NBER WORKING PAPER SERIES

DYNAMIC PREFERENCE “REVERSALS” AND TIME INCONSISTENCY

Philipp Strack
Dmitry Taubinsky

Working Paper 28961
http://www.nber.org/papers/w28961

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021, Revised July 2022

We thank Jim Andreoni, Ned Augenblick, Doug Bernheim, Ben Enke, Mira Frick, Yoram Halevy, Paul Heidhues, Ryota Iijima, Eric Koepcke, David Laibson, Peter Maxted, Xiaosheng Mu, Pietro Ortoleva, Doron Ravid, Alex Rees-Jones, Charlie Sprenger, Ted O’Donoghue, Tomasz Strzalecki, Can Urgun, and seminar participants at the University of Zurich, Stanford Institute for Theoretical Economics, Toulouse, Toronto, Harvard, Princeton, Berkeley, Chicago, and Cornell for helpful comments. We thank the Alfred P. Sloan Foundation for funding. Afras Sial and Laila Voss provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Dynamic Preference “Reversals” and Time Inconsistency
Philipp Strack and Dmitry Taubinsky
NBER Working Paper No. 28961
June 2021, Revised July 2022
JEL No. C9,D0,D9

ABSTRACT

Time inconsistency leads people to revise earlier plans, which has motivated empirical designs attempting to document such choice revisions. We study identification of time inconsistency in designs where an agent’s preferences are elicited in advance at time 0, and then again later at time 1, after they might have received additional decision-relevant information. We show that for single-peaked preferences, the only data that rejects time-consistent expected utility maximization is when an agent’s time-1 ranking between a pair of alternatives is the reverse of their time-0 ranking with probability one. We establish variations of this result under a variety of other assumptions. However, such patterns of choice are rarely observed in practice. To facilitate more robust identification, we present results about special conditions under which the degree of time inconsistency can be estimated.

Philipp Strack
Yale University
philipp.strack@gmail.com

Dmitry Taubinsky
University of California, Berkeley
Department of Economics
530 Evans Hall #3880
Berkeley, CA 94720-3880
and NBER
dmitry.taubinsky@berkeley.edu

In his seminal work on myopia and dynamic inconsistency, Strotz (1955) posed the following question about an individual choosing "a plan of consumption for a future period of time": "If he is free to reconsider his plan at later dates, will he abide by it or disobey it?" A fundamental intuition arising from his work is that individuals who do not discount future consumption at a constant rate will have time-inconsistent preferences and often choose to revise their consumption plans. For example, individuals might exhibit *present focus*, as in the quasi-hyperbolic discounting model, and thus revise their plans toward more immediately gratifying alternatives over time (e.g., Laibson, 1997, O'Donoghue and Rabin, 1999).

This raises a natural question that we answer in this paper: If an analyst observes individuals who tend to revise their consumption plans in a certain *systematic* direction, while being subject to random taste shocks, when can the analyst infer that the individuals are time-inconsistent, and when can the analyst quantify the degree of time inconsistency?

A number of influential empirical studies that utilize what we call *revision designs* have assumed that the presence of any type of systematic choice reversal implies time inconsistency. A classic example is the study by Read and van Leeuwen (1998), which is often cited as "a canonical example of a preference reversal" (Ericson and Laibson, 2019). Read and van Leeuwen find that when planning a week in advance, 50 percent of individuals choose a healthy over an unhealthy snack, but this fraction declines to 20 percent when individuals are given a surprise opportunity to revise their plans a week later. The implicit assertion in the conclusions drawn by Read and van Leeuwen is that *if* individuals systematically revise their plans toward some types of alternatives over others, *then* they must have time-inconsistent preferences. Economists and psychologists have since conducted numerous revision design studies, often with richer choice sets, in domains such as intertemporal allocation of work, entertainment choice, financial plan-making, opioid use, and nutrition.[1,2] In contrast to standard experiments on take-up of commitment contracts and other more qualitative tests of time inconsistency, a key motivation for revision designs has been the important goal of obtaining point estimates of time-preference parameters.[3] Additionally, revision designs

[1]Work allocation: Augenblick, Niederle, and Sprenger (2015), Abebe, Caria, and Ortiz-Ospina (2021), Andreoni et al. (2020), Barton (2015), Corbett (2016), Imas, Kuhn, and Mironova (forthcoming), Kölle and Wenner (2019), Augenblick and Rabin (2019), Fedyk (2018), see Imai, Rutter, and Camerer (2021) for a review; Financial plan-making: Kuchler and Pagel (2021); Entertainment: Read, Loewenstein, and Kalyanaraman (1999), Milkman, Rogers, and Bazerman (2009), Bartos et al. (forthcoming); Opioid use: Badger et al. (2007); Nutrition choice: Sadoff, Samek, and Sprenger (2019).

[2]Typical designs elicit preferences from an identical choice set at two different points in time in an incentive-compatible way by informing individuals that their initial preferences and revised preferences both have a positive probability of determining their outcomes. Which preference is implemented is determined after both are elicited.

[3]See, e.g., Augenblick, Niederle, and Sprenger (2015), Andreoni et al. (2020), Augenblick and Rabin (2019), Augenblick (2018).

were intended to identify time inconsistency without the assumption that people are fully aware of it.[4]

This paper provides formal results about the choice patterns in revision designs that can and cannot reject time consistency. Our model considers a data set of an agent's revealed ordinal preferences over a finite set of alternatives at two different points in time: an advance choice stage (time 0) and a revision stage (time 1). To encompass the different empirical designs used in practice, we allow the analyst to observe the preferences completely or incompletely. For example, in some designs, researchers elicit only the most preferred alternative; in other designs agents choose from different budget sets, which allows researchers to almost completely observe preferences. At time 1, the agent's preferences are state-dependent; e.g., the agent's rankings over different food items might depend on their level of hunger. To consider a best-case scenario for identification, we suppose that the data set includes the exact distribution of the agent's time-1 preferences.[5] We say that an agent's choices can be rationalized by time-consistent expected utility maximization (TC-EU) if there is an information structure and utility function that rationalizes the observed distribution of the agent's time-1 preferences, and the agent's time-0 ranking of alternatives is consistent with the (objective) expectation of time-1 utilities.[6]

In the first part of the paper, we show that TC-EU can rationalize a variety of different data that feature systematic choice reversals. We begin with numerical examples that capture several types of prominent designs. In these examples, we show that inferences about time inconsistency are highly sensitive to assumptions about how information is revealed to the agent over time.

Our first formal result, Theorem 1, shows that when the choice set is single-dimensional, time-1 preferences are single-peaked, and when at least the time-0 preference is completely

[4]Empirical work estimating both time inconsistency and people's sophistication about it typically finds that people are partially, but not fully, aware of their time inconsistency. See, e.g., DellaVigna and Malmendier (2006), Acland and Levy (2015), Augenblick and Rabin (2019), Chaloupka, Levy, and White (2019), Bai et al. (forthcoming), Carrera et al. (2022). The preferences-over-menu elicitations that are proposed in the decision theory literature (e.g., Gul and Pesendorfer, 2001, Dekel and Lipman, 2012, Ahn and Sarver, 2013), however, implicitly require full sophistication.

[5]In practice, data sets do not have more than several observations of an individual's propensity to revise their choices. A typical assumption that facilitates identification is that individuals who make the same choice in time 0 are homogeneous in their preferences and economic environments, and any differences in time-1 choices are due to independent realizations of time-1 taste shocks.

[6]To be clear, belief-based biases that generate behavior resembling time inconsistency, such as the *planning fallacy* (Kahneman and Tversky, 1982, Buehler, Griffin, and Peetz, 2010, Brunnermeier, Papakonstantinou, and Parker, 2008), overoptimism (Browning and Tobacman, 2015, Breig, Gibson, and Shrader, 2021) or other misperceptions of the time-1 decision environment (Sadoff, Samek, and Sprenger, 2019), are violations of TC-EU in our framework as the time-0 preference is not derived by taking the (correct) expectation over time-1 utilities. Our results thus imply that such biases can also not be identified from revision designs.

observed, the data can be rationalized by TC-EU as long as there is no *simple dominance violation*—i.e., alternatives x_1 and x_2 such that x_1 is preferred to x_2 in time 0, but where x_2 is preferred to x_1 *with probability 1* at time 1. In the general case where preferences are incompletely observed, Theorem 1 states that the data can be rationalized by TC-EU under the more general requirement that there is no *cyclic dominance violation*—i.e., a collection of alternatives $x_1, \ldots, x_k$ such that x_{j+1} is preferred to x_j either at time 0 or with probability 1 at time 1, and such that also x_1 is preferred to x_k either at time 0 or with probability 1 at time 1. Theorem 2 shows that under the additional assumption that the analyst knows the utility functions to be strictly concave in both time 0 and time 1, the data are consistent with TC-EU if and only if there is no simple dominance violation and the time-0 ranking of alternatives is single-peaked.

As neither of these patterns are empirically observed, Theorems 1 and 2 imply that, in line with our numerical examples, existing data sets can not reject the hypothesis that the agent is a time-consistent expected utility maximizer. Our findings caution against the use of revision designs to identify time-consistency if the environment does not admit additional structural properties that facilitate identification, which we discuss in Section 4.

We view Theorems 1 and 2 as our main results that cover a range of environments of economic interest. In addition, provide results for preferences that are not single-peaked. We show by example that TC-EU can be rejected for non-single-peaked preferences even when there are no cyclic dominance violations. The reason is that there may be *stochastic dominance violations*—which we define to be the case where lottery L is revealed to stochastically dominate lottery L' according to time-1 preferences, but where L' stochastically dominates L according to time-0 preferences. Proposition 1 shows that a data set is consistent with TC-EU if and only if there are no stochastic dominance violations. Proposition 3 generalizes this result to the case where the analyst obtains (or assumes) additional cardinal information by directly observing (or assuming) an agent's preferences over a set of lotteries. Additionally, we provide sufficient conditions for when data sets do not exhibit stochastic dominance violations.

In Section 4 we present a set of conditions under which the degree of time inconsistency can be identified. Roughly speaking, point identification requires the existence of at least one dimension of consumption, such as money "far in the future," such that the agent's beliefs about the marginal utility of consumption along this dimension do not change between time 0 and time 1. We argue that the conditions for identification are plausible in recent work such as Augenblick and Rabin (2019), Augenblick (2018), and Fedyk (2018).

Because our results about the difficulty of identification arise due to random taste shocks and other decision-relevant information being realized between time 0 and time 1, we guage the importance of this in Section 5.1. We show that data from recent work suggests that only a small fraction of the variance in time-1 choice revisions can be explained by stable individual differences in time preferences, which suggests a quantitatively large role for random taste shocks or the revelation of other decision-relevant information.

Section 5.2 addresses tests that attempt to differentiate between time preferences and random taste shocks by utilizing augmented designs that link choice in revision designs to decisions about take-up of commitment contracts. We show that combining revision design data with commitment contract take-up decisions does not mitigate the non-identification issues.

Our research question belongs to a small and important theoretical literature that studies what types of data sets can be rationalized by expected utility maximization with Bayesian learning. Shmaya and Yariv (2016) characterize data sets that can reject EU preferences in experimental designs where the agent might believe that the amount of information provided to them reveals information about an underlying state and the analyst observes which state subjects deem most likely. In the context of optimal stopping problems, Heidhues and Strack (2020) show that time preferences can not be identified from observing the distribution of times when an agent completes a task. De Oliveira and Lamba (2021) characterize, for general dynamic decision problems, what sequences of choices can be rationalized by EU preferences if the analyst *knows* the agent's utility function. They find that a distribution over actions can be rationalized if and only if the agent could not improve their expected payoff by deviating and changing their actions.

The rest of this paper proceeds as follows. Section 1 presents numerical examples. Section 2 presents the formal model. Section 3 presents our main results about the types of data sets that are consistent with TC-EU. Section 4 presents results about economic environments where it is possible to identify the degree of time inconsistency. Section 5 discusses generalizations, robustness, and connections to other related literatures in economics. Section 6 concludes. Proofs are relegated to the appendix.

1 Motivating Examples

1.1 Intertemporal Allocation of Consumption

To illustrate the difficulty of identifying time inconsistency from a revision data set, consider the following stylized example based on research designs that study people's intertemporal allocation of effort.[7] The insights from this example apply equally to empirical work on intertemporal allocation of consumption.

An agent has to complete a task that requires one unit of effort. They can decide what fraction $x \in [0, 1]$ of effort to complete at time 1 and what fraction of effort $1-x$ to complete at time 2. The agent is first asked to decide on the division of effort at time 0 and then given the chance to revise their decision at time 1. Truth-telling is incentive-compatible at both times because each decision is implemented with positive probability, as is typical in such experimental designs (e.g., Imai, Rutter, and Camerer, 2021). At time 0 and time 1, respectively, the agent chooses x to minimize:

$$\begin{aligned} \text{Time 0:} \quad & \mathbb{E}_0 \left[\beta\theta_1 c(x) + \beta\theta_2 c(1-x)\right] \\ \text{Time 1:} \quad & \mathbb{E}_1 \left[\theta_1 c(x) + \beta\theta_2 c(1-x)\right] \end{aligned}$$

where $\mathbb{E}_t[\cdot]$ denotes the expectation given the agent's time-t information. These preferences correspond to the commonly-assumed quasi-hyperbolic preferences.[8]

The analyst observes that the agent divides the effort equally between the two periods when deciding at time 0. However, at time 1 the agent instead allocates an average of 45% to period 1, with a standard deviation of 12%. The observed distribution of the ratio of efforts $x/(1-x)$ follows a log-normal distribution.[9] The analyst knows that the cost of effort is $c(x) = x^\gamma$, for a known value of $\gamma > 1$. Assuming that the analyst knows the cost of effort facilitates identification, but we show that nevertheless little can be inferred about the time inconsistency parameter β.

The challenge for identification arises because there are multiple plausible assumptions about how information is revealed to the agent that all fit the data exactly. In Table 1 below, we consider seven different assumptions about information revelation, all of which perfectly match the analyst's data set, but which produce significantly different estimates of the present focus parameter β. Appendix C provides formal mathematical calculations for each of these different sets of assumptions about information revelation.

Rows 1 and 2 consider the case where both θ_1 and θ_2 are independently and identically distributed, and are both revealed at time 1. Rows 3 and 4 consider the case where θ_1 and

[7] For experimental studies using similar designs see, e.g., Augenblick, Niederle, and Sprenger (2015), Barton (2015), Corbett (2016), Kölle and Wenner (2019), Andreoni et al. (2020), Abebe, Caria, and Ortiz-Ospina (2021), Imas, Kuhn, and Mironova (forthcoming); and Imai, Rutter, and Camerer (2021) for a review.

[8] We normalize the "exponential discount factor" δ to 1, which is without loss of generality as it can be included in θ_2.

[9] The allocations of 50 and 45 percent, respectively, are roughly in line with the data in Augenblick, Niederle, and Sprenger (2015).

Table 1: Implied time inconsistency under different information revelation assumptions

	Distribution of shocks	*Information* *time 0*	*time 1*	γ	Estimated β
1	iid		θ_1, θ_2	2	0.82
2	iid		θ_1, θ_2	3	0.67
3	independent	θ_1	θ_1, θ_2	2	0.93
4	independent	θ_1	θ_1, θ_2	3	1.11
5	independent	θ_2	θ_1, θ_2	2	0.72
6	independent	θ_2	θ_1, θ_2	3	0.41
7	independent		θ_1	2	0.72
8	independent		θ_1	3	0.41
9	mult. random walk		θ_1, θ_2	2	0.93
10	mult. random walk		θ_1, θ_2	3	1.11
11	mult. AR(1), α=1.5		θ_1	2	1.53
12	mult. AR(1), α=1.5		θ_1	3	8.17
13	mult. AR(1), α=0.5		θ_1	2	0.56
14	mult. AR(1), α=0.5		θ_1	3	0.15

θ_2 are distributed independently, with nothing learned at time 0 and only θ_1 learned at time 1. Rows 5-8 instead consider the case where the information obtained between time 0 and time 1 is θ_1. Rows 9 and 10, like rows 1 and 2, assume that θ_1 and θ_2 are learned at time 1, but make the alternative assumption that their joint distribution follows a multiplicative random walk, where $\theta_2 = \theta_1 \times \varepsilon$, with ε log-normally distributed and independent of θ_1. Rows 11-14 make the same assumptions as rows 3 and 4 about when θ_1 and θ_2 are learned, but instead assume that their joint distribution follows a multiplicative AR(1) process, where $\log(\theta_2) = \alpha \log(\theta_1) + \log(\varepsilon)$ and ε is log-normally distributed and independent of θ_1.

Table 1 shows that inferences about β are highly sensitive to equally-plausible assumptions about the agent's learning process. Existing empirical work analyzing data sets analogous to this example utilizes a reduce-form regression model—sometimes referred to as the "intertemporal Euler equation" (Augenblick, Niederle, and Sprenger 2015, Imai, Rutter, and Camerer 2021)—that corresponds to the assumptions in rows 1 and 2. The estimating equation utilized in the literature can be written as

$$(\gamma - 1) \log \left(\frac{x^1}{1 - x^1} \right) = (\gamma - 1) \log \left(\frac{x^0}{1 - x^0} \right) + \log(\beta) + \varepsilon,$$

where ε is normally distributed with $\mathbb{E}[\varepsilon] = 0$. To see that it corresponds to the assumptions in rows 1 and 2, note that more generally, the optimal effort at time t satisfies

$$(\gamma - 1) \log \frac{x^t}{1 - x^t} = \log \frac{\mathbb{E}_t[\theta_2]}{\mathbb{E}_t[\theta_1]} + \log(\beta) \mathbf{1}_{t=1} .$$

where $\mathbf{1}_{t=1}$ is an indicator that the decision is made at time $t = 1$. If, as in rows 1 and 2, θ_2 and θ_1 are iid and log-normal, and if they are both learned at time 1, then $\log \frac{\mathbb{E}_1[\theta_2]}{\mathbb{E}_1[\theta_1]}$ is normally distributed with mean 0, while $\log \frac{\mathbb{E}_0[\theta_2]}{\mathbb{E}_0[\theta_1]} = 0$. However, the learning processes in the other rows are inconsistent with the reduced-form linear model that the literature has previously used to estimate β. Moreover, Table 1 makes clear that the assumptions in rows 1 and 2 cannot be used to obtain either lower or upper bounds on β. Other assumptions can lead to significantly lower values of β, or to significantly higher values of β consistent with *future focus*.[10] While this example is stylized, our formal results show that the inability to identify time inconsistency is not a consequence of any of the special features of this example, but a general feature of revision designs.

1.2 Food choice

Inspired by Read and van Leeuwen (1998) and related studies,[11] consider an agent who chooses between a healthy snack and an unhealthy snack. In the Read and van Leeuwen (1998) experiment, participants choose between a healthy and an unhealthy snack at time 0, to be delivered at time 1 (after seven days). Then, at time 1, participants are given a surprise opportunity to revise their time-0 choice. On average (collapsing across individuals and conditions), subjects at time 0 choose the healthy snacks approximately 50 percent of the time, but at time 1 choose healthy snacks approximately 20 percent of the time.[12] Moreover, Read and van Leeuwen report that switching to the unhealthy snack was far more common than switching to the healthy snack.

[10]In Appendix B, Augenblick, Niederle, and Sprenger (2015) consider the alternative assumption that there is uncertainty on the *curvature* of the cost of effort functions, and that the shocks to time 1 and 2 cost of effort functions are perfectly correlated. Under these assumptions, they find that uncertainty generates moderate upward bias in their estimate of the present focus parameter β. As our results show, other forms of taste shocks—in particular those that affect the cost in both periods differentially— can generate either upward or downward bias in estimates of present focus. Relatedly, our additional analysis in Supplementary Appendix E.2.1 shows that Augenblick, Niederle, and Sprenger's supplementary data on commitment contract take-up does not substantially narrow the set of time preferences consistent with their data.

[11]Empirical designs with similar structures include the food-delivery field experiment of Sadoff, Samek, and Sprenger (2019), Read, Loewenstein, and Kalyanaraman's (1999) study of choice between high-brow and low-brow video rentals, and Milkman, Rogers, and Bazerman's (2009) quasi-experimental extension of Read, Loewenstein, and Kalyanaraman (1999).

[12]These summary statistics are reported in Cohen et al. (2020).

To obtain the results of Read and van Leeuwen with time-consistent preferences, consider a population of agents with the following preferences. At time 1, the agents feel gorged with probability 20 percent, in which case they crave cleansing healthy food, so that the utility difference between the unhealthy and healthy snack equals -5. With probability 80 percent the utility difference between the unhealthy and healthy snack equals 1. At time 0 it is thus optimal for the agents to choose the healthy snack if they do not know the time-1 state, as $0.8 \times 1 + 0.2 \times (-5) < 0$. Now suppose that 38 percent of the agents do not know the time-1 state at time 0, while the remaining 62 already know the state at time 0. Thus, $0.38 + 0.62 \cdot 0.2 = 50$ percent choose the healthy snack at time 0. However, only 20 percent choose the healthy snack at time 1. Moreover, the direction of revisions is asymmetric: 38 percent switch from choosing the healthy snack at time 0 to choosing the unhealthy snack at time 1, but no one switches from choosing the unhealthy snack to choosing the healthy snack.

