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 November 2024

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 November 2024
JEL No. C9,D0,D9

ABSTRACT

We study whether it’s possible to identify time-inconsistent preferences in empirical designs where preferences are elicited in advance at time 0, and then again at time 1, after the agent receives additional information. For single-peaked preferences, time-consistency is rejected only when the time-1 ranking between a pair of alternatives is always the reverse of the time-0 ranking. We establish variations and generalizations of this result. Since such stark reversals are rarely observed, choice-revision designs require stronger identification assumptions than perhaps previously appreciated. But we show that time-inconsistency is identifiable in environments where preferences over alternatives can be “priced-out.”

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. Consider an analyst who observes individuals who tend to revise their consumption plans in a certain *systematic* direction, while being subject to random taste shocks, such as ex-ante unpredictable variation in time constraints, appetite, fatigue, or additional consumption opportunities. 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, approximately 50 percent of individuals choose a healthy over an unhealthy snack, but this fraction declines to approximately 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. Especially in the last ten years, 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

[1] Work allocation: Augenblick et al. (2015), Abebe et al. (2021), Andreoni et al. (forthcoming), Barton (2015), Corbett (2016), Imas et al. (2022), Kölle and Wenner (2023), Augenblick and Rabin (2019), Fedyk (forthcoming), see Imai et al. (2021) for a review; Financial plan-making: Kuchler and Pagel (2021); Entertainment: Read et al. (1999), Milkman et al. (2009), Bartos et al. (2021); Opioid use: Badger et al. (2007); Nutrition choice: Sadoff et al. (2019).

[2] Typical designs elicit preferences from an identical choice set at two different points in time 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. This is incentive compatible for expected utility preferences that are linear in probabilities.

of time-preference parameters.[3] Additionally, revision designs 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. In the first part of the paper, we show that rejecting time consistency typically requires much more stark patterns of choice than perhaps previously appreciated by the literature. We illustrate this with numerical examples that capture several types of prominent designs concerning food choice and allocation of effort or consumption over time. 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. In the second part of the paper, however, we show that a subset of revision designs—namely, those that monetize valuations for different consumption opportunities, as in Augenblick and Rabin (2019), Carrera et al. (2022), and several other papers—can provide robust estimates of time inconsistency under plausible assumptions.

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 different amounts of information about the agent's preferences. For example, in some designs, researchers elicit only the most preferred alternative; in other designs agents choose from different budget sets, which gives researchers more information. 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]

[3]See, e.g., Augenblick et al. (2015), Andreoni et al. (forthcoming), Augenblick and Rabin (2019), Augenblick (2018). See also Halevy (2015) for a critical discussion.

[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 et al. (2019), Bai et al. (2021), 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 et al., 2010, Brunnermeier et al., 2008), overoptimism

Our Theorem 1, which formalizes and generalizes the numerical examples in Sections 1.1 and 1.2, 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 observed, the data can be rationalized by TC-EU as long as there is no *sure direct preference reversal*—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 *sure indirect preference reversal*—i.e., a collection of alternatives $x_1, \dots, x_k$ 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 sure direct preference reversal 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, most 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 5.

We view Theorems 1 and 2 as our main results that cover a range of environments of economic interest. In addition, we 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 sure indirect preference reversals. The reason is that there may be *sure lottery preference reversals*—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 sure lottery preference reversals. 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 sure lottery preference reversals.

In Section 5 we present a set of conditions under which the degree of time inconsistency can be identified, formalizing the numerical examples in Section 1.3. Roughly speaking, point

(Browning and Tobacman, 2015, Breig et al., 2024) or other misperceptions of the time-1 decision environment (Sadoff et al., 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.

identification can be obtained augmenting standard designs to elicit the agent's valuations of different alternatives using a numeraire commodity, such as future money. The key condition on this numeraire is that the agent's beliefs about the marginal utility of consuming the numeraire do not change between time 0 and time 1. This condition makes it possible to make cardinal comparisons about the strength of the agent's preferences across different states of the world, or, in effect, to monetize the taste shocks. This, in turn, makes it possible to more precisely differentiate between how much of the difference in the agent's choices between time 0 and time 1 is due to taste shocks versus time inconsistency. Recent designs that make use of this strategy include both laboratory studies (Augenblick and Rabin, 2019, Augenblick, 2018, Fedyk, forthcoming) and field studies (Acland and Levy, 2015, Chaloupka et al., 2019, Carrera et al., 2022, Allcott et al., 2022).

We end Section 5 by discussing several practical considerations for utilizing revision designs in experimental work. First, we discuss how to quantitatively evaluate the prevalence of random taste shocks, and show that recent empirical work suggests that these are quantitatively very meaningful. Second, we also address tests the attempts 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 can be a useful source of evidence, but it does not fully mitigate the non-identification issues.

Related Literature There seems to be fairly broad agreement that the identification of time inconsistent preferences (in the lab or in the field) poses a formidable challenge. Echenique et al. (2020), Blow et al. (2021), Echenique and Tserenjigmid (2023) consider the identification of exponential discounting and quasi-hyperbolic discounting from observed choice over consumption streams. Levy and Schiraldi (2020) present conditions for which time preferences are identified in dynamic discrete choice problems with history-dependent choice sets.[7] 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. Oliveira and Lamba (2023) 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.

