

Online Appendix to Daily Labor Supply and Adaptive Reference Points

Neil Thakral
Brown University

Linh T. Tô
Harvard University

May 2019*

Abstract

This document contains appendix material for Thakral and Tô (2019).

*Thakral: Department of Economics, Brown University, Box B, Providence, RI 02912 (email: neil_thakral@brown.edu). Tô: Department of Economics, Harvard University, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138 (email: linhto@fas.harvard.edu).

Contents

A	Data and Summary Statistics	2
B	Elasticity Estimates	5
C	Model and Simulation	6
C.1	A Model of Daily Labor Supply	6
C.2	Simulation Exercise	8
D	Reduced-Form Analysis and Robustness	11
D.1	Variation in Earnings and Instrumental-Variables Approach	11
D.2	Learning	12
D.3	Liquidity Constraints	13
D.4	Experience	14
D.5	Measurement of Hours	15
E	Additional Results for Structural Estimation	17
E.1	Calibration of Risk Aversion	17
E.2	Jump at the Target	17
E.3	Separating Gain-Loss Utility from Consumption Utility	18
E.4	Identification of Structural Parameters	19
E.5	Alternative Specification of the Income Target	20
E.6	Model with Hours Targeting	22
E.7	Stochastic Reference Points and Diminishing Sensitivity	23
F	Salience Model	24
F.1	Salience	24
F.2	Adaptive Reference Points	27

A Data and Summary Statistics

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: 0.30 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 define a shift as a sequence of consecutive trips that are not more than six hours apart from each other (Haggag and Paci, 2014). We remove shifts with trips that have been flagged with errors or shifts that 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, we restrict our sample to shifts in which minutes of work during the first and second

hours do not fall above 55 or below 10 minutes, respectively, so that the distribution of minutes worked in those hours is more comparable across hours, thus minimizing situations in which drivers work one planned trip at the beginning of a shift without continuing to work a regular shift, and we restrict our sample to shifts that stay within the five boroughs in NYC, consisting of 94 percent of the remaining shifts.

The local linear regressions following Equation (TT) (e.g., column 1 of Table 1) puts nonzero weight on 2,256,799 trips from 2,111,589 shifts and 37,164 drivers. The regressions following Equations (F-1) to (F-3) (e.g., columns 2 to 4 of Table 1) takes a two-fifteenths sample of trips to get 15,952,627 trips in 5,040,269 shifts for 37,512 drivers.

Appendix Table 1 provides summary statistics at the trip level and at the shift level for all 127 million NYC taxi trips in 2013 in the cleaned data. Over 85 percent of all trips start and end in Manhattan, and the median ride takes 10 minutes. Cabdrivers earn a median fare of about \$9.50, with 90 percent of fares falling below \$22. We observe tips for the 55 percent of fares paid using a credit card.

A driver collects an average of \$2.48 in tips per trip, but there is substantial variation in the rate of tipping as the scatterplot in Appendix Figure 4 shows. Given any fare between the minimum fare of \$2.50 and \$60, the figure shows the density of tips between \$0.50 and \$20. We see higher concentrations at round numbers and fixed fractions of the fare, between 10 percent and 35 percent. The vertical line at \$54 fares corresponds to trips between Manhattan and JFK international Airport, where we also see substantial variation in tips. 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.¹

¹For 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.

B 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.

Appendix Figure 3 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 Appendix Table 2 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 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.

C Model and Simulation

C.1 A Model of Daily Labor Supply

We present a neoclassical model of intertemporal utility maximization with time-separable utility following Farber (2005) and formulate testable predictions about daily labor-supply decisions.

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

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

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. See the online appendix of Fehr and Goette (2007) for a derivation following Browning et al. (1985).

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. As in Farber (2005, 2008, 2015), this discrete-choice formulation can be thought of as representing a reduced form of a forward-looking dynamic optimization model based on hours worked so far on the shift, expectations about future earnings possibilities, and other variables that could affect preferences for work. 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. After completing t trips in h_{int} hours, driver i decides to end shift n if $v(h_{int})$ exceeds the value $v(h_{int+1}) + \varepsilon_{int}$ of continuing to work for one more trip, where the error terms ε_{int} are independently drawn from a distribution F . We define $d_{int}^* = v(h_{int+1}) + \varepsilon_{int} - v(h_{int})$ as the latent value of continuing for another trip and let $d_{int} = \mathbf{1}_{\{d_{int}^* < 0\}}$ indicate the decision to stop working.

A key prediction of the model is that there are no daily income effects: cumulative daily earnings $y_{int} := y_n(h_{int})$ 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

$$\Pr(d_{int} = 1) = f(h_{int}) + \gamma(h_{int})y_{int} + X_{int}\beta(h_{int}) + \mu_i(h_{int}) + \varepsilon_{int}, \quad (\text{TT})$$

where $f(\cdot)$ represents the baseline hazard; X consists of controls that can potentially be related to variation in earnings opportunities from continuing to work, such as location, time, and weather; and μ absorbs differences in drivers' baseline stopping tendencies. This additive hazards model (Aalen, 1989) has an intuitive interpretation that the effect of each control variable can be expressed in monetary-equivalent terms in the driver's optimization problem. In a proportional hazards model, by contrast, the effect of a unit increase in a covariate would be multiplicative with respect to the hazard rate.

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.

C.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 Appendix C.1, where d_{int} denotes the decision to stop working, y_{int} denotes cumulative earnings, and h_{int} 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:

$$\Pr(d_{int} = 1) = \sum_j [(\alpha_j h_{int} + \gamma_j y_{int} + \mu_{i,j}) \mathbf{1}_{\{h_{int} \in H_j\}}] + \epsilon_{int} \quad (\text{TT})$$

$$\Pr(d_{int} = 1) = \Phi(\alpha h_{int} + \gamma y_{int} + \mu_i) \quad (\text{F-1})$$

$$\Pr(d_{int} = 1) = \Phi\left(\sum_j \alpha_j \mathbf{1}_{\{h_{int} \in H_j\}} + \gamma y_{int} + \mu_i\right) \quad (\text{F-1}^*)$$

$$\Pr(d_{int} = 1) = \Phi\left(\sum_j \alpha_j \mathbf{1}_{\{h_{int} \in \hat{H}_j\}} + \sum_j \gamma_j \mathbf{1}_{\{y_{int} \in \hat{Y}_j\}} + \mu_i\right) \quad (\text{F-2})$$

$$\Pr(d_{int} = 1) = \Phi\left(\sum_j \alpha_j \mathbf{1}_{\{h_{int} \in \hat{H}_j\}} + \sum_{j,\ell} \delta_{j,\ell} \mathbf{1}_{\{h_{int} \in \hat{H}_j\}} \mathbf{1}_{\{y_{int} \in \hat{Y}_\ell\}} + \mu_i\right). \quad (\text{F-3})$$

These correspond to the analogous equations in the main text, except with

the control variables omitted. We take H_j to partition the shift into 10-minute intervals so that Equation (TT) above corresponds to the analogous specification in the main text with uniform weights over a 10-minute window of time during the shift as Section 2.1 discusses. Equation (F-1) resembles the probit model from Farber (2005) and Crawford and Meng (2011), in which income and hours are constrained to enter linearly. Equation (F-1*) relaxes the constraint by allowing for a non-parametric relationship between hours and the probability of stopping. Equation (F-2) 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-2) and the more flexible specification in Farber (2015) corresponds to Equation (F-3) (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.

Appendix Table 3 reports the estimated income effects in both simulations from Equation (TT), from Equations (F-1) to (F-3) estimated as linear probability model, and from Equations (F-1) to (F-2) estimated as a probit model. We refer to the specification in Farber (2005) as F-2a, and the specification in Farber (2015) as F-2b. 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-2), or Equation (F-3) 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-1*), 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 $x \in (0, 1)$. Appendix Figure 5 shows this by plotting the distribution of p -values from the non-parametric specification, and we find a uniform distribution as expected.