2 Model

Preference data set There is an agent who has a preference over a finite set of alternatives X at time 0 and time 1.[13] Their preference at time 0 is deterministic and denoted by $\preceq^0$. Their preference $\preceq^1$ at time 1 is a random draw from $(\preceq^1_1, \ldots, \preceq^1_n)$, with $n < \infty$. We denote by $(f_1, \ldots, f_n)$ the strictly positive probabilities (or frequencies) associated with each realization. A *data set*

$$(\preceq^0, \preceq^1_1, \ldots, \preceq^1_n, f_1, \ldots, f_n)$$

consists of a time-0 preference $\preceq^0$ and a probability distribution over time-1 preferences, which we compactly write as $(\preceq^1, f)$.

In practice, some revision design experiments allow the analyst to observe preferences (almost) completely, while others provide much more limited information. For example, experimental designs can provide more detailed information about preferences by studying choice from multiple budget sets. We leave our model general enough to cover these possibilities by assuming that the preferences can be either complete, i.e. the ranking of any two alternatives is observed, or incomplete.[14] We assume that if two alternatives $x, y \in X$ are related by the (potentially incomplete) preference $\preceq$, then the analyst observes whether the agent is indifferent or prefers one of the alternatives strictly; i.e., either $x \sim y$, or $x \prec y$, or

[13] We make the assumption that the choice set is finite to avoid technicalities and streamline the presentation. We put no bound on the size of the choice set. In empirical applications the choice set is necessarily finite.

[14] We do not require $\preceq^0$ or any of the $\preceq^1_j$ to relate the same pairs of alternatives, or even the same number of alternatives.

$y \prec x$.[15] Note that if strict preferences cannot be observed, then any data set is trivially consistent with TC-EU, where the utility function assigns the same value to each alternative.

Time-Consistent Expected Utility Preferences The stochasticity in observed time-1 preferences might result from the agent receiving information at the beginning of time 1 about payoff-relevant aspects of the decision—such as how busy they will be in the future or whether they crave sweet or savory foods—that is not resolved until time 1. The analyst does not observe these states directly and only knows the distribution over the different rankings that result from variation in time-1 states.

To model this formally we consider a time-consistent expected-utility (TC-EU) agent who evaluates alternatives according to the utility function

$$u : X \times \Omega \to \mathbb{R}$$

that depends on the chosen alternative $x \in X$ and state $\omega \in \Omega$. The states capture taste shocks or information that arrives between time 0 and time 1, and the agent does not know the state at time 0, but observes it at time 1. Our convention is to denote the state by a subscript and the alternative as an argument, so that $u_\omega(x)$ denotes the utility of alternative x in state ω.

Without loss, we can assume that there as many states as realizations of the time-1 preference, with each realization of the time-1 preference corresponding to a state, so that $\Omega = \{1, \ldots, n\}$.[16] The TC-EU agent prefers alternative x over y in state ω at time 1 if and only if x has a higher associated utility; i.e., for all $x, y \in X$, $\omega \in \Omega$,

$$y \preceq^1_\omega x \quad \Leftrightarrow \quad u_\omega(y) \leq u_\omega(x)\,. \tag{1}$$

At time 0 the TC-EU agent prefers x over y if and only if the expected utility of x exceeds

[15] Formally, $\preceq$ is a preorder where we interpret $x \preceq y$ and $y \preceq x$ as indifference $x \sim y$, and $x \preceq y$ but not $y \preceq x$ as a strict preference for u over x. The preference between x,y is unobserved if neither $x \preceq y$ nor $y \preceq x$. If the ranking of any two alternatives is observed then $\preceq$ is a complete preorder.

[16] To see that assuming that each state corresponds to an observed preference profile is without loss, note that if we have a TC-EU representation (u, Ω, F), consisting of a utility u, a state space Ω, and a prior F that is consistent with the ordinal preferences $(\preceq^0, \preceq^1)$, then without loss of generality we can associate each set Ω_k of states that leads to a preference profile $\preceq^1_k$ with the newly defined state k. We define a utility function on this new state space as the conditional expectation $\tilde{u}_k(x) = \frac{\int_{\Omega_k} u_\omega(x) dF}{\int_{\Omega_k} dF}$ and obtain a new EU representation with the desired state space $\{1, \ldots, n\}$.

the expected utility of y; i.e., for all $x, y \in X$,

$$y \preceq^0 x \iff \sum_{\omega \in \Omega} f_\omega u_\omega(y) \leq \sum_{\omega \in \Omega} f_\omega u_\omega(x)\,. \tag{2}$$

Definition 1 (Consistency with TC-EU)**.** A data set $(\preceq^0, \preceq^1, f)$ is consistent with TC-EU if there exists a utility function $u : X \times \Omega \to \mathbb{R}$ that satisfies (1) and (2).

In words, consistency with TC-EU means that there exists a state-dependent utility function that is consistent with the observed time-1 preference $\preceq^1_\omega$ in each state ω, and such that the expectation of this utility function is consistent with the observed time-0 preference $\preceq^0$. Our definition of TC-EU requires the agent to correctly understand the distribution of states. Thus, time-inconsistent behavior generated by belief-based biases such as the *planning fallacy*[17] or other forms of overoptimism[18] is not compatible with our definition of TC-EU.

Remark 1. *Our assumption of deterministic time-0 preferences is without loss: if the analyst observes the time-1 choices conditional on the agent's time-0 choice, our results apply to the conditional choice distribution. If the analyst observes only the marginal distributions of preferences, she has strictly less information and identification of time-inconsistency becomes harder, which means that all our non-identification results apply.*[19]

Remark 2. *Our assumptions facilitate identification of time-inconsistent behavior, as the analyst observes the exact distribution of time-1 preferences. This corresponds to the limit case where the analyst either observes the same agent's behavior in exactly the same situation infinitely often, or to the limit case where the analyst observes infinitely many agents with identical preferences in the exactly same informational environment (but with independent realizations of states at time 1). Empirical work typically attempts to approximate one of these limit cases, usually by imposing some type of homogeneity assumption. Thus, our results can be seen as showing that rejecting TC-EU is difficult even with these arguably strong assumptions.*

[17] See e.g. Kahneman and Tversky (1982), Buehler, Griffin, and Peetz (2010), Brunnermeier, Papakonstantinou, and Parker (2008).

[18] See e.g. Browning and Tobacman (2015), Breig, Gibson, and Shrader (2021).

[19] While we do not formally introduce random time-0 choice in the body of the paper as it adds no insight and complicates notation, we recognize its empirical relevance and present the model with random time-0 choice in Supplementary Appendix E.

3 Rejecting Time Consistency

3.1 Single-Peaked Preferences

We first consider the case where the alternatives are real numbers $X \subset \mathbb{R}$, and the time-1 preferences are single-peaked, defined as follows:

Definition 2 (Single-Peaked Preference)**.** A potentially incomplete preference $\preceq$ is single-peaked if for any alternatives $x < y < z$ either $x \prec y$ or $z \prec y$.

The random preference $\preceq^1$ is single-peaked if each possible realization $(\preceq^1_\omega)_{\omega\in\Omega}$ is single-peaked. The single-peaked property is natural in environments where agents choose how to allocate consumption or effort over time, as in the numerical example in Section 1.1. If the utility from leisure is concave in each period, then the agent's utility function will be concave and thus single-peaked in the share of work done in time 1. The single-peaked property also mechanically applies to binary choice sets, as in the example in Section 1.2. We consider the slightly stronger assumption of concavity in the next subsection.

Our definition of single-peaked preferences applies to incomplete preferences. For example, suppose that it is known that an agent's most preferred alternative is some $x^* \in X$. The single-peakness implies that moving further left or right of x^* leads to less attractive alternatives, but it provides no information about how the agent compares alternatives to the left of x^* against those to the right of x^*.

A natural requirement on a data set is that the agents' choices do not contradict themselves: if at time 1 the agent *always* prefers some alternative over another then they should also prefer that alternative at time 0. The next definition formalizes this idea.

Definition 3 (Simple Dominance Violations)**.** A data set $(\preceq^0, \preceq^1, f)$ exhibits simple dominance violations if there exist $x, y \in X$ such that $x \preceq^0 y$, $y \preceq^1_\omega x$ for all $\omega \in \Omega$, and either $x \prec^0 y$ or $y \prec^1_\omega x$ for some $\omega \in \Omega$.

For example, if the analyst observes that the agent *always* prefers the unhealthy over the healthy snack at time 1, then the agent should also prefer the unhealthy snack at time 0 if they are time-consistent. It follows immediately from the definition that any data set that is consistent with TC-EU cannot exhibit simple dominance violations. Theorem 1 below presents a converse of that statement when the time-0 preference is complete and the time-1 preference is single-peaked.

We also introduce a generalization of simple dominance violations to provide a more general characterization of incomplete preferences. To simplify notation, we denote by $\preceq^1_*$ the preorder that is generated by agreement of the agent's preferences in all states ω: $x \preceq^1_* y \Leftrightarrow x \preceq^1_\omega y$ for all ω.

Definition 4 (Cyclic Dominance Violations)**.** A data set $(\preceq^0, \preceq^1, f)$ exhibits cyclic dominance violations if there exists a sequence of alternatives $x_1, x_2, \ldots, x_k$ that are alternatingly ranked by the orders $\preceq^0$, $\preceq^1_*$,

$$x_1 \preceq^0 x_2 \preceq^1_* x_3 \preceq^0 \ldots \preceq^0 x_k \preceq^1_* x_1,$$

with at least one relation strict.[20]

Clearly, a simple dominance violation is a cyclic dominance violation with a cycle of length 2. And as Lemma 3 in Appendix A.1 shows, cyclic dominance violations imply simple dominance violations when the time-0 preference is complete. In general, however, data sets can exhibit cyclic dominance violations without exhibiting simple dominance violations (see Example 1 below). The absence of a cyclic dominance violation is trivially necessary for a data set to be compatible with TC-EU and our next result shows that it is indeed sufficient.

Theorem 1 (Consistency with TC-EU)**.** *Consider $X \subset \mathbb{R}$ and let $(\preceq^0, \preceq^1, f)$ be a data set with strict and single-peaked time-1 preferences.*

(i) The data set is consistent with TC-EU if and only if it exhibits no cyclic dominance violations.

(ii) If $\preceq^0$ is complete, then the data set is consistent with TC-EU if and only if it exhibits no simple dominance violations.

Theorem 1 implies that it could be difficult to reject TC-EU using revision designs. For example, if at time 0 the agent chooses to allocate a fraction $x = 1/2$ of resources to time 1 (and the remainder $x = 1/2$ to time 2), then TC-EU is rejected if and only if at time 1 the agent *always* revises to allocate more resources to time 1, i.e. $x > 1/2$.

How demanding of a test this is depends on the randomness of taste shocks at time 1. In a deterministic environment (i.e., $|\Omega| = 1$) with a rich choice set, an agent with dynamically inconsistent time preferences will often exhibit dominance violations. However, the more variability there is in an agent's time-1 preference, the more unlikely dominance violations become, even if the agent does have dynamically inconsistent time preferences. Thus, Theorem 1 suggests that while revision designs can provide discerning tests of time inconsistency in deterministic environments, they are fairly uninformative about whether an agent is time-consistent or time-inconsistent when there is significant stochasticity in time-1

[20] We only need to consider cycles where the order is generated between alternations between $\preceq^1_*$ and $\preceq^0$ because due to transitivity we can always remove elements that are bounded from above and below in the same order. Any such cycle has an even number of elements.

choices.[21] We are not aware of any revision design study where the data contain a simple or cyclic dominance violation.

Remark 3. *The consistency conditions in Theorem 1 do not involve the probabilities f. Thus, Theorem 1 also applies to the case where the analyst simply knows that each time-1 preference profile $\preceq_j^1$ occurs with positive probability. An analogous comment applies to the other theorems in this section.*

We end this section with an example of a data set that admits a cyclic but not simple dominance violation.

Example 1. There are four alternatives $X = \{1, 2, 3, 4\}$ and only one state, with single-peaked time-1 preference $\preceq^1$ in that state. The time-1 (incomplete) preference is $1 \prec^1 2$ and $4 \prec^1 3$, while the time-0 (incomplete) preference is $2 \prec^0 4$ and $3 \prec^0 1$. Since each preference relates a different pair of alternatives, there is no simple dominance violation. However, no utility function is consistent with both the time-0 and time-1 preferences, as the time-1 preference would imply that $u(1) < u(2)$, the time-0 preference would imply that $u(2) < u(4)$, the time-1 preference would also imply that $u(4) < u(3)$, which would then imply that $u(1) < u(3)$, violating the time-0 preference. The cycle here is $1 \prec^1 2 \prec^0 4 \prec^1 3 \prec^0 1$.

3.2 Concave Utilities

An additional plausible restriction is that the agent's time-1 utility is concave in each state. For example, it is natural to assume that each period, the utility from consumption is concave, or that the cost of effort is convex. If the agent decides what share of resources x to allocate to time 1, and what share $1-x$ to allocate to time 2, and if $u_\omega(x) = v_\omega^1(x) + v_\omega^2(1-x)$, with v^1 and v^2 both concave, then u_ω will be concave as well. We say that a data set is consistent with *concave TC-EU* if is consistent with TC-EU for a strictly concave utility function.

It is immediate that any concave utility function leads to single-peaked preferences. Moreover, because the expectation of a strictly concave function is itself strictly concave, it follows that the time-0 utility must be concave whenever all time-1 utilities are concave. This immediately implies that a necessary condition for a data set to be consistent with concave

[21] Theorem 1 and related results about ordinal preference data sets can be generalized from TC-EU to recursive preferences that nest TC-EU, such as the Epstein and Zin (1989) preferences. To see this, let $\preceq^0$ be complete and note that any time-0 preference that is a monotonic (but not necessarily linear) function of time-1 utilities cannot generate simple dominance violations. Conversely, as TC-EU is a recursive preference, any data set consistent with TC-EU is also consistent with recursive preferences. Theorem 1 thus implies that if a strict, single-peaked data set is consistent with some monotone recursive preference it is also consistent with TC-EU, and hence TC-EU imposes little restrictions on the data. We thank Yoram Halevy and Faruk Gul for pointing this out.

TC-EU is that the time-0 preference must be single-peaked. Our next result shows that this condition, together with no *simple* dominance violations, is also sufficient, and thus provides a complete characterization of all data sets that are consistent with concave TC-EU.

Theorem 2. *If $X \subset \mathbb{R}$, a data set $(\preceq^0, \preceq^1, f)$ is consistent with concave TC-EU if and only if (i) time-0 and time-1 preferences are single-peaked and (ii) the data exhibit no simple dominance violations.*

We note that the additional restriction that the time-0 preference is also single-peaked implies that lack of simple dominance violations is enough to guarantee consistency with TC-EU; it is not necessary to consider cyclic dominance violations more generally in this case. In Example 1, the time-0 preference is not single-peaked, because single-peakness would require that $1 \preceq^0 2$ and $2 \preceq^0 3$ if $2 \prec^0 4$, which is inconsistent with $3 \prec^0 1$.

3.3 A Sketch of the Proofs of our Main Results

We next sketch the idea behind the proofs of Theorems 1 and 2. We invite readers not interested in the mathematical argument to skip this section. First, define $U^0 \subset \mathbb{R}^{|X|}$ to be the set of utilities consistent with $\preceq^0$ and $\bar{U}^1 \subset \mathbb{R}^{|X|}$ to be the set of *expected* utilities consistent with the random preference $\preceq^1$. The sets $U^0, \bar{U}^1$ have an empty intersection if and only if the data set $(\preceq^0, \preceq^1, f)$ cannot be explained by TC-EU. The sets $U^0, \bar{U}^1$ are convex, relatively open cones that are closed under the addition of constants. Thus, they have an empty intersection if and only if there exists a hyperplane that properly separates U^0 and $\bar{U}^1$; i.e., there is a vector $p \in \mathbb{R}^{|X|}$ and a constant $c \in \mathbb{R}$ such that $p \cdot u^0 \geq c \geq p \cdot u^1$ for all $u^0 \in U^0$ and $u^1 \in \bar{U}^1$, with one of the inequalities strict. Due to the structure of $U^0, \bar{U}^1$, this hyperplane must pass through the origin (i.e., $c = 0$) and the entries of p must sum to 0. Moreover, we show that that the same inequality must remain strict for *all* $u^0 \in U^0$ and $u^1 \in \bar{U}^1$. These properties imply that we can write p as the difference of two component-wise non-negative vectors whose entries sum to 1, and thus correspond to lotteries over alternatives. These lotteries have the property that one is preferred under any utility representation of the time-0 preference $\preceq^0$, while the other is preferred under any utility representation of the time-1 preference $\preceq^1$, with one of these preferences strict (we state variants of this insight as Propositions 1 and 3). The goal, then, is to establish that the existence of these lotteries, together with the assumptions of Theorems 1 or 2, implies cyclic or simple dominance violations, respectively.

For clarity, we sketch the remainder of the argument only for the case where all preferences are complete, and where there is a unique least-preferred alternative at both time 0 and in every state at time 1. Without these restrictions, the argument is substantially more involved, and makes use of supplementary results on aggregation of incomplete preorders in Appendix A.1. In the complete preference case, the remainder of the proof establishes that if $\bar{U}^1$ is generated by either single-peaked or concave utilities, then the existence of two such lotteries implies that one can find two *degenerate* lotteries with the same property, which then correspond to a simple dominance violation. Mathematically, this is the main step of the argument and might be of independent interest: It corresponds to establishing a condition under which if a separating hyperplane exists between two cones, there also exists a hyperplane where p has only two non-zero entries equal to $+1$ and -1.

We prove this result by induction over the number of alternatives. It is trivially true for two alternatives. For any number of alternatives larger than 2 we establish that the alternative that is least preferred according to the time-0 preference is such that (i) it either dominates another alternative according to the time-1 preference or (ii) its corresponding coefficient in p must equal zero. Case (i) implies a separating hyperplane with only two non-zero entries. Case (ii) reduces the problem to one with one less alternative, which then completes the proof by the induction hypothesis.

To establish the claim in case (ii) above, we show that if a given alternative x does not dominate any other alternative according to the time-1 preferences, and preferences are either single-peaked or concave, there exists time-1 utility representations that leaves the expected utility of any other alternative unchanged, and decreases the expected utility of x by an arbitrary amount. As x is chosen to be the least preferred element according to the time-0 preference, one can also arbitrarily decrease its utility in a representation of the time-0 preferences. But because the hyperplane separates U^0 and $\bar{U}^1$, the coefficient corresponding to x in the hyperplane p must then be zero.

3.4 Other Types of Ordinal Preference Data

We have thus far characterized consistency with TC-EU for data sets where the time-1 preference is single-peaked. A natural question is whether the results generalize to non-single-peaked preferences. The example below shows that they do not.

Example 2. There are 6 alternatives $X = \{1, 2, 3, 4, 5, 6\}$ and two states $\Omega = \{1, 2\}$. All preferences are strict and given by

$$\begin{gathered}
1 \prec^0 2 \prec^0 3 \prec^0 4 \prec^0 5 \prec^0 6 \\
2 \prec^1_1 3 \prec^1_1 6 \prec^1_1 1 \prec^1_1 4 \prec^1_1 5 \\
4 \prec^1_2 1 \prec^1_2 2 \prec^1_2 5 \prec^1_2 6 \prec^1_2 3\,.
\end{gathered}$$

It is easy to check that the preferences given in Example 2 exhibit no simple or cyclic dominance violations. However, the preferences are not consistent with TC-EU. In both states, the time-1 preference implies that the individual strictly prefers a uniform lottery over $\{1, 3, 5\}$ to a uniform lottery over $\{2, 4, 6\}$ in each state, since the former first-order stochastically dominates the latter. However, the time-0 preference implies that the individual strictly prefers a uniform lottery over $\{2, 4, 6\}$ to a uniform lottery over $\{1, 3, 5\}$.[22]

To formalize more general necessary and sufficient conditions for consistency with TC-EU, recall that a complete ordinal preference $\preceq$ over X induces an incomplete preference over lotteries through first-order stochastic dominance. Formally, denote by $L(x)$ the probability assigned to x by the lottery $L \in \Delta(X)$. For $x_1 \preceq x_2 \preceq \ldots \preceq x_{|X|}$ the lottery L is dominated by L' if for all $r \in \{1, \ldots, |X|\}$,

$$\sum_{s=1}^{r} L(x_s) \geq \sum_{s=1}^{r} L'(x_s) \,.$$

Strict dominance—denoted by $\prec$—holds if dominance holds, but not equality. We can also generalize the stochastic dominance order to incomplete preferences: $L \preceq L'$ if and only if first-order stochastic dominance holds for all completions of the incomplete preference.

Example 2 motivates a stronger necessary and sufficient condition for consistency with TC-EU that involves lotteries:

Definition 5 (Stochastic Dominance Violations)**.** A data set $(\preceq^0, \preceq^1, f)$ exhibits a stochastic dominance violation if there exist lotteries $L, L' \in \Delta(X)$ such that $L \preceq^0 L'$, $L' \preceq^1_\omega L$ for all $\omega \in \Omega$ and either $L \prec^0 L'$ or $L' \prec^1_\omega L$ for some $\omega \in \Omega$.

Since the lotteries can be degenerate, any simple dominance violation is also a stochastic dominance violation. Similarly, a cyclic dominance violation also implies a stochastic dominance violation.[24] Our next result establishes that absence of stochastic dominance violations characterises TC-EU preferences.