[7] Our paper is also broadly related to the which literature focuses on econometric techniques for identifying "behavioral" models. Barseghyan et al. (2013) and Barseghyan et al. (2021) develop techniques for identifying risk preferences in the presence of probability distortions and limited consideration. Rees-Jones and Wang (2022) develop techniques for identifying reference points.

While the focus of this paper is on dynamic preference reversals, there is also an important literature on static preference reversals, where participants choose between rewards at different points in time, and the research question is how the discount rate between two periods t and $t+1$ changes with t (see, e.g., Cohen et al. 2020 for a review). Halevy (2015) clarifies that such studies test the stationarity assumption, while revision designs test the time invariance assumption. There is also an active conversation about the potential confounding role of uncertainty in these designs (Halevy, 2005, 2008, Andreoni and Sprenger, 2012, Chakraborty et al., 2020).

Another active literature studies how (non-exponential) time discounting may arise from biases arising from perceptual and cognitive mechanisms (e.g., Gabaix and Laibson, 2022, Enke and Graeber, 2023, Enke et al., forthcoming). Many of these models imply that choice is noisy, which further complicates identification. Data sets that we show reject TC-EU could instead be due to these alternative mechanisms, rather than actual time-inconsistent preferences. In fact, to the extent that these alternative models of noisy perception lead people to have time-0 and time-1 information sets that are consistent with our baseline assumptions, our results provide a general characterization of how such models could rationalize time-inconsistent-like behavior.

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 Online Appendix C we flesh out connections to technical results in the literatures on social choice, random utility models, and dynamically consistent preferences over acts.

The rest of this paper proceeds as follows. Section 1 presents numerical examples that illustrate our main results, both about the difficulty of identifying time inconsistency and about paths forward. Section 2 presents the formal model. Section 3 presents our main results about the types of data sets that are consistent with TC-EU, and Section 4 sketches the proof of the main results. Section 5 presents results about economic environments where it is possible to identify the degree of time inconsistency. Section 6 concludes. Proofs are relegated to the Appendix.

1 Motivating Examples

In this section, we provide a series of examples that illustrate the main ideas of our paper. The first two examples illustrate the difficulty of identifying time inconsistency. The last two examples illustrate the types of designs that can be used to identify time inconsistency.

1.1 Intertemporal Allocation of Consumption

To illustrate the difficulty of identifying time inconsistency from a revision design, consider the following stylized example based on research designs that study people's intertemporal allocation of effort.[8] 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 of effort $x \in [0, 1]$ 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. With expected utility preferences, 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 et al., 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.[9]

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 0.45 to period 1, with a standard deviation of 0.. The observed distribution of the ratio of efforts $x/(1-x)$ follows a log-normal distribution.[10] 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, and they 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.[11]

[8] For experimental studies using similar designs see, e.g., Augenblick et al. (2015), Barton (2015), Corbett (2016), Kölle and Wenner (2023), Andreoni et al. (forthcoming), Abebe et al. (2021), Imas et al. (2022); and Imai et al. (2021) for a review.

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

[10] The allocations of 50 and 45 percent, respectively, are roughly in line with the data in Augenblick et al. (2015).

[11] Online Appendix B provides formal mathematical calculations for each of these different sets of as-

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 independent case where θ_1 is known at time 0 and θ_2 is learned only at time 1. Rows 5 and 6 instead consider the case where θ_2 is known at time 0 the information obtained between time 0 and time 1 is θ_1. Rows 7 and 8 consider the case where θ_1 and θ_2 are distributed independently, with nothing learned at time 0 and only θ_1 learned at time 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 7 and 8 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.

	Distribution of shocks	*Information*		γ	Estimated
		time 0	*time 1*		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

Table 1: Implied time inconsistency under different information revelation assumptions

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

sumptions about information revelation.

"intertemporal Euler equation" (Augenblick et al. 2015, Imai et al. 2021)—that corresponds to the assumptions in rows 1 and 2. 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*.[12] 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 the experiment of Read and van Leeuwen (1998) and related studies,[13] consider an agent who chooses 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.[14]

To rationalize the results of Read and van Leeuwen with time-consistent preferences, consider agents with the following preference: At time 1, the agents feel gorged with probability 20%, in which case they crave healthy food, so that the utility difference between the unhealthy and healthy snack equals -5. With probability 80% 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% 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, $0.38 + 0.62 \cdot 0.2 = 50\%$ choose the healthy snack at time 0. However, only 20% choose the healthy snack at time 1. Moreover, the direction of revisions is asymmetric: $38\% \times 0.8 = 30\%$ 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.

[12] In Appendix B, Augenblick et al. (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.

[13] Empirical designs with similar structures include the food-delivery field experiment of Sadoff et al. (2019), Read et al.'s (1999) study of choice between high-brow and low-brow video rentals, and Milkman et al.'s (2009) quasi-experimental extension of Read et al. (1999).

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

1.3 Identifying Time Inconsistency Using Willingness-to-Pay Designs

We now provide two examples, the first one hypothetical, and the second one based on recent experiments, where time inconsistency can be identified and quantified. The approach in both examples is to "price out" the utility from each alternative.

Example 1 (Read and van Leeuwen with Money)**.** As in the example in Section 1.2, suppose that the agent chooses between a healthy snack and an unhealthy snack. As before, the agent either feels normal (with probability 80%) or gorged (with probability 20%). 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.[15] 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, a time-consistent expected utility maximizer should have a WTP for the healthy snack at time 0 the equals 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 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.

