

Dynamic Choice under Ambiguity

Marciano Siniscalchi*

October 28, 2010

Abstract

This paper analyzes dynamic choice for decision makers whose preferences violate Savage's Sure-Thing principle [40], and therefore give rise to violations of dynamic consistency. The consistent-planning approach introduced by Strotz [46] provides one way to deal with dynamic inconsistencies; however, consistent planning is typically interpreted as a solution concept for a game played by “multiple selves” of the same individual.

The main result of this paper shows that consistent planning under uncertainty is fully characterized by suitable behavioral assumptions on the individual's preferences over decision trees. In particular, knowledge of *ex-ante* preferences over trees is sufficient to characterize the behavior of a consistent planner. The results thus enable a fully decision-theoretic analysis of dynamic choice with dynamically inconsistent preferences.

The analysis accommodates arbitrary decision models and updating rules; in particular, no restriction need be imposed on risk attitudes and sensitivity to ambiguity.

1 Introduction

In a dynamic-choice problem under uncertainty, a decision maker (DM henceforth) acquires information gradually over time, and takes actions in multiple periods and information scenar-

*Economics Department, Northwestern University, Evanston, IL 60208; marciano@northwestern.edu. I have greatly benefited from many extensive conversations with Peter Klibanoff and Nabil al-Najjar. I also thank the Coeditor, Bart Lipman, and two anonymous referees, as well as Pierpaolo Battigalli, Eddie Dekel, Larry Epstein, Paolo Ghirardato, Alessandro Lizzeri, Fabio Maccheroni, Alessandro Pavan and Larry Samuelson for their valuable insights, and Simone Galperti for excellent research assistance. All errors are my own.

ios. The basic formulation of expected utility (EU) theory instead concerns a “reduced-form,” atemporal environment, wherein preferences are defined over maps from a state space Ω to a set of prizes X (“acts”). Thus, in order to analyze dynamic choice problems, it is necessary to augment the atemporal EU theory with assumptions about the individual’s preferences at different decision points. The standard assumption is of course *Bayesian updating*: if the individual’s initial beliefs are characterized by the probability q , her beliefs at any subsequent decision point h are given by the conditional probability $q(\cdot|B)$, where the event B represents the information available to the individual at h . Together with the assumption that the individual’s risk preferences do not change, Bayesian updating ensures that the DM’s behavior satisfies a crucial property, *dynamic consistency* (DC): the course of action that the individual deems optimal at a given decision point h , on the basis of the preferences she holds at h , is also optimal when evaluated from the perspective of any earlier decision point h' (and conversely, if h is reached with positive probability starting from h'). This implies in particular that backward induction or dynamic programming can be applied to identify optimal plans of action.

Bayesian updating and the DC property are intimately related to the cornerstone of Savage’s axiomatization of EU, namely his Postulate P2; Section 2 discusses this tight connection and provides references. Sensitivity to *ambiguity* (Ellsberg [7]), or to the common-ratio or common consequence effects (Allais [2]; Starmer [45]), and other manifestations of non-EU risk attitudes, typically lead to violations of Savage’s Postulate P2. As a consequence, violations of DC are to be expected when such preferences are employed to analyze dynamic choice problems; again, Section 2 elaborates on this point, and provides illustrative examples. These violations of DC, and ways to address them, are the focus of the present paper.

Whenever a conflict arises among preferences at different decision points, additional assumptions are required to make clear-cut behavioral predictions (Epstein and Schneider [9], p. 7). One approach, introduced by Strotz [46] in the context of deterministic choice with changing time preferences and tastes, is to assume that the DM adopts the strategy of *Consistent Planning* (CP). In Strotz’s own words, at every decision point, a consistent planner chooses “the best plan among those that [s]he will actually follow” ([46], p. 173).

Formally, CP is a refinement of backward induction that incorporates a specific tie-breaking rule. Informally, CP reflects the intuitive notion that the DM is *sophisticated*: that is, she holds

correct “beliefs” about her own future choices. The problem with this intuitive notion is that, of course, “beliefs” about future choices cannot be observed directly; they also cannot be elicited on the basis of the DM’s initial and/or conditional preferences *over acts*.

