

Estimation of a Latent Reference Point: Method and Application to NYC Taxi Drivers

Marco Stenborg Petterson*

November 2021

Job Market Paper

The paper is updated often. Check here for the latest version.

Abstract

I use a dynamic discrete choice model with a latent variable to flexibly estimate reference-dependent utility models. The structure and evolution of the reference point are estimated directly from observational data. I apply the model to the daily labor-supply choices of NYC taxi drivers and use a Bayesian estimation approach. I find that rational expectations are an important determinant of the reference point but do not fully explain its evolution. The reference point adjusts asymmetrically, responding more to positive income shocks than to negative ones. The reference point also has an important transitory component: a shock to the reference point dissipates within hours. I use the estimated model to analyze the welfare effects of a guaranteed hourly wage floor for gig economy workers.

*Brown University, Department of Economics. Email: marco_stenborg_petterson@brown.edu I am grateful to Jesse Shapiro, Andriy Norets, and Neil Thakral for their guidance and advice. This project benefited from discussions with Zeky Murra-Anton, Tommaso Coen, Toru Kitagawa, Giacomo Rubbini, Susanne Schennach, Vincent Starck, Valeria Zurla, and seminar participants at Brown University. Financial support from Orlando Bravo Center for Economic Research and Morgan Stanley endowment at Brown University is gratefully acknowledged. The findings and conclusions expressed are solely those of the author and do not represent the views of funding sources.

1 Introduction

Reference-dependent preferences are widely used in economics. They rationalize many behaviors observed in the data by introducing a reference point in the individuals' utility around which marginal utility changes. Reference-dependent preferences have been used to explain consumer choice (Kőszegi and Rabin, 2006), the housing market (Genesove and Mayer, 2001), the stock market (Barberis, Huang, and Thaler, 2006), and electoral competition (Panunzi, Pavoni, and Tabellini, 2020). Camerer et al. (1997) propose a model with income targeting to explain the labor supply choices of taxi drivers in New York City. The idea is that drivers quit working upon reaching a target level of earnings, their reference point. In a different application, DellaVigna et al. (2017) propose a job search model with reference-dependent preferences, where the recent income level of unemployed individuals is the reference point. Search effort of newly unemployed individuals increases in anticipation of a cut in the level of unemployment benefits and, conditional on being unemployed, the effort in job search decreases again after a few periods. This behavior is consistent with reference-dependent preferences with a dynamic reference point.

Despite its importance for understanding economic behavior, much remains unknown about the nature of the reference point itself. As the reference point is a latent state, it is not directly observed. Empirical research with field data often posits a particular form for the reference point, such as a backward-looking average of past income (e.g., DellaVigna et al. 2017), or a forward-looking expectation of future income (e.g., Kőszegi and Rabin 2006, Thakral and Tô 2021). More direct evidence on the nature and evolution of the reference point comes from laboratory studies, many of which ask participants to make a sequence of choices between binary lotteries. Questions studied in the recent literature include the role of forward-looking expectations vs. backward-looking experiences (Heffetz and List 2014, Gneezy et al. 2017), whether the reference point adjusts asymmetrically to positive vs. negative shocks (Arkes et al. 2008, Arkes et al. 2010), and whether changes to the reference point are persistent or transitory (Thakral and Tô 2021).

In this paper, I estimate the evolution of reference points using a dynamic discrete choice model with the reference point as a latent variable using only observational data. I apply the model to the labor supply decision of NYC taxi drivers. Formulating the reference point as a state variable in a dynamic choice model allows one to capture the possibility of it being dynamic and study its evolution through the transition matrix of states, which lies at the core of dynamic choice models. I provide a flexible specification of the transition matrix for the unobservable state that allows me to

study the formation process of the reference point. I also allow the evolution of the reference point to differ across individuals, thus permitting the study of individual heterogeneity in the formation and evolution of the reference point.

I ground my analysis in recent developments in the identification and estimation of dynamic discrete choice models with persistent unobserved heterogeneity. Hu and Shum (2012) and Connault (2016) provide non-parametric identification results of models related to the structural model I present in this paper. Connault (2016) presents a generic identification result for a set of models, called hidden Rust models, where observed state variables and unobserved state variables follow a joint Markov process. The model I use in this paper is a special case of hidden Rust models.

I apply the approach to the daily labor supply decisions of NYC taxi drivers. Following the literature¹, I use a dynamic optimal stopping model applied to an agent with reference-dependent preferences. The reference point enters the utility affecting both the level of utility and the marginal utility of income before and after reaching the reference point. My main analysis uses the specification of the utility function from Farber (2015), but I also consider one from Farber (2008) in an extension. The daily labor supply model assumes that each day is identical to the previous one.² The observable state variables are the cumulative hours worked and the cumulative income earned within each day. The reference point is introduced as a latent state variable, observable by the agent but not the econometrician. I flexibly parametrize the transition matrix of the unobservable state variable to draw conclusions on its formation process.

Persistent unobserved heterogeneity implies a correlation over time of the unobserved component. Therefore, each evaluation of the likelihood function involves computing multidimensional integrals to integrate out the contribution of the unobserved component. The dimension of the integral grows with the number of time periods and with the number of possible values of the unobserved state. I make computation feasible by adopting particle filtering (e.g., Blevins 2016) and related methods. I adopt a Bayesian framework to estimate the model, which, as shown by Andrieu, Doucet, and Holenstein (2010) and Flury and Shephard (2011), allows me to do likelihood-based inference using simulations.

The data for the application is trip-level data provided by the NYC Taxi and Limousine. I have access to all the anonymized trip sheets of 2013 for all NYC taxi drivers. Although anonymized, I can follow the drivers throughout the year and study their daily labor supply decisions. The

¹See Buchholz, Shum, and Xu (2018), Thakral and Tô (2021)

²For a discussion of this assumption see Camerer et al. (1997) and Thakral and Tô (2021).

variables I use are the fare earned at the end of each trip by the driver and the start and end times of each trip. This data allows me to properly define the shifts and construct the relevant observable state variables within each shift.

The estimated transition matrix allows me to tackle the open questions in the literature mentioned above, distinguish between behavioral models, and gain insight into the formation and evolution of reference points. To do this, I use a method similar to impulse response functions by first simulating the average path of observable state variables and of the reference point given the estimated structural parameters. I then introduce a shock in the observable state variables and study the effect on the reference point. I find that the reference point depends on its past values and the current values of the observable components. Shocks to the reference point are fairly transitory, in line with the findings of Thakral and Tô (2021). Furthermore, the reference point adjusts asymmetrically to positive and negative shocks to the observable state variables, reacting much more prominently to positive shocks than negative ones. The asymmetry in responses to shocks is in line with the findings of Arkes et al. (2008) and Arkes et al. (2010) from surveys and incentivized choices in the laboratory. I also find that rational expectations drive a significant portion, but not all, of the variation of reference points in this setting. This finding corroborates the experimental literature on expectation-based reference-dependent preferences that test the theories developed following the seminal paper of Kőszegi and Rabin (2006) while highlighting that rational expectations do not fully explain the observed behavior. In the last part of the paper, I discuss some implications of my results for welfare analysis by considering the introduction of a minimum wage for taxi drivers, such as the one considered for gig economy workers in many countries.³

This paper contributes to the literature on reference-dependent preference by providing tools to tackle some of the questions raised by O'Donoghue and Sprenger (2018). I contribute a methodology to derive the reference point directly from observational data that allows for a flexible specification of the reference point and its formation process and a portable estimation framework. Being able to derive the drivers and the evolution path of the reference point directly from observational data is a novel approach with respect to the experimental strategy taken by most of the literature. Baillon, Bleichrodt, and Spinu (2020), for example, recover which reference point agents use in decisions under risk by using an experimental design with subjects choosing between lotteries and fixing six different possible rules the reference point must follow. Baucells, Weber, and Welfens (2011) use experiments in financial context to understand the formation of reference points over

³See, for example, this LA Times article

time. They then parametrically characterize the reference point as a weighted average and show that reference points are not recursive. In contrast to these studies, which test prespecified theories of reference point formation using experimentally generated data, I estimate a flexible model of the reference point as a latent variable using observational data. My approach permits testing a wide range of theories about the reference point based on a single set of estimates and could be applied in observational settings other than the one I study.

This work takes a different stance with respect to Buchholz, Shum, and Xu (2018) to the problem of estimating the daily labor supply of NYC taxi drivers. I explicitly model the reference point as an unobserved state variable and draw conclusions from its transition matrix. Buchholz, Shum, and Xu (2018) consider a utility function that depends on the income earned and hours worked but do not model the unobserved component. They instead estimate the model semi-parametrically and conclude that both neoclassical and behavioral wage responses are present in the data. I estimate the model at an individual level I also find that some individuals are heavily influenced by the reference point while others are not. Gentry and Pesendorfer (2018) consider a dynamic framework with price reference with an unobserved binary state variable, the interpretation of the unobserved state variable is whether the consumer is attentive, i.e., actively shopping for the product category or not. Though related, these analyses do not coincide with the model considered in this paper as the reference point is not explicitly modeled and analyzed.

The paper also contributes to the growing literature on gig economy workers, such as Jackson (2019) and Chen et al. (2019), by exploring a key determinant of their labor supply decisions and providing insights on how to account for possible reference-dependent preferences I provide a welfare analysis of introducing a minimum hourly wage and provide considerations along the lines of the recent work on welfare analysis with reference-dependent preferences by Reck and Seibold (2021).

The paper is structured as follows. Section 2 introduces the dynamic discrete choice model with a latent state. Section 3 provides information on the institutional context and describes the data. Section 4 provides the details underlying the estimation. Section 5 presents the results and implications of the estimated model. Section 6 concludes.

2 Structural model

Consider an infinite horizon dynamic discrete choice model. In its most general formulation, we can set up the agent's problem in the following way

$$V(s_t, x_t, \varepsilon_t) = \max_{a_t} \mathbb{E} \left[\sum_{j=0}^{\infty} \beta^j u(s_{j+t}, x_{j+t}, \varepsilon_{j+t}, a_{j+t}) \mid s_t, x_t, \varepsilon_t, a_t \right]$$

where V is the agent's lifetime utility from time t onwards, a is the control variables, i.e. the actions taken by the agent, while s and x are the state variables of the problem. The state variable s is observable both by the agent and the econometrician, while x is a state variable that is observed by the agent but not the econometrician. Finally, ε is an unobservables that is observed by the agent and not by the econometrician but its realizations are independent over time. In this paper I assume that a is a binary variable taking values $a_t \in \{0, 1\}$. Both s_t and x_t are discrete random variables taking values in \mathcal{S} and \mathcal{X} respectively.

