find that the income effect persists. To address the possibility that a driver is more likely to end a shift when a trip

[16] All of the results are unchanged if we instead partition the shift into 30-minute intervals.

Table 1.2: *Stopping model estimates: Income effect at 8.5 hours*

	(1) Baseline	(2) IV: GPS distance	(3) IV: Odometer distance
Dependent variable: Indicator for stopping			
Controlling for			
Hours	0.1442	0.2993	0.4516
	(0.0351)	(0.0459)	(0.0498)
& Drivers	0.6632	0.9145	0.9974
	(0.0282)	(0.0373)	(0.0354)
& Location	0.0930	0.1898	0.1789
	(0.0270)	(0.0360)	(0.0334)
& Time	0.3787	0.3987	0.4141
	(0.0266)	(0.0361)	(0.0330)
& Weather	0.3798	0.3983	0.4137
	(0.0266)	(0.0361)	(0.0330)
Mean stopping probability:	13.627%		
Number of drivers:	36, 900		

Note: This table reports in each cell an estimate of the percentage-point increase in the probability of ending a shift at 8.5 hours when cumulative earnings is 10 *percent higher. Each row specifies an additional set of controls and provides results under three different estimation strategies. Column (1) presents baseline estimates from Equation* (1.1)*, with H_j partitioning the shift into 10-minute intervals. Columns (2) and (3) instrument for cumulative earnings based on cumulative distance using GPS coordinates and odometer miles, respectively. All specifications control flexibly for minutes spent working, including indicators for the number of minutes with passengers in each hour. Location controls consist of neighborhood fixed effects and an indicator for being in the zip code where the cab must be returned interacted with hour of the day. Time controls include an interaction between the hour of day and day of week, the week of the year, and an indicator for federal holidays. Weather controls consist of the indicators for precipitation, wind speed, and temperature in the minute that a trip ends. Drivers denotes fixed effects for 36,900 unique driver's license numbers. Standard errors reported in parentheses are adjusted for clustering at the driver level.*

ends in a convenient location (e.g., near the driver's home, or near a location where the cab can be transferred to another driver), the specification in row 3 includes indicators for the 195 Neighborhood Tabulation Areas (NTA) in NYC where a trip may end and an indicator for being in the zip code where the cab must be returned interacted with hour of the day.[17] Figures 1.1 and 1.4 show a systematic pattern that many shifts end in the early evening, coinciding with a period of higher average wages, which reflects an institutional feature of the market that AM-shift drivers transfer shared cabs to PM-shift drivers and are therefore unable to serve passengers during that time. The estimates in row 4 include an interaction between clock hour and day of week as well as indicators for week of year and federal holidays. Row 5 uses high-frequency data from five stations in the NYC area to account for variation in stopping due to weather conditions. Farber [2015] points out that rainfall reduces the number of cabs on the street due to added disutility of driving in the rain. More generally, since adverse weather conditions can affect the labor-supply decisions of cabdrivers, we include the following indicators measured in the minute when a trip ends: precipitation, wind speed on the Beaufort scale, temperature above 80 degrees Fahrenheit, temperature below 30 degrees Fahrenheit. Under the full set of controls, a 10 percent increase in cumulative earnings corresponds to about a 3 percent increase in the probability of ending a shift (0.4 percentage-point increase relative to a baseline stopping probability of 13.6 percent) at 8.5 hours. Across all specifications, the data reveal a clear pattern of significant labor-supply reductions in response to cumulative daily earnings.

Figure 1.6 shows using the full specification (with driver fixed effects and controls for location, time, and weather) that the magnitude of the income effect from Table 1.2, evaluated at 8.5 hours of work, persists throughout the shift. The figure plots the income effect and probability of stopping every thirty minutes over a five hour period, roughly corresponding to the 10th and 90th percentile of the distribution of stopping times. As the average stopping probability varies from 4 percent to 27 percent, the magnitude of the

[17] We map the GPS coordinates of the drop-off location of each trip to the NTA, which is an aggregation of the census tracts. Source: `http://www.nyc.gov/html/dcp/html/census/nyc_cff_faqs.shtml`.

Figure 1.6: *Stopping model estimates: Income effect throughout the shift*

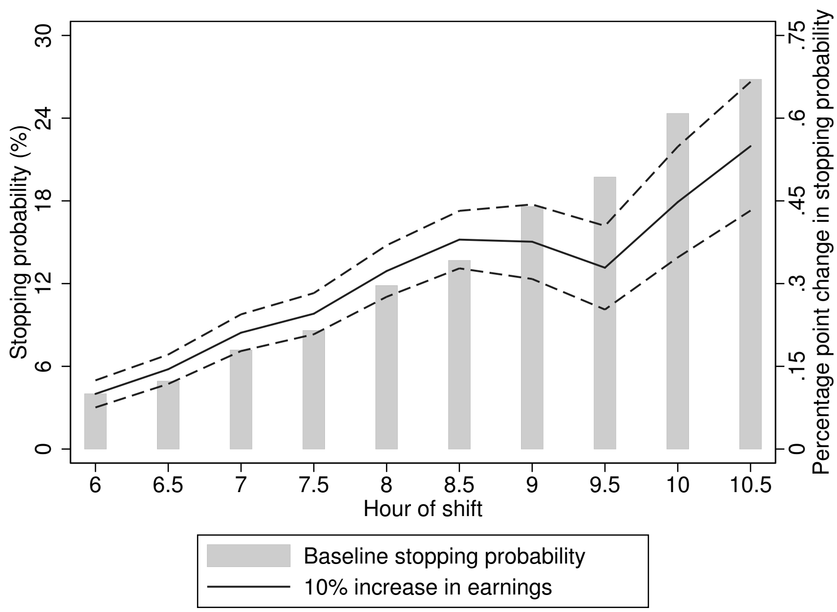

Note: The bars, corresponding to the scale on the left, show the probability that a driver ends a shift at the specified number of hours. The solid lines, corresponding to the scale on the right, depict the marginal effect of an additional 10 *percent in earnings on the probability of stopping at various times throughout the shift. Estimates obtain from Equation* (1.1) *with controls for location, time, and weather (see Table 1.2 for details) and fixed effects for 37,460 drivers. The dashed lines represent the 95-percent confidence interval, with standard errors adjusted for clustering at the driver level.*

income effect increases accordingly. Figure A.7 plots the percent change in the probability of stopping estimated on four separate groups of shifts: day-weekday shifts, day-weekend shifts, night-weekday shifts, and night-weekend shifts.[18] While the day shifts and the night-weekday shifts exhibit positive income effects consistent with our estimates from the full sample, the night-weekend shifts stands out with significant negative magnitudes. As Section A.3.1 points out, the trip-by-trip stopping model relies on the assumption that the option value of continuing to drive is sufficiently small (or that drivers ignore option value), and the pattern in Figure A.1 suggests that this assumption may not be reasonable for night-weekend shifts, when wages rise substantially and predictably over time. This observation explains a potential discrepancy with the results in Farber [2015] that hours of work are roughly unaffected by income during night shifts.

1.3.2 The Role of Timing

This section provides a test of the fungibility of money within a shift by evaluating whether daily labor supply responds to the *timing* of earnings. The previous section tests for daily income effects but implicitly makes a standard economic assumption of fungibility, which in this case entails that the effect of income on the probability of stopping at any point during the shift depends only on cumulative earnings and not on how recently the dollars are received within a shift. We relax the assumption that the probability of stopping does not depend on the *timing* of earnings by augmenting Equation (1.1) to express the probability of stopping as

$$\Pr(d_{i,n,t} = 1) = \sum_j \left[\left(f_j(h_{i,n,t}) + \gamma_{j,k} \sum_k y_{i,n,t,k} + X_{i,n,t}\beta_j + \mu_{i,j} \right) \mathbf{1}_{\left\{h_{i,n,t} \in H_j\right\}} \right], \qquad (1.2)$$

where $y_{i,n,t,k}$ denotes earnings accumulated in hour k of the shift. If money is fungible throughout the shift, then the impact of an additional dollar on the probability of ending a shift would be independent of when the dollar is earned (i.e., that $\gamma_{j,k}$ is independent

[18] Following Farber [2015], we define a day shift as one that starts between 4 AM and 10 AM and a night shift as one that starts between 2 PM and 8 PM. We classify night shifts on Friday and Saturday as well as day shifts on Saturday and Sunday as weekend shifts.

Figure 1.7: *Stopping model estimates: Income effect at 8.5 hours—Timing pattern*

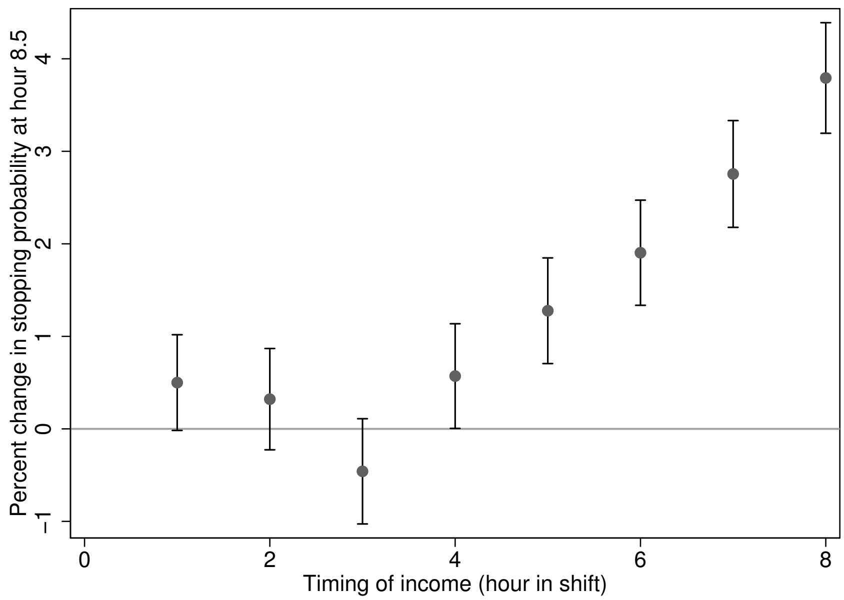

Note: The figure depicts the percent change in the probability of ending a shift at 8.5 hours in response to a \$10 increase in earnings accumulated at different times in the shift. Estimates obtain from Equation (1.2) *with controls for location, time, and weather (see Table 1.2 for details) and fixed effects for 36,900 drivers.*

of k). A hypothesis based on the pattern of positively autocorrelated hourly wages in Figure 1.5 would be that $\gamma_{j,k}$ is decreasing in k since greater recent earnings indicate higher continuation values.

Figure 1.7 plots the estimated effect of an additional \$10 earned at various times in the shift on the probability of ending a shift at 8.5 hours using the full set of controls. The estimates provide evidence against fungibility: the effect that an additional dollar has on labor supply depends on how recently the dollar is received within a shift. The effect of an additional dollar on the probability of stopping is greater if the dollar is earned more recently. The magnitude is substantial, with an additional dollar accumulated one hour earlier increasing the probability of ending a shift by seven times more than an additional dollar accumulated five hours earlier. Based on the positive autocorrelation in earnings

Figure 1.8: *Stopping model estimates: Income effect throughout the shift—Timing pattern*

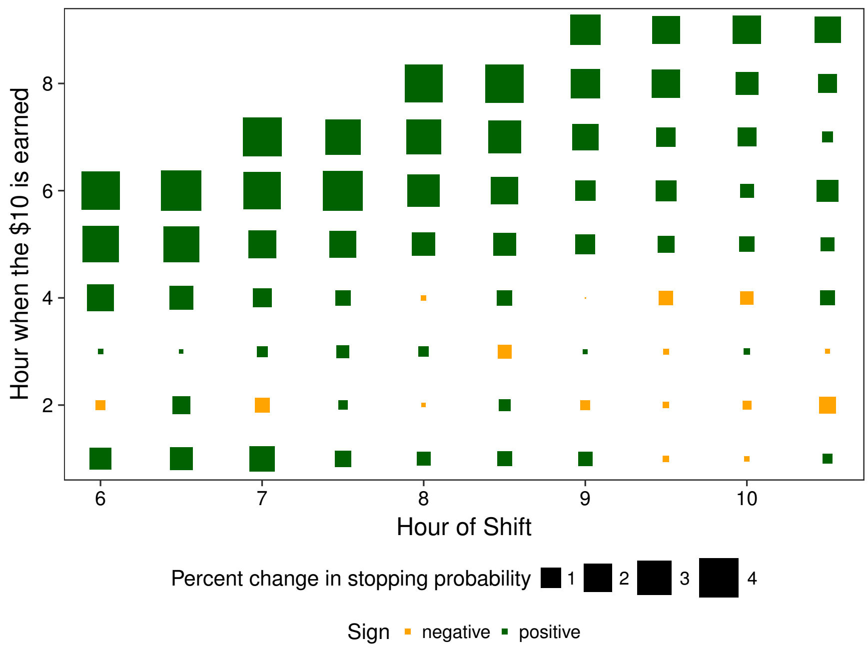

Note: The figure depicts the effect of an additional $10 in earnings accumulated at different times in the shift (vertical axis) on the probability of stopping at various times throughout the shift (horizontal axis). Each square has area proportional to the corresponding percent change in the probability of stopping. Estimates obtain from Equation (1.2) *with controls for location, time, and weather (see Table 1.2 for details) and fixed effects for 37,460 drivers.*

from Figure 1.5, higher recent earnings should, if anything, be associated with a higher value of continuing to work. However, holding total earnings fixed, the probability of ending a shift at the end of a given trip depends on the path of earnings throughout the shift in the opposite direction of this prediction: a driver is more likely to stop working after accumulating earnings at a higher rate toward the end of a shift. While additional earnings accumulated within the first three hours of the shift do not significantly impact the probability of stopping at 8.5 hours, drivers appear to respond strongly to income earned more recently in the shift by reducing labor supply.

Figure 1.8 shows that the violation of fungibility from Figure 1.7, evaluated at 8.5 hours of work, persists throughout the shift. The columns of the figure correspond to different

Figure 1.9: *Stopping model estimates: Income effect at 8.5 hours—Timing pattern by shift start hour*

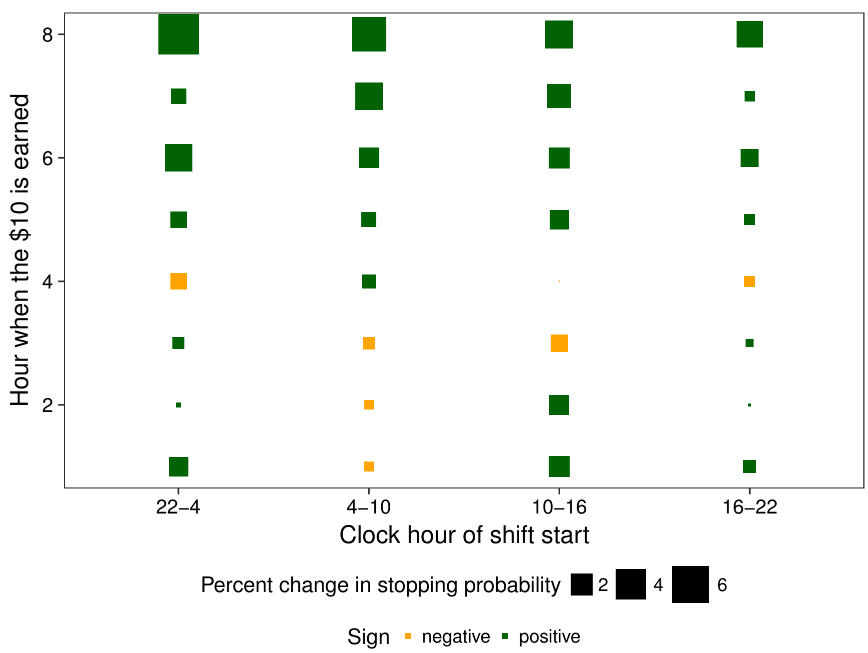

Note: The figure depicts the effect of an additional $10 in earnings accumulated at different times in the shift (vertical axis) on the probability of stopping at 8.5 hours for groups of shifts that start in different hours (horizontal axis). Each square has area proportional to the corresponding percent change in the probability of stopping. Estimates obtain from Equation (1.2) *with controls for location, time, and weather and driver fixed effects (see Table 1.2 for details).*

times during the shift between hour 6 and hour 11. Within a given column (i.e., fixing a point in time during the shift), each row depicts the effect of an additional $10 accumulated in a particular hour on the probability of stopping: additional earnings accumulated early in the shift do not substantially affect the probability of stopping, but drivers are consistently more likely to end a shift in response to more recent earnings. The result that the timing pattern holds throughout the shift suggests an interpretation based on recency. The estimates in Figure 1.9 further confirm that time-of-day effects do not drive these patterns. The figure separates shifts into four groups based on start hour: morning (4 AM–10 AM), afternoon (10 AM–4 PM), evening (4 PM–10 PM), and night (10 PM–4 AM). Each column depicts the effect of an additional $10 earned at various times in the shift on the probability of ending a shift at 8.5 hours for a different group of shifts, demonstrating a timing pattern consistent with our estimates from the full sample.

1.3.3 Discussion of Magnitudes

To provide a better understanding of the size of the income effect, consider the estimates in Table 1.2. A 10 percent increase in cumulative earnings (an average of $26.67) corresponds to a 2.8 percent increase in the probability of ending a shift at 8.5 hours under the baseline specification with the full set of controls. If the additional earnings arrive in the eighth hour of the shift, then the estimates in Figure 1.8 imply a 10.12 percent increase in the probability of stopping. For comparison, an additional 10 minutes of work (i.e., the median trip duration) increases the probability of ending a shift by 5.96 percent.

Most investigations of daily labor-supply decisions focus on the question of how workers adjust labor supply in response to wage changes. However, the few studies that examine daily income effects generally find that cumulative daily earnings do not affect labor supply.[19] Farber [2005] finds an insignificant effect of earnings, though the point estimate

[19] Farber [2005] analyzes income effects using observational data, while Andersen et al. [2014] and Dupas et al. [2016] use field experiments. Andersen et al. [2014] conduct a field experiment on vendors in an Indian open-air market and find that vendors do not adjust labor supply in response to an overpayment early in the day. As a possible resolution to the apparent inconsistency between the results of our observational study and their field experiment, our Figure 1.7 suggests that workers may adjust labor supply disproportionately in

implies that a 10 percent increase in cumulative earnings corresponds to an increase in the probability of ending a shift of 1.2 percent, which is less than half of the magnitude we find.[20]

We find not only that cumulative daily earnings influence labor-supply decisions but also that the income effect depends on recency. The impact of a dollar earned in the eighth hour of a shift is seven times greater than that of a dollar earned four hours earlier, which implies a violation of fungibility of money earned within a shift. Our results complement existing studies on mental accounting, which tends to focus on how consumers treat money from different sources, by demonstrating violations of fungibility based on earnings accrued at different times within a single day.[21]

1.3.4 Alternative Explanations

The evidence shows substantial changes in labor supply in response to small changes in wealth, and the magnitude of the reduction in labor supply depends on the timing of changes in wealth. This section addresses potential challenges to the modeling assumptions by considering the possibility that cumulative earnings are correlated with unobserved determinants of the stopping decision such as effort or fatigue, or that cumulative earnings convey information about future earnings opportunities. We also assess whether the relationship between cumulative earnings and stopping arises due to other factors such as option value, liquidity constraints, and inexperience.

---

response to recent changes in cumulative earnings, resulting in labor-supply responses that depend not only on the amount of income but also on the timing of income. Dupas et al. [2016] find evidence of income effects in a sample of Kenyan bicycle-taxi drivers who have daily cash needs and conduct a field experiment to evaluate the effect of unexpected windfalls on daily labor-supply decisions.

[20] Table 5 in Farber [2005] reports that an additional dollar increases the probability of ending a shift by 0.011 percentage points at 8 hours under the full set of controls. With a mean income of \$161.33, an additional 10 percent in earnings corresponds to a $16.13 \cdot 0.011 \approx 0.18$ percentage-point increase in the probability of stopping relative to a baseline of 14.67 percent. The income effect is neither significantly different from zero nor significantly different from the point estimate we find in Table 1.2.

[21] Hastings and Shapiro [2013], for example, find that households respond over 15 times more to a reduction in real income through an increase in the price of gasoline than an equivalent reduction in income.

Table 1.3: *Stopping model estimates: Income effect at 8.5 hours—Subsamples*

	(1) Night weekday	(2) Medallion owners	(3) Top decile experience
Panel A			
Income effect	0.3564	0.5421	0.4625
	(0.0473)	(0.1548)	(0.0805)
Panel B			
Income in hour 2	0.0725	-0.1175	-0.0130
	(0.0742)	(0.2351)	(0.1236)
Income in hour 4	0.0077	0.0282	0.3062
	(0.0717)	(0.2269)	(0.1284)
Income in hour 6	0.2645	0.2363	0.3309
	(0.0732)	(0.2389)	(0.1267)
Income in hour 8	0.3270	0.5714	0.5580
	(0.0752)	(0.2246)	(0.1335)

Note: Panel A reports estimates from Equation (1.1) *of the percentage-point increase in the probability of ending a shift at 8.5 hours when cumulative earnings is* 10 *percent higher. Panel B reports estimates from Equation* (1.2) *of the percentage-point change in the probability of ending a shift at 8.5 hours in response to a $10 increase in earnings accumulated at different times in the shift. The columns correspond to different sample restrictions: (1) trips on Friday and Saturday after 5* PM*, (2) cabdrivers who operate exactly one cab and no other driver shares that cab, and (3) the latest* 10 *percent of shifts for drivers with over* 100 *shifts. The control variables consist of the full set from Table 1.2. Standard errors reported in parentheses are adjusted for clustering at the driver level.*

Effort

We interpret the increase in the probability of ending a shift in response to higher cumulative earnings as evidence of a daily income effect exhibited by a reduction in labor supply. A potential concern with this conclusion is that labor supply consists of multiple dimensions, some of which are unobserved. If the reduction in hours coincides with an increase in the intensity of work, then the overall effect on labor supply would be unclear. We suggest three ways to address this: by using an instrumental-variable (IV) strategy, by constructing a proxy for effort, and by analyzing income effects and recency effects in the earlier hours of the shift.

First, columns (2) and (3) of Table 1.2 present IV estimates to address a possible correlation between cumulative daily earnings and unobserved determinants of the decision to end a shift such as effort or fatigue.[22] Since a cabdriver legally cannot refuse passengers because of destination, we instrument for cumulative earnings using the cumulative distance between pick-up and drop-off locations.[23] The magnitude of the income effect does not significantly change, measuring distance either by GPS coordinates (column 2) or by odometer miles (column 3).[24] Additionally, Section A.4.1 provides an alternative IV strategy that exploits variation in income due to tips. To the extent that tips may be indicative of effort, a larger magnitude of the IV estimate might suggest an upward bias in the magnitude of the income effect. However, the IV analysis results in somewhat smaller estimates, consistent with our interpretation of the effect of cumulative earnings on the probability of stopping as representing an income effect. Moreover, estimating Equation (1.1) using either the distance or tip instruments (see Figure A.4) produces the same timing pattern of the income effect as

[22] Although we do not report first-stage regressions, the F-statistics are sufficiently large for all of the IV specifications.

[23] Camerer et al. [1997] argue based on a survey that driving passengers to a specific destination requires less effort than driving while searching for potential passengers.

[24] Since we control for the amount of time spent riding with passengers, variation in distance can reflect differences in driving speed. To alleviate concerns about a possible correlation between driving speed and unobserved determinants of the decision to end a shift, Section A.4.1 shows that the results hold using trips that stay within the dense streets of Manhattan, where variation in driving speed plausibly arises due to traffic conditions unrelated to the driver's decisions to exert additional effort.

in Figure 1.7, suggesting that the income effect does not reflect recent effort expended.

Second, although our dataset does not contain a direct measure of effort, we use as a proxy how quickly a driver finds the next passenger. Drivers spend 38 percent of their working hours searching for passengers during shifts in the bottom decile of earnings, compared with only 35 percent of working hours during shifts in the top decile. Part of the relationship reflects the mechanical fact that drivers likely earn more money during shifts that have a higher fraction of the time riding with passengers. Despite this mechanical effect, the data show only a weak relationship between earnings and the fraction of time spent searching (correlation of -0.10), suggesting a small role, if any, for adjustments along the effort dimension.[25]

Third, Figures 1.6 and 1.8 show consistent and sizable income effects and recency effects *throughout* the shift, which poses a difficulty for fatigue-based explanations. An additional 10 percent in earnings corresponds to an increase in stopping probability of at least 2.5 percent with a significantly stronger response to recent earnings, even in the early hours of a shift. Furthermore, to the extent that drivers face an increasing marginal disutility of effort, we would expect much larger magnitudes of the income effect in the later hours of a shift (e.g., after working 10 hours compared to 8 hours) if effort poses a confound for estimating the effect of cumulative earnings. Using a non-parametric formulation mitigates the scope for the estimated income effect to reflect a response to fatigue, as Equation (1.1) allows for a driver-specific relationship between work hours and the probability of stopping and allows for a flexible relationship between the effect of work conditions on the probability of stopping.[26]

[25]In addition, Figure 1.2 provides suggestive evidence of low returns to effort in terms of tips. Although the figure does not condition on any trip characteristics, the mass of points at a fare of exactly $52 represents trips between Manhattan and JFK International airport and shows substantial variation in tips, ranging from $1 to $20.

[26]Specifically, note that all controls and driver fixed effects interact with a fine partition of the minutes in the shift. This allows us to account for the effects of fatigue due to work hours and driving conditions as well as to accommodate neoclassical patterns of behavior such as quitting after a target number of hours without incorrectly attributing these to effects of income.

Learning about future earnings

Another potential concern with interpreting the effect of earnings on stopping behavior is that accumulated earnings may convey additional information about future opportunities, either within the same shift or across shifts.

If higher cumulative earnings or higher recent earnings indicate lower expected earnings from continuing conditional on all the covariates, then the estimated relationship between earnings and quitting would overstate the income effect. The pattern in Figure 1.5, however, suggests the opposite.

Likewise if higher earnings correlate with plentiful opportunities on the next day, then drivers may engage in intertemporal substitution, quitting during times of high earnings to conserve energy for the next shift. The evidence in Section A.4.2 suggests a limited role for this channel, as earnings do not appear predictive of market conditions on subsequent days.

Option value

The model in Section A.3.1 posits that at the end of each trip, drivers decide whether to end the shift or continue working for one more trip. In practice, however, a driver who believes that the wage will rise later in the shift might decide to continue working, and driver who explicitly solves the dynamic optimization problem might appear to have a low probability of ending a shift in response to low cumulative daily earnings. Neglecting option value would be problematic if the rate of increase in the wage exceeds that of the monetary equivalent of the disutility of effort (see Section A.3.1).[27] Column (1) of Table 1.3 restricts the analysis to trips on Friday and Saturday after 5 PM, when the typical wage profile is nonincreasing (see Figure A.1), and reports a positive relationship between earnings and stopping that does not significantly differ from the estimate using the full sample.

[27] As Figures A.1 and 1.4 show, wages typically do not rise sufficiently rapidly to justify concerns about this assumption, with the exception of night-weekend shifts which we highlight in Figure A.7.

Liquidity constraints

Liquidity constraints often pose a challenge for identifying income effects. Johnson et al. [2006] and Parker et al. [2013], for example, find that household consumption exhibits excess sensitivity to small changes in wealth due to fiscal stimulus, but their results suggest an important role for liquidity constraints. Our work, by contrast, detects persistent income effects in labor-supply decisions at a high frequency, which limits the plausibility of an explanation based on liquidity constraints. Dupas et al. [2016] argue that bicycle-taxi drivers in Kenya set income targets as a commitment device to exert enough effort to meet daily needs, also pointing toward liquidity constraints. Such explanations in our setting would necessitate a consistent inability of NYC cabdrivers to smooth consumption across days.[28] The result that drivers react differently to earnings accrued over different hours of the shift would be particularly difficult to rationalize based on liquidity constraints.

Although such effects are less plausible in our setting, we replicate our analysis on a sample of drivers for whom liquidity constraints likely do not bind. Specifically, we estimate the stopping model restricted to owner-drivers, as such drivers possess enough borrowing power or wealth to purchase an independent medallion to operate a taxicab.[29] Although our data do not include information on ownership, we classify a driver as an owner-driver if (i) the driver operates exactly one cab, and (ii) no other driver shares that cab.[30] The estimates in column (2) of Table 1.3 suggest that liquidity constraints do not confound the income effects we observe.

[28]Camerer et al. [1997] argue that this seems unlikely because almost all lease-drivers pay their weekly fees in advance, and fleet drivers pay their daily fees at the end of the day or can pay late.

