Time Varying Time Preferences[*]

M. Steven Holloway[†]
University of Oregon

Connor T. Wiegand[‡]
University of Oregon

Dynamic inconsistency in intertemporal choice has long been considered a hallmark of non-exponential discounting. Recent work has challenged this view from a variety of perspectives, including the view that time variance –shifting preferences between measurement dates– can also explain apparent preference reversals. While a nascent literature identifies time-variance and demonstrates its role in explaining time-inconsistency, we lack both a model that allows time-variance to tractably interact with other properties of time preference, and a longitudinal study of sufficient depth to identify such a model. In this paper, we develop the nested exponential'' discount function which is general with respect to time-invariance, time-consistency, and stationarity. The function nests both exponential discounting and a version of present-biased discounting within its parameter space, enabling transparent model selection at both the aggregate and subject levels. We evaluate time-invariance and the performance of the nested exponential model in a 12-week longitudinal study featuring seven surveys. Our elicitations give us unprecedented precision in estimating dynamic inconsistency, non-stationarity, and time-variance. We find that subjects in our study exhibit significant decreasing patience over the course of the study, and that time-variance explains roughly 72% of time-inconsistent choices in our data. This does not mean our data are best-explained by exponential discounting plus preference drift: hyperbolicity is a key feature of our data, and it is well captured by the nested exponential function.

[*]IRB approval for this research was obtained from the University of Oregon, Study ID: STUDY00000768.

[†]Department of Economics, 1285 University of Oregon, Eugene, OR 97403, USA. E-mail: marcush@uoregon.edu. Web: www.mstevenholloway.page/home

[‡]Department of Economics, 1285 University of Oregon, Eugene, OR 97403, USA. E-mail: cwiegand@uoregon.edu

1 Introduction

A substantial body of evidence documents that economic preferences can be unstable.[4] Beyond shocks to preferences, decision-makers also learn about their preferences (either through feedback (Charness, Chemaya and Trujano-Ochoa, 2023) or the realization of uncertainty (Carrera *et al.*, 2022)) or "choose" preferences to better suit their environment (Bernheim *et al.*, 2021). Preference instability, shocks, and drift are especially problematic for researchers looking to estimate time preferences. Consider the framework of (Read and Leeuwen, 1998): an individual selects a healthy snack for their future self, but revises their choice to an unhealthy snack when the time for consumption arrives. Is this preference reversal best explained by a fixed feature of that individual's discount function (e.g. hyperbolic discounting), learning about preferences, or some other unobserved environmental factor? Halevy (2015) defines three properties of time preferences: time-consistency, stationarity, and time-invariance. Imagine a decision-maker evaluating a paycheck advance –which results in some distant future disutility– to purchase tickets to a concert. All-else equal, his time preferences violate *stationarity* if he takes an advance **today** to see a concert **tonight**, but does not take an advance **today** to see an equivalent concert **tomorrow**. His time preferences violate *time-consistency* if after deciding **yesterday** not to take an advance for **tonight's** concert, he wakes up **today**, changes his mind and takes the advance to see the concert **tonight**. His time preferences violate *time-invariance* if after deciding **yesterday** not to take an advance for **last night's** concert, he wakes up **today** and decides to take an advance for **tonight's** concert. Any two of these properties imply the third; if the behavioral economist theorizes that the decision-maker is not time-consistent, she assumes that the decision-maker is either not time-stationary or not time-invariant (or neither). Common models of discounting in behavioral economics all pair inconsistency with non-stationarity, while leaving time-invariance intact.

Share of sample:	Min. across approaches	Max. across approaches
Properties of choices	(1)	(2)
A) Invariant, consistent, stationary	35.04%	59.32%
B) Invariant only	6.83%	11.86%
A) + B)	43.58%	67.79%
C) Stationary only	16.95%	27.35%
D) Consistent only	5.08%	16.24%
E) None	4.27%	27.12%
C) + D) + E)	32.21%	56.42%

Table 1: Time preference property classification from citep()

Notes: Data are from Table 2 in citep(). There are a total of eight classifications presented in the table, representing the 2x2x2 treatment variation of stake size, strict vs. weak preferences, and whether subjects had full knowledge of the future aspects of the study when they made their initial choices.

[4]For example, traumatic events like the Great Depression (Malmendier and Nagel, 2011), exposure to violence (Voors *et al.*, 2012; Callen *et al.*, 2014), and natural disasters (Eckel, El-Gamal and Wilson, 2009; Bchir and Willinger, 2013; Cameron and Shah, 2013; Page, Savage and Torgler, 2014; Cassar, Healy and Kessler, 2017; Hanaoka, Shigeoka and Watanabe, 2018; Kettlewell *et al.*, 2018; Kuroishi and Sawada, 2019; Reynaud and Aubert, 2019; Bourdeau-Brien and Kryzanowski, 2020; Beine *et al.*, 2020; Kibris and Uler, 2021), short-term influences like happiness (Ifcher and Zarghamee, 2011), loss of control (Gneezy and Imas, 2014), and hunger (Ashton, 2015; Kuhn, Kuhn and Villeval, 2017), as well as environmental factors like rainfall (Jaramillo, LaFave and Novak, 2023) and political instability (Sepahvand, Shahbazian and Swain, 2019) all have demonstrated causal impacts on risk preferences.

Our goal in this paper is to provide a comprehensive treatment of time-variance in estimating time preferences. This entails a theoretical and structural discussion of how to model and allow for time invariance, and a 12-week longitudinal experiment that constructs a panel of intertemporal choices across seven elicitations. While the two-period studies in Halevy (2015), Janssens, Kramer and Swart (2017), and Harrison, Lau and Yoo (2023) allow for the impact of time-variance on experimental measures, they do not give the researcher much power to think about the source of time-variance and model it: is it noise, learning, a result of macroeconomic changes, a result of individual-level shocks to wealth, liquidity, stress, or something else entirely? Answering this question determines the scope of the time-variance problem. If time-variance is idiosyncratic noise, then what appears to be time-variance or inconsistency may instead arise from decision error. Failing to account for this would lead to individuals being incorrectly classified as certain types of discounters, but would not impact aggregate estimates of how a population discounts (besides decreasing precision). On the other hand, if time-variance is a systematic process wherein subjects learn about their preferences within the context of the novel experimental setting they encounter, or their time preferences react to correlated environmental factors, the problem is more pernicious. What is the risk associated with mistaking one kind of time-inconsistency for the other? The behavioral literature on "present-bias'' and associated policy interventions (e.g. commitment devices) are predicated on time-inconsistency as a pernicious effect of the opportunity for immediate gratification. If choices can be structured to avoid that opportunity, then the decision-maker can implement their stable long-run preference for saving more, eating healthier, etc. Time-inconsistent choices associated with time variance do not have the same interpretation. Consider a decision-maker that exhibits time-varying preferences such that their intertemporal allocations between rewards on two relative dates (e.g. the day of the choice and eight weeks after the choice) are getting less patient as time passes (meaning that the date of the choice and both reward dates move together in time). First, assume her preferences are time-stationary. If a researcher estimates her quasi-hyperbolic (β-δ) discount function via an observed violation of time-consistency, the researcher will conclude that she is present-biased ($\beta < 1$). Instead, assume her preferences are time-consistent. If her quasi-hyperbolic discount function is estimated via an observed violation of stationarity, the researchers will conclude that she is future-biased ($\beta > 1$).[5]. The first part of the paper is about why time variance is a problem for identification in time-invariant models, and then introduces models of intertemporal choice that allow time-varying preferences. We The second part of the paper presents our longitudinal experiment. Over the course of 12 weeks subjects take part in seven surveys, each featuring 6-10 incentivized time preference measurements. This design gives us six opportunities to observe time-inconsistency and time variance (for each measurement), and seven opportunities to observe non-stationarity. We then shift to analyses that allow for time-varying preferences. We identify a significant time trend in the aggregate discount factor consistent with *decreasing patience* over the period of the study. This trend stems from subject-level changes across time, not sample composition, suggesting systematic rather than idiosyncratic time variance. With time-variance accounted for, we can measure distinct inconsistent and non-stationary behavior at the subject-level; we find that inconsistency is more common and substantial than non-stationarity. The prevalence of individual violations of all three time preference properties leads us to classify subjects as time-variant at the upper end of the range Halevy (2015) suggests. Overall, time-variance explains 72% of the time-inconsistency in our sample. Finally, we structurally estimate the nested exponential discount function at both the aggregate and individual levels. The nested exponential model fits the data best, but adjusted for its extra parameters, it only slightly outperforms a version of the model constrained to be time-invariant only. In other words, classic "present-bias'' where time-inconsistency and non-stationarity always go together does a decent job at explaining aggregate behavior *so long as researcher uses a discount*

[5]We show this more formally in Section 2..1

function with some true hyperbolicity, rather than a discrete approximation. This model vastly outperforms all other two-parameter models, including the quasi-hyperbolic β-δ model.[6] At the individual level, we find that the full nested exponential model and the time-invariant only restricted model both best-describe about 40% of the sample. The remaining 20% fall into roughly equal-sized groups of exponential, stationary-only, and time-consistent-only discounters. Overall, we assign 54% of the sample to time-varying models of discounting, suggesting that understanding, measuring, and modelling preference shifts over time is crucial for accurately estimating time preferences.

2 Theory

We adopt notation from Halevy (2015) to describe the stationarity, time-consistency, and time-invariance properties of time preferences. Consider an agent in calendar time $\tau \geq 0$ evaluating bundles $(\varepsilon,\ t)$ and $(\psi,\ t')$, which deliver rewards $\varepsilon,\ \psi$ at time periods $t,\ t' \geq \tau$. We assume that the agent has complete and transitive preferences over all such bundles, according to a preference ordering that may depend on τ, $\succsim_\tau$.

Under classical consumer theory assumptions, including additive separability of utility across time periods, this preference relationship is represented by a *discount function*, $D(\tau,\ t)$, such that

$$(\varepsilon,\ t) \succsim_\tau (\psi,\ t') \iff D(\tau,\ t) \cdot \varepsilon \ \geq\ D(\tau,\ t') \cdot \psi^7 \tag{1}$$

$D(\tau,\ t)$ describes the agent's discount "factor" when evaluating, in period τ, a payoff in period t.

Definition 0.1. An agent's preferences satisfy **stationarity** if

$$D(\tau,\ \tau + \Delta_1)\ \varepsilon\ \geq\ D(\tau,\ \tau + \Delta_2)\ \psi \iff D(\tau,\ \tau' + \Delta_1)\ \varepsilon\ \geq\ D(\tau,\ \tau' + \Delta_2)\ \psi$$

An agent whose preferences are *stationary* will not exhibit preference reversals when the payoff time of all bundles is shifted by the same number of periods, all else equal.

Definition 0.2. An agent's preferences satisfy **time consistency** if

$$D(\tau,\ \tau + \Delta_1)\ \varepsilon\ \geq\ D(\tau,\ \tau + \Delta_2)\ \psi \iff D(\tau',\ \tau + \Delta_1)\ \varepsilon\ \geq\ D(\tau',\ \tau + \Delta_2)\ \psi$$

Agents whose preferences are *time-consistent* will not exhibit preference reversals when the same rewards, in the same payoff periods, are evaluated in different time periods.

Definition 0.3. An agent's preferences satisfy **time invariance** if

$$D(\tau,\ \tau + \Delta_1)\ \varepsilon\ \geq\ D(\tau,\ \tau + \Delta_2)\ \psi \iff D(\tau',\ \tau' + \Delta_1)\ \varepsilon\ \geq\ D(\tau',\ \tau' + \Delta_2)\ \psi$$

Agents whose preferences are *time-invariant* will not exhibit preference reversals when the evaluation period and payoff periods for all bundles are shifted by the same amount. We use TICS (time-invariance, consistency, stationarity) to refer joints to these three properties of time preferences.

Halevy (2015) proves the following proposition, linking these three properties together:

Proposition I. *The TICS properties form a binary. That is, if a preference relation satisfies any two of the properties, then it must imply the third.*

Thus, the TICS properties generate five groups that any agent's preferences may fall into: all, time-invariant only, stationary only, time-consistent only, none. The consequences of this binary relationship are not well-appreciated in the literature. Economists who wish to model time-

[6] Hyperbolicity is a key feature of our data; among one-parameter models, the true hyperbolic discount function outperforms exponential discounting.

[7] Without loss of generality, we replace $u(\varepsilon)$ and $u(\psi)$ with ε and ψ, under the assumption an instantaneous utility function representation exists.

inconsistent agents, for instance, must also depart from at least one of stationarity or time-invariance. The latter approach is typical; the hyperbolic and β-δ discount functions satisfy only time-invariance. There exists only a nascent literature on discount functions that do not satisfy time-invariance (Strulik, 2021), and on observed time-variant behavior (Halevy, 2015; Janssens, Kramer and Swart, 2017; DeJarnette, 2020; Imas, Kuhn and Mironova, 2022; Brownback, Imas and Kuhn, 2023; Harrison, Lau and Yoo, 2023). Our goal is to both offer new structural discount function that is fully general to the TICS properties and evaluate its time-variance properties in a longitudinal study.

Before proceeding, there are two terms that we will define for use in this paper that are convenient short-hands for phenomena we observe, but come with a lot of baggage in the literature: *(im)patience* and *present/future-bias*. In some papers these terms are closely linked (e.g. Prelec (2004), Attema *et al.* (2010), or Takeuchi (2011), where present-bias is considered decreasing impatience). However, it will be useful for us to keep them separate in order to use terms like increasing and decreasing patience to describe time variance. Specifically, a subject in our study exhibits decreasing patience if their discount factor for some reward delayed by $t - \tau$ periods is decreasing in τ: $D(\tau, \tau + \Delta) > D(\tau', \tau' + \Delta)$ for $\tau' > \tau$, $\Delta > 0$.

Present-bias is a term that either casually or formally refers to the phenomenon wherein an agent's relative preference for sooner rewards over later rewards is greatest when those rewards are immediately available (in the evaluation period). This is usually modeled discontinuously by economists, with extra discounting applied to all non-immediate rewards using the workhorse quasi-hyperbolic (or β-δ) discounting model (Laibson, 1997; O'Donoghue and Rabin, 1999), but is a continuous feature of the hyperbolic discounting model widely used outside of economics. Beyond ambiguity over whether the term applies only to discontinuous immediate gratification models, there is also a recent debate over whether the term "bias" is appropriate to describe an empirical phenomena that could stem from a variety of root causes (Bernheim and Taubinsky, 2018). In this paper we will use "present-bias" and "future-bias" to empirically classify non-stationary and time-inconsistent behavior. If shifting the payoff time of each bundle within a binary choice further into the future decreases the relative preference for the sooner bundle, we call this present-bias. Similarly, if shifting the evaluation time of a binary choice closer to its fixed payoff time periods increases the relative preference for the sooner bundle, we call this present-bias as well.

2..1 Measuring discounting properties

Consider a decision maker (DM) faced with a choice between an immediate reward at $t = \tau = 0$ and a reward on some future date $t = t_1 > 0$. Call a a scalar measure of how much their choice favors the reward at the later date, t_1, i.e. higher values of a result from more patient discounting of that delayed reward. In the same evaluation period, the DM also faces a choice between utility in two future time periods, $t_2 > 0$ and $t_2 + t_1$. Call b the corresponding measures of that choice. After time passes, such that $\tau = t_2$, the DM faces another similar choice: call c the measure of the choice between immediate utility at $t = \tau = t_2$ and future utility at $t_2 + t_1$. We call the collection of these three choices a "decision triangle." Using the values associated with any given triangle, $a - c$ is a measure of time-variance; how much a choice over fixed relative dates changes over time. $b - c$ measures inconsistency; how much a choice over fixed calendar dates changes over time. Finally, $b - a$ measures non-stationarity; how much a choice with a fixed delay changes in response to front-end delay. Combining these three calculations, we have that

$$\underbrace{a - c}_{\substack{\text{Time} \\ \text{Variance}}} = \underbrace{a - b}_{\text{Non-Stationarity}} + \underbrace{b - c}_{\text{Inconsistency}} \tag{2}$$

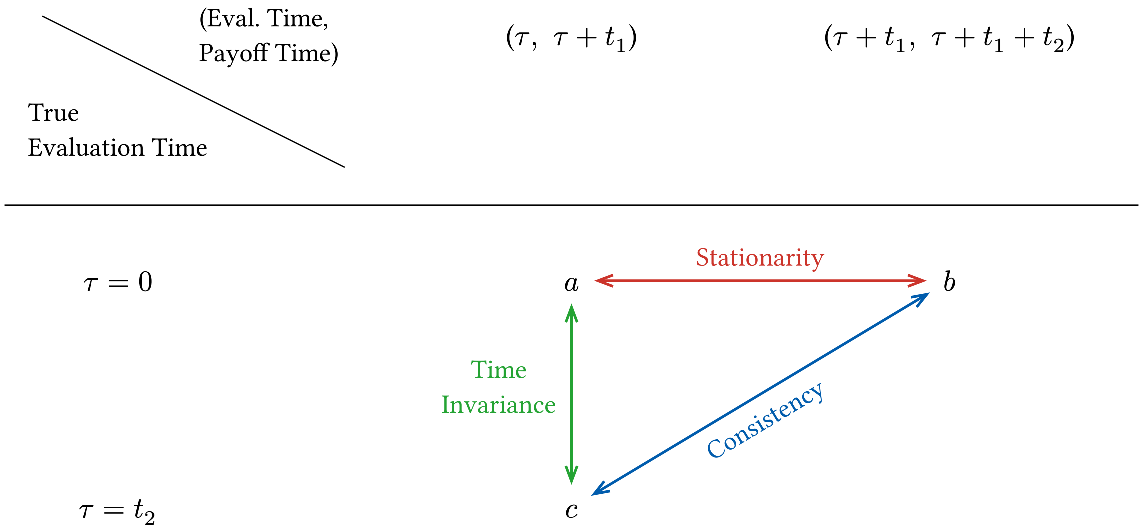

Figure 1: Time preference property decision triangle

With this framework, we can illustrate the pernicious empirical issues that arise when a researcher is not accounting for time-varying preferences. Suppose a researcher wishes to estimate parameters for an agent using the β-δ model, and they set up an experiment to test the agent's time-consistency. Assume the agent exhibits time-varying preferences in the form of decreasing patience, leading to $a - c < 0$. Assume also that the agent exhibits stationary preferences, so that $b - a = 0$. If this is the case, then the agent must exhibit time-inconsistent preferences such that $b - c = a - c < 0$. Thus, the researcher will conclude that the subject is present-biased and estimate $\beta < 1$. Now assume instead that the same time-varying preferences are present, but the agent is time-consistent and the researcher sets up an experiment to test the agent's stationarity. Since $b - c = 0$, it must be that the agent makes a non-stationary choice so that $a - b = a - c < 0$. Thus, the researcher will find the agent is future-biased, and estimate $\beta > 1$. Depending on research design, the same unobserved time-variance can lead to different incorrect conclusions about the agent's true behavioral response to an immediate gratification opportunity.