The literature on time-inconsistent preferences circumvents this difficulty by suggesting that CP is best viewed as a solution concept for *a game played by “multiple selves” of the same individual*. Strotz himself ([46, p. 179]) explicitly writes that “[t]he individual over time is an infinity of individuals”; see also Karni and Safra [28, pp. 392-393]), O’Donoghue and Rabin [36, p. 106], and Piccione and Rubinstein [37, p. 17]. However, at the very least, this interpretation represents “a major departure from the standard economics conception of the individual as the unit of agency” (Gul and Pesendorfer [21, p. 30]). It certainly does not clarify what it means for an *individual decision-maker* to adopt the strategy of consistent planning. It reinforces the perception that a sound, behavioral analysis of multi-period choice *requires* some form of dynamic consistency (Epstein and Schneider [9, p. 2]). Finally, it provides very little guidance as regards policy analysis.

This paper addresses these issues by providing a *fully behavioral* analysis of CP in the context of dynamic choice under uncertainty. In the spirit of the menu-choice literature initiated by Kreps [31], I assume that the individual is characterized by a single, ex-ante preference relation *over dynamic choice problems*, modeled as decision trees. I then show that:

- under suitable assumptions, conditional preferences can be derived from ex-ante preferences *over trees*, regardless of whether or not preferences *over acts* satisfy Savage’s postulate P2 (cf. Sec. 4.2 and Theorem 2);
- sophistication can be formalized as a behavioral axiom on preferences *over trees*, regardless of whether or not DC holds (cf. Sec. 4.3.2); and
- the proposed sophistication axiom, plus auxiliary assumptions, provides a behavioral rationale for CP (Theorems 3 and 4), again regardless of whether or not P2 or DC hold.

Three features of the analysis in this paper deserve special emphasis. First, the approach in this paper is “fully behavioral” in the specific sense that the implications of CP are entirely reflected in the individual’s ex-ante preferences over trees, which are observable.

Second, by providing a formal definition of sophistication that does not involve “multiple selves,” this paper provides a way to interpret this intuitive notion as a behavioral principle—but one that applies to preferences over trees, rather than acts. The analysis also indicates that seemingly minor differences in the way sophistication is formalized can have significant consequences in the context of choice under uncertainty; see Sec. 5.2.

Third, minimal assumptions are required on preferences *over acts*: the substantive requirements considered in this paper are imposed on preferences *over trees*. In particular, postulate P2 and hence DC play no role in the analysis. This allows for prior and conditional preferences that exhibit a broad range of attitude toward risk and ambiguity—a main objective of the present paper.

The main results in this paper do not restrict attention to any specific model of choice, or “updating rule.” However, to exemplify the approach taken here, Theorem 5 specializes Theorem 3 to the case of multiple-priors preferences (Gilboa and Schmeidler, [15]) and prior-by-prior updating. Furthermore, Sec. 4.4.2 leverages the framework and results in this paper to address what is often cited as a “paradoxical” implication of CP (e.g. Machina [34], Epstein and Le Breton [8]): a time-inconsistent, but sophisticated DM may forego freely available information, if by doing so she also limits her future options. The analysis in §4.4.2 shows that this behavior actually has a simple rationalization if preferences over trees, rather than just acts, are taken into account.

Organization of the paper. Sec. 2 illustrates the key issues by means of examples. Sec. 3 introduces the required notation and terminology. Sec. 4 presents the main results, the special case of multiple-priors preferences, and the application to value-of-information problems. Sec. 5 discusses the main results. Sec. 6 discusses the important connections with the existing, rich literature on dynamic choice under ambiguity, as well as work on menu choice, intertemporal choice with changing tastes, and dynamic choice with non-expected utility preferences.

2 Heuristic treatment

Savage’s P2 and DC. It was asserted above that Bayesian updating and DC are intimately related to Savage’s postulate P2; this implies that failures of DC are not pathological, but rather

the norm, when non-EU preferences are employed to analyze problems of choice under uncertainty. Savage himself provides an argument along these lines in [40, §2.7]; Ghirardato [13] formally establishes the equivalence of DC and Bayesian Updating with P2, under suitable ancillary assumptions. Proposition 1 in the present paper provides a corresponding, slightly more general equivalence result in the framework adopted here.¹ These results can be illustrated in simple examples that are also useful to describe the proposed behavioral approach to CP.

