

Did Harold Zuercher Have Time-Separable Preferences?*

Jay Lu, Yao Luo, Kota Saito, and Yi Xin

June 13, 2024

Abstract

This paper proposes an empirical model of dynamic discrete choice to allow for non-separable time preferences, generalizing the well-known Rust (1987) model. Under weak conditions, we show the existence of value functions and hence well-defined optimal choices. We construct a contraction mapping of the value function and propose an estimation method similar to Rust's nested fixed point algorithm. Finally, we apply the framework to the bus engine replacement data. We improve the fit of the data with our general model and reject the null hypothesis that Harold Zuercher has separable time preferences. Misspecifying an agent's preference as time-separable when it is not leads to biased inferences about structure parameters (such as the agent's risk attitudes) and misleading policy recommendations.

Keywords: dynamic discrete choice, Epstein-Zin preferences, preference for the timing of resolution of uncertainty

*Lu: Department of Economics, University of California, Los Angeles (email: jay@econ.ucla.edu); Luo: Department of Economics, University of Toronto (email: yao.luo@utoronto.ca); Saito: Division of the Humanities and Social Sciences, California Institute of Technology (email: saito@caltech.edu); Xin: Division of the Humanities and Social Sciences, California Institute of Technology (email: yixin@caltech.edu). Financial support from the NSF under awards SES-1919263 (Saito) and SES-1919275 (Lu) are gratefully acknowledged. We thank seminar and conference participants at Caltech, Georgetown University, Texas A&M, University of Arizona, UC Berkeley, UC Irvine, University of Colorado Boulder, and NASMES (2023). We appreciate the valuable discussions we had with Fedor Iskhakov, Peter Klibanoff, Robert Miller, Marciano Siniscalchi, Costis Skiadas, Tomasz Strazalecki and John Rust. In particular, Mitsuru Igami read an earlier version of the manuscript and offered helpful comments.

1 Introduction

It is well-known that in the standard model of expected utility with separable time preferences, an agent’s risk preference and her preference for intertemporal substitution are both captured by the curvature of her Bernoulli utility function. As a result, greater risk aversion is correlated with greater complementarity of consumption across different time periods. Risk and time however are two distinct phenomena and it would be natural to have different preferences over the two. Furthermore, time-separable models are widely recognized for imposing an unrealistic constraint on agents’ preferences regarding the timing of uncertainty resolution. As Rust (1994) pointed out, time-separable “expected-utility models imply that agents are indifferent about the timing of the resolution of uncertain events, whereas human decision-makers seem to have definite preferences over the time at which uncertainty is resolved.”

Models with non-separable time preferences such as Epstein and Zin (1989) allow for a clean separation of the two: the elasticity of intertemporal substitution is independent of risk aversion. Moreover, the models allow agents to have nontrivial preferences regarding the timing of the resolution of uncertainty, as suggested by the experimental literature (see, e.g., Nielsen (2020), Meissner and Pfeiffer (2022)). This feature makes them popular in fields such as macroeconomics and finance, where they have been used to explain phenomena such as the equity premium puzzle (see, e.g., Mehra and Prescott (1985) and Bansal and Yaron (2004)).

However, current empirical research has yet to incorporate non-separable preferences within the dynamic discrete choice (DDC) framework. Misspecifying an agent’s preference as time-separable when it is not can result in biased inferences about structural primitives and misleading policy recommendations. Our paper aims to bridge this gap. We begin by outlining a set of theoretical results crucial for empirical analysis in the DDC framework. Following this, we introduce an empirical model of dynamic discrete choice that incorporates non-separable time preferences, extending the well-known bus engine replacement model analyzed in Rust (1987). Our model includes the standard time-separable model, commonly employed in the DDC liter-

ature, as a special case. This is an important and empirically relevant feature as it allows us to statistically test whether the decision maker (Harold Zuercher in the bus engine replacement example) has separable time preferences or not.

While our model allows for more general and richer time preferences, we show that it can be solved and estimated similarly as the standard DDC model. In the paper, we first show that agents with our nonseparable time-preferences have well-defined optimal dynamic choices. Based on the Bellman equation with our nonseparable time-preferences, we first prove the existence of the value function. Specifically, we apply lattice theory and Tarski's fixed point theorem to show the existence of the smallest and largest fixed points of a mapping in the space of value functions. This allows us to use a simulated nested fixed point algorithm to find the smallest and largest fixed points for model estimation; if the two coincide, then the value function is unique. We also provide general sufficient conditions that ensure a contraction mapping and uniqueness of the value function. These conditions reduce to the standard $\beta < 1$ condition for separable preferences but we also consider specific parametrizations (CARA and CRRA) of Epstein-Zin preferences.

We apply our framework to the bus engine replacement data in Rust (1987). In his original work, Rust assumes that the manager Harold Zuercher is risk neutral with separable time preferences. At the beginning of each period, Harold makes a dynamic choice for each bus engine by trading off between an immediate lump sum cost of replacing it and higher maintenance costs for keeping it. In addition to allowing Harold to have non-separable time preferences, we generalize the Rust model in several other aspects. We assume that the manager may be risk-averse and may benefit from operating the bus (e.g., collecting passenger fares proportional to the additional mileage in each month).

We estimate the model with non-separable and separable time preferences, respectively, under the CARA parameterization. Our estimation results of the non-separable model suggest that the agent is likely to prefer late resolution of uncertainty. Because the model with separable preferences is nested in the non-separable model, comparing the likelihood from the two models is straightforward, and we are able to statistically

reject the null hypothesis that Harold has separable time preferences. Moreover, our estimates for the agent’s risk attitudes and the payoff parameters are significantly different when allowing for non-separable preferences. We find that, given the time-separable model, the risk preference parameter is overestimated, while the maintenance cost per mileage is underestimated.

To better interpret the model estimates, we compute the certainty equivalent given different model specifications. Specifically, we consider a counterfactual scenario in which a subsidy program assists the agent in smoothing revenue across periods. We solve for the monetary payoff that makes the agent (Harold) indifferent when the uncertainty about incremental mileage in each period is eliminated. Our general model with non-separable preferences implies that Harold would accept a sure payment of \$120 each month to stay indifferent. In contrast, the certainty equivalent under the misspecified model with separable preferences is \$61.6, which is 48.7% lower than the value implied by the general model.

Related Literature There are many empirical applications of DDC models. See Miller (1984) and Buchholz et al. (2021) for dynamic choices of workers, Pakes (1986) for patenting decisions, Crawford and Shum (2005), Hendel and Nevo (2006), and Gowrisankaran and Rysman (2012) for dynamic consumer demand. As far as we know, there is no existing work in the DDC literature that has incorporated non-separable time preferences. Our main methodological advance is in developing a feasible specification of DDC models with EZ preferences, arguably the most popular model of non-separable time preferences. Following Rust (1987), we propose a nested fixed point algorithm for estimating our model. Our proposed algorithm can be applied to estimate models with separable and non-separable preferences, and it therefore provides a useful tool for testing various model specifications.

The use of non-separable models has been common in the macrofinance literature since Kreps and Porteus (1978) and Epstein and Zin (1989). For example, Bansal and Yaron (2004) provide a unified explanation for several long-standing puzzles in asset markets using a recursive utility model. Epstein et al. (2014) explore the quantitative

implications of recursive utility on temporal resolution of uncertainty. We find that the risk preference parameter is significantly overestimated under the misspecified model with separable preferences, suggesting the importance of separating preferences for risk and intertemporal substitution.

In the theoretical literature, Frick et al. (2019) provide an axiomatic analysis of dynamic random utility, focusing on the history-dependency of choices and the type of randomness inherent in these models. Lu and Saito (2020) provide a theoretical model of dynamic stochastic choice in which the agent has non-separable preferences. They provide axiomatic foundations and also consider EZ preferences in an application focusing on its choice-theoretic implications. In contrast, our paper focuses on the DDC model as the workhorse model for analyzing dynamic decision processes in structural econometrics.

The rest of the paper is organized as follows. We describe EZ preferences and the dynamic discrete choice model setup in Section 2. Theoretical results regarding the existence and uniqueness of the value function are provided in Section 3. We discuss the estimation strategy and our empirical application in Section 4 and 5, respectively. Section 6 concludes. The proofs are delegated to the appendices.

2 Non-separable Time Preferences

2.1 Setup

In this section, we introduce non-separable time preferences in an infinite-horizon recursive framework. First, recall that under standard risk and intertemporal preferences, the utility function is given by

$$\mathbb{E}_c [(1 - \beta) u(c) + \beta \mathbb{E}_{v|c} [v]]. \quad (1)$$

Here, u is a strictly increasing utility function over a set \mathcal{C} of real-valued payoffs, $\beta \in [0, 1]$ is the discount factor and v is the future continuation value.¹ Let \mathcal{U} denote

¹Note that since the value function is recursive, this is equivalent to $u(c) + \beta v$ as the $1 - \beta$ term is a scalar that can be factored out.

the range of u . It is easy to see here that risk is additively-separable across time.²

Following Kreps and Porteus (1978) and Epstein and Zin (1989), we consider a generalization of standard preferences that allow for *non-separable time preferences*, i.e., where the risk is not additively-separable across time.³

We consider the following representation:

$$\mathbb{E}_c \left[\phi \left((1 - \beta) u(c) + \beta \phi^{-1} \left(\mathbb{E}_{v|c} [v] \right) \right) \right], \quad (2)$$

where ϕ is a strictly increasing aggregator function that is either concave or convex.⁴ Refer to Section 2.2 for parametric examples of the utility function (2), which will aid in understanding the general formulation presented in (2).