This is more complicated when considering an agent that exhibits violations of all three time preference properties. Consider a decision triangle, defined by the triple (a, b, c) of $(240, 235, 230)$. A stationarity violation exists, as $240 \neq 235$. The researcher would deem this choice future-biased, as an increase in payoff times for both bundles led the agent to be more impatient. A time-consistency violation exists, as $235 \neq 230$. The economist would deem this choice present-biased, as the passing of calendar time until the fixed payoffs led the agent to be more impatient. This confounds the estimation of β and may lead the economist to conclude that this agent is not present- or future-biased. While true on average, this is not a thorough explanation of this agent's preferences.

To account for time-varying preferences, the limited existing literature has used reduced-form approaches to correcting $b - c$ in light of $a \neq c$, or (less-likely given design constraints) correcting $a - b$ in light of $a \neq c$ (e.g. Janssens, Kramer and Swart (2017) or (Harrison, Lau and Yoo, 2023)). We offer an alternative: a time-varying structural model of intertemporal choice. The following section maps the five groups of time preference properties given by the TICS binary to structural discount functions and their properties.

2..2 Classes of Time Preferences

Time-invariant preferences are well represented in the literature. Discount functions that admit these preferences include the standard exponential, as well as the hyperbolic and β-δ models regularly used to model time-inconsistent decisions. These models all share a common property: They are a function of only the distance between payoff time and evaluation time, $t-\tau$. Since these models are only concerned with this distance, the $D(\tau,t)$ function is typically reduced to just $D(t)$, where τ is always 0, without loss of generality. This reliance on distance payoff time characterizes the necessary and sufficient condition for a function to admit time-invariance.

Proposition II. *A discount function admits time-invariant preferences if and only if* $D(\tau,\ t) = D(t-\tau)$.

See Appendix Section 8..2 for the proof.

We are only aware of one discounting functional form in the literature that violates this property. Strulik (2021) introduced a time-consistent hyperbolic discount function which violates the assumption of time-invariance:

$$D(\tau,\ t) \leq \left(\frac{1+\alpha\tau}{1+\alpha t}\right)^{\beta} \qquad \alpha > 0,\ \beta > 1. \tag{3}$$

Strulik (2021) designs the function for use in environmental resource models, relying on the increasing patience of the model to avoid extinction steady states. Importantly, commitment to consumption plans in his context requires that the discount function maintains time-consistency. Therefore, this model must be only time-consistent, as it is not time-invariant and so cannot be both time-consistent and stationary. How do we ensure that a discount function maintains time-consistency? Strotz (1956) suggests that the only discount function that can achieve time-consistency is the standard exponential. However, this result relies on the assumption of a discount function that is a function of only $t-\tau$. Strotz's result has since been generalized to prove instead that time-consistency of a discount function is equivalent to multiplicative separability in t and τ (Burness, 1976; Drouhin, 2020).

Proposition III. *A discount function $D(\tau,\ t)$ admits consistent preferences if and only if it can be written:*

$$D(\tau,\ t) = f(t)\cdot g(\tau)$$

Further, if the discount function can be written in this way, then $g(x) \equiv \frac{1}{f(x)}$.

See Appendix Section 8..2 for the proof.

Including the time-consistent hyperbolic model, three of the five possible time preference groups have been characterized by discount functions in the literature. To our knowledge, discount functions have not been developed to fit the two remaining groups: admitting only stationarity and admitting none of the TICS properties. In order to develop these functions in the next section, we present the necessary and sufficient conditions for a discount function to admit stationarity, which also, to our knowledge, has not been shown in the literature.

Proposition IV. *Let $\alpha \in \mathbb{R}^+$ and f be a C_1 function linear in $t-\tau$. A discount function admits stationary preferences if and only if it takes the form*

$$D(\tau,\ t) = \alpha^{-(t-\tau)f(\tau)}.$$

See Appendix Section 8..3 for the proof.[8]

With propositions II, III, and IV, we arrive at the following well-known result:

[8]Note that f may be trivial in τ. If $\alpha > 1$, we require that $f(\tau) < 0$ (equivalent to dropping the minus sign in the exponent).

Corollary. *The discount function $D(\tau,\ t)$ admits preferences exhibiting stationarity, time-consistency, and time invariance if and only if it is the exponential discounting function. That is, $D(\tau,\ t)$ can be written as a function of $t - \tau$ alone,*

$$D(t) = \delta^{\beta(t-\tau)}$$

2..3 Nested Exponential Discount Function

With necessary and sufficient conditions for a discount function to satisfy any of the TICS properties, we now have the tools to develop a discount function that can fit all five possible classes of time preferences. Two goals motivate our approach to developing such a function. First, the discount function should be fully general in its admittance of time preference properties. For instance, the β-δ model acts as an extension of the standard exponential, allowing for time-inconsistent and non-stationary preferences by setting $\beta \neq 1$. By extending the standard exponential in a similar manner, we can achieve each of the five time preference classes according to its parameter values. Second, a discount function that allows for time-varying preferences should be capable of exhibiting both increasing or decreasing patience depending on parameter values. The time-consistent hyperbolic model only allows for increasing patience, which is appropriate contextually. Allowing for decreasing patience, as well, however, expands the pool of time preference profiles we can accurately model.

$$D(\tau,\ t) = \delta^{\frac{(t^{\mu} - \tau^{\mu})^{\eta}}{\gamma^{\tau}}} \qquad \eta > 0,\ \gamma > 0,\ \mu > 0$$

This function, which we call the *nested exponential discount function*, can satisfy any of the five possible classes of time preference properties. To see this, suppose $\eta = \gamma = \mu = 1$. Then the nested exponential is identical to the standard exponential discount function. Now suppose $\eta = \gamma = 1$ while $\mu \neq 1$. Then the function is no longer a function of only $t - \tau$ nor does its exponent satisfy the conditions required to maintain stationarity. However, the function does remain multiplicatively separable in t and τ, granting it time-consistency. The argument follows similarly for when $\eta = \mu = 1$ and $\gamma \neq 1$. The function is not a function of only $t - \tau$ and is no longer be multiplicatively separable in t and τ. Its exponent is however can be expressed as $-(t - \tau) f(\tau)$, granting stationarity.

When $\gamma = \mu = 1$, but $\eta \neq 1$, the nested exponential is a discrete-time reformulation of the Weibull discount function, a less common behavioral model that achieves the same violations of the TICS properties as the hyperbolic and β-δ models. We have that the function is not multiplicatively separable in t and τ, the exponent is not of the form $-(t - \tau) f(\tau)$, but the function is still only a function of $t - \tau$. For the final class, satisfying no TICS properties, since any parameter not equalling one results in the violation of two properties, whenever two or more parameters are not equal to one, none of the properties are satisfied. Therefore, all five groups of time preference properties can be modeled with this one function.

This nested exponential function also allows for increasing or decreasing patience depending on the parameter values of η, γ, and μ. For instance, values of $\gamma < 1$ correspond to decreasing patience[9]. As τ increases, the relative distance between t and τ is amplified by a denominator that is increasingly less than 1, in turn increasing the exponent, and therefore decreasing the discount factor. The opposite holds for values of $\gamma > 1$. Values of $\mu < 1$ correspond to increasing patience. For any fixed $t - \tau = C$, the corresponding $t^{\mu} - \tau^{\mu}$ is decreasing in τ when $\mu < 1$, resulting in a discount factor increasing in τ. The opposite again holds for $\mu > 1$. Since η on its own cannot cause increasing or decreasing patience, it only acts to mitigate or amplify the effects of the other parameters when

[9]Assuming $\eta = \mu = 1$

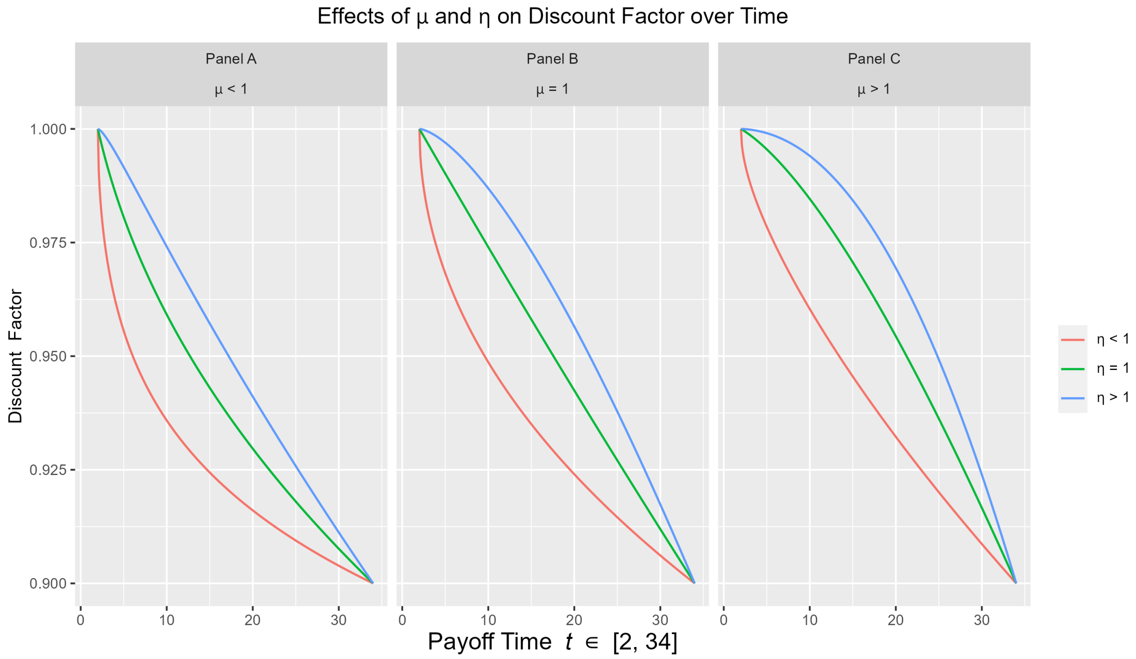

Figure 2: Nested exponential discount function by μ and η parameter values

Notes: the value of δ in each function is chosen such that $D(2, 34) = 0.9$. The evaluation time τ is two in each function, and the x-axis shows payoff dates, t, from 2 to 34. We set $\gamma = 1$ for all functions.

active at the same time. The exact condition that dictates the direction of patience of the nested exponential discount function is given by:

$$D(\tau, \tau + \Delta) > D(\tau', \tau' + \Delta) \iff \gamma - \exp\left[\eta\mu \frac{(\tau + C)^{\mu-1} - (\tau)^{\mu-1}}{(\tau + C)^{\mu} - (\tau)^{\mu}}\right] > 0 \tag{4}$$

Figures 2 and 3 demonstrate the effect each parameter has on the discount function in comparison to the standard exponential form (note that we have selected the δ parameter to force each function in Figure 2 to be equal at a common point, and that the scales of the vertical axes differ across figures for visual clarity). In Figure 2, the effect of η heavily influences the curvature of the discount function. Letting $\mu = 1 = \gamma$, when $\eta \neq 1$ the function exhibits present- or future-bias in the same way other invariant-only functions do. $\eta < 1$ results in present-bias, where the discount function decreases at a faster rate than the standard exponential for periods close to evaluation period, and then eventually at a slower rate.

The effect of μ can be intuited from the example of competing measurements of present- and future-bias presented in Section 2..1. Start by considering time preferences that are only time-consistent. If decreasing patience is present, it will result in future-bias when measuring discounting within evaluation period (a stationarity violation). Indeed, we see that the function where $\mu > 1$ results in a higher discount factor (relative to $\mu \leq 1$) for any delay length due to a smaller rate of change in response to an increase in delay length. The intensity of this effect is partially driven by evaluation period.

The effect of γ can be described using the same example. If time preferences are stationary, then changing patience (stemming from $\gamma \neq 1$) will result in time-inconsistency. This time-inconsistency occurs because $\gamma \neq 1$ results in higher or lower discount factors for the same delay length across evaluation period. In particular, when $\gamma < 1$, the curve is distorted downwards, so that discount factors are lower for all delay lengths greater than zero as τ increases. This results in a present-biased decision, as the rate of change of the discount rate is not only necessarily more rapid on the

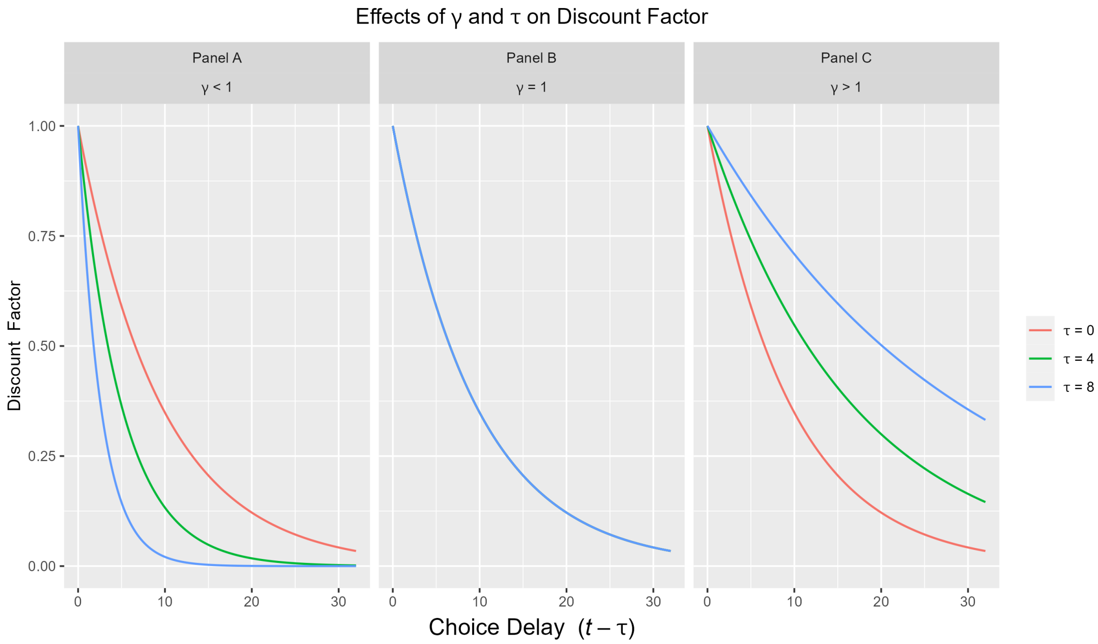

Figure 3: Nested exponential discount function by γ parameter value

Notes: for all functions we assume $\mu = \eta = 1$, and $\delta = 0.9$. The x-axis shows choice delay, $t - \tau$.

curve for the new evaluation period, but the decision is closer to evaluation period, where the rate of change is quickest.

When it comes to estimating time-invariance and the parameters of nested exponential discount function, as opposed to a reduced-form invariance correction, the longitudinal dimension of the data determines the precision and reliability of those estimates. Existing work by Halevy (2015) and Janssens, Kramer and Swart (2017) collects data on a single decision triangle (ignoring different prices or choice delays within a triangle).[10] We observe six triangles per subject over the course of a 12-week study. Time-variance need not be idiosyncratic subject-level noise at different dates; instead it can be systematic learning (about either the decisions themselves or preferences over them), changes to the macroeconomic environment, the psychological environment, or preference evolution. Indeed, to preview our results, we find evidence of both generally decreasing patience throughout our study, and an uptick in patience during finals week among our student sample.