Next, we extend Equation (TT) to allow for the probability of stopping to depend on the timing of earnings, analogous to Equation (1):

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

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

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

Appendix Table 4 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).

D Reduced-Form Analysis and Robustness

D.1 Variation in Earnings and Instrumental-Variables Approach

This section discusses the sources of variation in earnings. After controlling for time of day, location, and weather, about 80 percent of the residual variation in cumulative daily earnings can be accounted for by variation in distance (which influences the meter fare), tips, airport fares (fixed rate between Manhattan and JFK), and the number of trips (fixed cost per ride). A positive correlation between the residual variation and effort would challenge the interpretation of the effect of additional accumulated earnings on the probability of stopping as evidence of a daily income effect. Effort can arise in various forms. Drivers may increase effort in their interactions with riders to influence tips, or they may drive more aggressively to cover greater distance. They may also choose to pursue strategies to pick up different types of passengers, such as those who might be heading to the airport or those who might be heading a shorter distance.

First, although these factors may not be exogenous to the driver's stopping decision, we can assess the role of effort by analyzing whether these factors operate as intervening pathways connecting additional accumulated income to stopping decisions. We find that adding any subset of these variables as control variables in the reduced-form stopping model does not change the estimated magnitude of the relationship between accumulated earnings and quitting. If these sources of variation in earnings constitute different measures of effort and the quitting response to additional earnings operated through the effort channel, then we might expect that controlling for them would reduce the estimated effect of additional earnings on quitting.

Second, we isolate variation in income that arises due to differences in driving speed. We instrument for accumulated earnings in each hour using the average driving speed of other drivers' that pick up a passenger from the same NTA in the same quarter of an hour, which has a sufficiently large F-statistic for the first-stage regressions. We restrict to trips that stay within Manhattan, where variation in

driving speed plausibly arises due to traffic conditions unrelated to the driver’s decisions to exert additional effort. Appendix Table 9 presents both least-squares estimates and IV estimates. Even though the baseline estimates (Table 1) contain an extensive set of controls for location (195 NTA fixed effects), the least-squares estimates on the Manhattan subsample (Appendix Table 9 column 1) confirms 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). The IV estimates (Appendix Table 9 column 2) corroborate the finding that earnings in the most recent hour significantly increase the probability of ending a shift.

D.2 Learning

Within-day learning

The results in Figures 3 and 4 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 Appendix Figure 7).² Instead, a more plausible learning effect would bias the results away from finding stronger effects on stopping in 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 Appendix C.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

²We compute an autocorrelation using market wages, as driver-specific wages are endogenous to their stopping decisions.

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 Appendix Figure 8 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 that intertemporal substitution does not play an important role.

D.3 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. (2018) 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. 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. The result that drivers react differently to earnings accumulated 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.³ Although our data do not include information on ownership, we classify a

³The average price of an independent medallion in 2013 was approximately \$967,000. Source:

driver as an owner-driver if (i) the driver operates exactly one cab, and (ii) no other driver shares that cab. Our classification yields a subsample of owner-drivers, as we exclude those who lease to another driver.

The estimates in column (3) of Appendix Table 7 suggest that liquidity constraints do not confound the income effects we observe. In fact, owner-drivers exhibit stronger income effects. This is also consistent with Appendix Table 6, which shows stronger income effects for drivers with more flexibility in their hours decisions.

D.4 Experience

To investigate the effects of experience, we separate drivers into two groups: a less experienced group of 2,361 drivers who first appear in the data in April 2013 and drive at least 50 shifts until December 2013, and a more experienced group of 32,547 drivers who first appear in the data before April 2013.

Appendix Table 10 uses the across-driver definition of experience above and presents evidence against the hypothesis that our estimated positive effect of accumulated earnings on quitting represents the behavior of inexperienced drivers. In fact, the point estimates appear larger for the more experienced group of drivers.

Appendix Table 11 provides a complementary analysis using an within-driver definition of experience to analyze how the income effects vary with experience for the less experienced drivers. We define experience in three different ways: the number of months since the driver started, the number of shifts since the driver started, and an indicator variable which is equal to one for half of each driver’s shifts, those that occur later in the year. We do not find consistent evidence that income effects diminish as drivers begin to gain experience, though the estimates are noisy due to the small sample size. While we also find that the income effect and recency effect do not decrease over time for the more experienced group, we exclude them from our analysis in Appendix Table 11. Prior work documents significant learning among cabdrivers early in their experience profile, with productivity differences between new and experienced drivers vanishing after 17 to 62 shifts (Haggag et al., 2017). Thus,

http://www.nyc.gov/html/tlc/downloads/pdf/2014_taxicab_fact_book.pdf.

finding no effect for the more experienced group may mask learning effects that arise early in the experience profile.

D.5 Measurement of Hours

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

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.⁴

Taking breaks

Appendix Figure 9 classifies breaks as long periods of time without a passenger and presents estimates of the stopping model from Equation (1) 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. Appendix Figure 9 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 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.

⁴The 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.

E Additional Results for Structural Estimation

E.1 Calibration of Risk Aversion

This section obtains a conservative estimate of the degree of risk aversion over annual income implied by the daily income effect. We define utility over annual income as $u(I) = \frac{I^{1-\rho}}{1-\rho}$. Our structural estimates imply a decrease in the marginal utility of income of at least 50 percent from the beginning of a day to the end of the day (see Table 3). Taking a cabdriver's annual earnings to be \$25,000 and daily earnings to be \$150, we obtain $\frac{25150^{-\rho}}{25000^{-\rho}} = \frac{1}{2}$ or $\rho \approx 115$. This likely understates the degree of risk aversion required to rationalize our results. First, with average net earnings in our sample around \$280 and leasing fees around \$130, daily earnings would likely fall below \$150 due to gasoline costs. Second, net annual earnings likely exceed \$25,000. Third, the drop in the marginal utility of income is closer to 60 percent. Each of these factors would lead to an increase in our estimate of ρ .

E.2 Jump at the Target

This section provides suggestive evidence that drivers use an updated income target which incorporates more information from earlier in the shift. We estimate the reduced-form stopping model allowing for jumps at the shift-level expected income target I_0^r and targets that differentially incorporate gains and losses accumulated in the earliest several hours or latest several hours of the shift. Across specifications, we find that the probability of stopping significantly increases when income passes the target that updates more in response to earlier experiences.

In an analysis based on Table 3 from Crawford and Meng (2011), Appendix Table 12 presents the estimated effect of earnings or hours exceeding their targets on the probability of ending a shift. Let I_0 denote the daily-level income target from (Crawford and Meng, 2011). Let Δ^4 denote the difference between realized and expected earnings in the first four hours of the shift, and let Δ^8 denote the difference between realized and expected earnings in the first eight hours of the shift. Define updated reference point $R_{i,j} = I_0 + i \cdot \Delta^4 + j \cdot (\Delta^8 - \Delta^4)$. We consider various cases: $i, j \in \{0, 50, 100\}$,

$i, j \in \{0, 33, 66, 100\}$, $i, j \in \{0, 25, 50, 75, 100\}$, and $i, j \in \{0, 10, 25, 50, 75, 90, 100\}$. We denote $I^{early>late} = \sum_{i>j} R_{i,j}$ and $I^{early<late} = \sum_{i<j} R_{i,j}$.

Across all specifications, we find $I^{early>late}$ significantly exceeds both I_0 and $I^{early<late}$, providing evidence that points towards an updated reference point that incorporates earlier earnings to a greater extent than later earnings. By contrast, we find a significant effect for $I^{early<late}$ and I_0 in only one out of four specifications. In addition, we find a statistically significant coefficient on cumulative earnings in all specifications, suggesting that the significant effect of earnings may not arise entirely due to whether income is above or below the target. The effect of passing the updated income target $I^{early>late}$ on the probability of stopping ranges between one-third and three-fourths of the effect of passing the hours target, in line with the relative comparison between the income and hours targets in Table 3 of Crawford and Meng (2011), although we find smaller magnitudes for both.