Similarly to Rust's (1987) setup, the utility function of each agent i at each time t is

$$u(s_t, x_t, \varepsilon_t, a_t) = \begin{cases} u_1(s_t, x_t) + \varepsilon_{1t} & a_t = 1 \\ \varepsilon_{0t} & a_t = 0 \end{cases}$$

where ε_0 and ε_1 enter the utility function additively, are *iid* over time and between choices, and are distributed with type I extreme value distributions.

The state variables s and x follow a joint Markov process. Formally, $\Pr(s_t, x_t \mid s_{1:t-1}, x_{1:t-1}, a_{1:t-1}, \varepsilon_{1:t-1}) = \Pr(s_t, x_t \mid s_{t-1}, x_{t-1}, a_{t-1})$ ⁴.

With some very mild assumptions (see Bhattacharya and Majumdar, 1989; Rust, 1994) we can rewrite the problem in a recursive formulation

$$v_a(s, x, \varepsilon) = u(s, x, \varepsilon, a) + \beta E [V(s', x', \varepsilon') \mid s, x, \varepsilon, a]$$

$$V(s, x, \varepsilon) = \max_a \{v_a\}$$

The assumptions made so far on the support of the state variables and the utility function allow

⁴For any variable y I denote with $y_{1:t}$ all the realizations of y from period 1 to period t , equivalently $y_{1:t} = \{y_j\}_{j=1}^t$.

me to rewrite the problem even more conveniently using the following Bellman equation

$$v_a(s, x, \varepsilon) = u(s, x, \varepsilon, a) + \beta \sum_{s', x'} \left[\log \sum_{a'} e^{v_{a'}(s', x')} \right] q(s', x' | s, x, a)$$

Notice that if both s and x were observable, the model would mirror the assumptions and setup of Rust (1987).

I also impose some additional structure on the joint Markov process of the observed and unobserved state variables. I assume limited feedback between the observed state variables s and the unobserved state variables x . Formally

$$\Pr(s_t, x_t | s_{t-1}, x_{t-1}, a_{t-1}) = \Pr(x_t | s_t, x_{t-1}, a_{t-1}) \Pr(s_t | s_{t-1}, a_{t-1})$$

which implies that the unobserved state variable can only have an effect on the observed state variable through the choice variable in the previous period. In contrast, the observable state variable can have a contemporaneous effect on the unobserved state variable. Figure 1 the dependence structure of this model graphically.

If $a_{t-1} = 1$ then the values of both the observed and unobserved state variables are deterministically determined in period t . Both state variables return to their initial value which formally means $\Pr(s_t = 1 | a_{t-1} = 1, s_{t-1} = k) = \Pr(s_t = 1 | a_{t-1} = 1) = 1$ and $\Pr(x_t = 1 | a_{t-1} = 1, x_{t-1} = q, s_t = j) = \Pr(x_t = 1 | a_{t-1} = 1) = 1$. If instead $a_{t-1} = 0$ then we can denote the transition probability from observed states k to observed state j with θ_{jk}^S and the transition probability between unobserved states q and the unobserved state m given the observed state j with θ_{mqj}^X . Formally

$$\Pr(s_t = j | s_{t-1} = k, a_{t-1} = 0) = \theta_{jk}^S \tag{1}$$

$$\Pr(x_t = m | a_{t-1} = 0, x_{t-1} = q, s_t = j) = \theta_{mqj}^X \tag{2}$$

To simplify the estimation and reduce the computational burden, I impose some parametric assumptions on θ_{mqj}^X while maintaining a large degree of flexibility. In particular

Therefore, we can introduce some parametric simplifications that ease the identification require-

ments and reduce the computational requirements while maintaining a large degree of flexibility. A parametric specification can sometimes also provide some economic intuition of the underlying mechanism on the determinants of the transition probabilities between unobservable states. I model the evolution of the unobserved state variable as a multinomial choice made by Nature, where the choice probabilities of x_t depend on x_{t-1} and s_t .

$$\begin{aligned}
& \Pr(x_t, s_t | x_{t-1}, s_{t-1}, a_{t-1} = 0) \\
&= \Pr(x_t | s_t, x_{t-1}, s_{t-1}, a_{t-1} = 0) \Pr(s_t | x_{t-1}, s_{t-1}, a_{t-1} = 0) \\
&= \Pr(x_t | s_t, x_{t-1}, a_{t-1} = 0) \Pr(s_t | s_{t-1}, a_{t-1} = 0) \\
&= \frac{\exp(h(x_t, x_{t-1}, s_t))}{\sum_{x'} \exp(h(x', x_{t-1}, s_t))} \Pr(s_t | s_{t-1}, a_{t-1} = 0) \tag{3}
\end{aligned}$$

Notice that θ_{jk}^S is left to be non-parametrically identified directly from the data. The exact specification of $h(x_t, x_{t-1}, s_t)$ is presented in Section 2.1.

2.1 A model of reference-dependent daily labor supply

I now adapt the general model specified above to reference-dependent preferences. I also introduce some details and assumptions that are more pertinent to the application of this paper, where I estimate the evolution of a reference point underlying the daily labor supply decisions of NYC taxi drivers. The assumptions and details specific to the application considered in this paper can be easily adapted to other contexts of interest to the reader.

Let us start by defining what a day is, which is more precisely called a shift in the case of NYC taxi drivers. The agents work in shifts; each shift is composed of several trips. Each observation corresponds to a trip. A shift is then defined by consecutive periods in which $a_t = 0$ and the last trip of the shift has $a_t = 1$, we can think of $a_t = 1$ as the decision to end the shift. I define the n^{th} shift as the set of observations $\tau(n) = \{t : \sum_{j=1}^t a_j = n\}$. Consider then a generic variable y , denote by y_t the value of y at time t and denote with a capital letter Y_t the cumulative value of y up to time t from the beginning shift τ of which t is part of.⁵ Unless otherwise noted, I will assume that $Y_0 = 0$ for any variable y .

I follow Farber (2015) in defining the utility function. The observable states are composed of three main components: the cumulative income earned, the cumulative hours worked, and other

⁵Formally $Y_t = \sum_{j=0}^t y_j - \sum_{j=t'}^t y_j = \sum_{j=t'+1}^t y_j$ with $t' = \max\{l | a_l = 1, l < t\}$, where $a_t = 1$ signals the end of a shift.

observables the econometrician may want to include in the model, such as weather or indicator variables for rush hour traffic z_t . The cumulative income earned and hours worked, denoted respectively, W_t and H_t , are the income earned and the hours worked up to, and including, trip t from the start of shift τ of which t is part of. The unobservable state is $x_t = \xi_t^r$ and may take different interpretations depending on the assumed model.

The utility from ending a shift at period t is then defined as follows

$$\begin{aligned} u_1(W_t, H_t, z_t, \xi_t^r) &= (1 + \alpha)(W_t - g(W_t, H_t, z_t, \xi_t^r)) - \frac{\psi}{1+\eta} (H_t)^{1+\eta} \text{ if } W_t < g(W_t, H_t, z_t, \xi_t^r) \\ u_1(W_t, H_t, z_t, \xi_t^r) &= (1 - \alpha)(W_t - g(W_t, H_t, z_t, \xi_t^r)) - \frac{\psi}{1+\eta} (H_t)^{1+\eta} \text{ if } W_t \geq g(W_t, H_t, z_t, \xi_t^r) \end{aligned} \quad (4)$$

with

$$g(W_t, H_t, z_t, \xi_t^r) = f(W_t, H_t, z_t) + \xi_t^r$$

where $\alpha > 0$ is the change in the marginal utility of income upon reaching the reference point. Effort is measured in cumulative hours worked, the cost of effort is denoted by ψ , and η is related to elasticity. I assume $\eta > 0$, which implies an increasing marginal cost of effort.

$g(W_t, H_t, z_t, \xi_t^r)$ is the reference point at trip t . It is composed of two parts: $f(W_t, H_t, z_t)$, which is a prediction model that the econometrician can construct using observable variables, and ξ_t^r , which is the unobserved component of the reference point. A simple modeling choice may be to assume $f(W_t, H_t, z_t) = 0 \forall W_t, H_t, z_t$, so that ξ_t^r is the reference point.

The econometrician may instead be interested in understanding the possible impact of other variables in the formation of the reference point and assume a $f(W_t, H_t, z_t)$ that depends on the observable variable, ξ_t^r is then the unobserved component partialled out of the effect that the observable variables have on the formation of the reference point. In the main application of this paper, I will be using this second specification of the reference point with a function $f(W_t, H_t, z_t)$ defined in Section 3 and is a proxy for rational expectations.

A notable consequence of specifying a non-constant $f(W_t, H_t, z_t)$ is that the econometrician can insert any number of observable characteristics and test whether the added characteristic influences the transition matrix of the residual component. In principle, if the reference point is fully specified by the function $f(W_t, H_t, z_t)$ then ξ_t^r always be 0 and never move. The function $f(W_t, H_t, z_t)$ can hence help determine the drivers of the reference point and distinguish between theories on reference point formation. Increasing the number of observable characteristics included in the reference

point function comes at the expense of increasing the problem’s dimensionality. I address these concerns more precisely in Section 4.

Let us denote by $s : \mathbb{N}^S \rightarrow \mathbb{N}$, where $S = \dim(\mathcal{W}) \times \dim(\mathcal{H}) \times \dim(\mathcal{Z})$ ⁶, the bijective function that transforms each combination of observable states into a one-dimensional state⁷. Being bijective, allows us to reconstruct the state of the world at any given point in time and I denote $s_t = s(W_t, H_t, z_t)$.

We consider a time separable work utility model in which the problem resets every time the agent decides to stop. I consider the distribution of observables and unobservables to be known and deterministic at $t = 0$ and at each time t where $a_{t-1} = 0$. It is worth noticing that the nature of the observed state variables imposes some further structure on the transition matrix of the observed state variable

$$\Pr(s_t = j | s_{t-1} = k, a_{t-1} = 0) = \begin{cases} \theta_{jk}^S & \text{if } W_t \geq W_{t-1} \text{ and } H_t \geq H_{t-1} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

while I impose no additional structure on θ_{mqj}^X with respect to equation 3.

If $a_t = 1$ then $x_{t+1} = x_0$ and $s_{t+1} = s(0, 0, z_0)$, where x_0 and z_0 are fixed initial values. This is equivalent to the assuming the daily labor supply model to be time separable.

The decision to model the problem as being time separable implies that each shift starts from the same level of observable and unobservable variables. I think this assumption is reasonable in many applications, including NYC taxi drivers, the level of earned income and worked hours is always zero at the start of each shift. Thakral and Tô (2021) show that the assumption of time separability is realistic for NYC taxi drivers.

