

For Online Publication

Web Appendix F. Formal Analysis of the Axioms in the Experiment

In this appendix, we formalize the axioms we used in the experiment, and we prove that they (together with completeness and continuity, which we also assume) are jointly equivalent to expected-utility maximization. Most of the work in the formalization is in writing down notation because we need to distinguish between a number of different ways of framing the decision problem, including node-wise versus pairwise choices, choices made with versus without the rest of the decision tree shown, and simple versus compound lotteries. Wherever possible, we draw on notation and definitions from the decision theory literature, as highlighted below. Although our experiment is conducted in the context of a particular decision tree (Figure 2), the formalization of the axioms and proof of expected utility in this appendix are much more general.

Our framework and results draw most heavily on Kreps and Porteus (1978), a classic paper on dynamic choice under uncertainty. Relative to Kreps and Porteus, we extend the framework to distinguish between simple and compound lotteries (as in Segal 1990) and other framing differences (as in Salant and Rubinstein 2008; Bernheim and Rangel 2009), and we modify notation a bit to facilitate these extensions.

Our axiomatization is also related to previous work linking the Independence Axiom to axioms of dynamic choice. At the end of this appendix, we relate our axiomatization to one particular, well-known such axiomatization (Karni and Schmeidler 1991, as extended by Volij 1994).

Here are the elements of our framework, together with how those elements correspond to the decisions in our experiment:

- **Time** is a finite and discrete sequence of periods during which decisions may be made: $t = 0, 1, \dots, T$. In the experiment, there are two time periods: period 0 is age 35, and period 1 is age 50. (Although Figure 2 depicts the payoffs as occurring at age 65 and onward, we do not need to be so specific for modeling purposes; the payoffs are modeled as occurring after the period-1 choices.)
- A **payoff** is a real number z_t . The set of possible payoffs in period t is denoted $Z_t \subset \mathbb{R}$. Payoffs occur after an action is taken in period t (actions are defined below). In the master decision tree shown in Figure 2, $Z_0 = \emptyset$ because all of the payoffs are realized after age 50 (not at age 35), and $Z_1 = \{52000, 72000, 100000, 108000, 150000, 225000\}$.

Kreps and Porteus implicitly assume the Reduction of Compound Lotteries Axiom within a period, and as a result, they consider only **simple lotteries** $\Delta(Z_t)$. In the spirit of Segal (1990), we explicitly distinguish between simple lotteries and **compound lotteries**, which we denote by $\bigcup_{m=2}^{M'} \Delta^m(Z_t)$, where $M' = MT(T - 1)$ and M is a finite natural number. In our experiment, Figure 1B depicts a simple lottery, while Figure 1A depicts a compound lottery.

- Decisions are made at **decision nodes**. The set of decision nodes in period t is denoted N_t . (The nodes are linked to each other in a well-defined decision tree, defined below.)
- In each period, the individual chooses an **action** d_t . In period T , the action is a simple or compound lottery over payoffs: $D_T \subset \cup_{m=1}^M \Delta^m(Z_T)$. A **decision** is a set of possible actions available at a node. The decision at time t at node $n \in N_t$ is denoted $x_{t,n} \in X_t \subset 2^{D_t}$. In periods $t < T$, the action may involve both a choice over simple or compound payoff lotteries and a choice over simple or compound lotteries about which decision to face in the next period: $D_t \subset \cup_{m=1}^M \Delta^m(Z_t) \times \cup_{m=1}^M \Delta^m(X_{t+1})$. For example, in Figure 2, there is a single node at time 0, and the decision is {the degenerate lottery yielding 100000 in period 1, a 50-50 lottery over different decisions in period 1}. For expositional purposes, we label the first of these A and the second B. If B is chosen, then the individual ends up at either of two nodes at time 1. At each, the decision is between two simple lotteries (C vs. D at one node and E vs. F at the other).
- A **decision tree**, which we denote by $\Gamma = (N_t)_t$, describes for each period t a set of nodes $N_t \subset X_t$. Each node $n \in N_t$ is in fact a decision, and we will treat nodes as decisions throughout (and abuse notation accordingly). A decision tree must originate at a single node: $|N_0| = 1$. In addition, for each node $n \in N_t$ at period $t > 0$, then there must be a node $m \in N_{t-1}$ in the previous period that allows for an action that leads to n with positive probability: for any $n \in N_t$ in period $t > 0$, there exists $m \in N_{t-1}$ with $d \in m$ such that $d(n) > 0$.
- The **history** h at period t describes the set of decisions faced and actions taken up to that period (or equivalently, the sequence of nodes and the actions linking the nodes). The set of possible histories at period 0 is just the set of possible decisions faced in period 0, $H_0 := X_0$ (this set will contain more than one element if more than one decision tree is possible). In subsequent periods $t = 1, 2, \dots, T$, the set of possible histories is defined recursively as $H_t := H_{t-1} \times (D_{t-1} \times X_t)$. Note that both decisions and actions are included in the history because it is possible (although not true in any of the decision trees in our experiment) that more than one sequence of actions in earlier periods could lead to the same decision in period t . We define the set of possible histories at “period $T + 1$ ” in order to capture the action chosen in period T , $H_{T+1} := H_T \times D_T$. The set of all histories is denoted $H = \cup_{s=0}^{T+1} H_s$.
Because a history is just a sequence of decisions and actions, it can describe a sequence that is not feasible in some particular decision tree (for this reason, the decision tree is not specified in many of the axioms below; the axioms impose restrictions on behavior across possible decision trees). For a given decision tree Γ , we will denote by h_0 the initial history, which is precisely the initial decision faced by the decision maker in decision tree Γ . Conditional on a particular decision tree, a history defines a “subtree,” the decision tree that begins from that point.
In some cases, after an action is taken, more than one subsequent node has positive probability. For example, in Figure 2, after choosing B, there is a 50% chance of facing the C vs. D decision and a 50% chance of facing the E vs. F decision. We describe the history at that “chance node” (which is not technically a node according to our definition) as including the decision faced in period 0 (A vs. B) and the action taken in period 0 but not the decision faced in period 1.