E.3 Separating Gain-Loss Utility from Consumption Utility

Based on Crawford and Meng (2011), given choices to continue or stop working after each trip, we define

$$\begin{aligned} a_{1,it} &= \mathbf{1}_{\{I_{i,t+1} \leq I_{i,t}^r\}} (I_{i,t+1} - I_{i,t}^r) - \mathbf{1}_{\{I_{i,t} \leq I_{i,t}^r\}} (I_{i,t} - I_{i,t}^r) \\ a_{2,it} &= \mathbf{1}_{\{I_{i,t+1} > I_{i,t}^r\}} (I_{i,t+1} - I_{i,t}^r) - \mathbf{1}_{\{I_{i,t} > I_{i,t}^r\}} (I_{i,t} - I_{i,t}^r) \\ b_{1,it}(\nu) &= \mathbf{1}_{\{H_{i,t+1} \leq H_{i,t}^r\}} (H_{i,t+1}^{\nu+1} - (H_i^r)^{\nu+1}) - \mathbf{1}_{\{H_{i,t} \leq H_{i,t}^r\}} (H_{i,t}^{\nu+1} - (H_i^r)^{\nu+1}) \\ b_{2,it}(\nu) &= \mathbf{1}_{\{H_{i,t+1} > H_{i,t}^r\}} (H_{i,t+1}^{\nu+1} - (H_i^r)^{\nu+1}) - \mathbf{1}_{\{H_{i,t} > H_{i,t}^r\}} (H_{i,t}^{\nu+1} - (H_i^r)^{\nu+1}) \end{aligned}$$

and obtain likelihood functions of the form

$$\Phi \left((1 - \eta + \eta\lambda)a_{1,it} + a_{2,it} - (1 - \eta + \eta\lambda) \frac{\psi}{1 + \nu} b_{1,it}(\nu) - \frac{\psi}{1 + \nu} b_{2,it}(\nu) \right). \quad (2)$$

The general formulation allows for separate coefficients of loss aversion over income and hours, as in Appendix E.6 and column (2) of Appendix Table 14. The case $L_I = L_H = L$ corresponds to the loss-aversion model with the same coefficient of loss

aversion on each dimension of utility, as Kőszegi and Rabin (2006) suggest. The case $L_H = 0$ corresponds to having loss aversion only over income, as we discuss in the main text (e.g., Table 3).

We now discuss why the relative importance of gain-loss utility η and the degree of loss aversion λ are not separately identified. The choice to continue or stop does not change the reference point relative to which the decision maker assesses concurrent gains or losses (note that both $a_{1,it}$ and $a_{2,it}$ depend on the same I_t^r). In other words, the decision maker uses a reference point inherited from past behavior, as in the case of the choice-acclimating personal equilibrium (CPE) from Kőszegi and Rabin (2007). Thus, despite adding two parameters to Equation (3), the model consists of only one additional degree of freedom since behavior depends only on the ratio between utility from losses and gains, namely

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

As Barseghyan et al. (2013) note, although the model contains two parameters λ and η , these parameters always appear as the product $\eta(\lambda - 1)$ under CPE.

E.4 Identification of Structural Parameters

In a model without reference dependence, the first-order condition for the optimal stopping problem would state that a driver stops working if $\mathbb{E}[w] < \psi H^\nu$ or $\log w < \log \psi + \nu \log H$. This equation highlights that variation in wages and variation in work hours both contribute to the identification of ν . It also highlights that the two moments react differently to changes in the two parameters. For example, the marginal effect of an increase in w on ν depends on H , but not the effect of an increase in w on ψ . Adding reference dependence over income changes the marginal benefit of continuing (the left-hand side of the first-order condition above) when accumulated earnings exceeds the target. This observation clarifies why a larger effect of accumulated earnings on quitting would tend to increase the estimate of L but also

decrease the estimates of ψ and ν . The data appear to exhibit enough independent variation in the timing of income that the ratio between the effect of recent earnings to earlier earnings primarily identifies θ .

Appendix Figure 10 elucidates the link between the structural parameters and the sources of variation in the data. Each panel plots the relationship between a moment of the data and the parameters of the structural model. Each parameter varies within one standard deviation of its estimated value (i.e., one unit on the horizontal axis corresponds to one-tenth of a standard deviation).

We estimate the parameters jointly using maximum likelihood. We find a strong link between the structural parameters and the associated moments of the data, providing important sources of identification.

Panel (a) shows how increasing the disutility of effort ψ parameter would increase the effect of working for an additional 10 minutes on quitting. Increasing ν would also lead the model to predict a stronger relationship between work hours and quitting. Panel (b) shows how increasing the elasticity parameter ν would weaken the decrease in stopping probability in response to a 10 percent increase in expected wages. An increase in ψ would also lead the model to predict a weaker relationship between expected wages and quitting. Panel (c) shows how increasing the coefficient of loss aversion L would increase the effect of a \$10 increase in accumulated earnings on the probability of stopping. An increase in ψ or ν predicts the opposite. Panel (d) shows how increasing the speed of adjustment would weaken the difference between earnings accumulated at different times.

E.5 Alternative Specification of the Income Target

This section considers various alternative specifications of the reference point based on a one-period lag of expectations under different definitions of the lag. The specifications in Appendix Table 13 range from a reference point that does not update to a reference point that updates after each trip. The fixed reference point (column 1) represents a model in which the reference point updates sufficiently slowly that it does not vary. For estimation, we take the fixed reference point to equal the driver's average

daily earnings across all shifts, as Camerer et al. (1997) propose. The Crawford and Meng (2011) reference point (column 2) represents drivers' expectations at the end of the previous day. Since these reference points do not update within a shift, neither can account for differential quitting responses based on the within-day timing of earnings.

To allow for within-day updating, we consider a reference point defined as the expectations held at the end of the previous hour or the previous trip in columns (3) and (4). We compute these expectations based on a regression similar to Equation (1) with the outcome defined as shift total income. Our estimate of expected earnings therefore takes the driver's own behavioral responses into account, as in the Kőszegi and Rabin (2006) reference point based on rational expectations.

Appendix Figure 11 shows that the reference points based on lagged expectations of earnings predict a pattern of income effects that do not match the observed recency pattern in the data. Defining the reference point as the one-period lag of expectations produces a stark contrast between the most recently accumulated earnings and any earlier earnings. This occurs because the updated reference point incorporates fully the latter but not at all the former.

For the case of a reference point that represents expectations held at the end of the previous trip, we find a strong effect of earnings accumulated in the most recent hour but little effect of earnings accumulated in previous hours. Likewise, for the case of a reference point that represents expectations held an hour earlier, we find a strong effect of earnings accumulated in the past hour and little effect of earnings accumulated prior to that. Note that the penultimate hour shows an intermediate effect, which comprises an average of a strong effect for recent earnings and a weak effect for earlier earnings. This also occurs if we vary the definition of the lag. For example, if we define lagged expectations as the expectations held by the driver three hours earlier, we find equally strong effects of earnings in the most recent three hours (because the reference point has not adjusted to incorporate any of those additional earnings). We also find an intermediate effect in one of the hours in between and little effect of earnings in the earlier hours.

With the one-period lagged expectation formulation of the reference point, the

model predicts fungibility between all earnings accumulated since the time when the reference point updates. This occurs because the one-period lagged expectation is a discrete approximation to how the reference point adjusts. As the data instead show a gradual relationship between the timing of additional earnings and the probability of ending a shift, our results highlight that earlier lags remain important for explaining the recency pattern. The adaptive reference point, which consists of a weighted average of multiple lagged values of expectations, predicts such a gradual relationship as Figure 3 exhibits.