I assume the following structure for the function $h(x_t, x_{t-1}, s_t)$ of equation 3

$$h(x_t, x_{t-1}, s_t) = h_X(x_t, x_{t-1}) + h_S(x_t, s_t)$$

where

⁶Where \mathcal{W} , \mathcal{H} and \mathcal{Z} are the supports of the cumulative income earned per shift across all shifts, the cumulative hours worked across all shifts and the other state variables that may be included respectively.

⁷For simplicity, I define s such that the codomain is $\{1, \dots, S\}$

$$\begin{aligned}
h_X &= \gamma_1 \mathbb{1}\{x_t - x_{t-1} < 0\} |x_t - x_{t-1}| + \gamma_2 \mathbb{1}\{x_t - x_{t-1} > 0\} |x_t - x_{t-1}| \\
h_S &= \mathbb{1}\{f(s_t) \geq f(s_1)\} (\gamma_3 \mathbb{1}\{x_t - f(s_t) < 0\} |x_t - f(s_t)| + \gamma_4 \mathbb{1}\{x_t - f(s_t) > 0\} |x_t - f(s_t)|) + \\
&\quad \mathbb{1}\{f(s_t) < f(s_1)\} (\gamma_5 \mathbb{1}\{x_t - f(s_t) < 0\} |x_t - f(s_t)| + \gamma_6 \mathbb{1}\{x_t - f(s_t) > 0\} |x_t - f(s_t)|)
\end{aligned}$$

I do not impose any restrictions on any of the γ coefficients, so they may also be equal to zero, which would imply that characteristic does not influence the transition probabilities between states. Notice that I only need to focus on the case in which $a_{t-1} = 0$ since all other transitions are deterministic given the model.

I believe that the structure of the model assumed here, including the dependence structure of state variables described in equations 1 and 2, are very suited to the application of daily labor supply of NYC taxi drivers and gig economy workers. The workers considered do usually not have an influence on who the next customer will be, which naturally translates into an exogenous process of the observable state variable. Workers can though decide at any point in time to stop working and reset the problem to the following shift and this decision is influenced by both the observable state variables and the unobserved reference point.

3 Application

I apply the model outlined in Section 2.1 to the daily labor supply of NYC taxi drivers. I recover the evolution of the reference point and some insights into its determinants. I use trip-level provided by the NYC Taxi and Limousine Commission (TLC) for every fare served by NYC medallion taxicabs in 2013. These data are collected and transmitted electronically in accordance with the Taxicab Passenger Enhancements Project (TPEP). The data consists of all the trip sheets of the trips made by NYC medallion taxicabs in 2013 and provides detailed information on each trip's start and end time, the pick-up and drop-off locations, the fare charged, and the tips paid by credit card. Haggag and Paci (2014) and Farber (2015) provide further details about the data.

I mainly follow the variable construction of Thakral and Tô (2021) for my variables of interest. Shifts are not explicitly defined in the data. I, therefore, define a shift as a set of consecutive trips with less than 6 hours break between trips. Figure A4 shows the distribution of the length of the shifts in terms of hours in the sample. Because of the importance of properly defining

and characterizing the trips and shifts in this application, I discard any shift with contradictory information in the sample. Further distributions of interest are presented in Figures A3 and A2. In particular, Figure A3 provides information that I use to construct the values that the unobserved reference point may take. Figure A2 gives a general overview of the stopping probability at each trip, which will become important in interpreting the results of the analysis.

Notice that the restrictions that I impose on the observable and unobservable states imply that while the observable state variables, cumulative income earned and cumulative hours worked, may directly affect the unobserved state variables, the opposite is not true. One way to interpret this is that the taxi driver cannot choose the next customer depending on the level of the reference point or, equivalently, that the stream of customers is an exogenous process. I believe this assumption to be reasonable for this setting.

I set $f(s_t)$ to be the average end of shift income for taken over those shifts that contain s_t . Formally, define $A(s_t) = \{W^{T\tau} : \exists t' \in \tau \text{ s.t. } s_t = s(W_{t'}, H_{t'}, z_{t'})\}$ and $f(s_t)$ is then the average of the elements of $A(s_t)$. I choose this functional form because it is a proxy for a rational expectations model, which is the reference point assumed in the work of Kőszegi and Rabin (2006), and hence a natural starting point. Recall that Thakral and Tô (2021) find evidence against it in an application to NYC taxi drivers. To decide which possible values the reference point may take for an individual, I consider the difference between the maximum and the minimum end of shift income. This difference informs the values that ξ can take and creates a discrete grid of values that make it possible for any reference point to reach any other reference point at any observable state.

One way to solve the initial condition problem is to consider a reasonable reference point at the beginning of the shift. Several such points have been considered in the past. One example is the average income of the previous 30 shifts, similar to Crawford and Meng (2011). With more data available, one may also consider the income of the same period in the years before to avoid biases due to seasonality. In this application, I consider the initial value of the reference point to be the average end of shift income for the individual; this approach solves the initial condition problem by setting ξ_t^r at 0 at the beginning of each shift.⁸

⁸I also conduct a robustness check where splitting the sample randomly and using the average end-of-shift income of the other sample partition as the initial reference point. The results do not change qualitatively.

4 Estimation

In the specific application considered in this paper, the time dimension of the panel data is generally long, see Figure A5 I can therefore estimate the model separately for each individual while keeping good asymptotic properties. Furthermore, I adopt a Bayesian estimation approach with Normal priors centered at 0 on all parameters. I truncate the priors for α , ψ , and η to a positive support, following economic intuition on the meaning of these parameters. The Bayesian approach has the advantage of generally behaving better in highly nonlinear problems where maximum likelihood methods may get stuck in local maxima. I currently restrict the estimation to individuals with histories longer than 6,000 individuals. I plan to extend the analysis to a model pooling individuals rather than estimating it only at the individual level and also to include individuals with shorter histories in such exercise.

I follow Rust (1994) and proceed with an NFXP algorithm. The NFXP algorithm is composed of an inner loop and an outer loop. The inner loop solves the Dynamic Programming problem using a combination of value function iteration and Newton-Kantorovich algorithm, while the outer loop calculates the likelihood. The presence of a serially correlated unobservable implies computing multidimensional integrals of order equal to the length of the individual time series times the number of components of the unobserved state variable. I approximate this integral using a discrete equivalent of particle filtering. Andrieu, Doucet, and Holenstein (2010) and Flury and Shephard (2011) show unbiased likelihood approximations, like the one I use in this paper, inside Metropolis-Hastings result in exact posterior and I can therefore carry out likelihood-based inference using its simulations. I provide further details on the estimation in Appendix E.

The DP must be solved at each iteration of the algorithm. There is, therefore, much to gain by decreasing the number of parameters that need to be estimated or reducing the problem's dimensionality. I take steps in both of these directions with respect to the general model presented in Section 2.

The first difference with respect to the general model is that I do not include any additional state variables z_t in addition to the cumulative earnings W_t and cumulative hours worked H_t . As described in Section 2.1, it is possible to include additional state variables with respect to the variables that directly enter the agent's utility. These variables can be used to provide further flexibility to the reference point. Possible examples are current weather or a rush-hour indicator. I decide not to include them in the current model mainly because of limitations in computational

resources. Each value that z_t can take increases multiplicatively the dimension of the state space; this, in turn, increases the computational time to solve the dynamic program at each iteration of the estimation.

I now describe in more detail the likelihood and how it is calculated.

4.1 Likelihood

Given observations $\{a_{i,1:T_i}, s_{i,1:T_i}\}_{i=1}^N$ and $\theta = (\theta^X, \theta^S, \alpha, \psi, \eta)$ denoting the parameters characterizing the utility function and the transition matrix, the likelihood function is

$$L_N(\theta^X, \theta^S, \alpha, \psi, \eta) = \prod_{i=1}^N \int p(a_i, s_i, x_i | \theta) d(x_i)$$

where $a_i = \{a_{i,t}\}_{t=1}^{T_i}$, $s_i = \{s_{i,t}\}_{t=1}^{T_i}$, $x_i = \{x_{i,t}\}_{t=1}^{T_i}$ where T_i is the last t observed for agent i and

$$p(a_i, s_i, x_i | \theta) = \prod_{t=1}^{T_i} p(a_{i,t} | a_{i,t-1}, s_{i,t}, x_{i,t}; \theta) q(s_{i,t}, x_{i,t} | s_{i,t-1}, x_{i,t-1}, a_{i,t-1}; \theta)$$

In practice I use a particle filter⁹ to calculate the integral

$$\begin{aligned} & \int p(a_i, s_i, x_i | \theta) d(x_i) \\ & \approx \prod_{t=1}^{T_i} \frac{1}{M} \sum_{j=1}^M p(a_{i,t} | a_{i,t-1}, s_{i,t}, x_{i,t}^m) q(s_{i,t}, x_{i,t}^m | s_{i,t-1}, x_{i,t-1}^m, a_{i,t-1}; \theta) \end{aligned}$$

where $x_{i,t}^m \sim h(x_{i,t}^m | x_{i,t-1}^m, s_{i,t}, s_{i,t-1}, a_{i,t-1})$ are the particles, and $h(x_{i,t}^m | x_{i,t-1}^m, s_{i,t}, s_{i,t-1}, a_{i,t-1}) = \frac{q(s_{i,t}, x_{i,t}^m | s_{i,t-1}, x_{i,t-1}^m, a_{i,t-1}; \theta)}{\sum_{i_t^m} q(s_{i,t}, x_{i,t}^m | s_{i,t-1}, x_{i,t-1}^m, a_{i,t-1}; \theta)}$. M is the number of particles extracted, the higher the more precise the approximation. I assume $x_{i,0}^m = f(W_0^\tau, H_0^\tau)$, which is the same $\forall \tau$.

I assume that the unobserved variable is discrete in my application, which implies some advantages from a computational point of view. The particle filter in the discrete case can be simplified in calculation and becomes much faster to compute. By being discrete, I can enumerate in every period t all the values that the unobserved variable may take and calculate the probability of the unobserved process being in that state (see, e.g., Zucchini and MacDonald 2009). For long time series, though some numerical adjustments are needed to avoid underflow, I detail in Appendix E

⁹See Blevins (2016) for an overview on estimating dynamic models with particle filtering

how I avoid these computational issues.

5 Results

I apply the model to the NYC taxi driver data described in Section 3. I study the evolution and determinants of the reference point in two main ways. First, in Section 5.1, I use the estimation results to simulate the evolution of both observable and unobservable components for 100,000 shifts. To improve our understanding of the behavior and implication of reference points, I then conduct an analysis similar to impulse response functions by introducing shocks to the observable state variables and studying the evolution of the reference point. Secondly, in Section 5.2, I recover the most likely path the reference point takes during a shift for a random individual taken from the sample. I then conduct a heterogeneity analysis studying the distribution of the estimated parameters in the sample. Finally, I conclude with a welfare analysis of introducing a minimum hourly wage.