As we will see below, the curvature of ϕ captures the agent's attitudes towards how risk is resolved across different time periods. When ϕ is linear, this reduces to the standard time-separable preferences, which implies the indifference to the timing of resolution of uncertainty, otherwise, risk cannot be additively separated across time as in Equation (1). Note that although the above formula is recursive and assumes infinite periods, the functional form can be easily applied to finite time periods by calculating utilities from the last period via backwards induction.

The essential difference from separable preferences is the role of the aggregator ϕ , which allows for greater flexibility in modeling risk and time preferences. In fact, agents may exhibit non-trivial preferences over *when* risk is resolved. Consider the two lotteries shown in Figure 1. Both lotteries are the same in terms of probabilities and outcomes: they differ only in the timing of when uncertainty is resolved. In the left lottery P , the uncertainty about period 2 outcome (10 or 0) is resolved in period 1. In the right lottery Q , on the other hand, the uncertainty is resolved in period 2. If the agent always prefers a lottery where uncertainty is resolved in period 1 (e.g., P) to a lottery where uncertainty is resolved in period 2 (e.g., Q), then we say that


²To see this point, notice that because of the linearity of the model, (1) reduces to $(1 - \beta)\mathbb{E}_c[u(c)] + \beta\mathbb{E}_v[v]$, where the first expectation captures the risk over current consumption; while the second expectation captures the risk over future consumption.

³Due to the non-linearity of ϕ , the utility function (2) cannot be expressed as the additive form of two expectations, unlike the time-separable model described in footnote 2.1.

⁴Epstein and Zin (1989) consider an even more general formulation.

he *prefers early resolution of uncertainty*. If the agent always prefers a lottery where uncertainty is resolved in period 2 (e.g. Q) to a lottery where uncertainty is resolved in period 1 (e.g. P), then we say he *prefers late resolution of uncertainty*. If he is indifferent between the two lotteries, then we say that he is *indifferent to the timing of resolution of uncertainty*.

Figure 1: Early vs. late resolution of uncertainty



In our model, preferences for the timing of resolution of uncertainty are fully captured by the curvature of ϕ . To see why, consider the example above. Applying (2) to two-periods, the value of lottery P is given by

$$V(P) = \mathbb{E}[\phi((1 - \beta)u(5) + \beta z)]$$

where z is a random variable that takes on either $(1 - \beta)u(0)$ or $(1 - \beta)u(10)$. Note here that all uncertainty about z is resolved in the first period. On the other hand, the value of lottery Q is given by

$$V(Q) = \phi((1 - \beta)u(5) + \beta\phi^{-1}(\mathbb{E}[\phi(z)]))$$

Here, the uncertainty about z is not resolved until the second period, so the expectation is taken inside the function $\phi^{-1}(\cdot)$. If we define $\tilde{\phi}(z) := \phi((1 - \beta)u(5) + \beta z)$, then the agent prefers early resolution of uncertainty if $V(P) \geq V(Q)$ or

$$\begin{aligned} \mathbb{E}[\tilde{\phi}(z)] &\geq \tilde{\phi}(\phi^{-1}(\mathbb{E}[\phi(z)])) \\ \tilde{\phi}^{-1}(\mathbb{E}[\tilde{\phi}(z)]) &\geq \phi^{-1}(\mathbb{E}[\phi(z)]) \end{aligned}$$

which holds if $\tilde{\phi}$ is less concave than ϕ .

We can now generalize how the curvature of ϕ captures preference for early or late resolution of uncertainty. Recall that the Arrow-Pratt measure of a function ϕ is given by

$$A_\phi(z) := -\frac{\phi''(z)}{\phi'(z)}.$$

Proposition 1. *The agent prefers early (late) resolution of uncertainty iff for all $y \in \mathcal{U}$, $A_{\phi_y}(z) \leq A_\phi(z)$ (resp. $A_{\phi_y}(z) \geq A_\phi(z)$) where $\phi_y(z) := \phi((1 - \beta)y + \beta z)$.*

Proof. See Strzalecki (2013) and Stanca (2023).⁵ □

To see the condition in Proposition 1 more explicitly, note that

$$A_{\phi_y}(z) = \beta A_\phi((1 - \beta)y + \beta z)$$

so the agent prefers early resolution of uncertainty iff

$$\beta A_\phi((1 - \beta)y + \beta z) \leq A_\phi(z)$$

for all $y \in \mathcal{U}$. The condition for preference for late resolution is the same but with the inequality reversed.

Preferences over the timing of resolution of uncertainty can play an important role in agents' decision making process in dynamic environments. Different choices are not only associated with different levels of uncertainty, but are also related to when that uncertainty is realized over time. For example, in Rust's bus engine replacement problem, replacing an engine entails current (sure) cost but makes the future uncertainty of engine failure resolved *earlier*. On the other hand, not replacing the engine leads to the future uncertainty of engine failure resolved *later*.

There is no conclusive experimental evidence on whether individuals have preference for early or late resolution of uncertainty. In fact, agents' preference about the timing of resolution of uncertainty could be highly context-dependent. For example, when uncertainty is over positive consumption levels, a preference for early resolution may be more reasonable. Indeed, in the macro-finance literature, the assumption of

⁵In Strzalecki (2013), this result is presented in the context of ambiguity.

preference for early resolution of uncertainty is often imposed. On the other hand, in cases when payoffs can be very negative, individuals may prefer to avoid early resolution of uncertainty. As a heuristic example, consider a pronouncement of a fatal disease from a doctor. Some patients may prefer to learn the news later because this information can induce a great deal of disappointment, fear and anxiety. These negative emotions, while may not have instrumental values, could still directly affect their utility.⁶

2.2 Parametric Special Cases

We now present several parametrizations of non-separable time preferences that have been used in the literature.

2.2.1 CRRA Epstein-Zin

By far the most common parametric form of Epstein-Zin preferences is when both u and ϕ have CRRA (or power) functional forms. Here, the utility is given by $u(c) = c^{1-\rho}$ and the time aggregator is given by

$$\phi(z) = \left(z^{\frac{1}{1-\rho}}\right)^{1-\alpha} = z^{\frac{1-\alpha}{1-\rho}}$$

where $\alpha \leq 1$ and $\rho \leq 1$ are parameters that characterize intertemporal and risk preferences respectively. Under these assumptions, the value function (2) yields the following form:

$$\mathbb{E}_c \left[\left((1-\beta) c^{1-\rho} + \beta \mathbb{E}_{v|c} [v]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1-\alpha}{1-\rho}} \right]. \quad (3)$$

To see how ρ and α characterize intertemporal and risk preferences, first consider the case where there is no risk so we can remove the expectations in (2). Redefining a new value function as $\tilde{v} = \phi^{-1}(v)$, we obtain from (2) that the new value function is given by

$$(1-\beta) u(c) + \beta \tilde{v} = (1-\beta) c^{1-\rho} + \beta \tilde{v}$$

⁶Note that in the definition of a preference for the timing of resolution of uncertainty, the information must be non instrumental, i.e., the agent cannot do anything given the information.

where intertemporal preferences are determined solely by ρ . Here, ρ measures the elasticity of substitution, defined as the percentage change of the growth rate of consumption (i.e. c_{t+1}/c_t) to the percentage change of the intertemporal marginal rate of substitution (MRS):⁷

$$\frac{d \log (c_{t+1}/c_t)}{d \log MRS_{t,t+1}} = \rho^{-1}. \quad (4)$$

For risk preferences, consider a one-time shock that resolves within a single-period followed by zero consumption forever. It straightforward to see that the Bernoulli utility in (3) reduces to $c^{1-\alpha}$, which has standard constant relative risk aversion (CRRA) parameter α . The fact that both intertemporal and risk preferences can be so neatly separated here (which is not possible with standard separable preferences) speaks to its popularity in the literature.

Proposition 1 reduces to the following in this parametrization.

Corollary 1. *An agent with CRRA Epstein-Zin preferences prefers early (late) resolution of uncertainty iff $\rho \leq \alpha$ (resp. $\rho \geq \alpha$) iff ϕ is concave (resp. convex).*

Proof. Since $\phi(z) = z^{\frac{1-\alpha}{1-\rho}}$,

$$A_\phi(z) = \left(1 - \frac{1-\alpha}{1-\rho}\right) z^{-1},$$

$$A_{\phi_y}(z) = \beta \left(1 - \frac{1-\alpha}{1-\rho}\right) ((1-\beta)y + \beta z)^{-1}.$$

Thus, $A_{\phi_y}(z) \leq A_\phi(z)$ iff $\frac{1-\alpha}{1-\rho} \leq 1$ or $\rho \leq \alpha$. The result follows from Proposition 1 and noting that ϕ is concave (convex) iff $\rho \leq \alpha$ (resp. $\rho \geq \alpha$). \square

2.2.2 CARA Epstein-Zin

An alternative parameterization is when u has the CARA (or exponential) functional form. Here, the utility is given by $u(c) = u_\rho(c)$ where

$$u_\rho(c) := \rho^{-1} (1 - e^{-\rho c}) \quad (5)$$

⁷Recall that MRS is defined as $u'(c_{t+1})/u'(c_t)$.

and the time aggregator is given by

$$\phi(z) = u_\alpha(u_\rho^{-1}(z)) = \alpha^{-1} \left(1 - (1 - \rho z)^{\frac{\alpha}{\rho}} \right)$$

for parameters $\alpha \geq 0$ and $\rho \geq 0$ that characterize intertemporal and risk preferences. Under these assumptions, the value function (2) yields the following form

$$\mathbb{E}_c \left[\alpha^{-1} \left(1 - \left((1 - \beta) e^{-\rho c} + \beta (1 - \mathbb{E}_{v|c}[\alpha v])^{\frac{\rho}{\alpha}} \right)^{\frac{\alpha}{\rho}} \right) \right]. \quad (6)$$

These preferences have appeared in Skiadas (2009) and are exactly the CARA analog of CRRA Epstein-Zin preferences. One advantage of this formulation is that it allows for negative payoffs, which CRRA cannot accommodate. For example, in Rust (1987), the agent needs to pay a large cost when choosing to replace the bus engine resulting in net negative revenue that period.