Example 2 (Augenblick and Rabin 2019, Augenblick 2018, Fedyk forthcoming)**.** Augenblick and Rabin (2019) elicit willingness to work for various amounts of money. Suppose that preferences follow the quasi-hyperbolic discounting model and are given by

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

where $c(x)$ is the cost of x units of effort and z is monetary 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

[15] 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.

assumes that utility is quasi-linear in money.[16]

Now suppose that at time 0, the agent requires at least $10 to commit to 5 units of work, while at time 1 the average minimum payment to perform 5 units of work is $15. Then the modeling assumptions imply that

$$\beta \cdot \mathbb{E}_0[\theta c(5)] = \beta \cdot 10$$
$$\mathbb{E}_0[\theta c(5)] = \beta \cdot 15$$

from which it immediately follows that $\beta = 2/3$.

Examples 1 and 2 illustrate that one strategy to obtain point identification is to monetize agents' preferences over different consumption bundles, when the following conditions plausibly hold: (i) preferences are separable over money and these consumption bundles, (ii) the marginal utility of money does not vary with shocks to utility from the different consumption bundles, 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 Acland and Levy (2015), Chaloupka et al. (2019), Carrera et al. (2022) and Allcott et al. (2022).

More generally, identification is possible when there is a numeraire commodity that satisfies the three conditions listed above for money. For example, another form of a numeraire commodity might be consumption or effort that is at least one year in the future, while the time difference between time-0 and time-1 choices is only one week. In this case, it is plausibly that information received between time 0 and time 1 is unlikely to alter the agent's expected utility from the numeraire commodity in a year.

Interestingly, while fungibility of money makes identification of time preferences more difficult in designs where people choose between monetary amounts in different points in time (Ericson and Laibson, 2019), it does not create problems for the designs presented in Examples 1 and 2. In fact, to the extent that fungibility helps justify the assumptions made about utility from money in those examples, it aids identification.

[16] This is not crucial. Alternatively, the analyst could gather additional data on preferences over monetary lotteries to estimate a nonlinear utility function over money.

2 Model

There is an agent who has a preference over a finite set of alternatives X at time 0 and time 1.[17] 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, different types of revision designs provide different amounts of information about the agent's preferences. We leave our model general enough to cover a range of what the analyst can observe. Formally, this amounts to assuming that the preferences can be either complete, i.e. the ranking of any two alternatives is observed, or incomplete.[18] 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 $y \prec x$.[19] 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.

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 observes the distribution over the different rankings that result from variation in time-1 states. In practice, it is not possible to fully observe a given agent's distribution over rankings. Such a possibility corresponds to either the limit cases 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

[17]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.

[18]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.

[19]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 y 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.

type of homogeneity assumption. Of course, identification can only become more difficult if the distribution over rankings cannot be fully observed. Thus, the homogeneity assumptions are also crucial for identification.[20]

To formally model taste shocks that give rise to randomness over rankings, 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 are as many states as realizations of the time-1 preference, with each realization of the time-1 preference corresponding to a state, so that $\Omega = \{1, \ldots, n\}$.[21] 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 the expected utility of y; i.e., for all $x, y \in X$,

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

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

[20] In practice, because the validity of the time inconsistency estimates requires all of the theoretical assumptions to hold, empirical estimates of time inconsistency could in principle be biased even in cases where our theoretical results show that identification is possible.

[21] 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 $\{1, \ldots, n\}$.

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*[22] or other forms of overoptimism[23] is not compatible with our definition of TC-EU.[24]

3 Rejecting Time Consistency

3.1 General Preferences

We first consider the case where we impose no further restrictions on the agent's preferences. 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 the preference $\preceq$, which orders elements in ascending order $x_1 \preceq x_2 \preceq \ldots \preceq x_{|X|}$ 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 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.

Definition 2 (Sure Lottery Preference Reversal)**.** A data set $(\preceq^0, \preceq^1, f)$ exhibits a sure lottery preference reversal 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 $\omega \in \Omega$.

The absence of a sure lottery preference reversal is necessary for a data set to be consistent with TC-EU, as otherwise the analyst knows that the agent prefers a lottery L over L' at time 0, but at time 1 always prefers the lottery L' over L which can never happen under TC-EU preferences. Our next result establishes that absence of sure lottery preference reversals is also sufficient and thus 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 sure lottery preference reversals.*

[22]See e.g. Kahneman and Tversky (1982), Buehler et al. (2010), Brunnermeier et al. (2008).

[23]See e.g. Browning and Tobacman (2015), Breig et al. (2024).

[24]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.

Proposition 1 is "folk-wisdom" and similar results have appeared in other contexts, such as social choice.[25] The result of the proposition is, however, difficult to use to reject TC-EU because the space of possible lotteries is large. To obtain sharper results we will thus study environments that satisfy additional economically plausible restrictions.