- A **complete contingent action plan** is a function $p: N_t \rightarrow D_t$ that prescribes, for every period t and every node $n \in N_t$, a feasible decision $p(n) \in n$. We denote the set of possible complete contingent action plans in decision tree Γ by $\mathcal{P}(\Gamma)$, and we denote the set of possible complete contingent action plans in decision tree Γ starting at history $h \in H_t$ by $\mathcal{P}_h(\Gamma)$. (Kreps and Porteus do not define complete contingent action plans, but we will need them to formalize our frames and axioms.)

- The **temporal compound lottery** associated with complete contingent action plan p for the decision tree $\Gamma = (N_t)_t$, denoted $\Gamma^*(p)$, is a decision tree $(M_t)_t$ such that for all periods t and all nodes $m \in M_t$, there exists $n \in N_t$ such that $m = \{p(n)\}$. In words, the temporal compound lottery $\Gamma^*(p)$ is the probability tree (a trivial decision tree) derived by taking the original decision tree Γ and retaining only the single action at each node chosen by the complete contingent action plan p . (Kreps and Porteus call $\Gamma^*(p)$ a “temporal lottery,” but we include the word “compound” because, unlike Kreps and Porteus, our setup does not presuppose the Reduction of Compound Lotteries Axiom.) We denote the set of all temporal compound lotteries for decision tree Γ by $\mathcal{G}(\Gamma) = \Gamma^*(\mathcal{P}(\Gamma))$.

- Denote by $\mathcal{G}(\Gamma, t)$ the set of temporal compound lotteries for decision tree Γ whose uncertainty is completely realized by period t . That is, after period t all actions are degenerate probability distributions. The **temporal compound lottery resolving in period t** associated with complete contingent action plan p for the decision tree $\Gamma = (N_t)_t$ with $\Gamma^*(p) \in \mathcal{G}(\Gamma, t)$, which we denote by $\Gamma^t(p)$, is a decision tree $(M_t)_t$ such that (i) for all t , $|M_t| = 1$, (ii) for all $s \neq t$ and all $m \in M_s$, $|m| = 1$, and (iii) $m \in M_t$ is a compound lottery with the same uncertainty resolution and payoffs as $\Gamma^*(p)$. In words, the temporal compound lottery resolving in period t , $\Gamma^t(p)$, is derived by taking a temporal compound lottery that has all of its uncertainty realized in period t or earlier and putting all the uncertainty (and its resolution) in period t . Before and after period t , the temporal compound lottery resolving in period t is a degenerate probability tree, where each node has just one possible action, which leads to the next node. (Kreps and Porteus define a concept with a similar name, a “temporal lottery resolving in period t .” We define their concept below (an element of the set $\mathcal{G}_t(\Gamma, h)$, also defined below) without naming it, prior to stating Axiom KP5.1.)