[29]The average price of an independent medallion in 2013 was approximately $967,000. Source: `http://www.nyc.gov/html/tlc/downloads/pdf/2014_taxicab_fact_book.pdf`.

[30]Our classification yields a subsample of owner-drivers, as we exclude those who lease to another driver. See Farber [2015] for additional institutional details on regulations concerning medallions in NYC.

Experience

A hypothesis based on findings in related settings would be that the positive relationship between earnings and stopping reflects a failure to optimize by inexperienced drivers. Camerer et al. [1997] present evidence that more experienced drivers exhibit more positive wage elasticities of labor supply, which Farber [2015] corroborates. Recent work by Haggag et al. [2017] documents using the TPEP data from 2009 that productivity differences between new and experienced drivers vanish after 17 to 62 shifts (depending on the difficulty of the situations). Given that performance improves quickly with experience, drivers might also learn to supply labor more efficiently by ignoring daily earnings.

To consider the possibility of heterogeneity in income effects based on experience, column (3) of Table 1.3 restricts to the latest 10 percent of shifts in the sample for drivers with over 100 shifts. We find similar magnitudes of income effects as drivers gain more experience, with the full set of results that includes income effects at all deciles of experience in Table A.4.

Measurement error

At least two issues arise when measuring work hours in this setting. First, the data do not contain an explicit measure of break times. Second, the data do not distinguish between a driver who ends a shift immediately after dropping off their last passenger and a driver who spends time searching for another fare unsuccessfully. Section A.4.5 shows that accounting for the first issue does not change the magnitude of the income effect, and the second issue, if anything, biases the results against finding income effects.

1.4 Discussion

This paper documents violations of fungibility for money earned at different times. Contrary to the neoclassical prediction, drivers treat an additional dollar on one day as being different from a dollar on another day, resulting in a daily income effect whereby labor supply decreases in response to accumulated daily earnings. Moreover, drivers are more likely to stop working in response to earnings accumulated more recently within the same day. These facts taken together are inconsistent with a notion of income targeting in which drivers reduce labor supply after earning a particular amount each day. However, models of reference-dependent preferences can explain both facts with a reference level that adjusts within the day. As reference levels take time to adjust in response to recent changes in expectations, earlier experiences within the day become incorporated into the reference point, thereby moderating the income effect, while recent experiences induce stronger behavioral responses. Our findings provide field evidence for reference dependence and reference-point adjustment that persist in the face of experience and high stakes in a familiar, recurring setting.

Chapter 2

Asymmetric Reference Point Adjustment[1]

2.1 Introduction

A substantial body of work documents that reference points have an influence on behavior. While the predictions of reference dependence depend on how reference points are formed and how they adjust, reference points are typically assumed to coincide with rational expectations. This chapter investigates how reference points adjust. Our results indicate that reference points tend to adjust more readily in the direction of gains rather than losses, contrary to the canonical view of reference points based on rational expectations.

The previous chapter provides evidence of excess sensitivity of labor-supply decisions to daily earnings, with a magnitude and timing pattern that a neoclassical model of labor supply cannot explain. This chapter estimates and compares alternative models which can potentially generate the income effects that the data reveal. In Section 2.2, we present a neoclassical model of labor supply as well as three behavioral models: income targeting, loss aversion, and salience. Section 2.3 extends the loss aversion and salience models to incorporate reference point adjustment, which we demonstrate is crucial for understanding

[1] Co-authored with Linh T. Tô

the violation of fungibility shown in Section 1.3.2. The structural-estimation exercise in Section 2.4 yields two conclusions. First, the estimates reject the neoclassical model against several alternatives. Second, the models with adaptive reference points and asymmetric adaptation provide a better fit of the data than the corresponding models with static reference points.

Existing work that uses structural estimation to test a behavioral theory largely focuses on the objective of testing the null hypothesis of the neoclassical model, with a rejection of the neoclassical model interpreted as evidence in favor of a particular behavioral alternative. Section 2.5 departs from that paradigm by using the structural estimates to compare the models' implications with the goal of adjudicating between the behavioral mechanisms. First, we compare the magnitude and timing pattern of the income effect predicted by each of the models. Second, we derive and test an additional prediction about moments of the data that the models were not designed to fit, namely how the magnitude of the income effect changes around the reference point.

2.2 Behavioral Models of Daily Labor Supply

The neoclassical model in Section A.3.1 posits that the marginal utility of lifetime income—and hence labor supply—does not vary in response to small, within-day changes in wealth. We assume the objective function for a driver with earnings I_t and hours of work H_t takes the form

$$\begin{aligned} v(I_t, H_t) &= v_I(I_t) + v_H(H_t) \\ &= I_t - \frac{\psi}{1+\nu} H_t^{1+\nu}, \end{aligned} \tag{2.1}$$

where ψ parameterizes the disutility of work and ν is the elasticity parameter.

The stopping decision depends on the next trip's expected fare $\mathbb{E}_t[f_{t+1}]$ and duration $\mathbb{E}_t[h_{t+1}]$. A driver with earnings I_t and hours of work H_t at the end of trip t decides to end a shift if the driver expects that completing an additional trip results in lower utility:

$$\mathbb{E}_t[v(I_{t+1}, H_{t+1})] - v(I_t, H_t) + \varepsilon_t < 0, \tag{2.2}$$

where $I_{t+1} = I_t + \mathbb{E}_t[f_{t+1}]$, $H_{t+1} = H_t + \mathbb{E}_t[h_{t+1}]$, and ε_t are independent and normally distributed with mean zero and variance σ^2. While the objective function depends explicitly on cumulative daily earnings I_t, quasi-linearity implies that cumulative daily earnings do not affect the decision to end a shift.

To explain the results in Section 1.3, a model must allow for non-trivial within-day changes in the marginal utility of income. We analyze three different behavioral distortions that can create such income effects: income targeting, expectations-based loss aversion, and salience. Each of the models we consider makes an implicit assumption of narrow bracketing, that decision makers evaluate utility at the daily level.

Income targeting

This section formulates an ad-hoc model of daily income targeting in which the marginal utility of income declines substantially around the level of average daily earnings. As Camerer et al. [1997] suggest, a model in which drivers dislike falling short of their target more than they like exceeding it provides one possible explanation for the concavity in utility necessary for daily income effects.

The following objective function introduces a parameter $\alpha \geq 0$ to allow for a change in marginal utility at the target T:

$$v^{IT}(I_t, H_t) = v(I_t, H_t) + \alpha(I_t - T)\mathbf{1}_{\{I_t < T\}}. \tag{2.3}$$

We interpret this as a model of reference-dependent preferences with a fixed reference point.[2] The model implies a coefficient of loss aversion—the ratio between utility from losses and gains—of $1 + \alpha$. Farber [2015] tests the neoclassical model against this alternative

[2]DellaVigna [2009] describes a model of this form as a simplified version of prospect theory [Kahneman and Tversky, 1979] that incorporates reference dependence and loss aversion without diminishing sensitivity and probability weighting.

hypothesis.

Expectations-based loss aversion

The theory of reference-dependent preferences due to Kőszegi and Rabin [2006] provides the leading explanation in the literature for the mixed evidence on behavior in daily labor-supply decisions.[3] In their model, utility depends not only on a standard outcome-based consumption component but also on a gain-loss component which captures how decision makers assess choices relative to a reference point. Kőszegi and Rabin [2006] assume that rational expectations endogenously determine the reference point, consistent with laboratory experiments and field evidence in various settings.[4]

Our primary analysis of loss aversion involves a simplified version of the model from Kőszegi and Rabin [2006], in which the objective function of the driver takes the following form:

$$v^{LA}(I_t, H_t) = (1-\eta)v(I_t, H_t) + \eta \sum_{x \in \{I,H\}} n(x_t \,|\, x_t^r), \tag{2.4}$$

where I^r and H^r denote the reference levels for income and hours (i.e., the driver's expected earnings and hours for the shift), and the gain-loss utility is given by

$$n(x \,|\, x^r) = \big(\mathbf{1}_{\{x > x^r\}} + \lambda \mathbf{1}_{\{x < x^r\}}\big)\big(v_x(x) - v_x(x^r)\big),$$

where η determines the relative weight on gain-loss utility, and $\lambda \geq 1$ parameterizes the degree of loss aversion. This coincides with the neoclassical model when there is no difference in utility from gains and losses (i.e., $\lambda = 1$ or $\eta = 0$). Despite adding two parameters to Equation (2.1), the model consists of only one additional degree of freedom since η and λ are not separately identifiable: behavior depends only on the ratio between

---

[3] The survey by DellaVigna [2009] discusses this, but also see more recent work by Crawford and Meng [2011].

[4] For example, see lab evidence from Abeler et al. [2011], Gill and Prowse [2012], Karle et al. [2015], and Sprenger [2015] as well as field evidence from Post et al. [2008], Card and Dahl [2011], Ericson and Fuster [2011], and Pope and Schweitzer [2011].

utility from losses and gains, namely

$$\begin{aligned} L &= \frac{(1-\eta)+\eta\lambda}{(1-\eta)+\eta} \\ &= 1+(\lambda-1)\eta. \end{aligned}$$

This formulation makes two simplifying assumptions about the gain-loss component of utility. First, the reference levels represent a driver's point expectations for income and hours on a given shift, abstracting from stochasticity whereby the reference levels represent the full distribution of potential earnings and hours for that particular shift. Second, the piecewise-linear gain-loss function rules out diminishing sensitivity, the observation that decision makers experience smaller marginal changes in gain-loss sensations further away from their reference levels. While these assumptions follow the implementation from Crawford and Meng [2011], we relax both of them in Sections 2.4 and 2.5.

The Kőszegi and Rabin [2006] model features two important differences from the income-targeting model. First, drivers experience losses from working longer than their "hours target," analogous to the losses from earning less than their "income target." In other words, the separable gain-loss function indicates that drivers exhibit loss aversion not only over income but also over effort, with the same coefficient of loss aversion λ on both dimensions. Second, utility depends on expectations, as the reference points vary across drivers and across days. Section 2.3 discusses expectations and the specification of the reference point in more detail.

Salience

We adapt a model of salience based on Bordalo et al. [2015] to daily labor-supply decisions. The model combines two elements: (i) an evoked set determines the choice context and hence the salience of each attribute (income and hours), and (ii) decision makers place greater weight on the more salient attribute.[5]

[5]Bordalo et al. [2012] develop a theory of choice under risk in which decision makers overweight states that are more salient. Bordalo et al. [2013] extend this concept to riskless choice among goods with multiple

In this model, context influences decisions by distorting the relative weights that a driver places on income and leisure. A decision problem brings to mind an evoked set of options, each with an associated level of availability. The availability-weighted average of the options comprising the evoked set determines the normal levels of income and hours. The extent to which an attribute varies within the evoked set relative to the normal level determines the salience of that attribute. Drivers place greater weight on the more salient attribute—income or hours—of their decision problem. In describing the components of the model more formally, we start with the salience distortions, taking the normal levels of income and hours as given, and then address how to determine the normal levels.

The objective function consists of a weighted sum

$$v^S(I_t, H_t) = \sum_{x \in \{I,H\}} \frac{w(\sigma(x_t, x_t^n), \delta)}{\sum_{y \in \{I,H\}} w(\sigma(y_t, y_t^n), \delta)} v_x(x_t), \tag{2.5}$$

where the relative weight $w(\sigma(x, x^n), \delta)$ on the utility for a given attribute increases in the salience $\sigma(x, x^n)$ of that attribute, x^n denotes the normal level of the attribute, and $\delta \leq 1$ parameterizes the degree of distortion. We adopt the continuous salience weighting function from Bordalo et al. [2013]:

$$w(\sigma(x, x^n), \delta) = \frac{[1 + \sigma(x, x^n)]^{1-\delta}}{2},$$

where the case $\delta = 1$ embeds the neoclassical model without context dependence, and the salience function $\sigma(\cdot, \cdot)$ is a symmetric and continuous function that satisfies ordering and diminishing sensitivity conditions. The *ordering* condition requires that moving an attribute further apart from the normal level increases its salience. *Diminishing sensitivity* expresses the idea that increasing the normal level renders a given difference between an attribute and the normal level less salient.[6] Section 2.3 discusses the importance of these

attributes (e.g., quality and price), where consumers place more weight on more salient attributes, but take the evoked set as exogenous. Our formulation of salience follows Bordalo et al. [2015] which models the evoked set explicitly.

[6]Note that the salience model defines diminishing sensitivity relative to zero, whereas the loss aversion model refers to diminishing sensitivity relative to the reference point.

two properties for explaining the pattern of income effects. For a continuous and symmetric salience function satisfying these properties, Bordalo et al. [2012, 2013] suggest

$$\sigma(x, x^n) = \frac{|x - x^n|}{|x| + |x^n|},$$

which Hastings and Shapiro [2013] also use in empirical work.

To complete the description of the model, we discuss how to specify the evoked set and availability, which determine the normal levels of income and hours. We assume that the evoked set consists of the choices, stop or continue, and that the availability a_t of stopping at the end of trip t depends on how much more the driver must earn to reach the income target.[7] We then define the normal level of an attribute as the availability-weighted average of the level of that attribute from stopping or continuing:

$$I^n = a_t I_t + (1 - a_t) I_{t+1}$$

$$H^n = a_t H_t + (1 - a_t) H_{t+1}$$

where a_t increases in I_t, decreases in I_t^r, and lies between 0 and 1. Intuitively, as earnings accumulate up to and beyond the income target, the decision to stop becomes more typical and thus comes to the top of the driver's mind. For estimation, we assume that the availability of stopping corresponds to the predicted probability of ending a shift based on I_t and I_t^r from a logistic regression.[8] We isolate the channel by which earnings influences stopping behavior through I_t and I_t^r, though in principle other factors could could be included as well, which we discuss in Section 2.5.

The idea that choice context influences how decision makers weight different attributes of a decision problem appears in a number of recent economic models (e.g., salience [Bordalo et al., 2013], focusing [Kőszegi and Szeidl, 2013], and relative thinking [Bushong et al., 2016]) which take the choice context as a degree of freedom called the evoked set, consideration

[7] This assumption is based on Bordalo et al. [2015], who posit that availability is a map from past experiences and objective probabilities into a weight that reflects what comes to the decision maker's mind.

[8] We obtain similar results if we use different functional forms.

set, or comparison set. These models conceptually distinguish between the choice set and the evoked set but equate them when deriving predictions. If the normal levels were an unweighted average of the elements in the evoked set as in [Bordalo et al., 2013], then assuming that the two sets coincide would fail to produce the pattern of income effects from Section 1.3 because doing so imposes fungibility: in such a model, the behavioral distortion only depends on cumulative earnings, independent of the timing of earnings. One way to proceed would be to make ad hoc assumptions about which additional options enter the evoked set.[9] Instead, we assume the evoked set consists of the choices stop and continue, but we put structure on the components of the evoked set using a notion of availability based on Bordalo et al. [2015].

2.3 Adaptive Reference Points

In this section, we discuss how each of the behavioral models can potentially account for the evidence in Section 1.3. With the exception of the income-targeting model, which takes each driver's target to be fixed across all shifts, the predictions depend crucially on the reference level. As a starting point, we take each driver's reference level for each shift to be their rational expectations of income and hours for that particular shift. Under this view, the models of expectations-based loss aversion and salience yield daily income effects but treat money as fungible within the shift. By allowing for reference points that adjust within a shift, both models can generate the violations of fungibility necessary to explain the timing pattern of the income effect.

Consider the case of reference points based on rational expectations that represent drivers' steady-state beliefs about earnings and hours. Under this view, reference points can be thought of as drivers' targets for income and hours formed at the daily level, as in

[9]To explain observed gasoline-grade choice in 2006–2009, Hastings and Shapiro [2013] estimate a model of salience in which the evoked set consists of the current choices along with the grades of gasoline at the national mean prices from one week earlier. To improve fit, they propose an extended salience model in which the evoked set consists of the current choices along with the three grades of gasoline at prices \$1.00, \$1.10, and \$1.20.

Crawford and Meng [2011]: reference points vary across but not within shifts. We denote these ex ante expectations of shift-level earnings and hours by I_0^r and H_0^r. The model based on Kőszegi and Rabin [2006] generates daily income effects through loss aversion, whereby the marginal utility of income decreases once earnings exceeds the reference point. The same mechanism applies in the income-targeting model. The model based on Bordalo et al. [2015] generates daily income effects through the diminishing-sensitivity property of salience: the driver perceives the value of an additional fare less intensely at higher levels of daily earnings. Each of these models, despite generating income effects, treats money as fungible within the shift. Given the static reference point, the behavioral distortions depend on cumulative daily earnings in a shift, without scope for recent earnings to have a stronger influence on stopping decisions.

Suppose reference points adjust instantaneously to new information about how much the driver will earn by the end of the shift. In this case, if hourly earnings exhibit no substantial within-day autocorrelation, the driver updates expectations about earnings for the shift by immediately incorporating any difference between realized and expected earnings. Predictions of the Kőszegi and Rabin [2006] model approach those of the neoclassical model as reference-point adjustment becomes instantaneous. Intuitively, since rational expectations about daily earnings fully adjust, deviations from expectations no longer bring cumulative daily earnings closer to or further from the reference point. While the salience model continues to predict income effects due to diminishing sensitivity, no timing pattern emerges because instantaneous adjustment does not create a distinction between earnings at different times.[10]

Stronger income effects in response to recent earnings requires a slow-adjusting reference point. Letting Δ_t denote the difference between realized and expected earnings in trip t, we model the reference point in this case as a convex combination of the lagged reference point

[10]With a positive autocorrelation of hourly earnings as in Figure 1.5, higher recent earnings should if anything affect the reference point more strongly and hence affect the stopping decision less strongly. This implies a higher probability of ending a shift in response to *less* recent earnings.

and the expectation that obtains under full adjustment:

$$I_t^r = \theta I_{t-1}^r + (1-\theta)\left(I_0^r + \sum_{\tau=1}^{t} \Delta_\tau\right), \tag{2.6}$$

where $0 \leq \theta \leq 1$, with $\theta = 1$ corresponding to a static reference point and $\theta = 0$ corresponding to a reference point that adjusts instantaneously. The recursive formulation produces a reference point based on lagged expectations:

$$I_t^r = I_0^r + \sum_{\tau=1}^{t} (1 - \theta^{t+1-\tau})\Delta_\tau.$$

The expression highlights that the reference point for income incorporates less recent earnings to a greater extent, consistent with the idea that reference points take time to adjust in response to recent changes in expectations. This generates a violation of fungibility under both loss aversion and salience. Under loss aversion, the gain-loss component of utility depends on the difference between earnings and its reference level, and the reference point adjusts to a lesser extent in response to more recent earnings. The same applies to availability in the salience model, reflecting the intuition that recent earnings bring stopping closer to the top of the driver's mind, which makes leisure relatively more salient due to the ordering property of salience.[11] The qualitative predictions of both models under a slow-adjusting reference point corresponds to the following intuition about reacting to surprises: unexpected earnings constitutes a surprise, but surprises wear out over time so that quitting depends to a greater extent on recent earnings.

[11]When the availability of stopping increases, the normal levels I^n and H^n move away from the levels of these attributes under continuing (I_{t+1} and H_{t+1}). By the ordering property, this increases the salience of both income and hours, with the salience of hours increasing to a greater extent because of convexity in the disutility of work hours.

2.4 Structural Estimation

We use the data on stopping decisions to estimate the models via maximum likelihood under alternative specifications of the reference point.[12] Each behavioral model produces a stopping rule analogous to Equation (2.2), resulting in likelihood functions of the following form:

$$\sum \log \Phi\left(\frac{v(I_t, H_t) - \mathbb{E}_t[v(I_{t+1}, H_{t+1})]}{\sigma}\right). \tag{2.7}$$

As a benchmark, we consider the case of static reference points based on rational expectations, proxying for expectations using the sample average of income and hours by driver and day of week (excluding the current shift) as in Crawford and Meng [2011]. Although we estimate the parameters jointly for each of the models, the following intuition describes the main sources of identification. The disutility of effort ψ is primarily identified by variation in work hours, and the elasticity parameter ν is primarily identified by variation in expected wages from continuing. The behavioral parameter in each model (α for income targeting, L for expectations-based loss aversion, and δ for salience) is primarily identified by variation in cumulative earnings, which explains why we do not estimate a single model that nests and tries to distinguish multiple behavioral mechanisms. For the purpose of estimation, as in Section 1.3, we allow the parameters to be time-varying to avoid imposing any particular functional form on the relationship between work hours and quitting behavior. In what follows, we report estimates that reflect behavior after 8.5 hours of work using the same sample as in Table 1.2.

In Table 2.1, columns (2) to (4) report estimates of the three behavioral models corresponding to the objective functions from Equations (2.3) to (2.5). Each of these models nests the neoclassical model from Equation (2.1) with no income effects, reported in column (1). A likelihood ratio test in each case rejects the null hypothesis of the neoclassical model (i.e., $\alpha = 0$, $L = 1$, and $\delta = 1$). The results in column (3) relate to the main analysis in Crawford and Meng [2011], which also consists of a model of loss aversion with static

[12] Recall, however, that the income-targeting model depends on a fixed reference point, the driver's mean earnings.

Table 2.1: *Maximum likelihood estimates: Static reference points*

	(1) Baseline	(2) Income targeting	(3) Loss aversion	(4) Salience
Disutility of trip	0.0240	0.0332	0.0450	0.0706
	(0.0021)	(0.0023)	(0.0027)	(0.0040)
Error term distribution σ	0.0580	0.0644	0.0642	0.0688
	(0.0041)	(0.0049)	(0.0019)	(0.0054)
Income targeting $1 + \alpha$		1.3014		
		(0.0031)		
Loss aversion L			1.3578	
			(0.0025)	
Salience δ				0.3590
				(0.0224)
Log-likelihood	-493,669	-474,985	-451,678	-469,585
Likelihood ratio test: baseline		< 0.001	< 0.001	< 0.001

Note: This table presents maximum likelihood estimates of Equation (2.7) *for different specifications of the objective function under the restriction of static reference points. The estimation sample consists of over 1.2 million shifts from 36,900 drivers. The columns correspond to the objective functions in Equations* (2.1) *and* (2.3) *to* (2.5)*, and the rows correspond to parameters. The row labeled 'disutility of trip' reports a combination of the elasticity parameter ν and the disutility of effort ψ given by $\frac{\psi}{1+\nu}(H_{t+1}^{1+\nu} - H_t^{1+\nu})$. The last row contains the p-value from a likelihood ratio test which takes the neoclassical model in column (1) as the null hypothesis.*

Table 2.2: *Maximum likelihood estimates: Adaptive reference points*

	(1) Loss aversion	(2) Salience
Disutility of trip	0.0412	0.0694
	(0.0027)	(0.0041)
Error term distribution σ	0.0623	0.0702
	(0.0020)	(0.0010)
Loss aversion L	1.3273	
	(0.0022)	
Salience δ		0.3290
		(0.0227)
Adjustment θ	0.9548	0.9672
	(0.0012)	(0.0018)
Log-likelihood	-449,346	-465,204
Likelihood ratio test: $\theta = 1$	< 0.001	< 0.001

Note: This table presents maximum likelihood estimates of Equation (2.7) *for different specifications of the objective function under adaptive reference points. The estimation sample consists of over 1.2 million shifts from 36,900 drivers. The columns correspond to the objective functions in Equations* (2.1) *and* (2.3) *to* (2.5)*, and the rows correspond to parameters. See the note to Table 2.1 for additional details. The last row contains the p-value from a likelihood ratio test which takes the corresponding model with a static reference point as the null hypothesis.*

reference points.[13] The income targeting model in column (2) implies a smaller change in the marginal utility of income at the target level. To the extent that the income targeting model mis-specifies the reference point, we would expect to underestimate degree of loss aversion; however, the model does not consist of an hours target and hence may mis-attribute loss aversion over hours to income. The model of expectations-based loss aversion appears to provide the best statistical fit of the data based on the log-likelihood, though the models are not nested and thus cannot be directly compared using this criterion.

Relaxing the assumption of a static reference point, we estimate models of expectations-based loss aversion and salience with adaptive reference points. We allow the reference

[13]Our analysis uses fewer parameters than Crawford and Meng [2011] because they also estimate a set of control variables. We account for the control variables from Section 1.3 by residualizing the income variable and verified that this choice does not affect any of the results.

point to vary within a shift by introducing a new parameter, the adjustment term θ, and defining the reference level I^r according to Equation (2.6). This specification nests the static reference point above ($\theta = 1$) as well as a reference point that adjusts instantaneously ($\theta = 0$). Both extreme cases correspond to the fungibility of money within a shift, with $0 < \theta < 1$ indicating a violation of fungibility. The speed of adjustment is primarily identified by variation in the timing of income. Table 2.2 reports results consistent with a violation of fungibility, highlighting the importance of within-day adjustments. The estimate of θ differs significantly from 0 and 1 in both models, and a likelihood ratio test rejects the restriction to a static reference point. While θ appears close to 1 in both cases, its magnitude depends on the definition of a period. Estimating the speed of reference point adjustment at a lower frequency (e.g., defining a period as an hour instead of a trip) would result in a smaller magnitude of θ. The remaining parameters do not substantially differ from their counterparts in Table 2.1.

The structural estimates for both models provide empirical support for the notion of adaptive reference points. One interpretation of the evidence suggests a departure from rational expectations as the reference point. Rational expectations derive from a steady-state distribution of potential earnings, and a given shift constitutes a particular realization from that distribution. This view of rational expectations predicts a static reference point ($\theta = 1$), which the evidence rejects.[14] Under an alternative interpretation, rational expectations incorporate new information within a shift, which leads the reference point to adjust. While this does not require that expectations or beliefs adjust slowly to new information, the evidence does suggest that preferences do not instantaneously change, consistent with the idea from Kőszegi and Rabin [2006] that preferences depend on lagged expectations. Relaxing the assumption that the reference point adapts symmetrically, we estimate a model of expectations-based loss aversion with asymmetric reference point adjustment. We augment Equation (2.6) by allowing the speed of adjustment to differ for gains (θ_+) and

[14] The results in Section A.5 also reject a reference point that does not adjust (i.e., $\theta = 1$) if we allow for stochastic reference points that capture the distribution of potential earnings and hours.

Table 2.3: *Maximum likelihood estimates: Asymmetric adaptation*

Parameter	Asymmetric
Disutility of effort	0.0559
	(0.0077)
Elasticity	0.2412
	(0.0371)
Error term distribution σ	0.0611
	(0.0172)
Loss aversion λ	1.4578
	(0.0223)
Adjustment θ_+	0.8527
	(0.0122)
Adjustment θ_-	0.9454
	(0.0122)
Likelihood ratio test: $\theta_+ = \theta_-$	< 0.001

Note: This table presents maximum likelihood estimates of Equation (2.7) *for the objective function in Equation* (2.4) *under adaptive reference points with asymmetric adaptation. The estimation sample consists of over 1.2 million shifts from 36,900 drivers. The last row contains the p-value from a likelihood ratio test which takes the model with symmetric adaptation as the null hypothesis.*

losses (θ_-). The estimate of θ_+ differs significantly from that of θ_-, and a likelihood ratio test rejects the restriction to a reference point with symmetric adaptation. Our results in Table 2.3 indicate that reference points tend to adjust more readily in the direction of gains rather than losses, contrary to the canonical view of reference points based on rational expectations.

2.5 Comparison of Models

This section provides a more detailed comparison between the models of expectations-based loss aversion and salience by exploring their implications for explaining the patterns of stopping behavior in the data.

Magnitude of income effect

This section assesses whether drivers who behave according to the estimates in Table 2.2 exhibit a relationship between earnings and quitting that conforms to the results from Section 1.3. Figure 2.1 shows the timing pattern of the income effect under the estimated models of expectations-based loss aversion and salience with adaptive reference points. For each model, the figure plots the predicted effect of an additional $10 earned at various times in the shift on the probability of ending a shift at 8.5 hours, with the estimates from Figure 1.7 providing a benchmark for comparison. As Section 2.3 highlights, adaptive reference points in each case lead to a violation of fungibility, with stronger labor-supply reductions in response to more recent earnings.