As in CRRA Epstein-Zin, the parameters ρ and α characterize intertemporal and risk preferences. First, in the absence of risk, we can follow the same reasoning as for CRRA Epstein-Zin above by redefining a new value function $\tilde{v} = \phi^{-1}(v)$. We now obtain from (2) that the new value function is given by

$$(1 - \beta) u(c) + \beta \tilde{v} = (1 - \beta) \rho^{-1} (1 - e^{-\rho c}) + \beta \tilde{v}$$

so intertemporal preferences are determined solely by ρ . Here, in contrast to CRRA Epstein-Zin, ρ measures the responsiveness of the *difference* in consumption between two periods (i.e. $c_{t+1} - c_t$) to the percentage change of MRS:

$$\frac{d(c_{t+1} - c_t)}{d \log MRS_{t,t+1}} = \rho^{-1}. \quad (7)$$

For risk preferences, as before, consider a one-time shock that resolves within a single-period followed by zero consumption forever. It straightforward to see that the Bernoulli utility in (6) has Arrow-Pratt coefficient

$$\frac{(1 - \beta) \alpha + \beta e^{\rho c} \rho}{(1 - \beta) + \beta e^{\rho c}}.$$

This is a convex combination between α and ρ where the weight on α is decreasing in

β and consumption c . Note that holding all else constant, risk aversion is increasing in α . Unlike CRRA Epstein-Zin preferences, while intertemporal preferences depend solely on ρ , risk preferences here depend on both α and ρ .

Proposition 1 reduces to the following in this parametrization.

Corollary 2. *An agent with CARA Epstein-Zin preferences prefers early (late) resolution of uncertainty iff $\rho \leq \alpha$ (resp. $\rho \geq \alpha$) iff ϕ is concave (resp. convex).*

Proof. Since $\phi(z) = \alpha^{-1} \left(1 - (1 - \rho z)^{\frac{\alpha}{\rho}} \right)$,

$$A_{\phi}(z) = \frac{\alpha - \rho}{1 - \rho z},$$

$$A_{\phi_y}(z) = \frac{\beta(\alpha - \rho)}{1 - \rho((1 - \beta)y + \beta z)}.$$

Thus, $A_{\phi_y}(z) \leq A_{\phi}(z)$ iff $\rho \leq \alpha$ (noting that $y \leq \rho^{-1}$ given the range of $u = u_{\rho}$). The result follows from Proposition 1 and noting that ϕ is concave (convex) iff $\rho \leq \alpha$ (resp. $\rho \geq \alpha$). \square

2.2.3 Separable Preferences

Finally, both CRRA and CARA Epstein-Zin preferences nest separable preferences as special cases. When $\alpha = \rho$ in either case, the time aggregator becomes linear, i.e. $\phi(z) = z$, and the value function becomes

$$(1 - \beta) \mathbb{E}_c[u(c)] + \beta \mathbb{E}_v[v] \tag{8}$$

where u is either of the CRRA or CARA functional form. Furthermore, when $\alpha = \rho \rightarrow 0$ under either, $u(c) = c$ so this reduces to the separable risk-neutral case. This is the specification commonly adopted in many empirical studies, including the engine replacement model in Rust (1987).

The fact that separable preferences, i.e. $\alpha = \rho$, is nested in our more general parametrizations allows us to statistically test whether the agent has separable preferences. While separability provides computational convenience, misspecifying the agent's preference as additively time separable when it is not can lead to systematic

biases in inferences of the structural primitives of the model. We see this for the case of CARA Epstein-Zin preferences in Section 5.

3 Dynamic Discrete Choice: General Setup and Theoretical Results

We now apply the general non-separable preferences defined in the previous section to a general dynamic discrete choice framework. We provide two key theoretical results. The first result pertains to the existence of a value function, while the second addresses its uniqueness.

3.1 General Setup

Denote the set of states by \mathcal{X} and sets of actions by \mathcal{D} . We assume \mathcal{D} is finite. The timing of our setup is as follows. At the beginning of period t , the state x_t realizes and then an idiosyncratic shock ε_t realizes. The agent then makes a choice $d_t \in \mathcal{D}$. Finally, the uncertainty about another variable Δ_t , which affects the current period consumption is realized. The agent consumes the realized consumption that period, which is given by

$$c(d_t, x_t, \Delta_t, \varepsilon_t) = \pi(d_t, x_t, \Delta_t) + \varepsilon_t(d_t), \quad (9)$$

where $\pi(d_t, x_t, \Delta_t)$ represents the deterministic consumption. We assume that $\varepsilon_t(d_t)$ is a random shock for alternative d_t . The vector of random shocks $\varepsilon_t = (\varepsilon_t(1), \dots, \varepsilon_t(|\mathcal{D}|))$ is drawn from a cumulative distribution function $G(\cdot)$.

We represent the agent's optimization problem using the Bellman's equation as follows.

$$V(x_t, \varepsilon_t) = \max_{d \in \mathcal{D}} \left\{ \mathbb{E}_{\Delta_t | d, x_t} \left[\phi \left((1 - \beta) u(c(d, x_t, \Delta_t, \varepsilon_t)) + \beta \phi^{-1}(\mathbb{E}_{x_{t+1}, \varepsilon_{t+1} | d, x_t, \Delta_t} [V(x_{t+1}, \varepsilon_{t+1})]) \right) \right] \right\}. \quad (10)$$

The first expectation is due to uncertainty in the current-period payoff shock Δ_t , and

the second expectation is due to uncertainty in the future state $(x_{t+1}, \varepsilon_{t+1})$. We follow the dynamic discrete choice literature to assume that (1) the random shock ε 's are independent over time and (2) conditional on the current state x_t and choice d_t , the future state x_{t+1} is independent of the unobserved state ε_t . Given these assumptions, the distribution of future state $(x_{t+1}, \varepsilon_{t+1})$ does not depend on ε_t . We therefore omit ε_t in the second expectation $\mathbb{E}_{x_{t+1}, \varepsilon_{t+1} | d_t, x_t, \Delta_t}$.

From the Bellman's Equation in (10), we define the ex-ante value function as

$$V(x_{t+1}) = \int V(x_{t+1}, \varepsilon_{t+1}) G(\varepsilon_{t+1}). \quad (11)$$

Note that $V(x_{t+1})$ represent the agent's expected value at the beginning of period $t + 1$ before the realization of the future shock ε_{t+1} . If the agent chooses d_t , the value she receives can be represented as

$$\begin{aligned} & v(d_t, x_t, \varepsilon_t) \\ &= \mathbb{E}_{\Delta_t | d_t, x_t} \left[\phi \left((1 - \beta) u(c(d_t, x_t, \Delta_t, \varepsilon_t)) + \beta \phi^{-1}(\mathbb{E}_{x_{t+1} | d_t, x_t, \Delta_t} [V(x_{t+1})]) \right) \right]. \end{aligned} \quad (12)$$

Note that $v(d_t, x_t, \varepsilon_t)$ is essentially the "choice-specific" value function defined in the dynamic discrete choice literature. The value functions in Equation (12) are non-separable in random shocks ε_t even if we assume that those shocks enter the current payoff additively as in Equation (9). In the special case when ϕ is linear and the agent is risk-neutral, ε_t is additively separable from other expressions in the value function and this reduces to the standard discrete choice setup.

The decision maker's optimal choice is defined as

$$d_t^* = \arg \max_{d \in \mathcal{D}} v(d, x_t, \varepsilon_t),$$

which leads to a Conditional Choice Probability (CCP) of choosing action d_t

$$p(d_t | x_t) = \int \mathbf{1} \left\{ d_t \in \arg \max_{d \in \mathcal{D}} v(d, x_t, \varepsilon_t) \right\} dG(\varepsilon_t). \quad (13)$$

With the non-separability of random shocks, the CCP in Equation (13) and the ex-ante value function in Equation (11) do not have closed-form solutions even if we

assume that ε 's follow extreme value distributions.

Combining Equations (11) and (12), we can define a value function iteration operator T in the space of the ex-ante value functions. To simplify notation, define

$$r(d_t, x_t, \Delta_t, \varepsilon_t, V) := \phi \left((1 - \beta) u(c(d_t, x_t, \Delta_t, \varepsilon_t)) + \beta \phi^{-1} \left(\mathbb{E}_{x_{t+1}|d_t, x_t, \Delta_t} V(x_{t+1}) \right) \right).$$

Then the operator T is defined as

$$T(V)(x_t) = \int \max_{d_t \in \mathcal{D}} \left\{ \int r(d_t, x_t, \Delta_t, \varepsilon_t, V) dF_{d_t, x_t}(\Delta_t) \right\} dG(\varepsilon_t), \quad (14)$$

where $F_{d_t, x_t}(\Delta_t)$ represents the cumulative distribution function of Δ_t given the current state x_t and the agent's choice d_t . The ex-ante continuation value V is obtained as a solution of the functional equation $T(V) = V$.

Remark 1. *When ϕ is a linear function, this model reduces to a standard dynamic discrete choice model with time-separable preferences and the choice-specific value function becomes (assuming $\phi(x) = x$ for simplicity)*

$$v(d_t, x_t, \varepsilon_t) = (1 - \beta) \mathbb{E}_{\Delta_t|d_t, x_t} [u(c(d_t, x_t, \Delta_t, \varepsilon_t))] + \beta \mathbb{E}_{x_{t+1}|d_t, x_t} [V(x_{t+1})]. \quad (15)$$

Note that in this case, the current-period payoff shocks Δ_t are completely separated across time.