Proposition 1 (Consistency with TC-EU)**.** *A data set $(\preceq^0, \preceq^1, f)$ is consistent with TC-EU if and only if it exhibits no stochastic dominance violations.*

[22]Example 2 is minimal in the following sense: For all data sets with two states ($|\Omega| = 2$) and fewer than 6 alternatives ($|X| < 6$), consistency with TC-EU is ensured when there are no simple dominance violations (i.e., the conclusion of Theorem 1 holds for non single-peaked data with $|\Omega| = 2, |X| < 6$). Furthermore, Example 2 is the only data set (up to relabeling of the states) with $|\Omega| = 2, |X| = 6$ where Theorem 1 does not hold.[23] But there are many other examples of data sets not consistent with TC-EU and not exhibiting simple dominance violations when $|\Omega| > 2$ or $|X| > 6$.

[24]Following the notation of Definition 4, let L' be a uniform lottery over the alternatives with even indices in the cycle, and let L be a uniform lottery over the alternatives with odd indices in the cycle. Then $L \preceq^0 L'$ and $L' \preceq^1_\omega L$ for all ω, with at least one of the preferences strict.

Proposition 1 follows from more familiar separation arguments that serve as Lemmas for our main results in Theorems 1 and 2.[25] Although it is possible to reject TC-EU without cyclic dominance violations for certain types of non-single-peaked preferences, the proposition shows that, loosely speaking, rejecting TC-EU is not *that* much easier for other classes of preferences. Practically, the proposition is most useful for interpreting empirical designs where the alternatives are not naturally ordered along a single dimension; e.g., designs where individuals might choose between three or more different food options. Because in such designs the choice set is unlikely to be large, checking for stochastic dominance violations is not difficult in these applications.

Designs where the elements are not naturally ordered are often analyzed with convenient parametric models of stochastic discrete choice, such as the Luce model. Below, we provide a sufficient (but not necessary) condition that may be easier to check in some applications than the necessary and sufficient condition in Proposition 1—and that speaks directly to a key property of commonly used discrete choice models. The sufficient condition extends our simple dominance consistency condition to a consistency condition over sets.

Proposition 2. *The data set* $(\preceq^0, \preceq^1, f)$ *is consistent with TC-EU if one of the following holds:*

(i) For each x *there is a* ω *such that for all* $y \succeq^0 x$ *we do not have* $x \succeq^1_\omega y$*.*

(ii) For each x *there is a* ω *such that for all* $y \preceq^0 x$ *we do not have* $x \preceq^1_\omega y$*.*

As an illustration, suppose that each alternative is the most preferred alternative with positive probability—a *positivity* condition that holds for many standard stochastic discrete choice models. Positivity holds, for example, in models with a random effect that follows a Type-1 extreme value distribution—such as the one used by Sadoff, Samek, and Sprenger (2019) to estimate the degree of dynamic inconsistency in food choice. In this case, all alternatives that are lower-ranked according to the time-0 preference will also be lower-ranked according to the preferences in the state where it is the most preferred alternative. Thus, the second sufficient condition of the above proposition holds.

[25] An immediate conclusion obtained by combining Theorems 1 and 2 and Proposition 1 is the following equivalence: A data set $(\preceq^0, \preceq^1, f)$ with strict single-peaked time-1 and either complete or single peaked time-0 preferences admits a stochastic dominance violation if and only if it admits a simple dominance violation. Thus, intuitively, data sets composed of single-peaked preferences cannot violate dynamic consistency in a complicated way without also violating it in a simpler way. We thank Xiaosheng Mu and Pietro Ortoleva for pointing out this equivalence as an alternative perspective on our results.

3.5 Cardinal Preferences Data

Proposition 1 follows from a more general result that applies to cases where some cardinal information is known about time-0 or time-1 preferences. For example, some cardinal information can be acquired by observing preferences over a set of lotteries, allowing the analyst to draw some conclusions about the curvature of the utility functions. Some types of cardinal information might also be assumed. For example, the analyst might use supplementary data on risk aversion, the elasticity of labor supply, or the elasticity of intertemporal substitution to decide what is a "reasonable" degree of curvature.

Formally, we associate each cardinal preference with a vector $u_\omega \in \mathbb{R}^{|X|}$. We consider a data set (U^0, U^1) where it is known that the time-1 cardinal preferences in state ω satisfy $u_\omega(\cdot) \in U^1_\omega \subseteq \mathbb{R}^{|X|}$, and we set $U^1 \subseteq \mathbb{R}^{|X|\times|\Omega|}$ to be the product of the sets U^1_ω. We make three assumptions about U^1: (i) U^1_ω is convex; (ii) U^1_ω is open relative to its affine hull; (iii) if $u_\omega \in U^1_\omega$ then $\lambda_0 + \lambda_1 u_\omega \in U^1_\omega$ for all $\lambda_0 \in \mathbb{R}$ and $\lambda_1 \in \mathbb{R}_{++}$. Assumption (iii) corresponds to the fact that cardinal information about utility functions can only be learned up to monotonic linear transformations. Assumption (ii) is an innocuous technical condition. Assumption (i) is arguably the strongest.

All three of these assumptions are satisfied for sets U^0 and U^1 of utility functions that represent the ordinal relations $\preceq^0$ and $\preceq^1$, respectively (see Lemma 7 in the Appendix). The following definition generalizes consistency with TC-EU.

Definition 6 (Consistency with TC-EU)**.** A data set (U^0, U^1, f) is consistent with TC-EU if there exist utility functions $u^1 \in U^1$ and $u^0 \in U^0$ such that for all $x \in X$,

$$u^0(x) = \sum_{\omega\in\Omega} f_\omega u^1_\omega(x)\,.$$

Intuitively, U^0 captures the information that the analyst has about time-0 preferences and U^1 captures the information that the analyst has about time-1 preferences. A data set (U^0, U^1, f) is consistent with TC-EU if there is a way of picking a utility function consistent with the time-1 information such that the induced time-0 expected utility function is consistent with the information the analyst has about time-0 utility.

To characterize the preferences that are consistent with TC-EU, it will be necessary to consider lotteries over the alternatives, as we have done in the previous subsection. For any lottery $L \in \Delta(X)$ and utility function $u : X \to \mathbb{R}$, we define the associated expected utility $u(L) = \sum_{x\in X} L(x)u(x)$.

Definition 7 (Dominance with Respect to U^1)**.** We say that L' weakly dominates L if $u^1_\omega(L') \geq u^1_\omega(L)$ for all $u^1 \in U^1$ and $\omega \in \Omega$. We say that L' strictly dominates L if it weakly dominates L for each $u^1 \in U^1$, and there exists ω such that $u^1_\omega(L') > u^1_\omega(L)$.

The above definition is equivalent to first-order stochastic dominance whenever U^1_ω is the set of all utility functions consistent with a given ordinal preference $\preceq^1_\omega$. We generalize our definition of stochastic dominance violations accordingly.

Definition 8. A data set (U^0, U^1, f) exhibits stochastic dominance violations if there exist lotteries $L, L' \in \Delta(X)$ such that for every $u^0 \in U^0$ either

(i) L' weakly dominates L with respect to U^1 and $u^0(L) > u^0(L')$, or

(ii) L' strictly dominates L with respect to U^1 and $u^0(L) \geq u^0(L')$.

This definition facilitates the following generalization of Proposition 1:

Proposition 3. *A data set (U^0, U^1, f) is consistent with TC-EU if and only if it exhibits no stochastic dominance violations.*

4 Estimating Time Inconsistency

4.1 Data Sets with Identification of Time Inconsistency

We now propose a set of restrictions under which identification is possible. As a motivating example illustrating the intuition, consider the following hypothetical variation of the experiment by Read and van Leeuwen:

Example 3 (Read and van Leeuwen with Money)**.** Suppose again, as in the example in Section 1.2, that the agent chooses between a healthy snack and an unhealthy snack. As before, there are two time-1 states, where the agent either feels normal ($\omega = 1$) or gorged ($\omega = 2$), with the corresponding probabilities $f_1 = 0.8$ and $f_2 = 0.2$. But suppose now that the experimenter elicits—both at time 0 and at time 1—the maximal amount of money (to be received later at "time 2") that a person is willing to forego to receive their preferred option.[26] Given the small amounts of money involved, the experimenter assumes that the agent's preferences are (approximately) quasi-linear in the monetary amounts varied in the experiment, and that the marginal utility from money does not vary with the hunger state. Under these assumptions, TC-EU implies that the WTP for the healthy snack at time 0 must equal the average WTP at time 1.

Concretely, suppose that the agent has a WTP of \$1 for the healthy snack over the unhealthy snack at time 0. At time 1, the agent prefers the unhealthy snack by \$1 with

[26]For the purpose of this example, assume that both time-0 and time-1 decisions concern time-2 money, to eliminate any potential issues with money discounting. As we discuss later, the assumptions of this example are most likely to be satisfied when time 2 is reasonably far away from time 1.

probability 0.8, and prefers the healthy snack by \$5 with probability 0.2. Thus, because the agent has an average WTP for the healthy snack of $0.2 \times 5 - 0.8 \times 1 = \$0.20 < \$1$ at time 1, their behavior is inconsistent with TC-EU. Relative to the healthy snack, the agent values the unhealthy snack more at time 1 than at time 0, which might be explained by the pull of immediate gratification, as arising from models such as quasi-hyperbolic discounting.

We next formalize this idea for general environments, starting with two definitions.

Definition 9. Preferences $(\preceq^0, \preceq^1, f)$ on $X \subseteq Y \times Z$ have an *additively separable representation* (h^0, h^1, g^0, g^1) if there exist $h^0, h^1_\omega : Y \to \mathbb{R}$ and $g^0, g^1_\omega : Z \to \mathbb{R}$ such that for all states ω,

$$
\begin{aligned}
u^0(y,z) &= h^0(y) + g^0(z) \\
u^1_\omega(y,z) &= h^1_\omega(y) + g^1_\omega(z)
\end{aligned}
$$

are consistent with $\preceq^0$ and $\preceq^1_\omega$, respectively.[27]

Definition 10. A preference $\preceq$ over X is *responsive* if there exists a reference alternative y° such that for every pair $(y, y^\circ) \in Y^2$ there exists a pair $(z, z^\circ) \in Z^2$ such that $(y, z) \sim (y^\circ, z^\circ)$. An agent's preferences in a data set $(\preceq^0, \preceq^1, f)$ are responsive if $\preceq^0$ and $\prec^1_\omega$ are responsive for all ω, with the same reference alternative y°.

As Example 3 illustrated, the assumptions of separability and responsiveness correspond to the case where the agent's preference for receiving an alternative $y \in Y$ instead of $y^\circ \in Y$ can be "priced out" in units of Z.[28] The key additional assumption needed to use the priced-out valuations to point identify time preferences is that g is state-independent; i.e., the agent's valuations of alternatives in Z are state-independent. In the context of Example 3, this assumption amounts to the assumption that the agent's marginal utility of money does not vary with their appetite. The proposition below formalizes the general case.

Proposition 4. *If preferences in the data set $(\preceq^0, \preceq^1, f)$ are responsive, additively separable, and $g \equiv g^0 \equiv g^1_1 \equiv \ldots \equiv g^1_{|\Omega|}$ is known, then the following is point-identified:*

$$
\frac{h^0(y) - h^0(y^\circ)}{\sum_\omega f_\omega [h^1_\omega(y) - h^1_\omega(y^\circ)]} = \frac{g(z^\circ_{0,y}) - g(z_{0,y})}{\sum_\omega f_\omega [g(z^\circ_{\omega,y}) - g(z_{\omega,y})]}\,, \tag{3}
$$

where $(y, z_{0,y}) \sim^0 (y^\circ, z^\circ_{0,y})$ and $(y, z_{\omega,y}) \sim^1_\omega (y^\circ, z^\circ_{\omega,y})$ for all $\omega \in \Omega$.

[27] To be clear, we use $h^1 = (h^1_\omega)_{\omega\in\Omega}$ and $g^1 = (g^1_\omega)_{\omega\in\Omega}$ to denote vectors of possible time-1 utilities.

[28] Practically, a responsive data set can be easily generated using standard multiple price list or Becker–DeGroot–Marschak (BDM) techniques. The analyst simply needs to elicit how much money an agent is willing to forego to obtain their preferred option, and ensure enough range in monetary amounts to elicit the agent's maximum willingness to pay.

The right-hand-side of (3) is observable in the data set because g is known and the alternatives $z_{0,y}, z_{0,y^\circ}, z_{\omega,y}, z_{\omega,y^\circ}$ that make the agent indifferent between y and y° are observed in a responsive data set. Intuitively, the ratio on the right-hand-side captures how the agent's valuation of alternatives in Y, measured in units of Z, changes over time (on average). For a TC-EU agent, $h^0(y) = \sum_\omega f_\omega \, h^1_\omega(y)$ for all y, and thus the expression in (3) must equal 1.

An identification strategy corresponding to the logic of Proposition 4 is utilized in Augenblick and Rabin (2019), Augenblick (2018), and Fedyk (2018). We summarize the main idea in the example below:

Example 4 (Augenblick and Rabin 2019)**.** Augenblick and Rabin elicit willingness to work for various amounts of money. Suppose that preferences follow the quasi-hyperbolic discounting model and are given by

$$\begin{aligned} u^0(y,z) &= \left[\sum_\omega f_\omega \beta \, h^1_\omega(y)\right] + \beta \, g(z) \\ u^1(y,z) &= \quad h^1_\omega(y) + \beta \, g(z) \end{aligned} \tag{4}$$

where y is work at time 1 and z is compensation for this work, paid out later. Here, the analyst assumes that the marginal utility of money is independent of the marginal costs of effort in the experiment. Also, given the small stakes, the analyst has good reason to believe that utility is quasi-linear in money and thus sets $g(z) = z$.[29] As $h^0 \equiv \sum_\omega \beta f_\omega h^1_\omega$, we have that

$$\beta = \frac{h^0(y) - h^0(y^\circ)}{\sum_\omega f_\omega [h^1_\omega(y) - h^1_\omega(y^\circ)]}$$

and Proposition 4 implies that because g is known, β is point-identified. The intuition is that quasi-hyperbolic discounting implies that estimated willingness-to-accept for an extra unit of work must be higher by a proportion $^1/_\beta$, on average, in time 1 versus time 0.

As Examples 3 and 4 illustrate, one strategy to obtain point identification is to monetize agents' preferences over alternatives in Y, under conditions that ensure that it is plausible to assume that (i) preferences are separable over money and Y, (ii) the marginal utility of money does not vary with shocks to utility from alternatives in Y, and (iii) the analyst can estimate a utility function over money, either by assuming quasi-linearity or by using preferences over lotteries to estimate curvature. Although this set of conditions is restrictive, it is possible to plausibly approximate these conditions in the field, as has been done by Chaloupka, Levy, and White (2019), Carrera et al. (2022) and Allcott et al. (2022), building on the ideas in

[29] Alternatively, assume that the analyst gathers additional data on preferences over monetary lotteries to estimate the curvature of g.

DellaVigna and Malmendier (2004) and especially Acland and Levy (2015).

Of course, having a monetary domain is neither necessary nor sufficient to achieve point identification. In general, what is important is that at time 1, the agent does not update their expectation of utility over Z. This is most likely to be satisfied when (i) time 0 and time 1 are "close together," (ii) Y corresponds to consumption events realized at time 1, and (iii) Z corresponds to time-2 consumption events that are "far away" from time 1. For example, if time 1 is one week away from time 0, but Z corresponds to consumption or effort that is one year away, then it is unlikely that any information could be revealed between time 0 and time 1 that would alter the agent's expectation of utility from Z. This suggests strategies for point identification that don't involve money.[30] On the other hand, the conditions of Proposition 4 are unlikely to be satisfied if Y and Z correspond to large sums of money received at time 1 and time 2, with time 1 sufficiently far away from time 0 such that an agent subject to serially correlated liquidity and income shocks may plausibly update their beliefs at time 1 about their time-2 marginal utility of money.

4.2 Partial Identification of Quasi-hyperbolic Discounting

We next show provide a set identification result in a model where preferences are additively separable and responsive, but not independent along any of the dimensions. For expositional clarity, we focus on the quasi-hyperbolic discounting model with multiplicative taste shocks:

Definition 11 (Quasi-hyperbolic Discounting with Multiplicative Shocks)**.** A data set $(\preceq^0, \preceq^1, f)$ is consistent with quasi-hyperbolic discounting with multiplicative taste shocks if it is consistent with utilities of the form

$$\begin{aligned} u^0(y,z) &= \sum_{\omega\in\Omega} f_\omega \left(\theta^1_\omega h(y) + \theta^2_\omega g(z)\right) \\ u^1(y,z) &= \theta^1_\omega\, h(y) + \beta\, \theta^2_\omega\, g(z)\,. \end{aligned} \tag{5}$$

Note that if preferences can be represented as in (5), then there exists an additively separable representation (h^0, h^1, g^0, g^1) of the preferences where the function g^1 does not depend on the state and equals g^0. To obtain this representation, simply set $h^0 \equiv \frac{\sum_\omega f_\omega \theta^1_\omega}{\sum_\omega f_\omega \theta^2_\omega} h$,

[30] For a concrete example, consider a modification of the example in Section 1.1, but suppose now that time 2 corresponds to work that must be completed in one year. In this modified design, it is plausible to assume that the agent does not update their beliefs about the time-2 effort costs between time 0 and time 1, and thus that the expected cost of time-2 effort is constant across all states of the world. Thus, if the analyst is justified in assuming separable effort costs over time, so that choices from multiple budget sets identify time-2 effort costs, then Proposition 4 implies that β is identified. Thus, the convex time budget approach can be used to identify time preferences as long as the time-2 consumption/effort events are sufficiently far in time that the agent does not update their beliefs about time-2 utility between time 0 and time 1.

$h^1_\omega \equiv \frac{\theta^1_\omega}{\beta\theta^2_\omega} h$, $g^0 = g$ and $g^1 = \beta g$.

To obtain intuition for the types of inferences that can be made about the parameter β given data consistent with quasi-hyperbolic discounting with multiplicative taste shocks, consider an identification strategy that (wrongly) assumes no taste shocks and assumes instead that all differences in time 1 are due to variation in the time-preference parameter β. That is, if at time 0 the agent is indifferent between $(y, z_{0,y}) \sim^0 (y^\circ, z^\circ_{0,y})$ and at time 1 is indifferent between $(y, z_{\omega,y}) \sim^1_\omega (y^\circ, z^\circ_{\omega,y})$, then

$$\begin{aligned} h(y) + \quad g(z_{0,y}) &= h(y^\circ) + \quad g(z^\circ_{0,y}) \\ h(y) + \hat{\beta}_\omega g(z_{\omega,y}) &= h(y^\circ) + \hat{\beta}_\omega g(z^\circ_{\omega,y})\,. \end{aligned}$$

Rearranging these equations yields that

$$\hat{\beta}_\omega = \frac{g(z^\circ_{0,y}) - g(z_{0,y})}{g(z^\circ_{\omega,y}) - g(z_{\omega,y})}\,. \tag{6}$$

We use the "hat" notation in the definition above because $\hat{\beta}_\omega$ can also be thought of as a "noisy" estimate of the true present focus parameter β, which is a common statistic to report in empirical studies. An immediate corollary of Proposition 4 is that if utility over the Z dimension is state-independent, then β is point-identified and given by

$$\beta = \frac{1}{\sum_{\omega\in\Omega} f_\omega \hat{\beta}_\omega^{-1}}\,. \tag{7}$$

Without the state independence assumption, the range of the distribution of $\hat{\beta}_\omega$ identifies the range of possible values of β consistent with the data. This result provides a formal generalization of the numerical example in Section 1.1.

Proposition 5. *A responsive data set $(\preceq^0, \preceq^1, f)$ is consistent with quasi-hyperbolic discounting with multiplicative taste shocks if and only if*

(i) it has an additively separable representation (h^0, h^1, g^0, g^1) where for all $\omega \in \Omega$, $g^0 = g^1_\omega = g$ and $\frac{h^1_\omega(y)}{h^0(y)}$ is non-negative and constant in $y \in Y$, and

(ii) $\beta \in \left(\min_\omega \hat{\beta}_\omega, \max_\omega \hat{\beta}_\omega\right)$ for $\hat{\beta}_\omega = \frac{g(z^\circ_{0,y}) - g(z_{0,y})}{g(z^\circ_{\omega,y}) - g(z_{\omega,y})}$.

The range of $\hat{\beta}_\omega$ is estimable, and is frequently reported, which facilitates the application of Proposition 5. In practice, because many experiments feature only one revision observation per person, $\hat{\beta}_\omega$ is obtained as an individual-level estimate of present focus—we discuss this in greater detail in Appendix E. To illustrate what distributions of $\hat{\beta}$ look like in practice, the left panel of Figure 1 presents estimates $\hat{\beta}$ reported in Augenblick, Niederle, and Sprenger (2015), while the right panel presents estimates reported in Augenblick and Rabin (2019). Under the assumption of state-independent preferences over one dimension made in Proposition 4, the population (average) present focus is obtained from the average of these point estimates.[31] In Section 4.1 we have argued that the assumptions of Proposition 4 are plausible for Augenblick and Rabin (2019), but in Section 1.1 we have illustrated that they are less plausible for Augenblick, Niederle, and Sprenger (2015). Instead, under the assumptions of Proposition 5, the population β in Augenblick, Niederle, and Sprenger (2015) can be any value between the minimum and maximum of the support of the estimated $\hat{\beta}$. As we show in Appendix E.2, this implies that the data is consistent with any β in the range $[0.66, 1.12]$, even after 20 percent of the most extreme observations of the estimated distribution of $\hat{\beta}$ are excluded. This range includes values of population β that imply both future focus and significant present focus, consistent with the types of illustrative calculations we presented in Section 1.1.