5.1 Impulse response functions

To improve our understanding of the implications of the estimates of the model, I conduct an exercise in the same spirit as impulse response functions. First, I consider a random driver among the ones with more than 6,000 observations in the dataset. Secondly, I use the driver's Maximum a Posteriori (MAP) estimate of the parameters and derive the implied CCPs and transition matrix. I report the MAP estimates I use in Table A1. Using the derived objects, I simulate 100,000 shifts and their path and plot the average at each trip for all the shifts that have not yet ended. To get the impulse response component of this exercise, I introduce an exogenous shock to the evolution of the observable states of the shifts at a given point in time or a given trip and simulate again the evolution of both the observable and the unobservable component from the point of the shock onwards. Introducing a shock allows us to study its effect on the instant reference point, the persistence of the instantaneous effect on the reference point, and the consequences of this change on the stopping probability. In each figure, I plot the evolution of the average cumulative income, the average reference point on the left y-axis, and the cumulative stopping probability on the right axis. It is worth noticing that since the observable state variables are cumulative, I cannot introduce a directly negative shock to the observable, as a cumulative variable cannot decrease. I proxy a negative shock to income by forcing the driver to earn 0\$ for one or more trips. The number of

trips depends on the specific shock I am introducing.

In Figure 2, I study the introduction of a positive and a negative shock on the first trip after the 8-hours mark on each simulated shift. I consider the 8-hours mark to be the first trip within each shift for which the observable cumulative hours worked are 8 hours; the shock is introduced in the following trip. Figure 2 Panel (A) considers a positive income shock of 25\$, which is also the level at which the cumulative income has been discretized. We can see this implies an immediate positive shock to the reference point as well. It is worth noticing that the shock to the reference point is larger than the shock to the observable income; the shock to the reference point is about 31\$. The reference point then decreases slightly, keeping a difference of about 20\$, as the level of observables continues growing at the same pace as the original pattern but at a higher level. The reference point then settles at a higher level and follows a similar trend to the original reference point. Figure 2 Panel (B) considers a negative shock in income with respect to the original simulated path at the same point of the shift and the same magnitude as Panel (A). More specifically, recalling that the level of earned income cannot decrease, I keep the level of earned income the same for two additional trips. We can interpret this shock as a trip in which the driver made no money. We can see that the change in the observable income and its subsequent pattern have almost no effect on the reference point, which then becomes lower for the following periods but not by the same magnitude as the shock. In Figure A8, I keep the level of earned income fixed until the difference in the reference points is of the same magnitude as the initial difference in reference points for a positive shock of 25\$. The difference in observable income is about 52\$, i.e., more than double the positive shock of Figure 2 Panel (A).

Figure 3 follows the same structure as Figure 2, but at a different point of the simulated shifts, the shocks are introduced at the second trip of the shift. We notice that the instant effect on the reference point is very small, and the divergence in later periods, which we also notice to be small, is likely to be driven by differences in the income level rather than an effect of the early shock to income. The pattern confirms the findings of Thakral and Tô (2021), who show that the key income-related determinant of the decision to stop working is recent earnings. I can draw the same conclusions while not parametrizing the evolution of the reference point to be a fixed linear function of expectations.

Figure 4 considers shocks that take into account their effect on the observable component of the reference point, $f(s_t)$. In Panel (A) of Figure 4, I introduce a positive shock of 25\$ that increases the observed component of the reference point by the same amount. In Panel (B) of Figure 4, the

positive shock of 25\$ does not influence the reference point's observed component, which is the same as the original value. We notice a stark difference in the two shocks' effects; Panel (A) shows an increase of 25\$ while we see no change in Panel (B). I take this as evidence that the rational expectations component plays a big role in the evolution of the reference point. Notice that, in contrast to the figures presented so far, Figure 4 simulates the evolution of the cumulative earned income and of the reference point only from the point of the shock onwards. The values before the shock are observed in the data for the cumulative earned income, while the reference point is recovered using a variation of the Viterbi algorithm. I explain more precisely how I find the reference point up to the shock in the next section.

5.2 Most likely path

To estimate the most likely path of the reference point during the shift, I consider a variation of the Viterbi algorithm, often used in the Hidden Markov Model literature (for further details, see Zucchini and MacDonald 2009). The main idea is that the unobserved state has a transition probability, $\Pr(x_t|x_{t-1}, a_{t-1} = 0)$ and an emission probability $P(s_t|x_t)$, i.e. the probability of observing s_t given that the unobserved state is x_t . These two objects allow us to recover the most likely transition from one state to the next and the most likely path of the unobserved state.

In our case, the emission probabilities, $\Pr(s_t|x_t, a_{t-1} = 0)$ are recovered as follows

$$\begin{aligned} \Pr(s_t|x_t, a_{t-1} = 0) &= \frac{\Pr(a_{t-1} = 0, x_t, s_t)}{\Pr(x_t, a_{t-1} = 0)} \\ &= \frac{\Pr(x_t|s_t, a_{t-1} = 0) \Pr(s_t|a_{t-1} = 0) \Pr(a_{t-1} = 0)}{\Pr(x_t|a_{t-1} = 0) \Pr(a_{t-1} = 0)} \\ &= \frac{\sum_{x_{t-1}} \Pr(x_t|x_{t-1}, s_t, a_{t-1} = 0) \sum_{s'_{t-1}} \Pr(s_t|s'_{t-1}, a_{t-1} = 0) \Pr(x_{t-1}, s'_{t-1}, a_{t-1} = 0)}{\sum_{s'_t} \Pr(a_{t-1} = 0, x_t, s'_t)} \end{aligned}$$

while the transition probability $\Pr(x_t|x_{t-1}, a_{t-1} = 0)$ is

$$\Pr(x_t|x_{t-1}, a_{t-1} = 0) = \sum_{s_t} \Pr(x_t|s_t, x_{t-1}, a_{t-1} = 0) \sum_{s'_{t-1}} \Pr(s_t|s'_{t-1}, a_{t-1} = 0) \Pr(s'_{t-1}, x_{t-1}, a_{t-1} = 0)$$

all the terms are recoverable from the transition matrix of observable and unobservables or from the structure of the model described in Section 2. Figure 5 shows the most likely path of the reference point for an individual taken at random in the sample in two shifts: one in which the end-of-shift

income is higher than average and one in which the end-of-shift income is lower than average.

We notice that, in line with the findings of Section 5.1, the realizations of the observable income do not have an effect on the reference point at the beginning of the shift but only after about ten trips, which correspond to about 2.5 hours on average.

5.3 Distribution of parameters

As stated at the beginning of this section, the availability of long samples for several individuals allows me to estimate the model at an individual level. In turn, this allows me to recover the distribution in the sample of both the utility parameters and the parameters characterizing the transition matrix. I plot the kernel density distribution of the Maximum A Posteriori estimates of the sample in Appendix A. I restrict the estimation to individuals with at least 6,000 observations.

We can see there is some heterogeneity in the sample, which implies that allowing for heterogeneity at the individual level allows to uncover some mechanisms on the formation of the reference point and draw more appropriate policy implications. The heterogeneity shown in Figure A7 is in line with the findings of Buchholz, Shum, and Xu (2018). The distribution of α seems, in fact, to be bimodal, indicating that some individuals are “behavioral” in the sense that their marginal utility changes above or below the reference point, while for other drivers the MAP for α is close to 0 indicating that the marginal utility does not vary around the reference point and are therefore more “neoclassical”.

5.4 Welfare analysis

The persistence of a reference point, defined as how much the level of the reference point depends on its past values, has important consequences for welfare analysis in many settings. A recent paper by Reck and Seibold (2021) provides a conceptual framework and an empirical analysis of how to incorporate reference-dependent preferences in the case of the age of retirement. Another field in which the persistence of the reference point and, in general, reference-dependent preferences may be of particular importance is taxation. A more persistent reference point for taxation may imply a more severe welfare loss for individuals affected by a tax increase. Through the considerations of asymmetric updating of the reference point, the evidence gathered in my paper points in the direction that a tax cut would be rapidly absorbed in the formation of the reference point while a tax increase would not. Analyses of this kind can have significant consequences for policymaking

and, in general, welfare analysis of taxation reforms.

The city of London has recently imposed a wage floor for the drivers of the popular ride-sharing company Uber. The estimated model allows us to conduct some welfare analysis on the implications of such policy.¹⁰ The proposed minimum wage is \$12.11 per hour, which currently is less than the hourly wage of 1% of the shifts, as shown in Figure A6. If we do not allow for moral hazard, e.g., by forcing drivers to accept every customer, then the minimum wage will not affect the earnings and welfare of individuals. If, instead, we consider a situation in which drivers may refuse customers and consider a slightly higher minimum wage of about 20\$ per hour, this is likely to have an effect on both the reference point and the welfare of individuals.

6 Conclusions

In this paper, I estimate a dynamic discrete choice model with a latent variable to flexibly estimate reference-dependent utility models of the daily labor supply decisions of NYC taxi drivers. The analysis' first takeaway is that the reference point formation and evolution can be deduced directly from observational data not collected in a laboratory setting. Through counterfactual analysis, I provide evidence on the formation process of reference points that help distinguish between several theories in the literature. The first empirical result of the analysis is that the evolution of the reference point depends on its past values and the current values of observable components. In my application, the relationship between these two forces can vary during the shift. The second empirical result is that there is strong evidence that the reference point adjusts asymmetrically to positive and negative shocks to the observables: positive shocks have a much more prominent effect on the evolution of the reference point than negative ones. The finding is in line with the psychology literature on the reference point. I model the reference point with an observable component driven by rational expectations. Shocks to this component are a crucial element of the evolution of the reference point but do not fully explain the variation. This last observation implies that other theories are needed to fully characterize the reference point, and the presented methodology offers the tools to do so. Future research will adapt the methodology to other contexts, such as strategic interaction, to gain insights into other areas of economics where reference-dependent preferences are considered relevant.

¹⁰See, for example, this LA Times article

7 Figures

Figure 1: Illustration of a Joint Markov Process






Figure 2: Income shocks after 8 hours

Panel A: Positive 25\$ shock to cumulative income after 8 hours




Panel B: Negative shock to cumulative income after 8 hours




Notes: the plots illustrate the evolution of average earned income in a shift, average reference point, and cumulative stopping probability for 100,000 simulated shifts for an individual taken at random from the sample. The left axis measures the levels of earned income and the level of the reference point, which are represented by the dashed and solid lines, respectively. The right axis represents the cumulative stopping probability. As a counterfactual analysis, I introduce a shock in the evolution of cumulative income. The black lines represent the original evolution of the variables described, while the red line the evolution of the variable with the shock.