Remark 2. *When there is no current-period uncertainty (i.e., Δ_t does not enter the utility), the choice-specific value function can be simplified to*

$$v(d_t, x_t, \varepsilon_t) = \phi \left((1 - \beta) u(c(d_t, x_t, \varepsilon_t)) + \beta \phi^{-1} \left(\mathbb{E}_{x_{t+1}|d_t, x_t} [V(x_{t+1})] \right) \right). \quad (16)$$

Note that this does not reduce to the standard dynamic discrete choice setup. The reason is that there are uncertainties with respect to the future state $(x_{t+1}, \varepsilon_{t+1})$, and expectations with respect to these uncertainties are taken in between the aggregators ϕ^{-1} and ϕ . Whenever preferences are non-separable so ϕ is non-linear, Equation (16) differs from standard dynamic discrete choice model even if there is no current-period uncertainty.

3.2 Existence: Value Function

In this section, we establish existence of a value function. We show that the operator T defined in Equation (14) always has a fixed point under some mild assumptions. Moreover, we demonstrate how to obtain the largest and smallest fixed points.

Let $\underline{\pi}$ and $\bar{\pi}$ denote the smallest and largest possible value of profit π . Define

$$\begin{aligned} v^* &:= \max \{ \mathbb{E} [\phi (u (\bar{\pi} + \hat{\varepsilon}))], \phi (\mathbb{E} [u (\bar{\pi} + \hat{\varepsilon}))] \}, \\ v_* &:= \min \{ \mathbb{E} [\phi (u (\underline{\pi} + \varepsilon))], \phi (\mathbb{E} [u (\underline{\pi} + \varepsilon)]) \}, \end{aligned}$$

where $\hat{\varepsilon} := \max_{d \in \mathcal{D}} \varepsilon_d$. These are bounds for the largest and smallest possible values of continuation values. Suppose v^* and v_* are both finite and let \mathcal{V} be the set of value functions bounded between some $\underline{v} \leq v_*$ and $\bar{v} \geq v^*$, that is, the set of V such that $\underline{v} \leq V(x) \leq \bar{v}$ for all $x \in \mathcal{X}$.

We first show that $T : \mathcal{V} \rightarrow \mathcal{V}$ is well-defined (see Appendix A.1). We then define a complete lattice (\mathcal{V}, \geq) , where \geq is the partial order on \mathcal{V} such that $V \geq V'$ if $V(x) \geq V'(x)$ for all $x \in \mathcal{X}$.⁸ It is easy to see that T is monotonic, that is, $V \geq V'$ implies $T(V) \geq T(V')$.

Lemma 1. *T is monotonic.*

Proof. Follows from the fact that u and ϕ are all increasing functions so r is also increasing in V . □

Since T is monotonic, Tarski's fixed point theorem (see Aliprantis and Border (2006)) guarantees that the set of fixed points of T is non-empty and a complete lattice. Moreover, if we start iterating from the smallest and largest possible values of V , then we obtain the respective smallest and largest fixed points of T . We summarize this as follows.

Theorem 1. *Suppose v^* and v_* are finite and let \mathcal{V} be the set of functions bounded by some $\underline{v} \leq v_*$ and $\bar{v} \geq v^*$. Then $T : \mathcal{V} \rightarrow \mathcal{V}$ has a fixed point. Moreover, $\lim_n T^n(\underline{v})$ and $\lim_n T^n(\bar{v})$ are its smallest and largest fixed points respectively.*

⁸A partially ordered set (\mathcal{V}, \geq) is a complete lattice if every subset has both a supremum and an infimum in \mathcal{V} according to the partial order \geq .

Proof. See Appendix A.1. □

Existing results assume separable time preferences. Our contribution is demonstrating that for non-separable time preferences, we can apply Tarski's theorem as long as the T mapping is monotonic (i.e. increasing) in the value function.⁹ This also provides a way to compute the smallest and largest value functions by iterating T . Although this does not guarantee a unique fixed point, we can easily check by calculating the smallest and largest fixed points and seeing if they coincide.

How about specific parametrizations of non-separable preferences? Theorem 1 implies the following for CRRA Epstein-Zin preferences.

Corollary 3. *Suppose the agent has CRRA Epstein-Zin preferences. If $\mathbb{E}[\varepsilon]$ is finite, then T has a fixed point.*

Proof. Since both u and ϕ are both power functions, they both yield positive values so v_* must be finite. Next, note that

$$\mathbb{E}[u(\bar{\pi} + \hat{\varepsilon})] = \mathbb{E}[(\bar{\pi} + \hat{\varepsilon})^{1-\rho}] \leq \left(\bar{\pi} + \mathbb{E}\left[\max_{d \in \mathcal{D}} \varepsilon_d\right] \right)^{1-\rho}$$

where the inequality follows from the fact that $\rho \leq 1$. Since $\max_{d \in \mathcal{D}} \varepsilon_d \leq \sum_{d \in \mathcal{D}} \varepsilon_d$, it has finite expectation. The reasoning for $\mathbb{E}[\phi(u(\bar{\pi} + \hat{\varepsilon}))] = \mathbb{E}[(\bar{\pi} + \hat{\varepsilon})^{1-\alpha}]$ is symmetric, so v^* is finite and the result follows from Theorem 1. □

With CARA Epstein-Zin preferences, Theorem 1 implies the following.

Corollary 4. *Suppose the agent has CARA Epstein-Zin preferences. If $\mathbb{E}[e^{-t\varepsilon}]$ is finite for all $t \in \mathbb{R}$, then T has a fixed point.*

Proof. Since $\phi = u_\alpha \circ u_\rho^{-1}$ where u_α is CARA, ϕ is bounded above by α^{-1} . Thus, $v^* = \alpha^{-1}$ is finite. Next, note that

$$\mathbb{E}[u(\underline{\pi} + \varepsilon)] = \mathbb{E}[u_\rho(\underline{\pi} + \varepsilon)] = \rho^{-1} (1 - e^{-\rho \underline{\pi}} \mathbb{E}[e^{-\rho \varepsilon}])$$

which is finite given the assumption. The reasoning for $\mathbb{E}[\phi(u(\underline{\pi} + \varepsilon))] = \mathbb{E}[u_\alpha(\underline{\pi} + \varepsilon)]$ is symmetric, so v_* is finite and the result follows from Theorem 1. □

⁹Jia (2008) also uses Tarski's fixed point theorem but for a different problem (entry game). She also assumes separable time preferences.

The condition in Corollary 4 is weak; it is equivalent to saying that the moment generating function for $-\varepsilon$ always exists. This is true if ε is Normal or Extreme-Value distributed, which we note below.

Corollary 5. *Suppose the agent has CARA Epstein-Zin preferences and ε either has a Normal or Extreme Value Type I distribution. Then starting with the initial point $v^* = \alpha^{-1}$, $\lim_n T^n(v^*)$ converges to its largest fixed point.*

Proof. Note that both the Normal distribution and the Extreme Value Type I (Gumbel) distribution have well-defined (finite) moment generating functions. It is easy to see that the same holds for $-\varepsilon$. □

3.3 Uniqueness: Contraction Mapping

In this section, we present conditions under which T constitutes a contraction, ensuring the uniqueness of the value function. First, recall from Proposition 1 that $\phi_y(z) = \phi((1 - \beta)y + \beta z)$ and define

$$\psi_y(z) := \phi_y(\phi^{-1}(z)).$$

Recall that $\underline{\pi}$ and $\bar{\pi}$ are the smallest and largest possible value of profit π . Also recall that \underline{v} and \bar{v} are the bounds of the value function (which can be infinite). We now present our second theorem.

Theorem 2. *T is a contraction mapping if*

$$M := \mathbb{E} \left[\max_{d \in D} \sup_{\pi \in [\underline{\pi}, \bar{\pi}], z \in [\underline{v}, \bar{v}]} \psi'_{u(\pi + \varepsilon_d)}(z) \right] < 1 \tag{17}$$

Proof. See Appendix A.2. □

Theorem 2 implies that when the bound $M < 1$, the ex-ante value function $V(\cdot)$ is a unique fixed point of the contraction mapping T as defined in equation (14). For intuition, note that $\psi_y(z)$ captures the value given current period utility y and future continuation value z . For instance, if $\phi(z) = z$ as in the standard separable case,

then

$$\psi_y(z) = (1 - \beta)y + \beta z.$$

In this case, $\psi_y(\cdot)$ is linear so $\psi'_y(z) = \beta$. This immediately yields the standard result that coincides with Rust's condition for a contraction mapping.

Corollary 6. *Suppose the agent has separable preferences. If $\beta < 1$, then T is a contraction mapping.*

In general, when the agent does not have separable preferences, the bounds on $\psi'_y(z)$ would depend on the curvature of ϕ . More explicitly, we have

$$\psi'_y(z) = \frac{\phi'_y(\tilde{z})}{\phi'(\tilde{z})}$$

where $\tilde{z} = \phi^{-1}(z)$. We thus have

$$\psi''_y(z) = \frac{\phi''_y(\tilde{z})}{\phi'(\tilde{z})^2} (A_\phi(\tilde{z}) - A_{\phi_y}(\tilde{z})).$$

From Proposition 1, this implies that $\psi'_y(\cdot)$ is increasing (decreasing) if the agent prefers early (resp. late) resolution of uncertainty. This implies the following.