3.2 Single-Peaked Preferences

We next 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 3 (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' preferences do not directly 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 4 (Sure Direct Preference Reversal)**.** A data set $(\preceq^0, \preceq^1, f)$ exhibits sure direct preference reversals 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

[25] For example, 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. 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.

they are time-consistent. It follows immediately from the definition that any data set that is consistent with TC-EU cannot exhibit sure direct preference reversals. 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 sure direct preference reversals 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 5 (Sure Indirect Preference Reversal)**.** A data set $(\preceq^0, \preceq^1, f)$ exhibits sure indirect preference reversals 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.[26]

Clearly, a sure direct preference reversal is a sure indirect preference reversal with a cycle of length 2. And as Lemma 3 in Appendix A.1 shows, sure indirect preference reversals imply sure direct preference reversals when the time-0 preference is complete. In general, however, data sets can exhibit sure indirect preference reversals without exhibiting sure direct preference reversals (see Example 3 below). The absence of a sure indirect preference reversal 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 sure indirect preference reversals.

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

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

[26]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.

(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 sure direct preference reversals. However, the more variability there is in an agent's time-1 preference, the more unlikely sure preference reversals 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 may be uninformative about whether an agent is time-consistent or time-inconsistent when there is significant stochasticity in time-1 choices.[27] We are not aware of any revision design studies where the data contain a sure direct or sure indirect preference reversal.

Remark 1. *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 $\preceq^1_j$ occurs with positive probability. An analogous comment applies to the other theorems in this section.*

Remark 2. *An immediate conclusion obtained by combining Theorems 1 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 sure lottery preference reversal if and only if it admits a sure direct preference reversal. 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.*[28]

To illustrate the importance of the conditions in our Theorem 1, we end this section with an example of a data set with incomplete preferences that admits an indirect but not sure direct preference reversal, and an example of a data set where preferences are not single-peaked, where there is a sure lottery, but not a sure indirect preference reversal.

[27] 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 sure direct preference reversals. 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.

[28] We thank Xiaosheng Mu and Pietro Ortoleva for pointing out this equivalence as an alternative perspective on our results.

Example 3. 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 sure direct preference reversal. 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 sure indirect preference reversal is $1 \prec^1 2 \prec^0 4 \prec^1 3 \prec^0 1$.

A natural question is whether a sure indirect preference reversal is necessary for a data with non-single-peaked preferences that is not consistent with TC-EU. The example below shows that they do not.

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

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

It is easy to check that the preferences exhibit no sure direct or sure indirect preference reversals. 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 $\{1, 3, 5\}$.[29]

3.3 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,

[29]Example 4 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 sure direct preference reversals (i.e., the conclusion of Theorem 1 holds for non single-peaked data with $|\Omega| = 2, |X| < 6$). Furthermore, Example 4 is the only data set (up to relabeling of the states) with $|\Omega| = 2, |X| = 6$ where Theorem 1 does not hold. We verified this using a computer program that checks for each configuration of preferences if there is a sure direct preference reversal and solves the linear programming problems that corresponds to checking if there is an TC-EU representation. But there are many other examples of data sets not consistent with TC-EU and not exhibiting sure direct preference reversals when $|\Omega| > 2$ or $|X| > 6$.

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 $u_\omega(x) = v^1_\omega(x) + v^2_\omega(1-x)$, 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 it 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 TC-EU is that the time-0 preference must be single-peaked. Our next result shows that this condition, together with no sure direct preference reversal, 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 sure direct preference reversals.*

We note that the additional restriction that the time-0 preference is also single-peaked implies that the lack of sure direct preference reversals is enough to guarantee consistency with TC-EU; it is not necessary to consider sure indirect preference reversals more generally in this case.[30]

3.4 Preference Data Consistent with Stochastic Choice

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 sure preference reversal 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 $x \preceq^1_\omega y$.

[30] In Example 3, 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 $3 \prec^0 1$.

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. 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.

3.5 Cardinal Preferences Data

Proposition 1 can be generalized to 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 sure lottery preference reversals accordingly.

Definition 8. A data set (U^0, U^1, f) exhibits sure lottery preference reversals 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 sure lottery preference reversals.*

4 A Sketch of the Proof of Our Main Results

We next sketch a proof of Theorems 1 and 2. To simplify exposition, we focus on the case where all preferences are complete and strict, and skip over some technical details. Complete proofs for the general case can be found in the Appendix. Recall that we define $U^0 \subset \mathbb{R}^{|X|}$ to be the set of utilities consistent with $\preceq^0$, and we define $\bar{U}^1 \subset \mathbb{R}^{|X|}$ to be the closure of the set of *expected* utilities consistent with the random preference $\preceq^1$.

By Proposition 1, there exists two lotteries L, L', represented by vectors of probabilities of the corresponding alternatives, such that $p \cdot u^0 \geq 0 \geq p \cdot u^1$ for $p = L - L'$ and all $u^0 \in U^0, u^1 \in \bar{U}^1$, with one of the inequalities strict. Because the entries of p sum to zero, a sure direct preference reversal corresponds to a vector $p \in \mathbb{R}^{|X|}$ where only two of the entries are non-zero.

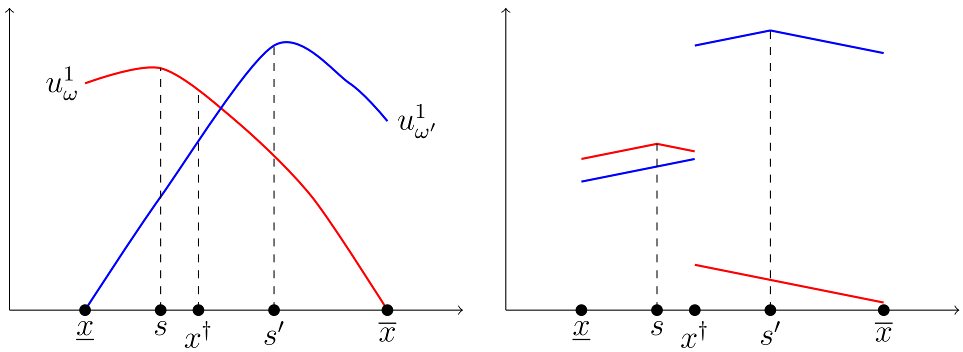

Figure 1: Illustration of the preferences in state ω, ω'.