Figure 3: Income shocks after the second trip of the shift

Panel A: Positive 25\$ shock to cumulative shift income after the second trip of the shift




Panel B: Negative shock to cumulative shift income after the second trip of the shift




Notes: the plots illustrate the evolution of average earned income in a shift, average reference point, and cumulative stopping probability for 100,000 simulated shifts for an individual taken at random from the sample. The left axis measures the levels of earned income and the level of the reference point, which are represented by the dashed and solid lines, respectively. The right axis represents the cumulative stopping probability. As a counterfactual analysis, I introduce a shock in the evolution of cumulative income. The black lines represent the original evolution of the variables described, while the red line the evolution of the variable with the shock.

Figure 4: Income shocks considering rational expectations

Panel A: Positive 25\$ shock altering rational expectations component of reference point




Panel B: Positive 25\$ shock leaving rational expectations component of reference point unchanged




Notes: the plots illustrate the evolution of average earned income in a shift, average reference point, and cumulative stopping probability for 100,000 simulated shifts for an individual taken at random from the sample. The left axis measures the levels of earned income and the level of the reference point, which are represented by the dashed and solid lines, respectively. The right axis represents the cumulative stopping probability. As a counterfactual analysis, I introduce a shock in the evolution of cumulative income. The black lines represent the original evolution of the variables described, while the red line the evolution of the variable with the shock.

Figure 5: Most likely path of the reference point

Panel A: Most likely path of the reference point for below-average end of shift income



Panel B: Most likely path of the reference point for above-average end of shift income



Notes: the plots illustrate the evolution of two specific shifts observed in the data for the same individual considered in Figures 2 and 3. Panel (A) is a shift where the end of shift income is lower than average, Panel (B) is a shift where the end of shift income is higher than average. The dotted lines are the observed cumulative income within the shift. The solid line is the most likely path of the reference point obtained through the Viterbi algorithm.

References

- Abeler, Johannes et al. (2011). “Reference points and effort provision”. In: *American Economic Review* 101.2, pp. 470–92.
- Aguirregabiria, Victor et al. (2020). “A dynamic structural model of virus diffusion and network production: A first report”. In.
- Andrieu, Christophe, Arnaud Doucet, and Roman Holenstein (2010). “Particle markov chain monte carlo methods”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72.3, pp. 269–342.
- Arkes, Hal R et al. (2008). “Reference point adaptation: Tests in the domain of security trading”. In: *Organizational Behavior and Human Decision Processes* 105.1, pp. 67–81.
- (2010). “A cross-cultural study of reference point adaptation: Evidence from China, Korea, and the US”. In: *Organizational Behavior and Human Decision Processes* 112.2, pp. 99–111.
- Baillon, Aurélien, Han Bleichrodt, and Vitalie Spinu (2020). “Searching for the reference point”. In: *Management Science* 66.1, pp. 93–112.
- Banerji, Abhijit and Neha Gupta (2014). “Detection, identification, and estimation of loss aversion: Evidence from an auction experiment”. In: *American Economic Journal: Microeconomics* 6.1, pp. 91–133.
- Barberis, Nicholas, Ming Huang, and Richard H Thaler (2006). “Individual preferences, monetary gambles, and stock market participation: A case for narrow framing”. In: *American economic review* 96.4, pp. 1069–1090.
- Baucells, Manel, Martin Weber, and Frank Welfens (2011). “Reference-point formation and updating”. In: *Management Science* 57.3, pp. 506–519.
- Bhattacharya, Rabi N and Mukul Majumdar (1989). “Controlled semi-markov models-the discounted case”. In: *Journal of Statistical Planning and Inference* 21.3, pp. 365–381.
- Blevins, Jason R (2016). “Sequential Monte Carlo methods for estimating dynamic microeconomic models”. In: *Journal of Applied Econometrics* 31.5, pp. 773–804.
- Bowman, David, Deborah Minehart, and Matthew Rabin (1999). “Loss aversion in a consumption–savings model”. In: *Journal of Economic Behavior & Organization* 38.2, pp. 155–178.
- Bruneel-Zupanc, Christophe Alain (2021). “Discrete-continuous dynamic choice models: identification and conditional choice probability estimation”. In.


- Buchholz, Nicholas, Matthew Shum, and Haiqing Xu (2018). “Flexible Work Hours and Dynamic Labor Supply: Evidence from Taxi Drivers”. In: *Available at SSRN 2748697*.
- Bunting, Jackson (2021). “Essays on the Econometrics of Dynamic Discrete Choice”. PhD thesis. Duke University.
- Camerer, Colin et al. (1997). “Labor supply of New York City cabdrivers: One day at a time”. In: *The Quarterly Journal of Economics* 112.2, pp. 407–441.
- Cerulli-Harms, Annette, Lorenz Goette, and Charles Sprenger (2019). “Randomizing endowments: An experimental study of rational expectations and reference-dependent preferences”. In: *American Economic Journal: Microeconomics* 11.1, pp. 185–207.
- Chen, M Keith et al. (2019). “The value of flexible work: Evidence from uber drivers”. In: *Journal of political economy* 127.6, pp. 2735–2794.
- Connault, Benjamin (2016). “Hidden Rust Models”. In: *Working Paper*.
- Crawford, Vincent P and Juanjuan Meng (2011). “New york city cab drivers’ labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income”. In: *American Economic Review* 101.5, pp. 1912–32.
- Decker, Wolfram et al. (2017). *SINGULAR 4A computer algebra system for polynomial computations, 2014*.
- DellaVigna, Stefano (2018). “Structural behavioral economics”. In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 1. Elsevier, pp. 613–723.
- DellaVigna, Stefano et al. (2017). “Reference-dependent job search: Evidence from Hungary”. In: *The Quarterly Journal of Economics* 132.4, pp. 1969–2018.
- Farber, Henry S (2008). “Reference-dependent preferences and labor supply: The case of New York City taxi drivers”. In: *American Economic Review* 98.3, pp. 1069–82.
- (2015). “Why you can’t find a taxi in the rain and other labor supply lessons from cab drivers”. In: *The Quarterly Journal of Economics* 130.4, pp. 1975–2026.
- Flury, Thomas and Neil Shephard (2011). “Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models”. In: *Econometric Theory* 27.5, pp. 933–956.
- Genesove, David and Christopher Mayer (2001). “Loss aversion and seller behavior: Evidence from the housing market”. In: *The quarterly journal of economics* 116.4, pp. 1233–1260.
- Gentry, Matthew L and Martin Pesendorfer (2018). “Price Reference Effects in Consumer Demand”. In.

- Gill, David and Victoria Prowse (2012). “A structural analysis of disappointment aversion in a real effort competition”. In: *American Economic Review* 102.1, pp. 469–503.
- Gneezy, Uri et al. (2017). “The limits of expectations-based reference dependence”. In: *Journal of the European Economic Association* 15.4, pp. 861–876.
- Haggag, Kareem and Giovanni Paci (2014). “Default tips”. In: *American Economic Journal: Applied Economics* 6.3, pp. 1–19.
- Heffetz, Ori and John A List (2014). “Is the endowment effect an expectations effect?” In: *Journal of the European Economic Association* 12.5, pp. 1396–1422.
- Hotz, V Joseph and Robert A Miller (1993). “Conditional choice probabilities and the estimation of dynamic models”. In: *The Review of Economic Studies* 60.3, pp. 497–529.
- Hu, Yingyao and Susanne M Schennach (2008). “Instrumental variable treatment of nonclassical measurement error models”. In: *Econometrica* 76.1, pp. 195–216.
- Hu, Yingyao and Matthew Shum (2012). “Nonparametric identification of dynamic models with unobserved state variables”. In: *Journal of Econometrics* 171.1, pp. 32–44.
- Jackson, Emilie (2019). “Availability of the gig economy and long run labor supply effects for the unemployed”. In: *Work. Pap., Stanford Univ., Stanford, CA Google Scholar Article Location*.
- Kasahara, Hiroyuki and Katsumi Shimotsu (2009). “Nonparametric identification of finite mixture models of dynamic discrete choices”. In: *Econometrica* 77.1, pp. 135–175.
- Kőszegi, Botond and Matthew Rabin (2006). “A model of reference-dependent preferences”. In: *The Quarterly Journal of Economics* 121.4, pp. 1133–1165.
- Marzilli Ericson, Keith M and Andreas Fuster (2011). “Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments”. In: *The Quarterly Journal of Economics* 126.4, pp. 1879–1907.
- O’Donoghue, Ted and Charles Sprenger (2018). “Reference-dependent preferences”. In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 1. Elsevier, pp. 1–77.
- Panunzi, Fausto, Nicola Pavoni, and Guido Tabellini (2020). “Economic Shocks and Populism”. In.
- Petterson, Marco Stenborg, David G Seim, and Jesse M Shapiro (2020). *Bounds on a Slope from Size Restrictions on Economic Shocks*. Tech. rep. National Bureau of Economic Research.
- Post, Thierry et al. (2008). “Deal or no deal? decision making under risk in a large-payoff game show”. In: *American Economic Review* 98.1, pp. 38–71.
- Reck, Daniel and Arthur Seibold (2021). “The Welfare Economics of Reference Dependence”. In.

- Rust, John (1987). “Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher”. In: *Econometrica: Journal of the Econometric Society*, pp. 999–1033.
- (1994). “Structural estimation of Markov decision processes”. In: *Handbook of econometrics* 4, pp. 3081–3143.
- Sarpietro, Silvia (2021). “Dynamic Unobservable Heterogeneity: Income Inequality and Job Polarization”. In.
- Schennach, Susanne (2021). “Measurement Systems”. In: *Journal of Economic Literature*.
- Smith, Alec (2019). “Lagged beliefs and reference-dependent utility”. In: *Journal of Economic Behavior & Organization* 167, pp. 331–340.
- Thakral, Neil and Linh T. Tô (Aug. 2021). “Daily Labor Supply and Adaptive Reference Points”. In: *American Economic Review* 111.8, pp. 2417–43.
- Zucchini, Walter and Iain L MacDonald (2009). *Hidden Markov models for time series: an introduction using R*. Chapman and Hall/CRC.