3 Experimental Procedures

The details of our experimental design and analysis were pre-registered on AsPredicted.org, protocol #133180. We recruited subjects from introductory-level economics courses at the University of Oregon. A total of 178 subjects were invited to participate in the experiment, and 153 consented and completed the first survey. After initial recruitment, which occurred both in-person and over email, all communication with subjects was conducted online through the Qualtrics survey platform. While the frequent, short points of contact in our design lend themselves well to an online study, because of the delayed and risky incentives, we decided it was important to run the study in a sample where we could leverage our institutional reputation to establish trust, and where subjects would have in-person access to the study administrators in case of any problems or concerns. Participants were provided with a study email address, as well as the office location and phone number for the faculty study PI.

[10] Harrison, Lau and Yoo (2023) collects longitudinal data and uses structural assumptions to analyze them like a triangle.

The study consisted of seven surveys; one every two weeks for a total of twelve weeks. Every subject that completed the first survey earned $5, paid upon completion, and every subject that completed the final survey earned and additional $35 upon completion. All payments were made via instant money transfer on the subjects preferred platform.[11] The first survey featured detailed instructions on our protocol and the last featured an exit questionnaire. Both took subjects roughly 20 minutes to complete. Each intermediate survey featured only prize drawings, price lists, and a very brief questionnaire, each taking roughly 10 minutes to complete.

Within each survey, subjects encountered several 'multiple price lists' (MPLs) that compared two amounts of 'tickets' presented in two columns, where each column represents a prize-drawing for $300 in a specific time period, and each row represents a choice between a (tickets, time period) bundle. Each ticket corresponded to an increment of 1 in 100,000 in odds within a drawing. The left column always offered tickets in a drawing that would take place earlier than the drawing in the right column. For all subjects, in every price list, the left column featured a sequence of 22 rows with 202 to 244 tickets by increments of two, whereas the right column featured a fixed 242 tickets in every row. This design is an adaptation of the multiple-lottery-list elicitation of `citep`. See Appendix Figure 14 for an example price list.

We denote surveys by the number of weeks from the start of the study that they occur (to correspond to the t and τ variables that the discounting models will be based on). In weeks zero and two (the first and second surveys) we offered subjects six choices: three "choice delay" lengths between drawings –two, four, and eight weeks– crossed with two "front-end delays" between the evaluation date and the earlier drawing –zero and two weeks. This design constructs the decision triangles outlined in Section 2..1; when a subject responds to sequential surveys, we can observe three decision triangles (one for each choice delay), allowing us to detect violations of any of three time preference properties. Starting in week four, we phased in much longer choice delays of 16 and 32 weeks, at first only with the two-week front-end delay, but then in week six, without the front-end delay. Week 12 was the final week that involved subject choices, and we did not elicit choices with the two week delay, as we would be unable to observe whether they revised those choices in the future. Figure 4 outlines the price lists faced by subjects in each survey.

The goal of any price list elicitation is to observe a single switching point between columns, such that the switch point approximately identifies at which row –as the value of one of the column's rewards changes monotonically– a subject is indifferent between the bundle in the left column and the bundle in the right column. If the subject switches more than once, this confounds that identification. A single switching point in tickets was enforced by our elicitation device, a bright yellow sliding bar that separated choices that the subject preferred the earlier drawing from the choices where subjects preferred the later drawing. Subjects were instructed to place the sliding bar *between* the last choice the preferred in the right column and the first choice they preferred in the left column. Specifically, we instructed them to put every row where they prefer the later drawing above the bar, and every row where they prefer the earlier drawing below the bar. The bar itself expressed their switch point in words, which changed when the bar was dragged to a new location. Subjects could leave the bar at the very top –indicating they always prefer the earlier drawing– but they had to actively click and release it to do so.[12] Subjects could not advance through a price list for at least six seconds in order to limit random or thoughtless choices.

All drawings had a positive probability of being realized, but only *one price list per earlier-drawing date* was selected to count. For example, in week zero, with equal probability, one of the price lists

[11]Subjects chose from Venmo, PayPal, CashApp, and Zelle, and we pre-screened potential participants for whether they were comfortable using one of these platforms to receive payment.

[12]We thank Antonia Krefeld-Schwalb for the idea, and the source code to help us build this tool.

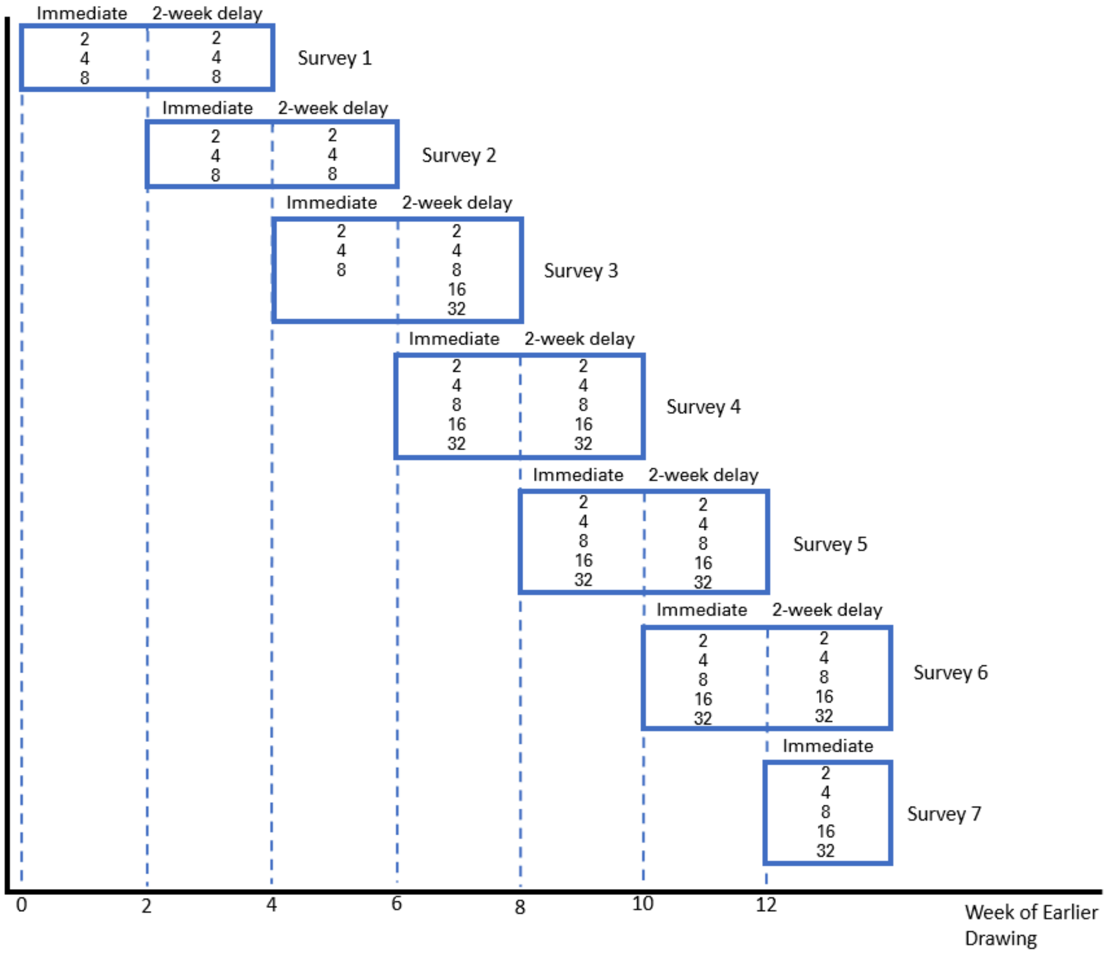

Figure 4: Study timeline

with an immediate earlier drawing was randomly selected, and then one of 22 rows was randomly selected from that list, again with equal probability. If the subject picked the immediate drawing in that row, the drawing took place at the end of the week zero survey. If the subject picked the later drawing in that row, the drawing took place *at the beginning* of the survey taking place in that week. The week zero price lists with a two-week front end delay were "saved for later": when subjects returned in week two and made another set of choices featuring an earlier drawing on that date, the price list that counted for that earlier date was randomly selected from the set of week zero and week two choices.[13]

Prize drawings were implemented with the selection of a random number from 1 to 100,000. If the random number was less than or equal to the number of tickets a subject has in a drawing, they would win. Winning automatically triggered an email to a study administrator to enable an immediate transfer of the $300 prize. Three subjects won drawings, including one in the first survey.[14] Many drawings took place after the survey portion of the study ended; in these cases, subjects received emails with a link to participate in the drawing with no data collection. Before participating in incentivized price lists in week zero, subjects had to fill out an example list (however they wanted), correctly use the yellow slider bar to illustrate a specified preference, choose a random drawing number that would allow one hypothetical subject to win and another to lose based on their choices, and answer a question about the independence of choices across price lists. Each

[13]First, we randomly determined with 50-50 chance from which week the choice-that-counted would come from, and then we randomly selected the price list from that week.

[14]As of this draft. The four most-delayed drawings have yet to occur.

survey after week zero included a refresher question that subjects had to correctly answer prior to making their incentivized choices.

Finally, while the data are not the subject of the current paper, each survey concluded with a brief questionnaire on recent life events. Subjects were asked about financial and psychological shocks they may have experienced since the last survey. The same questions were asked every survey, except for one rotating question that assessed risk preferences, cognition, and other measures of time preferences not captured by our main elicitation. The final survey also included an exit questionnaire that asked subjects about their understanding of the survey, demographics and socioeconomic status, expectations about the future, subjective time horizon, patience, and impulsivity.

The first survey was distributed at 7am on a Wednesday in the middle of the university term. All subsequent surveys were also distributed at 7am on Wednesday, covering the end of that term, finals week, and much of the subsequent term.[15] Each survey was open for a 26 hour window, from 7:00am on Wednesday to 9:00am on Thursday. Subjects were allowed to miss one out of the seven surveys and were barred from further participation after their second missed survey.

3..1 Attrition

While a concern in any longitudinal study, the length of time this study covers coupled with seven attempted points of contact, further coupled with the concern that attrition may be non-random with respect to discounting, made attrition a substantial ex-ante concern. For instance, if those that leave the study are less patient than those who remain –because the study features heavily delayed incentives– then the full sample may exhibit increasing-patience time variance in the aggregate despite no individual-level time-variance. Given resource constraints on offering overwhelming incentives and time constraints in hunting down missing individuals to complete a survey within a 26-hour window, we opted for an approach that 1) made completing each survey fast and easy, with multiple reminders and direct, individualized participation links (as opposed to requiring login information), 2) allowed us to carefully test for any selective attrition, and 3) allowed us to analyze time-variance at the subject level.

Of the 153 subjects that completed the first survey, 97 (63%) completed the last survey. 79 (52%) of subjects completed every survey. Throughout the results section of the paper, we present data and findings from the full sample of 153 subjects and 6,367 price lists, but we discuss and reference results from the "balanced" sample of 79 subjects throughout, with corresponding tables in the Appendix. None of the main results of the paper depends on which sample we use. In general, the balanced sample results are what we would expect if subjects that attrit make noisier choices and removing them reduces measurement error. We formally test for differential attrition in Section 5, and find no significant differences in the discount factor at any point in the study between those who are about to attrit and those who will remain (see Figure 9).

One mitigating factor for concerns about attrition is our focus on subject-level results. The study is designed to allow us to observe the time preference properties exhibited by subjects and estimate subject-level parameters of discount functions. For all subjects that completed at least two subsequent surveys –141 in total, 92% of the initial sample– our experiment provides sufficient data to inform our research question.

[15] The term length is 10 weeks, so we were unable to avoid an end of term somewhere in the study. Between having a survey in finals week or spring break, we selected finals week under the assumption that students would at least be on their computers, even if they were busy.

4 Time-invariant Behavior

We present our results in three sections. Here, we describe discounting behavior in our study from a time-invariant perspective, to establish the comparability of our data with other elicitation efforts in the literature. In the next section, we consider time variance from a reduced form perspective, including a consideration of attrition. Finally, we take a structural approach to time variance, returning to the motivating issue of accurately modeling subjects' axiomatic discounting.

4..1 Descriptive statistics

We first establish sample-wide discounting. Without discounting apparent in the aggregate, it would be unclear whether our elicitation method is inducing subjects to make intertemporal tradeoffs. We measure subject i's implied discount factor in price list j in survey week w, d_{ijw} as the midpoint of their switch interval divided by 242 (the fixed ticket allocation to the later lottery).[16] For example, if a subject prefers 242 tickets in the later draw to 230 tickets in the sooner draw, but 232 tickets in the sooner draw to 242 tickets in the later draw, then $d = \frac{231}{242} = 0.9545$.

Across every survey and price list for subjects that completed every survey[17], the average d is 0.9560 (corresponding to a switch point between 230 and 232 tickets). This is significantly different from one ($p < 0.0001$), consistent with a positive discounting over the tickets in our study.[18] For each week the delay increases, d decreases by 0.07pp (p = 0.0013), and when the sooner draw is immediate, d decreases by 0.24pp (p = 0.0007), consistent with mild but statistically significant present bias.[19] Assuming exponential discounting, the average *weekly* discount factor across the balanced sample is 0.9920 ,which corresponds to an *annual* discount rate of 51.84%, although we will show in Section 4..2 that once adjusted for present bias, the rate is considerably lower. Figure 5 shows the implied discount functions for our data, keeping choices where the sooner lottery is immediate separate from choices where it is delayed by two weeks.

At every delay length, subjects discount the future more when the sooner draw is immediate, confirming clear present bias in the data. Another notable takeaway from Figure 5 is that when the sooner draw is immediate, subjects treat two-week, four-week, and eight-week delay lengths very similarly. They discount future tickets in all cases, but at a similar rate. When the sooner draw is two weeks in the future, this is not the case, and discounting always responds to delay length.

Given the importance of structural estimation in the literature and this paper, it is important to consider the identification properties of our elicitation. The goal of any price list elicitation is to identify tight bounds on the set containing the indifference point between two options. When subjects never switch (always choose the later draw) or switch immediately (always choose the sooner draw), that set is only bounded below or above. 91% of choices in our study are *interior* switch points, and only one subject failed to deliver a single interior switch point. As such, we prioritize structural estimation strategies that assume a price list switch point identifies a point of *equality* between the sooner and later options. We consider the robustness of our estimates to a set-identification approach in the next section.

[16] When a subject never switches –always preferring the later draw– we assume a switch midpoint of 245. When a subject switches immediately –always preferring the sooner draw– we assume a switch midpoint of 201. In Section 4..2 we find that allowing for a censoring process at the endpoints instead has no impact on our estimates.

[17] From this point forward referred to as the balanced sample, and used as the main sample of interest in our analysis

[18] Two-tailed t-test.

[19] OLS regression of d on the delay and an indicator variable for an immediate sooner draw, standard errors clustered at the subject level.

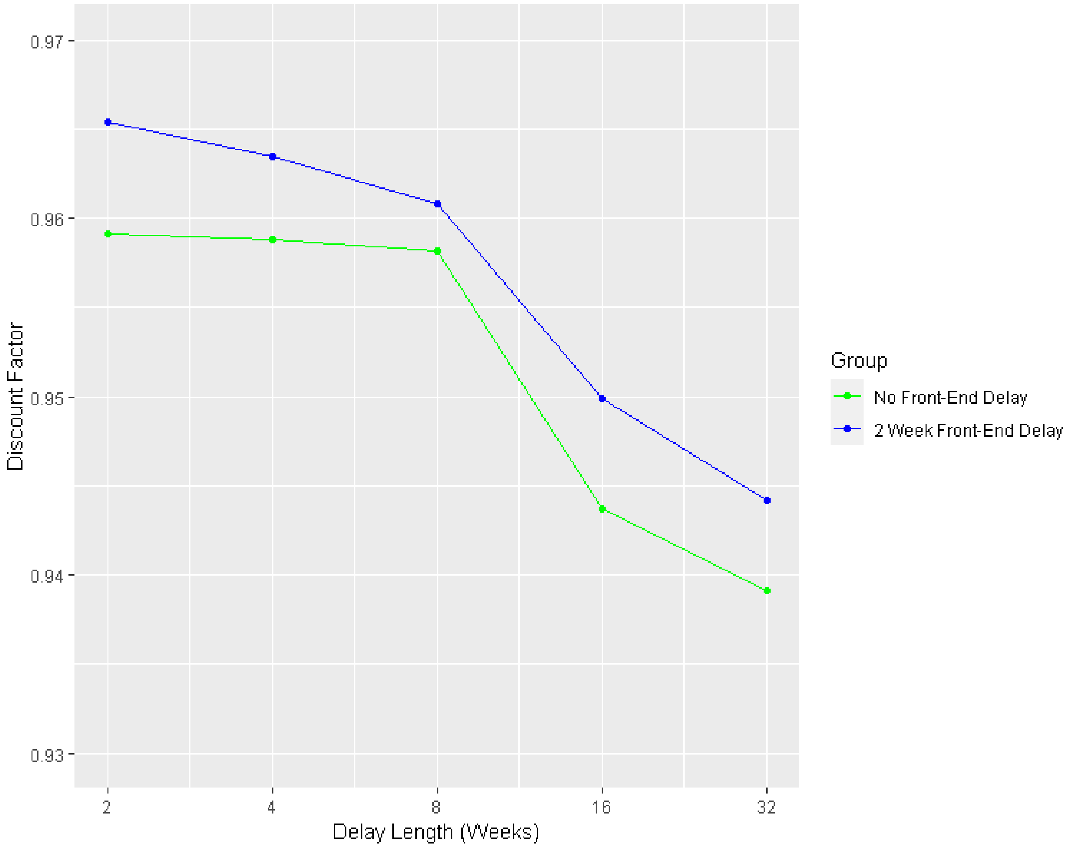

Figure 5: Average discount factor by choice delay

4..2 Aggregate discount parameters

A common application of time preference elicitations in the literature is to structurally estimate discounting parameters –typically an exponential δ parameter, and often a quasi-hyperbolic present-bias β parameter as well. Following our notation from Section 2 where the discount factor, $D(\tau, t)$, is a function of payoff period t and evaluation period τ, the quasi-hyperbolic form is $D(\tau, t) = \beta^{(\mathbb{1}(t>\tau)\}} \cdot \delta^{t-\tau}$. For our sooner and later options within each price list, we call t^S and t^L the payoff dates of the respective lotteries.[20] Assuming d identifies the point of indifference between the sooner and later draws in the price list, we have that

$$\beta^{(\mathbb{1}(t^S>\tau)\}} \cdot \delta^{t^S-\tau} \cdot (X-1) = \beta \cdot \delta^{t^L-\tau} \cdot 242 \Rightarrow d = \frac{\beta \cdot \delta^{t^L-\tau}}{\beta^{\mathbb{1}(t^S>\tau)} \cdot \delta^{t^S-\tau}} = \beta^{(\mathbb{1}(t^S=\tau)\}} \cdot \delta^{t^L-t^S}, \quad (5)$$

where X is the number of sooner tickets in the switching row (and thus $d = \frac{(X-1)}{242}$). We log-linearize this equality to obtain the regression equation

$$\ln\left(d_{ijw}\right) = \mathbb{1}\left(t_j^S = \tau_j\right) \cdot ln(\beta) + \left(t_j^L - t_j^S\right) \cdot ln(\delta) + \varepsilon_{ijw}, \quad (6)$$

which we estimate with standard errors clustered at the subject level. Estimates of δ and β from this approach are shown in column (1) of Table 6. Column (2) shows the corresponding estimates from an interval regression (a generalized Tobit model), which uses only the known bounds of each switch interval, including the one-sided bounds as the extremes. Column (3) shows estimate from a non-linear least squares (NLS) regression applied to equation (5). NLS will be our preferred technique for estimating the time-variant discount functions, which cannot be neatly linearized, so we include it here as well to show that it produces nearly identical estimates to the other models.