Although both models predict income effects that are qualitatively consistent with the pattern in Figure 1.7, the models make different predictions about the magnitude. The salience model from Equation (2.5) produces income effects that largely coincide with the reduced-form estimates. However, the loss-aversion model from Equation (2.4) leads to magnitudes that consistently exceed the income effects from the data. Re-estimating the model after relaxing the simplifying assumptions from Section 2.2 by allowing for stochastic reference points or curvature in the gain-loss function does not change this result (see

Figure 2.1: *Stopping model estimates: Income effect at 8.5 hours—Data and models*

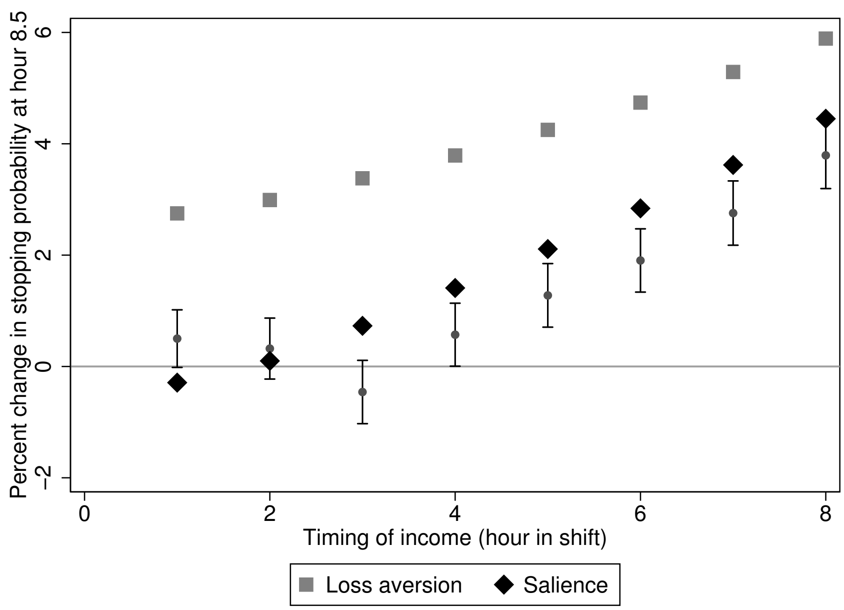

Note: The figure compares the predicted income effects from models of expectations-based loss aversion and salience estimated in Table 2.2 with the income effects estimated using Equation (1.2)*. The confidence interval displays the estimates from Figure 1.7 of the percent change in the probability of ending a shift at 8.5 hours in response to a $10 increase in earnings accumulated at different times in the shift. The gray square and black diamond represent the predictions of the loss-aversion model and the salience model, respectively.*

Table 2.4: *Maximum likelihood estimates: Loss aversion over income and hours*

	Estimates
Disutility of trip	0.0309
	(0.0030)
Error term distribution σ	0.0557
	(0.0015)
Loss aversion over income L_I	1.1037
	(0.0016)
Loss aversion over hours L_H	1.8628
	(0.0094)
Adjustment θ	0.8949
	(0.0022)
Log-likelihood	-437,900
Likelihood ratio test: $L_I = L_H$	< 0.001

Note: This table presents maximum likelihood estimates of Equation (2.7) *for the model of expectations-based loss aversion under adaptive reference points. The estimation sample consists of over 1.2 million shifts from 36,900 drivers. The objective function modifies Equation* (2.4) *by allowing for a separate coefficient of loss aversion on each dimension (income and hours). See the note to Table 2.1 for additional details. The last row contains the p-value from a likelihood ratio test which takes the model with a single coefficient of loss aversion as the null hypothesis.*

Section A.5).

As a possible explanation for the large estimated magnitudes, note that the Kőszegi and Rabin [2006] model assumes a constant coefficient of loss aversion across different dimensions of utility. To test whether this affects the results, we relax the assumption implicit in Equation (2.4) to allow for a different coefficient of loss aversion on each dimension (L_I for income and L_H for hours). The assumption of a "universal gain-loss function," designed to avoid introducing additional degrees of freedom, does not appear consistent with the estimates in Table 2.4. Although the estimates reveal a significant degree of loss aversion on both dimensions, a likelihood ratio test rejects the restriction that $L_I = L_H$. The degree of loss aversion over income is significantly smaller, consistent with the claim from Farber [2015] that the patterns in the data suggest a larger role for reference dependence in daily hours compared to income. The coefficients of loss aversion in Tables 2.1 and 2.2 attempt to simultaneously fit behavior over both of these dimensions, thereby overstating the degree of loss aversion over income.

This observation highlights the importance of analyzing the reduced-form implications of estimated models. A likelihood ratio test would simply lead to the conclusion that the behavioral models provide a better fit for the data than the neoclassical model. This says little about the ability of the models to explain the patterns of interest, especially those that are not directly matched. The exercise of checking what the parameter estimates imply about the underlying behavior that the model intends to capture leads to richer conclusions beyond statistical measures of fit.

While much of the motivation for studying daily income targeting stems from the observation that under some circumstances workers appear to work less when they can earn more, various forces naturally push in the opposite direction, each of which poses challenges for uncovering the role of reference dependence in earnings. Higher expected wages lead drivers to work more both in the neoclassical model and in the reference-dependent model [Kőszegi and Rabin, 2006]. Moreover, as Section 2.3 highlights, reference-point adjustment modulates the effects of loss aversion. Finally, if money constitutes news about future utility rather than contemporaneous consumption utility, then loss aversion over income can play a less pronounced role [Kőszegi and Rabin, 2009]. We address the first by analyzing behavior through the stopping model and the second by explicitly evaluating the role of timing. Under the view that our estimate of loss aversion over income is attenuated by news utility, the results provide an even stronger demonstration of reference dependence, though a better understanding of the relative importance of loss aversion over monetary outcomes remains a topic for future research.

Change in behavior at the reference point

This section considers an additional implication of the behavioral models from Section 2.2. Loss aversion and salience posit different mechanisms by which the income effect arises.

Figure 2.2: *Stopping model estimates: Income effect at 8.5 hours—Distance from income and hours targets*

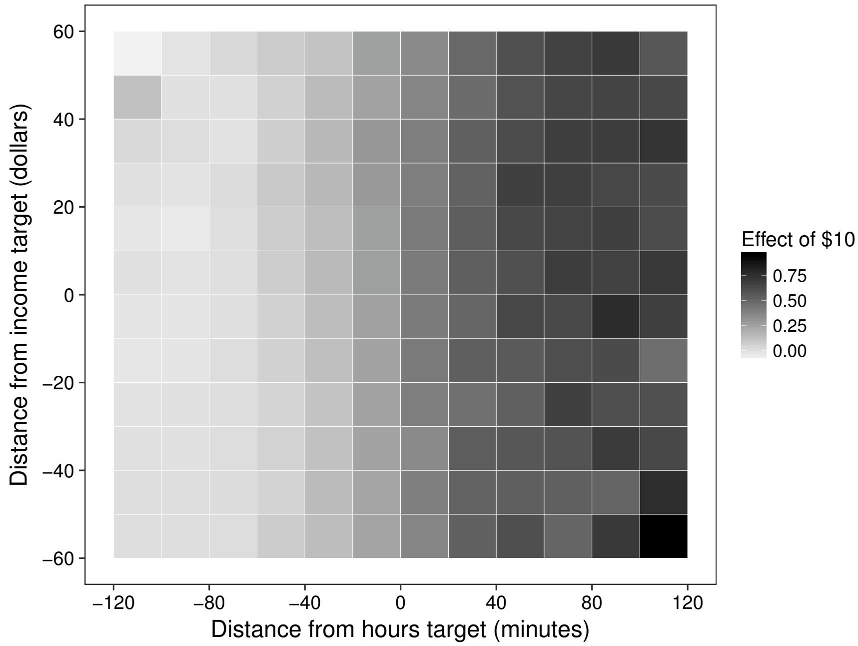

Note: The figure depicts the percentage-point increase in the probability of ending a shift at 8.5 hours when cumulative earnings is $10 higher. Estimates obtain from Equation (1.1) *extended to include interactions between income, twelve 20-minute intervals based on distance from the hours target, and twelve $10 intervals based on distance from the income target, with controls for location, time, and weather (see Table 1.2 for details) and fixed effects for 36,900 drivers. A darker shade represents a larger magnitude.*

Loss aversion invokes a sharp change in behavior near the reference levels. The model in Section 2.2 predicts, holding fixed income and hours, that the effect of income on the probability of stopping changes significantly when income is below the target compared to when income is above the target. The model of salience in Section 2.2, on the other hand, predicts a smooth change as income varies. The income effect arises through diminishing sensitivity, leading the driver to undervalue additional fares at higher levels of income, without a substantial change near the income target.

To directly test these predictions about how reference points influence stopping behavior, we re-estimate the marginal effect of earnings on the probability of stopping at 8.5 hours by extending Equation (1.1) to allow the income effect to vary based on distance from the income and hours target. We consider the set of trips that fall within $60 of the income target and within 2 hours of the hours target and categorize trips into 144 groups: twelve 20-minute intervals based on distance from the hours target interacted with twelve $10 intervals based on distance from the income target. Figure 2.2 depicts the effect of an additional $10 for each of these groups, with rows representing distance from the income target and columns representing distance from the hours target. Within each column, holding fixed the distance from the hours target but varying the distance from the income target, the income effect appears relatively stable. The magnitude of the income effect increases gradually with distance from the hours target but does not vary with distance from the income target.

Analyzing stopping behavior near the target helps to evaluate the simplifying assumptions in the model of loss aversion from Section 2.2. Once the level of income exceeds its target, the effect of an additional dollar on the probability of stopping under the assumption of a piecewise-linear gain-loss function becomes zero, which does not hold in the data. Models based on prospect theory [Kahneman and Tversky, 1979, Kőszegi and Rabin, 2006] typically posit a value function that exhibits diminishing sensitivity (convexity in losses and concavity in gains), but convexity in losses would predict a negative marginal effect of earnings on the probability of ending a shift below the income target. Relaxing the assumption that I^r and H^r denote the driver's point expectations for a given shift, a stochastic reference point [Kőszegi and Rabin, 2006] which represents the full distribution of a driver's potential earnings for that particular shift yields a smoother relationship between the income effect and distance from the target. For a driver who experiences gain-loss utility relative to the full distribution of potential earnings, additional income constitutes a partial mitigation of loss whether accumulated daily earnings fall above or below the target level. Finding positive income effects that do not change with distance from the income target therefore suggests an important role for stochastic reference points.[15]

The relationship in Figure 2.2 also clarifies the role of the assumptions in the salience

[15] The estimates in Section A.5 of the loss-aversion model with stochastic reference points show a substantial improvement in the log-likelihood compared to the model with point expectations as the reference level.

model from Section 2.2. In the salience model, the availability of stopping increases continuously as income accumulates up to and beyond the income target. For simplicity, our model of salience assumes that availability depends on income but not hours. Since the normal levels of both income and hours depend on availability, the marginal effect of earnings on the probability of stopping in the salience model varies with distance from the income target but not distance from the hours target. The pattern in the data that the magnitude of the income effect increases as the driver approaches or passes the hours target points toward hours as a key determinant of availability.

2.6 Discussion

The mechanisms explored in our setting may also be relevant for labor supply decisions more generally, especially with the rise of alternative work arrangements [Katz and Krueger, 2016]. A model of adaptive reference points could explain why cumulative earnings from fares may create a daily income effect while surge pricing generates the opposite behavior, as drivers' reference points for income quickly adjust upward to reflect higher anticipated earnings from future trips due to surge pricing. Further research can investigate the role of reference dependence and asymmetric adaptation in flexible work relationships as self-employment and contract work are becoming increasingly prevalent.

More broadly, the mechanisms behind the violations of fungibility could have implications for the design of government tax and transfer systems. For example, marginal propensities to consume may respond to the timing of income through institutional features such as tax withholding [Shapiro and Slemrod, 1995, Feldman, 2010]. Violations of fungibility via the timing of income could also help create and understand effective policy instruments for encouraging retirement savings or providing fiscal stimulus [Shefrin and Thaler, 1988, Souleles, 1999]. Future work can explore the influence of adaptive reference points and asymmetric adaptation in these as well as other field settings.

Chapter 3

The Public Housing Allocation Problem

3.1 Introduction

The efficacy of the public sector depends not only on the supply of public services but also on the design of systems for provision. While the former receives considerable attention from policymakers, poor design can entail substantial welfare losses. The market for public housing exemplifies this phenomenon. Tenants of public housing throughout the world—from approximately 1.2 million households in the United States to about one-third of all households in Hong Kong and Romania—express little choice in their place of residence, resulting in suboptimal allocations [Lui and Suen, 2011, Schwartz, 2010, Soaita, 2012].

Consistent with the idea of spatial mismatch, public-housing residents are less mobile than their private-housing counterparts, are less likely to work near where they live, and are more likely to move farther away from their original place of residence conditional on moving [Lui and Suen, 2011]. To the extent that housing is misallocated, designing improved allocation mechanisms can lead to substantial welfare gains. In the private-housing market, Glaeser and Luttmer [2003] argue that rent controls contribute to inefficient allocations, but a corresponding analysis of the market for public housing would require a model of how tenants are matched with apartments.

This paper models public-housing allocation as a matching problem in which units arrive stochastically over time and are assigned to applicants on a waiting list. Under this framework, we investigate the design of strategy-proof allocation mechanisms (i.e., those that are not subject to strategic manipulation) and explore whether these mechanisms can satisfy additional efficiency and fairness properties. Although there does not exist a mechanism that achieves these properties ex-post, we introduce a new mechanism—the Multiple-Waitlist Procedure (MWP)—which satisfies these properties ex-ante. Under MWP, applicants begin on a centralized waiting list in order of priority; a unit that arrives is offered to the agent at the top of the centralized waiting list, who can accept the offer or commit to a site-specific waiting list for another unit of their choice.

Mechanisms generally used in practice fail to satisfy desirable properties. Most public-housing agencies employ a take-it-or-leave-it procedure to allocate units to households on a priority-ordered waiting list.[1] A simple example with two buildings, two applicants, and two periods illustrates the possibility of unfair and inefficient allocations. Assume that household 1 has higher priority and that in period $k \in \{1, 2\}$ a unit in building b_k becomes available with certainty. Further assume that household 1 strictly prefers to wait for b_2 and that household 2 prefers b_1 to b_2. The take-it-or-leave-it mechanism assigns b_1 to household 1 and b_2 to household 2. The allocation is "unfair" in the sense that the higher-priority applicant prefers the assignment of the lower-priority applicant (i.e., the mechanism fails to eliminate justified envy); additionally, since household 2 also prefers the allocation of household 1, the mechanism is inefficient.[2] MWP satisfies both of these properties by giving applicants the opportunity to decline an offer and join a First-In/First-Out (FIFO) waiting list for the building of their choice.

To evaluate the performance of MWP in practice, we use a structural model of household

[1]Table 3.1 in Section 3.4.1 provides further details about public-housing agencies in the US and the procedures for allocating housing.

[2]These issues persist under a common variant of the take-it-or-leave-it mechanism that permits a household to decline some number of offers without forfeiting its position in the waiting list, as Section 3.4.1 explores.

preferences for public housing due to Geyer and Sieg [2013] and find that a counterfactual change in the allocation mechanism achieves substantial welfare gains. We estimate an arrival process for housing units using aggregate data from the Housing Authority of the City of Pennsylvania, and we obtain a sample of households eligible for public housing in Pittsburgh with preferences given by the Geyer-Sieg estimates using data from the Survey of Income and Program Participation. We simulate allocations under MWP as well as various alternative mechanisms that are currently being used to allocate public housing, some of which have support in the literature, to compare ex-post welfare. As a lower bound, we find that changing the allocation mechanism in Pittsburgh to MWP improves welfare by an average of $6,429 per household. Thus, MWP performs well ex-post, attaining 75 percent of an idealized benchmark defined as the maximum possible welfare gain that could be achieved if the realization of the arrival process were known in advance. Moreover, MWP increases the benefit of public housing by almost 20 percent without affecting its cost.

While several studies find evidence of the misallocation of private housing [Glaeser and Luttmer, 2003, Wang, 2011], empirical work related to public-housing allocation is more limited.[3] Recent work by Van Ommeren and Van der Vlist [2016] estimates the marginal willingness to pay for public-housing characteristics in Amsterdam and finds little welfare loss due to distortions in housing supply; they argue that deadweight loss mainly results from the inefficient match between households and apartments. Our paper complements these approaches by exploring the design aspect of public-housing allocation.[4] Geyer and Sieg [2013] develop an equilibrium framework for estimating household preferences for public housing under supply-side restrictions, which our paper uses to analyze welfare.

Our framework for studying the design of public-housing allocation mechanisms consists of a model of matching in a dynamic setting. Standard models from matching theory, such

[3]As Geyer and Sieg [2013] note, this is largely because "public housing agencies are not willing to disclose detailed micro-level data on wait lists."

[4]Recent work by Galiani et al. [2015] also studies the design of housing-assistance programs by using estimates from a structural model of neighborhood choice to simulate the effects of counterfactual housing-voucher policies.

as the canonical "house allocation problem" due to Hylland and Zeckhauser [1979], typically describe static situations; problems such as public-housing allocation, however, involve the feature that units arrive stochastically over time.[5] We develop a model that is based on existing static frameworks yet is rich enough to shed light on the operation of basic concepts from market design such as stability and efficiency in an environment with uncertainty. Instead of analyzing a fully dynamic problem, which would result in technical complications that are not relevant for the motivating application, we introduce a simple framework that maintains the key aspects of public-housing allocation. Specifically, our model incorporates the following features: (i) units are allocated dynamically as they arrive over time, (ii) there is uncertainty about the availability of the units, and (iii) applicants have preferences over waiting times. The robust policy recommendation that emerges from this framework involves households choosing among a set of waiting lists rather than only choosing to accept or reject units after they arrive. Leshno [2015] also studies an allocation problem with stochastic arrival and looks at a utilitarian objective in a simplified setting, while the present paper takes an axiomatic approach to a more general setting. Related work on dynamic assignment problems by Schummer [2016] and Bloch and Cantala [2017] analyzes selectivity thresholds for agents to accept or reject objects. While these papers also mention the allocation of public housing as a motivating example, they abstract away from the problem of eliciting waiting costs. The mechanism-design problem in the present paper, by contrast, permits heterogeneity in preferences over object types as well as over the amount of time spent waiting for an allocation. Furthermore, our results apply for any underlying stochastic process that governs the arrival of units.

A growing literature in market design uses simulations for welfare analysis, though much of this work focuses on the school-choice problem and involves randomly generated preference data.[6] Our work is among the first to use estimated preferences from data on

[5] For a survey treatment of the house allocation problem, see Sönmez et al. [2011].

[6] See, for example, Erdil and Ergin [2008], Dur [2011], Hafalir et al. [2013], Morrill [2013], Abdulkadiroğlu et al. [2015b], and Kesten and Ünver [2015].

real-world assignment procedures to quantify welfare gains due to adopting alternative mechanisms.[7] The counterfactual simulations are based on a model of preferences for public housing by Geyer and Sieg [2013], and we discuss in detail how their assumptions and our implementation all lead to an underestimate of the welfare gains from changing existing public-housing allocation mechanisms to MWP.

The paper is organized as follows. Section 3.2 describes the dynamic allocation problem and introduces various properties. Section 3.3 characterizes the Multiple-Waitlist Procedure and discusses the main theoretical results and extensions. Section 3.4 applies the framework to the allocation of public housing and uses estimates from a structural model to evaluate welfare. Section 3.5 concludes. Proofs can be found in Chapter B.

3.2 A Model of Public-Housing Allocation

Consider the problem that a public-housing agency faces in assigning units from apartment buildings that vary based on location to applicants on a long waiting list. Applicants have heterogeneous preferences over apartment buildings, as they may prefer to live closer to their respective workplaces, and heterogeneous waiting costs.[8] The public-housing agency ranks applicants based on priority; housing authorities typically rank applicants based on a coarse point system with ties broken based on waiting time, resulting in a strict priority ordering [Greely, 1977]. A unit must be allocated in the period in which it arrives. Although the period in which a given unit becomes available is not known in advance, the distribution of waiting times is known [Kaplan, 1986]. With this example in mind, we now proceed to introduce the components of the model more formally.

A *public-housing allocation problem* is a five-tuple $\langle A, B, \succ_B, \succ_A, \pi \rangle$, where A is a set

[7] In the literature on school choice, some recent papers estimate preferences using data from school districts: He [2012] in Beijing, China; Agarwal and Somaini [2014] in Cambridge, MA; Calsamiglia et al. [2016] in Barcelona, Spain; and Abdulkadiroğlu et al. [2015a] in New York City, NY.

[8] As a starting point, assume that units within each apartment building are identical. In practice, apartment buildings contain units of different sizes (i.e., as measured by the number of bedrooms); however, public-housing agencies administer separate waiting lists for units of different sizes.

of applicants, B is a finite set of buildings; $\succ_B = (\succ_b)_{b \in B}$ is a profile of the buildings' strict priority relations over the set of applicants; $\succ_A = (\succ_a)_{a \in A}$ is a profile of applicants' preference relations over building-time pairs $B \times \mathbb{R}_+$; and $\pi = \left(\pi_{b,\hat{t}}(\cdot \,|\, h^t)\right)_{b \in B, t, \hat{t} \in \mathbb{N}}$ is an arrival process that specifies, conditional on the history h^t, a probability distribution over the number of units in building b that arrive at time $\hat{t}$. A *history* is a map $h^t \colon B \times \{\tilde{t} \in \mathbb{N} : \tilde{t} \leq t\} \to \mathbb{N}$ that specifies the number of units in each building that have arrived in the previous periods. We make no specific assumptions about the underlying stochastic process which governs the arrival of units but will find it convenient to denote by $\tau_{b,t}(r)$ the expected waiting time of the r^{th} unit in building b to become available (counted from the beginning of time), conditional on the history h^t.

To denote applicant a being matched with the r^{th} unit to arrive in building b, we write $\mu(a) = \langle b, r \rangle$, where we refer to the map $\mu^t \colon A \to (B \times \mathbb{N}) \cup \{\emptyset\}$ as an *assignment in period* t and we refer to the collection $\mu = (\mu^1, \mu^2, \dots)$ of assignments in each period as an *assignment*. Agent a can be thought of as being on a waiting list for building b if the r^{th} unit is not yet available: given the history h^t of the arrival process, the unit is expected to arrive in $\tau_{b,t}(r)$ periods. If $\tau_{b,t}(r) = 0$, then this waiting list is degenerate, so the applicant receives the unit immediately in period t. We use the word *allocation* to refer to an agent's realized assignment, i.e., the applicant receives an allocation when the assigned unit arrives or becomes available.

A unit cannot be unmatched after the period in which it becomes available and cannot be reallocated in the future. The assumption that units must be allocated irreversibly upon arrival mirrors the corresponding assumptions in Gershkov and Moldovanu [2009] for the case of agents arriving sequentially. More importantly, this assumption captures a realistic institutional feature of public-housing allocation [Navarro, 2015].

To simplify the exposition, we assume that preferences are dynamically consistent and that waiting is costly. Preferences satisfy the *dynamic-consistency condition* if $(b, t) \succ_a (b', t')$ implies $(b, t + \hat{t}) \succ_a (b', t' + \hat{t})$ for every $\hat{t} > 0$. Preferences satisfy the *costly-waiting condition* if $t < t'$ implies $(b, t) \succ_a (b, t')$. These assumptions yield for each applicant a constant per-period waiting cost. Equivalently, the preference relation over $B \times \mathbb{R}_+$ reduces to a ranking over buildings and a vector in $\mathbb{R}_+^{|B|}$: for each building b, this representation encodes the greatest number of periods that the applicant is willing to wait before receiving a unit in her most-preferred building rather than receiving a unit in building b immediately. In addition, we assume that applicants are *risk neutral* with respect to preferences over waiting times. After stating the main result, Section 3.3.1 discusses how these assumptions can be relaxed.

Recall that preferences are defined over building-time pairs: $(b, t) \succ_a (b', t')$ if and only if applicant a prefers to receive a unit in building b in period t over a unit in building b' in period t'. However, an assignment $\mu^t(a)$ consists of a building $b \in B$ and a unit indexed by $r \in \mathbb{N}$ (i.e., the r^{th} unit that becomes available in building b). An agent a who is assigned $\mu^t(a)$ in period t expects to wait $\tau_{b,t}(r)$ periods for the unit to become available. The assumption of risk neutrality with respect to waiting times implies that applicants evaluate assignments based on expected waiting times: in particular, $\mu^t(a) = \langle b, r \rangle$ is preferred to $(\mu')^{t'}(a) = \langle b', r' \rangle$ if and only if $(b, t + \tau_{b,t}(r)) \succ_a (b', t' + \tau_{b',t'}(r'))$.

An *allocation mechanism* φ is a procedure that uses reported preferences, the exogenous priority orderings, and the history to choose an assignment μ^t in each period t. Let θ'_a denote the reported preferences of applicant $a \in A$, and let θ'_{-a} be the profile of reported preferences of all applicants except a. An allocation mechanism induces a preference-revelation game in which the set of players is A, the strategy space for player a is the set of preferences Θ, and each player $a \in A$ has true preference $\theta_a \in \Theta$.

We say that a mechanism φ is *strategy-proof* if deviation from truthful preference revelation is not profitable along any possible arrival history. Various authors have emphasized the desirability of strategy-proofness because of fairness (agents who lack information or sophistication are not at a disadvantage) and robustness (the equilibrium does not depend on beliefs about other agents' information or preferences). Another justification for strategy-proofness in our dynamic setting is that the social planner may make costly investments (though not modeled here) based on reported preferences.[9]

Next we define a property that can be interpreted as a form of fairness. An assignment μ *eliminates justified envy* or is *free of justified envy* if an applicant who prefers an alternate assignment does not have higher priority than the applicant to whom the other unit is assigned. Legal scholars posit this principle in the context of public-housing allocation, noting that "selection… would be by priority of application… to meet certain 'entitlements' arising out of a sense of fairness" [Greely, 1977]. A literature in market design that studies allocation problems with priorities, starting with the student placement model of Balinski and Sönmez [1999], employs this property as an analogue of stability. In our dynamic setting, whether an applicant prefers an alternate assignment depends on the timing of the match and the information available at the time. We say that an applicant *a envies* another applicant a' if given the information available at the time when a was matched a would have preferred the assignment of a'; we say that this envy is *justified* if a has higher priority than a' for the building with which a' is matched.

Definition 1. An assignment μ *eliminates justified envy* if whenever a is assigned $\langle b, r\rangle$ in period t and a' is assigned $\langle b', r'\rangle$ in period t', we have

$$\big(b', t + \tau_{b',t}(r')\big) \succ_a \big(b, t + \tau_{b,t}(r)\big) \implies a' \succ_{b'} a.$$

An ex-post variant of this no-envy condition would state that an applicant who prefers another unit to her own, given the realized arrival times of their respective units, cannot have higher priority than the applicant to whom the other unit is assigned.

An assignment μ is *efficient* if any reassignment μ' that one agent strictly prefers would make another agent strictly worse off (where, as before, we consider the information available at the time of the match). We refer to this as an ex-ante notion of efficiency because agents only take into account information that is available at the time when they are matched

[9] A housing agency may use reported preferences to determine where to construct a new building. In a separate context, Abdulkadiroğlu et al. [2009] point out that lack of demand as determined by reported student preferences under a strategy-proof mechanism contributed to the closing of an unpopular New York City high school in 2006.

and are evaluating their expected (rather than realized) arrival times.

Definition 2. An assignment μ is *efficient* if for any feasible assignment $\mu' \neq \mu$ for which there exists some a (who is assigned $\langle b_a, r_a \rangle$ in period t_a under μ and is assigned $\langle b'_a, r'_a \rangle$ in period t'_a under μ') such that $\left(b'_a, t_a + \tau_{b'_a, t_a}(r'_a)\right) \succ_a \left(b_a, t_a + \tau_{b_a, t_a}(r_a)\right)$, there exists some $a' \in A$ such that

$$\left(b_{a'}, t_{a'} + \tau_{b_{a'}, t_{a'}}(r_{a'})\right) \succ_{a'} \left(b'_{a'}, t_{a'} + \tau_{b'_{a'}, t_{a'}}(r'_{a'})\right)$$

An alternate notion of efficiency would be ex-post efficiency: any way of redistributing units that have already arrived (among the agents to whom they are assigned) cannot strictly improve the allocation of one applicant without making another strictly worse off.

We say that a mechanism satisfies a given property if the equilibrium allocations in the induced preference-revelation game satisfy that property. The following result shows that no allocation mechanism satisfies the ex-post variants of the properties introduced above.

Proposition 1. *There does not exist an allocation mechanism that is ex-post efficient or ex-post free of justified envy.*

The proof in Section B.1 consists of a simple example to demonstrate the impossibility of designing an allocation mechanism that guarantees these ex-post properties in an environment with stochastic arrival because the realization of the arrival process may be such that neither can possibly hold. Motivated by this result, our theoretical analysis in the next section focuses on the ex-ante properties defined above. Nevertheless, the empirical analysis in Section 3.4 uses ex-post welfare as a criterion for comparing mechanisms.

3.3 Multiple-Waitlist Procedure

3.3.1 Common priorities

We begin by considering the case that the buildings share a common priority ordering, i.e., that there exists $\succ_*$ such that $\succ_* = \succ_b$ for all $b \in B$. We introduce the Multiple-Waitlist Procedure (MWP) and characterize the matching that results from this mechanism.