Figure 1: Distribution of estimates $\hat{\beta}$ from two studies

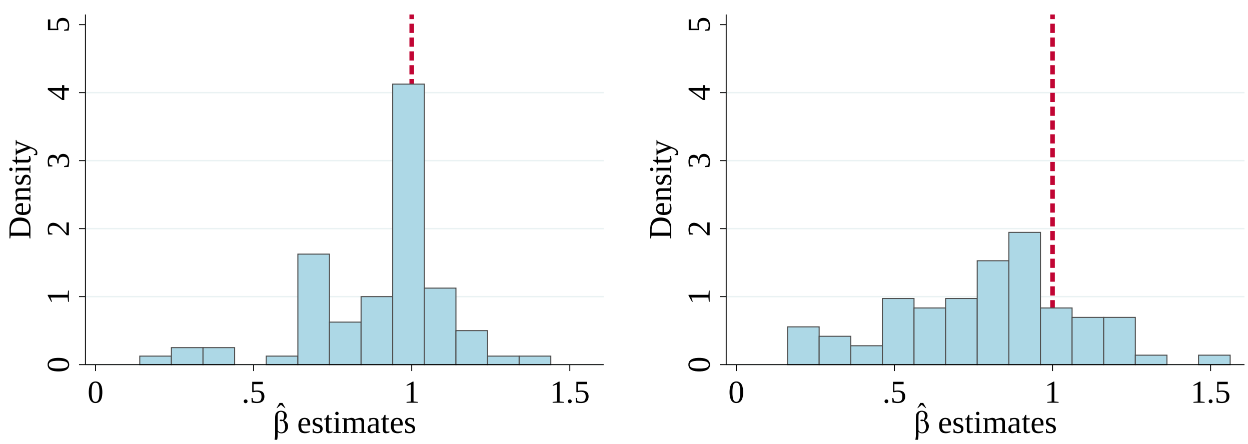

Notes: The left panel presents the histogram of of individual-level estimates $\hat{\beta}$ using the first experimental block in Augenblick, Niederle, and Sprenger (2015), as reported in the top panel of their Figure 6. The right panel presents the histogram of individual-level estimates $\hat{\beta}$ in Augenblick and Rabin (2019), as reported in the left panel of their Figure 6. The formatting of the histogram from Augenblick, Niederle, and Sprenger (2015) is slightly altered to match the formatting choice of the histogram in Augenblick and Rabin (2019).

[31]See Appendix E.2 for more details.

5 Discussion

5.1 Are Taste Shocks Important?

In principle, the analyst could simply assume that no individuals face uncertainty between time 0 and time 1, and thus that all differences in time-1 choices between individuals making the same time-0 choice are due solely to individual differences in time preferences.Both intuition and evidence suggest that this assumption is unrealistic. For example, Read and van Leeuwen (1998) experimentally show that variation in individuals' satiation has a large effect on their preferences for healthy versus unhealthy foods at time 1.

The strong assumption of no random taste shocks also implies that because any difference between time-0 and time-1 choice is due to (stable) time preferences, these differences should be stable within individuals over time. To test this assumption, the analyst needs to estimate individuals' time-0 and time-1 preferences on two (or more) separate occasions, $j \in \{1, 2\}$. If on both occasions the analyst obtains an identical measure of time preference for each individual (e.g., $\hat{\beta}_i^1 = \hat{\beta}_i^2$ for all agents i in the quasi-hyperbolic discounting model), then the assumption is satisfied. For example, Sadoff, Samek, and Sprenger (2019) include two sets of decisions to study the stability of their measure of time inconsistency. Sadoff, Samek, and Sprenger find correlations between 0.2 and 0.33 in agents' propensity to switch toward unhealthier foods at time 1, which suggests that stable individual differences explain $0.2^2 = 0.04$ to $0.33^2 = 0.11$ of the variance in time-1 revisions of time-0 choices, leaving random taste shocks and other forms of updating to explain the remaining variation. Similarly, as we show in Appendix E.1, stable individual differences in time preferences can explain at most 2 percent of the variation in time-1 revisions of time-0 preferences in the Augenblick, Niederle, and Sprenger (2015) data.

5.2 Commitment

Although demand for choice set restrictions cannot be used to point identify time-preference parameters (see, e.g., Carrera et al., 2022),[32] in some cases it serves as a useful correlate of time preferences that can help refine their set identification. For example, Augenblick, Niederle, and Sprenger (2015) find that their measure of demand for commitment relates negatively to their estimates of the present-focus parameter β, and thus provides evidence of some stable individual differences in time preferences. In Appendix E.2 we illustrate

[32] Commitment take-up is a coarse measure that might lead to false negatives in tests of time inconsistency because uncertainty and thus demand for flexibility reduce demand for choice set restrictions (Heidhues and Kőszegi, 2009, Laibson, 2015, Carrera et al., 2022), and can also deliver false positives because of noise in take-up decisions (Carrera et al., 2022).

how this finding can be used in conjunction with Proposition 5 to refine the set of time preferences that could be consistent with individuals' full set of choices. Appendix E.2 also illustrates the simple but important point that augmenting revision design data with data on commitment contract take-up does not restore point identification of time-preference parameters, for largely the same reasons that point identification on the full experimental sample was not possible in the first place.

5.3 Heterogeneity

Our results concern a data set in which the analyst observes an agent's time-0 preference and the *full* distribution of time-1 preferences. In practice, data sets are less rich, which implies that our results are a "best case" for identification. A more typical data set consists of a large population of individuals who each make a single choice at time 0 and at time 1.

One approach to analyzing such data sets is to assume that all individuals who make the same time-0 choice are homogeneous both in preferences and the economic environment, and consequently treat variation in time-1 decisions as due to realization of uncertainty. Under this assumption, the set of all individuals who make the same time-0 choice constitutes the kinds of data sets that we study in this paper.

Even without this homogeneity assumption, however, most of our key results still apply, as we show in Appendix B and further illustrate in Appendix E.2.

5.4 Costly Self Control

We assume the domain of TC-EU preferences to be the set of alternatives X, which rules out models of *costly self-control* (see e.g., Gul and Pesendorfer, 2001, Fudenberg and Levine, 2006). Thus, data sets that we show to reject TC-EU could be consistent either with time-inconsistent preferences or with costly self-control preferences.

5.5 Other Related Literature

Related mathematical results Propositions 1 and 3 are related to existing results in the literatures on social choice, dynamically-consistent preferences over acts, and random utility models. Our main results in Theorems 1 and 2 do not, to our knowledge, resemble existing mathematical results. In Appendix D we flesh out connections to technical results in the literatures on social choice, random utility models, and dynamically consistent preferences over acts.

"Static" preference reversals and money discounting. While the focus of this paper is on dynamic preference reversals, there is also an important literature on static preference reversals. A large literature studies how people trade off between the size of a monetary reward and its delay, under the assumption that time-dated payments and time-dated utils are interchangeable (see, e.g., Cohen et al. 2020 for a review). A smaller literature conducts such studies with real consumption events, such as juice squirts (McClure et al., 2007). These studies are an important and complementary source of evidence, although some of the results concerning monetary rewards could be due to violations of stationarity rather than time consistency (Halevy, 2015), or due to uncertainty changing with delay (Halevy, 2005, 2008, Andreoni and Sprenger, 2012, Chakraborty, Halevy, and Saito, 2020).

While there is an important ongoing conversation about whether time-dated monetary rewards should be treated interchangeably with time-dated utils (see, e.g., Ericson and Laibson 2019 and Cohen et al. 2020 for reviews), our results suggest one possible additional reason why studies such as those of Augenblick, Niederle, and Sprenger (2015) estimate fewer preference reversals for money versus actual consumption: utility over money may be subject to fewer random taste shocks.[33]

6 Conclusion

Our general characterization of consistency with TC-EU shows that in typical revision designs, it is difficult to identify the degree of time inconsistency. The difficulty arises from random taste shocks or other arrival of information, which are particularly plausible in more complex and economically consequential field settings. However, we have also provided guidance on the types of economic environments where the assumptions required for point identification are plausible. Thus, while identification of time inconsistency may be more difficult than initially intuited, it is certainly theoretically and empirically feasible.

Of course, our results do not imply that nothing can be learned from data sets where we show that it is not possible to formally reject TC-EU. For example, it is unlikely to be mere coincidence or file-drawer bias that in most circulated papers, the systematic reversals tend

[33]There are a number of other papers (see Imai, Rutter, and Camerer 2021 for a review) that follow Andreoni and Sprenger (2012) in applying the convex time budget (CTB) approach to time-dated monetary rewards. Our paper largely focuses its discussion and examples on revision designs with real consumption events because of the ambiguities highlighted by the ongoing conversation about interpreting preferences over time-dated monetary rewards. But given an interpretation, our results apply to those designs as well. For example, if the monetary rewards are treated as real consumption events, but the subjects are not plausibly subject to liquidity shocks that change their marginal utility of money, then our results imply that monetary CTBs constitute deterministic environments that allow point identification of time preferences.

to be toward more immediately gratifying options.[34] Just as proper Bayesian scientists reservedly update about causal relationships from all well-measured associations—even when the associations are not produced by experimental or quasi-experimental techniques—we think it is appropriate to carefully update from all revision design data. Correlational analyses that link choice revisions to supplementary proxies of time inconsistency or observable determinants of taste shocks (see, e.g., Augenblick, Niederle, and Sprenger, 2015, Sadoff, Samek, and Sprenger, 2019) can bolster the updating. At the same time, by formally studying identification in a general theoretical framework, this paper clarifies just how strong the assumptions for (point) identification of time inconsistency have to be, and helps identify the most theoretically robust designs. We hope that this will help further the important agenda of measuring time inconsistency.

References

Abbring, Jaap H. and Øystein Daljord. 2020. "Identifying the Discount Factor in Dynamic Discrete Choice Models." *Quantitative Economics* 11 (2):471–501.

Abebe, Girum, Stefano Caria, and Esteban Ortiz-Ospina. 2021. "The Selection of Talent: Experimental and Structural Evidence from Ethiopia." *American Economic Review* 111 (6):1757–1806.

Acland, Dan and Matthew R. Levy. 2015. "Naiveté, Projection Bias, and Habit Formation in Gym Attendance." *Management Science* 61 (1):146–160.

Ahn, David and Todd Sarver. 2013. "Preference for Flexibility and Random Choice." *Econometrica* 81 (6):341–361.

Aliprantis, Charalambos D. and Kim C. Border. 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3 ed.

Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman. 2022. "Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending." *Review of Economic Studies* 89 (3):1041–1084.

Alós-Ferrer, Carlos, Ernst Fehr, and Nick Netzer. 2021. "Time will tell: Recovering preferences when choices are noisy." *Journal of Political Economy* 129 (6):1828–1877.

Andreoni, James, Michael Callen, Karrar Hussain, Muhammad Yasir Khan, and Charles Sprenger. 2020. "Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan." NBER Working Paper No. 22019.

Andreoni, James and Charles Sprenger. 2012. "Estimating Time Preferences from Convex Budgets." *American Economic Review* 102 (7):3333–3356.

Augenblick, Ned. 2018. "Short-Term Discounting of Unpleasant Tasks." *Working Paper* .

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks." *The Quarterly Journal of Economics* 130 (3):1067–1115. URL `https://econpapers.repec.org/article/oupqjecon/v_3a130_3ay_3a2015_3ai_3a3_3ap_3a1067-1115..htm`.

[34]Based on their meta-analysis, Imai, Rutter, and Camerer (2021) suggest that selective reporting is modest in revision designs studying effort allocation tasks using the convex time budget approach.

Augenblick, Ned and Matthew Rabin. 2019. "An Experiment on Time Preference and Misprediction in Unpleasant Tasks." *The Review of Economic Studies* 86 (3):941–975. URL `https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdy019/4996235`.

Badger, Gary J., Warren K. Bickel, Louis A. Giordano, Eric A. Jacobs, George Loewenstein, and Lisa Marsch. 2007. "Altered States: The Impact of Immediate Craving on the Valuation of Current and Future Opioids." *Journal of Health Economics* 26 (5):865–876.

Bai, Liang, Benjamin Handel, Ted Miguel, and Gautam Rao. forthcoming. "Self-Control and Demand for Preventive Health: Evidence from Hypertension in India." *Review of Economics and Statistics* .

Barton, Blake. 2015. "Interpersonal Time Inconsistency and Commitment." *Working Paper* .

Bartos, Vojtech, Michal Bauer, Julie Chytilová, and Ian Levely. forthcoming. "Psychological Effects of Poverty on Time Preferences." *The Economic Journal* .

Block, Henry David, Jacob Marschak et al. 1959. "Random orderings and stochastic theories of response." Tech. rep., Cowles Foundation for Research in Economics, Yale University.

Breig, Zachary, Matthew Gibson, and Jeffrey G. Shrader. 2021. "Why Do We Procrastinate? Present Bias and Optimism." *Working Paper* .

Browning, Martin and Jeremy Tobacman. 2015. "Discounting and Optimism Equivalances." *Working Paper* .

Brunnermeier, Markus K., Filippos Papakonstantinou, and Jonathan A. Parker. 2008. "An Economic Model of the Planning Fallacy." *Working Paper* .

Buehler, Roger, Dale Griffin, and Johanna Peetz. 2010. "Chapter One - The Planning Fallacy: Cognitive, Motivational, and Social Origins." In *Advances in Experimental Social Psychology*, vol. 43, edited by Mark P. Zanna and James M. Olson. Academic Press, 1–62.

Carrera, Mariana, Heather Royer, Mark Stehr, Justin Sydnor, and Dmitry Taubinsky. 2022. "Who Chooses Commitment? Evidence and Welfare Implications." *The Review of Economic Studies* 89 (1):1205–1244.

Carroll, Gabriel. 2010. "An Efficiency Theorem for Incompletely Known Preferences." *Journal of Economic Theory* 145 (6):2463–2470.

Chakraborty, Anujit, Yoram Halevy, and Kota Saito. 2020. "The Relation between Behavior under Risk and over Time." *American Economic Review: Insights* 2 (1):1–16.

Chaloupka, Frank J., Matthew R. Levy, and Justin S. White. 2019. "Estimating Biases in Smoking Cessation: Evidence from a Field Experiment." *Working Paper* .

Clark, Stephen A. 1996. "The random utility model with an infinite choice space." *Economic Theory* 7 (1):179–189.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White. 2020. "Measuring Time Preferences." *Journal of Economic Literature* 58 (2):299–347.

Corbett, Colin. 2016. "Preferences for Effort and Their Applications." Unpublished.

De Oliveira, Henrique and Rohit Lamba. 2021. "Rationalizing dynamic choices." *Available at SSRN 3332092* .

Dekel, Eddie and Barton L. Lipman. 2012. "Costly Self-Control and Random Self-Indulgence." *Econometrica* 80:1271–1302.

DellaVigna, Stefano and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence." *The Quarterly Journal of Economics* 119 (2):353–402. URL `+http://dx.doi.org/10.1162/003355304323024098`.

———. 2006. "Paying Not to Go to the Gym." *American Economic Review* 96 (3):694–719.

Epstein, Larry G. and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4):937–969.

Ericson, Keith M. and David Laibson. 2019. "Intertemporal Choice." In *Handbook of Behavioral Economics*, vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, chap. 1. Elsevier, 1–67.

Fedyk, Anastassia. 2018. "Asymmetric Naivete: Beliefs about Self-Control." *Working Paper* .

Fishburn, P. C. 1998. "Stochastic Utility." *Handbook of Utility Theory* 1.

Frick, Mira, Ryota Iijima, and Tomasz Strzalecki. 2019. "Dynamic Random Utility." *Econometrica* 87 (6):1941–2002.

Fudenberg, Drew and David K. Levine. 2006. "A Dual-Self Model of Impulse Control." *American Economic Review* 96 (5):1449–1476.

Fudenberg, Drew and Tomasz Strzalecki. 2015. "Dynamic logit with choice aversion." *Econometrica* 83 (2):651–691.

Ghirardato, Paolo. 2002. "Revisiting Savage in a conditional world." *Economic Theory* 20:83–92.

Gul, Faruk and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control." *Econometrica* 69:1403–1435.

———. 2006. "Random expected utility." *Econometrica* 74 (1):121–146.

Halevy, Yoram. 2005. "Diminishing Impatience: Disentangling Time Preference from Uncertain Lifetime." *Working Paper* .

———. 2008. "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect." *American Economic Review* 98 (3):1145–162.

———. 2015. "Time Consistency: Stationarity and Time Invariance." *Econometrica* 83 (1):335–352.

Heidhues, Paul and Botond Kőszegi. 2009. "Futile Attempts at Self-Control." *Journal of the European Economic Association* 7 (2):423–434. URL `https://academic.oup.com/jeea/article-lookup/doi/10.1162/JEEA.2009.7.2-3.423`.

Heidhues, Paul and Philipp Strack. 2020. "Identifying present-bias from the timing of choices." *American Economic Review* .

Imai, Taisuke, Tom A. Rutter, and Colin F. Camerer. 2021. "Meta-Analysis of Present-Bias Estimation Using Convex Time Budgets." *The Economic Journal* 131:1788–1814.

Imas, Alex, Michael Kuhn, and Vera Mironova. forthcoming. "Waiting to Choose: The Role of Deliberation in Intertemporal Choice." *American Economic Journal: Microeconomics* .

Jackson, Matthew O. and Leeat Yariv. 2015. "Collective Dynamic Choice: The Necessity of Time Inconsistency." *American Economic Journal: Microeconomics* 7 (4):150–78.

Kahneman, Daniel and Amos Tversky. 1982. "Intuitive Prediction: Biases and Corrective Procedures." In *Judgment Under Uncertainty: Heuristics and Biases*, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky. Cambridge University Press, 414–421.

Kölle, Felix and Lukas Wenner. 2019. "Time-Inconsistent Generosity: Present Bias Across Individual and Social Contexts." *Working Paper* .

Kuchler, Theresa and Michaela Pagel. 2021. "Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown." *Journal of Financial Economics* 139 (2):359–388.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2):443–478.

———. 2015. "Why Don't Present-Biased Agents Make Commitments?" *American Economic Review* 105 (5):267–272.

Levy, Matthew and Pasquale Schiraldi. 2021. "Identification of Intertemporal Preferences from

Choice Set Variation." *Working Paper* .

Magnac, Therry and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70 (2):801–816.

Mahajan, Aprajit, Christian Michel, and Alessandro Tarozzi. 2020. "Identification of Time-Inconsistent Models: The Case of Insecticide Treated Nets." *Working Paper* .

McClure, Samuel M., Keith M. Ericson, David I. Laibson, George Loewenstein, and Johanthan D. Cohen. 2007. "Time Discounting for Primary Rewards." *Journal of Neuroscience* 27:5796–5804.

McFadden, Daniel and Marcel K Richter. 1990. "Stochastic rationality and revealed stochastic preference. Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz."

McLennan, Andrew. 2002. "Ordinal Efficiency and the Polyhedral Separating Hyperplane Theorem." *Journal of Economic Theory* 105 (2):435–449.

Milkman, Katherine L., Todd Rogers, and Max H. Bazerman. 2009. "Highbrow Films Gather Dust: Time-Inconsistent Preferences and Online DVD Rentals." *Management Science* 55 (6):1047–1059.

Millner, Antony. 2020. "Nondogmatic Social Discounting." *American Economic Review* 110 (3):760–75.

O'Donoghue, Ted and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review* 89 (1):103–124.

Read, Daniel, George Loewenstein, and Shobana Kalyanaraman. 1999. "Mixing Virtue and Vice: Combining the Immediacy Effect and the Diversification Heuristic." *Journal of Behavioral Decision Making* 12 (4):257–273.

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

Sadoff, Sally, Anya Samek, and Charles Sprenger. 2019. "Dynamic Inconsistency in Food Choice: Experimental Evidence from a Food Desert." *Review of Economic Studies* 87 (4):1–35.

Shmaya, Eran and Leeat Yariv. 2016. "Experiments on Decisions under Uncertainty: A Theoretical Framework." *American Economic Review* 106 (7):1775–1801. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20120978`.

Strotz, Robert H. 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies* :165–180.

Strzalecki, Tomasz. 2021. "Lecture notes in Decision Theory."

Appendix

Table of Contents

A Proofs

The appendix proceeds as follows: Section A.1 derives several results on the aggregation of incomplete preorders. Section A.2 derives several results on the separation of finite-dimensional cones generated by sets of utility functions. In Section A.3 we use these results to formalize the proof sketches in the body of the paper.

A.1 Results on Aggregating Incomplete Preorders

Recall that we defined a new preorder $\preceq^1_*$ that ranks one alternative y weakly higher than an alternative x if the agent ranks that alternative higher in all states ω

$$x \preceq^1 y \Leftrightarrow x \preceq^1_\omega y \text{ for all } \omega \in \Omega\,.$$

We thus denote by $\preceq^1_*$ the preorder that is generated by agreement of the preorders $\preceq^1_\omega$ in the different states ω. We define another binary relation $\trianglelefteq^*$ such that y is weakly preferred to x if it is either preferred according to the time-0 preference or according to all time-1 preferences

$$x \trianglelefteq^* y \quad \Leftrightarrow \quad x \preceq^0 y \text{ or } x \preceq^1_* y\,. \tag{8}$$

We define $\triangleleft^*$ to be the asymmetric component of $\trianglelefteq^*$. We note that $\trianglelefteq^*$ need not to be transitive and define $\trianglelefteq$ to be the smallest transitive closure of $\trianglelefteq^*$. We define $\sim_\trianglelefteq$ to be the symmetric part of $\trianglelefteq$, and $\triangleleft$ to be the asymmetric part of $\trianglelefteq$.