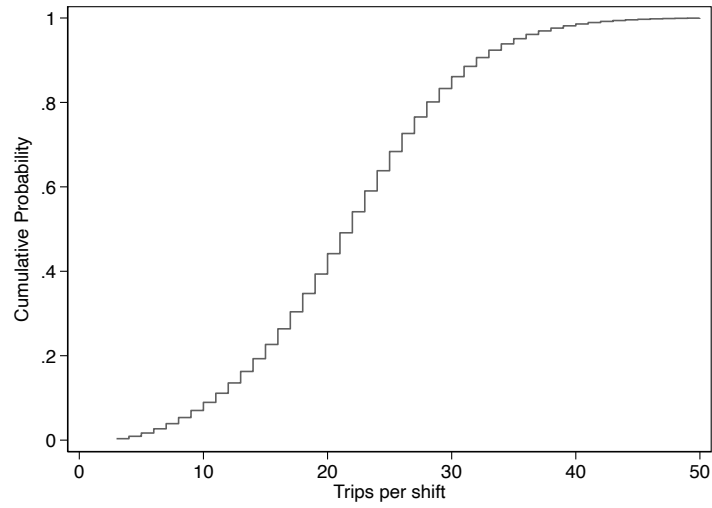
A Additional Tables and Figures

Figure A1: Distribution of fares




Notes: The plot illustrates the distribution of the fares paid for each trip in the sample.

Figure A2: Cumulative distribution of number trips in a shift



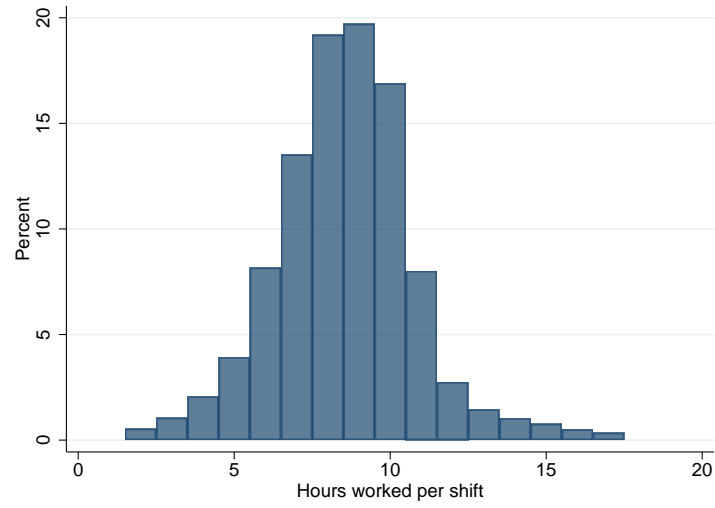
Notes: The plot illustrates the cumulative probability of a shift ending before a specific trip in the sample.

Figure A3: Distribution of end of shift earnings



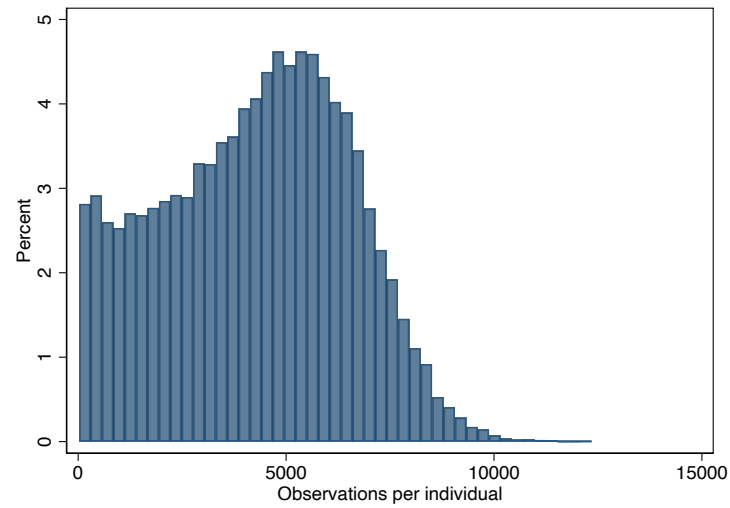
Notes: The plot illustrates the distribution of income earned by drivers at the end of a shift in the sample.

Figure A4: Distribution of the length of shifts in hours



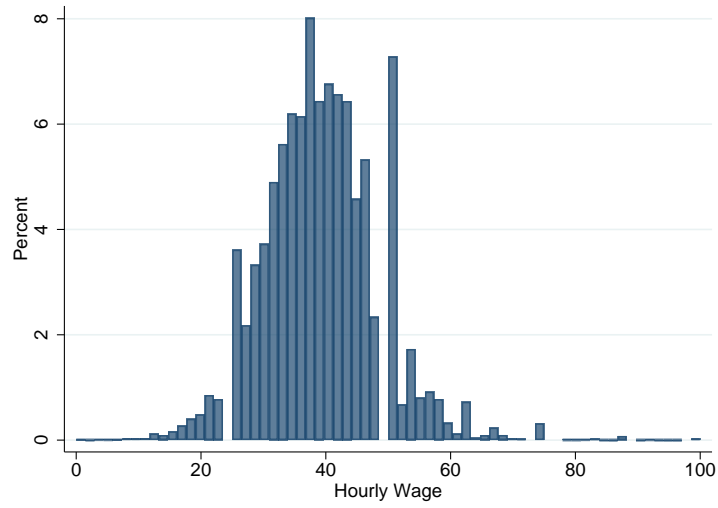
Notes: the plot illustrates the distribution of the length of shifts measured in hours in the sample.

Figure A5: Distribution of the number of observations per individual



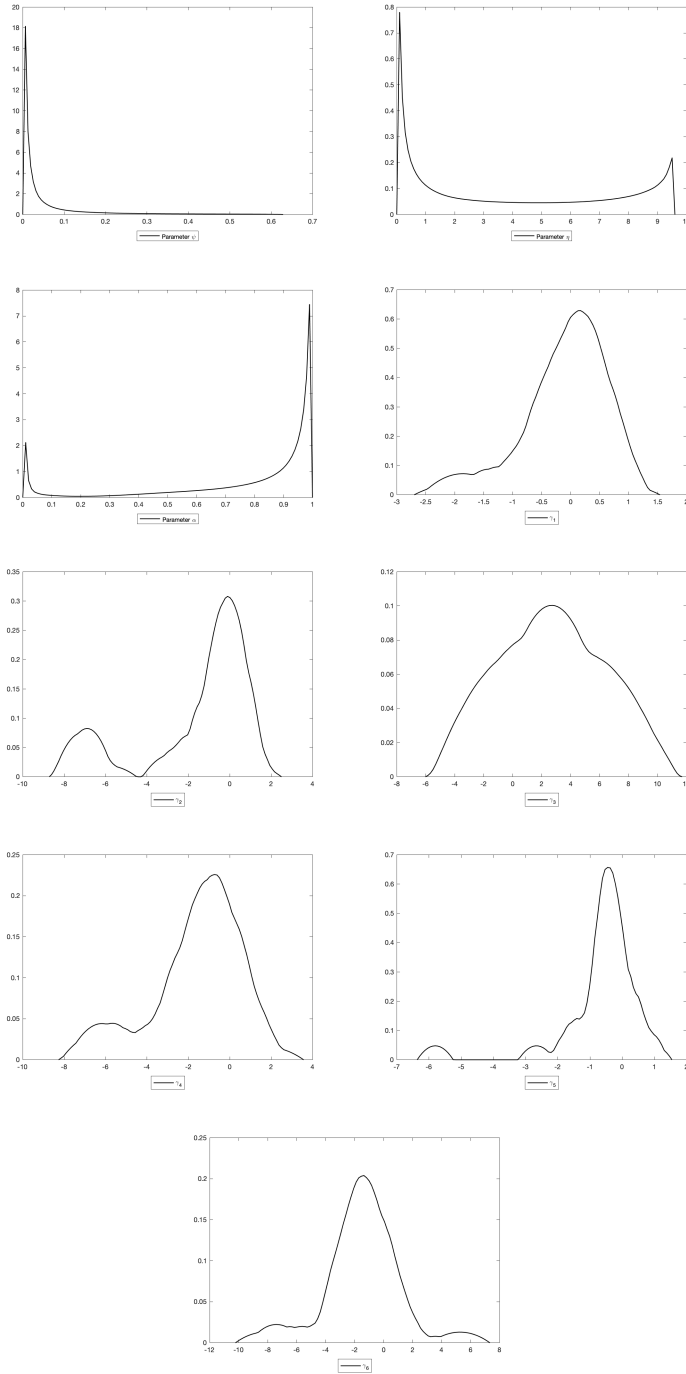
Notes: the plot illustrates the distribution of the number of observations for each individual in the sample.

Figure A6: Distribution of hourly wage per shift



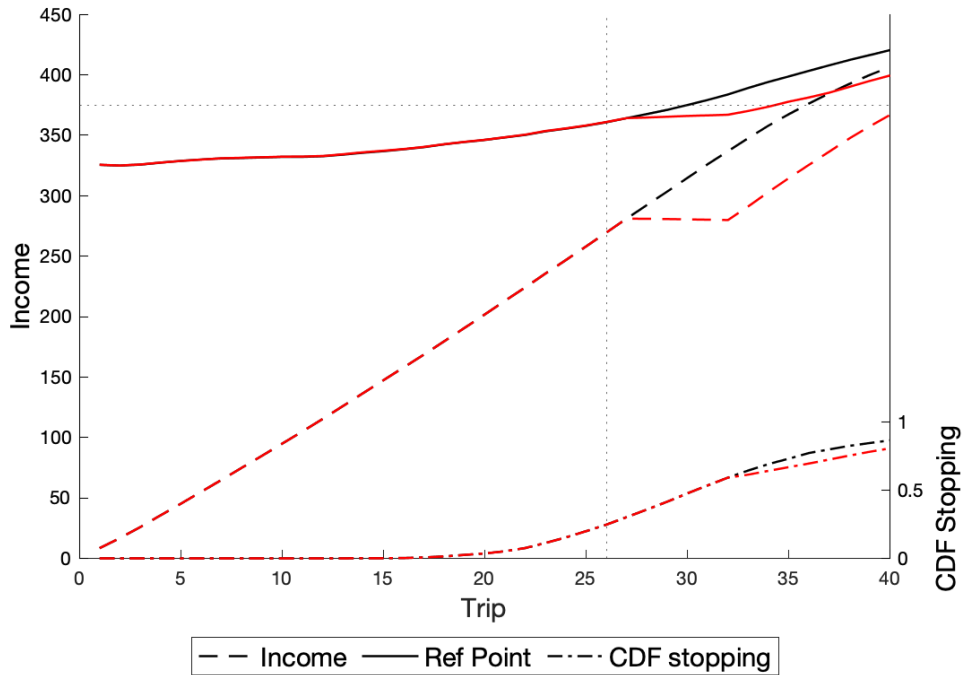
Notes: the plot illustrates the distribution of the hourly wage for all shifts in the sample.

Figure A7: Plots of kernel density of MAP from the sample



Notes: the figures represent the kernel density, estimated with an Epanechnikov kernel, of the maximum a posteriori (MAP) estimates of the individual-level analysis conducted in the sample.