Figure 3.1: *Multiple-Waitlist Procedure*

Step 0	Proceed to the next period.
Step 1	If there are no more available units this period: return to Step 0. Otherwise: choose a unit randomly from the set of units that are available this period.
Step 2	If the building in which the unit belongs has a non-empty FIFO queue: allocate the unit to the applicant at the top of the queue, and return to Step 1. Otherwise: offer the unit to the applicant at the top of the centralized waiting list.
Step 3	If the applicant accepts: allocate the unit to the applicant, and return to Step 1. Otherwise: the applicant chooses a different building and joins the associated FIFO queue to receive the next available unit in that building; return to Step 1.

Under MWP, all applicants begin on a centralized waiting list, and associated with each building there is a separate *First-In/First-Out* (FIFO) queue. A unit that becomes available in a given building belongs to the applicant at the top of the queue for that building. If the queue is empty, then the unit is offered to the applicant with the highest priority on the centralized waiting list. Given information about the distribution of arrival times, the applicant can either accept the offer or opt to join the FIFO queue for the next available unit in a different building of the applicant's choice. Note that if a' has higher priority than a, then a' receives an assignment (i.e., a unit or a place on some waiting list) before a does. Figure 3.1 describes MWP more formally.

The following proposition characterizes the main properties of this mechanism.

Proposition 2. *MWP satisfies the following properties: (i) strategy-proofness, (ii) efficiency, and (iii) elimination of justified envy.*

The following arguments summarize the formal proof in Section B.2. Note that an applicant receives an assignment after reaching the top of the centralized waiting list once a unit becomes available. The period when an applicant is matched depends on the choices of higher-priority applicants (and on the realization of the arrival process) but not on the agent's own reported preferences (and not on the choices of lower-priority applicants). In particular, since the priority ordering is fixed and independent of the agents' strategies, no agent can obtain a match sooner by misreporting preferences. Moreover, since the assignment is chosen to maximize the applicant's reported preference ordering (given the choices of the agents who have already been assigned, but independent of the choices of the agents who have not yet been assigned), the applicant cannot gain by deviating from truth telling. This argument suggests not only that MWP is strategy-proof but also that it satisfies a stronger condition, namely *obvious strategy-proofness*, which implies that the equilibrium prediction extends to agents with certain cognitive limitations [Li, 2015].[10]

The observation that the match of each applicant maximizes reported preferences (given the information available at the time of the match) also implies that all units assigned to lower-priority applicants are available but not chosen. Since priority orderings are common across buildings and no higher-priority applicant prefers the assignment of a lower-priority applicant, the allocation is efficient and free of envy. The results continue to hold after relaxing several assumptions in the framework as follows.

Risk neutrality An applicant who has risk-neutral preferences with respect to waiting times evaluates units based on expected arrival times and would have no use for additional moments of the distribution. A social planner can (i) implement MWP as a direct mechanism by eliciting preferences ex-ante, or (ii) disclose the expected arrival times and implement MWP as an indirect mechanism by making offers as in Figure 3.1. This holds more generally when the applicants have homogeneous risk preferences that the planner knows, as long as the planner can disclose certainty equivalents under the indirect mechanism. A planner with information about the applicants' risk preferences and the full distribution of waiting times can implement the direct mechanism, and a planner that can disclose the distribution of waiting times can implement the indirect mechanism without knowing the applicants'

[10] A mechanism is obviously strategy-proof if it has an equilibrium in obviously dominant strategies; truth telling is obviously dominant if the best-possible outcome from deviating is no better than the worst-possible outcome from reporting truthfully. Li [2015] notes that these notions also apply to the case that the set of outcomes of the mechanism consists of lotteries, so obvious strategy-proofness applies in our setting where an outcome is a unit in a building (which may consist of a distribution of waiting times). Under this view, MWP continues to implement the desired outcome when agents correctly perceive randomization by nature but use simplified mental representations of the other agents' strategies.

risk preferences.

Costly waiting The mechanism as described in Figure 3.1 moves an applicant who declines an offer from the top of the centralized waiting list to the end of a FIFO queue, thereby allocating the applicant with the next available unit in the chosen building. However, MWP can accommodate richer time preferences by allowing an applicant who prefers to wait longer (e.g., if moving costs change over time) to choose any unoccupied position in the queue.[11] In other words, the applicant can select any unassigned unit within the chosen building rather than the next available unassigned unit.

Static set of applicants The model captures the idea of an overloaded waiting list.[12] Under MWP, the centralized waiting list consists of the fixed priority-ordered set of applicants taken as a primitive of the model. Without any modification, the mechanism applies under an alternative formulation that involves an arrival process for applicants. The newly arrived applicants can be of any priority level and correspondingly can be added anywhere in the centralized waiting list. The mechanism likewise allows for departures from the centralized waiting list of applicants at any priority level. Insofar as the applicants' strategies do not influence such waiting-list dynamics, MWP satisfies all of the same properties since the procedure once an agent reaches the top of the centralized waiting list remains the same.

Static priorities We can interpret the exogenous priority ordering as reflecting the social planner's objective determined by verifiable information about the applicants. These priorities may change over time, for example, due to changes in the applicants' circumstances such as disability status or homelessness. The mechanism can incorporate any such changes, as the discussion above suggests for the case of applicants arriving or departing from the waiting list, by modifying the order of the applicants on the centralized waiting list. This

[11] In the context of public-housing allocation, for example, a household with an existing lease elsewhere may prefer a later allocation due to early-termination costs.

[12] See Leshno [2015] for a related allocation problem with two kinds of objects and an overloaded waiting list.

approach requires the planner to remain committed to allocating units that have not yet arrived to agents on the FIFO waiting lists. An alternative that also preserves the properties of the mechanism would be to allow applicants to become *critical* exogenously and alter the mechanism so that critical types receive an allocation immediately. The social planner can therefore accommodate situations such as domestic abuse or natural disasters by allowing such applicants to bypass the FIFO queues. The mechanism proceeds by adjusting the expected waiting times accordingly to account for the arrival process of critical agents.

Identical units The mechanism instantiates a separate FIFO queue for each building, where a "building" refers to an exogenously specified collection of identical units. If the classification of units into buildings were too fine, then an applicant who is indifferent between the units in two separate buildings would choose the building with the lower expected waiting time. Similarly, if the classification were too coarse (e.g., with unobserved heterogeneity in the applicants' preferences for units within a building), the applicant at the top of the centralized waiting list could still choose a FIFO queue to maximize expected utility, with the expectation taken over the distribution of unit quality in addition to waiting time.

We have shown that MWP satisfies several desirable properties and continues to do so in a more general setting. MWP satisfies additional regularity conditions discussed more formally in Section B.3: non-bossiness (no applicant can change another's allocation without changing her own allocation) and neutrality (the allocation does not depend on the labeling of the units). The next result complements our characterization of MWP by describing the sense in which MWP is unique. The uniqueness result does not rely on any of the efficiency or fairness properties introduced in Section 3.2.

Recall that a unit must be allocated in the period when it arrives. Under MWP, applicants at the top of the centralized waiting list receive offers (and move to their chosen FIFO queues if they decline) until no units available in the current period remain unallocated. Proposition 2 also applies to extended versions of this procedure that may place additional applicants from the top of the centralized waiting list in FIFO queues, even after allocating all units available in the current period. The following proposition states that under the regularity conditions, this class of extended multiple-waitlist procedures consists of all strategy-proof allocation mechanisms.

Proposition 3. *Any strategy-proof, non-bossy, and neutral allocation mechanism is an extended multiple-waitlist procedure.*

Section B.3 contains a formal proof of this result. Under the alternative procedures, the planner makes commitments before receiving information—even though waiting is feasible—that the planner may not have made after receiving the information. At an extreme lies the serial dictatorship: all applicants on the overloaded waiting list sequentially choose units in the initial period based on expected arrival times. Among this class of strategy-proof mechanisms, MWP uses the most information from the arrival process and requires the fewest promises about future allocations. Accordingly, MWP is most conducive to the more general setting highlighted above, which allows for time-varying priorities, risk preferences, and changes in the set of agents over time.

3.3.2 Heterogeneous priorities

Although our empirical analysis in Section 3.4 focuses on a case with common priority orderings, applications of our framework may not satisfy this restriction. Housing authorities often assign higher priority to elderly, disabled, or homeless applicants, as well as victims of natural disasters or domestic abuse. In some cases, however, priorities may differ across apartment buildings. For example, certain communities may be designated for specific groups such as seniors. For the general case that priority orderings are not common across object types, we provide a necessary and sufficient condition for the existence of a strategy-proof allocation mechanism that satisfies the ex-ante efficiency and fairness properties. The existence of such a mechanism depends on a property that can be interpreted as a measure of similarity between the priority orderings. Suppose a_1 has strong priority over a_3 at building b_1 in the sense that there is some a_2 such that $a_1 \succ_{b_1} a_2 \succ_{b_1} a_3$. We say that the priority orderings are *acyclic* if this implies a_1 has priority over a_3 at all buildings.

By defining the priority ordering over the set of applicants A, the framework implicitly assumes complete priority lists, i.e., that all applicants are acceptable to all buildings. The assumption is without loss of generality for the case of common priority orderings, as an unacceptable applicant can effectively be removed from the set A. With heterogeneous priority orderings, however, there may be applicants who are only acceptable to some buildings. In this case, an acyclic priority ordering satisfies the following condition: if a_1 has strong priority over a_3 at some building where a_3 is acceptable, then a_1 has priority over a_3 at any building where a_1 is acceptable. This property, adapted from Ergin [2002], is formalized as follows.

Definition 3. The collection of priority orderings $\succ_B$ contains a *cycle* if there exist $a_1, a_2, a_3 \in A$ and $b_1, b_2 \in B$ such that $a_1 \succ_{b_1} a_2 \succ_{b_1} a_3 \succ_{b_1} \emptyset$ and $a_3 \succ_{b_2} a_1 \succ_{b_2} \emptyset$. The collection $\succ_B$ is *acyclic* if it does not contain a cycle.

A complete priority ordering is acyclic if and only if the following property holds: for any applicant, there is no more than one other applicant who has higher priority at some buildings but lower priority at other buildings. In the context of public-housing allocation, the acyclicity condition can be satisfied by an agency that administers multiple programs which share a ranking over applicants but have some discretion on final priorities based on interviews.[13] The acyclicity condition is more permissive when priority orderings differ in terms of which applicants are acceptable. For example, the case of a public-housing agency that uses a common point scale to determine priorities but restricts access to certain buildings (e.g., senior housing) satisfies the acyclicity condition. We will show that a generalized version of MWP satisfies the properties in Proposition 2 for acyclic priority orderings.

The generalized Multiple-Waitlist Procedure is similar to MWP except that the order of the applicants' turns may be switched when one applicant prefers a building which

[13]See, for example, the allocation procedure in the District of Columbia.

Figure 3.2: *Generalized Multiple-Waitlist Procedure*

Step 0	Proceed to the next period.
Step 1	If there are no more available units this period: return to Step 0. Otherwise: choose a unit randomly from the set of units that are available this period.
Step 2	If the building in which the unit belongs has a non-empty FIFO queue: allocate the unit to the applicant at the top of the queue, and return to Step 1. Otherwise: offer the unit to the applicant (among those who have not yet been assigned a unit) with the highest priority for that building.
Step 3	If the applicant accepts: allocate the unit to the applicant, and return to Step 1. Otherwise: the applicant specifies a different building and requests to join the associated FIFO queue.
Step 4	If another applicant has the highest priority at the specified building: move that applicant to the top of the priority list for every building, and return to Step 2. Otherwise: the applicant joins the specified FIFO queue to receive the next available unit in that building; return to Step 1.

gives another applicant higher priority. In particular, an applicant who refuses an offer and prefers to join a waiting list for a different building must wait for the applicant (among those who have not yet been assigned a unit) with the highest priority at that particular building to choose first. In the language of Abdulkadiroğlu and Sönmez [1999], one way to describe this mechanism would be "you request my building—I get your turn." Figure 3.2 provides a more formal description of the generalized MWP.

The following result, shown formally in Section B.4, characterizes the generalized MWP under acyclic priority orderings.

Proposition 4. *If the priority orderings are acyclic, then the generalized MWP satisfies (i) strategy-proofness, (ii) efficiency, and (iii) elimination of justified envy.*

By exhibiting a mechanism that satisfies the three properties for acyclic priority orderings, this result demonstrates constructively that the acyclicity condition is sufficient for the existence of a mechanism that satisfies the desired properties. Additionally, the acyclicity condition is necessary for the existence of such a mechanism: to show this, Section B.4 provides an example of a deterministic arrival process (in an environment with two buildings that have cyclic priorities over three applicants) such that every possible allocation exhibits inefficiency or justified envy. The following proposition summarizes these findings.

Proposition 5. *There exists a strategy-proof allocation mechanism that is efficient and free of justified envy if and only if the priority orderings are acyclic.*

Recall from the framework that a mechanism assigns a particular unit in a particular building to each applicant. A natural question that arises is whether the characterization holds more generally for *stochastic allocation mechanisms*, i.e., mechanisms that assign a lottery over units to each applicant. Under this more general class of mechanisms, the concept of elimination of justified envy from Definition 1 does not apply since lotteries do not have priority orderings over applicants.

To address this issue, we use the notion of *strong stability* first introduced by Roth et al. [1993] and adapted by Kesten and Ünver [2015] in the context of school-choice lotteries.[14] We say that an applicant a' *strongly envies* another applicant a if there is a unit r in building b such that a can be assigned to $\langle b, r\rangle$ with positive probability while a' can be assigned to a less desirable unit (for her) than $\langle b, r\rangle$ with positive probability, and we say that this strong envy is *justified* if a' has higher priority than a for building b. A stochastic allocation mechanism is *strongly stable* if it eliminates justified strong envy.[15]

As the following result demonstrates, the characterization in Proposition 5 extends to stochastic allocation mechanisms.

Proposition 6. *There exists a strategy-proof stochastic allocation mechanism that satisfies efficiency and strong stability if and only if the priority orderings are acyclic.*

Even when considering the more general class of stochastic allocation mechanisms, the proof in Section B.5 shows that the presence or absence of cycles in the priority orderings fully characterizes the possibility for a strategy-proof allocation mechanism to satisfy

[14]Kesten and Ünver [2015] refer to this notion as *ex-ante stability* and provide a more general formulation that allows for weak priorities.

[15]Under non-stochastic allocation mechanisms, strong stability coincides with the elimination of justified envy because only a single unit is assigned to each agent with positive probability.

efficiency and the elimination of justified envy.

Due to the incompatibility between efficiency and the elimination of justified envy in the absence of acyclic priority orderings, we suggest strategy-proof allocation mechanisms for arbitrary priority orderings that satisfy each of these criteria separately (see Thakral [2015] for more detail).

A simple variation of MWP that satisfies efficiency would be to choose any ordering of the applicants and apply MWP as if this ordering were the common priority ordering. The ordering can be dynamically constructed, e.g., by choosing the applicant with the highest priority at the building that becomes available. This class of procedures produces efficient allocations: no applicant would benefit from an alternate allocation since each chooses her most-preferred unit at the time of assignment. However, these procedures do not eliminate justified envy since an applicant may choose a place on a waiting list for a building at which her priority is low.

A modified version of MWP can achieve the elimination of justified envy by constructing the priority ordering dynamically (as described above) but restricting the set of waiting lists that an applicant may join. A procedure that only allows an applicant to join a given queue if the applicant has top priority at the associated building satisfies the no-envy condition since any household that envies the applicant necessarily has lower priority. The mechanism permits inefficiencies since applicants who prefer units in buildings at which their respective priorities are low may benefit from switching their assignments.

3.4 Comparison of Allocation Mechanisms

We begin by evaluating the theoretical properties of mechanisms that are used to assign public housing and then proceed to investigate public-housing allocation mechanisms empirically. Using estimated preferences for public housing from a structural model due to Geyer and Sieg [2013], we find substantial ex-post welfare gains from changing the existing public-housing allocation mechanism to the Multiple-Waitlist Procedure described in Section 3.3.

3.4.1 Existing mechanisms

Public-housing applicants throughout the world express little choice in their place of residence under the most widely used allocation mechanisms. Applicants receive offers after units become available and can refuse only a limited number of times before forfeiting their position on the waiting list. This section provides institutional details on public-housing allocation in the United States (see Thakral [2017] for more details) and characterizes existing mechanisms.

A *Public Housing Authority* (PHA) is a state-run or locally-run entity that administers federal housing assistance programs. There are about 3,300 such agencies in the United States with approximately 1.2 million households living in public housing. The US Department of Housing and Urban Development (HUD) authorizes and funds PHAs and suggests two types of procedures that a PHA may use to allocate units, both of which involve finding the highest-priority applicant who is willing to accept the available unit rather than refuse. Under *Plan A*, the PHA offers a unit that becomes available to the applicant with the highest priority; if the applicant refuses, then the applicant is removed from or placed at the bottom of the waiting list. Under *Plan B*, an applicant who refuses a unit receives another offer, up to a limit of two or three total offers [Devine et al., 1999]. More generally, letting $k \in \mathbb{N}$ denote the maximum total number of units that the PHA offers to a given applicant, we refer to these procedures as PHA-k mechanisms. The top panel of Table 3.1 shows the number of housing agencies using these procedures. PHA-1 corresponds to Plan A, the take-it-or-leave-it procedure in which no applicant can refuse an offer without losing their position on the waiting list.

As mentioned in the introduction, the take-it-or-leave-it mechanism can lead to unfair and inefficient allocations. This observation applies more generally to the entire class of PHA-k mechanisms.

Proposition 7. *The PHA-k mechanisms are inefficient and do not eliminate justified envy.*

Section B.6 provides two separate proofs of this result, each consisting of an example of a deterministic arrival process for which the mechanism results in justified envy and

Table 3.1: *Distribution of allocation procedures by size of housing agency*

	Small	Medium	Large	Extra-large	Total
PHAs making 1 or 2 offers	430	114	71	3	618
PHAs making more than 2 offers	446	96	40	7	589
	Small	Medium	Large	Extra-large	Total
Centralized waiting list	929	179	117	9	1,234
Non-centralized waiting list	293	70	14	7	384

Note: Data are from the Division of Program Monitoring and Research, US Department of Housing and Urban Development, 1998. The top frame shows the relationship between housing-agency size and the maximum number of offers an applicant can receive before being removed from the waiting list. The bottom frame shows the relationship between housing-agency size and waiting-list method. Small: between 100 and 500 units. Medium: between 500 and 1,250 units. Large: between 1,250 and 6,600 units. Extra-large: 6,600 units or more. From the universe of over 3,100 housing agencies, those that operate fewer than 100 units are excluded, leaving a set of agencies that accounts for 94 percent of all public-housing units.

inefficiency. The first example contains only two buildings, and an applicant who refuses an offer from one building may receive another offer from the same building. The second example, in which the number of buildings depends on k, applies even if the mechanism requires that an applicant who refuses an offer does not receive another offer from the same building. Note from the top panel of Table 3.1, which displays the number of PHAs using the PHA-k mechanism for $k \in \{1, 2\}$ and for $k > 2$, that k is typically small in practice.

The first-come-first-served (FCFS) allocation mechanism can be thought of as a PHA-k mechanism with $k \to \infty$. Although not widespread in practice, this mechanism appears in several theoretical papers [Su and Zenios, 2004, Bloch and Cantala, 2017, Schummer, 2016]. Declining an offer creates a positive externality by reducing other agents' waiting time, but the agent at the top of the waiting list prefers to receive an allocation sooner due to private waiting costs.[16] Van Ommeren and Van der Vlist [2016] present empirical evidence from Amsterdam that the FCFS mechanism for allocating public housing produces an inefficient matching.

[16] As Su and Zenios [2004] argue, FCFS allocation mechanisms lead to an "inherent inefficiency [because of their] inability. . . to contain the externalities generated by [applicants'] self-serving behavior."

The properties of the PHA-k mechanisms also depend on the format of the waiting list. Under a *centralized* waiting list system, the applicant with the highest priority can be offered a unit in any building that becomes available. A *non-centralized* waiting list is either *site-specific* or *sub-jurisdictional*, depending on whether an applicant can only be assigned a unit in a particular building or in a group of buildings. The bottom panel of Table 3.1 shows the prevalence of centralized waiting lists. MWP uses both types of waiting lists: the PHA offers one unit to a household on the centralized waiting list and moves them to their chosen site-specific waiting list if the offer is refused. An alternative mechanism that places applicants on two (or more) site-specific waiting lists and then makes take-it-or-leave-it offers would fail to be strategy-proof. Consider an applicant with low waiting costs and a strong preference for a particular building: to maximize the chance of receiving a unit in that building, the applicant may report as her second choice a popular building with a long waiting list (even if the popular building lies at the bottom of her true preference ordering). Details about how the waiting lists are constructed can affect applicants' incentives, and procedures that are generally used in practice for assigning applicants to non-centralized waiting lists tend not to be strategy-proof.[17]

A potential justification for the use of the take-it-or-leave-it mechanism and its variants might be that these mechanisms act as screening devices, but several facts about housing policy suggest otherwise. First, the application process itself functions as a screening device: applications are costly to fill out, housing authorities verify the information reported on applications to determine priorities, and many housing authorities conduct in-person interviews of households that approach the top of the waiting list.[18] Second, some housing authorities use mechanisms that offer no additional screening benefit: FCFS allocation mechanisms, for example, do not remove applicants from the waiting list after declining

[17]The New York City Housing Authority, which uses sub-jurisdictional waiting lists (one for each of the five boroughs), asks applicants to report their first and second borough choice and explicitly advises applicants to "select their first borough choice carefully."

[18]See HUD [2009] for more details.

an offer.[19] Third, housing authorities typically aim to minimize the length of time required to fill a vacant unit: this is perhaps the most likely justification for using procedures in which each applicant receives a limited number of offers. MWP requires that the housing authority makes only a single offer to each applicant and thus performs equally as well as the take-it-or-leave-it procedure (and better than its variants) by this measure.

3.4.2 Estimation of welfare gains

Although the theoretical results establish that MWP is ex-ante efficient, a question that remains is whether the choice of an allocation mechanism affects ex-post welfare in real-world settings. We address this question by investigating the allocation of public housing using a sample of 215 eligible households in Pittsburgh, PA from the 2001 Survey of Income and Program Participation (SIPP) collected by the US Census Bureau. The model of household preferences consists of a standard random-utility specification as in Geyer and Sieg [2013]. Our goal is to quantify the welfare gains from changing the allocation procedure by simulating arrival processes and matchings under counterfactual mechanisms.

Buildings

Buildings are classified by size (small, fewer than 40 units; medium, between 40 and 100 units; or large, more than 100 units) and community type (family, senior, or mixed). In practice, any applicant can reside in any building. The 34 buildings operated by the Housing Authority of the City of Pittsburgh (HACP) in 2001 fall within six categories based on size and community type: family large, family medium, family small, mixed, senior large, and senior small. Buildings have a common priority ordering over applicants.

Table 3.2 displays the number of units in each building category that became available over a five-year period. From these data we estimate a binomial arrival process.[20] In any

[19] The City of Toronto adopted a FCFS allocation mechanism in July 2014. Similar mechanisms are also used in Britain and the Netherlands.

[20] Kaplan [1984] argues that the length of time that a household lives in public housing follows an exponential distribution; in a discrete-time setting (where the length of time in public housing follows a geometric

Table 3.2: *Arrivals by building type*

Building category	Number of buildings	Number of arrivals
Family large	12	677
Family medium	6	144
Family small	1	24
Mixed	4	300
Senior large	3	59
Senior small	8	191

Note: This table shows the number of units that arrive for each building category from June 2001 to June 2006. Data on the number of arrivals are from the HACP, as reported in Table 3 of Geyer and Sieg [2013].

given building, a unit arrives with an estimated probability of approximately 0.17 each week.[21]

Applicants

The utility of applicant $i \in A$ living in building $j \in B$ at time $t \in \mathbb{N}$ is given by

$$u_{i,j,t} = \beta \log y_{i,t} + \gamma_j + \kappa x_i - c\mathbf{1}_{\{d_{i,t} \neq d_{i,t-1}\}} + \varepsilon_{i,j,t}, \tag{3.1}$$

where $y_{i,t}$ denotes net income; γ_j is a building-category fixed effect; x_i is a vector of demographic characteristics, namely indicators for female, nonwhite, senior, and children; $d_{i,t} \in B \cup \{0\}$ denotes the residence and c is a moving-cost parameter; and $\varepsilon_{i,j,t}$ captures idiosyncratic tastes for public housing. We normalize the utility of living in private housing ($j = 0$) to be

$$u_{i,0,t} = \log y_{i,t} - c\mathbf{1}_{\{d_{i,t} \neq d_{i,t-1}\}} + \varepsilon_{i,0,t}. \tag{3.2}$$

distribution), the "moveout process" follows a binomial distribution.

[21] The results in this section are unaffected if we simulate a daily or monthly arrival process.

Following McFadden [1973], we assume that the idiosyncratic components are independently and identically distributed according to a standard type-I extreme-value distribution.[22]

A household that lives in private housing does not necessarily *prefer* to live there over public housing. Due to supply-side restrictions, households exhibit preferences for public housing by joining the waiting list.[23] A simple logit demand model based on observed choices would therefore fail to capture the reality of strong preferences for public housing. Geyer and Sieg [2013] develop an equilibrium framework that incorporates rationing and excess demand to model public-housing allocation. They identify the structural parameters of the utility function above using household exit behavior.[24] Table 3.3 reproduces estimates of the structural parameters from Geyer and Sieg [2013], which are based on household-level data from the HACP and the SIPP.[25]

Using a sample of low-income households eligible for public housing in Pittsburgh from the SIPP, we construct preferences based on the structural model above.[26] Since equations (3.1) and (3.2) describe the utility of living in a given building in a given period, a complete description of preferences requires assumptions about the discount factor and the match duration. Given our objective of measuring the welfare gains from using a mechanism that provides agents more choice at the cost of additional waiting time, we make conservative assumptions which cause our estimates to understate the true gains. A low discount factor increases the cost of waiting for a more preferred building, and a low match duration decreases the benefit of being matched with a more-preferred building.

---

[22] Geyer and Sieg [2013] find that using a nested logit specification designed to account for correlation in unobserved preferences among public-housing communities does not improve the fit of the model.

[23] A household in Pittsburgh typically waits between 14 and 22 months for a unit [Geyer and Sieg, 2013].

[24] Identification relies on the assumption of voluntary exit. The fact that housing authorities do not evict "over-income households" (i.e., those that exceed income thresholds) supports this assumption.

[25] The estimates suggest that minorities and female-headed households with children exhibit stronger preferences for public housing than other households. The fact that the coefficient β on log income is less than one suggests that an increase in income makes public housing less desirable than the outside option, consistent with the fact that a household residing in public housing pays 30 percent of its income as rent.

[26] Annual gross income in an eligible household must fall below 80 percent of the Area Median Income (AMI).

Table 3.3: *Parameter estimates for public-housing preferences*

Parameter	Mean	Standard error
Income	0.329	(0.028)
Moving cost	3.186	(0.017)
Demographics		
Nonwhite, nonsenior	1.222	(0.071)
White, senior	0.209	(0.113)
Nonwhite, senior	1.000	(0.101)
Children	−0.315	(0.123)
Female	0.053	(0.061)
Female, senior	−0.174	(0.094)
Female, children	0.426	(0.130)
Fixed effects		
Family large	4.217	(0.254)
Family medium	4.848	(0.261)
Family small	4.604	(0.277)
Mixed	4.394	(0.260)
Senior large	4.626	(0.263)
Senior small	4.907	(0.258)

Note: This table reports estimates of the parameters from the utility function in equations (3.1) *and* (3.2). *Parameter estimates are from the model with supply-side restrictions in Table 10 of Geyer and Sieg [2013], estimated using household-level data from the HACP and the SIPP.*

We assume that applicants maximize a discounted sum of utilities with a weekly discount factor of $\delta \approx 0.99$. Moreover, every match lasts for three years, so the applicant resides in private housing in all periods except for the three years immediately after moving into public housing. Section 3.4.3 provides further discussion on these assumptions.

Mechanisms

For each simulated arrival process, we determine the allocations that would result under each of the following mechanisms: MWP, PHA-1, PHA-2, and PHA-∞. During the sample period, the existing housing-allocation procedure is the take-it-or-leave-it mechanism (PHA-1).

MWP allows applicants to choose which site-specific waiting list to join based on information about the arrival process. Since the arrival time of the r^{th} unit in a building follows a negative binomial distribution, we use closed-form expressions to compute the expected utilities from joining waiting lists.

Under PHA-1, each applicant receives only one offer and faces the straightforward decision problem of accepting a unit if and only if the outside option provides lower utility.