[20]Note that t^L is always greater than τ.

When price lists feature larger intervals or do not enforce a single switch point, researchers often take a binary-choice approach to these data at the choice-row level, and use maximum likelihood to estimate the discounting parameters. We take this approach with a logistic choice model, and present estimates in column (4).

Model:	OLS	Interval Regression	NLS	Logistic Choice
	(1)	(2)	(3)	(4)
δ	0.9980 (0.0003)	0.9979 (0.0003)	0.9981 (0.0003)	0.9983 (0.0003)
β	0.9711 (0.0032)	0.9712 (0.0033)	0.9723 (0.0031)	0.9778 (0.0031)
r (annual discount rate)	0.1098 (0.0150)	0.1146 (0.0174)	0.1066 (0.0148)	0.0901 (0.0145)
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\beta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$

Figure 6: Aggregate discount parameter estimates from quasi-hyperbolic model *Notes:* standard errors are clustered at the subject level. δ is the weekly exponential discount factor, and $r = \delta^{-52} - 1$ is the annual discount rate. The data consist of 4,345 price lists from 79 subjects.

Estimates for β and δ are nearly invariant to estimation method. Across the four models, our estimate of annual discount rate ranges from 9% to 11%. This is consistent with the lower end of estimates from the time preference elicitation literature, for example Andersen *et al.* (2008) and Andersen *et al.* (2014) find estimates of 10% and 9% in a nationally-representative sample of Danes, respectively. Andreoni and Sprenger (2012) estimates an annual rate of 37% using the Convex Time Budget technique with a US undergraduate population, and (Kuhn, Kuhn and Villeval, 2017) estimates a rate between 20% and 24% using the same technique among French undergraduates. Our estimates for β are all approximately 0.97, corresponding to present-bias that is roughly equivalent to an increase in delay length of 15 weeks. The estimates are precise; all are statistically significantly different from one with $p < 0.0001$.[21] Using the full sample of subjects, we obtain nearly identical results, shown in Appendix Table 7.

4..3 Individual discount parameters

When a time preference elicitation is bundled into the design of a larger research study, it is expected to deliver an individual-specific measure of discounting that is either a mediating variable or an outcome of interest. Using the OLS technique from equation (6), we successfully estimate individual-specific δ and β parameters for every subject in the sample. Figure 7 shows the distributions of estimates of r, the annual interest rate, and β.

The median r estimate is 9.69%, with 87% of subjects exhibiting positive discounting. The median estimate for β estimates is 0.9822. There is limited evidence of future-bias, with only 19% of subjects assigned $\beta > 1$[22] Overall, we conclude that our time preference elicitation was successful, delivering estimated utility parameters for all participants, with both the distributions of individual parameter

[21]The literature is mixed on whether incentivized discounting studies show evidence of present bias, with Andreoni and Sprenger (2012); Andersen, Harrison, Lau and Rutström (2014); Augenblick, Niederle and Sprenger (2015); Andreoni, Kuhn and Sprenger (2017); Kuhn, Kuhn and Villeval (2017) finding either no or minimal present bias over money in a laboratory setting, Augenblick, Niederle and Sprenger (2015); Imas, Kuhn and Mironova (2022) finding present bias over effort, Aycinena *et al.* (2022) finding future bias over large stakes in Guatemala, and Belot, Kircher and Muller (2021) finding present bias over lottery tickets.

[22]Results for the full sample can be found in Appendix Figure 16. Medians and proportions are qualitatively similar.

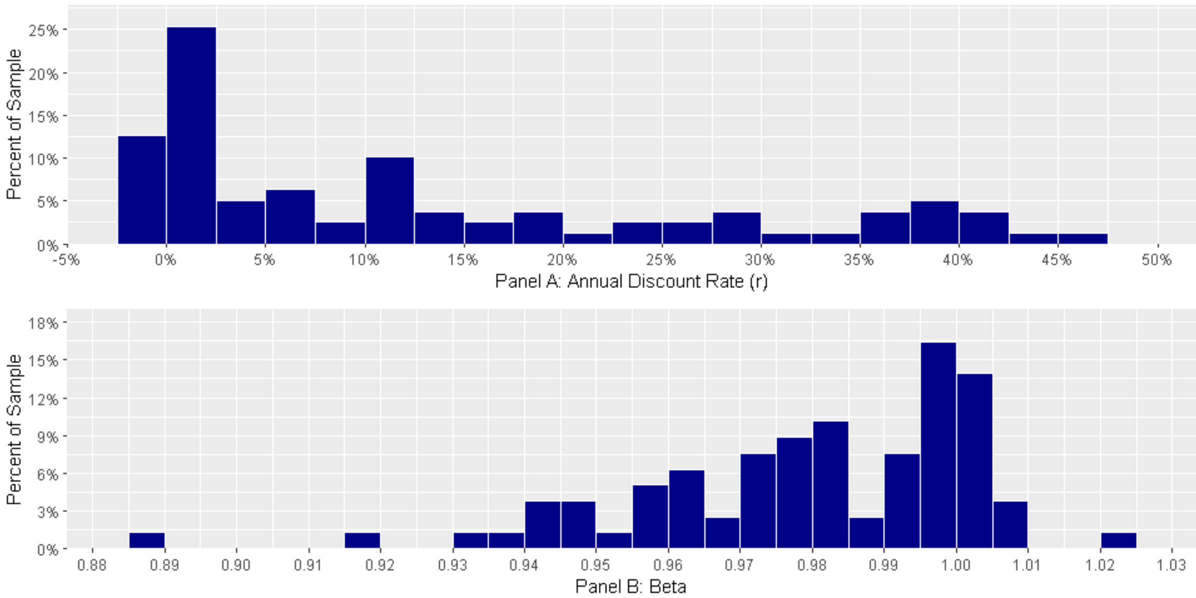

Figure 7: Distributions of individual estimates from the β-δ model

Figure 8: Time invariance as measured by survey-week fixed effects

estimates and aggregate estimates well within the ranges established by previous work. This establishes the platform from which we will examine time-varying properties of discounting.

5 Time-varying Behavior

5..1..1 Aggregate Time Variance

We now exploit the longitudinal nature of our data. In this section, we take a na"{i}ve reduced form approach before building up to structural models of time-variance in the next section. We start by augmenting the regressions we use to obtain descriptive statistics of discounting behavior in Section 4..1 with survey-week fixed effects. Specifically, we estimate

$$d_{ijw} = \mu_w + \lambda_1 \cdot \left(\left(t_j^L - t_j^S\right) - 8\right) + \lambda_2 \cdot \left(1 - \mathbb{1}\left(t_j^S = \tau_j\right)\right) + \eta_{ijw} \tag{7}$$

where we use $\left(\left(t^L - t^S\right) - 8\right)$ as the choice delay variable, and $\left(1 - \mathbb{1}\left(t^S = \tau\right)\right)$ as the front-end-delay variable so that the values of μ correspond to average discounting when comparing immediate tickets to tickets in eight weeks. Standard errors (η) are clustered at the subject level. Figure 8 shows these week fixed effects and their 95% confidence intervals.

We highlight three aspects of the data in Figure 8. First, subjects exhibit *decreasing patience* as the study progresses: we see substantially more discounting in week 12 (future utility discounted by 5.6%) compared to week zero (future utility discounted by 3.9%). In terms of an annualized (exponentially) discount rate, this is akin to a shift from 29% to 46% over the course of our study. If we replace μ_w in Equation (7) with a linear time trend, $\mu \cdot w$, we estimate that μ is negative (-0.0022) and statistically different than zero (p = 0.0008). Second, the time path of discounting is very similar for the full sample and for the sub-sample that completed all seven surveys, especially from week four (the third survey) onward. If anything, using the full sample appears to attenuate the trend of decreasing patience we observe among those who complete all the surveys. If we estimate linear time variance for the full sample we estimate that μ is −0.0013 (p = 0.0172). Third, we anticipated a detectable effect of final exam week in our study, and indeed there is a trend-break in survey week four: subjects place a higher value on future tickets in that week. The average discount factor in that survey is higher than in the previous survey, despite the overall decreasing trend (p = 0.0524). This is an example of time variance that is likely attributable to the influence of external factors.[23]

We estimate additional linear time trends models to check for time-variance that interacts with either the delay length between payoff periods in a choice, or whether the sooner date is immediate. Results are in Table 2[24]. We find robust evidence of across-the-board time variance, at a magnitude of 0.22 percentage points per week. There are no large or statistically significant interactions between the linear time trend and either choice delay or front-end delay in either sample.

	(1)	(2)	(3)	(4)	(5)
Constant	0.9554 (0.0043)	0.9698 (0.0046)	0.9697 (0.0047)	0.9697 (0.0047)	0.9697 (0.0048)
Choice delay ($(t^L - t^S) - 8$)	-0.0007^{***} (0.0002)	-0.0005^{***} (0.0002)	-0.0006 (0.0003)	-0.0005^{***} (0.0002)	-0.0006 (0.0003)
Front-end delay ($1 - \mathbb{1}(t^S = \tau)$)	0.0048^{***} (0.0014)	0.0020 (0.0014)	0.0020 (0.0013)	0.0022 (0.0029)	0.0022 (0.0029)
Survey week (w)		-0.0022^{***} (0.0006)	-0.0022^{***} (0.0005)	-0.0022^{***} (0.0006)	-0.0022^{***} (0.0006)
$w \cdot \left((t^L - t^S) - 8\right)$			-0.0000 (0.0000)		0.0000 (0.0000)
$w \cdot \left(1 - \mathbb{1}(t^S = \tau)\right)$				-0.0000 (0.0005)	-0.0000 (0.0005)

Table 2: Linear-time-trend estimates of time variance

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$. Coefficients are from linear models with standard errors are clustered at the subject level. The data consist of 4,345 price lists from 79 subjects.

A potential driver of differences in findings between the balanced and full samples is that measurements of time variance can be influenced by attrition. While subjects have a small likelihood of winning a prize each time they take a survey, the guaranteed, substantive payment for participating in the study comes after all surveys have been completed. When it comes time to participate in survey week two, that delayed reward is ten weeks in the future, and it gets two weeks closer in each successive survey. If impatient subjects attrit because they do not value the completion payment, we should expect to see a large initial decrease in discounting from survey weeks zero to two (because week zero is incentivized on its own), and then a stable pattern thereafter.[25] This is not the pattern of time variance we find in Figure 8. If selective patience-based attrition is affecting our results, it is masking even more decreasing patience than we observe. However, we do not find any relationship between discounting and attrition in our study. Figure 9 shows how the discount factor within each survey correlates with whether a subject does not complete the following survey. None of the estimates are statistically significant (p = 0.771, 0.541, 0.354, 0.503, 0.508, and 0.597, for Weeks 0-10 respectively), and the sign is not consistent across surveys.[26]

[23]This effect is driven not by subjects who fail to complete the finals-week survey, but it is stronger for subjects who will eventually attrit *after* completing it. This is why the trend break is more notable for the full sample.

[24]Corresponding estimates for the full sample are in Appendix Table 8

[25]Purely from a patience perspective, if week two is worth completing for the reward in ten weeks, then week four is worth it for the reward in eight weeks.

[26]Estimates correspond to equation (7) estimated at the survey level, with the survey fixed effect replaced with an indicator for whether a subject returns to participate in the next survey.

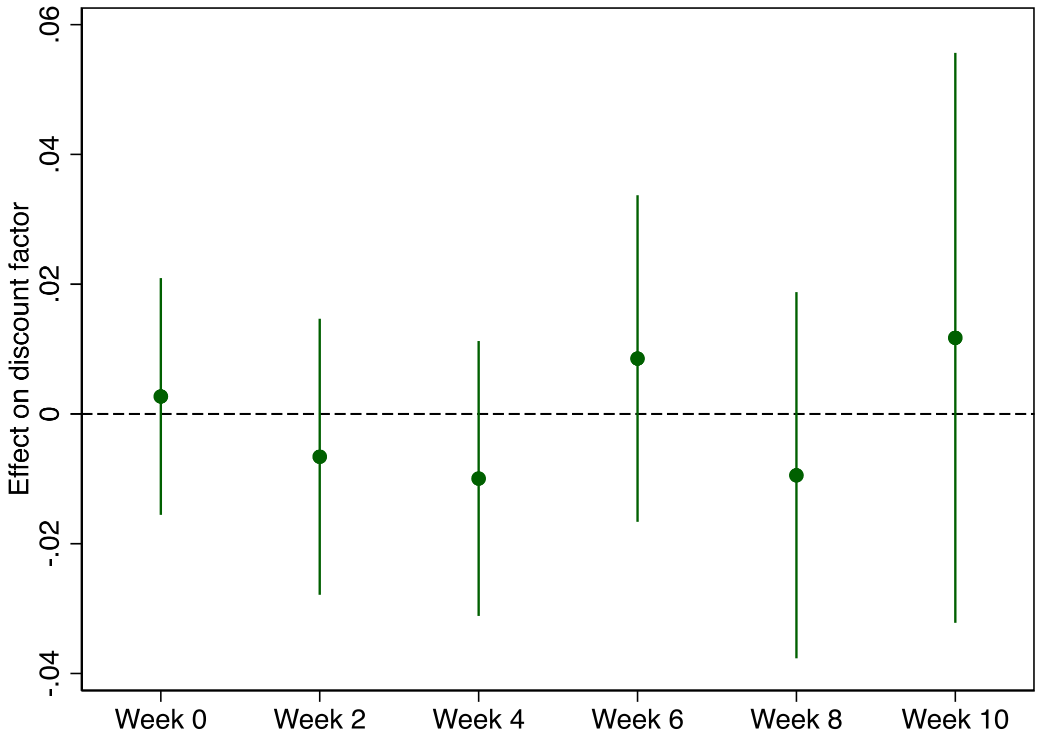

Figure 9: Relationship between subsequent survey attrition status and discounting, point estimates and 95% confidence intervals

Notes: the plot shows the coefficients and 95% confidence intervals from a regression of the within-survey discount factor on an indicator for whether a subject fails to complete the subsequent survey. We control for the delay-length and front-end delay of each choice so that the discount factor corresponds to a tradeoff between the present and eight weeks in the future.

5..1..2 Subject-level Time Variance

A key strength of our data collection is our ability to observe subject-level changes in discounting over time. A reduced-form approach to this is to modify Equation (7) to feature subject-week fixed effects, μ_{iw}, rather than pooled week fixed effects. Figure 10 below shows the distribution of these fixed effects by subject across weeks. Across all survey weeks, modal discounting looks very stable, while the right tail expands consistently throughout the study.

While these distributions do not track individual subjects, it suggests a group of time-invariant discounters and a group of decreasing-patience discounters. To confirm this, we measure individual time trends using linear, individual-specific survey week time trends, replacing μ_w in Equation (7) with $\rho_i \cdot w$, while also allowing for fixed level differences across individuals using subject fixed effects, μ_i. Figure 11 plots the ρ_i coefficients along with their standard errors[27]. We can reject time-invariance for 57% of subjects, with 37% and 20% exhibiting decreasing and increasing patience, respectively. Note that conditional on exhibiting time variance, the magnitude of decreasing patience is larger on average than the magnitude of increasing patience, leading to the overall trends shown in Figures 8 and 10.