E.6 Model with Hours Targeting

The Kőszegi and Rabin (2006) model posits that decision makers experience loss aversion relative to each dimension of utility. Thus, their formulation predicts that drivers experience losses from working longer than their “hours target,” analogous to the losses from earning less than their “income target.”

Following Crawford and Meng (2011), the objective function of the driver takes the form

$$v^{LA}(I_t, H_t) = (1 - \eta)v(I_t, H_t) + \eta \sum_{x \in \{I, H\}} n_x(x_t | x_t^r), \quad (3)$$

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), η determines the relative weight on gain-loss utility, and the gain-loss utility is given by

$$n_x(x | x^r) = \left(\mathbf{1}_{\{x > x^r\}} + \lambda_x \mathbf{1}_{\{x < x^r\}} \right) (v_x(x) - v_x(x^r)),$$

where $\lambda_x \geq 1$ parameterizes the degree of loss aversion over each dimension of utility. Based on the discussion in Appendix E.3, we define the coefficient of loss aversion as $L_x = 1 - \eta + \eta\lambda_x$.

Appendix Table 14 column (1) presents results from a specification with a common coefficient of loss aversion for income and hours ($L_I = L_H$). Column (2) relaxes the assumption of a constant coefficient of loss aversion to allow for a different coefficient of loss aversion on each dimension (L_I for income and L_H for hours). The estimates

reveal a significant degree of loss aversion and similar magnitude for both dimensions. The speed of adjustment is similar to (but slightly faster than) the estimate in Table 3 for the model without loss aversion over hours ($L_H = 0$).

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). 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.

E.7 Stochastic Reference Points and Diminishing Sensitivity

The model of loss aversion in Section 3.1 makes two simplifying assumptions: first, utility is piecewise linear in gains and losses, and second, the targets I^r and H^r represent point expectations.

Allowing for diminishing sensitivity corresponds to an objective function that exhibits convexity in losses and concavity in gains. We use the power function

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

where we follow Hastings and Shapiro (2013) by calibrating the parameter ζ to 0.88 (Tversky and Kahneman, 1992). Table 4 column (4) re-estimates the loss-aversion model with diminishing sensitivity. This specification produces a similar estimated coefficient of loss aversion and speed of adjustment of the reference point. Overall we find that relaxing the simplifying assumptions in the loss-aversion model does not change the conclusions about the importance of loss aversion and adaptive reference points.

A stochastic reference point would consist of the distribution of earnings and hours for each shift. We approximate this by defining the stochastic reference point to be a normal distribution with a mean that updates as in Equation (5) and a variance given by the parameter ς . Table 4 column (5) reports an estimated standard deviation of

the income target of about \$38. This specification yields a lower disutility of effort and higher elasticity parameter. We also find a greater degree of loss aversion and correspondingly faster speed of adjustment of the reference point.

F Saliency Model

F.1 Saliency

We adapt a model of saliency 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 saliency of each attribute (income and hours), and (ii) decision makers place greater weight on the more salient attribute.⁵

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 saliency 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 saliency 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

$$v^S(I_t, H_t) = \sum_{x \in \{I, H\}} \frac{w(\sigma(x_t, x_t^n), \delta)}{\sum_{y \in \{I, H\}} w(\sigma(y_t, y_t^n), \delta)} v_x(x_t), \quad (4)$$

where the relative weight $w(\sigma(x, x^n), \delta)$ on the utility for a given attribute increases

⁵Bordalo 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 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 saliency follows Bordalo et al. (2015) which models the evoked set explicitly.

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 Appendix F.2 discusses the importance of these 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

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

⁷This 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.

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. 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. We isolate the channel by which earnings influences stopping behavior through I_t and I_t^r , though additional factors could be included as well.

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, 2012), and relative thinking (Bushong et al., 2016)) which take the choice context as a degree of freedom called the evoked set, consideration 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 2 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.⁸ 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).

⁸To 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.

F.2 Adaptive Reference Points

In this section, we discuss how the salience model can potentially account for the evidence in Section 2. 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.

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. 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.⁹

Stronger income effects in response to recent earnings requires a slow-adjusting reference point. 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

⁹With a positive autocorrelation of hourly earnings as in Appendix Figure 7, 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.

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.¹⁰ 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.

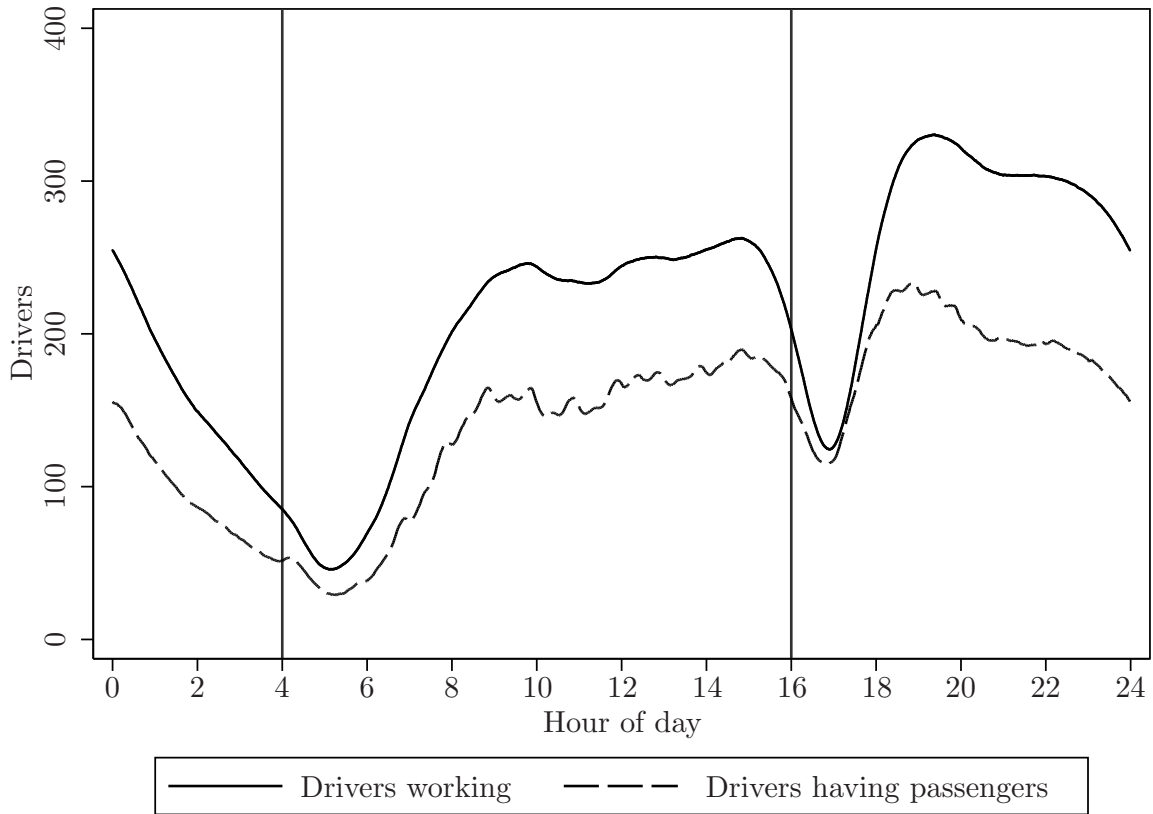
References

- Barseghyan, Levon, Francesca Molinari, Ted O’Donoghue, and Joshua C. Teitelbaum**, “The Nature of Risk Preferences: Evidence from Insurance Choices,” *American Economic Review*, 2013, *103* (6), 2499–2529.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Salience Theory of Choice Under Risk,” *Quarterly Journal of Economics*, 2012, *127* (3), 1243–1285.
- , –, and –, “Salience and Consumer Choice,” *Journal of Political Economy*, 2013, *121* (5), 803–843.
- Browning, Martin, Angus Deaton, and Margaret Irish**, “A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle,” *Econometrica*, 1985, *53* (3), 503–543.
- Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein**, “A Model of Relative Thinking,” *Mimeo*, 2016.
- Hastings, Justine S. and Jesse M. Shapiro**, “Fungibility and Consumer Choice: Evidence from Commodity Price Shocks,” *Quarterly Journal of Economics*, 2013, *128* (4), 1449–1498.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles**, “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 2006, *96* (5), 1589–1610.