PHA-2 gives applicants the opportunity to refuse an offer once. This decision depends on the applicant's beliefs about the households on the waiting list that have higher priority. We assume that the applicant knows how many households are ahead on the waiting list but does not know their preferences; instead, the applicant knows the distribution of preferences.[27] We use numerical approximations, drawing from the distributions of preferences and future arrivals, to compute the expected utility of refusing an offer.

PHA-∞ (FCFS) allows applicants to refuse an unlimited number of offers. We do not attempt to characterize equilibrium behavior under this mechanism; Bloch and Cantala [2017] study a model in which agents have private values and ultimately focus on equilibrium behavior under the two-agent case due to difficulties in providing a general characterization.

[27]The applicant also knows that any households that are ahead on the waiting list will accept the next offer if and only if the unit is acceptable (since they have already declined an offer).

For a given applicant, a state can be described by the number of households that are ahead in the waiting list and the number of offers that each of them has already refused. Any subset of the households ahead in the waiting list can still be on the waiting list by the time the applicant gets another offer, and each refusal reveals some additional information about the preferences of those households. This yields a number of states that is exponential in the size of the waiting list, which creates difficulties in trying to approximate equilibrium behavior. One way to proceed would be to introduce some heuristics by which applicants make such complex decisions. However, the resulting estimates would only provide a lower bound for welfare under the mechanism, as the heuristics might be too simplistic. Instead of assuming that applicants make suboptimal decisions, we make generous assumptions about the environment to obtain an upper bound for welfare under this mechanism. Specifically, we assume that applicants know the realization of the arrival process in advance and have complete information about the preferences of other households.[28] We denote the resulting upper bound on welfare by PHA-max.

Welfare

Our counterfactual simulations provide evidence that the welfare gain from changing the public-housing allocation mechanism to MWP is substantial. We convert the difference in utilities between each mechanism and PHA-1 to monetary values by computing the equivalent variation (EV), i.e., the transfer that the applicant would have to receive when public housing is assigned by PHA-1 that would give the applicant the same lifetime utility as the assignment under the new mechanism. Table 3.4 reports the average of the present-discounted values of these payments for PHA-2, PHA-max, and MWP.

Our estimates suggest that a change from PHA-1 to MWP improves the welfare of the average applicant who receives a housing assignment by an amount equivalent to a

[28]This assumption could in theory lead to an underestimate of welfare if lower-priority applicants have stronger preferences for public housing, but this concern does not arise in our setting: first, priorities are randomly drawn in each simulation; second, the model incorporates limited heterogeneity in household preferences, as Section 3.4.3 discusses in more detail.

Table 3.4: *Welfare gain relative to PHA-*1 *mechanism*

Mechanism	Mean $\overline{\text{EV}}$	95% Confidence Interval
PHA-2	\$3,866	[\$3,589, \$4,158]
MWP	\$6,429	[\$6,153, \$6,705]
PHA-max	\$6,779	[\$6,503, \$7,053]
Ex-post optimum	\$8,505	[\$8,212, \$8,800]

*Note: This table contains the results of 100 counterfactual simulations. In each simulation, we compute the average across all applicants of the equivalent variation ($\overline{EV}$) of changing the allocation mechanism from PHA-*1*. The second column reports the mean of present-discounted values (with an annual discount factor of* 0.6*). The final column provides lower and upper bounds of the bootstrapped 95-percent confidence interval (based on 10,000 replications). PHA-*max *denotes an upper bound for welfare under PHA-∞ in which applicants know the realization of the arrival process in advance and have complete information about the preferences of other households. The bottom row displays the maximum possible gain that could be achieved by a social planner that knows the complete realization of the arrival process in advance.*

one-time transfer payment of between \$6,100 and \$6,700.[29] Providing applicants with some choice in the allocation process by using PHA-2 rather than PHA-1 achieves 60 percent of this gain. By allowing applicants to express additional choice, MWP provides further improvement and falls within about 5 percent of PHA-max. Based on these estimates, we conclude that MWP outperforms PHA-∞. Recall that PHA-max is an idealized bound which assumes that applicants have complete information about other households' preferences and perfect foresight about the realization of the stochastic arrival process, which cannot be attained in practice. Although a precise prediction about welfare under PHA-∞ would require assumptions about applicants' strategies, the difficulty in characterizing equilibrium behavior suggests that real-world decision making under PHA-∞ would be far from optimal.

Note that the welfare measure only accounts for matches and ignores the costs of processing refusals and leaving units vacant. In our simulations, the next agent on the waiting list receives an offer immediately after a refusal. In practice, however, housing

[29]We exclude applicants at the top of the centralized waiting list (before more than one applicant accepts an offer in the present period), as these applicants may benefit from facing empty site-specific waiting lists when our simulation begins.

authorities typically provide three to five days for applicants to respond to an offer, which leads to longer vacancies. Due to underutilization, a mechanism such as PHA-2 or PHA-∞ that makes multiple offers to each applicant produces lower social welfare than our estimates suggest.

3.4.3 Interpretation of magnitudes

Given that a typical Pittsburgh household living in public housing earns less than $15,000 annually, MWP substantially improves ex-post welfare.[30]

As one way to assess the economic significance of the welfare gain, we compare our result with estimates of the benefits and costs of public housing. We measure the benefit of living in public housing using the equivalent variation, i.e., the transfer that each applicant would have to receive in private housing to give the applicant the same utility as the public-housing assignment under PHA-1. According to this measure, living in public housing is equivalent to a $14,412 increase in annual income. The gain from changing the allocation mechanism to MWP corresponds to an annual transfer payment of $2,572, which represents an 18 percent increase in the benefit of public housing. Moreover, the change in the allocation mechanism increases the cost effectiveness of public housing by 14 percentage points.[31]

As another way to evaluate the magnitude of the gain, we express the estimates relative to commuting times. The welfare gain corresponds to a daily transfer payment of $9, or equivalently 105 minutes at the minimum hourly wage ($5.15 from 2001 to 2006). According to the 2000 Census, a low-income worker in Pittsburgh spends on each working day an average of 46 minutes commuting by car or 83 minutes commuting by public transit. In that sense, the gain from changing the allocation mechanism exceeds the gain from eliminating

[30]The median total annual income (which includes means-tested transfers) of households eligible for public housing in our sample is $14,184. Households living in public housing earn an average of $9,082 in the HACP data as reported by Geyer and Sieg [2013].

[31]Using data from New York City, Olsen and Barton [1983] estimates a resource cost for providing a public-housing unit—consisting of loan payments for initial development costs, property taxes, operating costs—of $18,049 (inflated to 2006 dollars using the Consumer Price Index, and adjusted for administrative costs as pointed out by Olsen [2003]).

commuting times.

Finally, we compare the welfare gain to the maximum possible gain that could be achieved by a social planner that knows the complete realization of the arrival process in advance. As the bottom panel of Table 3.4 shows, the ex-post optimal allocation would result in an average welfare gain of $8,505, which provides an upper bound for the maximum possible gain that any mechanism can reach. MWP attains 75 percent of this perfect-foresight optimal-allocation benchmark.

We interpret our welfare estimates as providing a lower bound on the gains from changing the allocation mechanism to MWP, as discussed below.

First, the assumptions about timing in the simulations lead to an underestimate of the welfare gains. Note that households in our simulations spend three years in public housing, compared to almost seven years according to the HACP data as reported by Geyer and Sieg [2013]. A low match duration decreases the benefit of being matched with a more preferred building. If households with higher utility from public housing spend more time living in public housing, then using a fixed match duration further understates the average gains. Next note that the weekly discount factor of about 0.99 corresponds to an annual discount factor of 0.6, which falls below most estimates in the literature. A low discount factor increases the cost of waiting for a more preferred building. Both assumptions about timing lead to smaller estimated gains from a mechanism that allows applicants to wait longer for better matches.

Second, the measure of welfare places equal weight on all households. This might not be the most realistic welfare objective because of the fact that housing authorities assign priorities to each applicant. Under an objective that gives more weight to higher-priority applicants, the welfare gains are even larger.

Third, the structural model understates the extent of heterogeneity in preferences. Note that the model only includes six building-category fixed effects instead of a full set of building-specific fixed effects. Moreover, equation (3.1) does not contain any interaction between household characteristics and the building-category fixed effects, which implies that building categories are vertically differentiated and households on average only differ in terms of waiting costs. Some forms of heterogeneity depend on persistent idiosyncratic factors: applicants may prefer to live closer to their respective workplaces, allowing them to reduce commuting time and travel costs [Lui and Suen, 2011]. A more realistic model would consist of more heterogeneity in preferences from serially correlated error terms. Applicant characteristics also contribute to heterogeneity: households with children may prefer units that are located near better schools, and seniors may prefer units that have access to certain amenities. Using a model that incorporates these aspects, the estimated gain from using an allocation mechanism that improves match quality would be even larger.

Finally, the magnitude of our welfare estimate does not arise due to assumptions about unobserved heterogeneity in preferences. The idiosyncratic preference component can only affect the decision to accept an offer versus join a waiting list—but not which waiting list to join—because households can only condition on the realization of the idiosyncratic shock for the present period. Intuitively, the situation when the idiosyncratic term can affect decisions arises only when a household is on the margin of accepting an offer, when the difference in utilities is smallest. Additional simulations confirm that the estimates are not sensitive to assumptions about functional forms. Taking the extreme approach of entirely omitting idiosyncratic preferences, the point estimate for the welfare gain of MWP still lies within the 95 percent confidence interval presented in Table 3.4 at approximately the lower bound. Therefore we conclude that the estimated welfare gains are not driven by the extent of unobserved heterogeneity imposed by the structural model. To the extent that households' preferences in the real world do exhibit significant unobserved heterogeneity that the model does not capture, this further suggests that our results underestimate the actual welfare gains.

3.5 Discussion

We conclude with a discussion of practical matters pertaining to MWP in the context of public-housing allocation.

First, with over 1.2 million households living in public housing in the US alone, the overall gains from improved matching are substantial. Using a sample of households eligible for public housing in Pittsburgh, we find a lower bound to the welfare gains from changing the most commonly used take-it-or-leave-it allocation mechanisms to MWP of \$6,429 per applicant.[32] Although our theoretical analysis does not focus on ex-post welfare maximization, which may require specific assumptions about the arrival process, MWP performs well relative to the ex-post optimal allocation.

Second, MWP allows applicants to express choice without creating additional delays. Expanding choice by allowing applicants to refuse multiple offers creates delays in filling vacancies, which in practice results in welfare losses. The media often publicizes the average number of days that units remain vacant because vacancies lead to losses in rent, which imposes a burden on taxpayers when the federal government has to subsidize housing authorities for repairs and other operating costs. As a recent example, rent losses due to vacancies in New York City exceed \$8 million [Navarro, 2015]. Under MWP, households choose among a set of waiting lists rather than accept or reject units after they arrive. By only requiring housing authorities to make a single offer to each applicant, MWP provides both choice and utilization.

Third, households face a simple decision problem under MWP, which is strategy-proof. In the direct mechanism, an applicant must report (i) a ranking over buildings, and (ii) for each building, the number of periods the applicant would be willing to wait to receive a unit in her most-preferred building rather than receiving a unit in that building immediately. In the indirect mechanism, an applicant must select among a set of units associated with expected waiting times given information about the arrival process, which a housing authority can provide [Kaplan, 1986]. Existing mechanisms, by contrast, require difficult computations for agents to behave optimally.

Finally, although the model involves various stylized assumptions, the general lessons

[32]These measures do not account for objectives such as racial or economic integration of housing projects; Kaplan [1987] discusses how priority-assignment policies can achieve such goals.

apply in more realistic situations. The main results hold when we relax assumptions about time or risk preferences as well as in the presence of unobserved heterogeneity in applicants' preferences over units. Moreover, if additional applicants can join the waiting list or the priority ordering can change over time, a housing authority can implement MWP without any modification.

References

Atila Abdulkadiroğlu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, October 1999.

Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review*, 99(5):1954–1978, 2009.

Atila Abdulkadiroğlu, Nikhil Agarwal, and Parag A. Pathak. The welfare effects of coordinated assignment: Evidence from the NYC HS match. Working Paper 21046, National Bureau of Economic Research, March 2015a.

Atila Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda. Expanding "choice" in school choice. *American Economic Journal: Microeconomics*, 7(1):1–42, 2015b.

Johannes Abeler, Armin Falk, Lorenz Goette, and David Huffman. Reference points and effort provision. *The American Economic Review*, 101(2):470–492, 2011.

Nikhil Agarwal and Paulo Somaini. Demand analysis using strategic reports: An application to a school choice mechanism. Working Paper 20775, National Bureau of Economic Research, December 2014.

Sumit Agarwal, Mi Diao, Jessica Pan, and Tien Foo Sing. Labor supply decisions of singaporean cab drivers. *Mimeo*, 2015.

Eric J Allen, Patricia M Dechow, Devin G Pope, and George Wu. Reference-dependent preferences: Evidence from marathon runners. *Management Science*, 2016.

Steffen Andersen, Alec Brandon, Uri Gneezy, and John A. List. Toward an understanding of reference-dependent labor supply: Theory and evidence from a field experiment. Working Paper 20695, National Bureau of Economic Research, 2014.

Orley Ashenfelter, Kirk Doran, and Bruce Schaller. A shred of credible evidence on the long-run elasticity of labour supply. *Economica*, 77(308):637–650, 2010.

Michel Balinski and Tayfun Sönmez. A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84(1):73–94, 1999.

Nicholas Barberis, Ming Huang, and Tano Santos. Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53, 2001.

Levon Barseghyan, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. The nature of risk preferences: Evidence from insurance choices. *American Economic Review*, 103(6):2499–2529, 2013.

Francis Bloch and David Cantala. Dynamic assignment of objects to queuing agents. *American Economic Journal: Microeconomics*, 9(1):88–122, 2017.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285, 2012.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843, 2013.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Memory, attention and choice. *Mimeo*, 2015.

Martin Browning, Angus Deaton, and Margaret Irish. A profitable approach to labor supply and commodity demands over the life-cycle. *Econometrica*, 53(3):503–543, 1985.

Benjamin Bushong, Matthew Rabin, and Joshua Schwartzstein. A model of relative thinking. *Mimeo*, 2016.

Meghan R. Busse, Devin G. Pope, Jaren C. Pope, and Jorge Silva-Risso. The psychological effect of weather on car purchases. *The Quarterly Journal of Economics*, 2014.

Caterina Calsamiglia, Chao Fu, and Maia Güell. Structural estimation of a model of school choices: the Boston mechanism vs. its alternatives. *Mimeo*, 2016.

Colin Camerer, Linda Babcock, George Loewenstein, and Richard Thaler. Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2): 407–441, 1997.

David Card and Gordon B. Dahl. Family violence and football: The effect of unexpected emotional cues on violent behavior. *The Quarterly Journal of Economics*, 126(1):103–143, 2011.

Tom Chang and Tal Gross. How many pears would a pear packer pack if a pear packer could pack pears at quasi-exogenously varying piece rates? *Journal of Economic Behavior & Organization*, 99:1–17, 2014.

Yuan K. Chou. Testing alternative models of labour supply: Evidence from taxi drivers in singapore. *The Singapore Economic Review*, 47(01):17–47, 2002.

Michael Conlin, Ted O'Donoghue, and Timothy J Vogelsang. Projection bias in catalog orders. *The American Economic Review*, pages 1217–1249, 2007.

Vincent P. Crawford and Juanjuan Meng. New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review*, 101(5):1912–1932, 2011.

Stefano DellaVigna. Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–372, 2009.

Stefano DellaVigna, John A. List, and Ulrike Malmendier. Testing for Altruism and Social Pressure in Charitable Giving. *The Quarterly Journal of Economics*, 127(1):1–56, feb, 2012.

Stefano DellaVigna, Attila Lindner, Balázs Reizer, and Johannes F. Schmieder. Reference-dependent job search: Evidence from hungary. *Quarterly Journal of Economics*, 2017.

Deborah J. Devine, Lester Rubin, and Robert W. Gray. *The Uses of Discretionary Authority in the Public Housing Program: A Baseline Inventory of Issues, Policy, and Practice*. US Department of Housing and Urban Development, Office of Policy Development and Research, 1999.

Kirk Doran. Are long-term wage elasticities of labor supply more negative than short-term ones? *Economics Letters*, 122(2):208–210, 2014.

Pascaline Dupas, Jonathan Robinson, and Santiago Saavedra. The daily grind: Cash needs, labor supply and self-control. 2016.

Umut Dur. Tie-breaking is still a problem: How to find the pareto efficient and fair matching in the school choice problem. *Mimeo*, 2011.

Aytek Erdil and Haluk Ergin. What's the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review*, 98(3):669–689, 2008.

Haluk I. Ergin. Efficient resource allocation on the basis of priorities. *Econometrica*, 70(6): 2489–2497, November 2002.

Keith M. Marzilli Ericson and Andreas Fuster. Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments. *The Quarterly Journal of Economics*, 126(4):1879–1907, 2011.

Henry S. Farber. Is tomorrow another day? the labor supply of new york city cabdrivers. *Journal of Political Economy*, 113(1):46–82, 2005.

Henry S. Farber. Reference-dependent preferences and labor supply: The case of new york city taxi drivers. *American Economic Review*, 98(3):1069–1082, 2008.

Henry S. Farber. Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, page qjv026, 2015.

Ernst Fehr and Lorenz Goette. Do workers work more if wages are high? evidence from a randomized field experiment. *The American Economic Review*, 97(1):298–317, 2007.

Naomi E. Feldman. Mental accounting effects of income tax shifting. *Review of Economics and Statistics*, 92(1):70–86, 2010.

Sebastian Galiani, Alvin Murphy, and Juan Pantano. Estimating neighborhood choice models: Lessons from a housing assistance experiment. *American Economic Review*, 105 (11):3385–3415, 2015.

David Genesove and Christopher Mayer. Loss aversion and seller behavior: Evidence from the housing market. *The Quarterly Journal of Economics*, 116(4):1233–1260, 2001.

Alex Gershkov and Benny Moldovanu. Learning about the future and dynamic efficiency. *American Economic Review*, 99(4):1576–1587, 2009.

Judy Geyer and Holger Sieg. Estimating a model of excess demand for public housing. *Quantitative Economics*, 4(3):483–513, 2013.

David Gill and Victoria Prowse. A structural analysis of disappointment aversion in a real effort competition. *American Economic Review*, 102(1):469–503, 2012.

Edward L. Glaeser and Erzo F. P. Luttmer. The misallocation of housing under rent control. *American Economic Review*, 93(4):1027–1046, 2003.

Lorenz Goette, David Huffman, and Ernst Fehr. Loss aversion and labor supply. *Journal of the European Economic Association*, 2(2-3):216–228, 2004.

Hank Greely. The equality of allocation by lot. *Harvard Civil Rights-Civil Liberties Law Review*, 12:113, 1977.

Michael D. Grubb. Dynamic nonlinear pricing: Biased expectations, inattention, and bill shock. *International Journal of Industrial Organization*, 30(3):287–290, 2012.

Faruk Gul. A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, pages 667–686, 1991.

Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363, 2013.

Kareem Haggag and Giovanni Paci. Default tips. *American Economic Journal: Applied Economics*, 6(3):1–19, 2014.

Kareem Haggag, Brian McManus, and Giovanni Paci. Learning by driving: Productivity improvements by new york city taxi drivers. *American Economic Journal: Applied Economics*, 9(1):70–95, 2017.

Justine S. Hastings and Jesse M. Shapiro. Fungibility and consumer choice: Evidence from commodity price shocks. *The Quarterly Journal of Economics*, 128(4):1449–1498, 2013.

Yinghua He. Gaming the Boston school choice mechanism in Beijing. *Mimeo*, 2012.

Martin Huber, Michael Lechner, and Conny Wunsch. The performance of estimators based on the propensity score. *Journal of Econometrics*, 175(1):1–21, 2013.

HUD. A guide to interviewing for owners of HUD-subsidized multifamily housing projects. Multifamily Housing Case Studies, August 2009.

A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, pages 293–314, 1979.

David S. Johnson, Jonathan A. Parker, and Nicholas S. Souleles. Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610, 2006.

David Jonason and Joel Wållgren. Labor supply of stockholm cab drivers: Revisiting the evidence. *Mimeo*, 2013.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

Edward H. Kaplan. *Managing the Demand for Public Housing*. PhD thesis, Massachusetts Institute of Technology, 1984.

Edward H. Kaplan. Tenant assignment models. *Operations Research*, 34(6):832–843, 1986.

Edward H. Kaplan. Analyzing tenant assignment policies. *Management Science*, 33(3):395–408, 1987.

Heiko Karle, Georg Kirchsteiger, and Martin Peitz. Loss aversion and consumption choice: Theory and experimental evidence. *American Economic Journal: Microeconomics*, 7(2): 101–120, 2015.

Lawrence F Katz and Alan B Krueger. The rise and nature of alternative work arrangements in the united states, 1995-2015. *Mimeo*, 2016.

Onur Kesten and M. Utku Ünver. A theory of school choice lotteries. *Theoretical Economics*, 10(2):543–595, 2015.

Botond Kőszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006.

Botond Kőszegi and Matthew Rabin. Reference-dependent consumption plans. *American Economic Review*, 99(3):909–936, 2009.

Botond Kőszegi and Adam Szeidl. A model of focusing in economic choice. *The Quarterly Journal of Economics*, 128(1):53–104, 2013.

David Laibson, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. Estimating discount functions with consumption choices over the lifecycle. *Mimeo*, 2015.

Jacob D. Leshno. Dynamic matching in overloaded waiting lists. *Mimeo*, 2015.

Shengwu Li. Obviously strategy-proof mechanisms. SSRN Scholarly Paper ID 2560028, Social Science Research Network, November 2015.

George Loewenstein, Ted O'Donoghue, and Matthew Rabin. Projection bias in predicting future utility. *The Quarterly Journal of Economics*, 118(4):1209–1248, 2003.

Hon-Kwong Lui and Wing Suen. The effects of public housing on internal mobility in Hong Kong. *Journal of Housing Economics*, 20(1):15–29, March 2011.

Daniel McFadden. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, pages 105–142. Academic Press: New York, 1973.

Ender Faruk Morgul and Kaan Ozbay. Revisiting labor supply of new york city taxi drivers: Empirical evidence from large-scale taxi data. *Transportation Research Board*, 2014.

Thayer Morrill. Making efficient school assignment fairer. *Mimeo*, 2013.

Mireya Navarro. New York City public housing units remain empty unnecessarily, audit finds. *The New York Times*, June 2015.

Quang Nguyen and PinngSun Leung. Revenue targeting in fisheries: The case of hawaii longline fishery. *Environment and Development Economics*, 2015.

Terrance Odean. Are investors reluctant to realize their losses? *The Journal of Finance*, 53(5): 1775–1798, 1998.

Edgar O. Olsen. Housing programs for low-income households. *NBER*, pages 365–442, January 2003.

Edgar O. Olsen and David M. Barton. The benefits and costs of public housing in New York City. *Journal of Public Economics*, 20(3):299–332, April 1983.

Jonathan A. Parker, Nicholas S. Souleles, David S. Johnson, and Robert McClelland. Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–53, 2013.

Devin G. Pope and Maurice E. Schweitzer. Is tiger woods loss averse? persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, 101(1): 129–157, 2011.

Thierry Post, Van Den Assem, Martijn J, Guido Baltussen, and Richard H. Thaler. Deal or no deal? decision making under risk in a large-payoff game show. *American Economic Review*, 98(1):38–71, 2008.

Alvin E. Roth, Uriel G. Rothblum, and John H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, November 1993.

James Schummer. Influencing waiting lists. *Mimeo*, 2016.

Heather L. Schwartz. Housing policy is school policy. *A Century Foundation Report*, 2010.

Matthew D. Shapiro and Joel Slemrod. Consumer response to the timing of income: Evidence from a change in tax withholding. *The American Economic Review*, 85(1):274–283, 1995.

Hersh M. Shefrin and Richard H. Thaler. The behavioral life-cycle hypothesis. *Economic Inquiry*, 26(4):609–643, 1988.

Adriana Mihaela Soaita. Strategies for *In Situ* home improvement in Romanian large housing estates. *Housing Studies*, 27(7):1008–1030, October 2012.

Changcheng Song. An experiment on reference points and expectations. *Mimeo*, 2016.

Tayfun Sönmez, M Utku Ünver, et al. Matching, allocation, and exchange of discrete resources. *Handbook of Social Economics*, 1:781–852, 2011.

Nicholas S. Souleles. The response of household consumption to income tax refunds. *The American Economic Review*, 89(4):947–958, 1999.

Charles Sprenger. An endowment effect for risk: Experimental tests of stochastic reference points. *Journal of Political Economy*, 123(6):1456–1499, 2015.

Tess M Stafford. What do fishermen tell us that taxi drivers do not? an empirical investigation of labor supply. *Journal of Labor Economics*, 33(3):683–710, 2015.

Stephen M Stigler. Do robust estimators work with real data? *The Annals of Statistics*, pages 1055–1098, 1977.

Xuanming Su and Stefanos Zenios. Patient choice in kidney allocation: The role of the queueing discipline. *Manufacturing & Service Operations Management*, 6(4):280–301, 2004.

Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, September 1999.

Neil Thakral. Matching with stochastic arrival. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 343–343, New York, NY, USA, 2015. ACM.

Neil Thakral. The public-housing allocation problem: Theory and evidence from Pittsburgh. *Mimeo*, 2016.

Neil Thakral. The design of public-housing policies. *Mimeo*, 2017.

Neil Thakral and Linh T Tô. Daily labor supply and adaptive reference points. *Mimeo*, 2018.

Jos N. Van Ommeren and Arno J. Van der Vlist. Households' willingness to pay for public housing. *Journal of Urban Economics*, 92:91–105, 2016.

Shing-Yi Wang. State misallocation and housing prices: Theory and evidence from China. *American Economic Review*, 101(5):2081–2107, 2011.

Appendix A

Appendix to Chapters 1 and 2

A.1 Data

As in Haggag and Paci [2014], we first process the data by dropping data errors including those resulting from electronic tests.

1. If the drop-off time is before the pick-up time in a trip, then we swap the drop-off time and pick-up time: 0.01 percent of the trips.
2. If the same driver or the same car's drop-off time is after the pick-up time of a subsequent trip, then we set the drop-off time to be equal to the pick-up time of the subsequent trip: 0.06 percent of the trips.

We flag trips that have any of the inconsistencies outlined below:

1. Trips that have distance of zero: 0.65 percent of the trips.
2. Trips that have ride duration of zero: 0.3 percent of the trips.
3. Trips with payment type recorded as "No Charge" or "Dispute": 0.42 percent of trips.
4. When fare is too high or too low compared to distance and time and locations: 0.16 percent of the trips.
5. When fare is too low compared to distance and time and locations: 0.25 percent of the trips.

6. Trip time as indicated by the pick-up and drop-off timestamps and the recorded ride duration do not match: 0.20 percent of the trips.
7. Outliers with trip durations longer than 3 hours or trip distance longer than 100 miles: 0.65 percent of the trips.
8. Trip between Manhattan and an airport in under 5 minutes: 0.07 percent of the trips.
9. Trip between Manhattan and JFK International airport in under 10 miles: 0.06 percent of the trips.
10. When a ride lasts fewer than ten seconds, or fewer than one minute and costs over $10: 0.68 percent of the trips.
11. When a ride lasts fewer than ten seconds, or fewer than one minute and costs over $10: 0.49 percent of the trips.
12. When average speed during a trip exceeds 80 miles per hour: 0.39 percent of the trips.
13. Trips belonging to truncated shifts (those that start before the first day or end after the last day of the year): 0.09 percent of the trips.

We then remove shifts with trips that have been flagged with errors or shifts that are are outliers:

1. Trips in the same shift but with more than one cars: 0.40 percent of the shifts.
2. Shifts that are longer 18 hours: 1.12 percent of the shifts.
3. Shifts that are shorter than two hours: 2.31 percent of the shifts.
4. Shifts by drivers with under 100 rides on record (may be electronic tests sent by TLC or the vendors): 0.12 percent of the shifts.
5. Shifts with fewer than three trips: 1.5 percent of the shifts.

As many of the analyses require that trips are in successive order, we remove the whole shift when one trip is questionable, and the process reduces our sample by approximately 24 percent. After cleaning out shifts, we remove drivers with under ten shifts, or an additional 0.1 percent of the observations. We are left with a sample of 127 million observations from over 37,000 drivers in over 5.8 million shifts. Finally, in many of the analyses, we restrict our sample to shifts that stay within the five boroughs in NYC, consisting of 94 percent of the remaining shifts.