5..2 Inconsistency and Non-Stationarity

In Section 2, we illustrated how time-varying behavior implies either non-stationary or inconsistent behavior (or both). We have also empirically established both decreasing patience and inconsistent/

[27] $N = 72$. 7 subjects had insufficient variation to estimate coefficients

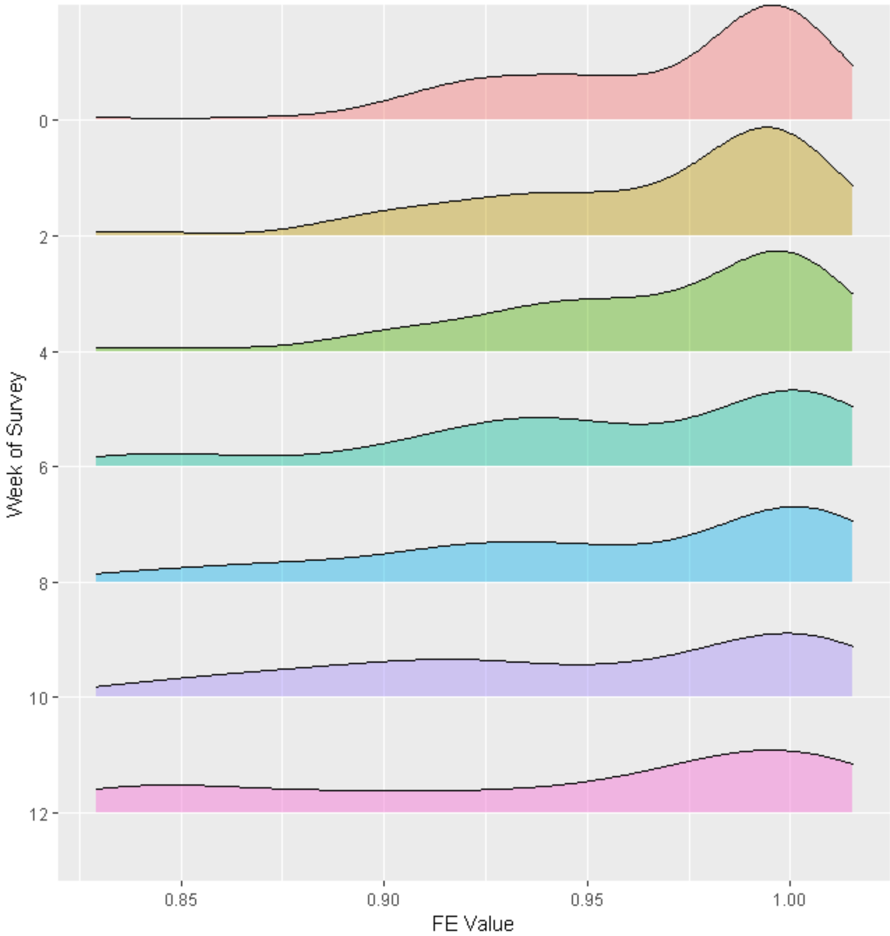

Figure 10: Distributions of subject-week fixed effects, by survey week

Notes: the plot shows the distribution of subject-week fixed effects from a regression of the discount factor on those fixed effects and controls for the delay-length and front-end delay of each choice so that the discount factor corresponds to a tradeoff between the present and eight weeks in the future.

non-stationary behavior in the data –we have not yet drawn a distinction between the two because under the assumption of time-invariance, they are constrained to be the same. In this section, we allow them to differ and explore the connection between the two; how much observed inconsistency and non-stationarity is explained by time-variance?

We start by separately estimating non-stationarity and inconsistency. Equation (7), models the impact of choice delay (t_j) and front-end delay ($t_j - \tau_j$) with survey-week fixed effects. We adjust these variables such that the constant term describes the discount factor for an eight-week choice delay without a front-end delay. The fixed effects limit the model to within-survey variation in $t_j - \tau_j$, corresponding to a stationarity violation. In columns (1) and (2) of Table 3, we present the survey-week fixed effect models, with and without allowing an interactive effect between choice delay and front-end delay, respectively. We do not estimate a statistically significant relationship between a *within-survey* front-end delay and the discount factor ($p = 0.1406$ and $p = 0.1036$, respectively), although the coefficient is positive, qualitatively consistent with a present-bias. The estimates in column (2) show that for shorter choice delays, we do not observe significantly more stationarity-violating present bias, although we do in the full sample. Results in Appendix Table 9.

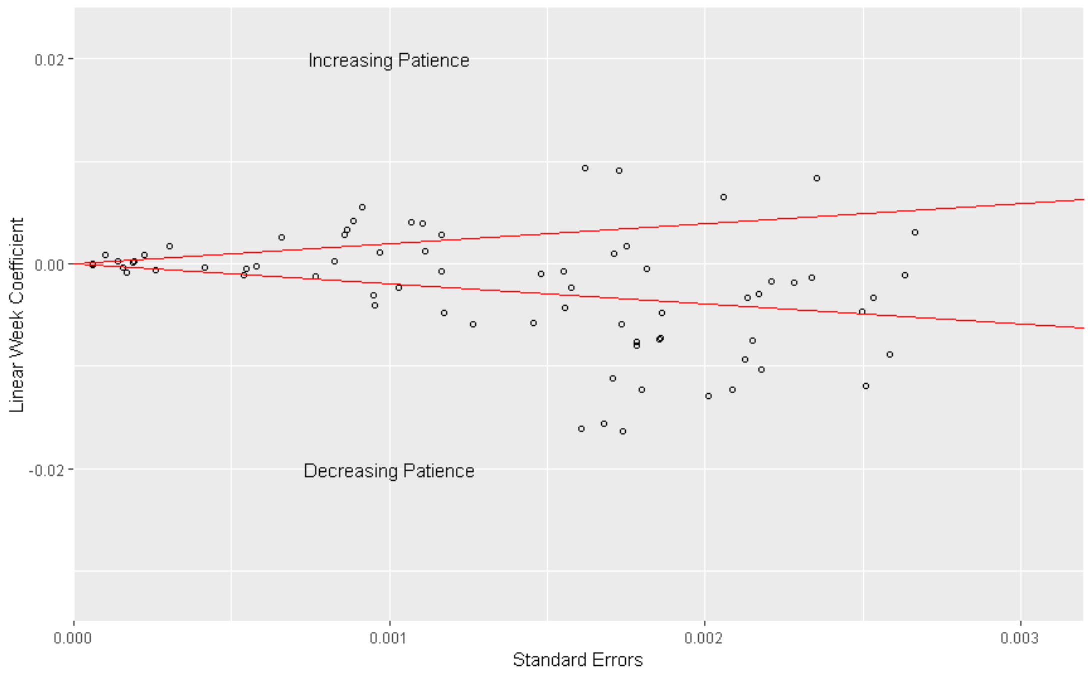

Figure 11: Subject-level estimates of time variance

Violation type:	Stationarity		Consistency	
	(1)	(2)	(3)	(4)
Constant	0.9698 (0.0046)	0.9698 (0.0046)	0.9698 (0.0046)	0.9698 (0.0046)
Choice delay $((t^L - t^S) - 8)$	-0.0005^{***} (0.0002)	-0.0005^{**} (0.0002)	-0.0005^{***} (0.0002)	-0.0005^{**} (0.0002)
Front-end delay $\left(1 - \mathbb{1}(t^S = \tau)\right)$	0.0020 (0.0014)	0.0064^{***} (0.0013)	0.0065^{***} (0.0016)	0.0063^{***} (0.0016)
$\left((t^L - t^S) - 8\right) \cdot \left(1 - \mathbb{1}(t^S = \tau)\right)$		-0.0001 (0.0001)		-0.0001 (0.0001)
Survey-week FEs	Y	Y	N	N
Drawing-week FEs	N	N	Y	Y

Table 3: Separate estimates of non-stationarity and inconsistency

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors are clustered at the subject level. The data consist of 4,345 price lists from 79 subjects.

To estimate inconsistency instead, we replace the week effects with *drawing-week* fixed effects in columns (3) and (4) of Table 3. The drawing week of a decision is the week the *earlier* drawing takes place. For example, decisions in the first survey with a two-week front-end delay have the same drawing week as decisions in the second survey without the front-end delay. We find that making a choice two weeks in advance results in a significantly higher discount factor ($p < 0.0001$ and $p = 0.0001$, respectively). These findings are consistent with Section 2, where time-inconsistent present bias is a direct result of decreasing patience under the assumption of stationarity.

Given that inconsistency and non-stationarity differ in their magnitude and frequency, time-variance must play a role in explaining these phenomena. We follow Janssens, Kramer and Swart (2017) in asking how many of the instances of inconsistency and non-stationarity in our study can be attributed to time-variance. Referring to the decision triangle language of Section 2, given any level of time invariance ($a - c$), we can calculate whether any observed inconsistency ($b - c$) is exactly predicted by, in excess of, or less than what would be predicted assuming stationarity holds ($a = b$). We can do the same for any observed non-stationarity as well. Specifically, for all decision triangles where inconsistency exist ($b \neq c$), we calculate $\frac{a-c}{b-c}$. This is effectively the percentage of time inconsistency that is attributable to time variance. *Time variance explains 72% of the inconsistency in our sample.* While inconsistency is typically attributed to a behavioral present-bias (when $c < b$) in decisions, this result suggests that a large part of it could be explained by decreasing patience ($c < a$). For all decision triangles where non-stationarity is present ($a \neq b$), we construct a similar ratio of $\frac{a-c}{a-b}$. This ratio describes the percentage of non-stationarity that can be attributed to time variance. Time variance explains 56% of non-stationarity in our sample.

5..3 Time Preference Classification

With evidence of time variance, inconsistency, and non-stationarity in our sample, we now follow Halevy (2015) and attempt to classify each subject into one of the five possible combinations of time preference properties illustrated in Table 1. We start by classifying each applicable decision within a triangle as either time-invariant or time-varying. This requires that subjects must have attended at least two sequential surveys to be included in this analysis. Out of 153 total subjects, this applies to 141[28]. These 141 subjects account for 6,267 of the 6,367 price lists collected in our study. containing 2,506 decision triangles. Of these 141 subjects, 130 subjects (92%) make a time-varying decision at some point in the study. Of the 2,506 decision triangles, 1,500 (60%) include a time-varying decision[29]. Both the percentage of subjects and percentage of choices that are time-varying exceed the findings of Halevy (2015), that found a maximum of 56.42% of subjects exhibiting time-varying preferences. While the longer longitudinal dimension of our study contributes to that difference at the subject level, this should not be the case at the triangle level. We return to this issue later in this section.

We count the occurrences of inconsistent and non-stationary decisions in a similar manner to the time-varying decisions. While the same 141 subjects and 6,267 price lists generate 2,710 decision triangles where b and c the test for inconsistency is possible and 2,782 decision triangles where the test for non-stationarity is possible, we restrict our analysis to the 2,506 decision triangles that overlap with the preceding time-invariance analysis. We find that 129 subjects (91%) make an inconsistent decision at some point, with 1,448 (58%) of triangles being inconsistent. Similarly, 127 subjects (90%) make a non-stationary decision at some point, with 1,276 (51%) triangles in total being non-stationary. Both proportions exceed those from Halevy (2015) again, where at most 66.33% of subjects exhibited inconsistent preferences and at most 55.22% exhibited non-stationary preferences.

Table 10 shows the classification of subjects by property in our sample in column (1) of Panel A. Only eight subjects (10%) exhibit invariant preferences. Of these eight, seven (9% of the sample) make consistent and stationary decisions. However, these nine subjects display constant discounting, not reacting to differences in choice delay or front-end delay either. Thus, among subjects who reacted to any stimuli, all violated a time preference property. 69 (87%) subjects violate each property at some point, two (2%) maintain only stationarity, and none maintain only consistency[30]

A key caveat to these statistics is that the prevalence of violations in our sample can be, in part, attributed to the number of opportunities we offer to record an error in decision-making within a 22-

[28]This eliminates ten subjects who only responded to the first survey and two more subjects who only responded to the first and third.

[29]This proportion is very similar if we also consider time-variance in the choices made *with* a front-end delay. In this case, 94% of subjects display time-varying behavior in 59% of triangles.

[30]These results are qualitatively invariant to the inclusion of the full sample, seen in Appendix Table 10.

row price list. As we will show in the next section, true errors in decision-making can be down-weighted by the underlying trend of one's decisions. By contrast, similar studies with only two points of contact have only one comparison to draw conclusions from, meaning what could be an error in decision-making may instead be treated as an out-right violation of one of the properties. Our study design allows for formal modeling and testing of each property at the individual level. However, we first address this issue with two ad-hoc approaches that allow us to continue to compare our results to other studies. First we allow for a margin of error that spans a standard deviation (14 tickets) of the decisions we observe (pm 6 tickets due to the two-ticket gap between each choice row). If a subject's choice is within six tickets on either side of their previous choice, we treat the decision as if it did not violate the tested time preference property.[31] Results are in column (2) of Panel A in Table 4. Second, instead of aggregating to the subject-level, we treat each decision triangle as a distinct observation. Results are in column (1) of Panel B, and results using both modifications are in column (2) of Panel B.

Margin of Error:	None	$\pm$ 6 tickets
Properties of choices	(1)	(2)
Panel A: Subject level		
A) Invariant, consistent, stationary	8.86%	20.25%
B) Invariant only	1.27%	0%
A) + B)	10.13%	20.25%
C) Stationary only	2.53%	3.80%
D) Consistent only	0%	0%
E) None	87.34%	75.95%
C) + D) + E)	89.87%	79.75%
Panel B: Triangle level		
A) Invariant, consistent, stationary	37.45%	63.40%
B) Invariant only	6.07%	7.01%
A) + B)	43.52%	70.41%
C) Stationary only	15.24%	14.40%
D) Consistent only	7.86%	7.96%
E) None	33.39%	7.23%
C) + D) + E)	56.48%	29.59%

Table 4: Subject- and decision-triangle-level classification of time preference properties

Notes: Data 79 subjects that completed all surveys. Column(1) assumes any deviation from is a time preference property violation. Column(2) assumes any deviation greater than six tickets from an analogous choice is a time preference violation, capturing a range of a standard deviation around a choice ($\frac{\sigma}{2} \approx 7$, which due to interval of two tickets between choice rows corresponds to equality within six tickets).

Using the margin of error at the subject level, 16 subjects (20%), exhibit time-invariant behavior. All 16 of these subjects are also consistent and stationary. The increase in the size of this group comes mostly from a reduction in size of those that violate every property, now 60 subjects (76%). Three subjects maintain only stationarity (4%) and still none maintain only consistency (1%) using this window of error. Taking the decision-triangle level of analysis approach instead we find that around 43% of decision triangles maintain time-invariance, while 56% do not. This is close to the minimum

[31] One problem with adding a margin of error to this accounting exercise is that it is possible for only one property to be violated, even though theoretically, when two hold, the third is guaranteed. When this is the case, we classify a decision triangle or subject as satisfying all time preference properties.

classification from Halevy (2015), which has 44% of subjects exhibiting time-invariance. Our most conservative estimates of time-varying behavior come from combining the decision triangle sorting with the margin of error. This method shows 70% of decision triangles maintain time-invariance, which is close to the maximum classification from Halevy (2015), which has 68% of subjects exhibiting time-invariance. Overall, these results suggest that observing a single decision triangle per individual understates the potential for time-varying behavior. Defined strictly, time invariance is a universal feature of subjects who responded to our stimuli.

6 Estimating Time-varying Discount Functions

Beyond establishing the prevalence of time variance and its role in explaining time-inconsistency and non-stationarity in our sample, our goal in this paper is to develop a structural toolkit for estimating time variance as a part of the discount function, and formally testing whether time preference properties hold at the individual level. Recall the "nested exponential" discount function we introduced in Section 2:

$$D(\tau, t) = \delta^{\frac{(t^{\mu} - \tau^{\mu})^{\eta}}{\gamma^{\tau}}}.$$

Assuming $\delta \neq 1$, when all three (none) of the other parameters differ from one, the function fulfills none (all) of the three time preference properties. If any pair of μ, η, and γ equal one, then only one of three properties holds. In the case where all properties holds the function nests classical exponential discounting, and in the case where only invariance holds ($\mu = \gamma$ = 1, $\eta \neq 1$, e.g. the "behavioral" scenario of present bias) the functions nests something akin to Weibull discounting.

We estimate parameters for the nested exponential discount function at the aggregate sample level using the non-linear least squares method, separately allowing for each possible class of time preference properties. Results and descriptive statistics of these estimations are presented in Table 5. Column (1) enforces exponential discounting ($\mu = \eta = \gamma$ = 1, thus $D(\tau, t) = \delta^{t-\tau}$). The weekly discount factor of 0.9973 translates into an annual discount rate (r_0) of 15%.