- The **temporal simple lottery resolving in period t** associated with complete contingent action plan p for the decision tree $\Gamma = (N_t)_t$ with $\Gamma^*(p) \in \mathcal{G}(\Gamma, t)$, denoted $\Gamma_{simple}^t(p)$, is the temporal compound lottery resolving in period t with the compound lottery replaced by its reduced simple lottery. (Since Kreps and Porteus implicitly assume the Reduction of Compound Lotteries Axiom in their setup, they do not define temporal simple lotteries separately from temporal compound lotteries.)

We need to go beyond Kreps and Porteus’s framework in order to capture frames where the individual is restricted to thinking about some part of the decision tree (as in Figures 3A and 3C) or where the individual is making pairwise choices between complete contingent action plans (as in Figures 2 and 3D) rather than making nodewise decisions. As far as we know, there is no existing decision-theoretic framework for describing these frames. Therefore, following

Salant and Rubinstein (2008) and Bernheim and Rangel (2009), we adopt the modeling trick of simply using a number ϕ to index what we call **context frames**:

No context ($\phi = 0$): The individual chooses an action while restricted to thinking about the current decision and its future consequences, without any knowledge of the history up to the current decision (or paths foregone). The frame “Single Action in Isolation” is the only $\phi = 0$ frame in the experiment.

Full context, nodewise decisions ($\phi = 1$): The individual is aware of the entire decision tree and chooses an action at each node. The $\phi = 1$ frames in the experiment are “Single Action with Backdrop,” “Two Contingent Actions with Backdrop,” and “Complete Contingent Action Plan.”

Full context, pairwise decisions ($\phi = 2$): The individual is aware of the entire decision tree and makes pairwise decisions between complete contingent action plans. The $\phi = 2$ frames in the experiment are “Pairwise Choices Between Complete Strategies,” “Pairwise Choices Between Compound Lotteries,” and “Pairwise Choices Between Reduced Simple Lotteries.”

To allow for the possibility that the individual violates the axioms we study in the experiment, we define the individual’s **preference relation** \succsim_h^ϕ as a function of both the history $h \in H_t$ (as in Kreps and Porteus) at some period t and also the context frame ϕ . The relations \succ_h^ϕ and \sim_h^ϕ are derived from \succsim_h^ϕ in the usual way. These preference relations are meant to capture revealed preferences; for example, the choices made in our experiment.

In the experiment, the **nodewise frames** require choices of an action at one or more nodes, whereas **pairwise frames** require a choice between two complete contingent action plans. Correspondingly, we allow for preference relations that rank actions (context frames 0 and 1) *and* preference relations that rank complete contingent action plans (context frames 1 and 2); our framework does not require that these two types of preference relations agree. Since it will be clear from context which preference relation is being described, we use the same notation \succsim_h^ϕ for both.

Although actions are defined above as probability distributions over payoffs and subsequent decisions, we will omit writing out the probability distribution when the action yields some subsequent decision and/or payoff with certainty. For example, in the statement of Axiom KP3.1 below, “ $(z, x) \succsim_h^1 (z, x')$ ” means that the action yielding payoff z and subsequent decision x with certainty is preferred to the action yielding payoff z and subsequent decision x' with certainty.

Axiom 0 is a preliminary assumption.

Axiom 0 (Completeness, Transitivity, and Continuity): For both preference relations over actions and over complete contingent action plans, and for all $t, h \in H_t$, and ϕ the preference relation \succsim_h^ϕ is complete, transitive, and continuous.

While we test transitivity directly in our experiment, we assume completeness and continuity.

Axiom 1 states that the individual has the same preferences in the frame “Single Action in Isolation” (Figure 3A) as in the frame “Single Action with Backdrop” (Figure 3B). Since the

former is the context frame $\phi = 0$ and the latter is $\phi = 1$, the axiom simply says that preferences over actions are same across these two context frames.

Axiom 1 (Irrelevance of Background Counterfactuals): Let $h \in H_t$ with $d, d' \in D_t$. Then $d \succ_h^0 d'$ if and only if $d \succ_h^1 d'$.

Axiom 2 states that the individual has the same preferences over actions in the frame “Single Action with Backdrop” (Figure 3B) as in the frame “Two Contingent Actions with Backdrop” (Figure 3C). To formalize the latter frame in our framework, we think of the individual as making the C vs. D decision at the “chance node,” where it is still uncertain whether the individual will face the C vs. D decision or the E vs. F decision.