¹⁰When 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.


- Kőszegi, Botond and Matthew Rabin**, “Reference-Dependent Risk Attitudes,” *American Economic Review*, 2007, *97* (4), 1047–1073.
- **and Adam Szeidl**, “A model of focusing in economic choice,” *Quarterly Journal of Economics*, 2012, *128* (1), 53–104.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland**, “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 2013, *103* (6), 2530–53.
- Stigler, Stephen M.**, “Do robust estimators work with real data?,” *Annals of Statistics*, 1977, pp. 1055–1098.
- Huber, Martin, Michael Lechner, and Conny Wunsch**, “The performance of estimators based on the propensity score,” *Journal of Econometrics*, 2013, *175* (1), 1–21.
- Thakral, Neil and Linh Tô**, “Daily Labor Supply and Adaptive Reference Points,” *Mimeo*, 2018.

Appendix Figure 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.

Appendix Figure 2: Pattern of wages throughout the day



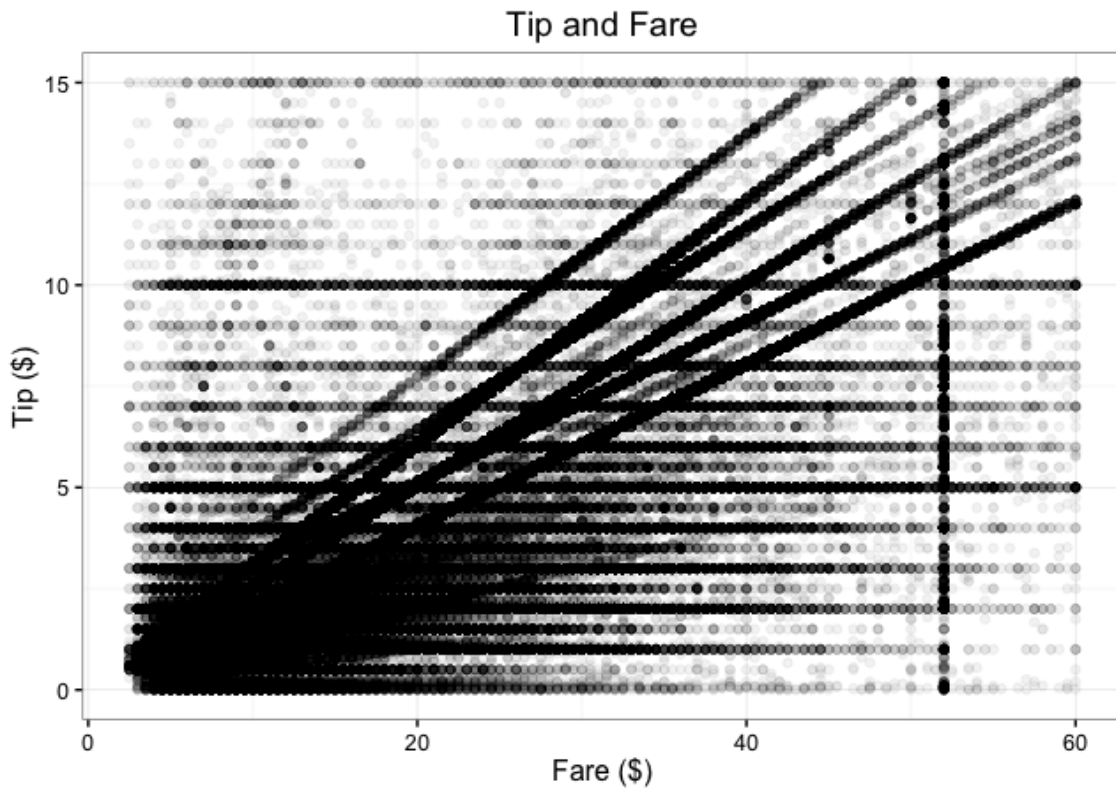
Note: The figure depicts the average market wage every minute throughout the day from hour 0 to hour 23. The market wage in each minute is the average of the per-minute wages of all drivers working during that minute, where a driver's per-minute wage is the ratio of the fare (not including tips) to the 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.

Appendix Figure 3: Pattern of wages: Weekday versus weekend




Note: 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.

Appendix Figure 4: 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.

Appendix Figure 5: Simulated stopping model: Distribution of p values for non-parametric specification



Note: 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.

Appendix Figure 6: Stopping model estimates: Timing pattern of income effect at 8.5 hours—Day versus night




(a) Day shifts (start hour 4AM–10AM)

(b) Night shifts (start hour 2PM–8PM)


Note: The figures depict the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift, during day shifts and night shifts separately. Estimates obtain from Equation (1) with the full set of controls (see Table 1 for details).

Appendix Figure 7: Autocorrelation of residualized hourly market wage



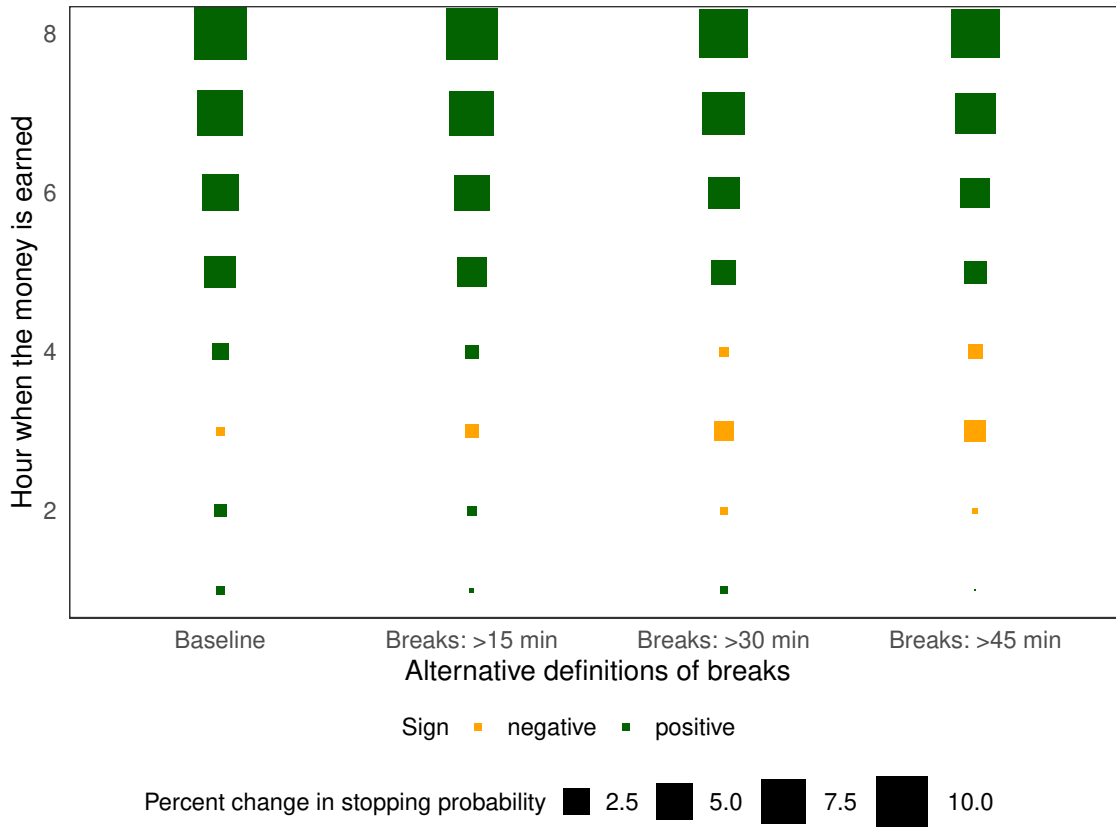
Note: The figure depicts the autocorrelation of hourly market wages indexed by hour of the calendar year 2013. The hourly market wage is the sum of the minute market wage in each hour, with the minute market wage computed as in Appendix Figure 2. The hourly market wage is residualized from a regression 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 Bartlett 95-percent confidence band.