Restrictions:	$\mu=\eta=\gamma=1$	$\mu=\gamma=1$	$\mu=\eta=1$	$\eta=\gamma=1$	None
Properties					
Invariant	✓	✓	X	X	X
Stationary	✓	X	✓	X	X
Consistent	✓	X	X	✓	X
	(1)	(2)	(3)	(4)	(5)
δ	0.9973	0.9779	0.9973	0.9787	0.9756
	(0.0003)	(0.0027)	(0.0004)	(0.0045)	(0.0039)
η	1	0.3855	1	1	0.4250
	.	(0.0326)	.	.	(0.0452)
γ	1	1	1.0012	1	0.9703
	.	.	(0.0164)	.	(0.0176)
μ	1	1	1	0.5141	0.7321
	.	.	.	(0.0494)	(0.1380)
Log-likelihood	$-18,149.6$	-17921.5	$-18,149.6$	$-18,068.5$	$-17,917.2$
AIC	$36,301.2$	$35,846.9$	$36,303.2$	$36,141.0$	$35,842.4$
r_0	0.1499	0.1082	0.1513	0.1786	0.0880
	(0.0177)	(0.0130)	(0.0250)	(0.0192)	(0.0138)
r_{12}	$=r_0$	$=r_0$	0.1491	0.1113	0.1178
			(0.0222)	(0.0149)	(0.0174)
PB_S	1	0.9755	1	0.9771	0.9791
	.	(0.0030)	.	(0.0052)	(0.0116)
PB_C	1	$=PB_S$	1.0001	1	0.9744
	.		(0.0007)	.	(0.0032)
$[H_0]:\delta=1$	$p<0.0001$	$p<0.0001$	$p<0.0001$	$p<0.0001$	$p<0.0001$
$[H_0]:\eta=1$	.	$p<0.0001$	.	.	$p<0.0001$
$[H_0]:\gamma=1$	.	.	$p=0.9441$	.	$p=0.0955$
$[H_0]:\mu=1$	.	.	.	$p<0.0001$	$p=0.0558$
$[H_0]:\mu=\gamma=1$	.	.	.	.	$p=0.1385$
$[H_0]:\mu=\eta=1$	.	.	.	.	$p<0.0001$
$[H_0]:\eta=\gamma=1$	.	.	.	.	$p<0.0001$
$[H_0]:\mu=\eta=\gamma=1$	.	.	.	.	$p<0.0001$
$[H_0]:r_0=r_{12}$	.	.	$p=0.9440$	$p<0.0001$	$p=0.1295$
$[H_0]:PB_S=1$	.	$p<0.0001$	.	$p<0.0001$	$p=0.0703$
$[H_0]:PB_C=1$	.	$p<0.0001$	$p=0.9443$	.	$p<0.0001$
$[H_0]:PB_S=PB_C$	.	.	.	.	$p=0.6666$

Table 5: Aggregate parameter estimates for the nested exponential model

Notes: standard errors are clustered at the subject level. $r_\tau = D(\tau,\tau+52)^{\{-1\}}-1$ is a measure of the annual discount rate. PB_S and PB_C are stationarity-based and consistency-based measures of present-bias. The data consist of 4,345 price lists from 79 subjects that completed all seven surveys. All estimates are from non-linear least squares regressions.

The model in column (2) allows for time-inconsistency and non-stationarity, but enforces time-invariance, our Weibull-like specification ($\mu = \gamma$ = 1, thus $D(\tau, t) = \delta^{(t-\tau)^{\eta}}$). This model fits the data much better than the exponential model, even taking its extra parameter into account, as shown by the decrease in the Akaike information criteria (AIC). This method for assessing model fit penalizes the more complex models for the higher degrees of freedom they have in fitting the data. The estimate of δ = 0.9779 (the weekly discount factor for a choice delay of one week) suggests much more rapid discounting of near-term future outcomes, but the "present-bias" parameter η flattens the discount function for longer choice delays. η = 0.3811 turns a choice delay of two weeks into an effective delay of 1.3 weeks, and a 52-week choice delay into an effective delay of 4.5 weeks, such that the annual discount rate is lower than in the exponential model: 11%.[32] Qualitatively, this means we estimate significant hyperbolicity in discounting, clearly rejecting $\eta = 1$ ($p < 0.0001$). To asses the magnitude of present-bias we estimate in this model, we construct measures PB_S and PB_C, which are discount factor ratios that correspond to measurements of the quasi-hyperbolic β through a stationarity violation or a consistency violation, respectively. Specifically,

$$\mathrm{PB}_S = \frac{D(0,8)}{D(0,0)} \Big/ \frac{D(0,10)}{D(0,2)}, \quad \text{and} \quad \mathrm{PB}_C = \frac{D(2,10)}{D(2,2)} \Big/ \frac{D(0,10)}{D(0,2)}$$

When invariance holds these measures are identical, and we estimate PB_S = PB_C = 0.9755, which is significantly different from one ($p < 0.0001$).[33]

The model in column (3) allows for time-variance and inconsistency, but enforces time-stationarity, the property specification that does not appear in the literature ($\mu = \eta$ = 1, thus $D(\tau, t) = \delta^{\frac{t-\tau}{\gamma^{\tau}}}$). While this model represents a modest improvement in fit than exponential discounting in column (1), the fit is much worse than the time-invariant model in column (2). Without access to hyperbolicity, the estimates attempt to fit both short and long choice delays with an initial (τ = 0) weekly discount factor of 0.9973, which translates to an annual rate (r_0) of 15%. Time variance due to $\gamma = 1.0012$ marginally reduces that weekly discount factor as the study progresses. The annual discount rate at the end of our study (τ = 12, termed r_{12} in the table), the annual discount rate falls to just under 15%. In an attempt to best-fit the data, this model estimates statistically insignificant *increasing-patience* time variance (p = 0.9440 for the test of $r_0 = r_{12}$, p = 0.9441 for the test of γ = 1), at odds with the reduced form results. A side effect of this is that the model predicts statistically insignificant *future-bias* (p = 0.9443 for the test of $\mathrm{PB}_C = 1$); this is a direct implication of increasing patience under the assumption of stationarity. These results are more pronounced, and statistically significant for the full sample. Results in Appendix Table 5.

The model in column (4) allows for time-variance and non-stationarity, but enforces time-consistency, the "time-consistent hyperbolic" property pattern studied in Strulik (2021), but in a more flexible form ($\eta = \gamma$ = 1, thus $D(\tau, t) = \delta^{t^{\mu} - \tau^{\mu}}$). This model lies in between the poor performance of the exponential and stationary-only models in columns (1) and (3), and the improved fit of the invariant-only model in column (2). The initial (τ = 0) weekly discount factor estimate is 0.9787 which captures substantial discounting of near-term future outcomes, with the hyperbolicity parameter μ flattening the discount function for longer choice delays, much like η in the invariant-only model. μ = 0.5141 turns the choice between zero and two weeks into an effective delay of 1.4 weeks, and the choice delay between zero and 52-weeks into an effective delay of 6.7 weeks, yielding

[32] Without the hyperbolic "flattening" of the discount factor, a weekly discount factor of 0.9778 exponentially implies an annual discount rate of 221%.

[33] While it lies outside the nested exponential model, it is worth pointing out that the data clearly reject the quasi-hyperbolic model in favor of this Weibull-like restricted model. Hyperbolicity is a key feature of our data; indeed the single parameter true hyperbolic discount function $D(\tau, t) = \frac{1}{1+\alpha \cdot (t-\tau)}$ fits substantially better than the single-parameter classic exponential model.

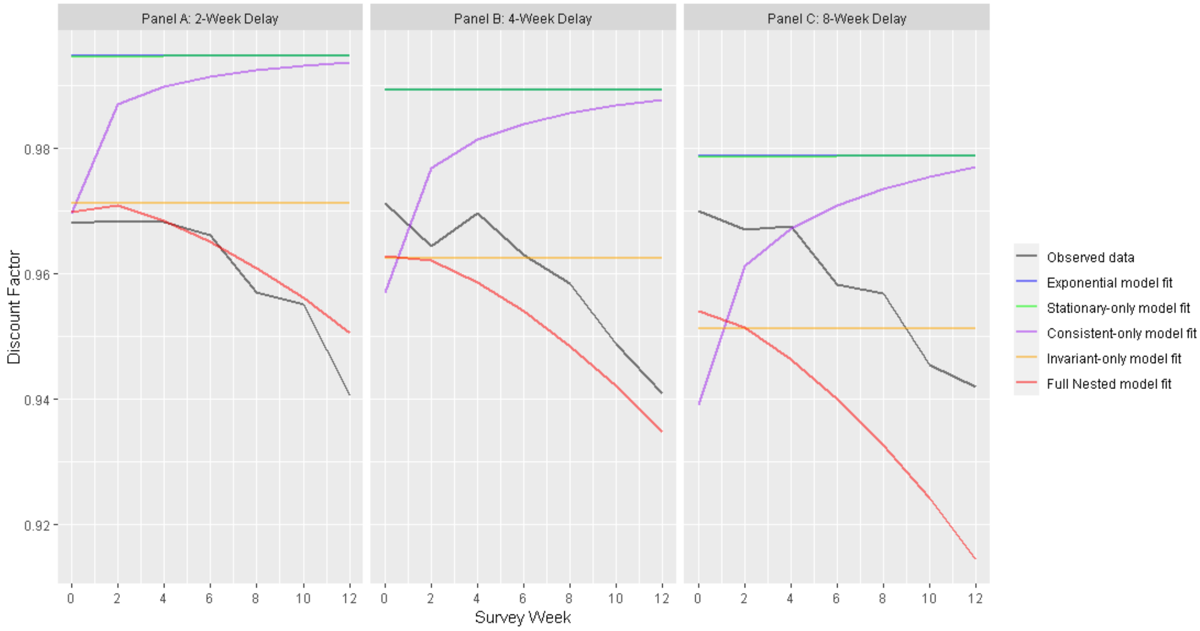

Figure 12: Observed and estimated discount factors by delay length

an initial ($\tau = 0$) estimate of the annual discount rate of 18%. Fitting hyperbolicity in a model constrained to be time-consistent forces it to predict *increasing-patience* time variance, just like the stationary-only model, however this model finds the increasing patiences statistically significant. The end-of-study ($\tau = 12$) annual discount rate predicted by this model is 11%, which we can reject is equal to the initial rate ($p < 0.0001$). The μ parameter acts on τ as well as t, thus $\mu < 1$ delivers both hyperbolicity and *increasing hyperbolicity* in τ. While the effective choice delay between zero and 52 weeks is 6.7 weeks, the effective choice delay between 12 and 64 weeks is 4.1 weeks. This model cannot deliver both decreasing patience and hyperbolicity, which explains its poor fit. As a side effect of increasing patience under the assumption of consistency, this model predicts statistically significant present-bias ($p < 0.0001$ for the test of $\text{PB}_S = 1$).

The model in column (5) puts no restrictions on the parameter values. It features the best overall fit of the data, but crucially, its AIC is nearly identical to that of the invariant-only model in column (2). We estimate an initial weekly discount factor of 0.9756, and an initial annual discount rate of 9%. This grows to 12% over the course of the study, reflecting the decreasing patience we identified in the reduced-form analysis, but we can only marginally reject that the two are the same ($p = 0.1295$). However, we can marginally reject the assumptions of the invariant-only model. We can reject $\gamma = 1$ ($p = 0.0955$) and that $\mu = 1$ ($p = 0.0558$). We can also nearly reject that $\gamma = \mu = 1$ ($p = 0.1385$). This model predicts a larger degree of consistency-violation present-bias ($\text{PB}_C = 0.9744 < 1$, $p = 0.0032$) then stationarity-violation present-bias ($\text{PB}_C = 0.9791 < 1$, $p = 0.0032$), although we cannot reject that they are equal in magnitude ($p = 0.6666$). We ultimately find that this fully estimated nested model, is best fitting. The relative likelihood of the invariant-only model to the fully estimated nested model is 0.11, whereas it is essentially zero for the other three models.[34]

Figure 12 shows the discount factors predicted by the estimated functions along with observed discount factors (averages of d from the data) by delay length (for choice delays of two, four, and eight weeks, in Panels A, B, and C, respectively) and week of survey (x-axis). Visualizing the model fit this way helps demonstrate how the nested and invariant-only functions performed better in

[34]This measure ranges from zero to one, and assesses the relative explanatory of two models. Using AIC, the relative likelihood is defined as $\exp\left(\frac{\text{AIC}_1 - \text{AIC}_2}{2}\right)$ where AIC_1 is better-fitting model (lower AIC).

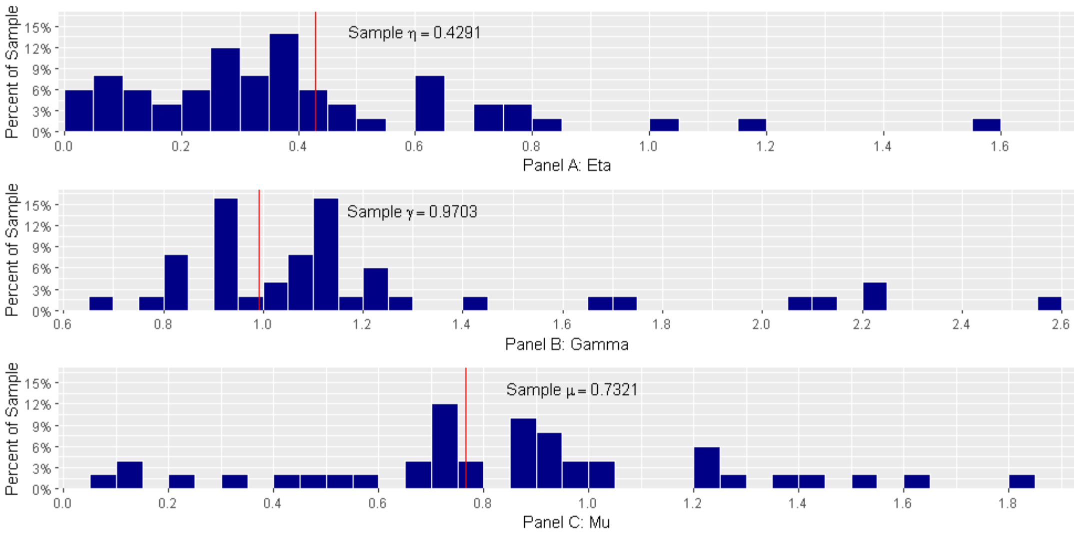

Figure 13: Individual parameter estimates for the unrestricted nested exponential model

Notes: 9 and 8 estimates, respectively, are suppressed from the right tails of the distributions in Panel (B) and Panel (C) for visual clarity

fitting the data, and what they still fail to capture. The key feature of all three panels is that only the fully estimated model can mimic the decreasing patience shown by subjects. The invariant-only model performs well despite not predicting decreasing patience because –like the nested model– it adapts well to changes in choice delay lengths. The other three models –exponential, stationary-only and consistent-only– not only cannot match the pattern of decreasing patience, but also change very little in response to changes in delay length, resulting in poor fit relative to the invariant-only and nested model.[35]

6..1 Subject-level estimates

The aggregate results suggests that marginal time-variance is present, but that when it comes to structurally fitting the preferences of a representative agent, decreasing patience is not as important as other features of the data –notably hyperbolicity– for a model to fit. If a researcher were to opt for an invariant-only discount function to model the preferences of this sample, predictions would be accurate in general, but could underestimate the level of discounting that would occur in future periods. However, even if a time-invariant model may be satisfactory to describe the sample, the literatures on time-preference estimation and the importance of present-bias are focused on individual heterogeneity in preferences. If 20% of people are quasi-hyperbolic discounters, we may not estimate a representative discount function with β significantly less than 1, but we still want to correctly identify that sample. In a time-varying world, this necessitates allowing for heterogeneity in time-variance as well.

We attempt to estimate the parameters of the nested exponential model for each of the 79 subjects that completed every survey, once again using the non-linear least squares method. Given lack of variation in responses for some subjects, we successfully estimate parameters for all models for 50 subjects. As a point of reference to aggregate results, we plot the distribution of parameter estimates for the fully estimated nested model in Figure 13. Again using AIC we determine which model best fits each individual in sample, with the results in Table 6.

Approximately 42% of subjects are best fit by time-invariant models. While this is consistent with the lower end of the range found by Halevy (2015) (44%), the distribution between exponential discounters and invariant-only discounters that we find is very different; no time-invariant discounters in our study are best fit by the standard exponential that also allows for consistency and

[35]These models fit the level well in the 16-week choice delay, but then badly over-react to the 32-week choice delay.

stationarity. The *lowest* estimate of that proportion in Halevy (2015) is 77%. Similar to the reduced-form methods of sorting used in the previous section, we find few subjects are best-described as being stationary-only or consistent-only. These results hold qualitatively in the full sample, where we successfully estimate parameters for all models for 91 subjects. One key difference is that 14% of time-invariant discounters, six subjects, are best represented by the standard exponential discount function. Analogous estimation and sorting results can be found in Appendix Figure 17 and Appendix Table 12

Properties of choices	
A) Invariant, consistent, stationary	0%
B) Invariant only	42%
A) + B)	42%
C) Stationary only	4%
D) Consistent only	6%
E) None	48%
C) + D) + E)	58%

Table 6: Subject-level structural classification of time preference properties
Notes: Data comes from 50 subjects for which we were able to obtain parameter estimates from all models and completed every survey.