Axiom 2 (Simple Actions = State-Contingent Actions): Let $h = (h_{t-1}, d_{t-1}, x_t) \in H_t$ such that $d_{t-1}(x_t) \in (0,1)$, and $d, d' \in x_t$. Then $d \succ_h^1 d'$ if and only if $d \succ_{(h_{t-1}, d_{t-1})}^1 d'$.

In Axiom 2, the history h occurs at the decision x_t where the individual is choosing between d and d' , while the history (h_{t-1}, d_{t-1}) refers to the “chance node” at which x_t is just one of the possible decisions the individual may face.

Axiom 3 states that the individual has the same preferences in the frame “Two Contingent Actions with Backdrop” (Figure 3C) as in the frame “Complete Contingent Action Plan” (Figure 1). As described above, the former frame is modeled as the individual making the C vs. D decision before knowing if she will face the C vs. D or the E vs. F decision. We model the latter frame as the individual choosing a complete contingent action plan.

Axiom 3 (Irrelevance of Counterfactual Choices): Let $h \in H_t$, $h' = (h, d_t, x_{t+1}) \in H_{t+1}$, and $d, d' \in x_{t+1}$. Then $d \succ_{(h, d_t)}^1 d'$ if and only if $p \succ_h^1 p'$ for any $p, p' \in \mathcal{P}_h(\Gamma)$ with $p = p'$ except $p(h') = d$ and $p'(h') = d'$.

To apply the axiom to the frames in our experiment, we impose $t = 0$ so that the complete contingent action plans p and p' are chosen in period 0, while the choice between actions d and d' is made at the “chance node” before the period-1 decision is known (as in the discussion of Axiom 2 above).

As we highlight in the paper (and in the name “Irrelevance of Counterfactual Choices”), Axiom 3 rules out anticipated regret and counterfactual reference points as influences on choices. In our experiment, when an individual chooses a complete contingent action plan (in the frame “Complete Contingent Action Plan”), the individual’s choice for C vs. D and E vs. F is in the context of having foregone the payoff from A. When the individual makes the C vs. D and E vs. F choices separately (in the frame “Two Contingent Actions with Backdrop”), it is less salient that the individual must have foregone the payoff from A to end up at those nodes. The axiom says that the choices in the two frames must agree.

Axiom 4 states that the individual has the same preferences in the frame “Complete Contingent Action Plan” (Figure 1) as in the frame “Pairwise Choices Between Complete Strategies” (Figure 3D).

Axiom 4 (Shift from Nodewise to Pairwise Choices): Let $h \in H_t$ with $p, p' \in \mathcal{P}_h(\Gamma)$. Then $p \succ_h^1 p'$ if and only if $p \succ_h^2 p'$.

This axiom simply says that individuals have the same preferences over complete contingent action plans when choosing the actions nodewise as when making pairwise choices between complete contingent action plans.

Axiom 5, which we call in the paper “Complete Strategies = Implied Lotteries,” has two subcomponents, which we separately write out here as Axioms 5A and 5B. These two are conceptually distinct. We combined them in the experiment because the experiment was already quite long and because, when looking at drafts of decision screens where they were separated, we worried that in reconsiderations involving just Axiom 5B, the differences in the timing of the resolution of uncertainty may have been too subtle for experimental participants to notice. As a result, the frame “Pairwise Choices Between Complete Strategies” (Figure 3E) differs from the frame “Pairwise Choices Between Compound Lotteries” (Figure 1A) in two ways: the latter replaces all decision nodes with the choice made according to some complete contingent action plan, and it pushes all the uncertainty from age 35 to age 50. The first of these changes does not matter according to Axiom 5A, and the second change does not matter according to Axiom 5B.

Axiom 5A (Complete Strategies = Implied Temporal Compound Lotteries): Let $h \in H_t$ with $p, p' \in \mathcal{P}_h(\Gamma)$. Then $p \succsim_h^2 p'$ if and only if $\Gamma^*(p) \succsim_h^2 \Gamma^*(p')$.

Axiom 5A says that an individual’s pairwise choices over complete contingent action plans coincide with their pairwise choices over the temporal compound lotteries derived from those plans.

Axiom 5B (Indifference to the Timing of the Resolution of Uncertainty in Pairwise Choices): Let $h \in H_t$ with $p, p' \in \mathcal{P}_h(\Gamma)$ and $\Gamma^*(p), \Gamma^*(p') \in \mathcal{G}(\Gamma, s)$ for some $s \geq t$. Then $\Gamma^*(p) \succsim_h^2 \Gamma^*(p')$ if and only if $\Gamma^s(p) \succsim_h^2 \Gamma^s(p')$.