Lemma 1. *If the data set admits no simple dominance violations then*

$$x \triangleleft^* y \quad \Leftrightarrow \quad x \prec^0 y \text{ or } x \prec^1_* y.$$

Proof. We first note that $x \triangleleft^* y$ if $x \trianglelefteq^* y$ and neither $y \preceq^1_* x$ nor $y \preceq^0 x$. We furthermore note that if there is no simple dominance violation, then $x \prec^0 y$ implies that we do not have $y \preceq^1_* x$, and $x \prec^1_* y$ implies that we do not have $y \preceq^0 x$. Hence, $x \prec^0 y$ or $x \prec^1_* y$ implies $x \triangleleft^* y$.

To see that the converse direction also holds, note that $x \triangleleft^* y$ implies that neither $y \preceq^0 x$ nor $y \preceq^1_* x$, and either $x \preceq^0 y$ or $x \preceq^1_* y$, which together implies that either $x \prec^0 y$ or $x \prec^1_* y$. □

We next translate the condition of no cyclic dominance violations in the data set ($\preceq^0, \preceq^1$ $, f$) into a condition on the induced order $\trianglelefteq^*$.

Definition 12 (Only Weak Cycles)**.** We say that $\trianglelefteq^*$ admits *only weak cycles* if $x_1 \trianglelefteq^* \cdots \trianglelefteq^* x_n \trianglelefteq^* x_1$ implies that $x_1 \sim_{\trianglelefteq^*} \cdots \sim_{\trianglelefteq^*} x_n$.

Lemma 2. *The following are equivalent:*

(i) The data set has no cyclic dominance violation.

(ii) The data set has no simple dominance violation and $\trianglelefteq^*$ satisfies the only weak cycles condition.

Proof. $(ii) \Rightarrow (i)$: Suppose that the elements $x_1, \ldots, x_k$ constitute a cyclic dominance violation. Without loss we can assume that the first inequality is strict (otherwise reorder the elements) and according to the time-0 order (the argument for the time-1 order is identical)

$$x_1 \prec^0 x_2 \preceq^1_* x_3 \preceq^0 \ldots \preceq^0 x_k \preceq^1_* x_1\,.$$

If there is no simple dominance violation, then Lemma 1 implies that

$$x_1 \lhd^* x_2 \unlhd^* x_3 \unlhd^* \ldots \unlhd^* x_k \unlhd^* x_1 \,,$$

and thus a violation of the only weak cycles condition.

$(i) \Rightarrow (ii)$: Suppose (ii) does not hold. If the data set has a simple dominance violation it also has a cyclic dominance violation, as any simple dominance violation is a cyclic dominance violation with a cycle of length 2. Thus, suppose that there is no simple dominance violation but that there is a non-weak cycle in $\unlhd^*$ involving $x_1, \ldots, x_k$. Without loss, assume that $x_k \lhd^* x_1$. Then Lemma 1 implies that there exist alternatives $x_1, \ldots, x_k$ such that:

(i) For all $j \leq k-1$, either $x_j \preceq^0 x_{j+1}$ or $x_j \preceq^1_* x_{j+1}$

(ii) Either $x_k \prec^0 x_1$ or $x_k \prec^1_* x_1$

Now if $x_j \preceq^0 x_{j+1} \preceq^0 x_{j+2}$, then $x_j \preceq^0 x_{j+2}$, and thus there is a non-weak cycle over the set $\{x_1, \ldots, x_k\} \setminus \{x_{j+1}\}$. A similar statement applies to three adjacent alternatives in a cycle related by $\preceq^1_*$. Thus, any non-weak cycle can be reduced to a non-weak cycle where no three adjacent alternatives are in increasing order according to $\preceq^0$ or according to $\preceq^1_*$; this non-weak cycle amounts to a cyclic dominance violation. □

Lemma 3. *If $\preceq^0$ is complete and the data set $(\preceq^0, \preceq^1, f)$ exhibits cyclic dominance violations then there exists a simple dominance violation.*

Proof. We prove this result by contraposition. Assume that there is no simple dominance violation. Then $\unlhd^*$ preserves the asymmetric part of $\preceq^0$ by Lemma 1. As $\preceq^0$ is by assumption complete, $\unlhd^*$ must be complete as well. As $\preceq^0$ is transitive, $\unlhd^*$ can only admit weak cycles. But then Lemma 2 implies that there is no cyclic dominance violation. □

Lemma 4. *Suppose that $\unlhd^*$ satisfies the only weak cycles condition. Then $x \lhd^* y \Rightarrow x \lhd y$.*

Proof. Suppose that $x \lhd^* y$. Then the only way for $y \unlhd x$ is if there is a cycle in $\unlhd^*$ involving x and y. Since $x \lhd^* y$, this cycle is not a weak cycle. □

Lemma 5. *Suppose that $\preceq^0$ and $\preceq^1$ are single-peaked, and that there are no simple dominance violations. Then there exists a set $X^* \subseteq \{\min X, \max X\}$ such that $x \unlhd x^*$ does not hold for any $x \in X \setminus X^*$ and $x^* \in X^*$, and such that $\min X \sim_{\unlhd} \max X$ if $X^* = \{\min X, \max X\}$.*

Proof. Consider the set of alternatives that is not strictly better than any other alternative according to the preorder $\unlhd$

$$\mathcal{X} = \{x \in X \colon \text{ there does not exist } y \in X \text{ with } y \lhd x\} \,.$$

We first argue that $\mathcal{X} \cap \{\min X, \max X\} \neq \emptyset$. If not, then there would exist $y \in X$ such that either $\min X \succ^0 y$ or $\min X \succ^1_* y$. Suppose first that $\min X \succ^0 y$. Single-peakness implies that if $y \neq \max X$ then $\min X \succ^0 y \succeq^0 \max X$, so that $\min X \succ^0 \max X$. Thus, $\min X \succ^0 \max X$ if $\min X \notin \mathcal{X}$. Now if $\max X \notin \mathcal{X}$, then a similar argument shows that $\min X \prec^1_* \max X$ (since $\min X \prec^0 \max X$ is impossible by transitivity), which implies a simple dominance violation. Thus, if $\min X \succ^0 y$, then there is a simple dominance violation. A symmetric argument also shows that $\min X \succ^1_* y$ implies a simple dominance violation.

Next, pick a maximal subset of $\mathcal{X}$, X^*, such that $X^* \cap \{\min X, \max X\} \neq \emptyset$, and such that $x \sim_{\trianglelefteq} x'$ for all $x, x' \in X^*$. Without loss, assume that $\min X \in X^*$. We now argue that $X^* \subseteq \{\min X, \max X\}$. If not, then there is a $y \in X \setminus \{\min X, \max X\}$ such that $y \sim_{\trianglelefteq^*} \min X$. Then, because we assumed no simple dominance violations, either

(i) there is some $y \in X \setminus \{\min X, \max X\}$ such that either $y \sim^0 \min X$ or $y \sim^1_* \min X$, or

(ii) $\min X$ and y are part of a cycle in $\trianglelefteq^*$ that is not a weak cycle

To see why (ii) must hold if (i) does not, note that if $\min X$ were part of a *weak* cycle that relates it to y through $\sim_{\trianglelefteq^*}$, and condition (i) did not hold for *any* $y' \in X \setminus \{\min X, \max X\}$, then there would need to be simple dominance violations to generate the indifferences in the weak cycle.

We next argue that in both cases (i) and (ii) above, either $\min X \succ^0 \max X$ or $\min X \succ^1_* \max X$. We then show that neither possibility is inconsistent with there being no simple dominance violations.

Case (i) If $y \sim^0 \min X$ then single-peakness implies that $y \succ^0 \max X$ and thus $\min X \succ^0 \max X$. An identical argument shows that $y \sim^1_* \min X$ implies that $\min X \succ^1_* \max X$.

Case (ii) If $\min X$ is part of a cycle in $\trianglelefteq^*$ that is not a weak cycle, then there is some $y \in X \setminus \{\min X, \max X\}$ such that either (i) $y \sim^0 \min X$ or (ii) $y \sim^1_* \min X$ or (iii) $y \prec^0 \min X$ or (iv) $y \prec^1_* \min X$. In the first two cases, we have already shown that $\min X \succ^0 \max X$ or $\min X \succ^1_* \max X$. In the second two cases, the first paragraph of this proof shows that either $\min X \succ^0 \max X$ or $\min X \succ^1_* \max X$.

Thus, if $X^* \subseteq \{\min X, \max X\}$ does not hold, then either $\min X \succ^0 \max X$ or $\min X \succ^1_* \max X$, and thus $\max X \trianglelefteq \min X$. But since $\min X \in X^*$, the definition of X^* requires that $\min X \sim_{\trianglelefteq} \max X$, and thus that $\max X \in X^*$. Moreover, since either $\min X \succ^0 \max X$ or $\min X \succ^1_* \max X$, the only way for $\min X \sim_{\trianglelefteq} \max X$ to hold in the absence of simple dominance violations is if there is a non-weak cycle involving both $\min X$ and $\max X$. But then, reasoning identical to Case (ii) above implies that we must have $\max X \succ^0 \min X$ or $\max X \succ^1_* \min X$, which implies a simple dominance violation. $\square$

A.2 Results on the Cones Generated from Sets of Utility Functions

As there are only finitely many alternatives we will throughout identify utilities with vectors in $\mathbb{R}^{|X|}$.

For any sets A and B, we define the Minkowsky sum $A + B := \{a + b \mid a \in A, b \in B\}$. We say that a set $A \subset \mathbb{R}^n$ is *constant shift invariant* if $a \in A$ implies that $a + (\lambda, \ldots, \lambda) \in A$ for all $\lambda \in \mathbb{R}$.

Lemma 6 (Cone Separation Lemma)**.** *Suppose that $A, B \subset \mathbb{R}^n$ are convex cones that are open relative to their affine hull and constant shift invariant with $A \cap B = \emptyset$. Then there exists a vector $p \in \mathbb{R}^n$ with $p \neq 0$ and $\sum_{i=1}^n p_i = 0$, such that for all $a \in A, b \in B$*

$$p \cdot a \geq 0 \geq p \cdot b$$

and one of the inequalities is strict for all $a \in A$ and $b \in B$.

Proof. As A, B are disjoint they can be properly separated by a hyperplane; i.e., there exists $p \in \mathbb{R}^n$ with $p \neq 0$ and $c \in \mathbb{R}$ such that for all $a \in A, b \in B$

$$p \cdot a \geq c \geq p \cdot b$$

with at least one inequality strict. As A is constant shift invariant, $a + (\lambda, \ldots, \lambda) \in A$ if $a \in A$, which implies that for all $\lambda \in \mathbb{R}$

$$p \cdot a + \lambda \sum_{i=1}^{n} p_i \geq c\,.$$

The above inequality can only hold for all $\lambda \in \mathbb{R}$ if $\sum_{i=1}^n p_i = 0$, which thus must hold.

Similarly, as A is a cone, $a \in A$ implies that $\lambda a \in A$ for all $\lambda > 0$ and hence

$$\lambda(p \cdot a) \geq c\,.$$

Taking the limit $\lambda \to 0$ yields that $0 \geq c$. Applying the same argument using that B is a cone yields that $c \geq 0$ and hence we have that $c = 0$.

By the proper separation there exists either $a \in A$ such that $p \cdot a \neq 0$ or $b \in B$ such that $p \cdot b \neq 0$. Consider the first case and assume that $a_1 \in A$ with $p \cdot a_1 > 0$. If no $a_2 \in A$ exists with $p \cdot a_2 = 0$ we have established that $p \cdot a > 0$ for all $a \in A$ and thus completed the proof. If $a_2 \in A$ exists with $p \cdot a_2 = 0$ consider another point $a_3 = a_2 + \epsilon(a_2 - a_1)$. As A is open relative to its affine hull, $a_3 \in A$ for ϵ small enough. However, note that

$$p \cdot a_3 = (1+\epsilon)(p \cdot a_2) - \epsilon(p \cdot a_1) = -\epsilon(p \cdot a_1) < 0\,.$$

This contradicts that $p \cdot a \geq 0$ for all $a \in A$ and thus implies that no $a_2 \in A$ with $p \cdot a_2 = 0$ can exist. Hence, we have established that $p \cdot a > 0$ for all $a \in A$.

The proof for the case where there exists a $b \in B$ with $p \cdot b < 0$ is analogous. This implies that one of the inequalities is always strict for all $a \in A$ and $b \in B$. This completes the proof. □

Lemma 7. *The set of utility functions consistent with a given (potentially incomplete) preference relation is open relative to its affine hull.*

Proof. Fix two utility functions u, v consistent with $\preceq$ and fix $\epsilon > 0$ such that

$$\epsilon < \frac{\min_{y,y' \colon u(y) > u(y')} u(y) - u(y')}{\max_{y,y' \colon v(y) > v(y')} v(y) - v(y')}\,.$$

We have that $u(x) > u(x')$ and $v(x) > v(x')$ implies that

$$[u(x) + \epsilon(u(x) - v(x))] - [u(x') + \epsilon(u(x') - v(x'))] \geq [u(x) - u(x')] - \epsilon[v(x) - v(x')] > 0\,.$$

Similarly, $u(x) = u(x')$ and $v(x) = v(x')$ implies that

$$u(x) + \epsilon(u(x) - v(x)) = u(x') + \epsilon(u(x') - v(x'))\,.$$

Hence, the utility $u(x) + \epsilon(u(x) - v(x))$ is also consistent with the preference $\preceq$ for every ϵ small enough and the set of utilities consistent with $\preceq$ is open relative to its affine hull (c.f. Aliprantis and Border, 2006, page 277). □

Lemma 8. *The set of strictly concave utility functions is open.*

Proof. A utility function u is strictly concave if for all $x, y, z \in X$

$$\frac{y-z}{y-z}u(x) + \frac{z-x}{y-z}u(y) < u(z)\,.$$

Note that if $v \in B_\epsilon(u)$ (i.e., an ϵ ball around u) we have that

$$\left[\frac{y-z}{y-z}v(x) + \frac{z-x}{y-z}v(y)\right] - v(z) \leq \left[\frac{y-z}{y-z}u(x) + \frac{z-x}{y-z}u(y)\right] - u(z) + 2\epsilon\,.$$

Thus we have that the utility v is strictly concave for

$$\epsilon < \frac{1}{2} \min_{x,y,z \in X} \left| \left[\frac{y-z}{y-z} u(x) + \frac{z-x}{y-z} u(y) \right] - u(z) \right| . \qquad \square$$

We also make use of the following straightforward properties of Minkowski sums of sets of utilities in our proofs. Define $\bar{U}^1 = \sum_\omega U^1_\omega$ and note that $\bar{U}^1$ is convex and open relatively to its affine hull if all $(U^1_\omega)_\omega$ are convex and open relative to their affine hull. Note also that if each of the U^1_ω are cones and constant shift invariant, then so is $\bar{U}^1$: if $\bar{u} \in \bar{U}$, then $(\lambda_0, \cdots, \lambda_0) + \lambda_1 \bar{u} \in \bar{U}$, for any $\lambda_1 > 0$ and any real λ_0. Finally, note that $v = \sum_\omega f_\omega u_\omega$ for some choices of $u_\omega \in U^1_\omega$ if and only if $v = \sum_\omega u_\omega$ for some choices of $u_\omega \in U^1_\omega$.

Lemma 9. *Let U^1 denote all utility functions consistent with the single-peaked time-1 preferences. Fix an alternative $m \in \{\min X, \max X\}$ such that for every $x \neq m$ there exists ω such that $x \preceq^1_\omega m$ does not hold. Then for every $v \in U^1$ there is $u \in U^1$ such that $\mathbb{E}\left[u_\omega(x)\right] = \mathbb{E}\left[v_\omega(x)\right] - \mathbf{1}_{x=m}\Delta$.*

Similarly, if U^1 denotes all concave utility functions consistent with the single-peaked time-1 preferences, then for every $v \in U^1$ there is $u \in U^1$ such that $\mathbb{E}\left[u_\omega(x)\right] = \mathbb{E}\left[v_\omega(x)\right] - \mathbf{1}_{x=m}\Delta$.

Proof. Without loss, assume that $m = \min X$. By assumption, there exists at least one state ω' such that $\max X \preceq^1_{\omega'} m$ does not hold. For this ω', it is then more generally true that $x \preceq^1_{\omega'} m$ does not hold for any $x \neq m$ by the definition of single-peaked preferences. To see this, note that if $x \neq \max X$ and $x \preceq^1_{\omega'} m$, then single-peakness requires that $\max X \preceq^1_{\omega'} x$, and thus that $\max X \preceq^1_{\omega'} \min X$.

Now define

$$u_\omega(x) = v_\omega(x) - \mathbf{1}_{x=m \text{ and } \omega=\omega'} \frac{\Delta}{f_{\omega'}}.$$

By construction, u_ω is identical to v_ω in all states $\omega \neq \omega'$. Moreover, because m is not weakly preferred to any other alternative $x \neq m$ in state ω', it follows that if $v_{\omega'}$ is consistent with $\preceq^1_{\omega'}$ then subtracting a positive constant from $v'_\omega(m)$ still preserves consistency with $\preceq^1_{\omega'}$ as well as single-peakness. Thus, $v \in U^1$ if $u \in U^1$, and $\mathbb{E}\left[u_\omega(x)\right] = \mathbb{E}\left[v_\omega(x)\right] - \mathbf{1}_{x=m}\Delta$ by construction.

Finally, note that if v_ω is concave for all ω, including $\omega = \omega'$, then subtracting a positive constant from $v_{\omega'}(m)$ also preserves concavity. This establishes the last part of the Lemma. $\square$

Lemma 10. *Assume that time-1 preferences are single-peaked and have no indifferences. Let U^1 denote all single-peaked utility functions with no indifferences that are consistent with the time-1 preferences. Fix an alternative $m \notin \{\min X, \max X\}$ such that for every $x \neq m$*

there exists ω such that $x \prec^1_\omega m$ does not hold. Then for every $v \in U^1$ there is $u \in U^1$ such that $\mathbb{E}\left[u_\omega(x)\right] = \mathbb{E}\left[v_\omega(x)\right] - \mathbf{1}_{x=m}\Delta$.

Proof. By definition, there exist states ω' and ω'' such that $\min X \prec^1_{\omega'} m$ and $\max X \prec^1_{\omega''} m$ do not hold. We will show that $u : X \times \Omega \to \mathbb{R}$ defined as below belongs to U^1 if $v \in U^1$:

$$u_\omega(x) = \begin{cases} v_\omega(x) + \Delta & \text{if } \omega \notin \{\omega', \omega''\} \\ v_\omega(x) - \mathbf{1}_{x \geq m} \frac{\Delta}{f_{\omega'}} + \Delta & \text{if } \omega = \omega' \\ v_\omega(x) - \mathbf{1}_{x \leq m} \frac{\Delta}{f_{\omega''}} + \Delta & \text{if } \omega = \omega'' \end{cases} .$$

First, note that if $\min X \prec^1_{\omega'} m$ does not hold, then $y \prec^1_{\omega'} m$ cannot hold for any $y < m$. Otherwise, if $y \prec^1_{\omega'} m$, then $\min X \prec^1_{\omega'} y$ by definition of single-peakness, and thus $\min X \prec^1_{\omega'} m$. Similarly, $m \prec^1_{\omega'} x$ cannot hold for any $x > m$. We now argue that $y \prec^1_{\omega'} x$ cannot hold for $y < m < x$. If it did, then $m \succ^1_{\omega'} y$ by single-peakness. But we have already shown that this cannot hold.

Together, this implies that $y \prec^1_{\omega'} x$ cannot hold for any $x \geq m$ and $y < m$ (and indifference cannot occur by the assumption of the lemma). Thus, subtracting a constant from $v_{\omega'}$ for all alternatives $x \geq m$ leads to another utility function compatible with the preference $\preceq^1_{\omega'}$. A symmetric argument implies that subtracting a constant from $v_{\omega''}$ for all alternatives $x \leq m$ leads to another utility compatible with the preference $\preceq^1_{\omega''}$.

Thus, the utility function u defined above belongs to U^1 if $v \in U^1$ (where we also us the obvious fact that adding Δ to the utility from all elements preserves inclusion in U^1).

Observe that

$$\begin{aligned} \mathbb{E}\left[u_\omega(x)\right] &= \mathbb{E}\left[v_\omega(x)\right] - \mathbb{E}\left[\mathbf{1}_{\omega=\omega'}\right]\mathbf{1}_{x \geq m}\frac{\Delta}{f_{\omega'}} - \mathbb{E}\left[\mathbf{1}_{\omega=\omega''}\right]\mathbf{1}_{x \leq m}\frac{\Delta}{f_{\omega''}} + \Delta \\ &= \mathbb{E}\left[v_\omega(x)\right] - \mathbf{1}_{x \geq m}\Delta - \mathbf{1}_{x \leq m}\Delta + \Delta \\ &= \mathbb{E}\left[v_\omega(x)\right] - \mathbf{1}_{x=m}\Delta , \end{aligned}$$

which completes the proof. □

Lemma 11. *Let $L_1, L_2 \in \Delta(|X|)$ be two lotteries over outcomes.*

(i) L_2 dominates L_1 with respect to U^1 if and only if $\bar{u}(L_2) \geq \bar{u}(L_1)$ for all $\bar{u} \in \bar{U}^1$.

(ii) L_2 strictly dominates L_1 with respect to U^1 if and only if $\bar{u}(L_2) > \bar{u}(L_1)$ for all $\bar{u} \in \bar{U}^1$.