Corollary 7. *Suppose the agent prefers early (late) resolution of uncertainty. If $\mathbb{E} \left[\sup_{d \in D, \pi \in [\underline{\pi}, \bar{\pi}]} \psi'_{u(\pi + \varepsilon_d)}(\hat{v}) \right] < 1$ where $\hat{v} = \bar{v}$ (resp. $\hat{v} = \underline{v}$), then T is a contraction mapping.*

Proof. Follows from Proposition 1 and Theorem 2. □

Note that this bound can be further simplified if we know whether ϕ is convex or concave. For instance, if ϕ is convex, then $\psi'_y(z)$ is increasing in y so we can use $\pi = \bar{\pi}$ as the bound when evaluating M in Equation (17). The case for concave ϕ is symmetric.

How about specific parametrizations of non-separable preferences? Theorem 2 implies the following for CRRA Epstein Zin preferences.

Corollary 8. *Suppose the agent has CRRA Epstein-Zin preferences. If $\rho \leq \alpha$ and $\beta^{\frac{1-\alpha}{1-\rho}} < 1$, then T is a contraction mapping.*

Proof. Since $\phi(z) = z^{\frac{1-\alpha}{1-\rho}}$,

$$\begin{aligned}\psi'_y(z) &= \frac{\beta \left((1-\beta)y + \beta z^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1-\alpha}{1-\rho}-1}}{\left(z^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1-\alpha}{1-\rho}-1}} \\ &= \beta \left((1-\beta)yz^{\frac{1-\alpha}{1-\rho}} + \beta \right)^{\frac{1-\alpha}{1-\rho}-1}.\end{aligned}$$

If $\rho \leq \alpha$, then $\frac{1-\alpha}{1-\rho} - 1 \leq 0$. Since $y \geq 0$, this means

$$\psi'_y(z) \leq \beta(\beta)^{\frac{1-\alpha}{1-\rho}-1} = \beta^{\frac{1-\alpha}{1-\rho}}.$$

Thus $\beta^{\frac{1-\alpha}{1-\rho}} < 1$ ensures a contraction mapping. \square

Note that in the other $\rho > \alpha$ case, we can still use Theorem 2 to obtain a sufficient condition for a contraction mapping but the bound M in Equation (17) would depend on the error distribution. The reason we are able to obtain a simple sufficient condition for when the agent prefers early resolution of uncertainty ($\rho \leq \alpha$) is that the value function under CRRA Epstein-Zin is bounded below (since $u(c)$ has to be non-negative). The fact that values are bounded below and ϕ is concave given $\rho \leq \alpha$ (see Corollary 1), allows us to create bounds to ensure a contraction mapping.

With CARA Epstein-Zin preferences, Theorem 2 implies the following.

Corollary 9. *Suppose the agent has CARA Epstein-Zin preferences. If $\rho \geq \alpha$ and $\beta^{\frac{\alpha}{\rho}} < 1$, then T is a contraction mapping.*

Proof. Since $\phi = u_\alpha \circ u_\rho^{-1}$,

$$\begin{aligned}\psi'_y(z) &= \frac{\beta \left((1-\beta)(1-\rho y) + \beta(1-\alpha z)^{\frac{\rho}{\alpha}} \right)^{\frac{\alpha}{\rho}-1}}{\left((1-\alpha z)^{\frac{\rho}{\alpha}} \right)^{\frac{\alpha}{\rho}-1}} \\ &= \beta \left((1-\beta) \frac{1-\rho y}{(1-\alpha z)^{\frac{\rho}{\alpha}}} + \beta \right)^{\frac{\alpha}{\rho}-1}.\end{aligned}$$

Since $\rho \geq \alpha$, $\frac{\alpha}{\rho} - 1 \leq 0$ so

$$\psi'_y(z) \leq \beta \left(\beta^{\frac{\alpha}{\rho}-1} \right) = \beta^{\frac{\alpha}{\rho}}.$$

Thus $\beta^{\frac{\alpha}{\rho}} < 1$ ensures a contraction mapping. □

In other words, for an agent with CARA Epstein-Zin preferences and prefers late resolution of uncertainty (i.e. $\rho \geq \alpha$), T is a contraction mapping as long as β is sufficiently small (i.e. $\beta^{\frac{\alpha}{\rho}} < 1$). The fixed point would be unique. As with the CRRA Epstein-Zin case, we can still use Theorem 2 for the case when $\rho < \alpha$ but the bound M would again depend on the error distribution. Note that in contrast to CRRA Epstein-Zin, values are bounded above under CARA Epstein-Zin, so it is when the agent prefers late resolution of uncertainty ($\rho \geq \alpha$) where we can obtain simple sufficient conditions for a contraction mapping.

4 Estimation

We propose to use an estimation approach similar to Rust (1987)'s Nested Fixed Point algorithm to estimate the structural parameters in the dynamic model described in Section 3.1.¹⁰ In our model, the choice-specific value function is not additively separable in the idiosyncratic shock ε due to the nonlinear nature of the non-separable time preferences. Therefore, the well-known multinomial logit formula for the conditional choice probabilities and the McFadden's social surplus function are no longer applicable. We describe in this section how we obtain the conditional choice probabilities using simulation methods. Our nested fixed point algorithm contains an inner loop that solves a DDC problem for a given set of parameter values and an outer loop that maximizes the log-likelihood.

We focus on the estimation of two parameterizations of the model introduced in Section 2.2. Let $\theta = (\alpha, \rho, \theta_\pi)$ denote the vector of parameters to be estimated, where α and ρ are parameters in the utility function u and aggregator function ϕ in either the CRRA or CARA Epstein-Zin case. These two parameters capture the agent's risk

¹⁰Since Rust (1987), several estimation methods for standard DDC models became available, such as the CCP approach of Hotz and Miller (1993), the nested pseudo-likelihood approach of Aguirregabiria and Mira (2002) and their extensions to various settings. See Aguirregabiria and Mira (2010) for a survey of methodologies. Recently, Igami (2017) and Igami and Uetake (2020) extend the Nested Fixed Point algorithm to dynamic oligopoly games with sequential or stochastically alternating moves.

and intertemporal preferences. θ_π represents the parameters in the payoff function $\pi(d_t, x_t, \Delta_t)$. Throughout the estimation, we assume that the value of the discount factor β is known to the econometrician.¹¹ Let x_{it} and d_{it} denote the observable state and choice for agent $i = 1, 2, \dots, N$ at period $t = 1, 2, \dots, T$. Δ_{it} represents the realization of current period uncertainty. We derive the likelihood of the choice pattern and state transition process observed in the data:

$$LL(\theta) = \sum_i^N \log \left(Pr(d_{i1}|x_{i1}; \theta) \prod_{t=2}^T Pr(d_{it}|x_{it}; \theta) Pr(x_{it}|d_{it-1}, x_{it-1}, \Delta_{it-1}) \right). \quad (18)$$

Note that $Pr(d_{it}|x_{it}; \theta)$ represents the choice probability conditional on the observable state x_{it} and $Pr(x_{it}|d_{it-1}, x_{it-1}, \Delta_{it-1})$ represents the state transition probability, which can be estimated separately outside of the structural model.

For a given parameter θ , we solve the ex-ante value function $V(x_t)$ through an iterative procedure outlined as follows.

- (1) Initiate the iteration process with a chosen $V^0(x_t)$. Denote the value function at the r -th iteration as $V^r(x_t)$.
- (2) Given $V^r(x_t)$, compute the choice-specific value function $v(d_t, x_t, \varepsilon^s)$ for a random draw ε^s from G :

$$\mathbb{E}_{\Delta_t|d_t, x_t} \left[\phi \left((1 - \beta) u(c(d_t, x_t, \Delta_t, \varepsilon^s(d_t))) + \beta \phi^{-1}(\mathbb{E}_{x_{t+1}|d_t, x_t, \Delta_t}[V^r(x_{t+1})]) \right) \right].$$

- (3) Update $V^{r+1}(x_t)$ by taking the average over S simulation draws:

$$V^{r+1}(x_t) \approx \frac{1}{S} \sum_{s=1}^S \left(\max_{d \in \mathcal{D}} \left\{ v(d, x_t, \varepsilon^s) \right\} \right).$$

- (4) Repeat Steps (2)–(3) until convergence.

The contraction property proved in Theorem 2 guarantees the uniqueness of the value function. For our specific parameterizations of non-separable preferences, Corol-

¹¹The identification and estimation of the discount factor β in standard time-separable models is known to be difficult and often requires exclusion restrictions or the availability of terminating actions. See discussions in Fang and Wang (2015), Bajari et al. (2016), Komarova et al. (2018), Abbring and Daljord (2020), Schneider (2021).

laries 8 and 9 provide sufficient conditions for a contraction mapping when agents have certain preference for the temporal resolution of uncertainty under CRRA and CARA Epstein-Zin preferences, respectively. When the bound M in Theorem 2 is difficult to compute, we can check empirically whether the largest and the smallest fixed points coincide with each other. Theorem 1 guarantees that we obtain the largest (smallest) fixed point of T through the iterative procedure above if we start iterating from the largest (smallest) possible values of V . In all of our simulation exercises and empirical application, the fixed points are all unique.

At the point of convergence, the conditional choice probabilities $Pr(d_{it}|x_{it}; \theta)$ can also be computed numerically.

$$p(d_t|x_t) \approx \frac{1}{J} \sum_{j=1}^J \mathbf{1} \left\{ d_t = \arg \max_{d \in \mathcal{D}} v(d, x_t, \varepsilon^j) \right\}.$$

In the outer loop of the algorithm, we search for the value of $\theta \in \Theta$ to maximize the log-likelihood function in Equation (18), i.e., the estimator for the structural parameter is $\hat{\theta} = \max_{\theta \in \Theta} LL(\theta)$.