Axiom 5B says that an individual’s pairwise choices over the temporal compound lotteries do not depend on whether the uncertainty occurs (and is realized) at different periods in the probability tree or all in a future period $s \geq t$ after all of the uncertainty would have otherwise been realized. (Axiom 5B is related to an axiom studied by Kreps and Porteus, KP5.1, which we discuss and compare to Axiom 5B below.)

Axiom 6 is a Reduction of Compound Lotteries Axiom in which all decisions are made in the current period, but all uncertainty and payoffs are realized in a future period $s \geq t$.

Axiom 6 (Reduction of Compound Lotteries): Let $h \in H_t$ with $p, p' \in \mathcal{P}_h(\Gamma)$ and $\Gamma^*(p), \Gamma^*(p') \in \mathcal{G}(\Gamma, s)$ for some $s \geq t$. Then $\Gamma^s(p) \succsim_h^2 \Gamma^s(p')$ if and only if $\Gamma_{simple}^s(p) \succsim_h^2 \Gamma_{simple}^s(p')$.

In the remainder of this appendix, we will show that Axioms 0-6 are necessary and sufficient for an individual’s preferences to be represented by a von Neumann-Morgenstern utility function that is the same in each context frame $\phi \in \{0,1,2\}$. Clearly, if an individual’s preferences can be represented by a von Neumann-Morgenstern utility function that is the same in each context frame, then Axioms 0-6 are satisfied. We therefore focus on proving sufficiency of Axioms 0-6.

To prove sufficiency, we will use a result from Kreps and Porteus, which relies on three axioms we re-state here using our notation.

Axiom KP2.3 (Substitution): For all $t, h \in H_t$, and $d, d' \in D_t$ such that $d \succsim_h^1 d'$, $\alpha d + (1 - \alpha)d'' \succsim_h^1 \alpha d' + (1 - \alpha)d''$ for all $\alpha \in (0,1)$ and $d'' \in D_t$.

This first axiom, their ‘‘Substitution Axiom,’’ is a version of the Independence Axiom defined for decision trees (as opposed to single, static decisions as in the traditional formulation) that holds separately for each time period and history.

The second of their axioms, their Axiom 3.1, states that if a payoff today and a decision tomorrow are preferred to the same payoff today and an alternative decision tomorrow, then once the payoff is realized, the preference order over decisions is maintained.

Axiom KP3.1 (Temporal Consistency): For all $t, h \in H_t$, $z \in Z_t$, and $x, x' \in X_{t+1}$, $(z, x) \succsim_h^1 (z, x')$ if and only if for all $d' \in x'$, there exists $d \in x$ such that $d \succsim_{(h, x \cup x')}^1 d'$.

In our statement of the axiom, $(h, x \cup x')$ refers to a history where the payoff (z) has occurred but either decision x or x' may still occur.

The third of their axioms, their Axiom 5.1, is an axiom about indifference to the timing of the resolution of uncertainty. It states that if two lotteries have the same distribution of payoffs, then it does not matter to the individual when the uncertainty is realized. To be specific, suppose there are no payoffs or resolutions of any uncertainty before period t , and the individual may receive either of two temporal compound lotteries whose uncertainty is realized beginning in period t . Then the individual is indifferent to being told in period t versus in period $t - 1$ which of the two lotteries they receive. To state the axiom formally, define $\mathcal{G}_t(\Gamma, h)$ to be the set of all temporal compound lotteries derived from complete contingent action plans in decision tree Γ such that no uncertainty is resolved before period t with identical histories up to $h \in H_t$. Also, for $t \geq 1$ and $k \leq t$, $h = (h_k, h_{k \rightarrow t}) \in H_t$, $q, q' \in \mathcal{G}_t(\Gamma, h)$, and $\alpha \in [0,1]$, define $(k, \alpha; q, q') \in \mathcal{G}_t(\Gamma, h_k)$ to be the temporal compound lottery formed by α -mixing q and q' in period k .

Axiom KP5.1 (Indifference to the Timing of the Resolution of Uncertainty in Nodewise Choices): For all $t \geq 1$, $h \in H_t$, $\alpha \in [0,1]$ and $q, q' \in \mathcal{G}_t(\Gamma, h)$, $(t, \alpha; q, q') \sim_{h_0}^1 (t - 1, \alpha; q, q')$.