An example. An urn contains 90 amber, blue and green balls; in the following, I shall consider different assumptions about what the DM knows regarding its composition. A single ball will be drawn; denote the corresponding state space by $\Omega = \{\alpha, \beta, \gamma\}$, in the obvious notation. At time 0, without knowing the prevailing state, the DM can choose a “safe” action s that yields a prize of $\frac{1}{2}$ if the ball is amber or blue, and $x \in \{0, 1\}$ otherwise; alternatively, the DM can choose to place a “contingent” bet c . In this case, the DM receives x if the ball is green, and can place a bet on amber (a) or blue (b) at time 1 otherwise. The situation is depicted in Fig. 1: solid circles denote decision points, and empty circles denote points where “Nature moves,” or more properly reveals information to the DM.




Figure 1: A dynamic decision problem; $x \in \{0, 1\}$.

Given the state space Ω and prize space $X = \{0, \frac{1}{2}, 1\}$, the atemporal choice environment corresponding to the decision problem under consideration consists of all acts (functions) $h \in X^\Omega$. Suppose first that the DM knows the composition of the urn, and that she has risk-neutral EU preferences; her beliefs $q \in \Delta(\Omega)$ reflect the composition of the urn. Thus, in the atemporal

¹ All versions of this argument incorporate the assumptions of *consequentialism* and (with the exception of Prop. 1 in this paper) *reduction*; the discussion of these substantive hypotheses is deferred till Sec. 3.3.

setting, the DM evaluates acts $h = (h_\alpha, h_\beta, h_\gamma) \in X^\Omega$ according to the functional $V(h) = E_q[h]$.

Now, as described above, augment this basic preference specification by assuming that the DM updates her beliefs q in the usual way. At the second decision node, she then conditionally (weakly) prefers a to b if and only if $q(\{\alpha\}|\{\alpha, \beta\}) \geq q(\{\beta\}|\{\alpha, \beta\})$. This is of course equivalent to $q(\{\alpha\}) \geq q(\{\beta\})$, which is the restriction on ex-ante belief that ensures that, from the point of view of the initial node, the course of action “ c then a ” is weakly preferred to “ c then b .” This is an instance of DC: the ex-ante and conditional rankings of the actions a and b coincide. In turn, this provides a rationale for the use of backward induction: the plans of action available to the DM at the first decision node are “ c then a ,” “ c then b ” and “ s ”; but one of the two “ c ” plans can be eliminated by first solving the choice problem at the second node. A simple calculation then shows that “ s ” is never strictly preferred, regardless of the ratio of blue vs. green balls. Hence, for instance, if $q(\{a\}) > q(\{b\})$, then “ c then a ” is the unique optimal plan.²

To provide a concrete illustration of the relationship between DC and P2, recall that the assumptions of ex-ante EU preferences and Bayesian updating delivered two conclusions: (i) the ranking of a vs. b at the second decision is the same as the ranking of “ c then a ” vs. “ c then b ” at the first decision node; furthermore, (ii) the ranking of a vs. b at the second node is independent of the value of x . Now assume that the modeler does *not* know that ex-ante preferences conform to EU, nor that conditional preferences are derived by Bayesian updating; however, he does know that (i) and (ii) hold. Clearly, the modeler is still able to conclude that the ranking of “ c then a ” and “ c then b ” must also be independent of x , so that

$$(1, 0, 0) \succcurlyeq (0, 1, 0) \iff (1, 0, 1) \succcurlyeq (0, 1, 1), \quad (1)$$

where \succcurlyeq denotes the DM’s preferences over acts. This is an implication of Savage’s Postulate P2 (cf. [40, p. 23], or Axiom 4.2 in §4.1 below). In other words, as claimed, Eq. (1) is also a *necessary condition for Dynamic Consistency in Fig. 1*.

² As per footnote 1, this argument incorporates the substantive assumptions of *Consequentialism* and *Reduction* (see Sec. 3.3). In the tree of Fig. 1, the relevant aspect of Consequentialism is the fact that the ranking of a vs. b at the second decision node is independent of the value of x ; Reduction instead implies that the choice of c followed by, say, a is evaluated by applying the functional $V(\cdot)$ to the associated mapping from states to prizes, i.e. $(1, 0, x)$. I maintain both assumptions in this Introduction; the formal results in the body of the paper allow for arbitrary departures from Reduction.