Before moving on to the empirical example, we briefly discuss identification of the model parameters. With the introduction of non-separable time preferences, our model is much richer than the standard additively-separable model. Most importantly, the value functions in our model are non-separable in the random shocks ε_t . Existing identification results for time separable models (e.g., Magnac and Thesmar, 2002; Arcidiacono and Miller, 2020) rely on the inversion from CCPs to value functions given additively separable random shocks (Hotz and Miller, 1993), and thus do not directly apply to our model.

Our empirical model imposes specific parametric assumptions on u and ϕ . These parameters affect the agent's risk and intertemporal preferences in non-trivial ways so deriving generic identification results is not a straightforward exercise. All model parameters (i.e., θ) jointly determine the conditional choice probabilities $p(d_t|x_t)$ in Equation (13). Matching the model-implied CCPs with their empirical counterparts for various x_t , we essentially characterize θ as the solution to a system of nonlinear

equations. Given the high nonlinearity of the dynamic model, one possibility is to at least locally identify θ given certain rank conditions (see Lewbel (2019) for more detailed discussions and examples on local identification). Since the main purpose of our paper is to provide a general framework of dynamic discrete choice models that allow for non-separable time preferences, we leave a thorough investigation of identification for future research.


5 Empirical Application: Optimal Engine Replacement Revisited

We apply our model to the dataset of bus engine replacement decisions originally studied in Rust (1987). The manager (Harold Zuercher) makes a dynamic choice for each bus engine by trading off between an immediate lump sum cost of replacing it and higher maintenance costs for keeping it at the beginning of each period. We extend the Rust model to allow for non-separable time preferences, agents being risk-averse, and earning revenues from operating the bus.

5.1 Setup

The observable state variable x_t represents the accumulated mileage at the beginning of period t . Let ε_t denote the unobserved payoff shocks and Δ_t denote the incremental mileage realized within period t . The agent decides whether or not to replace the engine upon observing (x_t, ε_t) . Let $d_t = 1$ represent the case where the engine is replaced, and $d_t = 0$ otherwise. The timeline of the model is shown in Figure 2.

Figure 2: Timeline of the model



We first specify the agent's payoff function

$$\pi(d_t, x_t, \Delta_t) = \begin{cases} \theta_d \Delta_t - RC & \text{if } d_t = 1 \\ \theta_d \Delta_t - \theta_x x_t & \text{if } d_t = 0 \end{cases}. \quad (19)$$

In Equation (19), RC represents the replacement cost and $\theta_d \Delta_t$ represents the revenue collected by the agent within a month for the incremental mileage Δ_t . When the bus engine is not replaced, the agent has to pay a maintenance cost proportional to the accumulated mileage, i.e., $\theta_x x_t$. The realized consumption agent has in period t is given by

$$c(d_t, x_t, \Delta_t, \varepsilon_t) = \pi(d_t, x_t, \Delta_t) + \sigma \varepsilon_t(d_t), \quad (20)$$

where σ is the standard deviation of the random shock and $\varepsilon_t(d_t)$ is the shock that the agent receives for alternative $d_t \in \{0, 1\}$. Note that $\varepsilon_t = (\varepsilon_t(1), \varepsilon_t(0))$ and it is randomly drawn from a cumulative distribution function $G(\cdot)$. Note that in this empirical application, the payoff agent receives might be negative. For example, if the agent chooses to replace the engine, he needs to pay an immediate replacement cost. We adopt the CARA Epstein-Zin preferences described in Section 2.2.2 as it allows for negative payoffs.

Applying the CARA parameterization to the choice-specific value function defined in Equation (16), we have

$$v(d_t, x_t, \varepsilon_t) = \frac{1}{\alpha} \left[1 - \mathbb{E}_{\Delta_t | d_t, x_t} \left\{ (1 - \beta) \exp \left(-\rho(\pi(d_t, x_t, \Delta_t) + \sigma \varepsilon_t(d_t)) \right) + \beta \left(1 - \alpha V(x_{t+1}) \right)^{\frac{\rho}{\alpha}} \right\}^{\frac{\alpha}{\rho}} \right],$$

where $V(x_{t+1})$ is the ex-ante value function given by

$$V(x_{t+1}) = \int \left(\max \left\{ v(0, x_{t+1}, \varepsilon), v(1, x_{t+1}, \varepsilon) \right\} \right) dG(\varepsilon),$$

and the next period accumulated mileage x_{t+1} is updated following:

$$x_{t+1} = (1 - d_t)x_t + \Delta_t.$$

Intuitively, if the agent chooses to replace the engine, i.e., $d_t = 1$, the current accumulated mileage is reset to zero, so the future mileage equals to the incremental mileage Δ_t . If the agent does not replace the engine, the next-period accumulated mileage is increased by Δ_t . Note that in this engine replacement example, the future state variable x_{t+1} is uniquely determined once d_t , x_t and Δ_t are realized. In other words, there is no uncertainty with respect to x_{t+1} conditional on (d_t, x_t, Δ_t) . We therefore omit the expectation $\mathbb{E}_{x_{t+1}|d_t, x_t, \Delta_t}$ in front of $V(x_{t+1})$. The probability that an engine is replaced conditional on the observed state x_t is

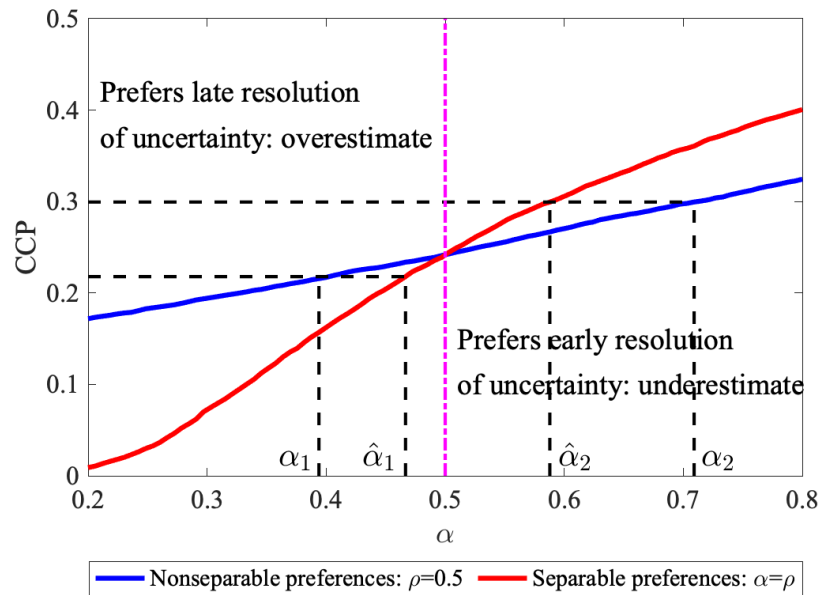
$$p(d_t = 1|x_t) = \int \mathbf{1}\left(v(1, x_t, \varepsilon_t) > v(0, x_t, \varepsilon_t)\right) dG(\varepsilon_t).$$

Based on the empirical model described above, we conduct a numerical exercise to compare the conditional choice probabilities under separable and non-separable preferences. In Figure 3, the red and blue curve represent the probabilities of replacing the engine when the agent has separable and non-separable time preferences, respectively. Throughout the exercise, we fix the intertemporal preference parameter ρ to be 0.5 and change the value of risk preference parameter α . Note that when $\alpha = \rho = 0.5$, the two models are equivalent, resulting in the same conditional choice probabilities (i.e., the red and blue curves intersect at $\alpha = 0.5$). When the agent prefers late resolution of uncertainty (i.e., $\rho > \alpha$), if we misspecify the agent's preference to be time-separable and use the observed CCP to estimate the agent's risk preference, we would over-estimate α . Conversely, we would under-estimate α if the agent prefers early resolution of uncertainty (i.e., $\rho < \alpha$). This exercise highlights the consequences of misspecifying the agent's preference to be time-separable when it is in fact not. Importantly, the direction of the bias depends on whether the agent prefers early or late resolution of uncertainty.

5.2 Estimation Results

In the estimation, we discretize mileage into 130 intervals of length 3000 (miles) and assume that ε 's are drawn from a standard normal distribution. To identify the

Figure 3: Misspecifying nonseparable preferences with separable preferences: biased estimates of the risk preference parameter α



Note: In this numerical exercise, we set $RC = 3$, $\theta_d = 3$, $\theta_x = 0.5$, $\sigma = 2$, $\beta = 0.9$, and $\rho = 0.5$. α takes values from $[0.2, 0.8]$. For simplicity, we allow the incremental and accumulated mileages to take 3 values: 0, 1, and 2. The state transition probabilities are specified as: $Pr(\Delta_t = 0|x_t = 0) = 0$, $Pr(\Delta_t = 1|x_t = 0) = 0.5$, $Pr(\Delta_t = 2|x_t = 0) = 0.5$; $Pr(\Delta_t = 0|x_t = 1) = 0.2$, $Pr(\Delta_t = 1|x_t = 1) = 0.6$, $Pr(\Delta_t = 2|x_t = 1) = 0.2$; $Pr(\Delta_t = 0|x_t = 2) = 0.6$, $Pr(\Delta_t = 1|x_t = 2) = 0.4$, $Pr(\Delta_t = 2|x_t = 2) = 0$.

scale of the unobserved random shock, we fix the RC to be 8 based on the cost of engine replacement reported by Harold Zurcher (see Rust (1987) Table III). We set the discount factor $\beta = 0.9$ and use 2,500 simulation draws when approximating value functions in the iteration process. To summarize, the structural primitives to be estimated include: parameters for risk-aversion and intertemporal preferences, α and ρ ; coefficient for revenue and maintenance cost, θ_d and θ_x ; and standard deviation of the idiosyncratic shock, σ .