Appendix Figure 8: Autocorrelation of residualized daily market wage



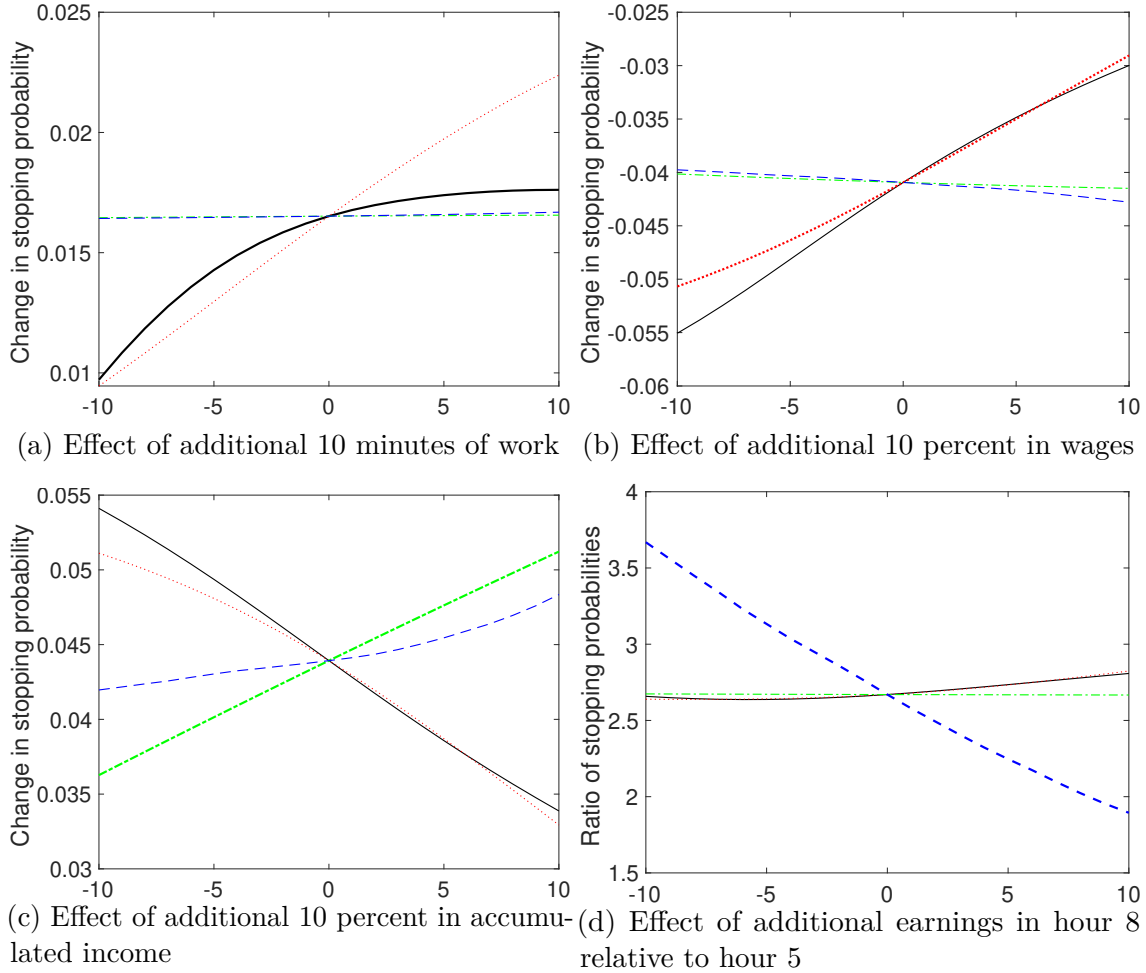
Note: The figure depicts the autocorrelation of daily market wages indexed by day of the year in 2013. The daily market wage is the sum of the minute market wage in each calendar day, with the minute market wage computed as in Appendix Figure 2. The daily market wage is residualized from a regression 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 Bartlett 95-percent confidence band.

Appendix Figure 9: Stopping model estimates: Income effect at 8.5 hours—Timing pattern with alternative definitions of breaks




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 from Figure 3, 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.

Appendix Figure 10: Relationship between Moments and Parameters



Note: Each panel plots the relationship between a moment of the data and the four key structural parameters: disutility of effort ψ (solid black line), elasticity parameter ν (dotted red line), coefficient of loss aversion L (dash-dot green line), and speed of adjustment θ (dashed blue line). Each parameter varies within one standard deviation of its estimated value. Panel (a) plots the relationship between the parameters and the change in the probability of stopping in response to working for an additional 10 minutes. Panel (b) plots the relationship between the parameters and the change in the probability of stopping in response to a 10 percent increase in the wage from continuing. Panel (c) plots the relationship between the parameters and the change in the probability of stopping in response to a \$10 increase in accumulated earnings. Panel (d) plots the relationship between the parameters and the effect of a \$10 increase in earnings accumulated in the most recent hour relative to three hours earlier on the probability of stopping.

Appendix Figure 11: Stopping model estimates: Income effect at 8.5 hours—Data and lagged-expectation reference point model



Note: The figure compares the income effects estimated using Equation (1) with the predicted income effects from specification of reference points based on the lagged expectation estimated in columns (3) and (4) of Appendix Table 13. The confidence interval displays the estimates from Figure 3 of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. The gray squares and black diamonds represent the predictions of the expectations-based reference point model with a lag of an hour and a trip, respectively.

Appendix Table 1: Trip-level and shift-level summary statistics

	Mean	25p	50p	75p
Ride duration in minutes	12.6	6.1	10	16
Wait duration in minutes	11.3	2	5	12
Fare in dollars	12.2	6.5	9.5	14
Percent tip	19.3	15.4	20	21.7
	Mean	25p	50p	75p
Number of shifts per driver	163	100	170	276
Number of trips in a shift	21.7	16	22	27
Shift total hours	8.6	7.2	8.7	10.1
Shift ride hours	4.5	3.5	4.5	5.5
Fraction break time	.16	0	.13	.25
Shift income in dollars	270	214	267	324
Shift income (day shifts)	240	195	245	289
Shift income (night shifts)	274	219	273	330
Hourly wage in dollars	31.6	27.7	31.7	35.6
Hourly wage (day shifts)	29.8	26.5	29.9	33.2
Hourly wage (night shifts)	33.8	30.0	34.1	37.9

Note: The top panel reports summary statistics at the trip level for all 127 million NYC taxi trips in 2013 in the cleaned data. The bottom panel reports summary statistics at the shift level. 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 55% of trips with credit card as the payment type. Shift total hours is the number of hours between the pickup time of the first passenger of the shift and the dropoff time of the last passenger of the shift. Shift ride hours is the sum of the duration of each fare during the shift. Breaks are defined as in Appendix Figure 9. Income and hourly wage represent gross fares, not including tips, gasoline costs, or leasing fees.

Appendix Table 2: Wage elasticity estimates: Weekday versus weekend

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

Note: 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.