Thus, if there are only two alternatives to choose from, i.e. $|X| = 2$, the existence of a sure lottery preference reversal implies the existence of a sure direct preference reversal. We prove the result by induction over the number of alternatives in $|X|$.

Denote by $x^\dagger$ the worst alternative at time 0: $x^\dagger \prec^0 x$ for all $x \neq x^\dagger$. If $x^\dagger \succ^1_\omega x$ for all ω, then there is a sure direct preference reversal. Thus, it is without loss to focus on the case where $x^\dagger$ does not dominate any other alternative according to the time-1 preference. Denote by $\underline{x}, \overline{x}$, respectively, the smallest and largest elements of X according to the natural order on the reals. For this argument, suppose that $x^\dagger$ is not equal to $\underline{x}$ or $\overline{x}$. As $x^\dagger$ does not dominate $\underline{x}$ or $\overline{x}$ at time 1, there exists at least one state ω where $x^\dagger \prec^1_\omega \underline{x}$ and at least one state ω' where $x^\dagger \prec^1_{\omega'} \bar{x}$.

As preferences are single peaked, that implies that the peak s of $\prec^1_\omega$ must be to the left of $x^\dagger$ in state ω and the peak s' of $\prec^1_{\omega'}$ must be to the right of $x^\dagger$. This is illustrated in the left panel of Figure 1. The right panel illustrates another utility representation of the ordinal (single-peaked) preferences in states ω, ω' given by

$$u^1_\omega(x) = -\epsilon|x - s| - \delta \mathbf{1}\{x \geq x^\dagger\}$$
$$u^1_{\omega'}(x) = -\epsilon|x - s'| + \delta \mathbf{1}\{x > x^\dagger\}\,.$$

The sum of these utility functions is given by

$$u^1_\omega(x) + u^1_{\omega'}(x) = -\epsilon(|x - s| + |x - s'|) - \delta \mathbf{1}\{x = x^\dagger\}\,.$$

Because $\bar{U}^1$ is the closure of expected utilities consistent with $\preceq^1$, and we can choose ϵ arbitrarily small, it follows that $-\delta \mathbf{1}\{x = x^\dagger\} \in \bar{U}^1$.

At the same time, if $u^0 \in U^0$, then so is $u^0 - \delta \mathbf{1}\{x = x^\dagger\}$ for any $\delta \geq 0$, because $x^\dagger$ is the least preferred element at time 0. Consequently, for any $u^0 \in U^0, u^1 \in \bar{U}^1$ and any $\delta \geq 0$ we have that the separating hyperplane p satisfies

$$p \cdot u^0 - \delta p(x^\dagger) \geq 0 \geq p \cdot u^1 - \delta p(x^\dagger)\,.$$

This equation can only be satisfied for all $\delta \geq 0$ if $p(x^\dagger) = 0$. Thus, because $p(x^\dagger) = 0$ we have found a separating hyperplane which involves only $|X| - 1$ elements and we can thus remove $x^\dagger$ from the problem and the resulting data set with $|X| - 1$ alternatives will still not be compatible with TC-EU preferences.

Finally, if $x^\dagger$ is either $\underline{x}$ or $\overline{x}$, it follows from the same argument that $p(x^\dagger) = 0$ as $x^\dagger$ must already be the least preferred option in some state at time 1. We have thus completed the induction step.

5 Estimating Time Inconsistency

In the prior section we have shown that many patterns of choice—including those previously described as evidence of time inconsistency—can, in fact, be rationalized with time consistent preferences. In this section, we formalize the identification approaches illustrated by the examples in Section 1.3.

5.1 Data Sets with Identification of Time Inconsistency

To formalize our results about identification, we start 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,

$$u^0(y,z) = h^0(y) + g^0(z)$$
$$u^1_\omega(y,z) = h^1_\omega(y) + g^1_\omega(z)$$

are consistent with $\preceq^0$ and $\preceq^1_\omega$, respectively.[31]

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°.

[31] 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.

Examples 1 and 2 illustrated how the 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.[32] 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 1, 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 $\omega \in \Omega$.

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.

To see how this proposition generalizes Example 2, take $g(z) = \beta z$, $h^0(y) = -\beta\mathbb{E}[\theta_\omega]c(y)$, and $h^1_\omega(y) = -\theta_\omega c(y)$. The reference level of effort is y°, and $y = y^\circ + 5$ is 5 additional units of work. At time 0, the agent requires $z_{0,y} - z^\circ_{0,y}$ money to complete this additional effort, and at time 1, in state ω, the agent requires $z_{\omega,y} - z^\circ_{\omega,y}$ money to complete this additional effort. The goal is to identify $\beta = \frac{h^0(y) - h^0(y^\circ)}{\sum_\omega f_\omega[h^1_\omega(y) - h^1_\omega(y^\circ)]}$. This is given by the ratio of the willingness to accept money for effort at time 0, $g(z^\circ_{0,y}) - g(z_{0,y}) = 10$, to the average willingness to accept money for effort at time 1, $\sum_\omega f_\omega[g(z^\circ_{\omega,y}) - g(z_{\omega,y})] = 15$. Thus, $\beta = 2/3$.