A.2 Elasticity Estimates

In this section, we discuss how the wage profile can potentially lead to mechanical biases of wage-elasticity estimates. To illustrate how elasticities can be biased in the positive direction, consider a hypothetical driver who supplies labor inelastically, with some noise around the optimal stopping time. If the wage increases throughout each shift, then a regression of log hours on log wages will have a positive coefficient on log wages since longer shifts mechanically have higher average earnings.

We provide suggestive evidence that this bias might be present in the setting of cabdrivers in NYC by estimating elasticities on subgroups of shifts with different wage patterns. To estimate wage elasticities, we follow the approach by Camerer et al. [1997] and Farber [2015], regressing the logarithm of the total working hours in a shift on the logarithm of the average earnings per hour in that shift, with time controls (indicators for day of week, week of year, and federal holidays) and driver fixed effects. We also use average market wage of a non-overlapping sample of drivers to instrument for a driver's wage.

Figure A.1 displays the pattern of average wages on weekdays and on weekends. Though the patterns are similar across AM shifts, they diverge significantly between 10 PM and 1:30 AM when the average wage is rising for weekend shifts but falling for weekday shifts. As around half of the cabdrivers who work during the PM shift stop during this period, this distinction may have a nontrivial impact on the elasticity estimates. For each type of shift (day or night), we estimate the wage elasticity of weekday shifts and weekend shifts separately, restricting the sample to drivers who appear in both groups to avoid compositional differences in responsiveness to wage changes. The estimates in Table A.1 confirm that while wage elasticities across weekdays and weekends are similar for AM shifts, they are substantially higher during weekends for PM shifts, consistent with the pattern of increasing average wages on weekend nights. Instrumental-variable estimates would imply

Figure A.1: *Pattern of wages: Weekday versus weekend*

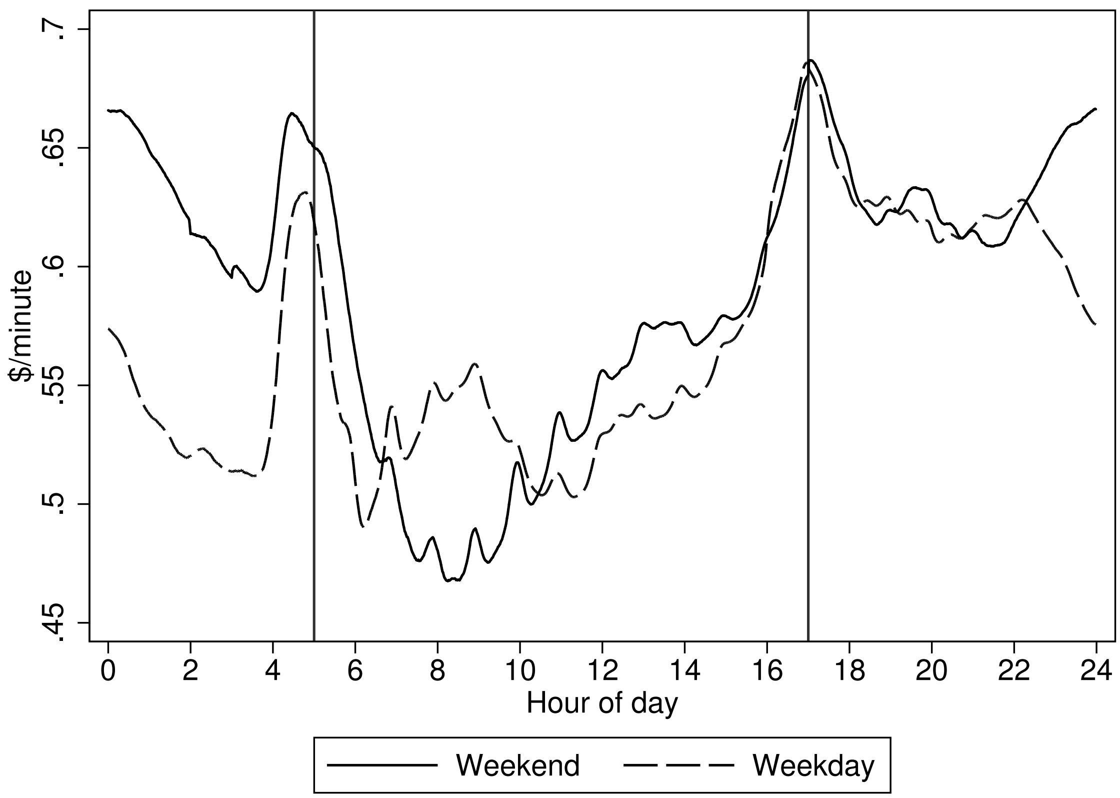

The figure depicts the average market wage every minute throughout the day, separated into weekdays and weekends. The market wage in each minute is the average of the per-minute wages of all drivers working during that minute, where a driver's per-minute wage is the ratio of the fare (not including tips) to the number of minutes spent searching for or riding with passengers for their current trip. Weekend is defined as 5 PM Friday through 5 PM Sunday. Weekday is defined as 5 PM Sunday through 5 PM Friday.

Table A.1: *Wage elasticity estimates: Weekday versus weekend*

Shift	Time	OLS	IV
Day	Weekday	-0.1636	0.2632
		(0.0031)	(0.0074)
Day	Weekend	-0.1214	0.0076
		(0.0040)	(0.0293)
Night	Weekday	-0.3546	0.3067
		(0.0027)	(0.0083)
Night	Weekend	-0.1418	1.3085
		(0.0041)	(0.0260)

Each cell presents elasticity estimates from a regression of log hours on log wages, with time controls (indicators for day of week, week of year, and federal holidays) and driver fixed effects (21,244 for night-shift drivers and 18,569 for day-shift drivers). Day shifts start between 4 AM and 10 AM, and night shifts start between 2 PM and 8 PM. Weekend shifts consist of night shifts on Friday and Saturday as well as day shifts on Saturday and Sunday. For each shift type (day or night), the sample consists of drivers who appear in both the weekday and weekend group. The IV column instruments for wages using the average hourly wage from a non-overlapping sample of 2,108 drivers on the same day. Standard errors reported in parentheses are adjusted for clustering at the driver level.

striking differences in behavior between weekdays and weekends for PM shifts, with an elasticity of 0.3067 on weekends and 1.3085 on weekdays, and no such difference in behavior for AM shifts. As the direction of these estimates coincides with predictions based on daily wage patterns, our results suggest that within-day variation in wages can lead to biases in elasticity estimates.

A.3 Model and Estimation

A.3.1 A Model of Daily Labor Supply

We begin by presenting a neoclassical model of intertemporal utility maximization with time-separable utility. We then formulate testable predictions about daily labor-supply decisions and proceed to evaluate these predictions using data on the labor supply of NYC cabdrivers.

An individual maximizes lifetime utility given by

$$U = \sum_{n=0}^{N} \rho^n u(c_n, h_n),$$

where ρ is the discount factor, c_n is consumption in period n, h_n is hours worked in period n, and $u(\cdot)$ is a per-period utility function which is increasing in consumption, decreasing in hours worked, and concave in both arguments. The lifetime budget constraint is given by

$$\sum_{n=0}^{N} (1+r)^{-1}(y_n(h_n) - p_n c_n) = 0,$$

where p_n denotes the price of consumption, r denotes the interest rate, and daily earnings $y_n(\cdot)$ is an increasing function of labor supply. The first-order conditions for this intertemporal maximization problem equate the marginal utility of lifetime income with the marginal utility of consumption and the marginal disutility of effort per unit of wage.

The problem of maximizing lifetime utility is equivalent to that of maximizing a static one-period objective function

$$v(h_n) = \lambda y_n(h_n) - g(h_n, \lambda p_n), \tag{A.1}$$

where the monetary equivalent of the disutility of effort $g(\cdot)$ is convex and λ is the lifetime marginal utility of income along the optimal path.[1] Taking one period to be a shift, we model the decision of a driver to continue working or to stop working at the end of each trip. To evaluate whether to continue working, drivers must form expectations about the additional income earned from and the additional time spent on a prospective trip.

A driver decides to end a given shift when the disutility of effort for completing an additional trip outweighs the expected fare.[2] As in Farber [2015], the stopping model can be thought of as a reduced form of a forward-looking dynamic optimization model based on expected future earnings opportunities, hours worked so far, and income earned so far. After completing t trips in $h_{i,n,t}$ hours, driver i decides to end shift n if $v(h_{i,n,t})$ exceeds the value $v(h_{i,n,t+1}) + \varepsilon_{i,n,t}$ of continuing to work for one more trip, where the error terms $\varepsilon_{i,n,t}$ are independently drawn from a distribution F. We define $d^*_{i,n,t} = v(h_{i,n,t+1}) + \varepsilon_{i,n,t} - v(h_{i,n,t})$ as the latent value of continuing for another trip and let $d_{i,n,t} = \mathbf{1}_{\left\{d^*_{i,n,t} < 0\right\}}$ indicate the decision to stop working. A key prediction of the model is that there are no daily income effects: cumulative daily earnings $y_{i,n,t} := y_n(h_{i,n,t})$ do not affect the decision to end a shift. Taking a reduced-form approximation for the value of continuing, we test the prediction that daily income effects are inconsequential by expressing the probability that driver i ends shift t at trip n non-parametrically as

$$\Pr(d_{i,n,t} = 1) = \sum_j \Big[\big(f_j(h_{i,n,t}) + \gamma_j y_{i,n,t} + X_{i,n,t}\beta_j + \mu_{i,j} \big) \mathbf{1}_{\left\{h_{i,n,t} \in H_j\right\}} \Big] + \epsilon_{i,n,t}, \tag{1.1}$$

where X consists of controls for location, time, and weather which can potentially be related to variation in earnings opportunities from continuing to work; $f(h)$ is a function of work hours; μ absorbs differences in drivers' baseline stopping tendencies; and H_j partitions the minutes of the shift into intervals to allow a time-varying relationship between each of the

[1]See the online appendix of Fehr and Goette [2007] for a derivation following Browning et al. [1985].

[2]This assumption follows Farber [2005, 2008] and Crawford and Meng [2011], who suggest that not explicitly modeling option value is behaviorally reasonable. Moreover, for convex disutility of effort, this trip-by-trip stopping rule is consistent with maximizing the static objective function as long as the wage rate $y'_n(h_n)$ does not increase too rapidly.

covariates and the probability of stopping. The model predicts that $\gamma_j = 0$ for all j, i.e., that the decision to end a shift is unrelated to cumulative daily earnings.

A.3.2 Simulation Exercise

To evaluate various approaches for estimating stopping behavior, we conduct a set of empirical Monte Carlo studies [Stigler, 1977, Huber et al., 2013]. The data for our simulations consists of a sample of over 3 million trips from 1,000 drivers. The first set of simulations considers stopping decisions that do not depend on earnings. The second set of simulations considers stopping decisions that depend on cumulative daily earnings but not on the timing of earnings. We find that the non-parametric approach in the present paper produces the expected result across all of the simulations, whereas alternative approaches from the literature may yield a significant positive or negative effect of earnings (and the timing of earnings) on stopping.

We follow the notation in Section A.3.1, where $d_{i,n,t}$ denotes the decision to stop working, $y_{i,n,t}$ denotes cumulative earnings, and $h_{i,n,t}$ denotes the number of hours driver i has worked at the end of t trips in shift n. Letting Φ denote the standard normal cumulative distribution function, we consider regression equations of the following forms:

$$\Pr(d_{i,n,t} = 1) = \sum_j \left[\left(\alpha_j h_{i,n,t} + \gamma_j y_{i,n,t} + \mu_{i,j} \right) \mathbf{1}_{\left\{ h_{i,n,t} \in H_j \right\}} \right] + \epsilon_{i,n,t} \tag{TT}$$

$$\Pr(d_{i,n,t} = 1) = \Phi(\alpha h_{i,n,t} + \gamma y_{i,n,t} + \mu_i) \tag{F-1}$$

$$\Pr(d_{i,n,t} = 1) = \Phi\left(\sum_j \alpha_j \mathbf{1}_{\left\{ h_{i,n,t} \in H_j \right\}} + \gamma y_{i,n,t} + \mu_i \right) \tag{F-2}$$

$$\Pr(d_{i,n,t} = 1) = \Phi\left(\sum_j \alpha_j \mathbf{1}_{\left\{ h_{i,n,t} \in \hat{H}_j \right\}} + \sum_j \gamma_j \mathbf{1}_{\left\{ y_{i,n,t} \in \hat{Y}_j \right\}} + \mu_i \right) \tag{F-3}$$

$$\Pr(d_{i,n,t} = 1) = \Phi\left(\sum_j \alpha_j \mathbf{1}_{\left\{ h_{i,n,t} \in \hat{H}_j \right\}} + \sum_{j,\ell} \delta_{j,\ell} \mathbf{1}_{\left\{ h_{i,n,t} \in \hat{H}_j \right\}} \mathbf{1}_{\left\{ y_{i,n,t} \in \hat{Y}_\ell \right\}} + \mu_i \right). \tag{F-4}$$

Equation (TT) corresponds to Equation (1.1) in Thakral and Tô [2018], excluding the control variables, with H_j partitioning the shift into 10-minute intervals. Equation (F-1) resembles the probit model from Farber [2005] and Crawford and Meng [2011], in

Table A.2: *Simulated stopping model: Decision rule independent of income*

	Simulation 1: stop at 9.5		Simulation 2: stop at $\bar{H}_i$	
	Effect of 20% increase in income	p-value: income coefs. = 0	Effect of 20% increase in income	p-value: income coefs. = 0
Non-parametric model				
TT	-0.0041	0.1979	-0.0058	0.9368
	(0.0030)		(0.0043)	
Linear probability model				
F-1	-0.0083	0.0000	-0.0269	0.0000
	(0.0005)		(0.0006)	
F-2	0.0005	0.2329	-0.0116	0.0000
	(0.0004)		(0.0006)	
F-3a		0.0017		0.0000
F-3b	0.0436	0.0000	-0.0202	0.0000
	(0.0056)		(0.0110)	
F-4	-0.0054	0.0000	0.0112	0.0000
	(0.0052)		(0.0186)	
Probit model				
F-1	-0.0081	0.0000	-0.0317	0.0000
	(0.0005)		(0.0009)	
F-2	0.0004	0.2467	-0.0207	0.0000
	(0.0003)		(0.0013)	
F-3a		0.0303		0.0000
F-3b	0.0185	0.0000	-0.0189	0.0000
	(0.0029)		(0.0114)	

Each row corresponds to a different regression equation defined in Section A.3.2. The row F-3a corresponds to Equation (F-3) *with the partitions* $\hat{H}$ *and* $\hat{Y}$ *over hours and income defined as in Farber [2005], and the row F-3b corresponds to Equation* (F-3) *with the partitions defined as in Farber [2015]. Simulation 1 denotes a stopping rule in which all drivers end their shifts after exceeding 9.5 hours with some noise in the stopping decision prior to that. Simulation 2 denotes a stopping rule in which all drivers end their shifts after exceeding a driver-specific quantity of hours with some noise in the stopping decision prior to that. The top panel reports estimates from Equation* (TT) *of the percentage-point increase in the probability of ending a shift at 8.5 hours when cumulative earnings is* 20 *percent higher. The middle panel reports results from Equations* (F-1) *to* (F-4) *estimated as a linear probability model. The bottom panel reports results from the probit models in Equations* (F-1) *to* (F-4)*. The income effect (columns 1 and 3) reports the estimated effect of a 20% increase in cumulative daily earnings on the probability of ending a shift after working 8.5 hours and earning $300. The income effect for the model in F-3a mechanically does not predict any effect of income on the probability of stopping after earning $300 because of how the partition is defined. The p-value (columns 2 and 4) presents the result of an F-test (Panels A and B) or* χ^2*-test (Panel C) of the null hypothesis that the income-related coefficients are jointly zero. The test imposes 1 restriction for Equations* (F-1) *and* (F-2)*, 9 restrictions for Equation* (F-3)*, 72 restrictions for Equation* (F-4)*, and 59 restrictions for Equation* (TT) *in Simulation 1, and 78 restrictions for Equation* (TT) *in Simulation 2.*

Table A.3: *Simulated stopping model: Decision rule independent of timing of income*

	Simulation 3: Pr(stop) increases in income
	Effect of 20% increase in income
TT	
Income in hour 1	0.0474
	(0.0207)
Income in hour 2	0.0243
	(0.0218)
Income in hour 3	-0.0012
	(0.0223)
Income in hour 4	0.0102
	(0.0225)
Income in hour 5	0.0215
	(0.0227)
Income in hour 6	0.0241
	(0.0212)
Income in hour 7	0.0285
	(0.0208)
Income in hour 8	0.0157
	(0.0228)
p-value: income coefs. = 0	0.0000
p-value: Equality of income coefs.	0.8464

Each row reports the estimated percentage-point change in the probability of ending a shift at 8.5 hours in response to a \$60 increase in earnings accumulated at different times during the shift from Equation (TT∗). *Simulation 3 denotes a stopping rule in which all drivers end their shifts after exceeding 9.5 hours, prior to which the drivers probability of ending a shift is an increasing function of cumulative daily earnings but does not depend on the timing of those earnings. The penultimate row presents the result of an F-test of the null hypothesis that the income-related coefficients are jointly zero. The last row tests the null hypothesis that the* $\gamma_{j,k}$ *coefficients in Equation* (TT∗) *are independent of k (for every j).*

Figure A.2: *Simulated stopping model: Distribution of p values for non-parametric specification*

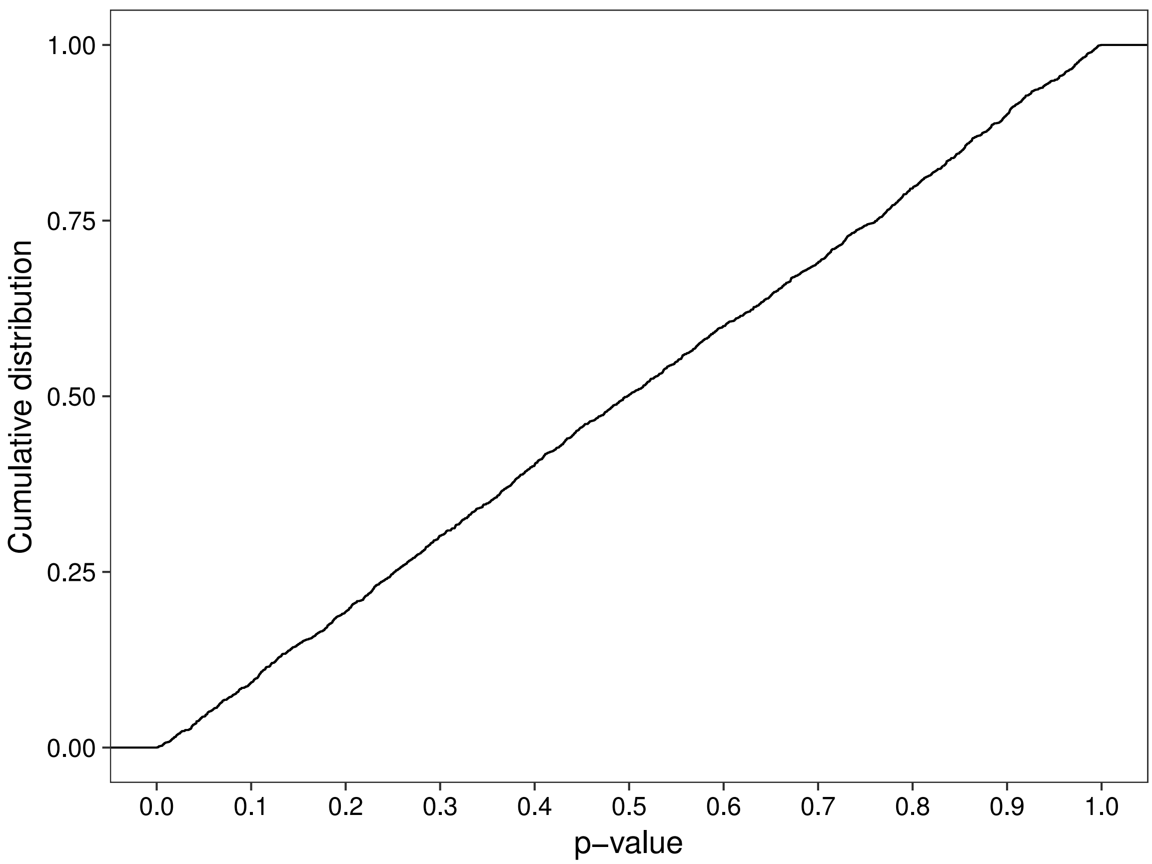

The figure depicts the results from using Equation (TT) *to test the null hypothesis that income has no effect on stopping decisions in Simulation 1 and Simulation 2 repeated 1,000 times each. Simulation 1 denotes a stopping rule in which all drivers end their shifts after exceeding 9.5 hours with some noise in the stopping decision prior to that. Simulation 2 denotes a stopping rule in which all drivers end their shifts after exceeding a driver-specific quantity of hours with some noise in the stopping decision prior to that. The curve represents the cumulative distribution of p-values.*

which income and hours are constrained to enter linearly. Equation (F-2) relaxes the constraint by allowing for a non-parametric relationship between hours and the probability of stopping. Equation (F-3) corresponds to the alternative specification in Farber [2005] when we take $\hat{H}$ and $\hat{Y}$ to partition the shift at $\{180, 360, 420, 480, 540, 600, 660, 720\}$ minutes and $\{25, 50, 75, 100, 125, 150, 175, 200, 225\}$ dollars, respectively. The main specification in Farber [2015] corresponds to Equation (F-3) and the more flexible specification in Farber [2015] corresponds to Equation (F-4) (both estimated as linear probability models) when we take $\hat{H}$ and $\hat{Y}$ to partition the shift at $\{180, 360, 420, 480, 540, 600, 660, 720, 780\}$ minutes and $\{100, 150, 200, 225, 250, 275, 300, 350, 400\}$ dollars, respectively.

We consider the following stopping rules in which decisions do not depend on earnings:

Simulation 1: End the shift with certainty at the end of a trip if hours exceeds 9.5, and stop with independent probability 0.05 at the end of any given trip that ends before 9.5 hours.

Simulation 2: Driver i ends the shift with certainty at the end of a trip if hours exceeds a driver-specific level of hours $\bar{H}_i$, and stops with independent probability 0.05 at the end of any given trip that ends before $\bar{H}_i$ hours, where we define $\bar{H}_i$ as one less than the mean hours across all of driver i's shifts in the data.

Table A.2 reports the estimated income effects in both simulations from Equation (TT), from Equations (F-1) to (F-4) estimated as linear probability model, and from Equations (F-1) to (F-3) estimated as a probit model. We refer to the specification in Farber [2005] as F-3a, and the specification in Farber [2015] as F-3b. Equation (TT) produces the expected result in both simulations that we cannot reject the null hypothesis that the income-related coefficients are jointly zero. Using Equation (F-1), Equation (F-3), or Equation (F-4) leads to the incorrect conclusion in both simulations that income significantly influences the probability of stopping, even though the data are generated precisely so that income has no effect. By controlling flexibly for hours in Equation (F-2), we cannot reject the null hypothesis of no income effects in Simulation 1 since the probability of ending a shift as a function of hours is generated to be identical across drivers; in Simulation 2, however, we incorrectly reject the null hypothesis.

We repeat the exercise 1,000 times for each simulation. While the parametric specifications overwhelmingly produce false positives by incorrectly rejecting the null hypothesis of no income effects, Equation (TT) rejects this null hypothesis at the x percent significance level about in about x percent of simulations for all $x \in (0, 1)$. Figure A.2 shows this by plotting the distribution of p-values from the non-parametric specification.

Next, we extend Equation (TT) to allow for the probability of stopping to depend on the timing of earnings, analogous to Equation (1.2):

$$\Pr(d_{i,n,t} = 1) = \sum_j \left[\left(\alpha_j h_{i,n,t} + \gamma_{j,k} \sum_k y_{i,n,t,k} + \mu_{i,j} \right) \mathbf{1}_{\{h_{i,n,t} \in H_j\}} \right] + \epsilon_{i,n,t} \qquad \text{(TT}*\text{)}$$

We simulate a stopping rule in which income does affect stopping decisions, but the timing of income is irrelevant (i.e., money is fungible).

Simulation 3: End the shift with certainty at the end of a trip if hours exceeds 9.5, and stop with independent probability $0.05 \cdot y_{i,n,t}$ at the end of any given trip that ends before 9.5 hours, where $y_{i,n,t}$ denotes cumulative daily earnings.

Table A.3 reports the estimated effects of earnings in each hour from Equation (TT). An F-test rejects the hypothesis that the income-related coefficients are jointly zero (i.e., $\gamma_{j,k} = 0$ for all j, k) but fails to reject the hypothesis that the timing of income is irrelevant (i.e., $\gamma_{j,k_1} = \gamma_{j,k_2}$ for all j, k_1, k_2).

A.4 Robustness

A.4.1 Effort

Section 1.3.4 discusses IV estimates for the income effect, instrumenting for earnings with the cumulative distance between pick-up and drop-off locations (in GPS distance or odometer miles).

As an alternative estimation strategy, we instrument for cumulative daily earnings using cumulative tips from credit-card transactions. Figures A.3 and A.4 reproduce Figures 1.6

Figure A.3: *Stopping model IV estimates: Income effect throughout the shift*

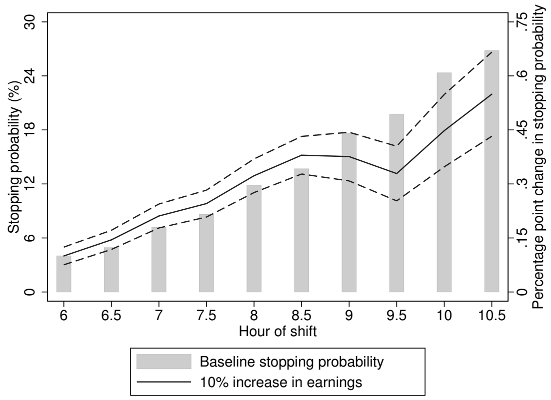

(a) *Baseline*

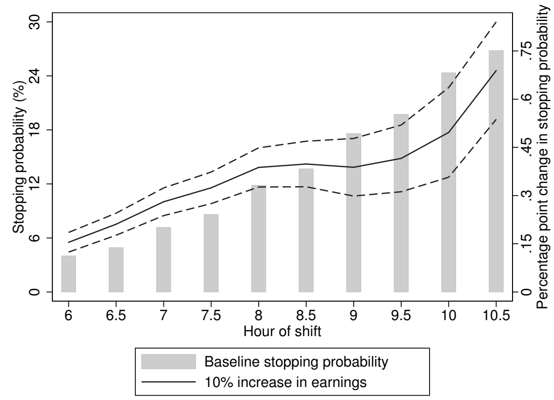

(b) *IV: GPS distance*

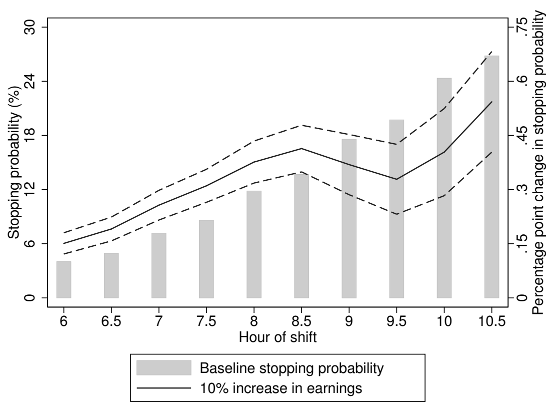

(c) *IV: Odometer distance*

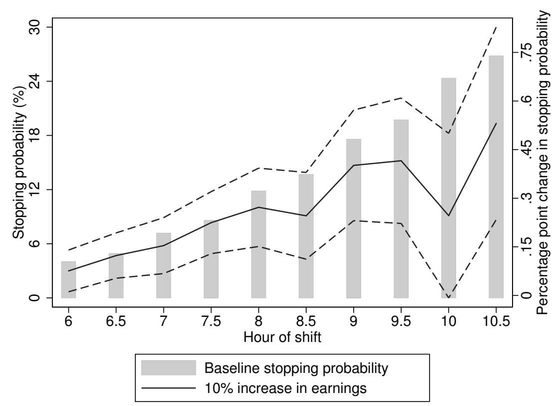

(d) *IV: Tips*

The bars, corresponding to the scale on the left, show the probability that a driver ends a shift at the specified number of hours. The solid lines, corresponding to the scale on the right, depict the marginal effect of an additional 10 *percent in earnings on the probability of stopping at various times throughout the shift. Estimates obtain from Equation* (1.1) *with controls for location, time, and weather (see Table 1.2 for details) and fixed effects for 37,460 drivers. The dashed lines represent the 95-percent confidence interval, with standard errors adjusted for clustering at the driver level.*

Figure A.4: *Stopping model IV estimates: Income effect at 8.5 hours—Timing pattern*

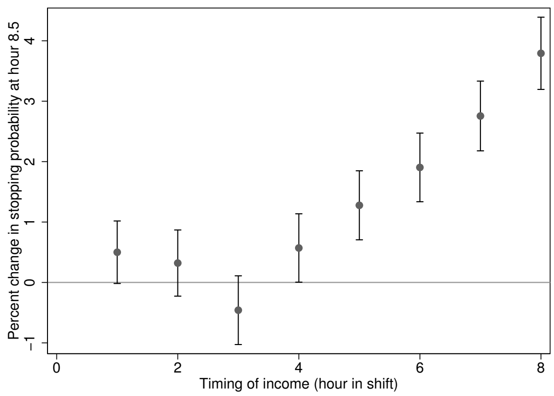

(a) *Baseline*

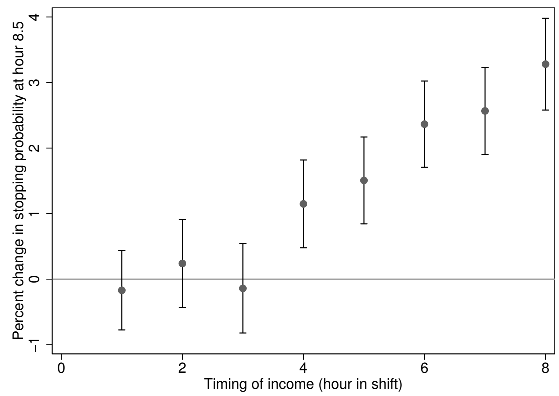

(b) *IV: GPS distance*

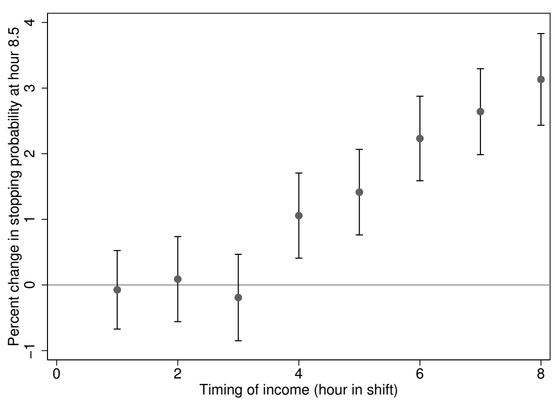

(c) *IV: Odometer distance*

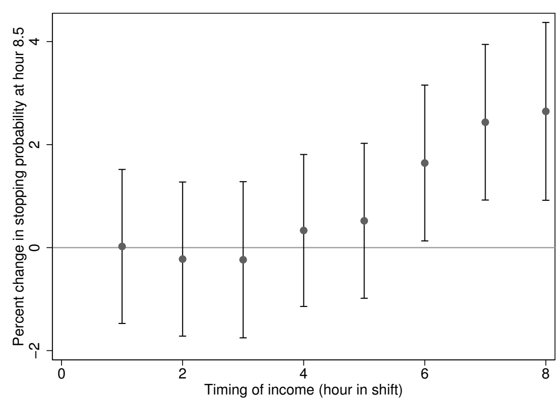

(d) *IV: Tips*

The figure depicts the percent change in the probability of ending a shift at 8.5 hours in response to a \$10 increase in earnings accumulated at different times in the shift. Estimates obtain from Equation (1.2) *with controls for location, time, and weather (see Table 1.2 for details) and fixed effects for 36,900 drivers.*

and 1.7, respectively, using the three IV strategies and find consistent income effects throughout the shift as well as the timing pattern across all specifications. The results in Figures A.3d and A.4d show smaller point estimates of the income effect, consistent with our interpretation of the effect of cumulative earnings on the probability of stopping as representing an income effect.