7 Discussion

We use a 12-week longitudinal study to detect violations of three related time preference properties across seven surveys. We find that in all but our most conservative classification approach, the largest single group of subjects by adherence to the TICS properties is "none". A minimum of 30% and a maximum of 90% of subjects exhibit time-variance, a pattern of behavior that confounds the identification of time preferences in traditional time-invariant models. Average discounting exhibits systematically decreasing patience over the course of our study. Without adjusting for time-variance, this would've led us to dramatically overstate the degree of behavioral time-inconsistency in our study (e.g. present-bias). Indeed, 72% of the time-inconsistency we observe in our data is explained by time variance. This is yet another reason that the term "present-bias" to describe decision-making may be ill-founded when not accounting for time-varying behavior.

Beyond highlighting this measurement issue, we developed a new discount function –the nested exponential– that allows researchers to estimate a model that is fully general with respect to the TICS properties, yet tractable enough to deliver subject-level parameter estimates in a simple non-linear least squares estimation. This model is easy to pick up and estimate in a longitudinal research study that features all three choices in a decision triangle, but also it can be used in a calibration to explore the robustness of estimates of time-inconsistency or non-stationarity derived from only two of the three choices.

The presence and acknowledgment of time-varying preferences has potential to impact policy, particularly relating to commitment mechanisms and behavioral nudges. If a policy is meant to reinforce commitment to a consumption plan, take-up of the commitment mechanism may be lower than expected when time-varying preferences exist and are internalized by the policy target. If the target understands that their preferences are changing over time, committing to a previous decision would be welfare-damaging. Similarly, the effect of behavioral nudges may be overstated if the target has time-varying preferences that coincide with the direction the policy is meant to push them.

A deep exploration of the sources of time-varying behavior is left for future work using the data we have collected. Past studies on this topic focus on forces of risk, uncertainty, and liquidity-constraint as likely causes of this behavior. More thorough analysis of these potential factors, as well as the impact of demographics, socioeconomic status, cognition, and macroeconomic changes, will help in determining in what contexts time-varying behavior may be present and to what extent it affects outcomes or can be controlled for.

Bibliography

Andersen, S. *et al.* (2008) “Eliciting Risk and Time Preferences,” *Econometrica*, 76(3), pp. 583–618

Andersen, S. *et al.* (2014) “Discounting behavior: A reconsideration,” *European Economic Review*, 71, pp. 15–33

Andreoni, J. and Sprenger, C. (2012) “Estimating Time Preferences from Convex Budgets”, p. 50–51

Andreoni, J., Kuhn, M. A. and Sprenger, C. (2017) “Measuring time preferences: A comparison of experimental methods,” *Journal of Economic Behavior and Organization*, 116, pp. 451–464

Ashton, L. (2015) “Hunger Games: Does Hunger Affect Time Preferences?”

Attema, A. E. *et al.* (2010) “Time-Tradeoff Sequences for Analyzing Discounting and Time Inconsistency,” *Management Science*, 56(11), pp. 2015–2030

Augenblick, N., Niederle, M. and Sprenger, C. (2015) “Working Over Time: Dynamic Inconsistency in Real Effort Tasks,” *Quarterly Journal of Economics*, 130(3), pp. 1067–1115

Aycinena, D. *et al.* (2022) “Intertemporal choice experiments and large-stakes behavior,” *Journal of Economic Behavior and Organization*, 196, pp. 484–500

Bchir, M. and Willinger, M. (2013) “Does the exposure to natural hazards affect risk and time preferences? Some insights from a field experiment in Perú”

Beine, M. *et al.* (2020) “Shaking things up: on the stability of risk and time preferences”

Belot, M., Kircher, P. and Muller, P. (2021) “Eliciting Time Preferences When Income and Consumption Vary: Theory, Validation and Application to Job Search”

Bernheim, B. D. and Taubinsky, D. (2018) “Behavioral public economics,” *Handbook of behavioral economics: Applications and Foundations 1*, 1, pp. 381–516

Bernheim, B. D. *et al.* (2021) “A Theory of Chosen Preferences,” *American Economic Review*, 111(2), pp. 720–754

Bourdeau-Brien, M. and Kryzanowski, L. (2020) “Natural disasters and risk aversion,” *Journal of Economic Behavior and Organization*, 177, pp. 818–835

Brownback, A., Imas, A. and Kuhn, M. A. (2023) “Behavioral Food Subsidies,” *The Review of Economic and Statistics* [Preprint]

Burness, H. S. (1976) “A Note on Consistent Naive Intertemporal Decision Making and an Application to the Case of Uncertain Lifetime,” *The Review of Economic Studies*, 43(3), pp. 547–549

Callen, M. *et al.* (2014) “Violence and risk preference: Experimental evidence from Afghanistan,” *American Economic Review*, 104(1), pp. 123–148

Cameron, L. and Shah, M. (2013) “Risk-taking behavior in the wake of natural disasters”

Carrera, M. *et al.* (2022) “Who Chooses Commitment? Evidence and Welfare Implications,” *Review of Economic Studies*, 89, pp. 1205–1244

Cassar, A., Healy, A. and Kessler, C. V. (2017) “Trust, Risk, and Time Preferences after a Natural Disaster: Experimental Evidence from Thailand,” *World Development*, 94, pp. 90–95

Charness, G., Chemaya, N. and Trujano-Ochoa, D. (2023) “Learning your own risk preferences,” *Journal of Risk and Uncertainty*, 67, pp. 1–19

DeJarnette, P. (2020) “Temptation over time: Delays help,” *Journal of Economic Behavior and Organization*, 177, pp. 752–761

Drouhin, N. (2020) “Non-stationary additive utility and time consistency,” *Journal of Mathematical Economics*, 86, pp. 1–14

Eckel, C., El-Gamal, M. and Wilson, R. (2009) “Risk loving after the storm: A bayesian-network study of Hurricane Katrina evacuees,” *Journal of Economic Behavior and Organization*, 69(2), pp. 110–124

Gneezy, A. and Imas, A. (2014) “Poverty traps and the effect of helplessness on impatience and risk seeking: Evidence from the lab and field”

Halevy, Y. (2015) “Time Consistency: Stationarity and Time Invariance,” *Econometrica*, 83(1), pp. 335–352

Hanaoka, C., Shigeoka, H. and Watanabe, Y. (2018) “Do Risk Preferences Change? Evidence from the Great East Japan Earthquake,” *American Economic Journal: Applied Economics*, 10(2), pp. 298–330

Harrison, G. W., Lau, M. I. and Yoo, H. I. (2023) “Constant Discounting, Temporal Instability and Dynamic Inconsistency in Denmark: A Longitudinal Field Experiment”

Ifcher, J. and Zarghamee, H. (2011) “Happiness and time preference: The effect of positive affect in a random-assignment experiment,” *American Economic Review*, 101(7), pp. 3109–3129

Imas, A., Kuhn, M. A. and Mironova, V. (2022) “Waiting to choose: The role of deliberation in intertemporal choice,” *American Economic Journal: Microeconomics*, 14(3), pp. 414–440

Janssens, W., Kramer, B. and Swart, L. (2017) “Be patient when measuring hyperbolic discounting: Stationarity, time consistency and time invariance in a field experiment,” *Journal of Development Economics*, 126, pp. 77–90

Jaramillo, P., LaFave, D. and Novak, L. (2023) “Extreme Weather Events Impact Risk Tolerance and Time Preferences”

Kettlewell, N. *et al.* (2018) “Civil War, Natural Disaster and Risk Preferences: Evidence from Sri Lankan Twins”

Kibris, A. and Uler, N. (2021) “The Impact of Exposure to Armed Conflict on Risk and Ambiguity Attitudes”

Kuhn, M. A., Kuhn, P. and Villeval, M. C. (2017) “Decision-environment effects on intertemporal financial choices: How relevant are resource-depletion models?,” *Journal of Economic Behavior and Organization*, 137, pp. 72–89

Kuroishi, Y. and Sawada, Y. (2019) “On the Stability of Preferences: Experimental Evidence from Two Disasters”

Laibson, D. (1997) “Golden eggs and hyperbolic discounting,” *Quarterly Journal of Economics*, 112(2), pp. 443–477

Malmendier, U. and Nagel, S. (2011) "Depression babies: Do macroeconomic experiences affect risk-taking?," *Quarterly Journal of Economics*, 126(1), pp. 373–416

O'Donoghue, T. and Rabin, M. (1999) "Doing it now or later," *American Economic Review*, 89(1), pp. 103–124

Page, L., Savage, D. A. and Torgler, B. (2014) "ariation in Risk Seeking Behaviour Following Large Losses: A Natural Experiment," *European Economic Review*, 71, pp. 121–131

Prelec, D. (2004) "Decreasing Impatience: A Criterion for Non-stationary Time Preference and "Hyperbolic" Discounting," *The Scandinavian Journal of Economics*, 106(3), pp. 511–532

Read, D. and Leeuwen, B. van (1998) "Predicting Hunger: The Effects of Appetite and Delay on Choice," *Organizational Behavior and Human Decision Processes*, 76(2), pp. 189–205

Reynaud, A. and Aubert, C. (2019) "Does flood experience modify risk preferences? Evidence from an artefactual field experiment in Vietnam," *The Geneva Risk and Insurance Review*, 45, pp. 36–74

Sepahvand, M. H., Shahbazian, R. and Swain, R. B. (2019) "Does revolution change risk attitudes? Evidence from Burkina Faso"

Strotz, R. H. (1956) "(1955-1956). Myopia and inconsistency in dynamic utility maximization," *The Review of Economic Studies.-1955-1956*, 23(3), p. 165–166

Strulik, H. (2021) "Hyperbolic discounting and the time-consistent solution of three canonical environmental problems," *Journal of Public Economic Theory*, 23(3), pp. 462–486

Takeuchi, K. (2011) "Non-parametric test of time consistency: Present bias and future bias," *Games and Economic Behavior*, 71(2), pp. 456–478

Voors, M. *et al.* (2012) "Violent conflict and behavior: A field experiment in Burundi," *American Economic Review*, 102(2), pp. 941–964

8 Appendix I

8..1 Tables and Figures

Model :	OLS	Interval Regression	NLS	Logistic Choice
	(1)	(2)	(3)	(4)
δ	0.9980 (0.0002)	0.9980 (0.0002)	0.9981 (0.0002)	0.9984 (0.0002)
β	0.9712 (0.0024)	0.9716 (0.0025)	0.9724 (0.0024)	0.9778 (0.0024)
r (annual discount rate)	0.1074 (0.0120)	0.1104 (0.0136)	0.1047 (0.0118)	0.0877 (0.0115)
$H_0 : \delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
$H_0 : \beta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$

Table 7: Aggregate discount parameter estimates from quasi-hyperbolic model, full sample

Notes: standard errors are clustered at the subject level. δ is the weekly exponential discount factor, and $r = \delta^{-52} - 1$ is the annual discount rate. The data consist of 6,367 price lists from 153 subjects.

	(1)	(2)	(3)	(4)	(5)
Constant	0.9561	0.9640	0.9643	0.9643	0.9646
	(0.0033)	(0.0034)	(0.0036)	(0.0035)	(0.0036)
Choice delay $((t^L - t^S) - 8)$	-0.0006^{***}	-0.0005^{***}	-0.0003	-0.0005^{***}	-0.0003
	(0.0002)	(0.0002)	(0.0003)	(0.0002)	(0.0003)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0034^{***}	0.0019	0.0018	0.0011	0.0010
	(0.0013)	(0.0012)	(0.0012)	(0.0021)	(0.0021)
Survey week (w)		-0.0013^{**}	-0.0013^{**}	-0.0014^{**}	-0.0014^{**}
		(0.0005)	(0.0005)	(0.0005)	(0.0005)
$w \cdot ((t^L - t^S) - 8)$			-0.0000		-0.0000
			(0.0000)		(0.0000)
$w \cdot (1 - \mathbb{1}(t^S = \tau))$				0.0002	0.0001
				(0.0004)	(0.0004)

Table 8: Linear-time-trend estimates of time variance, full sample

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$. Coefficients are from linear models with standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects.

Violation type:	Stationarity		Consistency	
	(1)	(2)	(3)	(4)
Constant	0.9614	0.9613	0.9617	0.9619
	(0.0032)	(0.0032)	(0.0031)	(0.0031)
Choice delay $((t^L - t^S) - 8)$	-0.0005^{***}	-0.0004^{**}	-0.0005^{***}	-0.0005^{***}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0007	0.0010	0.0042^{***}	0.0044^{***}
	(0.0012)	(0.0012)	(0.0014)	(0.0015)
$((t^L - t^S) - 8) \cdot (1 - \mathbb{1}(t^S = \tau))$		-0.0002^{**}		-0.0001
		(0.0001)		(0.0001)
Survey-week FEs	Y	Y	N	N
Drawing-week FEs	N	N	Y	Y

Table 9: Separate estimates of non-stationarity and inconsistency, full sample

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects.

Margin of Error:	None	± 6 tickets
Properties of choices	(1)	(2)
Panel A: Subject level		
A) Invariant, consistent, stationary	6.38%	20.57%
B) Invariant only	1.42%	0%
A) + B)	7.80%	20.57%
C) Stationary only	3.55%	7.09%
D) Consistent only	2.13%	1.42%
E) None	86.52%	70.92%
C) + D) + E)	92.2%	79.43%
Panel B: Triangle level		
A) Invariant, consistent, stationary	33.88%	61.61%
B) Invariant only	6.26%	7.30%
A) + B)	40.14%	68.91%
C) Stationary only	15.20%	15.20%
D) Consistent only	8.34%	8.30%
E) None	36.31%	7.58%
C) + D) + E)	59.86%	31.09%

Table 10: Subject- and decision-triangle-level classification of time preference properties, full sample

Notes: Data 141 subjects that completed at least two consecutive surveys. Column(1) assumes any deviation from is a time preference property violation. Column(2) assumes any deviation greater than six tickets from an analogous choice is a time preference violation, capturing a range of a standard deviation around a choice ($\frac{\sigma}{2} \approx 7$, which due to interval of two tickets between choice rows corresponds to equality within six tickets).

Restrictions:	$\mu=\eta=\gamma=1$	$\mu=\gamma=1$	$\mu=\eta=1$	$\eta=\gamma=1$	None
Properties					
Invariant	✓	✓	X	X	X
Stationary	✓	X	✓	X	X
Consistent	✓	X	X	✓	X
	(1)	(2)	(3)	(4)	(5)
δ	0.9973	0.9778	0.9963	0.9756	0.9735
	(0.0002)	(0.0020)	(0.0004)	(0.0032)	(0.0027)
η	1	0.3811	1	1	0.4291
	.	(0.0257)	.	.	(0.0408)
γ	1	1	1.0460	1	0.9918
	.	.	(0.0168)	.	(0.0168)
μ	1	1	1	0.4801	0.7664
	.	.	.	(0.0342)	(0.1207)
Log-likelihood	$-26,484.8$	$-26,134.2$	$-26,474.4$	$-26,312.6$	$-26,131.9$
AIC	$52,971.7$	$52,272.3$	$52,952.7$	$52,629.1$	$52,271.7$
r_0	0.1481	0.1065	0.2113	0.1792	0.1033
	(0.0140)	(0.0102)	(0.0264)	(0.0145)	(0.0124)
r_{12}	$=r_0$	$=r_0$	0.1183	0.1058	0.1064
			(0.0174)	(0.0115)	(0.0142)
PB_S	1	0.9755	1	0.9735	0.9806
	.	(0.0023)	.	(0.0038)	(0.0109)
PB_C	1	$=PB_S$	1.0025	1	0.9737
	.		(0.0011)	.	(0.0023)
$H_0:\delta=1$	$p<0.0001$	$p<0.0001$	$p<0.0001$	$p<0.0001$	$p<0.0001$
$H_0:\eta=1$	.	$p<0.0001$	.	.	$p<0.0001$
$H_0:\gamma=1$	.	.	$p=0.0070$	.	$p=0.6248$
$H_0:\mu=1$	.	.	.	$p<0.0001$	$p=0.0548$
$H_0:\mu=\gamma=1$	.	.	.	.	$p=0.0531$
$H_0:\mu=\eta=1$	.	.	.	.	$p<0.0001$
$H_0:\eta=\gamma=1$	.	.	.	.	$p<0.0001$
$H_0:\mu=\eta=\gamma=1$	.	.	.	.	$p<0.0001$
$H_0:r_0=r_{12}$	.	.	$p=0.0059$	$p<0.0001$	$p=0.8682$
$H_0:PB_S=1$	.	$p<0.0001$	.	$p<0.0001$	$p=0.0743$
$H_0:PB_C=1$	.	$p<0.0001$	$p=0.0200$	.	$p<0.0001$
$H_0:PB_S=PB_C$	.	.	.	.	$p=0.5267$

Table 11: Aggregate parameter estimates for the nested exponential model, full sample

Notes: standard errors are clustered at the subject level. $r_\tau = D(\tau,\tau+52)^{\{-1\}}-1$ is a measure of the annual discount rate. PB_S and PB_C are stationarity-based and consistency-based measures of present-bias. The data consist of 6,367 price lists from 153 subjects. All estimates are from non-linear least squares regressions.

Properties of choices	
A) Invariant, consistent, stationary	6.52%
B) Invariant only	39.13%
A) + B)	45.65%
C) Stationary only	6.52%
D) Consistent only	7.61%
E) None	40.22%
C) + D) + E)	54.35%

Table 12: Subject-level structural classification of time preference properties, full sample

Notes: Data comes from 92 subjects for which we were able to obtain parameter estimates from all models.