5.2 Partial Identification of Quasi-hyperbolic Discounting

Importantly, additive separability and responsiveness are not by themselves enough to achieve point identification, even with additional parametric assumptions about the nature of time inconsistency and taste shocks. This illustrates that having a state-independent numeraire commodity, as in Proposition 4, is crucial for point identification.

[32] 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.

Specifically, we focus on the quasi-hyperbolic discounting model with multiplicative taste shocks, which generalizes our example in Section 1.1.

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

$$\begin{aligned} u^0(y,z) &= \beta \sum_{\omega\in\Omega} f_\omega \left(\theta^1_\omega h(y) + \theta^2_\omega g(z)\right) \\ u^1_\omega(y,z) &= \theta^1_\omega\, h(y) + \beta\, \theta^2_\omega\, g(z)\,. \end{aligned} \tag{4}$$

This captures the Section 1.1 example if h corresponds to time-1 effort cost and g corresponds to time-2 effort cost.

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

$$\hat\beta_\omega = \frac{g(z^\circ_{0,y}) - g(z_{0,y})}{g(z^\circ_{\omega,y}) - g(z_{\omega,y})}\,. \tag{5}$$

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

$$\beta = \frac{1}{\sum_{\omega\in\Omega} f_\omega \hat\beta_\omega^{-1}}\,. \tag{6}$$

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, formalized in Proposition

5, provides a generalization of the numerical examples in Section 1.1.

To state the proposition, we note that if preferences can be represented as in (4), then there exists an additively separable representation (h^0, h^1, g^0, g^1) (as in Definition 9) of the preferences where the function g^1 does not depend on the state and equals g^0. To obtain this representation, simply set $h^0 \equiv \frac{\sum_\omega f_\omega \theta^1_\omega}{\sum_\omega f_\omega \theta^2_\omega} h$, $h^1_\omega \equiv \frac{\theta^1_\omega}{\beta \theta^2_\omega} h$, $g^0 = g$ and $g^1 = \beta g$.

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 in experiments, 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. Thus, Proposition 5 suggests a different way of treating individual-level estimates of $\hat{\beta}_\omega$ than is typically done in experiments. Instead of interpreting the distribution of $\hat{\beta}_\omega$ literally as the distribution of time preferences, Proposition 5 illustrates that this distribution can instead be consistent with a homogeneous present focus parameter β that can take on any value in the support of the distribution of $\hat{\beta}_\omega$.

5.3 Practical Considerations for Experimental Designs

5.3.1 Quantifying the Importance of Taste Shocks

Our theoretical results clarify that taste shocks tend to impede identification of time inconsistency. Both intuition and evidence suggest that taste shocks are quantitatively important. 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.

One way of quantifying the relevance of taste shocks is 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., $\hat{\beta}^1_i = \hat{\beta}^2_i$ for each agent i in the quasi-hyperbolic discounting model), then that suggests that there are no taste shocks. In that case, the $\hat{\beta}_i$ can be interpreted as individual-level estimates of the present focus parameter. On the other hand, a low correlation between $\hat{\beta}^1_i$ and $\hat{\beta}^2_i$ is consistent with prevalent taste shocks.

5.3.2 Commitment

Although demand for choice set restrictions cannot be used to point identify time-preference parameters (see, e.g., Carrera et al., 2022),[33] in some cases it serves as a useful correlate of time preferences that can help refine their set identification. For example, Augenblick et al. (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. Such additional covariates can be productively used in conjunction with Proposition 5 in the following simple way: apply the Proposition to the set of people choosing a commitment contract and separately to the set of people not choosing the commitment contract. This should refine the identified sets.

5.3.3 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 discuss in Online Appendix A.

6 Conclusion

Our mathematical and numerical results show 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.

[33]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).

The identification challenge posed by taste shocks and learning might also offer a potential explanation for several empirical regularities that have been extensively examined in the time preference literature (e.g., Cohen et al., 2020). Evidence suggests that individual estimates of time preference parameters exhibit significant heterogeneity, domain-specificity, and correlations with personal characteristics such as age and cognitive ability. Moreover, some studies have shown that time preference parameters estimated in non-monetary domains tend to be smaller than those estimated in monetary domains. Some of these regularities might be due to how the prevalence of taste shocks varies across individuals and domains, which is a potentially fruitful avenue for additional empirical investigation.

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 time consistency. For example, it is unlikely to be mere coincidence or file-drawer bias that in most circulated papers, the systematic reversals tend 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 (e.g., Augenblick et al., 2015, Sadoff et al., 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.

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

References

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

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

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

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

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

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

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

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

_ , Michael Callen, Karrar Hussain, Muhammad Yasir Khan, and Charles Sprenger, "Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan," *Journal of the European Economic Association*, forthcoming.

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

_ and Matthew Rabin, "An Experiment on Time Preference and Misprediction in Unpleasant Tasks," *The Review of Economic Studies*, 2019, *86* (3), 941–975.