Proof. If L_2 weakly dominates L_1 then $u_\omega(L_2) \geq u_\omega(L_1)$ for all $\omega \in \Omega, u \in U^1$. Thus, $\sum_\omega u_\omega(L_2) \geq \sum_\omega u_\omega(L_1)$ for all u in U^1 and hence weak dominance implies $\bar{u}(L_2) \geq \bar{u}(L_1)$ for all $\bar{u} \in \bar{U}^1$. The argument for strict dominance is analogous.

For the opposite direction observe that $\bar{u}(L_2) \geq \bar{u}(L_1)$ for all $\bar{u} \in \bar{U}^1$ implies that for all $u \in U^1$ we have $\sum_\omega u_\omega(L_2) \geq \sum_\omega u_\omega(L_1)$. Note that if $u_\omega \in U^1_\omega$ then $\alpha_\omega u_\omega \in U^1_\omega$ for all $\alpha_\omega > 0$. Thus, for all $\alpha \in \mathbb{R}^{|\Omega|}_{++}$, we have that

$$\sum_\omega \alpha_\omega u_\omega(L_2) \geq \sum_\omega \alpha_\omega u_\omega(L_1)\,.$$

Choosing $\alpha_\omega = \mathbf{1}_{\omega=\tilde{\omega}} + \epsilon \mathbf{1}_{\omega\neq\tilde{\omega}}$ for $\epsilon > 0$ yields that

$$u_{\tilde{\omega}}(L_2) - u_{\tilde{\omega}}(L_1) \geq \epsilon \left(\sum_{\omega\neq\tilde{\omega}} u_\omega(L_1) - u_\omega(L_2) \right)\,.$$

Taking the limit $\epsilon \to 0$ yields that for each state $\tilde{\omega}$ we have that $u_{\tilde{\omega}}(L_2) \geq u_{\tilde{\omega}}(L_1)$, and thus that L_2 weakly dominates L_1. This establishes part (i) of the Lemma. Furthermore, note that if $\bar{u}(L_2) > \bar{u}(L_1)$, and if $u_\omega(L_2) \geq u_\omega(L_1)$ for all ω, then the inequality must be strict for at least one ω, which establishes part (ii) of the Lemma. □

A.3 Proofs of Theorems and Propositions in the Paper

To simplify exposition we will say that a set is *relatively open* if it is open relative to its affine hull.

Proof of Theorem 1. Let $\bar{U}^1 = \sum_\omega U^1_\omega$ be the set of all utility functions that can be rationalized as sums of single-peaked utility functions representing the time-1 preferences $\preceq^1_\omega$. Let U^0 be the set of all utility functions that are consistent with $\preceq^0$. We will show that if U^0 and $\bar{U}^1$ do not intersect and the data set exhibits no violation of the only weak cycles condition, then there is a simple dominance violation. By Lemma 2 this implies a cyclic dominance violation; hence, if U^0 and $\bar{U}^1$ do not intersect, there must be a cyclic dominance violation.

Because the set of all utility functions consistent with a given (potentially incomplete) preference relation is relatively open by Lemma 7, the set of single-peaked utility functions is relatively open, and the intersection of relatively open sets is relatively open, it follows that U^0 and each U^1_ω are relatively open. Because the Minkowski sum of relatively open sets is relatively open it follows that $\bar{U}^1$ is relatively open.

By Lemma 6 there exists a separating hyperplane $p \in \mathbb{R}^{|X|}$ with $p \neq 0$ and $\sum_{x\in X} p_x = 0$ (throughout, we denote by p_x the entry of p corresponding to alternative $x \in X$) such that for all $u^0 \in U^0, u^1 \in \bar{U}^1$

$$p \cdot u^0 \geq 0 \geq p \cdot u^1 \tag{9}$$

and at least one of the inequalities is strict for all u^0, u^1.

We next argue that this implies the existence of another hyperplane $\tilde{p} \in \mathbb{R}^{|X|}$ that satisfies the same properties and only two non-zero entries, which is equivalent to the existence of a simple dominance violation and thus the statement of the theorem. We prove this by induction, showing that if the statement holds for a set with $|X| - 1$ objects then it holds for a set with $|X|$ objects.

Induction hypothesis We first prove the result for two alternatives $|X| = 2$. In this case, by definition p is either $(+\alpha, -\alpha)$ or $(-\alpha, +\alpha)$ for some $\alpha > 0$, and thus the result holds.

Induction step We next prove the induction step. Assume that the result holds whenever the number of alternatives is not more than $|X| - 1$. We consider the preorder $\trianglelefteq$ defined above.

Consider the set of alternatives the is not strictly better than any other alternative according to the preorder $\trianglelefteq$

$$\mathcal{X} = \{x \in X \colon \text{ there does not exist } y \in X \text{ with } y \triangleleft x\}.$$

Note that $|\mathcal{X}| > 0$ because $\trianglelefteq$ is by construction transitive. Clearly, the elements in $\mathcal{X}$ are either related by indifference $\sim_{\trianglelefteq}$ or unrelated. Pick a maximal subset of this set X^* such that $x \sim_{\trianglelefteq} x'$ for all $x, x' \in X^*$. By definition for all $x' \notin X^*$ and $x \in X^*$ we have that $x' \trianglelefteq x$ can not hold as otherwise there exists an element $y \triangleleft x' \trianglelefteq x$ which contradicts that $x \in X^*$ as $\trianglelefteq$ is transitive.

Fix any $x^* \in X^*$. As $x \preceq^0 x^*$ implies $x \trianglelefteq^* x^*$, and thus $x \trianglelefteq x^*$ we have that $x \preceq^0 x^*$ implies $x \in X^*$. By the same argument $x \preceq^1_* x^*$ implies $x \in X^*$.

Thus, if $x^* \in X^*$ and $x \notin X^*$, either $x^* \prec^0 x$, or x is unrelated to x^* by $\preceq^0$. Thus, if $u^0 \in U^0$ then $u^0 - \lambda \mathbf{1}_{x \in X^*} \in U^0$ for all $\lambda \geq 0$, as we can always make alternatives that are either least preferred or unranked relative to others worse without violating the ranking of the alternatives implied by $\preceq^0$. Thus, for every $u^0 \in U^0$ we have that

$$0 \leq p \cdot (u^0 - \lambda \mathbf{1}_{x \in X^*}) = p \cdot u_0 - \lambda \sum_{x \in X^*} p_x \,.$$

Taking $\lambda \to \infty$ implies that

$$\sum_{x \in X^*} p_x \leq 0. \tag{10}$$

Because time-1 preferences are strict, it follows that if there are no simple dominance violations, then for any $x \in X^*$ and any $x' \in X$, there is a state ω' such that $x \succeq^1_{\omega'} x'$ does not hold. Otherwise it would have to be that $x' \preceq^1_* x$, which implies $x' \prec^1_* x$ because we rule out indifference; thus Lemma 1 implies $x' \lhd^* x$, and Lemma 4 $x' \lhd x$, which contradicts $x \in X^*$. If there is a simple dominance violation then we have reached our desired contradiction. So assume that there is not, so that for any $x \in X^*$ and any $x' \in X$, there is a state ω' such that $x \succeq^1_{\omega'} x'$ does not hold.

Then, Lemmas 9 and 10 imply that if $x^* \in X^*$ and $u^1 \in \bar{U}^1$, then $u^1 - \lambda \mathbf{1}_{x=x^*} \in \bar{U}^1$ for all $\lambda \geq 0$. Thus, for every $u^1 \in \bar{U}^1$ we have that

$$0 \geq p \cdot (u^1 - \lambda \mathbf{1}_{x=x^*}) = p \cdot u_1 - \lambda \mathbf{1}_{x=x^*} p_{x^*} \,.$$

Taking $\lambda \to \infty$ implies that $p_{x^*} \geq 0$ for every $x^* \in X^*$. Together with (10), this implies that $p_{x^*} = 0$.

Thus, there exists a vector $p \in \mathbb{R}^{|X|-|X^*|}$ such that $\sum_{x \in X \setminus X^*} p_x = 0$ with $p \neq 0$ such that (9) is satisfied on the set $X \setminus X^*$. As the preferences are single-peaked on $X \setminus X^*$ and this set contains only $|X| - |X^*|$ alternatives, there exists a vector $p \neq 0$ with $\sum_{x \in X} p_x = 0$ and only two non-zero entries, satisfying (9) on that set of alternatives. As this vector corresponds to a simple dominance violation this completes the proof. □

Proof of Theorem 2. As in the proof of Theorem 1, we will show that if U^0 and $\bar{U}^1$ do not intersect, then there must be a simple dominance violation. Because the set of all utility functions consistent with a given (potentially incomplete) preference relation is relatively open by Lemma 7, the set of strictly concave utility functions is relatively open by Lemma 8, and the intersection of relatively open sets is relatively open, it follows that U^0 and U^1_ω are relatively open. Because the Minkowski sum of relatively open sets is relatively open it follows that $\bar{U}^1$ is relatively open.

By Lemma 6 there exists a separating hyperplane $p \in \mathbb{R}^{|X|}$ with $p \neq 0$ and $\sum_{x \in X} p_x = 0$ such that equation (9) is satisfied for all $u^0 \in U^0, u^1 \in \bar{U}^1$, with at least one of the inequalities strict. As in the proof of Theorem 1, we show that this implies the existence of another hyperplane $\tilde{p} \in \mathbb{R}^{|X|}$ that consists of only two non-zero entries, which is equivalent to the existence of a simple dominance violation and thus the statement of Theorem 2. We prove this by induction, showing that if the statement holds for a set with $|X| - 1$ objects then it holds for a set with $|X|$ objects.

As in the proof of Theorem 1, the statement holds trivially when $|X| = 2$. We next prove the induction step and assume that the result holds whenever the number of alternatives is less than $|X| - 1$. If there is a simple dominance violation then the statement of the theorem holds. If there is not, then Lemma 5 implies that there exists a set $X^* \subseteq \{\min X, \max X\}$ such that if $x \in X \setminus X^*$ and $x^* \in X^*$, then $x \unlhd x^*$ cannot hold, and such that $x \sim_{\unlhd} x'$ for

$x, x' \in X^*$.

As in the proof of Theorem 1, if $u^0 \in U^0$ then $u^0 - \lambda \mathbf{1}_{x \in X^*} \in U^0$ for all $\lambda \geq 0$, as we can always make alternatives that are either least preferred or unranked relative to others worse without violating the ranking of the alternatives implied by $\preceq^0$.

Now for every $u^0 \in U^0$ we have that

$$0 \leq p \cdot (u^0 - \lambda \mathbf{1}_{x \in X^*}) = p \cdot u_0 - \lambda \sum_{x \in X^*} p_x \,.$$

Taking $\lambda \to \infty$ implies that

$$\sum_{x \in X^*} p_x \leq 0. \tag{11}$$

We divide the remainder of the proof into five cases. We rely on the induction step in the last three cases, which imply a vector $p \in \mathbb{R}^{|X|-|X^*|}$ such that $\sum_{x \in X \setminus X^*} p_x = 0$ with $p \neq 0$ such that (9) is satisfied on the set $X \setminus X^*$.

Case (i): Suppose that $X^* = \{\min X, \max X\}$ and either $\min X \prec^1_* \max X$ or $\max X \prec^1_* \min X$. In the first case, this implies that $\max X \preceq^0 \min X$; otherwise, we would have $\min X \trianglelefteq \max X$, which violates the assumption that $X^* = \{\min X, \max X\}$. However, if $\max X \preceq^0 \min X$ and $\min X \prec^1_* \max X$, then there is a simple dominance violation, which establishes the claim. The second case follows analogously.

Case (ii): Suppose that $X^* = \{\min X, \max X\}$, $\min X \sim^1_* \max X$ and the time-0 preference relates $\min X$ and $\max X$. Then as the time zero preference relates $\min X$ and $\max X$ it must also be that $\min X \sim^0 \max X$ if there is not a simple dominance violation.

Thus, we can just identify the alternatives $\min X$ and $\max X$ with each other to arrive at a problem with $|X| - 1$ alternatives. Formally, for any p satisfying (9), note that p' given by $p'_{\min X} = 0$ and $p'_{\max X} = 2p_{\max X}$ also satisfies (9) but belongs to $\mathbb{R}^{|X|-1}$. By the induction hypothesis, there must thus exist a simple dominance violation on $X \setminus \{\min X\}$.

Case (iii): Suppose that $X^* = \{\min X, \max X\}$, $\min X \sim^1_* \max X$ and the time-0 preference does not relate $\min X$ and $\max X$. Note that by construction, if $x \in X \setminus X^*$ and $x^* \in X^*$, then either $x^* \prec^0 x$ or x^* is unrelated to x. Thus, if the time-zero preference does not relate $\min X$ and $\max X$, then for any $x^* \in X^*$ and *any* $x \neq x^*$, $x \preceq^0 x^*$ cannot hold. Because of this, $u^0 \in U^0$ implies that $u^0 - \lambda \mathbf{1}_{x=x^*} \in U^0$ for all $\lambda \geq 0$ and each $x^* \in X^*$, as we can always make alternatives that are either least preferred or unranked relative to others worse without violating the ranking of the alternatives implied by $\preceq^0$. Thus, for every $u^0 \in U^0$ and $x^* \in X^*$ we have that

$$0 \leq p \cdot (u^0 - \lambda \mathbf{1}_{x=x^*}) = p \cdot u^0 - \lambda \mathbf{1}_{x=x^*} p_{x^*}$$

and taking $\lambda \to \infty$ implies that $p_{x^*} \leq 0$ for each $x^* \in X^*$.

Now because $\min X \sim^1_* \max X$ and they do not dominate any other alternative in X, we can identify them with each other in $\bar{U}^1$, and thus Lemma 9 implies that if $u^1 \in \bar{U}^1$ then

$$u^1 - \lambda \mathbf{1}_{x \in X^*} \in \bar{U}^1$$

Thus, for every $u^1 \in \bar{U}^1$ we have that

$$0 \geq p \cdot (u^1 - \lambda \mathbf{1}_{x \in X^*}) = p \cdot u_1 - \lambda \sum_{x \in X^*} p_x \, .$$

Taking $\lambda \to \infty$ implies that $\sum_{x \in X^*} p_x \geq 0$, which implies that $p_x = 0$ for each $x \in X^*$.

Case (iv): Suppose that X^* has only one element x^*. Since $x \preceq^1_* x^*$ cannot hold for any $x \neq x^*$, Lemma 9 implies that if $u^1 \in \bar{U}^1$ then $u^1 - \lambda \mathbf{1}_{x=x^*} \in \bar{U}^1$ for all $\lambda \geq 0$. Thus, for every $u^1 \in \bar{U}^1$ we have that

$$0 \geq p \cdot (u^1 - \lambda \mathbf{1}_{x=x^*}) = p \cdot u_1 - \lambda \mathbf{1}_{x=x^*} p_{x^*} \, .$$

Taking $\lambda \to \infty$ implies that $p_{x^*} \geq 0$ for every $x^* \in X^*$. Together with (11), this implies that $p_{x^*} = 0$.

Case (v): Suppose that $X^* = \{\min X, \max X\}$, and that neither $\min X \preceq^1_* \max X$ nor $\max X \preceq^1_* \min X$. Then application of Lemma 9 as in Case (iv) implies that $p_{x^*} \geq 0$ for each $x^* \in X^*$. Together with (11), this again implies that $p_{x^*} = 0$ for all $x^* \in X^*$.

Completing the proof in Cases (iii), (iv), and (v): In all of these cases, there exists a vector $p \in \mathbb{R}^{|X|-|X^*|}$ such that $\sum_{x \in X \setminus X^*} p_x = 0$ with $p \neq 0$ such that (9) is satisfied on the set $X \setminus X^*$. As the preferences are single-peaked on $X \setminus X^*$ and this set contains only $|X| - |X^*|$ alternatives, there exists a vector p with $|p_x| = 1$ or $p_x = 0$ $\forall x \in X$, satisfying (9) on that set of alternatives. As this vector corresponds to a simple dominance violation this completes the proof. □

Proof of Proposition 2. First note that we can restrict to the case where preferences are complete. This is because if the preferences are incomplete and satisfy one of the two conditions, then there exist completions of the preferences that satisfy one of the conditions. And the incomplete preferences can of course be rationalized by TC-EU if their completions can be.

We shall prove the proposition under the assumption that condition 1 holds. The proof for condition 2 uses identical arguments that start with the most preferred time-0 alternatives rather than the least preferred time-0 alternatives.

As in the proof of Theorem 1, Lemma 7 implies that U^0, U^1_ω and $\bar{U}^1$ can be shown to be relatively open. We shall show that under the conditions of the Proposition, if $p \in \mathbb{R}^{|X|}$ satisfies (i) $\sum_{x \in X} p_x = 0$ and (ii) $p \cdot u^0 \geq 0 \geq p \cdot u^1$ for all $u^0 \in U^0, u^1 \in \bar{U}^1$, then $p \equiv 0$.This then implies that $U^0 \cap \bar{U}^1 \neq \emptyset$, which is the statement of the Proposition.

We shall prove this by induction, showing that if the statement holds for a set with no more than $|X| - 1$ objects, then it holds for a set with $|X|$ objects.

Let X^* denote the set of least preferred elements of the time-0 preference. Observe that if $u^0 \in U^0$ then $u^0 - \lambda \mathbf{1}_{x \in X^*} \in U^0$ for all $\lambda \geq 0$: we can always make least preferred alternatives worse without changing the ranking of the alternatives. Thus, for every $u^0 \in U^0$ we have that

$$0 \leq p \cdot (u^0 - \lambda \mathbf{1}_{x \in X^*}) = p \cdot u_0 - \lambda \sum_{x \in X^*} p_x \,. \tag{12}$$

Taking $\lambda \to \infty$ implies that $\sum_{x \in X^*} p_x \leq 0$.

Now by the assumption of the proposition, for each $x^* \in X^*$ there exists a state ω^* in which x^* is the least preferred element. We have that if $u^1 \in U^1$ then $u^1 - \frac{\lambda}{f_{\omega^*}} \mathbf{1}_{x = x^* \text{ and } \omega = \omega^*} \in U^1$ for all $\lambda > 0$. This implies that if $u \in \bar{U}^1$ then $u - \lambda \mathbf{1}_{x = x^*} \in \bar{U}^1$ for all $\lambda \geq 0$. Thus, for every $u \in \bar{U}^1$, we have that

$$0 \geq p \cdot (u - \lambda \mathbf{1}_{x = x^*}) = p \cdot u - \lambda p_{x^*} \,.$$

Taking $\lambda \to \infty$ implies that $p_{x^*} \geq 0$. Together with the condition that $\sum_{x \in X^*} p_x \leq 0$, this implies that $p_{x^*} = 0$ for all $x^* \in X^*$.

Now when $|X| = 2$, the above implies that at least one of the elements of p must equal 0, which implies that all elements of p must equal zero by the condition that $\sum_{x \in X} p_x = 0$.

When $|X| > 2$, the above implies that at least one of the elements of p must equal 0. Say that this element corresponds to an element x^*, and note that the condition of the proposition still applies to preferences on the set $X \setminus \{x^*\}$. Thus, if the result holds for sets of size $|X| - 1 \geq 2$, it must hold for sets of size $|X|$. □

Proof of Proposition 3. Assume that $U^0 \cap \bar{U}^1 = \emptyset$ and the data set can thus not be rationalized by TC-EU. By Lemma 6, there exists $p \in \mathbb{R}^{|X|}$ with $p \neq 0$ and $\sum_{x \in X} p_x = 0$ such that $p \cdot u^0 \geq 0 \geq p \cdot u$ for $u^0 \in U^0$ and all $u \in \bar{U}^1$, such that at least one inequality is strict for all $u^0 \in U^0$ and $u \in \bar{U}^1$. Define $L_1, L_2 \in \mathbb{R}^{|X|}$

$$L_1(x) = \frac{\max\{p_x, 0\}}{\sum_{\tilde{x} \in X} \max\{p_{\tilde{x}}, 0\}}$$

$$\text{and } L_2(x) = \frac{\max\{-p_x, 0\}}{\sum_{\tilde{x} \in X} \max\{-p_{\tilde{x}}, 0\}} \,.$$

Note that by definition, the entries of L_1, L_2 are non-negative and sum up to one, which implies that L_1, L_2 are well-defined lotteries over the alternatives X. Furthermore, we have that

$$0 = \sum_{\tilde{x} \in X} p_{\tilde{x}} = \sum_{\tilde{x} \in X} \max\{p_{\tilde{x}}, 0\} - \sum_{\tilde{x} \in X} \max\{-p_{\tilde{x}}, 0\} \,.$$

This implies that

$$L_1(x) - L_2(x) = \frac{\max\{p_x, 0\}}{\sum_{\tilde{x} \in X} \max\{p_{\tilde{x}}, 0\}} - \frac{\max\{-p_x, 0\}}{\sum_{\tilde{x} \in X} \max\{-p_{\tilde{x}}, 0\}} = p_x \frac{1}{\sum_{\tilde{x} \in X} \max\{p_{\tilde{x}}, 0\}} \,.$$

Thus, $p \cdot u^0 \geq 0$ implies that $u^0(L_1) \geq u^0(L_2)$, and $p \cdot u^0 > 0$ implies that $u^0(L_1) > u^0(L_2)$. Similarly, $p \cdot u \leq 0$ for all $u \in \bar{U}$ implies that $\bar{u}(L_2) \geq \bar{u}(L_1)$ for all $u \in \bar{U}$ and $p \cdot u < 0$ for all $u \in \bar{U}$ implies that $\bar{u}(L_2) > \bar{u}(L_1)$ for all $u \in \bar{U}$. Thus, by Lemma 11, for all $u^0 \in U^0$, either (i) $u^0(L_1) \geq u^0(L_2)$ and L_2 strictly dominates L_1 with respect to U^1 or (ii) $u^0(L_1) > u^0(L_2)$ and L_2 weakly dominates L_1 with respect to U^1. Hence, according to Definition 8, the data set exhibits a stochastic dominance violation if $U^0 \cap \bar{U}^1 = \emptyset$; i.e., if the data set (U^0, U^1, f) is inconsistent with TC-EU then it exhibits a stochastic dominance violation.