In Table 1, we present the estimation results for four model specifications with standard errors provided in the parenthesis. We first estimate a model with non-separable time preferences, where we impose no restrictions on the values of α and ρ . Second, we restrict that $\alpha = \rho$, which leads to a model with separable time preferences and risk aversion. Next, we consider adding revenue to the original Rust model. Finally, for comparison purposes we also estimate the original Rust model¹².

Note that in the Rust model with revenue, the choice-specific value function is define as

$$v(d_t, x_t, \varepsilon_t) = \mathbb{E}_{\Delta_t|d_t, x_t} \left[(1 - \beta) \left(\pi(d_t, x_t, \Delta_t) + \sigma \varepsilon_t(d_t) \right) + \beta V(x_{t+1}) \right].$$

This is nested in the non-separable case where $\alpha = \rho \rightarrow 0$. The original Rust model further removes the revenue from the payoff function π (or in other words, restricting $\theta_d = 0$), so that the choice-specific value function is reduced to:

$$v(d_t, x_t, \varepsilon_t) = (1 - \beta) \left(\pi(d_t, x_t) + \sigma \varepsilon_t(d_t) \right) + \beta \mathbb{E}_{\Delta_t|d_t, x_t} [V(x_{t+1})].$$

From the first column of Table 1, we find that $\rho > \alpha$, suggesting that the agent is likely to prefer late resolution of uncertainty. Comparing the log-likelihood between Columns (1) and (2), we are able to reject the hypothesis that the agent has a separable preference at a 10% significant level (the likelihood ratio test statistic = 2.747, p -value = 0.097). When allowing for non-separable time preferences, our estimate for α is smaller. This is consistent with our simulation exercises: when agents prefer late

¹²In Rust (1987), the distribution of ε is assumed to be Type I extreme values, mainly for computational convenience. We use normal distribution for the shocks ε .

resolution of uncertainty, misspecifying the agent’s preference to be time-separable results in an overestimate of the risk preference parameter α .

Our estimates for the payoff parameters θ_d and θ_x are also significantly different when allowing for non-separable preferences. In particular, we underestimate the maintenance cost parameter θ_x , while overestimate the revenue parameter θ_d under time-separable preferences. This result highlights that misspecifying agents’ preferences can lead to biased inferences about structural primitives, and potentially misleading policy recommendations. Comparing Columns (3)-(4), we find that adding revenue to the original rust model has a minimal impact.

Table 1: Comparing estimation results for different models

	(1)	(2)	(3)	(4)
	Nonseparable	Separable	Rust - rev	Rust-orig
θ_d	0.0526 (0.0073)	0.1019 (0.0829)	0.0001 (0.0351)	
θ_x	0.1077 (0.0122)	0.0329 (0.0022)	0.0208 (0.0009)	0.0208 (0.0014)
σ	1.6070 (0.0491)	1.5436 (0.0623)	1.4883 (0.0542)	1.4770 (0.0566)
α	0.1023 (0.5087)	0.1457 (0.0095)		
ρ	0.5555 (0.0200)			
LL	-299.4404	-300.8139	-301.4402	-301.6273

The parameters that capture agents’ risk and time preferences (α and ρ) are important for predicting choices under uncertainty in dynamic settings and for making policy recommendations. To better interpret the estimated values of α and ρ , we compute the certainty equivalent given different model specifications. In the engine replacement model, the agent faces uncertainty about the incremental mileage in each period, which not only affects the current-period revenue, but also affects future maintenance costs. We compute the monetary payoff that makes the agent indifferent when the uncertainty about incremental mileage in each period is removed. Specifically, we assume that in a counterfactual scenario there is a subsidy program that helps the agent to smooth revenue across periods. The agent receives C each period no matter what the incremental mileage is. This is similar to the “capacity payment”

often seen for power plants in the electricity industry and pipelines in the natural gas industry.¹³

We compute the certainty equivalent using estimates from both separable and non-separable models. We find that agents with separable preferences are willing to accept a sure payment of \$120 each month to stay indifferent. However, for agents with non-separable time preferences, the calibrated subsidy payment is \$61.6, almost a half of what an agent with separable time preferences would require. This result is intuitive, because when the agent prefers late resolution of uncertainty, the original setting where there is uncertainty in revenue is less unfavorable. This exercise confirms that misspecifying agents' preferences as time-separable when it is in fact not can give rise to misleading policy implications.

6 Conclusion

In this paper, we propose an empirical model of dynamic discrete choice with Epstein-Zin preferences, generalizing the well-known bus engine replacement model studied in Rust (1987). Based on the Bellman equation with Epstein-Zin preferences, we prove the existence of the value function and provide conditions under which it is the unique fixed point of a contraction mapping. We propose a simulated nested fixed point algorithm for model estimation based on the theoretical results. Since Epstein-Zin preferences include separable expected utility as a special case, our framework allows us to test whether an agent has separable time preferences or not.

We apply the framework to the bus engine replacement data and show that we can reject the null hypothesis that Harold has time-separable preferences. Based on our estimates, the agent is likely to prefer late resolution of uncertainty. Misspecifying the agent's preference to be time separable when it is in fact not leads to a biased inference of agents' risk and intertemporal preferences as well as other key structural parameters. This has non-trivial consequences for making policy evaluations.

¹³Regulators have adopted capacity payment mechanisms to encourage capacity investment and promote system reliability. See details of capacity payment mechanisms in electricity markets in Pfeifenberger et al. (2009).

The tools we develop in this paper can be easily applied to analyze many empirical dynamic discrete choice models in industrial organization, environmental, health, and labor economics while allowing for more general non-separable time preferences. For example, a recent empirical literature studies the impact of uncertainty on economic outcomes such as oil drilling, firm investment, technology adoption, and exit decisions (see Kellogg, 2014; Dorsey, 2019; Handley and Li, 2020; Gowrisankaran et al., 2012). In these applications, preferences for the temporal resolution of uncertainty would naturally play an important role in firms' dynamic responses to policy changes, such as resolving policy uncertainty early. Allowing for more general non-separable time preferences could lead to different welfare implications.

A Proofs

To reduce notational burden, we drop the subscript t in the following proofs.

A.1 Proof of Theorem 1

Fix some $\underline{v} \leq v_*$ and $\bar{v} \geq v^*$. We first show that for any $v \in \mathcal{V}$, if $v \leq \bar{v}$, then $T(v) \leq \bar{v}$. Note that

$$\begin{aligned} T(v)(x) &\leq \int \max_{d \in D} \phi \left((1 - \beta) u(\bar{\pi} + \varepsilon_d) + \beta \phi^{-1}(\mathbb{E}[\bar{v}]) \right) dG(\varepsilon) \\ &= \int \phi \left((1 - \beta) u \left(\bar{\pi} + \max_{d \in D} \varepsilon_d \right) + \beta \phi^{-1}(\bar{v}) \right) dG(\varepsilon) \\ &= \mathbb{E} \left[\phi \left((1 - \beta) u(\bar{\pi} + \hat{\varepsilon}) + \beta \phi^{-1}(\bar{v}) \right) \right] \end{aligned}$$

First, if ϕ is convex, then

$$\begin{aligned} T(v)(x) &\leq \mathbb{E} \left[(1 - \beta) \phi(u(\bar{\pi} + \hat{\varepsilon})) + \beta \bar{v} \right] \\ &\leq (1 - \beta) \mathbb{E} [\phi(u(\bar{\pi} + \hat{\varepsilon}))] + \beta \bar{v} \\ &\leq (1 - \beta) v^* + \beta \bar{v} \leq \bar{v} \end{aligned}$$

as desired. Now, if ϕ is concave, then

$$\begin{aligned} T(v)(x) &\leq \phi \left(\mathbb{E} \left[(1 - \beta) u(\bar{\pi} + \hat{\varepsilon}) + \beta \phi^{-1}(\bar{v}) \right] \right) \\ \phi^{-1}(T(v)(x)) &\leq (1 - \beta) \mathbb{E} [u(\bar{\pi} + \hat{\varepsilon})] + \beta \phi^{-1}(\bar{v}) \\ &\leq (1 - \beta) \phi^{-1}(v^*) + \beta \phi^{-1}(\bar{v}) \\ &\leq \phi^{-1}(\bar{v}) \end{aligned}$$

as desired.

For the case of $v \geq \underline{v}$ implies $T(v) \geq \underline{v}$, note that if $v \geq \underline{v}$, then

$$\begin{aligned} T(v)(x) &\geq \int \max_{d \in D} \phi \left((1 - \beta) u(\underline{\pi} + \varepsilon_d) + \beta \phi^{-1}(\mathbb{E}[\underline{v}]) \right) dG(\varepsilon) \\ &\geq \mathbb{E} \left[\phi \left((1 - \beta) u(\underline{\pi} + \varepsilon) + \beta \phi^{-1}(\underline{v}) \right) \right] \end{aligned}$$

where the second inequality follows from the fact that ε_d are identically distributed.

The rest of the proof follows symmetrically as the argument above.

Thus, $T : \mathcal{V} \rightarrow \mathcal{V}$ is well-defined. Let \geq be the partial order on \mathcal{V} where $v \geq \tilde{v}$ if $v(x) \geq \tilde{v}(x)$ for all $x \in \mathcal{X}$. Note that for any subset $\mathcal{U} \subset \mathcal{V}$, we can define

$$v_{\wedge}(x) := \inf_{v \in \mathcal{U}} v(x)$$

which is the greatest lower bound of \mathcal{U} . Symmetrically, we can define its least upper bound so (\mathcal{V}, \geq) is a complete lattice. Moreover, it is straightforward to see that T is monotonic, that is, $v \geq \tilde{v}$ implies $T(v) \geq T(\tilde{v})$. Thus, by Tarski's fixed point theorem, the set of fixed points of T is non-empty and also a complete lattice.