_ , Muriel Niederle, and Charles Sprenger, "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *The Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.

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

Bai, Liang, Benjamin Handel, Ted Miguel, and Gautam Rao, "Self-Control and Demand for Preventive Health: Evidence from Hypertension in India," *The Review of Economics and Statistics*, 2021, *103* (5), 835–856.

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum, "The Nature of Risk Preferences: Evidence from Insurance Choices," *American Economic Review*, 2013, *103* (6), 2499–2529.

_ , Maura Coughlin, Francesca Molinari, and Joshua C. Teitelbaum, "Heterogeneous Choice Sets and Preferences," *Econometrica*, 2021, *89* (5), 2015–2048.

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

Bartos, Vojtech, Michal Bauer, Julie Chytilová, and Ian Levely, "Psychological Effects of Poverty on Time Preferences," *The Economic Journal*, 2021, *131* (638), 2357–2382.

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

Blow, Laura, Martin Browning, and Ian Crawford, "Non-parametric analysis of time-inconsistent preferences," *The Review of Economic Studies*, 2021, *88* (6), 2687–2734.

Breig, Zachary, Matthew Gibson, and Jeffrey G. Shrader, "Why Do We Procrastinate? Present Bias and Optimism," *Working Paper*, 2024.

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

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

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

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

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

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

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

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

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

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

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

DellaVigna, Stefano and Ulrike Malmendier, "Paying Not to Go to the Gym," *American Economic Review*, 2006, *96* (3), 694–719.

Echenique, Federico and Gerelt Tserenjigmid, "Revealed preferences for dynamically inconsistent models," *arXiv preprint arXiv:2305.14125*, 2023.

_, Taisuke Imai, and Kota Saito, "Testable implications of models of intertemporal choice: Exponential discounting and its generalizations," *American Economic Journal: Microeconomics*, 2020, *12* (4), 114–143.

Enke, Benjamin and Thomas Graeber, "Cognitive Uncertainty," *The Quarterly Journal of Economics*, 2023, *138* (4), 2021–2067.

_, _, and Ryan Oprea, "Complexity and Hyperbolic Discounting," *Journal of the European Economic Association*, forthcoming.

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

Ericson, Keith M. and David Laibson, "Intertemporal Choice," in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics*, Vol. 2, Elsevier, 2019, chapter 1, pp. 1–67.

Fedyk, Anastassia, "Asymmetric Naivete: Beliefs about Self-Control," *Management Science*, forthcoming.

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

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

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

Gabaix, Xavier and David Laibson, "Myopia and Discounting," 2022. NBER Working Paper No. 23254.

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

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

_ and _, "Random expected utility," *Econometrica*, 2006, *74* (1), 121–146.

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

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

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

Heidhues, Paul and Botond Kőszegi, "Futile Attempts at Self-Control," *Journal of the European Economic Association*, 2009, *7* (2), 423–434.

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

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

Imas, Alex, Michael Kuhn, and Vera Mironova, "Waiting to Choose: The Role of Deliberation in Intertemporal Choice," *American Economic Journal: Microeconomics*, 2022, *14* (3), 414–440.

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

Kahneman, Daniel and Amos Tversky, "Intuitive Prediction: Biases and Corrective

Procedures," in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, 1982, pp. 414–421.

Kölle, Felix and Lukas Wenner, "Is Generosity Time-Inconsistent? Present Bias across Individual and Social Contexts," *The Review of Economics and Statistics*, 2023, *105* (3), 683–699.

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

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

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

Levy, Matthew and Pasquale Schiraldi, "Identification of Intertemporal Preferences in History-dependent Dynamic Discrete Choice Models," *Working Paper*, 2020.

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

Mahajan, Aprajit, Christian Michel, and Alessandro Tarozzi, "Identification of Time-Inconsistent Models: The Case of Insecticide Treated Nets," *Review of Economic Studies*, forthcoming.

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

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

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

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

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

Oliveira, Henrique De and Rohit Lamba, "Rationalizing dynamic choices," *Available at SSRN 3332092*, 2023.

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

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

Rees-Jones, Alex and Ao Wang, "An Approach to Testing Reference Points," 2022.

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

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *The*

Review of Economic Studies, 1955, pp. 165–180.
Strzalecki, Tomasz, “Lecture notes in Decision Theory,” November 2021.

A Appendix

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

$$x \trianglelefteq^* y \quad \Leftrightarrow \quad x \preceq^0 y \text{ or } x \preceq^1_* y\,. \tag{7}$$

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 sure direct preference reversals 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 sure direct preference reversal, 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 sure indirect preference reversals 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 sure indirect preference reversal.

(ii) The data set has no sure direct preference reversal and $\trianglelefteq^*$ satisfies the only weak cycles condition.

Proof. $(ii) \Rightarrow (i)$: Suppose that the elements $x_1, \ldots, x_k$ constitute a sure indirect preference reversal. 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 sure direct preference reversal, then Lemma 1 implies that

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

and thus a violation of the only weak cycles condition.