Figure A.5 restricts the analysis to trips that stay within Manhattan and demonstrates that the income effect persists. To understand the effect of cumulative earnings on stopping, this restriction serves several purposes. First, even though the baseline estimates contain an extensive set of controls for location (195 NTA fixed effects and an indicator for being in the zip code where the cab must be returned interacted with hour of the day), the estimates from this subsample ensure that income effects do not only appear when drivers end a trip in one of the outer boroughs (e.g., near the garage where they return the cab or their home). Second, given the dense streets of Manhattan, the variation in earnings due to distance based on differences in driving speed plausibly arises due to traffic conditions unrelated to the driver's decisions to exert additional effort.

A.4.2 Learning

Within-day learning

The results in Figures 1.7 and 1.8 suggest that drivers react differently to money earned in different hours of the shift, which we interpret as a violation of fungibility. Differences in behavior due to timing of payment could also result from learning. To explain the timing pattern with a learning story, we would have to assume that drivers tend to ignore recent experiences in the market and instead rely on earnings earlier in the shift to predict future opportunities. The data do not support the view that more recent market conditions are less relevant for predicting future market conditions (see Figure 1.5).[3] Instead, a more plausible learning effect would bias the results away from finding stronger effects on stopping in

[3]We compute an autocorrelation using market wages, as driver-specific wages are endogenous to their stopping decisions.

Figure A.5: *Stopping model IV estimates: Income effect at 8.5 hours—Timing pattern for Manhattan trips*

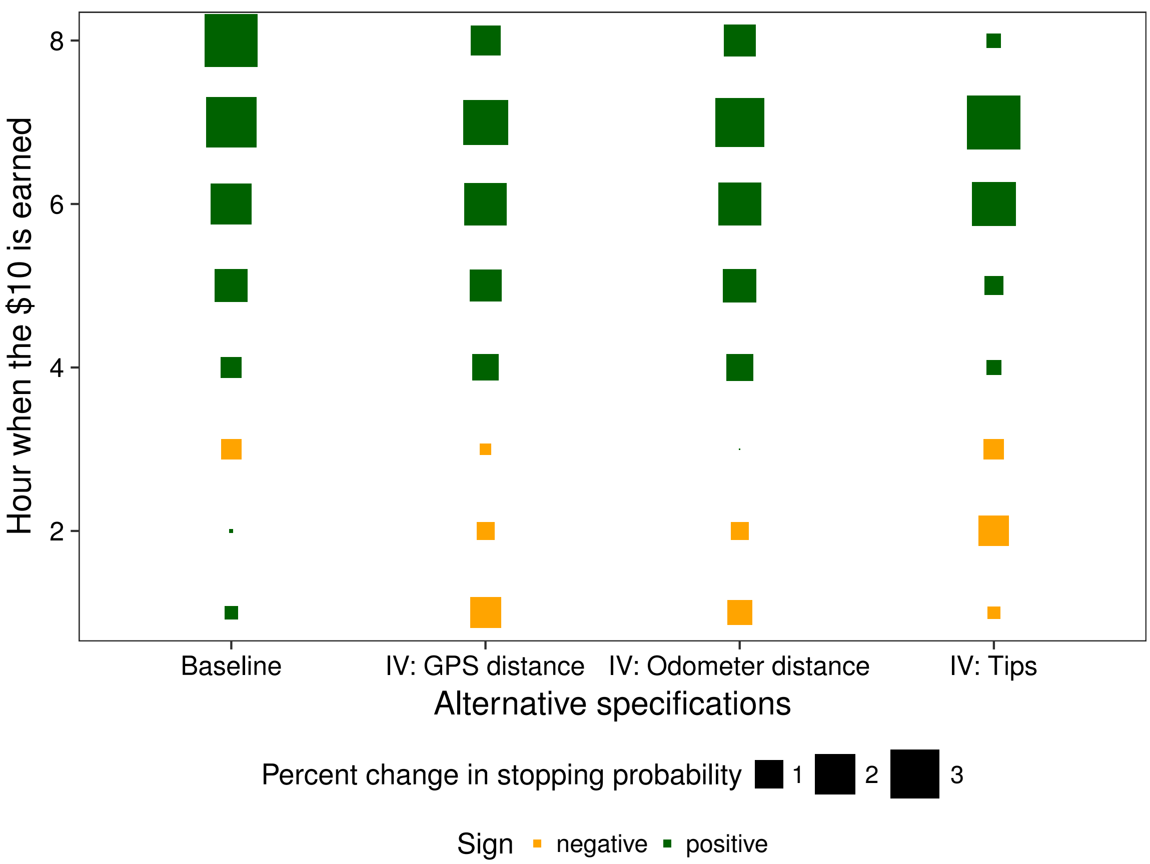

The figure depicts the effect of an additional $10 in earnings accumulated at different times in the shift (vertical axis) on the probability of stopping at 8.5 hours for trips that start and end within Manhattan under various estimation strategies. Estimates obtain from Equation (1.2) *with controls for location, time, and weather and driver fixed effects (see Table 1.2 for details). Each square has area proportional to the corresponding percent change in the probability of stopping.*

response to more recent earnings. Insofar as within-day learning influences behavior, the estimated violation of fungibility understates the true effect.

Across-day learning

In the model from Section A.3.1, earnings on one day convey no information about earnings on another day, so intertemporal optimization is equivalent to maximization of a static one-period objective function. If higher earnings correlate with plentiful opportunities on the next day, then a driver may decide to work less on one day to conserve energy to work more on the next day. The insignificant autocorrelation in the transitory component to daily wages in Figure A.6 suggests little scope for this type of intertemporal substitution to drive the relationship between earnings and quitting.

In a related setting, Agarwal et al. [2015] find that daily income distributions for Singaporean cabdrivers are independent of income shocks in the previous days, which suggests a limited role for intertemporal substitution.

A.4.3 Option value

The trip-by-trip stopping model relies on the assumption that the option value of continuing to drive is sufficiently small (or that drivers ignore option value). The pattern in Figure A.1 suggests that this assumption may not be reasonable for night-weekend shifts, when wages rise substantially and predictably over time, but may be quite reasonable for night-weekday shifts, when the typical wage profile is nonincreasing. Figure A.7 plots the percent change in the probability of stopping estimated on four separate groups of shifts (day-weekday shifts, day-weekend shifts, night-weekday shifts, and night-weekend shifts). We see a significant negative effect of earnings on quitting for night-weekend shifts, for which the assumptions of the stopping model likely do not apply, and significant positive income effects for day shifts and night-weekday shifts, for which the assumptions of the stopping model likely do apply.

Figure A.6: *Autocorrelation of residualized daily market wage*

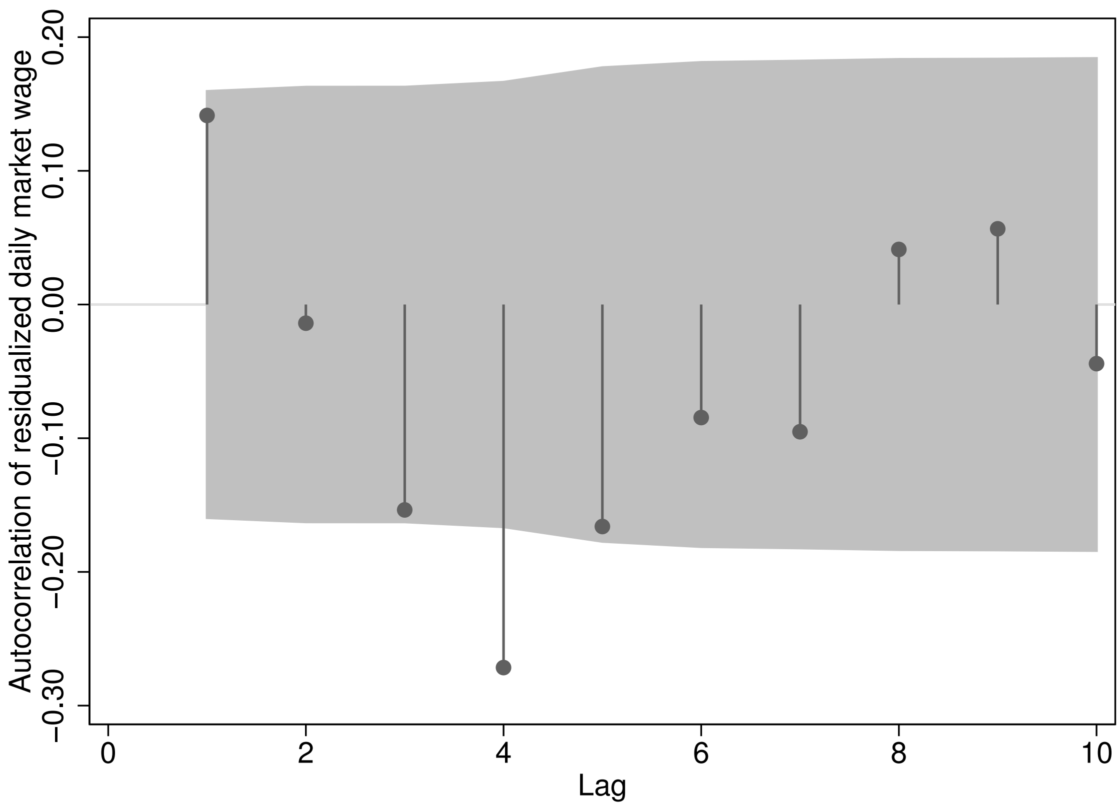

The figure depicts the autocorrelation of daily market wages indexed by day of the year in 2013. The daily market wage is the sum of the minute market wage in each calendar day, with the minute market wage computed as in Figure 1.4. The daily market wage is residualized from a regression on a set of time and weather effects: day of week, week of year, an indicator for federal holidays, an indicator for whether it rains during that day, and indicators for high (over 80 degrees Fahrenheit) and low (under 30 degrees Fahrenheit) average daily temperature. The shaded region denotes a 95-percent confidence band.

Figure A.7: *Stopping model estimates: Income effect by shift type*

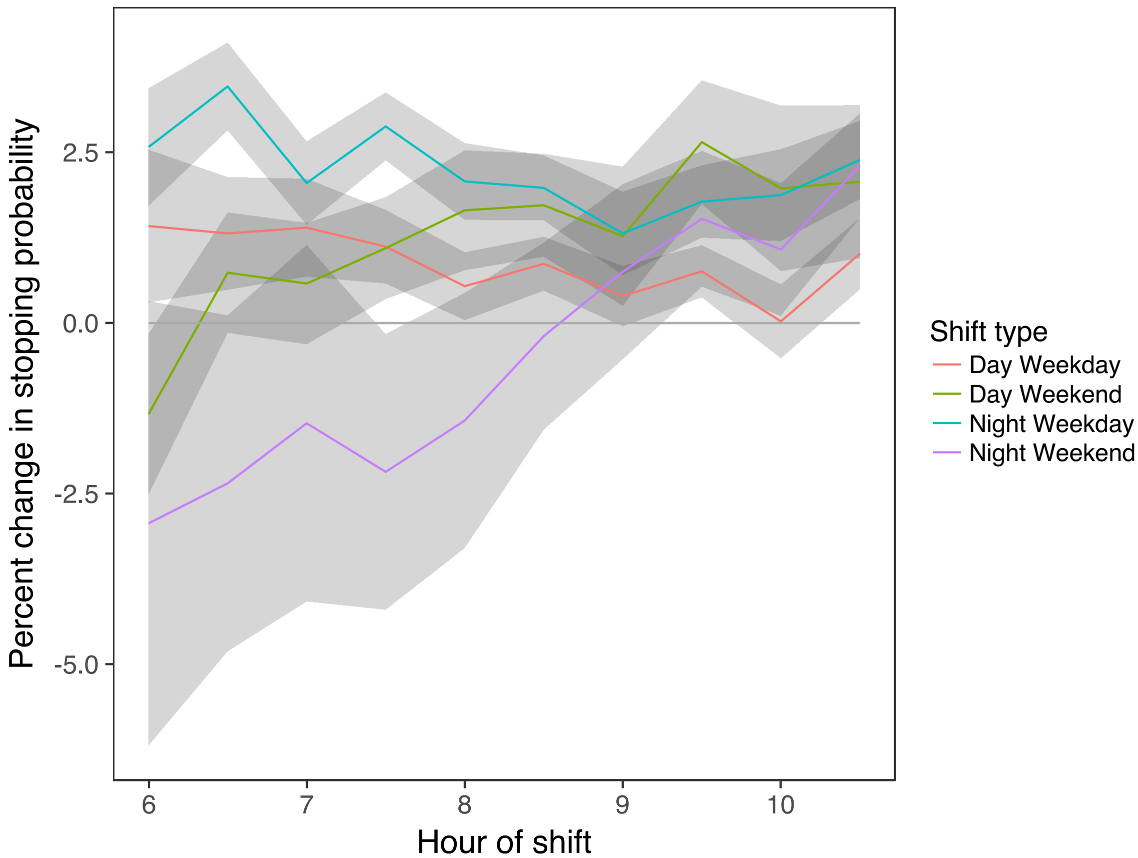

The figure depicts the percent change in the probability of stopping at various times throughout the shift in response to a 10 percent increase in cumulative earnings. Each line represents estimates of Equation (1.1) (see Figure 1.6 for details) restricted to the corresponding group of shifts: day-weekday, day-weekend, night-weekday, and night-weekend. Day shifts start between 4 AM and 10 AM, and night shifts start between 2 PM and 8 PM. Weekend shifts consist of night shifts on Friday and Saturday as well as day shifts on Saturday and Sunday.

A.4.4 Experience

To investigate the effects of experience, Table A.4 restricts the sample to drivers with over 100 shifts and separates each driver's shifts into deciles based on the date. The first row corresponds to the first 10 percent of each driver's shifts, while the last row corresponds to the last 10 percent of each driver's shifts. The estimates show consistent magnitudes of the income effects as well as violations of fungibility across all levels of experience. For comparison, Haggag et al. [2017] document significant learning among cabdrivers in a relatively short time horizon, with productivity differences between new and experienced drivers vanishing after 17 to 62 shifts.

A.4.5 Measurement Error

Observability of shift ending

Our empirical approach reveals a decrease in labor supply under the assumption that all shifts end as soon as the driver drops off the last passenger. However, the data do not distinguish between a driver who ends a shift immediately after dropping off their last passenger and a driver who spends time searching for another fare unsuccessfully. The conclusion that drivers respond to higher cumulative earnings with a reduction in labor supply might be overstated if drivers spend relatively more time searching before quitting in high-income shifts.

Explaining the patterns in our data by the fact that drivers may spend unrecorded amounts of time searching before quitting would require that finding a passenger is more difficult at the end of a shift in which the driver earns more. As noted in the discussion of unobserved effort in Section 1.3.4, however, a high-income shift is more likely to be one in which the driver generally spends less time searching for passengers. The negative correlation between the share of working hours spent searching for passengers and total earnings in a shift provides suggestive evidence that drivers are unlikely to spend relatively more time searching for a passenger before ending a shift when earnings are high. As an

Table A.4: *Stopping model estimates: Income effect at 8.5 hours—Within-driver experience*

	(1)	(2)			
	Overall	Hour 2	Hour 4	Hour 6	Hour 8
0–10% shifts	0.2883	0.2984	-0.1003	0.0515	0.1307
	(0.1056)	(0.1558)	(0.1655)	(0.1633)	(0.1723)
10–20% shifts	0.1874	-0.0933	-0.2133	0.0919	0.5044
	(0.1085)	(0.1516)	(0.1568)	(0.1568)	(0.1622)
20–30% shifts	0.4576	-0.1167	0.1844	0.3713	0.4914
	(0.1077)	(0.1528)	(0.1590)	(0.1590)	(0.1673)
30–40% shifts	0.4760	0.0993	-0.1107	0.1451	0.4835
	(0.1085)	(0.1588)	(0.1662)	(0.1597)	(0.1655)
40–50% shifts	0.3911	0.2189	0.0388	0.1822	0.7382
	(0.1131)	(0.1618)	(0.1719)	(0.1711)	(0.1781)
50–60% shifts	0.1848	-0.0180	0.0700	0.1722	0.4549
	(0.1152)	(0.1674)	(0.1751)	(0.1758)	(0.1840)
60–70% shifts	0.3531	0.2247	0.1595	0.2413	0.5915
	(0.1130)	(0.2247)	(0.1595)	(0.2413)	(0.5915)
70–80% shifts	0.3097	0.0959	0.0022	0.5388	0.4347
	(0.1162)	(0.1647)	(0.1727)	(0.1678)	(0.1765)
80–90% shifts	0.5603	0.0479	0.3023	0.4168	0.6786
	(0.1150)	(0.1638)	(0.1698)	(0.1689)	(0.1783)
90–100% shifts	0.4625	-0.0130	0.3062	0.3309	0.5580
	(0.0805)	(0.1236)	(0.1284)	(0.1267)	(0.1335)

Specification (1) reports estimates from Equation (1.1) *of the percentage-point increase in the probability of ending a shift at 8.5 hours when cumulative earnings is* 10 *percent higher. Specification (2) reports estimates from Equation* (1.2) *of the percentage-point change in the probability of ending a shift at 8.5 hours in response to a $10 increase in earnings accumulated at different times in the shift. Each row corresponds to a different level of experience, with the last row denoting shifts with the greatest experience for each driver. The control variables consist of the full set from Table 1.2. Standard errors reported in parentheses are adjusted for clustering at the driver level.*

alternative measure of the difficulty of searching at the end of a shift, we use the amount of time that the driver spent searching for the last passenger. Indeed we find a similar pattern: drivers spend an average of 11.2 minutes searching for their last passenger among shifts in the bottom decile of earnings, compared with only 10.1 minutes among shifts in the top decile.[4]

The evidence suggests that the income effect does not emerge from the fact that our dataset does not report the amount of time that a driver spends working at the end of a shift. If anything, drivers may spend relatively more minutes working after dropping off the last passenger on a low-wage shift, which would imply that the reduction in labor supply that we observe in response to higher cumulative earnings underestimates the true income effect.[5]

Taking breaks

Figure A.8 classifies breaks as long periods of time without a passenger and presents estimates of the stopping model from Equation (1.2) with additional controls for minutes spent on break. Farber [2005] uses the following thresholds to classify waiting times as breaks: 30 minutes between Manhattan fares; 60 minutes between non-airport, non-Manhattan fares; 90 minutes between airport fares. We test the sensitivity of the income effect by uniformly adjusting the thresholds of waiting time for defining breaks by 15 minutes in either direction. Figure A.8 verifies that the results remain unchanged using these definitions of breaks.

Instead of directly controlling for break time as in Farber [2005] (e.g., if taking breaks constitutes an outcome of earnings), we can re-estimate the stopping model using breaks as the dependent variable. A decrease in the probability of taking a break in response to

[4]The result is also similar if we measure the difficulty of finding a passenger after a given trip by computing the number of minutes spent searching averaged across all drivers whose trips end in the same minute.

[5]The concern that shift ending times are unobservable might be more relevant for elasticity-based analyses of daily labor-supply decisions, since the fact that drivers spend more unrecorded minutes searching for passengers during shifts with lower average wages could bias elasticity estimates in the positive direction.

Figure A.8: *Stopping model estimates: Income effect at 8.5 hours—Timing pattern with alternative definitions of breaks*

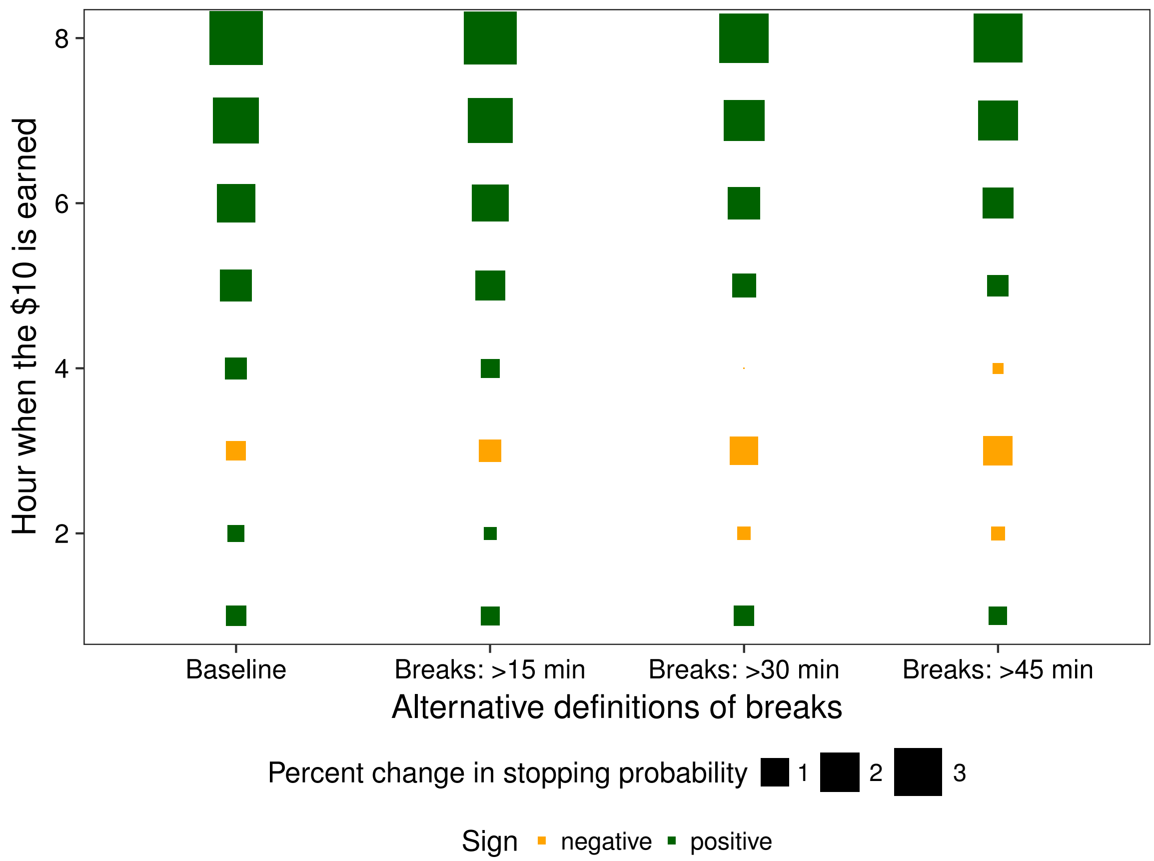

The figure depicts the effect of an additional $10 in earnings accumulated at different times in the shift (vertical axis) on the probability of stopping at 8.5 hours from Figure 1.7, with controls for break time under various definitions of breaks (horizontal axis). The first column replicates the baseline specification, which controls for minutes spent working, including indicators for the number of minutes with passengers in each hour. The second column uses the following minimum thresholds to classify time spent without a passenger as breaks: 15 minutes between Manhattan fares; 45 minutes between non-airport, non-Manhattan fares; 75 minutes between airport fares. The third column uses the following thresholds: 30 minutes between Manhattan fares; 60 minutes between non-airport, non-Manhattan fares; 90 minutes between airport fares. The fourth column uses the following thresholds: 45 minutes between Manhattan fares; 75 minutes between non-airport, non-Manhattan fares; 105 minutes between airport fares. Each square has area proportional to the corresponding percent change in the probability of stopping.

Table A.5: *Maximum likelihood estimates: Loss aversion—Stochastic reference points*

	Estimates
Disutility of trip	0.0348
	(0.0022)
Error term distribution σ	0.0578
	(0.0031)
Loss aversion over income L_I	1.1164
	(0.0280)
Loss aversion over hours L_H	6.3827
	(0.0978)
Adjustment θ	0.9060
	(0.0141)
Log-likelihood	-350,205
Likelihood ratio test: $L_I = L_H = 1$	¡0.001
Likelihood ratio test: $\theta = 1$	¡0.001
Likelihood ratio test: $L_I = L_H$	¡0.001

This table presents maximum likelihood estimates of Equation (2.7) *for the objective function in Equation* (2.4) *with stochastic reference points. The estimation sample consists of over 1.2 million shifts from 36,900 drivers. See the note to Table 2.1 for additional details. The last three rows contain the p-value from likelihood ratio tests of the following null hypotheses: (i) the baseline model, (ii) a static reference point, and (iii) a single coefficient of loss aversion.*

additional earnings might lead to concerns that the stopping model incorrectly attributes the effect of hours worked to the effect of income, but the evidence points against this. We find that an additional 10 percent in earnings corresponds to an increase of 0.0072 to 0.0756 percentage points in the probability of taking a break at 8.5 hours. We find an increase in the probability of taking a break at earlier hours of the shift and no significant change in the probability of taking a break at later hours of the shift.