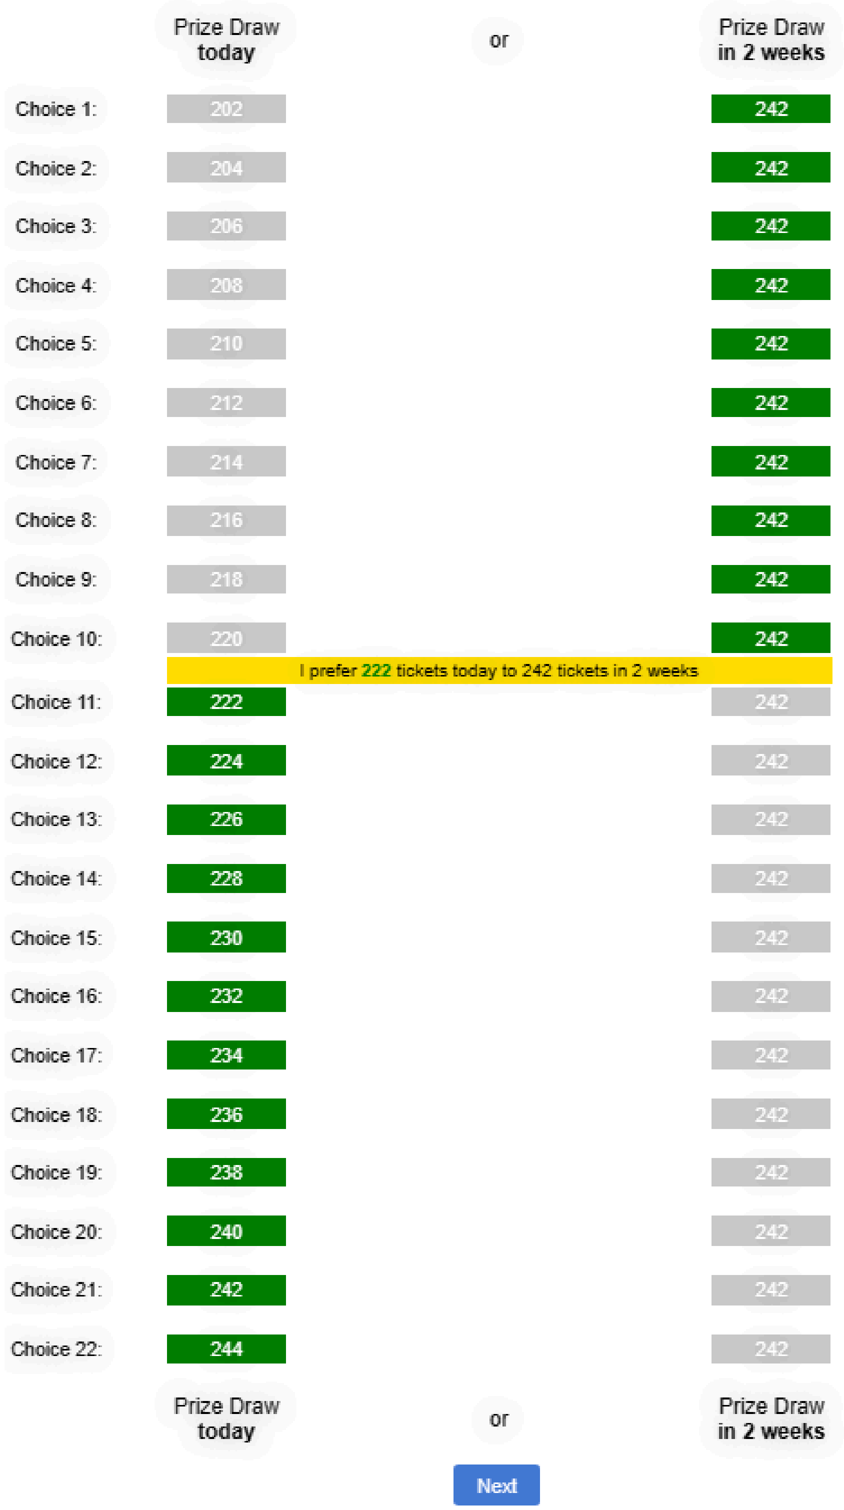

Figure 14: Example price list

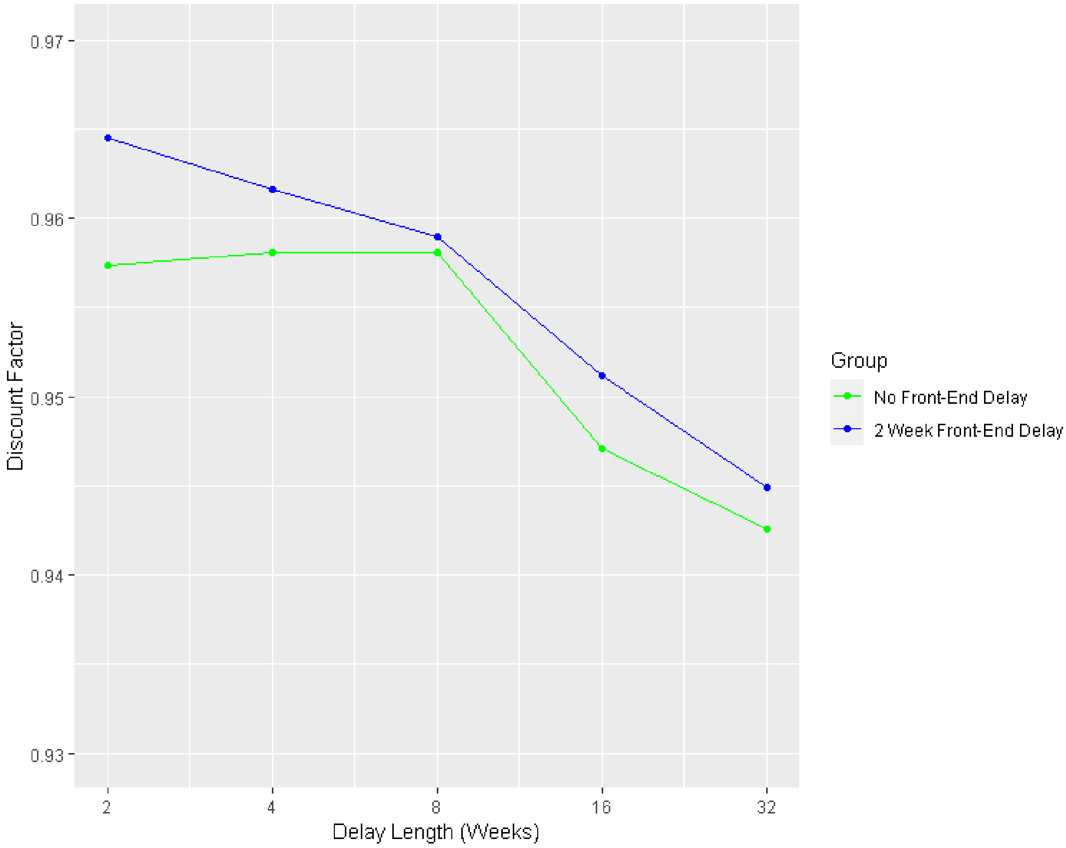

Figure 15: Average discount factor by choice delay, full sample

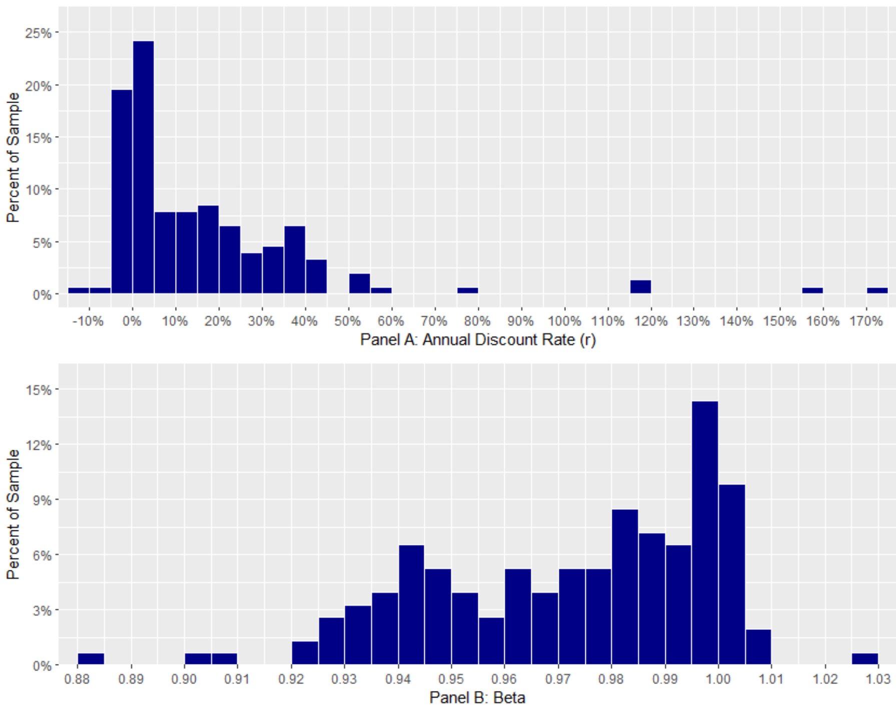

Figure 16: Distributions of individual estimates from the β-δ model, full sample

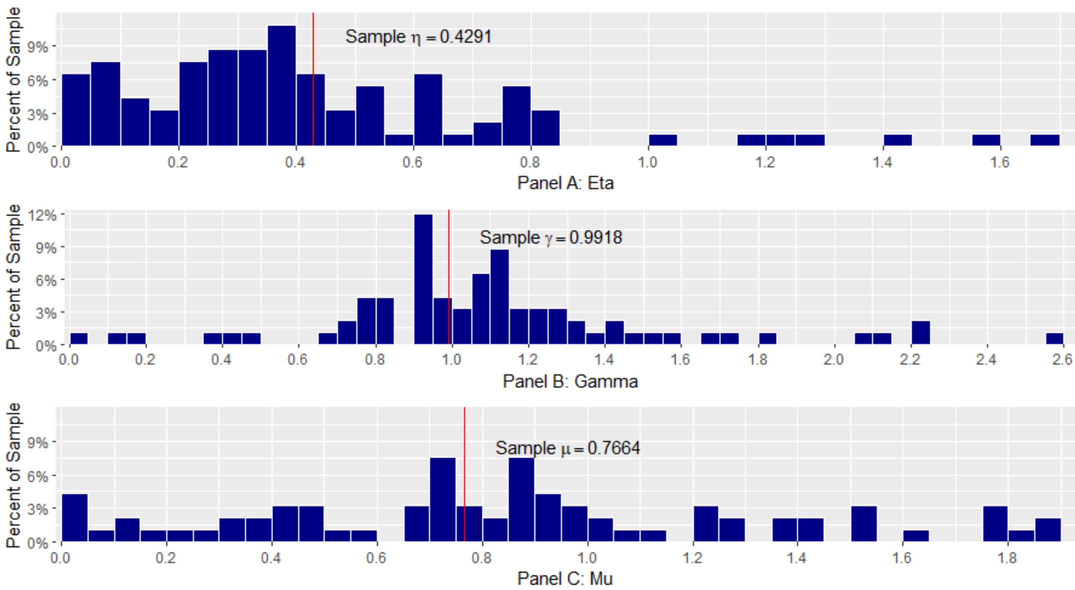

Figure 17: Individual parameter estimates for the unrestricted nested exponential model

Notes: 18 estimates are suppressed from the right tails of the distributions in Panel (B) and Panel (C) for visual clarity.

8..2 Proof of Proposition II

Proof. ($\Leftarrow$) Suppose a discount function $D(\tau, t)$ is only a function of $t - \tau$. Given bundles $(\varepsilon,\ t + \Delta_1) \sim_t (\psi,\ t + \Delta_2)$ which the agent is indifferent between at evaluation time t, we have

$$D(t,\ t + \Delta_1)\ \varepsilon = D(t, t + \Delta_2)\ \psi$$
$$\implies D(t + \Delta_1 - t)\ \varepsilon = D(t + \Delta_2 - t)\ \psi$$
$$\implies D(\Delta_1)\ \varepsilon = D(\Delta_2)\ \psi$$

Note that for any t', we can certainly write $\Delta_i = t' + \Delta_i - t'$ for $i = 1,\ 2$. Thus, the above equation implies

$$D(t' + \Delta_1 - t')\ \varepsilon = D(t' + \Delta_2 - t')\ \psi$$

By assumption, this is equivalent to

$$D(t',\ t' + \Delta_1)\ \varepsilon = D(t',\ t' + \Delta_2)\ \psi$$

which establishes Time Invariance.

($\Rightarrow$) Let $D(\tau, t)$ be a time-invariant function. Suppose there exists bundles $(\varepsilon,\ t + \Delta_1) \sim_t (\psi,\ t \Longleftarrow \Delta_2)$ which the agent is indifferent between at evaluation time t. According to the definition of Time Invariance, the agent's discounted utility must satisfy

$$D(t,\ t + \Delta_1)\ \varepsilon = D(t,\ t + \Delta_2)\ \psi$$

for <u>all</u> $t \geq 0$. Namely, let us choose $t = 0$. Since

$$D(0,\ \Delta_1)\ \varepsilon = D(0,\ \Delta_2)\ \psi$$
$$\implies D(0,\ t - \tau)\ \varepsilon = D(0,\ t' - \tau)\ \psi$$

$=\blacksquare$

8..3 Proof of Proposition IV

Proof. $\Leftarrow$ It is straightforward to verify that $D(\tau,\ t) = \alpha^{-(t-\tau)f(\tau)}$ is a stationary function. Consider bundles $(\varepsilon,\ t)$ and $(\psi,\ t + \Delta)$ which the discounter is indifferent between at time τ. Let $t' \geq \tau$ with $t' - t = k$. Exponent rules tell us

$$D(\tau,\ t')\ \varepsilon = \alpha^{-(t'-\tau)f(\tau)}\ \varepsilon$$
$$= \alpha^{-kf(\tau)}\ D(\tau,\ t)\ \varepsilon$$

The indifference between $(\varepsilon,\ t)$ and $(\psi,\ t+\Delta)$ further tells us that this equals $\alpha^{-kf(\tau)}\ D(\tau,\ t+\Delta)\ \psi$. Expanding upon the particular form of the discount, it is straightforward to refactor this expression into $D(\tau,\ t')\ \psi$. Thus, stationarity has been established.

($\Rightarrow$) Suppose $D(\tau,\ t)$ admits stationary preferences. The definition of stationarity implies that if an agent is indifferent between $(\varepsilon,\ t)$ and $(\psi,\ t+\Delta)$ when evaluated at time τ, then they are also indifferent between bundles $(\varepsilon,\ t')$ and $(\psi,\ t'+\Delta)$ when evaluated at time τ, for any $t' \geq \tau$. In terms of discounted utility, we can summarize this as

$$D(\tau,\ t)\ \varepsilon = D(\tau,\ t+\Delta)\ \psi$$

for all $t \geq \tau$. Taking the log of both sides, we obtain

$$\log[D(\tau,\ t)] + \log\varepsilon = \log[D(\tau,\ t+\Delta)] + \log\psi$$

Since this holds for arbitrary (sensible) t, we can take the derivative on either side with respect to t. Letting $\dot{D}$ denote the derivative of D with respect to payoff time, t, we obtain

$$\frac{\dot{D}(\tau,\ t)}{D(\tau,\ t)} = \frac{\dot{D}(\tau,\ t+\Delta)}{D(\tau,\ t+\Delta)}$$

Therefore, the rate of change of the discount function is independent of payoff time. In other words, given an arbitrary function $g(\cdot)$, stationarity implies that the rate of change of the discount function is solely a function of τ:

$$\frac{\dot{D}(\tau,\ t)}{D(\tau,\ t)} = g(\tau)$$

Review of an entry-level differential equations textbook reveals that one particular function which satisfies $h'(x) = k \cdot h(x)$ is the exponential function $h(x) = a \cdot e^{kx+b}$. The extension to our multivariate world is straightforward: one solution to the differential equation $\dot{D}(\tau,\ t) = g(\tau)D(\tau,\ t)$ is given by

$$D(\tau, t) = a \cdot e^{\ f_1(\tau)\ \cdot\ t + f_2(\tau)}$$

where $f_{1'} = g$ and a is a constant. Since we require $D(0,\ 0) = 1$, we can immediately say that $a = 1$. Furthermore, we require that $D(\tau,\ t) = 0$ whenever $\tau = t$, restricting our exponent to satisfy

$$f_1(\tau) \cdot \tau + f_2(\tau) = 0$$
$$\implies f_2(\tau)\ = -\tau \cdot f_1(\tau)$$

Consequently, we can drop the subscripts on f and our exponent can be reduced to $tf(\tau) - \tau f(\tau) = (t-\tau)f(\tau)$. That is, our stationary discount function is given by

$$D(\tau,\ t) = e^{(t-\tau)f(\tau)}$$

To allow for full generality, we replace e with some $\alpha > 0$. We would like D to be clearly decreasing in t, but this depends on the sign of $1-\alpha$ (as well as the sign of f). Since f is arbitrary, we can factor -1 out of it when $\alpha < 1$. As this is often the case, our function now takes the form

$$D(\tau,\ t) = \alpha^{-(t-\tau)f(\tau)}$$

Hence, we have shown that stationarity implies a discount function of this form. $\blacksquare$