$(i) \Rightarrow (ii)$: Suppose (ii) does not hold. If the data set has a sure direct preference reversal it also has a sure indirect preference reversal, as any sure direct preference reversal is a sure indirect preference reversal with a cycle of length 2. Thus, suppose that there is no sure direct preference reversal but that there is a non-weak cycle in $\trianglelefteq^*$ 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 sure indirect preference reversal. □

Lemma 3. *If $\preceq^0$ is complete and the data set $(\preceq^0, \preceq^1, f)$ exhibits sure indirect preference reversals then there exists a sure direct preference reversal.*

Proof. We prove this result by contraposition. Assume that there is no sure direct preference reversal. Then $\trianglelefteq^*$ preserves the asymmetric part of $\preceq^0$ by Lemma 1. As $\preceq^0$ is by assumption complete, $\trianglelefteq^*$ must be complete as well. As $\preceq^0$ is transitive, $\trianglelefteq^*$ can only admit weak cycles. But then Lemma 2 implies that there is no sure indirect preference reversal. $\square$

Lemma 4. *Suppose that $\trianglelefteq^*$ 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 $y \trianglelefteq x$ is if there is a cycle in $\trianglelefteq^*$ involving x and y. Since $x \lhd^* y$, this cycle is not a weak cycle. $\square$

Lemma 5. *Suppose that $\preceq^0$ and $\preceq^1$ are single-peaked, and that there are no sure direct preference reversals. Then there exists a nonempty set $X^* \subseteq \{\min X, \max X\}$ such that $x \trianglelefteq x^*$ does not hold for any $x \in X \setminus X^*$ and $x^* \in X^*$, and such that $\min X \sim_{\trianglelefteq} \max X$ 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 $\trianglelefteq$

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

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 sure direct preference reversal. Thus, if $\min X \succ^0 y$, then there is a sure direct preference reversal. A symmetric argument also shows that $\min X \succ^1_* y$ implies a sure direct preference reversal.

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 sure direct preference reversals, 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 sure direct preference reversals 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 sure direct preference reversals.

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 $y \in X \backslash \{\min X, \max X\}$ 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 sure direct preference reversals 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 sure direct preference reversal. □

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

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

□

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 $\mathbb{E}[u_\omega(x)] = \mathbb{E}[v_\omega(x)] - \mathbf{1}_{x=m}\Delta$.*

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 $\mathbb{E}[u_\omega(x)] = \mathbb{E}[v_\omega(x)] - \mathbf{1}_{x=m}\Delta$.

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.

□

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 sure direct preference reversal. By Lemma 2 this implies a sure indirect preference reversal; hence, if U^0 and $\bar{U}^1$ do not intersect, there must be a sure indirect preference reversal.

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 $u^0 \in U^0, u^1 \in \bar{U}^1$

$$p \cdot u^0 \geq 0 \geq p \cdot u^1 \tag{8}$$

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 sure direct preference reversal 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 $\trianglelefteq$

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

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 $y \lhd x' \trianglelefteq x$ 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

$$\sum_{x \in X^*} p_x \leq 0. \tag{9}$$

Because time-1 preferences are strict, it follows that if there are no sure direct preference reversals, 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 sure direct preference reversal 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 (9), 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 (8) 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 (8) on that set of alternatives. As this vector corresponds to a sure direct preference reversal this completes the proof. $\square$

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 sure direct preference reversal. 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 (8) 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 sure direct preference reversal 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 sure direct preference reversal 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 $x \trianglelefteq x^*$ cannot hold, and such that $x \sim_{\trianglelefteq} x'$ 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

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

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 (8) 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 sure direct preference reversal, 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 sure direct preference reversal.

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 (8), note that p' given by $p'_{\min X} = 0$ and $p'_{\max X} = 2p_{\max X}$ also satisfies (8) but belongs to $\mathbb{R}^{|X|-1}$. By the induction hypothesis, there must thus exist a sure direct preference reversal 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 (10), 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 (10), 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 (8) 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 (8) on that set of alternatives. As this vector corresponds to a sure direct preference reversal 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

$$0 \leq p \cdot (u^0 - \lambda \mathbf{1}_{x \in X^*}) = p \cdot u_0 - \lambda \sum_{x \in X^*} p_x \,. \tag{11}$$

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 $|X|$. $\square$

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 sure lottery preference reversal 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 sure lottery preference reversal.

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 (4)

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

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

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

which completes the proof. □

Online Appendix

Table of Contents

A 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 1 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.

B Mathematical Details for Section 1.1

Denote by $\mathrm{var}_0(\log \mathbb{E}_1[\theta_2]/\theta_1)$ 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*

$$\begin{aligned} \mathbb{E}_0[\theta_2] &= \mathbb{E}_0[\theta_1] \\ (\gamma-1)m &= \mathbb{E}_0\left[\log \frac{\mathbb{E}_1[\theta_2]}{\theta_1}\right] + \log(\beta) \\ (\gamma-1)^2 v &= var_0\left(\log \frac{\mathbb{E}_1[\theta_2]}{\theta_1}\right) . \end{aligned} \tag{12}$$

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

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

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 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 (12) 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 analyst's 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 ratio captures *both* the 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 (12) 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 (12) 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 (12) 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 (12) 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, (12) 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.

C 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 4 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

[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).

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 sure direct preference reversals 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 sure direct preference reversal" is, without the restriction to single-dimensional choice sets and single-peaked preferences, not sufficient to ensure the existence of an EU representation.

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 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

[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 (2020), Mahajan et al. (forthcoming).

et al. (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 5.2.[39]

[39] The literature on dynamic random utility (e.g. Fudenberg and Strzalecki, 2015, Frick et al., 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 et al. (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.