To aid in understanding the axiom, Figure A1 illustrates the two objects, $(t, \alpha; q, q')$ and $(t - 1, \alpha; q, q')$, for a simple example where $T = 2$, $t = 1$, $\alpha = 0.3$, and q and q' are 50-50 gambles with different payoffs. Panels (a) and (b) show q and q' , and panels (c) and (d) show $(t, \alpha; q, q')$ and $(t - 1, \alpha; q, q')$.

This axiom differs from our Axiom 5B in two main ways. First, we interpret Axiom KP5.1 to be an assumption about what we call nodewise choices (context frame $\phi = 1$), whereas Axiom 5B is an assumption about pairwise choices (context frame $\phi = 2$). Second, Axiom KP5.1 is a restriction on choices between realizing uncertainty in any period t versus in period $t - 1$, whereas Axiom 5B is a restriction on choices between realizing uncertainty in any combination of periods versus all in a future period $s \geq t$ after all of the uncertainty would have otherwise been realized. However, when put together with some of the other axioms, satisfying

either axiom implies that the individual is indifferent between realizing uncertainty in any two periods.

The key result from Kreps and Porteus, their Corollary 3, is the following (again, stated in our notation):

Theorem 1 (KP Corollary 3). Axioms 0, 6, KP2.3, KP3.1, and KP5.1 are necessary and sufficient for the individual's choices in context frame $\phi = 1$ to be represented by a von Neumann-Morgenstern utility function.

Our statement of the result differs slightly from that in Kreps and Porteus for three reasons. First, Kreps and Porteus define two axioms (their axioms 2.1 and 2.2) which jointly have the same content as our Axiom 0. Second, their framework assumes that lotteries within any period are simple lotteries (i.e., they implicitly assume the Reduction of Compound Lotteries Axiom within each period), while our framework explicitly allows for compound lotteries. For that reason, our statement of the result assumes Axiom 6 so that all compound lotteries can be replaced by their reduced simple lotteries. Third, their setup does not distinguish between the context frames $\phi \in \{0,1,2\}$, while ours does, and we interpret axioms KP3.1 and KP5.1 as applying to context frame $\phi = 1$. Thus, we state their conclusion as also applying to context frame $\phi = 1$.

Given this result, our key result is:

Theorem 2. Suppose Axioms 0 and 2-6 hold. Then Axioms KP2.3, KP3.1, and KP5.1 are also satisfied.

Proof. We first show that Axiom KP5.1 holds. Fix $t \geq 1$, $h = (h_{t-1}, h_{t-1 \rightarrow t}) \in H_t$, $\alpha \in [0,1]$ and $q, q' \in \mathcal{G}_t(\Gamma, h)$. Define $q_t \equiv (t, \alpha; q, q') \in \mathcal{G}_t(\Gamma, h) \subset \mathcal{G}_{t-1}(\Gamma, h_{t-1})$ and $q_{t-1} \equiv (t-1, \alpha; q, q') \in \mathcal{G}_{t-1}(\Gamma, h_{t-1})$, two α -mixtures of q and q' that differ only in whether the individual is told which of the two lotteries they receive. Note that $q_t, q_{t-1} \in \mathcal{G}(\Gamma, T)$. Since q_t and q_{t-1} differ only in the timing of the resolution of uncertainty, and since $\Gamma^T(q_t)$ and $\Gamma^T(q_{t-1})$ push this timing to period T , $\Gamma^T(q_t) = \Gamma^T(q_{t-1})$. That is, it must be that $\Gamma^T(q_t) \sim_h^2 \Gamma^T(q_{t-1})$ for all $h \in H$. In particular, indifference holds for h_0 . Applying Axioms 5B, 5A, and 4,

$$\Gamma^T(q_t) \sim_{h_0}^2 \Gamma^T(q_{t-1}) \stackrel{A5B}{\Leftrightarrow} \Gamma^T(q_t) \sim_{h_0}^1 \Gamma^T(q_{t-1}) \stackrel{A5A}{\Leftrightarrow} q_t \sim_{h_0}^2 q_{t-1} \stackrel{A4}{\Leftrightarrow} q_t \sim_{h_0}^1 q_{t-1}.$$

Thus Axiom KP5.1 holds.