Figure A8: Large negative income shock



Notes: the plots illustrate the evolution of average earned income in a shift, average reference point, and cumulative stopping probability for 100,000 simulated shifts for an individual taken at random from the sample. The left axis measures the levels of earned income and the level of the reference point, which are represented by the dashed and solid lines, respectively. The right axis represents the cumulative stopping probability. As a counterfactual analysis, I introduce a shock in the evolution of cumulative income. The black lines represent the original evolution of the variables described, while the red line the evolution of the variable with the shock.

Table A1: Estimates for individual in Impulse Response functions

Parameters	
ψ	0.425 [0.021,0.581]
η	0.0912 [0.017,1.608]
α	0.956 [0.949,0.961]
γ_1	-0.566 [-0.668,-0.523]
γ_2	-2.11 [-2.242,-2.042]
γ_3	0.364 [0.327,0.467]
γ_4	1.974 [1.904 ,2.103]
γ_5	-1.97 [-2.058 ,0.177]
γ_6	-0.79 [-1.13 ,-0.541]

Notes: the table reports the MAP estimates used in the Impulse Response figures.

B Alternative Specification

In this section I explore an alternative utility specification, which is also used in Farber (2008) and Thakral and Tô (2021). I define

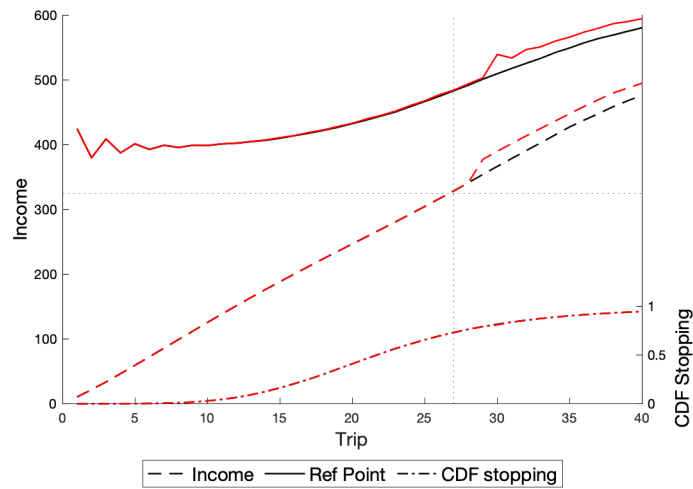
$$u_1(W_t, H_t, z_t, \xi_t^r) = (1 - \delta) \left(W_t - \frac{\psi}{1+\eta} (H_t)^{1+\eta} \right) + \delta \lambda (W_t - g(W_t, H_t, z_t, \xi)) \text{ if } W_t < g(W_t, H_t, z_t, \xi_t^r)$$
$$u_1(W_t, H_t, z_t, \xi_t) = (1 - \delta) \left(W_t - \frac{\psi}{1+\eta} (H_t)^{1+\eta} \right) + \delta (W_t - g(W_t, H_t, z_t, \xi)) \text{ if } W_t \geq g(W_t, H_t, z_t, \xi_t^r)$$

where λ is then the loss aversion coefficient while δ is the relative weight given to the two components of the utility. All other components have the same meaning as in Section 2.1.


I replicate the same impulse response functions and exercises as in Section 5 and find similar results. Figure A9 shows again the asymmetry in the response to positive and negative shocks. Figure A10, similar to Figure 3, shows that the effect of a shock at the beginning of the shift is minimal. Both specifications allow hence to draw the same conclusions on the reference point for NYC taxi drivers.

Figure A9: Income shocks after 8 hours

Panel A: Positive 25\$ shock to cumulative income after 8 hours




Panel B: Negative shock to cumulative income after 8 hours




Notes: the plots illustrate the evolution of average earned income in a shift, average reference point, and cumulative stopping probability for 100,000 simulated shifts for an individual taken at random from the sample. The left axis measures the levels of earned income and the level of the reference point, which are represented by the dashed and solid lines, respectively. The right axis represents the cumulative stopping probability. As a counterfactual analysis, I introduce a shock in the evolution of cumulative income. The black lines represent the original evolution of the variables described, while the red line the evolution of the variable with the shock.

Figure A10: Income shocks after the second trip of the shift

Panel A: Positive 25\$ shock to cumulative shift income after the second trip of the shift



Panel B: Negative shock to cumulative shift income after the second trip of the shift



Notes: the plots illustrate the evolution of average earned income in a shift, average reference point, and cumulative stopping probability for 100,000 simulated shifts for an individual taken at random from the sample. The left axis measures the levels of earned income and the level of the reference point, which are represented by the dashed and solid lines, respectively. The right axis represents the cumulative stopping probability. As a counterfactual analysis, I introduce a shock in the evolution of cumulative income. The black lines represent the original evolution of the variables described, while the red line the evolution of the variable with the shock.

C Data processing details

I follow a process similar to Haggag and Paci (2014) and Thakral and Tô (2021). Given the methodology's emphasis on the exact evolution of the observable characteristics and the stopping decisions, I am particularly keen on having shifts that are correct at every trip. Any inconsistency implies that the whole shift gets dropped. Also, given that I estimate the model at an individual level, I do not consider drivers that are rarely observed, as they would not provide much information. The data cleaning proceeded as follows:

1. If the payment type for one trip is "No Charge", "Dispute" or "Unknown" the shift is dropped from the sample
2. If the time of a trip is less than a minute, the trip is dropped from the sample
3. If the shift has less than three trips, it is dropped from the sample
4. Drivers with less than 100 trips in the sample are dropped
5. If a shift has more than 50 trips, it is dropped from the sample

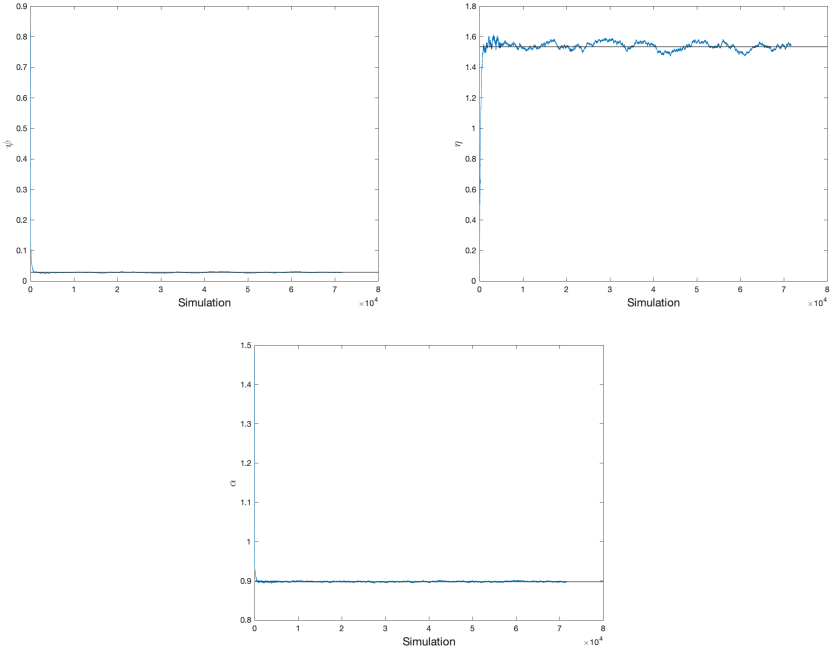
In general, if other inconsistencies, e.g., different trips taking place simultaneously, are found, then the whole shift with the inconsistency is dropped. I take a very conservative approach to data cleaning, but the correct definition of trips and shifts is essential given the level and type of analysis. Overall I reduce the sample by about 18 percent.

The cumulative income level within each shift, W_t , is constructed by first calculating the cumulative sum of income within a shift and then discretizing it at the 25\$ level. The cumulative hours worked, H_t , are constructed by first calculating the cumulative time worked within a shift and then discretizing this variable at the 1-hour level.

D Parameter Recovery

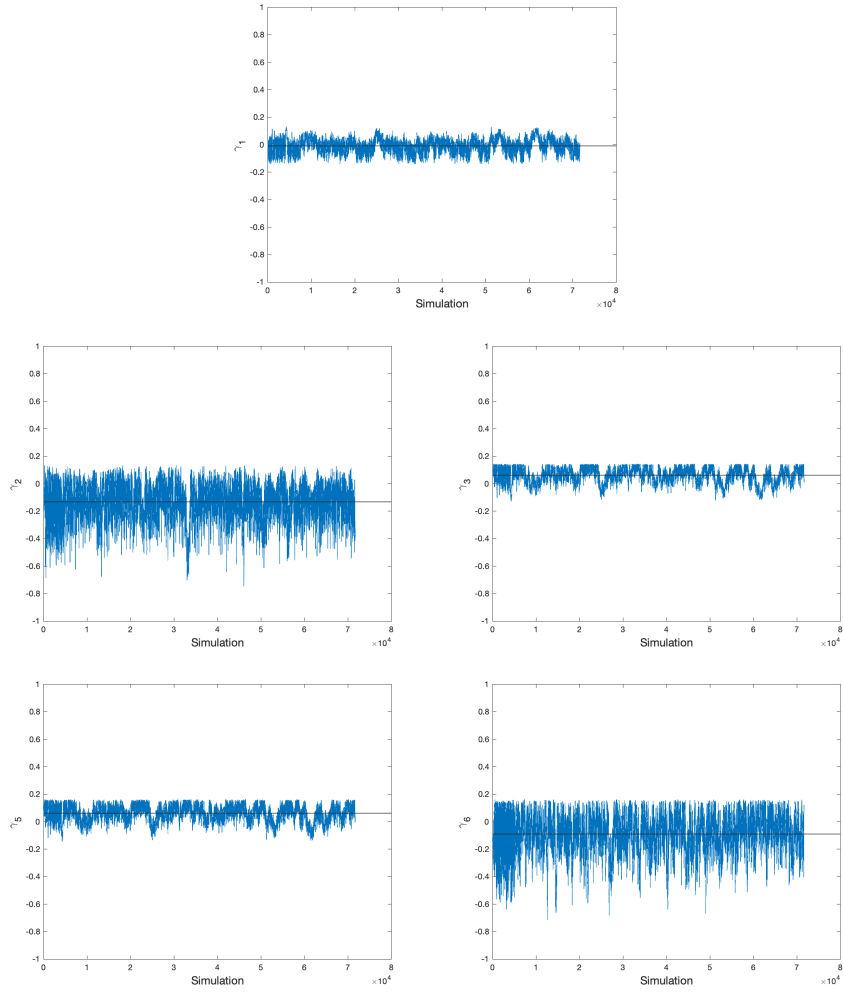
In this Appendix I report some plots that show that the model recovers the DGP parameters when run on estimated data. I simulate 50,000 draws from the model with parameters similar to the ones found in the estimation in the true sample and I see whether the true parameters are recovered. We can see from Figures A11 and A12 that this is the case

Figure A11: Plots of parameter recovery utility function parameters



Notes: the blue line represent the MCMC draws for the parameter. The black solid line represent the original value of the parameter used to simulate the data on which the model is estimated and convergence is tested.