Finally, define $v_0 = \underline{v}$ and $v_n = T(v_{n-1})$. Since $\underline{v} \leq v$ for all $v \in \mathcal{V}$, $v_0 \leq v_1$ so by induction, $v_n \leq v_{n+1}$ for all n . Since v_n is an increasing sequence, it converges to some v^* . Since T is continuous, we have

$$v^* = T(v^*) = \lim_n T^n(\underline{v})$$

To see why v^* is the smallest fixed point, suppose there is some other fixed point v^{**} . Since $v_0 \leq v^{**}$, we have

$$T^n(v_0) \leq T^n(v^{**}) = v^{**}$$

so taking limits, we obtain $v^* \leq v^{**}$. The case for \bar{v} is symmetric.

A.2 Proof of Theorem 2

Recall that

$$\psi_y(z) = \phi_y(\phi^{-1}(z))$$

so

$$T(v)(x) = \int \max_{d \in D} \mathbb{E}_{x,d} [\psi_y(\mathbb{E}_{d,x,\Delta}[v])] dG(\varepsilon)$$

where $y = u(\pi_{x,d,\Delta} + \varepsilon_d)$ depends on the realizations of x, d, Δ . We thus have

$$\begin{aligned} & T(v)(x) - T(\hat{v})(x) \\ &= \int \left(\max_{d \in D} \mathbb{E}_{x,d} [\psi_y(\mathbb{E}_{d,x,\Delta}[v])] - \max_{d \in D} \mathbb{E}_{x,d} [\psi_y(\mathbb{E}_{d,x,\Delta}[\hat{v}])] \right) dG(\varepsilon) \\ &\leq \int \max_{d \in D} \mathbb{E}_{x,d} [\psi_y(\mathbb{E}_{d,x,\Delta}[v]) - \psi_y(\mathbb{E}_{d,x,\Delta}[\hat{v}])] dG(\varepsilon) \end{aligned}$$

By the mean value theorem,

$$\psi_y(z_1) - \psi_y(z_2) = \psi'_y(z)(z_1 - z_2)$$

for some z strictly between z_1 and z_2 . Thus,

$$\begin{aligned} & T(v)(x) - T(\hat{v})(x) \\ &\leq \int \max_{d \in D} \mathbb{E}_{x,d} \left[\sup_{\pi \in [\underline{\pi}, \bar{\pi}], z \in [\underline{v}, \bar{v}]} \psi'_{u(\pi + \varepsilon_d)}(z) (\mathbb{E}_{d,x,\Delta}[v] - \mathbb{E}_{d,x,\Delta}[\hat{v}]) \right] dG(\varepsilon) \\ &= \int \max_{d \in D} \sup_{\pi \in [\underline{\pi}, \bar{\pi}], z \in [\underline{v}, \bar{v}]} \psi'_{u(\pi + \varepsilon_d)}(z) \mathbb{E}_{x,d}[v - \hat{v}] dG(\varepsilon) \\ &\leq \int \max_{d \in D} \sup_{\pi \in [\underline{\pi}, \bar{\pi}], z \in [\underline{v}, \bar{v}]} \psi'_{u(\pi + \varepsilon_d)}(z) dG(\varepsilon) \|v - \hat{v}\| \\ &= M \|v - \hat{v}\| \end{aligned}$$

where

$$M := \int \max_{d \in D} \sup_{\pi \in [\underline{\pi}, \bar{\pi}], z \in [\underline{v}, \bar{v}]} \psi'_{u(\pi + \varepsilon_d)}(z) dG(\varepsilon)$$

We thus have

$$\|T(v) - T(\hat{v})\| \leq M \|v - \hat{v}\|$$

If $M < 1$, then T is a contraction mapping.

References

- ABBRING, J. H. AND Ø. DALJORD (2020): “Identifying the discount factor in dynamic discrete choice models,” *Quantitative Economics*, 11, 471–501.
- AGUIRREGABIRIA, V. AND P. MIRA (2002): “Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models,” *Econometrica*, 70, 1519–1543.
- (2010): “Dynamic discrete choice structural models: A survey,” *Journal of Econometrics*, 156, 38–67.
- ALIPRANTIS, C. AND K. BORDER (2006): *Infinite dimensional analysis*, Springer.
- ARCIDIACONO, P. AND R. A. MILLER (2020): “Identifying dynamic discrete choice models off short panels,” *Journal of Econometrics*, 215, 473–485.
- BAJARI, P., C. S. CHU, D. NEKIPELOV, AND M. PARK (2016): “Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action,” *Quantitative Marketing and Economics*, 14, 271–323.
- BANSAL, R. AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *The Journal of Finance*, 59, 1481–1509.
- BUCHHOLZ, N., M. SHUM, AND H. XU (2021): “Semiparametric estimation of dynamic discrete choice models,” *Journal of Econometrics*, 223, 312–327.
- CRAWFORD, G. S. AND M. SHUM (2005): “Uncertainty and learning in pharmaceutical demand,” *Econometrica*, 73, 1137–1173.
- DORSEY, J. (2019): “Waiting for the courts: Effects of policy uncertainty on pollution and investment,” *Environmental and Resource Economics*, 74, 1453–1496.
- EPSTEIN, L. G., E. FARHI, AND T. STRZALECKI (2014): “How much would you pay to resolve long-run risk?” *American Economic Review*, 104, 2680–97.

- EPSTEIN, L. G. AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–969.
- FANG, H. AND Y. WANG (2015): “Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions,” *International Economic Review*, 56, 565–596.
- FRICK, M., R. IJIMA, AND T. STRZALECKI (2019): “Dynamic Random Utility,” *Econometrica*, 87, 1941–2002.
- GOWRISANKARAN, G., A. LANGER, AND W. ZHANG (2022): “Policy uncertainty in the market for coal electricity: The case of air toxics standards,” Tech. rep., National Bureau of Economic Research.
- GOWRISANKARAN, G. AND M. RYSMAN (2012): “Dynamics of consumer demand for new durable goods,” *Journal of political Economy*, 120, 1173–1219.
- HANDLEY, K. AND J. F. LI (2020): “Measuring the effects of firm uncertainty on economic activity: New evidence from one million documents,” Tech. rep., National Bureau of Economic Research.
- HENDEL, I. AND A. NEVO (2006): “Measuring the implications of sales and consumer inventory behavior,” *Econometrica*, 74, 1637–1673.
- HOTZ, V. J. AND R. A. MILLER (1993): “Conditional choice probabilities and the estimation of dynamic models,” *The Review of Economic Studies*, 60, 497–529.
- IGAMI, M. (2017): “Estimating the innovator’s dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998,” *Journal of Political Economy*, 125, 798–847.
- IGAMI, M. AND K. UETAKE (2020): “Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016,” *The Review of Economic Studies*, 87, 2672–2702.

- JIA, P. (2008): “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry,” *Econometrica*, 76, 1263–1316.
- KELLOGG, R. (2014): “The effect of uncertainty on investment: evidence from Texas oil drilling,” *American Economic Review*, 104, 1698–1734.
- KOMAROVA, T., F. SANCHES, D. SILVA JUNIOR, AND S. SRISUMA (2018): “Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models,” *Quantitative Economics*, 9, 1153–1194.
- KREPS, D. M. AND E. L. PORTEUS (1978): “Temporal resolution of uncertainty and dynamic choice theory,” *Econometrica*, 185–200.
- LEWBEL, A. (2019): “The identification zoo: Meanings of identification in econometrics,” *Journal of Economic Literature*, 57, 835–903.
- LU, J. AND K. SAITO (2020): “Repeated Choice,” Working paper.
- MAGNAC, T. AND D. THESMAR (2002): “Identifying dynamic discrete decision processes,” *Econometrica*, 70, 801–816.
- MEHRA, R. AND E. C. PRESCOTT (1985): “The equity premium: A puzzle,” *Journal of monetary Economics*, 15, 145–161.
- MEISSNER, T. AND P. PFEIFFER (2022): “Measuring preferences over the temporal resolution of consumption uncertainty,” *Journal of Economic Theory*, 200, 105379.
- MILLER, R. A. (1984): “Job matching and occupational choice,” *Journal of Political economy*, 92, 1086–1120.
- NIELSEN, K. (2020): “Preferences for the resolution of uncertainty and the timing of information,” *Journal of Economic Theory*, 189, 105090.
- PAKES, A. (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica: Journal of the Econometric Society*, 755–784.

- PFEIFENBERGER, J., K. SPEES, AND A. SCHUMACHER (2009): “A comparison of PJM’s RPM with alternative energy and capacity market designs,” *Prepared for PJM Interconnection, Inc. September*.
- RUST, J. (1987): “Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher,” *Econometrica*, 999–1033.
- (1994): “Structural estimation of Markov decision processes,” *Handbook of econometrics*, 4, 3081–3143.
- SCHNEIDER, U. C. (2021): “Identification of time preferences in dynamic discrete choice models: Exploiting choice constraints,” *Working Paper*.
- SKIADAS, C. (2009): *Asset pricing theory*, Princeton University Press.
- STANCA, L. (2023): “Recursive Preferences, Correlation Aversion, and the Temporal Resolution of Uncertainty,” *Working paper*.
- STRZALECKI, T. (2013): “Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion,” *Econometrica*, 81, 1039–1074.