We next show that Axiom KP3.1 holds. Fix $t \geq 1$, $h = (h_-, n) \in H_t$, and a node n in period t that gives payoff z and two different decisions, x and y : $n = \{(x, z), (y, z)\}$. From Axiom 0, the preferences \succsim_h^1 are complete and continuous, so there is a preference-maximizing element of y . Let $p, p' \in \mathcal{P}_h(\Gamma)$ such that $p(n) = (x, z)$, $p'(n) = (y, z)$, $p(n') = p'(n')$ for all $n' \neq n$, and conditional on facing the decision y , both p and p' select the preference-maximizing element of y . Applying Axioms 3-5A,

$$(x, z) \succsim_h^1 (y, z) \stackrel{A3}{\Leftrightarrow} p \succsim_{h_-}^1 p' \stackrel{A4}{\Leftrightarrow} p \succsim_{h_-}^2 p' \stackrel{A5A}{\Leftrightarrow} \Gamma^*(p) \succsim_h^2 \Gamma^*(p').$$

Define Γ' to be identical to Γ except that instead of giving a choice between x and y at node n , the choice is pushed back to the following period: $n = \{(x \cup y, z)\}$. Define $q \in \mathcal{P}_h(\Gamma')$ to be the same as p except that it makes the degenerate choice of $(x \cup y, z)$ at node n and then chooses x in the following period: at node m ,

$$q(m) = \begin{cases} (x \cup y, z) & \text{if } m = n \\ p(x) & \text{if } m = x \cup y. \\ p(m) & \text{otherwise} \end{cases}$$

Define $q' \in \mathcal{P}_h(\Gamma')$ analogously for p' . By construction, $\Gamma'^*(q) = \Gamma^*(p)$ and $\Gamma'^*(q') = \Gamma^*(p')$, so $\Gamma^*(p) \succsim_h^2 \Gamma^*(p') \Leftrightarrow \Gamma'^*(q) \succsim_h^2 \Gamma'^*(q')$. Applying Axioms 3-5A again,

$$\Gamma'^*(q) \succsim_h^2 \Gamma'^*(q') \stackrel{A5A}{\Leftrightarrow} q \succsim_h^2 q' \stackrel{A4}{\Leftrightarrow} q \succsim_h^1 q' \stackrel{A3}{\Leftrightarrow} p(x) \succsim_{(h,x \cup x')}^1 p'(y).$$

Putting all of this together, we have shown that $(x, z) \succsim_h^1 (y, z) \Leftrightarrow p(x) \succsim_{(h,x \cup x')}^1 p'(y)$. By construction, $p(x) \in x$ and p' selects the maximal element of y , so for every $d' \in y$, $p(x) \succsim_{(h,x \cup x')}^1 p'(y) \succsim_{(h,x \cup x')}^1 d'$. Hence Axiom KP3.1 holds.

Finally, we show that Axiom KP2.3 is satisfied. Fix $t \geq 1$, $h \in H_{t-1}$, $h' = (h, d_{t-1}, x_t) \in H_t$, and $d, d' \in x_t$ such that $d \succsim_{h'}^1 d'$. Let $\alpha \in (0,1)$ and $d'' \in D_t$. Define Γ^1 to be identical to Γ except at history h there is a unique action available, denoted d^1 , that leads to x_t with probability α and $\{d''\}$ with probability $1 - \alpha$. Since at x_t the reduced simple lotteries implied by any complete contingent action plans are identical, it follows from Axioms 2-6 that $d \succsim_{(h,d^1,x_t)}^1 d'$. Applying Axioms 2 and 3,

$$d \succsim_{h'}^1 d' \stackrel{A2}{\Leftrightarrow} d \succsim_{(h,d^1)}^1 d' \stackrel{A3}{\Leftrightarrow} p \succsim_h^1 p'$$

for all $p, p' \in \mathcal{P}_h(\Gamma^1)$ with $p = p'$ except $p(x_t) = d$ and $p'(x_t) = d'$. Applying Axioms 4-6, $\Gamma_{simple}^{1T}(p) \succsim_h^2 \Gamma_{simple}^{1T}(p')$.