Ambiguity, DC and CP. I now describe ambiguity-sensitive preferences that violate P2, and hence yield a failure of DC; see below for an analogous example based on the common-consequence effect. Assume that, as in the three-color-urn version of the Ellsberg paradox [7], the DM is only told that the urn contains 30 amber balls. Assume that she initially holds multiple-priors (a.k.a. “maxmin-expected utility,” or MEU) preferences (Gilboa and Schmeidler [15]), is risk-neutral for simplicity, and updates her beliefs prior-by-prior (e.g. Jaffray [26], Pires [38]) upon learning that the ball drawn is not green. Formally, her preferences over acts $h \in \mathbb{R}^\Omega$ conditional on either $F = \Omega$ or $F = \{\alpha, \beta\}$, are given by $V_F(h) = \min_{q \in C} E_q[h|F]$, where C is the set of all probabilities q on Ω such that $q(\{\alpha\}) = \frac{1}{3}$. Notice that such conditional preferences are independent of the value of x , as is the case for Bayesian updates of EU preferences.

Note first that, a priori (i.e. conditional on $F = \Omega$), this DM exhibits the modal preferences reported by Ellsberg [7]: she prefers a bet on amber to a bet on blue, but she also prefers betting on blue *or* green rather than amber *or* green. Therefore, the DM’s preferences violate Eq. (1), hence Savage’s postulate P2. Furthermore, conditional on $\{\alpha, \beta\}$, this DM prefers $(1, 0, x)$ to $(0, 1, x)$ regardless of the value of x , and hence will strictly prefer a to b .

If now $x = 1$, DC is violated: at the first decision node, the DM strictly prefers the plan “ c followed by b ” to “ c followed by a ,” but at the second node she strictly prefers a to b .

To resolve these inconsistencies, suppose that the DM adopts CP. The intuitive assumption of sophistication implies that, at the first decision node, the DM should *correctly anticipate her future choice of a* , regardless of the value of x . This is true despite the fact that, for $x = 1$, she really would like to commit to choosing b instead. Hence, when contemplating the choices c and s at the first decision node, the DM understand that she is really comparing the plan “ c then a ” to “ s ”. For $x = 0$, she will strictly prefer the former; but, for $x = 1$, she will strictly prefer the latter. This logic thus delivers unambiguous and coherent behavioral predictions.

A dynamic “common consequence” paradox (cf. Allais [2]). Violations of DC can also arise when preferences are probabilistically sophisticated but not EU; again, CP provides a way to deal with them. Suppose that one ball will be drawn from an urn containing 100 balls, numbered 1 through 100. Fig. 2 depicts the choice problem and payoffs, where “ M ” denotes one million (dollars), and “1...11,” “12..., 100” etc. refer to the number on the ball drawn.

The DM’s beliefs are uniform on $\Omega = \{1, \dots, 100\}$ at the initial node, and determined via




Figure 2: A dynamic Allais-type problem; $x \in \{0, 1M\}$.

Bayes' rule at the second; her preferences are of the rank-dependent EU form (Quiggin [39]), with quadratic distortion function. If $x = 1M$, the plan “ c followed by b ” is preferred to “ c followed by a ,” whereas the opposite holds if $x = 0$: this corresponds to the usual violations of the Independence axiom, and hence of P2. Furthermore, the DM strictly prefers a to b at the second decision node if $x = 1M$, so preferences are dynamically inconsistent. Nevertheless, CP again delivers well-defined behavioral predictions: if $x = 1M$, the DM will correctly anticipate choosing a at the second node, and hence, by a simple calculation, opt for s at the initial node.

Karni and Safra [27, 28] illustrate applications of CP with non-EU preferences under risk.

Behavioral analysis of CP. As noted in the Introduction, this paper provides a fully behavioral analysis of CP. To illustrate the key ingredients of the analysis, refer back to the decision tree in Fig. 1, and adopt a simplified version of the notation to be introduced in Sec. 3 (an analogous treatment can be provided for the tree in Fig. 2). Denote the original tree in Fig. 1 by f_x ; also denote by c_x and s_x the subtrees of f_x where c or, respectively, s is the only action available at the initial node. Finally, denote by ca_x and cb_x the subtrees of c_x where a or, respectively, b is the only action available at the second decision node; note that s_x , ca_x and cb_x can be interpreted as fully specified plans of action. Assume that, at time 0, the DM expresses the following strict preferences (\succ) and indifferences (\sim) over decision trees:

$$ca_0 \sim c_0 \sim f_0 \succ s_0 \succ cb_0 \quad \text{and} \quad cb_1 \succ f_1 \sim s_1 \succ ca_1 \sim c_1. \quad (2)$$

The preferences in Eq. (2) exhibit two key features. First, *preferences over plans are consistent with act preferences in the Ellsberg paradox*, and more generally with the assumed MEU preferences at the initial node. Specifically, $ca_0 \succ s_0 \succ cb_0$ and $ca_1 \prec s_1 \prec cb_1$ correspond to the DM's

ranking of the acts $(1, 0, x)$, $(0, 1, x)$ and $(\frac{1}{2}, \frac{1}{2}, x)$ for $x = 0, 1$ provided by the MEU utility index V_0 .