Appendix Table 3: 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 (0.0030)	0.1979	-0.0058 (0.0043)	0.9368
Linear probability model				
F-1	-0.0083 (0.0005)	0.0000	-0.0269 (0.0006)	0.0000
F-1*	0.0005 (0.0004)	0.2329	-0.0116 (0.0006)	0.0000
F-2a		0.0017		0.0000
F-2b	0.0436 (0.0056)	0.0000	-0.0202 (0.0110)	0.0000
F-3	-0.0054 (0.0052)	0.0000	0.0112 (0.0186)	0.0000
Probit model				
F-1	-0.0081 (0.0005)	0.0000	-0.0317 (0.0009)	0.0000
F-1*	0.0004 (0.0003)	0.2467	-0.0207 (0.0013)	0.0000
F-2a		0.0303		0.0000
F-2b	0.0185 (0.0029)	0.0000	-0.0189 (0.0114)	0.0000

Note: Each row corresponds to a different regression equation defined in Appendix C.2. The row F-2a corresponds to Equation (F-2) with the partitions \hat{H} and \hat{Y} over hours and income defined as in Farber (2005), and the row F-2b corresponds to Equation (F-2) 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-3) estimated as a linear probability model. The bottom panel reports results from the probit models in Equations (F-1) to (F-3). 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-2a 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-1*), 9 restrictions for Equation (F-2), 72 restrictions for Equation (F-3), and 59 restrictions for Equation (TT) in Simulation 1, and 78 restrictions for Equation (TT) in Simulation 2.

Appendix Table 4: 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)
<i>p</i> -value: income coefs. = 0	0.0000
<i>p</i> -value: Equality of income coefs.	0.8464

Note: 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*).

Appendix Table 5: Elasticity of stopping at 8.5 hours with respect to income: Day and night shifts

	(TT)	(F-1)	(F-2)	(F-3)
Day shifts (start hour 4AM–10AM)				
Controlling for				
Hours	0.1921 (0.0403)	0.1680 (0.0205)	0.9887 (0.0885)	0.6804 (0.1337)
& Drivers	0.6913 (0.0255)	0.3325 (0.0169)	1.1779 (0.0860)	0.8883 (0.1302)
& Time	0.8183 (0.0244)	0.4891 (0.0182)	0.9063 (0.0828)	0.6454 (0.1252)
& Location	0.3227 (0.0232)	-0.0877 (0.0178)	0.7969 (0.0815)	0.4974 (0.1229)
& Weather	0.3225 (0.0232)	-0.0881 (0.0178)	0.7974 (0.0815)	0.4975 (0.1229)
	(TT)	(F-1)	(F-2)	(F-3)
Night shifts (start hour 2PM–8PM)				
Controlling for				
Hours	0.7556 (0.0325)	1.2211 (0.0181)	-0.9291 (0.0612)	-1.1880 (0.1471)
& Drivers	0.5497 (0.0274)	1.4170 (0.0160)	-0.9085 (0.0604)	-1.1831 (0.1456)
& Time	0.8027 (0.0271)	0.6700 (0.0190)	0.1434 (0.0588)	-0.0049 (0.1395)
& Location	0.2704 (0.0263)	0.3110 (0.0189)	0.0090 (0.0579)	-0.1361 (0.1376)
& Weather	0.2702 (0.0263)	0.3153 (0.0190)	0.0088 (0.0579)	-0.1360 (0.1376)

Panel B: Comparison with previous estimates

This paper (day):	0.3225	(0.2770, 0.3679)
This paper (night):	0.2702	(0.2186, 0.3219)
Farber 2015 (day):	0.9456	
Farber 2015 (night):	0.0727	

Note: This table reports in each cell an estimate of the percent change in the probability of ending a shift at 8.5 hours in response to a 1 percent increase in cumulative earnings. The columns corresponds to the specifications in Equations (TT) to (F-3), respectively. The control variables consist of the full set from Table 1. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 6: Stopping model estimates: By variance of daily hours

	(1) < 10th percentile	(2) < 25th percentile	(3) > 75th percentile	(4) > 90th percentile
Panel A				
Cumulative income	0.2479 (0.0641)	0.2866 (0.0398)	0.3854 (0.0403)	0.5236 (0.0673)
Panel B				
Income in hour 2	0.0029 (0.1716)	-0.0251 (0.1137)	0.1082 (0.1411)	0.5117 (0.2441)
Income in hour 4	0.0268 (0.1933)	-0.0020 (0.1197)	0.0679 (0.1361)	-0.0784 (0.2322)
Income in hour 6	0.5402 (0.2062)	0.6222 (0.1246)	0.6296 (0.1322)	0.6042 (0.2258)
Income in hour 8	0.7941 (0.2212)	0.8954 (0.1350)	1.2187 (0.1430)	1.3157 (0.2482)

Note: Panel A reports estimates from Equation (TT) of the percent change in the probability of ending a shift at 8.5 hours in response to a 1 percent increase in cumulative earnings. Panel B reports estimates from Equation (1) of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. The columns restrict drivers based on the variance of their distribution of hours worked across days: (1) below the 10th percentile, (2) below the 25th percentile, (3) above the 75th percentile, and (4) above the 90th percentile. The control variables consist of the full set from Table 1. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 7: Stopping model estimates: Income effect at 8.5 hours—Robustness

	(1) Control individual- specific location effects	(2) Control previous-day income	(3) Medallion owners
Panel A			
Cumulative income	0.3551 (0.0246)	0.3070 (0.0185)	0.7440 (0.1896)
Previous-day income		-0.0284 (0.0151)	
Driver \times NTA	X		
Panel B			
Income in hour 2	-0.0261 (0.0809)	0.0629 (0.0601)	-0.1154 (0.5955)
Income in hour 4	0.0680 (0.0799)	0.0962 (0.0591)	0.1347 (0.5965)
Income in hour 6	0.4409 (0.0803)	0.4919 (0.0589)	1.2930 (0.5717)
Income in hour 8	1.1098 (0.0862)	1.0213 (0.0631)	1.0159 (0.5803)
Previous-day income		-0.0281 (0.0151)	
Driver \times NTA	X		

Note: The table uses the same sample and control variables as in Table 1. Panel A reports estimates from Equation (TT) of the percent change in the probability of ending a shift at 8.5 hours in response to a 1 percent increase in cumulative earnings. Panel B reports estimates from Equation (1) of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. Column (1) adds location-by-neighborhood fixed effects to the last row of Table 1 column (TT). Column (2) adds a control for previous-day income to the last row of Table 1 column (TT). Column (3) corresponds to a restricted sample of all 255 drivers who operate exactly one cab and no other driver shares that cab. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 8: Stopping model estimates: Timing pattern of income effect—
Alternative specifications for the stopping model

	(F-1)	(F-2)	(F-3)
Income in hour 2	0.1471 (0.0172)	0.7504 (0.0102)	0.7515 (0.0132)
Income in hour 4	-0.1765 (0.0190)	0.3216 (0.0166)	0.5027 (0.0220)
Income in hour 6	-0.0122 (0.0232)	0.7246 (0.0220)	0.7814 (0.0292)
Income in hour 8	1.0914 (0.0344)	1.2029 (0.0398)	1.2289 (0.0522)

Note: This table reports estimates of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. Columns (1)–(3) extend Equations (F-1) to (F-3), respectively, to add an effect of earnings in each hour $\sum_{\ell} \beta^{\ell} y_{int}^{\ell}$. The control variables consist of the full set from Table 1. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 9: Stopping model estimates: Income effect at 8.5 hours—Timing pattern for Manhattan trips

	(1) OLS	(2) IV
Income in hour 2	0.0044 (0.1082)	-0.2814 (0.8109)
Income in hour 4	0.1147 (0.1062)	1.0616 (0.9517)
Income in hour 6	0.5781 (0.1072)	-0.2716 (0.9909)
Income in hour 8	0.9935 (0.1124)	2.3673 (1.1844)

Note: This table reports estimates of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. As in Figure 3, estimates obtain from Equation (1) with the full set of controls. Both columns restrict the sample to trips that remain in Manhattan. The first column presents least-squares estimates, while the second column presents IV estimates, instrumenting for earnings with speed.