Now define Γ^2 to be identical to Γ except that at history h' , x_t is replaced by $x' \equiv \{\alpha d + (1 - \alpha)d'', \alpha d' + (1 - \alpha)d''\}$. Define $q, q' \in \mathcal{P}_h(\Gamma^2)$ by $q(x') = \alpha d + (1 - \alpha)d''$ and $q'(x') = \alpha d' + (1 - \alpha)d''$ and $q = p$ at all other decisions. By construction, $\Gamma_{simple}^{1T}(p) = \Gamma_{simple}^{2T}(q)$ and $\Gamma_{simple}^{1T}(p') = \Gamma_{simple}^{2T}(q')$. Thus $\Gamma_{simple}^{2T}(q) \succsim_h^2 \Gamma_{simple}^{2T}(q')$. Applying Axioms 2-6 in reverse we have $\alpha d + (1 - \alpha)d'' \succsim_{h'}^1 \alpha d' + (1 - \alpha)d''$. Thus Axiom KP2.3 is satisfied. ■

Given Theorem 2, we need Axioms 1 and 4 in order to extend the expected-utility representation to context frames $\phi = 0$ and $\phi = 2$.

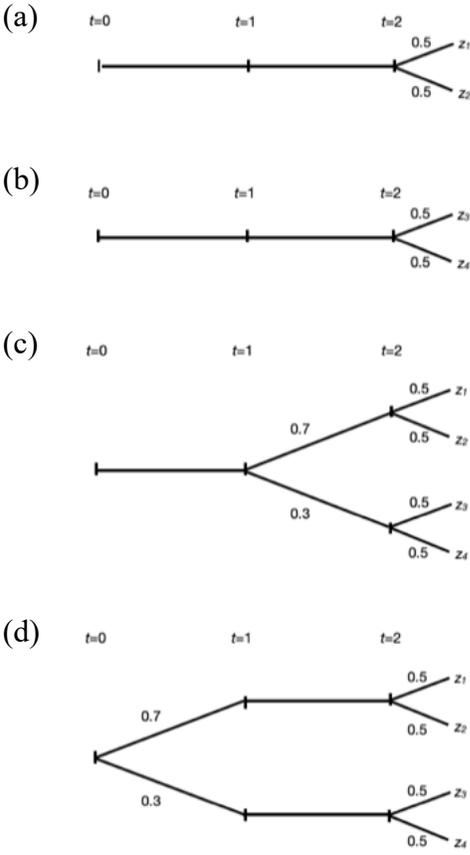
Corollary 1. Axioms 0-6 hold if and only if the individual's preference is represented by a single von Neumann-Morgenstern utility function that is the same in each context frame $\phi \in \{0,1,2\}$.

Proof. Theorems 2 and 1 together imply that the individual's choices in context frame $\phi = 1$ can be represented by a von Neumann-Morgenstern utility function. Axioms 1 and 4 then imply that the same von Neumann-Morgenstern utility function represents the individual's choices in context frames $\phi = 0$ and $\phi = 2$. ■

We end this appendix by relating our axiomatization to a well-known previous axiomatization linking the Independence Axiom to axioms of dynamic choice. Specifically, Karni and Schmeidler (1991) and Volij (1994), taken together, showed that if any two of

consequentialism (roughly: choices do not depend on the past), dynamic consistency (roughly: currently optimal behavior equals planned behavior), and Reduction of Compound Lotteries is assumed, then the third is equivalent to the Independence Axiom. Their setup is in terms of compound lotteries rather than decision trees, so theirs has no notion of making a choice at a “chance node.” As a result, our Axioms 2 and 3 are implicit in their setup. They assume a version of our Axiom 0. Naturally, their Reduction of Compound Lotteries Axiom corresponds to our Axiom 6. Although there is no exact correspondence, their other axioms roughly correspond to our other axioms (although some of the content of our other axioms is implicit in their setup): their dynamic consistency axiom captures a related intuition as our Axiom 4 (Shift from Nodewise to Pairwise Choices) but also captures some of our Axioms 5A (Complete Strategies = Implied Temporal Compound Lotteries) and 5B (Indifference to the Timing of the Resolution of Uncertainty in Pairwise Choices); and their consequentialism axiom captures both our Axioms 1 (Irrelevance of Background Counterfactuals) and 3 (Irrelevance of Counterfactual Choices).

Figure A1. Illustration of Axiom KP5.1



References

- Bernheim, B. Douglas, and Antonio Rangel.** 2009. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics* 124 (1): 51–104.
- Kreps, David M., and Evan L. Porteus.** 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory." *Econometrica* 46 (1): 185.
- Salant, Yuval, and Ariel Rubinstein.** 2008. "(A, f): Choice with Frames." *Review of Economic Studies* 75 (4): 1287–96.
- Segal, Uzi.** 1990. "Two-Stage Lotteries without the Reduction Axiom." *Econometrica* 58 (2): 349–77.