The opposite direction is immediate: Suppose that the data set is consistent with TC-EU and for $u^0 \in U^0, u^1 \in U^1$

$$u^0(x) = \sum_{\omega \in \Omega} f_\omega u^1_\omega(x) \,.$$

Then, $u^1_\omega(L_2) \geq u^1_\omega(L_1)$ for all ω implies that $u^0(L_2) \geq u^0(L_1)$, and thus the data set cannot exhibit dominance violations. □

Proof of Proposition 4. As preferences are responsive, there exists $z_{\omega,y}$ and $z^\circ_{\omega,y}$ known to the analyst such that

$$(y, z_{\omega,y}) \sim^1_\omega (y^\circ, z^\circ_{\omega,y}) \,.$$

As preferences are additive, this means that

$$h^1_\omega(y) + g(z_{\omega,y}) = h^1_\omega(y^\circ) + g(z^\circ_{\omega,y}) \,.$$

This implies that $h^1_\omega(y) - h^1_\omega(y^\circ) = g(z^\circ_{\omega,y}) - g(z_{\omega,y})$ for each ω, and thus that

$$\sum_\omega \left(h^1_\omega(y) - h^1_\omega(y^\circ)\right) = \sum_\omega \left(g(z^\circ_{\omega,y}) - g(z_{\omega,y})\right) .$$

By the same argument, there exist $z_{0,y}$ and $z^\circ_{0,y}$ such that

$$h^0(y) - h^0(y^\circ) = g(z^\circ_{0,y}) - g(z_{0,y}) \,.$$

Dividing the terms yields that

$$\frac{h^0(y)-h^0(y^\circ)}{\sum_\omega f_\omega(h^1_\omega(y)-h^1_\omega(y^\circ))}=\frac{g(z^\circ_{0,y})-g(z_{0,y})}{\sum_\omega f_\omega(z^\circ_{\omega,y}-z_{\omega,y})}\,.$$

As all terms on the right-hand side are observable to the analyst, it follows that the left-hand side is point-identified. □

Proof of Proposition 5. We first argue necessity of the conditions. Revealed additive separability is implied by additive separability, by simply setting $h^0\equiv\frac{\theta^1_\omega}{\theta^2_\omega}h$ and $h^1_\omega\equiv\frac{\theta^1_\omega}{\beta\theta^2_\omega}h$. To see that the condition on β is necessary, observe that for any representation (h,g) of preferences consistent with (5)

$$\begin{aligned}\beta&=\frac{(\sum_\omega f_\omega\theta^1_\omega)(h(y)-h(y^\circ))}{\sum_\omega f_\omega\theta^1_\omega\frac{1}{\beta}(h(y)-h(y^\circ))}=\frac{(\sum_\omega f_\omega\theta^1_\omega)/(\sum_\omega f_\omega\theta^2_\omega)(h(y)-h(y^\circ))}{(\sum_\omega f_\omega\theta^2_\omega)^{-1}(\sum_\omega f_\omega\theta^1_\omega\frac{1}{\beta}(h(y)-h(y^\circ)))}\\&=\frac{g(z^\circ)-g(z_{0,y})}{(\sum_\omega f_\omega\theta^2_\omega)^{-1}\sum_\omega f_\omega\theta^2_\omega(g(z_{\omega,y})-g(z^\circ_{\omega,y}))}=\frac{g(z^\circ)-g(z_{0,y})}{\sum_\omega\alpha_\omega(g(z_{\omega,y})-g(z^\circ_{\omega,y}))}\\&\in\left(\min_\omega\frac{g(z^\circ)-g(z_{0,y})}{g(z_{\omega,y})-g(z^\circ_{\omega,y})},\max_\omega\frac{g(z^\circ)-g(z_{0,y})}{g(z_{\omega,y})-g(z^\circ_{\omega,y})}\right)=(\min\hat\beta_\omega,\max\hat\beta_\omega)\,,\end{aligned}$$

where we define $\alpha_\omega=f_\omega\theta^2_\omega/(\sum_{\omega'}f_{\omega'}\theta^2_{\omega'})$.

We next prove sufficiency. By the assumptions of the proposition there exist weights $\alpha\in\Delta^{|\Omega|}$ such that

$$\beta=\frac{1}{\sum_{\omega\in\Omega}\alpha_\omega\hat\beta_\omega^{-1}}\,.$$

We define $\theta^1_\omega=\frac{\alpha_\omega}{f_\omega}\hat\beta^{-1}_\omega$, $\theta^2_\omega=\frac{\alpha_\omega}{f_\omega}$. We note that as $\frac{h^1_\omega(y)}{h^0(y)}=\frac{h^1_\omega(y^\circ)}{h^0(y^\circ)}$ by the assumptions of the proposition,

$$\hat\beta_\omega=\frac{g(z^\circ_{0,y})-g(z_{0,y})}{g(z^\circ_{\omega,y})-g(z_{\omega,y})}=\frac{h^0(y)-h^0(y^\circ)}{h^1_\omega(y)-h^1_\omega(y^\circ)}=\frac{h^0(y)-h^0(y^\circ)}{h^0(y)\frac{h^1_\omega(y)}{h^0(y)}-h^0(y^\circ)\frac{h^1_\omega(y^\circ)}{h^0(y^\circ)}}=\frac{h^0(y^\circ)}{h^1_\omega(y^\circ)}\,.$$

We thus have that

$$h^1_\omega(y)+g(z)=h^0(y)\frac{h^1_\omega(y^\circ)}{h^0(y^\circ)}+g(z)=h^0(y)\hat\beta^{-1}_\omega+g(z)=h^0(y)\frac{\theta^1_\omega}{\theta^2_\omega}+g(z)\,.$$

Consequently, the utility

$$\theta^1_\omega h^0(y)+\theta^2_\omega g(z)$$

represents the preference $\preceq^1_\omega$. Finally, we have that

$$\sum_\omega f_\omega \theta^1_\omega h^0(y) = h^0(y) \sum_\omega \alpha_\omega \hat{\beta}^{-1}_\omega = \frac{1}{\beta} h^0(y)$$
$$\sum_\omega f_\omega \theta^2_\omega g(z) = g(z) \sum_\omega \alpha_\omega = g(z)\,,$$

which completes the proof. □

B Heterogeneity

Theorems 1 and 2 and Propositions 1 and 3 continue to hold verbatim when it's possible to observe the joint distribution of time-0 and time-1 preferences. Because the only pattern that rejects time consistency with homogeneous preferences is a (stochastic) dominance violation, and because such a violation cannot be explained by heterogeneous preferences, a data set is consistent with TC-EU under homogeneous preferences if and only if it is consistent with TC-EU under heterogeneous preferences. To see this, fix a time-0 preference profile, allowing for individuals to be heterogeneous conditional on this time-0 preference. If the analyst observes a dominance violation (simple for Theorems 1 and 2, or stochastic for Propositions 1 and 3), then the analyst must conclude that all agents with that time-0 preference are time-inconsistent. If the analyst does not observe a dominance violation, then Theorems 1 and 2 and Propositions 1 and 3 imply that the data can be rationalized with a homogeneous time-consistent EU agent.

Our identification results in Propositions 4 and 5 are also straightforward to generalize to give a measure of *average* time inconsistency when there is some unobserved heterogeneity. To give a concrete example of applying Proposition 4 to a heterogeneous population, suppose that there there is a finite number of agent types making the same time-0 choice, with agents of each type having the same preferences and receiving independent draws from the same distribution of taste shocks. Suppose the analyst observes only a single realization of the time-1 preference for each agent. Then, the logic of Example 3 is still identical (under the maintained assumptions of that example): In expectation, the average time-1 WTP for the healthy over the unhealthy snack must equal the average time-0 WTP for the unhealthy over healthy snack. In Supplementary Appendix E.2 we provide additional examples applying Propositions 3 and 4 to heterogeneous populations, with applications to the data sets from Augenblick, Niederle, and Sprenger (2015) and Augenblick and Rabin (2019).

C Mathematical Details for Section 1.1

Denote by $\mathrm{var}_0(\log \mathbb{E}_1[\theta_2]/\mathbb{E}_1[\theta_1])$ the variance of $\log \mathbb{E}_1[\theta_2]/\theta_1$ conditional on time 0 information.

Lemma 12. *Suppose that the agent choose the allocation $x^0 = 1/2$ at time 0 and the (random) allocation x^1 at time 1. Suppose that $\log(x^1/(1-x^1))$ is normally distributed with mean m and variance v. Then*

$$\begin{aligned}\mathbb{E}_0[\theta_2] &= \mathbb{E}_0[\theta_1]\\ (\gamma-1)m &= \mathbb{E}_0\left[\log\frac{\mathbb{E}_1[\theta_2]}{\theta_1}\right] + \log(\beta)\\ (\gamma-1)^2 v &= var_0\left(\log\frac{\mathbb{E}_1[\theta_2]}{\theta_1}\right)\,.\end{aligned} \tag{13}$$

Proof. Taking first order conditions yields that the optimal effort at time 0 is satisfies

$$(\gamma-1)\log\frac{x^0}{1-x^0} = \log\frac{\mathbb{E}_0[\theta_2]}{\mathbb{E}_0[\theta_1]}\,.$$

The optimal effort at time 1 satisfies

$$(\gamma-1)\log\frac{x^1}{1-x^1} = \log\frac{\mathbb{E}_1[\theta_2]}{\theta_1} + \log(\beta)\,.$$

As $x^0 = 1 - x^0$, we have that $\mathbb{E}_0[\theta_2] = \mathbb{E}_0[\theta_1]$. Taking expectations yields that

$$(\gamma-1)m = (\gamma-1)\mathbb{E}_0\left[\log\frac{x^1}{1-x^1}\right] = \mathbb{E}_0\left[\log\frac{\mathbb{E}_1[\theta_2]}{\theta_1}\right] + \log(\beta)\,.$$

Furthermore, we we have that

$$(\gamma-1)^2 v = \mathrm{var}_0\left(\log\frac{\mathbb{E}_1[\theta_2]}{\theta_1}\right) \qquad \square\,.$$

With this Lemma in hand, we now provide calculations for how β is identified under the different structural assumptions listed in Table 1.

Rows 1 and 2: Independent θ_2, θ_1, revealed at $t = 1$ Suppose that θ_1, θ_2 are independent and $\log(\theta_1) \sim \mathcal{N}(\mu_1, \sigma_1^2)$, $\log(\theta_2) \sim \mathcal{N}(\mu_2, \sigma_2^2)$. Then (13) becomes

$$\begin{aligned}
\exp(\mu_1 + \sigma_1^2/2) = \mathbb{E}_0[\theta_1] = \mathbb{E}_0[\theta_2] &= \exp(\mu_2 + \sigma_2^2/2) \\
(\gamma - 1)m &= \mathbb{E}_0[\log(\theta_2/\theta_1)] - \log(\beta) = \mu_2 - \mu_1 + \log(\beta) \\
(\gamma - 1)^2 v &= \sigma_1^2 + \sigma_2^2 = \sigma_1^2(1 + \sigma_2^2/\sigma_1^2) \,.
\end{aligned}$$

The first equation and third equation imply that

$$\mu_2 - \mu_1 = 0.5(\sigma_1^2 - \sigma_2^2) = 0.5\sigma_1^2(1 - \sigma_2^2/\sigma_1^2) = 0.5(\gamma - 1)^2 v \frac{1 - \sigma_2^2/\sigma_1^2}{1 + \sigma_2^2/\sigma_1^2} \,.$$

Plugging into the second equation yields

$$\log(\beta) = (\gamma - 1)m - 0.5(\gamma - 1)^2 v \frac{1 - \sigma_2^2/\sigma_1^2}{1 + \sigma_2^2/\sigma_1^2} \,.$$

In the case of i.i.d. taste shocks, this reduces to

$$\log(\beta) = (\gamma - 1)m \,.$$

The analysts estimate thus depends σ_2^2/σ_1^2 which captures what the analyst assumes about how well informed the agent is at time 0 about their time 2 taste-shock relative to their time 1 taste-shock. This ration capture *both* he variance of the agents taste shocks as well as the quality of the information about the agents taste shocks. Setting $\sigma_2^2/\sigma_1^2 = \infty$ captures the case where the analyst assumes that the agent is uninformed about the time 1 preference at time 0. Setting $\sigma_2^2/\sigma_1^2 = 0$ captures the case where the analyst assumes that the agent is uninformed about the time 2 preference at time 0 (and recovers the result we obtained before. If the agent knows equally much about their time 1 and time 2 preferences at time 0, i.e. $\sigma_2^2/\sigma_1^2 = 0$ we get that

$$\log(\beta) = (\gamma - 1)m \,.$$

Rows 3 and 4: θ_1 learned at $t = 0$ and θ_2 learned in $t = 1$ If θ_1 is learned at time 0 but θ_2 is not, then θ_1 must always take on the value $\theta_1 = \mathbb{E}_1\theta_2$. Then (13) becomes

$$\begin{aligned}
\theta_1 &= \mathbb{E}_0[\theta_2] = \exp(\mu_2 + \sigma_2^2/2) \\
(\gamma - 1)m &= \mathbb{E}_0[\log(\theta_2/\theta_1)] - \log(\beta) = \mu_2 - \log(\theta_1) + \log(\beta) \\
(\gamma - 1)^2 v &= \sigma_2^2 \,.
\end{aligned}$$

The first equation and third equation imply that

$$\mu_2 - \log(\theta_1) = -0.5\sigma_2^2 = -.5(\gamma-1)^2 v\,.$$

Plugging this into the second equation yields that

$$\log(\beta) = .5(\gamma-1)^2 v + (\gamma-1)m$$

Rows 5 and 6: θ_2 learned in $t=0$ and θ_1 learned in $t=1$ If θ_2 is learned at time 0 but θ_1 is not, then θ_2 must always take on the value $\theta_2 = \mathbb{E}_1\theta_1$. Then (13) becomes

$$\begin{aligned}\theta_2 &= \mathbb{E}_0[\theta_1] = \exp(\mu_1 + \sigma_1^2/2)\\(\gamma-1)m &= \mathbb{E}_0[\log(\theta_2/\theta_1)] - \log(\beta) = \log(\theta_2) - \mu_1 + \log(\beta)\\(\gamma-1)^2 v &= \sigma_1^2\,.\end{aligned}$$

The first equation and third equation imply that

$$\mu_1 - \log(\theta_2) = -0.5\sigma_1^2 = -.5(\gamma-1)^2 v\,.$$

Plugging this into the second equation yields that

$$\log(\beta) = -.5(\gamma-1)^2 v + (\gamma-1)m$$

Rows 7 and 8: θ_1 learned in $t=1$ and θ_2 learned after $t=1$ Assume that $\log(\theta_1)$ is Normally distributed with mean μ and variance σ^2, and assume, without loss of generality, that $\mathbb{E}[\theta_1] = 1$. By assumption we have that $\mathbb{E}_0[\theta_2] = \mathbb{E}_1[\theta_2]$ and hence (13) becomes

$$\begin{aligned}1 &= \mathbb{E}_0[\theta_1] = \exp(\mu_1 + \sigma^1/2)\\(\gamma-1)m &= \log(\mathbb{E}_0[\theta_2]) - \mathbb{E}_0\left[\log(\theta_1)\right] + \log(\beta)\\(\gamma-1)^2 v &= \mathrm{var}_0(\log(\theta_1)) = \sigma_1^2\,.\end{aligned}$$

The first and third equation together imply that $\mu = -\sigma^2/2 = -0.5(\gamma-1)^2 v$ and hence

$$\log(\beta) = (\gamma-1)m - [\mu + \sigma^2/2] + \mu = (\gamma-1)m - \sigma^2/2 = (\gamma-1)m - 0.5(\gamma-1)^2 v\,.$$

Rows 9 and 10: Multiplicative random walk with θ_1, θ_2 both learned in $t=1$ Formally, $\theta_2 = \theta_1 \cdot \epsilon_1$, where $\log(\epsilon_1) \sim N(\mu, \sigma)$ is log-Normally distributed and θ_1, θ_2 are learned by the agent only at the beginning of time 1. Then (13) becomes

$$\begin{aligned}1 &= \mathbb{E}_0[\epsilon_1] = \exp(\mu + \sigma^2/2)\\(\gamma - 1)m &= \mathbb{E}_0[\log(\epsilon_1)] - \log(\beta) = \mu - \log(\beta)\\(\gamma - 1)^2 v &= \sigma^2 \,.\end{aligned}$$

The first and third equation together imply that $\mu = -\sigma^2/2 = -0.5(\gamma-1)^2 v$ and hence

$$\log(\beta) = (\gamma - 1)m + 0.5(\gamma - 1)^2 v \,.$$

Rows 11-14: Mulitplicative AR(1), with θ_1 learned in $t = 1$ and θ_2 learned after $t = 1$ Suppose that $\log(\theta_2) = \alpha \log(\theta_1) + \log(\varepsilon)$, where $\log(\theta_1) \sim N(\mu_1, \sigma_1^2)$ and $\log(\varepsilon) \sim N(\mu_\varepsilon, \sigma_\varepsilon^2)$. That is, $\log(\theta_1)$ and $\log(\theta_2)$ form an AR(1) process. The agent learns θ_1 at time 1 and ε at time 2. The scalar α can be regarded as a parametrization of how much is learned at about time-1 versus time-2 shocks at time 1. For example, $\alpha = 0$ means that nothing is learned about time-2 shocks at time 1, while $\alpha \to \infty$ captures the case where at time 1 the agent mostly learns about time-2 shocks.

Under this process, we have that $\log\theta_2|\theta_1 \sim N(\alpha\log\theta_1 + \mu_\varepsilon, \sigma_\varepsilon^2)$ and $\mathbb{E}_1[\log\theta_2|\theta_1] = \alpha\theta_1 + \mu_\varepsilon + \sigma_\varepsilon^2/2$. Thus, (13) becomes

$$\begin{aligned}\exp(\mu_1 + \sigma_1^2/2) = \mathbb{E}_0[\theta_1] = \mathbb{E}_0[\theta_2] &= \exp(\alpha\mu_1 + \alpha^2\sigma_1^2/2 + \mu_\varepsilon + \sigma_\varepsilon^2/2)\\(\gamma - 1)m &= \mathbb{E}_0[\alpha\log\theta_1 + \mu_\varepsilon + \sigma_\varepsilon^2/2 - \log\theta_1] + \log\beta = (\alpha - 1)\mu_1 + \mu_\varepsilon + \sigma_\varepsilon^2/2 + \log\beta\\(\gamma - 1)^2 v &= (\alpha - 1)^2\sigma_1^2 \,.\end{aligned}$$

The first equality implies that

$$(\alpha^2 - 1)\mu_1 + \mu_\varepsilon + \sigma_\varepsilon^2/2 = (1 - \alpha^2)\sigma_1^2/2.$$

The third equality implies that

$$\sigma_1^2 = (\gamma - 1)^2 v/(\alpha - 1)^2$$

and thus that

$$(1 - \alpha^2)\sigma_1^2/2 = 0.5(\gamma - 1)^2 v \frac{1 - \alpha^2}{(\alpha - 1)^2}$$

Plugging this into the expression for $(\gamma - 1)m$ yields

$$\begin{aligned}\log\beta &= (\gamma-1)m - 0.5(\gamma-1)^2 v\frac{1-\alpha^2}{(\alpha-1)^2}\\ &= (\gamma-1)m + 0.5(\gamma-1)^2 v\frac{1+\alpha}{\alpha-1}\end{aligned}$$

Now when $\alpha > 1$, β becomes arbitrarily large as α converges to 1 from the right. When $\alpha < 1$, β becomes arbitrarily small as α converges to 1 from the left.

D Relation to Other Technical Results

Social Choice. The ordinal efficiency welfare theorem (McLennan, 2002, Carroll, 2010) states that for any lottery that is Pareto efficient given a vector of ordinal preferences, there exist utility functions consistent with the ordinal preferences such that this lottery maximizes the sum of utilities. This result is mathematically equivalent to the special case of Proposition 3 where the analyst only observes the most preferred time-0 alternative.[35] The sharper and more interesting characterizations that we provide for single-peaked and concave preferences in Theorems 1 and 2 do not, to our knowledge, relate to any known results in the social choice literature—although they of course have implications for that literature. For example, they imply that for complete single-peaked preferences, it is not necessary to consider lotteries: an alternative is a maximand of some social welfare function as long as it is not Pareto dominated by any other alternative. Example 2 shows that this stronger conclusion fails for social choice problems without the single-peaked property.[36]

Dynamically Consistent Preferences over Acts. A literature in decision theory has studied the question of when preferences over acts are consistent with EU (e.g. Chapter 8.2 in Strzalecki, 2021). In this literature the analyst observes any decision-relevant state as well as preferences over acts. This contrasts with our setting where states are unobserved and only preferences over actions—i.e., constant acts—are observed by the analyst. For example, in the context of food choices, the assumption made in this literature would correspond to the analyst observing how hungry the agent is, what type of meal he had last, and whether or not it is a warm day, as well as preferences over strategies that specify at time 0 what the agent will eat in *each* of these observable states. Notably, such a data set—where states and preferences over strategies are observable—is much richer than the data sets collected in the preference reversal literature, which are our objects of study. This decision literature refers to the analogue of our no simple dominance violations condition on acts as "dynamic consistency" (Axiom 8.6 in Strzalecki, 2021). Imposed over acts this condition is much more restrictive and (together with consequentialism) implies that there is a subjective EU representation of the preference (Theorem 8.10 and Theorem 8.24 in Strzalecki, 2021 and Ghirardato, 2002). This is in contrast to our setting where we show that "no simple dominance violation" is, without the restriction to single-dimensional choice sets and single-peaked preferences, not sufficient to ensure the existence of an EU representation.