Figure A12: Plots of parameter recovery transition matrix parameters



Notes: the blue line represent the MCMC draws for the parameter. The black solid line represent the original value of the parameter used to simulate the data on which the model is estimated and convergence is tested.

D.1 Identification result

This section considers the identification results lying at the core of the model presented in Section 2.1. I model the observable and unobservable state variable, $\{s_t, x_t\}_{t=1}^T$, as a joint Markov process. I then impose some further assumptions

Assumption. *Markov process and limited feedback*

$$(i) \Pr(s_t, x_t | s_{1:t-1}, x_{1:t-1}, a_{1:t-1}) = \Pr(s_t, x_t | s_{t-1}, x_{t-1}, a_{t-1})$$

$$(ii) \Pr(s_t, x_t | s_{t-1}, x_{t-1}, a_{t-1}) = \Pr(x_t | s_t, x_{t-1}, a_{t-1}) \Pr(s_t | s_{t-1}, a_{t-1})$$

where for any variable y , $y_{1:t-1}$ stands for the set of all the realizations of y from period 1 to period $t - 1$. The first part of the assumption is satisfied under Markovian dynamic decision models (Rust 1994). The second part of the assumption mirrors Assumption 1 in Hu and Shum (2012) and rules out direct feedback of last period's unobserved state variable x_{t-1} on this period's observed state variable s_t . The only feedback possible is through the decision of the agent a_{t-1} .

The value of the reference point at time t is affected by the reference point's value at time $t - 1$, the value of the observed state at time t , and the decision of the agent at time $t - 1$.

We can exemplify the structure of this model with Figure 1

The identification problem can then be written as a system of equations that links the transition probabilities of the variables observable both by the agent and the econometrician to the law of motion of the unobserved state variable. Let us focus on the case where $a_t = a_{t-1} = 0$, all the other cases follow naturally. For each combination of $j, k \in \mathcal{S}$ I can write

$$\begin{aligned} & \Pr(a_t=0 | s=j, a_{t-1}=0, s_{t-1}=k) \\ &= \sum_m \Pr(a_t=0, x_t=m | s_t=j, a_{t-1}=0, s_{t-1}=k) \\ &= \sum_m \Pr(a_t=0 | s_t=j, x_t=m, a_{t-1}=0, s_{t-1}=k) \Pr(x_t=m | s_t=j, a_{t-1}=0, s_{t-1}=k) \\ &= \sum_m \Pr(a_t=0 | s_t=j, x_t=m) \sum_q [\Pr(x_t=m | s_t=j, a_{t-1}=0, s_{t-1}=k, x_{t-1}=q) \Pr(x_{t-1}=q | s_t=j, a_{t-1}=0, s_{t-1}=k)] \\ &= \sum_m P_{jm} \sum_q \left[\theta_{mqj}^X \frac{\Pr(a_{t-1}=0, s_{t-1}=k, x_{t-1}=q)}{\sum_g \Pr(a_{t-1}=0, s_{t-1}=k, x_{t-1}=g)} \right] \end{aligned} \tag{6}$$

Let us first focus on $\theta_{mqj}^X = \Pr(x_t = m | s_t = j, a_{t-1} = 0, x_{t-1} = q)$, probability of transitioning from $x_{t-1} = q$ to $x_t = m$ when the $s_t = j$. $\theta_{mqj}^X \forall m, q, j$ are the only unknowns in the system of equations as the other components are either observable or functions of θ_{mqj}^X and observable components. To achieve nonparametric identification of this component, we need to be able to solve this highly nonlinear system for θ^X and show that the solution is unique.

One way to solve the system of equations is to show that the system of equations can be rewritten as a fixed point problem and then show that the fixed point operator is a contraction. Banach's fixed point theorem then implies a unique fixed point, which implies that the solution to the system of equations is unique.

It is worth noticing that, since the system of equations represented by equation 6 is an identity as it only uses properties of probabilities, the existence of a solution is guaranteed. The only remaining

question is the uniqueness of such solution.

To understand the structure of the fixed point operator and how I obtain it from the system of equations above, I consider an example where $a \in \{0, 1\}$, $s \in \{A, B, C, D\}$ and $x \in \{\alpha, \beta, \gamma\}$.¹¹

Then equation 6 can be written as

$$\begin{aligned} \Pr(a_t=0|s_t=A, s_{t-1}=A, a_{t-1}=0) &= P_{A\alpha} \left(\theta_{\alpha\alpha A}^X \Gamma_{A\alpha} + \theta_{\alpha\beta A}^X \Gamma_{A\beta} + \theta_{\alpha\gamma A}^X \Gamma_{A\gamma} \right) + \\ &+ P_{A\beta} \left(\theta_{\beta\alpha A}^X \Gamma_{A\alpha} + \theta_{\beta\beta A}^X \Gamma_{A\beta} + \theta_{\beta\gamma A}^X \Gamma_{A\gamma} \right) + P_{A\gamma} \left(\theta_{\gamma\alpha A}^X \Gamma_{A\alpha} + \theta_{\gamma\beta A}^X \Gamma_{A\beta} + \theta_{\gamma\gamma A}^X \Gamma_{A\gamma} \right) \end{aligned}$$

which can in turn be rewritten as

$$P_{AA}^{00} - P_{A\gamma} = (P_{A\alpha} - P_{A\gamma}) \left(\theta_{\alpha\alpha A}^X \Gamma_{A\alpha} + \theta_{\alpha\beta A}^X \Gamma_{A\beta} + \theta_{\alpha\gamma A}^X (1 - \Gamma_{A\alpha} - \Gamma_{A\beta}) \right) + (P_{A\beta} - P_{A\gamma}) \left(\theta_{\beta\alpha A}^X \Gamma_{A\alpha} + \theta_{\beta\beta A}^X \Gamma_{A\beta} + \theta_{\beta\gamma A}^X (1 - \Gamma_{A\alpha} - \Gamma_{A\beta}) \right)$$

$$\text{where } \Gamma_{kq} = \frac{\Pr(a_{t-1}=0, s_{t-1}=k, x_{t-1}=q)}{\sum_g \Pr(a_{t-1}=0, s_{t-1}=k, x_{t-1}=g)}$$

We then notice that we can isolate the relevant θ^X on one side of the equations and obtain a fixed point operator, which I call F . Hence I rewrite $\theta^X = F(\theta^X, \theta^S, \alpha, \psi, \eta)$ and denote with J the jacobian matrix of F . We then get the following proposition

Proposition 1. *Let F be a fixed point operator associated with the system represented by equation 6 and denote by J its Jacobian with respect to θ^X . If the highest eigenvalue of Jacobian matrix is smaller than 1 then the system of equations has a unique solution, and the estimator for θ^X is identified.*

The proof follows from Banach's contraction theorem. It is important to note that the Proposition considers only θ^X and hence assumes that all other parameters can be identified separately. If it is possible to augment the system of equations with a fixed point operator also for the parameters of the utility function, then Proposition 1 also implies the identification of the utility parameters. Notice that θ^S can always be identified directly from the transitions observed in the data. Also, notice that Proposition 1 only gives a sufficient condition and that not all fixed-point operators equivalent to the system of equations may be a contraction. The proposition requires only the existence of one such operator.

An alternative approach is to apply directly the results of Connault (2016) and computing the singular region for which the system is not identified. There are several tools to do this, e.g.,

¹¹Notice that α , β , and γ are not linked to the values they take in the rest of the text but are used for notational ease in this example.

Decker et al. (2017). It is worth noticing though that while the system here considered could be viewed as polynomial, the way that θ^X enters the CCPs is not polynomial. The added nonlinearity implies some additional caution is needed in finding the singular region of the system of equations represented by equation 6.

E Details of Estimation Procedure

In this Appendix I detail exactly how I ran the estimation and how each iteration of the estimation works. The first step, since I am taking a Bayesian approach is to specify the priors:

- $\psi \sim \mathcal{N}(0, 10) \mathbb{1}\{\psi > 0\}$
- $\eta \sim \mathcal{N}(0, 10) \mathbb{1}\{\eta > 0\}$
- $\alpha \sim \mathcal{N}(0, 10) \mathbb{1}\{\alpha \geq 0\}$
- $\theta_1, \dots, \theta_6 \sim \mathcal{N}(0, 10)$

the truncation for ψ , η and η comes from economic intuition on the meaning of the parameters.

Then, to calculate the posterior, I use a Random Walk Metropolis-within-Gibbs sampling in which I sample each parameter separately for the first 10,000 iterations and then all the parameters together with a covariance matrix derived from the first 10,000 draws.

More in detail, at each iteration I draw one parameter (or all parameters) from a proposal distribution that has been tuned to have an acceptance rate of about 0.4. I solve the dynamic program with the new draw using a combination of value function iteration and Newton-Kantorovich. I use as initial value the solution of the previous iteration as it is likely to be close to the new solution. I then calculate the likelihood with the latest draw and the corresponding solution to the dynamic program. To calculate the likelihood, I have to marginalize out the contribution of the unobserved component and to do this I use a discrete filter. For long time series, to avoid underflow of the estimation procedure, I use the following normalization of the discrete filter¹²

$$\text{initialization} \begin{cases} \tilde{\pi}_1 & = \pi_1 = \Pr(a_0, s_0, x_0) \\ \log \rho_1 & = 0 \end{cases}$$

$$\text{iteration} \begin{cases} \pi_{t+1} & = \tilde{\pi}_t Q_{t+1} \\ \tilde{\pi}_{t+1} & = \frac{\pi_{t+1}}{\|\pi_{t+1}\|_1} \\ \log \rho_{t+1} & = \log \rho_t + \|\pi_{t+1}\|_1 \end{cases}$$

where $Q_{t+1,xx'} = \mathbb{P}(x_{t+1} = x', s_{t+1}, a_{t+1} \mid x_t = x, s_{t+1}, a_{t+1})$ i.e. the transition matrix of the unobserved state variable keeping fixed the observed state variables at the observed values.

¹²This normalization procedure follows Connault (2016).

After the likelihood calculation comes the Metropolis - Hastings step where the new draw, x^* , is accepted with probability

$$\mu = \min \left\{ 1, \frac{p(x^*)}{p(x^{j-1})} \right\}$$

where $p(x^*)$ is the likelihood with the draw x^* times its prior.

These steps are repeated for 500,000 iterations, which, by visual inspection, is often more than enough to reach convergence of the estimator.