A.5 Additional results for structural estimation

The model of loss aversion in Section 2.2 makes two simplifying assumptions: first, the targets I^r and H^r represent point expectations, and second, utility is piecewise linear in gains and losses. A stochastic reference point would consist of the distribution of earnings

Table A.6: *Maximum likelihood estimates: Loss aversion—With diminishing sensitivity*

	Estimates
Disutility of trip	0.0816
	(0.0082)
Error term distribution σ	0.0795
	(0.0031)
Loss aversion over income L_I	1.2763
	(0.0170)
Loss aversion over hours L_H	2.6971
	(0.0239)
Log-likelihood	-459,015
Likelihood ratio test: $L_I = L_H = 1$	¡0.001
Likelihood ratio test: $L_I = L_H$	¡0.001

This table presents maximum likelihood estimates of Equation (2.7) *for the objective function in Equation* (2.4) *with diminishing sensitivity in the gain-loss function. The estimation sample consists of over 1.2 million shifts from 36,900 drivers. See the note to Table 2.1 for additional details. The last two rows contain the p-value from likelihood ratio tests of the following null hypotheses: (i) the baseline model, and (ii) a single coefficient of loss aversion.*

and hours for each shift. We approximate this by using quartiles of the distribution of earnings and hours for each driver and day of the week. The results in Table A.5 show a larger coefficient of loss aversion but still an important role for adaptive reference points and a smaller degree of loss aversion over income. Allowing for diminishing sensitivity corresponds to an objective function that exhibits convexity in losses and concavity in gains. We use the power function

$$n(x \mid x^r) = \left(\mathbf{1}_{\{x > x^r\}} + \lambda \mathbf{1}_{\{x < x^r\}}\right)\left(v_x(x) - v_x(x^r)\right)^{\zeta},$$

where we follow Hastings and Shapiro [2013] by calibrating the parameter ζ to 0.88 [Kahneman and Tversky, 1979]. Table A.6 re-estimates the loss-aversion model with diminishing sensitivity. Overall we find that relaxing the simplifying assumptions in the loss-aversion model does not change the conclusions about the importance of adaptive reference points and the smaller degree of loss aversion over income.

Appendix B

Appendix to Chapter 3

B.1 Ex-post impossibility

Proof of Proposition 1. We will show by example that no mechanism can guarantee ex-post efficiency or ex-post elimination of justified envy for every arrival process.

Consider the following example with three periods (0, 1, and 2), three buildings (α, β, and γ), and three applicants (A, B, and C).

Assume that each building gives applicant A the highest priority. The applicants' preferences are given in Table B.1.

Let the arrival process be specified as follows. In each period, a unit becomes available with certainty: the unit that becomes available in period 0 is in building α; the unit that becomes available in period 1 is equally likely to be in building β or in building γ; and the unit that becomes available in period 2 is equally likely to be in the building from which no unit has become available yet or in building α.

Suppose applicant A is assigned building α in period 0. With probability $1/2$, building β becomes available in period 1. There are two cases to consider. First, suppose β is assigned to applicant B. Then with probability $1/2$, building γ arrives in period 2 and is allocated to applicant C. Notice that the allocation is ex-post inefficient because applicants A and B prefer to switch: since $(\beta, 1) \succ_{\text{A}} (\alpha, 1)$ and $(\alpha, 1) \succ_{\text{B}} (\beta, 1)$, we see that A and B prefer to leave

Table B.1: *Preferences for applicants* A, B, *and* C *in Proposition 1*

$\succ_A$	$\succ_B$	$\succ_C$
$(\beta, 0)$	$(\alpha, 0)$	$(\alpha, 0)$
$(\beta, 1)$	$(\alpha, 1)$	$(\alpha, 1)$
$(\gamma, 0)$	$(\gamma, 0)$	$(\gamma, 0)$
$(\alpha, 0)$	$(\alpha, 2)$	$(\alpha, 2)$
$(\beta, 2)$	$(\beta, 0)$	$(\beta, 0)$
$(\gamma, 1)$	$(\gamma, 1)$	$(\gamma, 1)$
$(\alpha, 1)$	$(\beta, 1)$	$(\beta, 1)$
$(\gamma, 2)$	$(\gamma, 2)$	$(\gamma, 2)$
$(\alpha, 2)$	$(\beta, 2)$	$(\beta, 2)$

Preferences for applicants A, B, *and* C *listed in order from most-preferred to least-preferred. Agent* A*'s preference can be generated by the utility function* $u_A(b, t) = f(b) - 3t$*, where* $f(\alpha) = 1$*,* $f(\beta) = 6$*,* $f(\gamma) = 2$*; applicant* B*'s and applicant* C*'s preference can be generated by the utility function* $u_B(b, t) = u_C(b, t) = g(b) - 3t$*, where* $g(\alpha) = 8$*,* $g(\beta) = 1$*,* $g(\gamma) = 3$.

their assigned buildings to switch with each other in period 1 (or any subsequent period). Second, suppose building β is assigned to applicant C when it becomes available in period 1. Then with probability 1/2, building γ arrives in period 2 and is allocated to applicant B. Again, the allocation is ex-post inefficient because applicants A and C prefer to switch.

The analysis is similar if applicant B or applicant C is assigned building α in period 0. Regardless of which applicant is assigned building α in period 0, there is always some realization of the arrival process in which a Pareto improvement can be found. The remaining cases of the argument are summarized in Table B.2.

In each case, there is justified envy since applicant A (who has the highest priority) prefers another applicant's allocation. □

B.2 Properties of MWP

Proof of Proposition 2. (i) Consider the strategy of applicant $a \in A$. We will show that truthful preference revelation is weakly dominant. Since the order in which allocations are made

Table B.2: *Example used in proof of Proposition 1*

Period 0		Period 1		Period 2		Switch	Envy
α	A	β	B	γ	C	A, B	A, B
α	A	β	C	γ	B	A, C	A, C
α	B	γ	A	β	C	A, C	A, C
α	B	γ	C	α	A	A, C	A, C
α	C	γ	A	α	B	A, B	A, B
α	C	γ	B	β	A	A, B	A, B

An example in which any assignment rule can lead to a violation of ex-post efficiency and ex-post elimination of justified envy. Buildings are denoted α, β, and γ. Agents are denoted A, B, *and* C, *and their preferences are given in Table B.1. Each building gives applicant* A *the highest priority. The penultimate column specifies which applicants could trade to obtain a Pareto improvement in each setting, and the last column specifies whether some applicant justifiably envies another.*

depends only on the priority ordering and *not* on the applicants' preferences, we restrict our attention to the step in which applicant a has the highest priority among those who are unmatched. In this step, we see that a is assigned her most-preferred unit among those that have not yet been assigned. Therefore there is no incentive to misreport preferences.

(ii) Let the allocation resulting from MWP be given by μ and consider a reallocation μ'. It suffices to show that there is some applicant who prefers the original allocation μ. Denote the set of applicants whose allocations differ across μ and μ' by $A' = \{a : \mu'(a) \neq \mu(a)\}$. Let a_0 denote the applicant who is assigned first among the applicants in A'. For any $a' \in A'$ with $a' \neq a_0$, the unit $\langle b_{a'}, r_{a'} \rangle$ is assigned at a later step than $\langle b_{a_0}, r_{a_0} \rangle$ is assigned. Since MWP is strategy-proof by (i), applicant a_0 strictly prefers the original allocation $\left(b_{a_0}, t_0 + \tau_{b_{a_0}, t_0}(r_{a_0})\right)$ over $\left(b_{a'}, t_0 + \tau_{b_{a'}, t_0}(r_{a'})\right)$ as desired.

(iii) Suppose applicant a' has higher priority than another applicant a. According to MWP, a' receives an assignment (i.e., a unit or a place on some waiting list) before a does. This implies that $\langle b_a, r_a \rangle$ is available when $\langle b_{a'}, r_{a'} \rangle$ is assigned. By strategy-proofness, we have that a' prefers $\langle b_{a'}, r_{a'} \rangle$ to the unit assigned to a. In other words, a' does not envy the lower priority applicant a. □

B.3 Uniqueness

For a given history h^τ of arrivals up to time τ, a preference profile $\succ$ over building-time pairs induces a profile of rankings $\succ' = (\succ^t)_t$ over units, where we denote the profile of period-t rankings by $\succ^t$. Given an arrival history and the rankings over units induced by reported preferences, a mechanism φ specifies an allocation $\mu^t_\succ$ in each period.

A mechanism φ is *non-bossy* if for any agent a and any pair of preference profiles $\succ$ and $\hat{\succ}$, we have $\mu_{\hat{\succ}_a, \succ_{-a}} = \mu_\succ$ whenever $\mu_{\hat{\succ}_a, \succ_{-a}}(a) = \mu_\succ(a)$. In a non-bossy mechanism, no agent can change the allocation without changing her own assignment.

Let $\succ^\tau$ be a ranking over units in period τ, and let π_τ be any permutation over the set of units. Define $\hat{\succ}^\tau = \pi_\tau \succ^\tau$ and $\hat{\succ}^t = \succ^t$ for $t \neq \tau$. We say that φ satisfies *within-period neutrality* if $\mu^\tau_{\hat{\succ}'}(a) \neq \emptyset$ implies $\mu^\tau_{\hat{\succ}'}(a) = \pi_\tau \mu^\tau_{\succ'}(a)$. Now let π be any permutation over the set of units, and define $\tilde{\succ}^t = \pi \succ^t$. We say that φ satisfies *across-period neutrality* if $\mu_{\tilde{\succ}'} = \pi \mu_{\succ'}$ whenever $\succ^t$ is independent of t. A mechanism is *neutral* if it satisfies both within- and across-period neutrality.

For any unit x and a history h^t, let $L^t_{\succ_a}(x) = \{y : x \succ^t_a y\}$ denote the lower contour set of x under ranking $\succ^t_a$. We say that x is *relatively better* under $\succ'_a$ than under $\succ_a$ conditional on history h^t if the lower contour set expands when the preference changes from $\succ_a$ to $\succ'_a$, i.e., if $L^t_{\succ_a}(x) \subseteq L^t_{\succ'_a}(x)$.

A mechanism φ is *monotonic* if $L^t_{\succ_a}(\mu_\succ(a)) \subseteq L^t_{\hat{\succ}_a}(\mu_\succ(a))$ implies $\mu^t_\succ = \mu^t_{\hat{\succ}}$. In other words, if for every agent the mechanism yields under preference profile $\succ$ an allocation in period t which is relatively better under $\hat{\succ}$, then the mechanism yields the same allocation under $\hat{\succ}$.

Proposition 3 states that any strategy-proof, non-bossy, and neutral mechanism is an extended multiple-waitlist procedure. As Thakral [2016] notes, the proof is related to that of an analogous result due to Svensson [1999] for a setting in which objects do not arrive stochastically.

Proof of Proposition 3. Denote the set of all units by X. Fix a history of arrivals, and let X_t

denote the set of units that arrive in period t. The proof uses the following lemmas, which will be proven afterward.

Lemma 1. *If φ is strategy-proof and non-bossy, then φ is monotonic.*

Lemma 2. *When all agents' preferences are identical, the allocation under a strategy-proof, non-bossy, and neutral mechanism is ex-ante Pareto efficient.*

First suppose the preferences are such that the induced rankings over units are common (i.e., identical across agents) and persistent (i.e., identical across time): $\succ_a^t = \succ_*$ for all a and all t. By Lemma 2, the assignment $\mu_\succ$ is efficient. Without loss of generality, label the agents so that $i < j$ whenever (i) $T(a_i) < T(a_j)$, i.e., whenever a_i is matched before a_j; or (ii) $T(a_i) = T(a_j) =: T$ and $\mu_\succ^T(a_i) \succ_* \mu_\succ^T(a_j)$, i.e., a_i and a_j are matched in the same period but the unit assigned to a_i is preferable. Since φ is neutral, for any profile $\succ'$ in which preferences are common and persistent, the order is preserved: $\mu_{\succ'}(a_1) \succ_* \mu_{\succ'}(a_2) \succ_* \cdots$.

Now consider an arbitrary preference profile $\rhd$. Given a sequence $t(j)$, let y_j be the most-preferred unit among $X \setminus \{y_{j'}\}_{j'<j}$ under ranking $\rhd_{a_j}^{t(j)}$. Choose the mapping t so that (i) $t(j) \geq t(j-1)$; and (ii) $t(j) > t(j-1)$ only if $\bigcup_{t<t(j)} X_t \subseteq \{y_{j'}\}_{j'<j}$. Define the ranking $\hat{\succ}_*$ by $y_j \hat{\succ}_* y_{j+1}$ for all j, and let $\hat{\succ}'$ be the corresponding profile of common and persistent rankings. By the result above from Lemma 2, we have that $\mu_{\hat{\succ}'}(a_j) = y_j$ for all j.

The above shows that under common and persistent rankings, a non-bossy and strategy-proof allocation mechanism is an extended multiple-waitlist procedure. The conditions on $t(\cdot)$ require that all elements of X_t (and perhaps more units) are assigned in period t. We will proceed to show that by imposing neutrality, this characterization extends to the arbitrary profile of rankings.

Suppose x satisfies $\mu_{\hat{\succeq}'}(a_j) \hat{\succ}_* x$. Then by construction $x \in X \setminus \{y_{j'}\}_{j'<j}$, so we have $\mu_{\hat{\succ}'}(a_j) \trianglerighteq_{a_j}^{t(j)} x$. Hence, by Lemma 1, we have $\mu_{\rhd'} = \mu_{\hat{\succ}'}$ as desired. □

Proof of Lemma 1. Consider a preference profile $\succ'$ such that for each agent a, the allocation $\mu_\succ(a)$ is relatively better under $\succ'$ than under $\succ$. To prove that φ is monotonic, it suffices to show that $\mu_{\succ'} = \mu_\succ$.

We start by showing that $\mu_{\succ'_a,\succ_{-a}} = \mu_\succ$ for any agent a. Since φ is strategy-proof, agent a reports her preference truthfully, so we have $\mu_\succ(a) \succeq_a \mu_{\succ'_a,\succ_{-a}}(a)$. Now since $\mu_\succ(a)$ is relatively better under $\succ'$,

$$\mu_\succ(a) \succeq'_a \mu_{\succ'_a,\succ_{-a}}(a).$$

Again by strategy-proofness, we have

$$\mu_{\succ'_a,\succ_{-a}}(a) \succeq'_a \mu_\succ(a).$$

The strategy-proofness condition thus implies that a is indifferent between $\mu_{\succ'_a,\succ_{-a}}$ and $\mu_\succ$. From non-bossiness we obtain the result that $\mu_{\succ'_a,\succ_{-a}} = \mu_\succ$.

Repeating the argument for each of the remaining agents gives the desired result that the mechanism yields the same allocation under $\succ'$. □

Proof of Lemma 2. Let $\succ_*$ denote the common preference. Without loss of generality, label the agents so that $i < j$ whenever (i) $T(a_i) < T(a_j)$, i.e., whenever a_i is matched before a_j; or (ii) $T(a_i) = T(a_j) =: T$ and $\mu^T_\succ(a_i) \succ^T_* \mu^T_\succ(a_j)$, i.e., a_i and a_j are matched in the same period but the unit assigned to a_i is preferable.

Suppose on the contrary that $\mu_\succ$ is not efficient. This implies that there exists a unit $\xi = \langle \xi_b, \xi_r \rangle$ such that (i) ξ is ranked higher in period t than the unit that a_i is assigned for some i, i.e., $\xi \succ^t_* \mu_\succ(a_i)$, where $t := T(a_i)$; and (i) ξ is not assigned to any agent in period t, i.e., $\xi \neq \mu^t_\succ(a)$ for any a.

Define $n = \min\{i : \xi \succ^t \mu^t_\succ(a_i) \neq \emptyset\}$ to be the agent who receives the most-preferred unit among those which are ranked below ξ but still allocated under $\mu_\succ$. Let $\hat{\succ}^t_*$ be a ranking over units that induces a ranking over units which coincides with that of $\succ^t_*$ except that $\mu_\succ(a_n) \hat{\succ}^t_* \xi$ (i.e., the order of ξ and $\mu_\succ(a_n)$ is switched). Denote by π_t the permutation that switches the rankings of ξ and $\mu_\succ(a_n)$ so that $\hat{\succ}^t = \pi_t \succ^t$.

The argument consists of two steps. First we will use Lemma 1 to show that the assignments under $\succ$ and $\hat{\succ}$ must be identical. Next, using the neutrality property, we will obtain a contradiction.

We begin by showing that $\mu^t_\succ(a)$ is relatively better under $\hat{\succ}^t$ than under $\succ^t$ in period

t for every a. Let x satisfy $\mu_{\succ}(a) \succeq_* x$. If $x = \mu_{\succ}(a_n)$ and $a = a_n$, then trivially we have $\mu_{\succ}(a_i) \hat{\succeq}_* x$, so suppose otherwise. By construction, note that (i) $\succ^t$ and $\hat{\succ}^t$ produce identical rankings over the units excluding ξ and $\mu_{\succ}(a_n)$, and (ii) $\mu_{\succ}(a)$ is either ranked higher than both ξ and $\mu_{\succ}(a_n)$ or neither of them. This implies $\mu_{\succ}(a) \hat{\succeq}^t_* x$. Since φ is monotonic by Lemma 1, we have that $\mu_{\succ} = \mu_{\hat{\succ}}$.

Now since φ satisfies within-period neutrality, we have $\mu^t_{\hat{\succ}}(a_n) = \xi$, which contradicts the result from Lemma 1. Therefore we conclude that $\mu_{\succ}$ is efficient. $\square$

B.4 Acyclicity and Generalized MWP

Proof of Proposition 4. (i) We will show that truthful preference revelation is weakly dominant for each applicant $a \in A$. Since the priority ordering is acyclic, there is at most one applicant $\hat{a}$ who has higher priority than a at some buildings but lower priority at the others. If there is no higher-priority applicant, then applicant a receives her most-preferred unit (among those that are available) when she reaches the top of the centralized waiting list. Otherwise, if there is a higher-priority applicant, then a receives her most-preferred unit unless $\hat{a}$ prefers the same unit and has higher priority for the building. In that case, a is assigned her most-preferred unit among those that remain after $\hat{a}$ selects a unit. Since the event that a receives her most-preferred unit among those that are available depends only on the priorities and the other applicants' stated preferences, there is no incentive for a to misreport preferences.

(ii) For any reallocation μ', it suffices to show that there is some applicant who prefers the original allocation μ resulting from the generalized MWP. Let a_0 denote the applicant who is offered a unit first among the set of applicants $A' = \{a : \mu'(a) \neq \mu(a)\}$ whose allocations differ across μ and μ'. There is at most one applicant $\hat{a}_0$ who has higher priority than a_0 at some buildings but lower priority at the others. If such $\hat{a}_0$ does not exist, or if a prefers a unit in a building at which $\hat{a}_0$ does not have higher priority, then the proof proceeds as in Proposition 2. Otherwise, $\hat{a}_0$ receives her most-preferred unit among those that are available at the time of assignment, which implies that $\hat{a}_0$ prefers the original allocation

$\left(b_{\hat{a}_0}, t_0 + \tau_{b_{\hat{a}_0},t_0}(r_{\hat{a}_0})\right)$ over $\left(b_{a'}, t_0 + \tau_{b_{a'},t_0}(r_{a'})\right)$ for any $a' \in A'$ as desired.

(iii) Suppose applicant a' has higher priority than another applicant a at the building b_a (where a is assigned), and consider the time at which a' receives an assignment. If a' prefers a unit in a building at which she has has the highest priority, then the proof proceeds as in Proposition 2: a' receives an assignment under generalized MWP before a does, so $\langle b_a, r_a \rangle$ is available when $\langle b_{a'}, r_{a'} \rangle$ is assigned, which means that a' prefers $\langle b_{a'}, r_{a'} \rangle$ to the unit assigned to a. Now suppose that a' prefers a unit in a building at which another applicant $\hat{a}'$ has higher priority.[1] In the case that $\hat{a}' \neq a$, the argument is the same as before because a' receives an assignment before a does. Otherwise, we have $\hat{a}' = a$, i.e., that a has higher priority than a' at some building. Since a' prefers a unit in a building at which a has higher priority, a receives an assignment before a' does. However, $\langle b_a, r_a \rangle$ was available to a' (since a' has higher priority at b_a by assumption) but not chosen, which implies that a' prefers $\langle b_{a'}, r_{a'} \rangle$. In all cases, a' does not envy the lower priority applicant a. □

Proof of Proposition 5. Proposition 4 implies that the acyclicity condition is sufficient for the existence of an allocation mechanism that is strategy-proof, efficient, and eliminates justified envy. We now show that if the priority orderings violate the acyclicity condition, then there does not exist a strategy-proof allocation mechanism that satisfies both efficiency and the elimination of justified envy.

If $\succ_B$ does not satisfy the acyclicity condition, then there exist buildings α and β such that the priority orderings form a cycle; that is, there exist applicants A, B, and C such that:

$$\text{A} \succ_\alpha \text{B} \succ_\alpha \text{C},$$

$$\text{and C} \succ_\beta \text{A}.$$

Let the applicants' preferences be as given in Table B.3. Consider the following deterministic arrival process: from building α, a unit becomes available in period 0 and another unit becomes available in period 2; from building β, a unit becomes available in period 1.

[1] As noted earlier, acyclicity implies that there is at most one such applicant.

Table B.3: *Preferences for applicants* A, B, *and* C *in Proposition 5*

$\succ_{\text{A}}$	$\succ_{\text{B}}$	$\succ_{\text{C}}$
$(\beta, 0)$	$(\alpha, 0)$	$(\beta, 0)$
$(\beta, 1)$	$(\alpha, 1)$	$(\alpha, 0)$
$(\alpha, 0)$	$(\beta, 0)$	$(\beta, 1)$
$(\beta, 2)$	$(\alpha, 2)$	$(\alpha, 1)$
$(\alpha, 1)$	$(\beta, 1)$	$(\beta, 2)$
$(\alpha, 2)$	$(\beta, 2)$	$(\alpha, 2)$

Preferences for applicants A, B, *and* C *listed in order from most-preferred to least-preferred. Agent* A*'s preference can be generated by the utility function* $u_{\text{A}}(b, t) = f(b) - 2t$, *where* $f(\alpha) = 1$ *and* $f(\beta) = 4$; *applicant* B*'s preference can be generated by the utility function* $u_{\text{B}}(b, t) = g(b) - 2t$, *where* $g(\alpha) = 4$ *and* $g(\beta) = 1$; *applicant* C*'s preference can be generated by the utility function* $u_{\text{C}}(b, t) = h(b) - 2t$, *where* $h(\alpha) = 1$ *and* $h(\beta) = 2$.

Suppose applicant A is assigned building α in period 0. If B is assigned β in period 1, then the allocation is inefficient because A and B prefer to switch. Likewise, if C is assigned β in period 1, then the allocation is inefficient because A and C prefer to switch.

The analysis is similar if applicant B or applicant C is assigned building α in period 0. Regardless of which applicant is assigned building α in period 0, either there is a Pareto improvement or some applicant justifiably envies another. The remaining cases of the argument are summarized in Table B.4.

This demonstrates the necessity of the acyclicity condition. □

B.5 Stochastic mechanisms

Proof of Proposition 6. Necessity can be demonstrated using the same preferences and arrival process as in the proof of Proposition 5. Assume $\succ_B$ does not satisfy acyclicity so that there exist buildings α and β, and applicants A, B, and C such that A $\succ_\alpha$ B $\succ_\alpha$ C and C $\succ_\beta$ A. The applicants' preferences are given in Table B.3. As before, the arrival process is deterministic: a unit in building α arrives in period 0; a unit in building β arrives in period 1; and another unit in building α arrives in period 2.

Table B.4: *Example used in necessity proof of Proposition 5*

Period 0		Period 1		Period 2		Switch	Envy
α	A	β	B	α	C	A, B	—
α	A	β	C	α	B	A, C	—
α	B	β	A	α	C	—	C, A
α	B	β	C	α	A	—	A, B
α	C	β	A	α	B	—	B, C
α	C	β	B	α	A	—	A, C

An example in which no assignment rule can satisfy Pareto efficiency and the elimination of justified envy when priority orderings violate acyclicity. Buildings are denoted α and β. Agents are denoted A, B, and C, and their preferences are given in Table B.3. Building α ranks applicant A above C, with B in between; but building β ranks C above A. The penultimate column specifies which applicants could trade to obtain a Pareto improvement in each setting, and the last column specifies whether some applicant justifiably envies another.

We begin by determining the assignment of applicant A. If A is assigned a positive probability of $(\alpha, 2)$, then A would justifiably lottery-envy any applicant who is assigned a positive probability of $(\alpha, 0)$. If A is assigned a positive probability of $(\beta, 1)$, then C must assigned positive probability of either $(\alpha, 0)$ or $(\alpha, 2)$: in the former case, B has justified lottery-envy towards C; and in the latter case, C has justified lottery-envy towards A. This leaves us with the conclusion that A must be assigned $(\alpha, 0)$ with certainty.

Now either B is assigned $(\alpha, 2)$ with certainty or B is assigned a positive probability of $(\beta, 1)$: in the former case, the assignment is inefficient since A and C would prefer to switch; in the latter case, B has justified lottery-envy towards C.

This establishes the necessity of acyclicity. Sufficiency follows from the same construction as in Proposition 4. □

B.6 PHA-k mechanisms

Proof of Proposition 7. Consider the following example with 2 buildings, $B = \{b_0, b_1\}$, each with one unit; $k+2$ applicants, $A = \{a_i\}_{i=0}^{k+1}$; and $k+2$ periods.

Assume that the buildings' common priority list ranks applicant a_i higher than applicant

Table B.5: *Preferences for each applicant a_i in Proposition 7*

$\succ_{a_0}$	$\succ_{a_1}$	$\cdots$	$\succ_{a_{k-1}}$	$\succ_{a_k}$	$\succ_{a_{k+1}}$
$(b_{k+1}, k+1)$	(b_k, k)	$\cdots$	$(b_2, 2)$	$(b_1, 1)$	$(b_0, 0)$
(b_k, k)	$(b_{k-1}, k-1)$	$\cdots$	$(b_1, 1)$	$(b_0, 0)$	$(b_{k+1}, k+1)$
$(b_{k-1}, k-1)$	$(b_{k-2}, k-2)$	$\cdots$	$(b_0, 0)$	$(b_{k+1}, k+1)$	(b_k, k)
$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$
$(b_2, 2)$	$(b_1, 1)$	$\cdots$	$(b_5, 5)$	$(b_4, 4)$	$(b_3, 3)$
$(b_1, 1)$	$(b_0, 0)$	$\cdots$	$(b_4, 4)$	$(b_3, 3)$	$(b_2, 2)$
$(b_0, 0)$	$(b_{k+1}, k+1)$	$\cdots$	$(b_3, 3)$	$(b_2, 2)$	$(b_1, 1)$

Preferences for applicants $\{a_i\}_{i=0}^{k+1}$ listed in order from most-preferred to least-preferred. The most-preferred unit for applicant a_i is in building b_{k-i+1} which arrives in period $k-i+1$. Each applicant a_i prefers the unit in building b_r over the unit in building b_{r-1} for all $r \not\equiv -i \pmod{k+2}$.

a_{i+1} for all i.

Units arrive deterministically: in periods 0 through k, a unit in building b_1 becomes available with certainty; in period $k+1$, a unit in building b_0 becomes available with certainty.

Applicant a_0 prefers the unit in building b_0 that arrives in period $k+1$, and all other applicants prefer units that arrive earlier.

Since applicant a_0 cannot attain any unit outside building b_1, she refuses no offers and receives a unit in building b_1 immediately. Applicants a_1 through a_{k+1} also do not refuse any offers, as they prefer units that arrive earlier.

Applicant a_0 prefers the unit in building b_1 and therefore justifiably envies a_{k+1}. Matching applicant a_0 with the unit in building b_0 and matching each of the remaining applicants with the room that arrives one period earlier would be a Pareto improvement. □

Proof of Proposition 7. Consider the following example with $k+2$ buildings, $B = \{b_i\}_{i=0}^{k+1}$, each with one unit; $k+2$ applicants, $A = \{a_i\}_{i=0}^{k+1}$; and $k+2$ periods.

Assume that the buildings' common priority list ranks applicant a_i higher than applicant a_{i+1} for all i. The applicants' preferences are given in Table B.5.

Units arrive deterministically: in period i, the unit in building i becomes available with certainty.

Applicant a_0 refuses k offers and receives a unit in building b_k, since all units that become available sooner are less desirable. For $i = 1, \ldots, k-1$, applicant a_i refuses $k-i$ offers; her first-choice unit (in building $k-i+1$) will be taken by applicant a_{i-1}, who has higher priority, so applicant a_i receives her second-choice unit (in building b_{k-i}). Likewise, applicant a_k accepts the offer of a unit in building b_0, since a_{k-1} will accept the unit in building 1. This leaves the unit in building b_{k+1} for applicant a_{k+1}. The allocation procedure concludes with each applicant receiving her second-choice unit.

Since it is possible to redistribute the units so that each applicant receives her most-preferred unit, the assignment is inefficient. Furthermore, the procedure fails to eliminate justified envy since applicant a_0 has the highest priority but prefers the unit assigned to applicant a_{k+1}. $\square$