Appendix Table 10: Stopping model estimates: Income effect at 8.5 hours—Across-driver experience

	(1)	(2)
	Less experienced drivers	More experienced drivers
Panel A		
Cumulative income	0.0695 (0.0669)	0.3183 (0.0164)
Panel B		
Income in hour 2	-0.1373 (0.2112)	0.0465 (0.0516)
Income in hour 4	-0.1461 (0.2200)	-0.0124 (0.0511)
Income in hour 6	0.4125 (0.2209)	0.5514 (0.0506)
Income in hour 8	0.5118 (0.2330)	0.9145 (0.0547)

Note: Panel A reports estimates from Equation (TT) of the percent change in the probability of ending a shift at 8.5 hours in response to a 1 percent increase in cumulative earnings. Panel B reports estimates from Equation (1) of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. The sample in column (1) restricts to the 2,361 drivers who first appear in the data in April 2013 and drive at least 50 shifts until December 2013. The sample in column (2) consists of the 32,547 drivers who first appear in the data before April 2013. The control variables consist of the full set from Table 1. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 11: Stopping model estimates: Income effect at 8.5 hours—Within-driver experience

	(1)	(2)	(3)
Panel A			
Cumulative income	-0.0118 (0.0880)	0.0306 (0.0781)	0.0766 (0.0680)
× experience	0.0253 (0.0186)	0.0007 (0.0007)	-0.0135 (0.0231)
Panel B			
Income in hour 2	0.1052 (0.3256)	0.0233 (0.3164)	-0.0420 (0.2616)
× experience	-0.0745 (0.0744)	-0.0026 (0.0039)	-0.1892 (0.3014)
Income in hour 4	-0.6026 (0.3287)	-0.4356 (0.3144)	-0.4509 (0.2693)
× experience	0.1402 (0.0770)	0.0047 (0.0040)	0.6249 (0.3227)
Income in hour 6	0.3580 (0.3201)	0.4002 (0.3094)	0.3710 (0.2605)
× experience	0.0146 (0.0741)	0.0001 (0.0037)	0.0752 (0.2984)
Income in hour 8	0.6871 (0.3377)	0.4480 (0.3267)	0.6603 (0.2768)
× experience	-0.0526 (0.0752)	0.0011 (0.0039)	-0.2964 (0.3006)

Note: The estimates in this table restrict to the sample of 2,361 drivers who first appear in the data in April 2013 and drive at least 50 shifts until December 2013. Panel A reports estimates from Equation (TT) of the percent change in the probability of ending a shift at 8.5 hours in response to a 1 percent increase in cumulative earnings. Panel B reports estimates from Equation (1) of the percent change in the probability of ending a shift at 8.5 hours in response to a \$26 increase in earnings accumulated at different times in the shift. Column (1) defines experience as the number of months since the driver started. Column (2) defines experience as the number of shifts since the driver started. Column (3) defines experience as an indicator variable which is equal to one for half of each driver's shifts, those that occur later in the year. The control variables consist of the full set from Table 1. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 12: Stopping model estimates: Income effect at 8.5 hours—Reduced-form model allowing jumps at the targets

	(1)	(2)	(3)	(4)
Above hours target	0.0393 (0.0019)	0.0391 (0.0019)	0.0392 (0.0019)	0.0391 (0.0019)
Above I_0	0.0059 (0.0019)	0.0028 (0.0023)	0.0025 (0.0027)	0.0024 (0.0041)
Above $I^{early>late}$	0.0127 (0.0024)	0.0176 (0.0033)	0.0219 (0.0042)	0.0298 (0.0069)
Above $I^{early<late}$	0.0031 (0.0023)	0.0078 (0.0032)	0.0065 (0.0041)	0.0064 (0.0069)
Cumulative income / 100	0.0065 (0.0027)	0.0060 (0.0028)	0.0061 (0.0028)	0.0061 (0.0028)

Note: This table reports the change in the probability of stopping at 8.5 hours in response to reaching targets for hours and income. The first row corresponds to expected daily hours as the hours target. The second row corresponds to expected daily earnings as the income target. The third row corresponds to the increment in the stopping probability from exceeding the updated targets for income that put more weight on income earned in hours 0–4 of the shift. The fourth row corresponds to the increment in the stopping probability from exceeding the updated targets for income that put more weight on income earned in hours 5–8 of the shift. Column (1) defines updated targets with weights 0, 50, or 100 on earnings in each half of the shift. Column (2) uses weights 0, 33, 66, and 100. Column (3) uses weights 0, 25, 50, 75, and 100. Column (4) uses weights 0, 10, 25, 50, 75, 90, and 100. Standard errors reported in parentheses are adjusted for clustering at the driver level.

Appendix Table 13: Maximum likelihood estimates: Alternative specifications of the reference point

	One-period lagged expectation			
	Fixed Ref. Pt. (1)	Day (2)	Hour (3)	Trip (4)
Disutility of effort ψ	0.2722 (0.1500)	0.1932 (0.1480)	0.2524 (0.1465)	0.2663 (0.1401)
Elasticity ν	0.8829 (0.2576)	0.7902 (0.2582)	0.9415 (0.2633)	0.8878 (0.2563)
Loss aversion L	2.7228 (0.3528)	1.9992 (0.2364)	2.8339 (0.4064)	2.6650 (0.4043)
Error term distribution σ	0.6659 (0.0856)	0.4077 (0.0914)	0.6982 (0.0987)	0.6616 (0.1077)
Mean log likelihood	-27,298	-27,056	-27,277	-27,312
Test $L = 1$	<0.001	<0.001	<0.001	<0.001

Note: This table presents maximum likelihood estimates of Equation (6) with the reference point given by the lagged expectation of earnings under different definitions of the lag. The estimation sample consists of over 37,000 drivers, using data from 100 subsamples of 150,000 trips each without replacement. Column (1) corresponds to a fixed reference point (sufficiently long lag that reference points do not adjust) equal to the driver's average earnings across all shifts. Column (2) corresponds to the case that the updated reference point is formed the previous day, as in Crawford and Meng (2011), which also appears in column (2) of Table 3. Column (3) and (4) correspond to updated reference points that are formed the previous hour and trip, respectively, estimated using Equation (1) with total fare earnings for the shift as the outcome variable. The mean log likelihood reports the average of the log likelihood across the subsamples. The last row contains p -values from likelihood ratio tests of the null hypotheses of the model without loss aversion.

Appendix Table 14: Maximum likelihood estimates: Loss aversion over income and hours

	(1)	(2)
Disutility of effort ψ	0.2980 (0.1000)	0.3793 (0.1045)
Elasticity ν	0.1544 (0.1545)	0.1700 (0.0888)
Loss aversion over income L_I	1.9180 (0.1488)	2.0484 (0.2969)
Loss aversion over hours L_H	1.9180 (0.1488)	1.8255 (0.1535)
Speed of adjustment θ	0.7417 (0.1226)	0.7849 (0.0679)
Error term distribution σ	0.2874 (0.0526)	0.3118 (0.0896)
Mean log likelihood	-26,835	-26,652
Test $L = 1$	<0.001	<0.001
Test $\theta = 1$	<0.001	<0.001

Note: This table presents maximum likelihood estimates of Equation (2) for the objective function Equation (3), which allows for loss aversion on both dimensions (income and hours), with the adaptive reference point given by Equation (5). The estimation sample consists of over 37,000 drivers, using data from 100 subsamples of 150,000 trips each without replacement. The mean log likelihood reports the average of the log likelihood across the subsamples. The last two rows contain p -values from likelihood ratio tests of the following null hypotheses: (i) the model without loss aversion, and (ii) the model with a static reference point.