[35] Specifically, this is the case for the more general version stated by Carroll (2010). The original version stated by McLennan (2002) imposes a more special structure.

[36] It is perhaps also worth clarifying that to our knowledge and understanding, our results do not have a mathematical connection to the literature on aggregation of time preferences (e.g., Jackson and Yariv, 2015, Millner, 2020).

Random Utility Models. In the literature on random utility models, the analyst observes the distribution of *optimal choices* from *all* choice sets at a single point in time (comparable to our time-1 data $(\preceq^1, f)$). The question is what can be learned about the agent's mean utility for the different alternatives. By contrast, we assume that the analyst observes the distribution of *preferences* over a choice set. This data can not be reconstructed from the optimal choices (Fishburn, 1998). The data sets we study, which are based on the types of experimental data collected in practice, are therefore richer. Allowing the analyst to observe the distribution over a complete ranking of all alternatives is equivalent to allowing the analyst to observe a joint distribution of preferred alternatives from all subsets in the random utility literature.[37] While our time-1 data is always consistent with EU, one needs additional conditions to ensure consistency with EU when only the marginal distribution of choices from subsets, but not the joined distribution is observed. A focus of the random utility literature has been to identify such conditions (Block, Marschak et al., 1959, McFadden and Richter, 1990, Clark, 1996, Gul and Pesendorfer, 2006).

A second difference is that the random utility literature typically makes the "positivity" assumption that each alternative is the most preferred one with positive probability. This is a strong assumption when combined with the assumption of single-peaked preferences, which are the main focus of our paper. Positivity and single-peakness together imply that the agent ranks the alternatives both in increasing and decreasing order with positive probability. Furthermore, a corollary of our Proposition 1 implies that this assumption is highly consequential, as it implies that the average utility cannot be identified without imposing additional structure on the preference shocks.[38] This generalizes the insight from Alós-Ferrer, Fehr, and Netzer (2021) who highlight a related identification issue in a setting where the analyst has less information and only observes the marginal distribution of preferences over binary choice sets. They propose to resolve it by inferring cardinal information from response times, which is similar to the additional choice dimension we propose in Section 4.2.[39]

[37] Formally, when observing the distribution f over strict rankings, one can infer the probability of choosing x from the set $M \subseteq X$ as $\sum_\omega f_\omega \mathbf{1}_{x \succeq^1_\omega y \forall y \neq x}$.

[38] There is also a thematic, but not mathematical, connection to identifying time preferences in dynamic discrete choice models. See, e.g., Magnac and Thesmar (2002), Abbring and Daljord (2020), Levy and Schiraldi (2021), Mahajan, Michel, and Tarozzi (2020).

[39] The literature on dynamic random utility (e.g. Fudenberg and Strzalecki, 2015, Frick, Iijima, and Strzalecki, 2019) studies questions that are further removed from ours. We are interested in settings where the agent makes the same choice repeatedly, while that literature studies when a sequence of dynamic choices can be rationalized if the agent's utility function and choice set can change over time. An exception is the case of Bayesian evolving beliefs discussed in Section 6.2 of Frick, Iijima, and Strzalecki (2019). Their Proposition 6 concerns a special case of their model which is similar to a special case of our Proposition 3 where preferences over some set of lotteries are observable. Similarly, our model is different from those analyzed in the literature on preferences for flexibility due to taste uncertainty, as in Ahn and Sarver (2013), where the agent chooses a menu at time-0 and then chooses from that menu at time-1.

E Additional Results on Implementation

In this appendix, we illustrate the additional tests that researchers can conduct to gauge the importance of random taste shocks in their data. We also illustrate how researchers can apply Propositions 3 and 4 to data sets featuring these taste shocks to point or set identify the quasi-hyperbolic discounting model. We provide these illustrations by reanalyzing the data from the important work of Augenblick, Niederle, and Sprenger (2015) and Augenblick and Rabin (2019), which generate ideal data sets for these types of analyses.

E.1 Gauging Uncertainty

In the Augenblick, Niederle, and Sprenger (2015) experiment, participants are asked to allocate effort—in the form of units of unpleasant tasks such as transcriptions—between weeks 2 and 3 of the experiment. Participants do this in week 1 (time 0) and week 2 (time 1). The week 1 preference is implemented with probability 10 percent, and the week 2 preference is implemented with probability 90 percent.

Individuals also complete this task a second time, in weeks 4-6. Consistent with Augenblick, Niederle, and Sprenger, we call weeks 1-3 block 1, denoted $j = 1$, and we call weeks 4-6 block 2, denoted $j = 2$.

Throughout this appendix, we follow Augenblick, Niederle, and Sprenger in limiting analysis to the "80 subjects who completed all aspects of the experiment with positive variation in their responses in each week." We use $\hat{\beta}_i^1$ to denote an individual-level estimate of person i's present focus from block 1, and we use $\hat{\beta}_i^2$ to denote an individual-level estimate of person i's present focus from block 2.

We find a statistically insignificant and directionally negative correlation between $\hat{\beta}_i^1$ and $\hat{\beta}_i^2$ of -0.08, with a 95% confidence interval of $[-0.29, 0.15]$. This rules out correlations larger than 0.15 with high confidence, and suggests that taste shocks and other forms of updating explain most of the variation in individuals' time-1 revisions of their time-0 choices. The upper bound of 0.15 on the correlation suggests that stable time preferences cannot explain more than about $0.15^2 \approx 2$ percent of the variance in time-1 revisions of time-0 preferences. Figure B1 presents the relationship between the two estimates of present focus in the real-effort task.

Another variable analyzed by Augenblick, Niederle, and Sprenger is an indicator for present focus: $PB_i^j = \hat{\beta}_i^j < 1$. We find that the correlation between PB_i^1 and PB_i^2 is 0.04, with a 95% confidence interval of $[-0.18, 0.26]$.

We also note that this low correlation is unlikely to be due to mere "noise" in the data. As Figure B2 shows, subjects' revised (time 1) allocations of effort are very similar to subjects'

Figure B1: Within-person stability of the individual-level estimates of present focus over effort

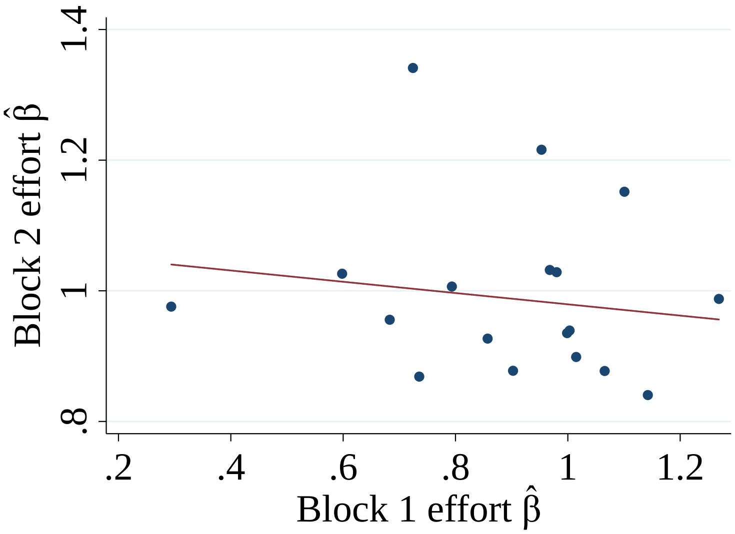

Notes: This figure presents a binned scatter plot of the relationship between $\hat{\beta}_i^2$, the block 2 individual-level estimates of present focus over effort (y-axis) and $\hat{\beta}_i^1$, the block 1 individual-level estimates of present focus over effort (x-axis).

initial (time 0) allocations of effort. The correlations between revised and initial effort share allocations are 0.76 and 0.83 in blocks 1 and 2, respectively. A regression of revised on initial allocations produces coefficients of 0.83 (cluster-robust SE 0.05) and 0.82 (cluster-robust SE 0.04) in blocks 1 and 2, respectively, and these decrease only modestly to 0.57 (cluster-robust SE 0.08) and 0.68 (cluster-robust SE 0.06) when including dummies for the five different task rates and the type of task (Tetris versus transcription). This persistence in allocation preferences across time is consistent with large individual differences in preferences over effort allocation. Sources of these differences could include individuals already knowing at time 0 whether they will be relatively more busy in week 2 or in week 3, or stable individual differences in the curvature of the effort cost function. In sum, these results suggest that stable individual differences can be well-measured in the Augenblick, Niederle, and Sprenger (2015) design, but that individual differences in the present focus parameter β are not the primary explanation for the variation in the differences between time 1 and time 0 preferences.

Figure B2: Revised vs. Initial Effort Allocations in Block 1 (left) and Block 2 (right)

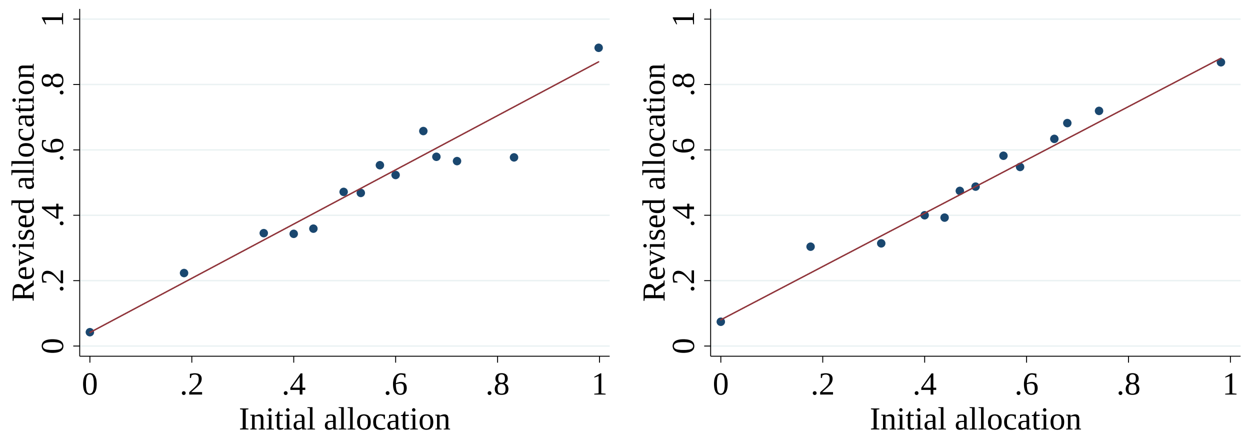

Notes: This figure presents binned scatter plots of the revised versus the initial allocations of effort, pooled over the five task rates and the two types of tasks. The first panel presents results for block 1 (weeks 1-3), while the second panel presents results for block 2 (weeks 4-6). The x-axes correspond to the average effort share designated for weeks 2 and 5 by subjects in weeks 1 and 4, respectively. The y-axes correspond to the average effort share designated for weeks 2 and 5 by subjects in weeks 2 and 4, respectively.

E.2 Identification

We next show how Propositions 4 and 5 can be used to draw inferences about present focus using data in Augenblick, Niederle, and Sprenger (2015) and Augenblick and Rabin (2019).

E.2.1 Set Identification in Augenblick, Niederle, and Sprenger (2015)

We begin by discussing Augenblick, Niederle, and Sprenger (2015). As different decisions are taken at time 0, we define a state (α, ω) to capture both the variation in time 0 and time 1 choices. We denote by $f_{\omega|\alpha}$ the probability of the time-1 state ω conditional on the time-0 state α. Suppose that the cost of effort is additively separable and isoelastic, as assumed in Augenblick, Niederle, and Sprenger (2015). Suppose, moreover, that taste shocks are multiplicative, as in the assumptions of Proposition 5; i.e.,

$$\begin{aligned} u^0(y,z) &= \delta\bar{\theta}^1_\alpha y^\gamma + \quad \delta^2\bar{\theta}^2_\alpha z^\gamma \\ u^1_\omega(y,z) &= \quad \theta^1_{\omega,\alpha} y^\gamma + \beta\delta\theta^2_{\omega,\alpha} z^\gamma \end{aligned}$$

where $\bar{\theta}^1_\alpha := \sum_\omega f_{\omega|\alpha}\theta^1_\omega$, $\bar{\theta}^2_\alpha := \sum_\omega f_{\omega|\alpha}\theta^2_\omega$. We assume that γ is known; Augenblick, Niederle, and Sprenger identify this parameter by observing how preferred allocations vary with the returns to time-2 effort.

The analysis in Supplementary Appendix E.1 above has shown that empirically, much of the variation in individual-level point estimate $\hat{\beta}_i$ is *not* due to variation in time preferences, but rather due to taste shocks. We thus begin by assuming that β is homogeneous in the population. Note that for each value of α that fixes a time-0 choice, $\hat{\beta}_{\omega,\alpha}$—as defined in equation (6) in the body of the paper—is the key object in Proposition 5. In terms of the assumed model primitives, $\hat{\beta}_{\omega,\alpha}$ is given by

$$\begin{aligned}\hat{\beta}_{\omega,\alpha} &= \frac{g(z^{\circ}_{0,\alpha,y}) - g(z_{0,y})}{g(z^{\circ}_{\omega,\alpha,y}) - g(z_{\omega,\alpha,y})} = \frac{\frac{1}{\sum_{\omega} f_{\omega|\alpha}\theta^2_{\omega,\alpha}}\left(\sum_{\omega\in\Omega} f_{\omega|\alpha}\theta^1_{\omega,\alpha}h(y) - \sum_{\omega\in\Omega} f_{\omega|\alpha}\theta^1_{\omega,\alpha}h(y^{\circ})\right)}{\frac{1}{\beta\theta^2_{\omega,\alpha}}\left(\theta^1_{\omega,\alpha}h(y) - \theta^1_{\omega,\alpha}h(y^{\circ})\right)} \\ &= \beta\frac{\bar{\theta}^1_{\alpha}/\bar{\theta}^2_{\alpha}}{\theta^1_{\omega,\alpha}/\theta^2_{\omega,\alpha}}\,.\end{aligned}$$

Similarly, taking first-order conditions for each realization of (α, ω) implies an equation analogous to equation (6) of Augenblick, Niederle, and Sprenger, except with β replaced by

$$\hat{\beta}_{\omega,\alpha} = \beta\frac{\bar{\theta}^1_{\alpha}/\bar{\theta}^2_{\alpha}}{\theta^1_{\omega,\alpha}/\theta^2_{\omega,\alpha}}\,.$$

Thus, the individual-level estimates $\hat{\beta}$ produced by Augenblick, Niederle, and Sprenger give the $\hat{\beta}_{\omega,\alpha}$ that Proposition 4 applies to. Intuitively, this is because under the assumption of isoelastic cost functions with equal curvature, Augenblick, Niederle, and Sprenger's budget set variation allows them to estimate g up to multiplicative taste shocks. Thus, their data set is informationally equivalent to the data set assumed in Proposition 4.

To facilitate identification of β assume that the distribution of $\frac{\bar{\theta}^1_{\alpha}/\bar{\theta}^2_{\alpha}}{\theta^1_{\omega,\alpha}/\theta^2_{\omega,\alpha}}$ conditional on α is independent of α. This assumption is satisfied, for example, if taste shocks are multiplicatively separable over time: $\theta^t_{\omega,\alpha} = \bar{\theta}^t_{\alpha} \times \epsilon^t_{\omega}$ for $t \in \{1, 2\}$. If one does not make such an assumption, more values of β are consistent with the data. This independence assumption implies that the distribution of $\hat{\beta}_{\omega,\alpha}$ does not depend on α. Proposition 5 then implies that $\beta \in [\min \hat{\beta}_{\omega,\alpha}, \max \hat{\beta}_{\omega,\alpha}]$.

Because Augenblick, Niederle, and Sprenger produce an estimate of $\hat{\beta}_{\omega,\alpha}$ for each individual when they estimate equation (6) individual by individual (under the homogeneity assumption), the distribution of $\hat{\beta}_{\omega,\alpha}$ is estimated by the empirical distribution of Augenblick, Niederle, and Sprenger's individual-level estimates $\hat{\beta}_i$, where i indexes the subjects.[40]

Using their estimates we obtain that the 90-percent intervals that exclude the 5 percent highest and 5 percent lowest estimates of $\hat{\beta}_i$ are [0.43,1.17] in block 1, and [0.57, 1.25] in block 2. The 80-percent intervals that exclude the 10 percent highest and 10 percent lowest

[40]Formally, the empirical distribution converges to the distribution of $\hat{\beta}_{\omega,\alpha}$ as the number of subjects goes to infinity.

estimates of $\hat{\beta}_i$ are [0.66,1.12] in block 1 and [0.69,1.14] in block 2.[41] Thus, the data are consistent both with future focus, as well as with present focus if one allows for taste shocks.

Of course, the complete homogeneity assumption may be too strong, even if the analysis surrounding Figure B1 above suggests that it might be a good approximation. Clearly, relaxing this assumption can only make identification harder. Augenblick, Niederle, and Sprenger (2015) report that individuals who desire a choice set restriction in week 4 have a lower estimate of $\hat{\beta}_i$. The average value of $\hat{\beta}_i$ is 0.86 and 0.95 in blocks 1 and 2, respectively, for individuals who want the choice set restriction; it is 0.96 and 1.05 in blocks 1 and 2, respectively, for individuals who do not want the choice set restriction. Thus, there are at least some individual differences in present focus. Our results are easily compatible with the incorporation of observed heterogeneity, as set identification is possible when splitting individuals by their demand for a commitment device in week 4 of the experiment.

Concretely, assume instead that present focus is homogeneous conditional on the decision to take up or not the commitment device. We can then apply Proposition 5 to each of the subgroups. We obtain the following identified sets if we base them on 80-percent intervals that exclude the 10 percent highest and 10 percent lowest estimates. For those who take up commitment, we infer that their β could belong to the set $[0.44, 1.11]$ based on block-1 data, and to the set $[0.66, 1.13]$ based on block-2 data. For those who do not take up commitment, we obtain $[0.73, 1.13]$ and $[0.84, 1.16]$ for blocks 1 and 2, respectively. In particular, note that it is possible that present focus is large for both those who do and don't take up commitment. It is also possible that present focus is small for both groups, and in fact that those who do not take up commitment contracts are future-focused. Both possibilities are realistic. As Carrera et al. (2022) show, uncertainty about the future erodes demand for commitment even among present-focused individuals, so that large present focus is possible even among those who do not take up commitment. On the other hand, Carrera et al. (2022) also show that there may be noise and confusion in commitment take-up, so that even individuals who are time-consistent might erroneously take up commitment contracts. Finally, note that individuals who are future-focused are still time-inconsistent, and thus might take up commitment contracts as well. Thus, theoretically, any value of β can be consistent with commitment take-up, and our identified set is consistent with this theoretical ambiguity.

[41] We report our results for the two blocks separately to facilitate immediate comparison to the results in Augenblick, Niederle, and Sprenger (2015), who focus on block 1 in the body of the paper and relegate separate analysis of block 2 to the appendix.

E.2.2 Point Identification in Augenblick and Rabin (2019)

Assume that there is no variation in θ^2 and that preferences over the second dimension are known. As the second consumption dimension in Augenblick and Rabin (2019) corresponds to small monetary amounts, this assumption is equivalent to assuming that taste shocks do not affect preferences over money and the agent is (approximately) risk neutral over small monetary amounts. We relax the assumption that β is homogeneous and allow for β to be a random variable (capturing variation of β in the population). We first observe that for a fixed value of β and a fixed time-0 choice captured by α, we have that

$$\frac{1}{\beta} = \sum_{\omega \in \Omega} f_{\omega|\alpha} \frac{1}{\hat{\beta}_{\omega,\alpha}} = \mathbb{E}\left[\frac{1}{\hat{\beta}_{\omega,\alpha}} \middle| \alpha, \beta\right].$$

The first equality we derived in equation (7) in the body of the paper and the second equality follows by definition because $f_{\omega|\alpha}$ is the probability of the time-1 state realization ω conditional on the time-0 state realization α. Then taking iterated expectations, we obtain the following immediate corollary of equation (7):

$$\mathbb{E}\left[\frac{1}{\beta}\right] = \mathbb{E}\left[\mathbb{E}\left[\frac{1}{\hat{\beta}_{\omega,\alpha}} \middle| \alpha, \beta\right]\right] = \mathbb{E}\left[\frac{1}{\hat{\beta}_{\omega,\alpha}}\right]. \tag{14}$$

Thus, the average over the observed estimates $\frac{1}{\hat{\beta}_i}$—i.e., present focus estimates produced from each individual's data—constitute an unbiased estimator of $\frac{1}{\beta}$. The logic above does not rely on multiplicative taste shocks. Nearly-identical reasoning can be used to derive (14) under the more general assumptions of Proposition 4.