The remaining preference rankings involve *non-degenerate trees*, and do *not* merely follow from the assumption of MEU preferences (even if augmented with prior-by-prior updating); rather, *they reflect the intuition behind sophistication* that is the focus of this paper. In particular, the indifference $c_1 \sim ca_1$ indicates that the DM does not value the *option* to choose b at time 1, when a is also available. This is not because she dislikes action b from the perspective of time 0: on the contrary, the ranking $cb_1 \succ ca_1$ suggests that she would like to *commit* to choosing b at time 1. Therefore, it must be the case that this DM correctly anticipates her future strict preference for a over b , and evaluates the tree c_1 accordingly.

I emphasize that this argument relies crucially upon the rankings of non-degenerate trees—e.g., $c_1 \sim ca_1$ in Eq. (2). Indeed, this pattern of preferences will constitute the *behavioral definition* of Sophistication in §4.3.2. More generally, the proposed approach leverages preferences over *trees* to elicit conditional preferences and analyze sophistication and related behavioral traits, just like the literature on menu choice leverages preferences over *menus* to investigate attitudes toward flexibility or commitment, as well as temptation and self-control (see §6).³

The preferences in Eq. (2) indicate how this particular DM resolves the conflict between her prior and posterior preferences. Furthermore, the rankings $f_0 \sim ca_0$ and $f_1 \sim s$ may be interpreted as the behavioral implications of Sophistication: if $x = 0$, the DM will choose c and plan to follow with a , and if $x = 1$, she will choose s instead—as predicted by CP.

It was just argued that, if the DM is *assumed* to strictly prefer a to b at the second decision node, then the prior preferences in Eq. (2) *reveal* that she is sophisticated. But, reversing one’s perspective, the following interpretation is equally legitimate: if the DM is *assumed* to be sophisticated, then the prior preferences in Eq. (2) *reveal* her ranking of a vs. b at the second decision node. To elaborate, as noted above, the rankings $cb_1 \succ ca_1 \sim c_1$ suggest that the DM *expects* to choose a rather than b at the second decision node; if the DM is *assumed* to be sophisticated, this expectation must be correct, so she must actually prefer a to b at that node. In this respect, the DM’s prior preference relation \succsim over trees, partially described in Eq. (2),

³Although I assumed that Reduction holds in this specific example, the notation and formal setup *allow* the DM to strictly rank two *plans* p, p' that can be reduced to the same *act*. This is orthogonal to the issue of sophistication; imposing Reduction throughout would neither simplify nor hamper the analysis. See §3.3.

provides all the information required to analyze behavior in this example.

Details. Certain subtle aspects of CP in the context of choice under uncertainty require further analysis, and are fully dealt with in the remainder of this paper. First, eliciting conditional preferences in general trees requires a more refined approach than the one just described; the details are provided in Sec. 4.2. Note that only a weak form of sophistication is required.

Second, ties must be handled with care. The Sophistication axiom in Sec. 4.3.2 is purposely formulated so as to entail no restrictions in case multiple optimal actions exist at a node. Instead, a separate axiom captures the tie-breaking assumption that characterizes CP.

Third, this “division of labor” is essential in the setting of choice under uncertainty. Sec. 5.2 shows that, under solvability conditions that are satisfied by virtually all known parametric models of non-EU preferences, strengthening the Sophistication axiom so as to deal with ties as well has an undesirable side effect: it imposes a version of P2 on preferences *over acts*, and hence, for instance, rules out the modal preferences in the Ellsberg example.

3 Decision Setting

Due to the approach taken in this paper, the notation for decision trees must serve two purposes. First, it must provide a rigorous *description* of dynamic-choice problem; second, it must allow a precise, yet relatively straightforward formalization of “*tree-surgery*” operations—pruning actions at a given node, replacing actions at a node with different ones, and more generally “composing” new trees out of old ones. The proposed description of decision trees will be relatively familiar;⁴ however, formally describing tree-surgery operations requires a level of detail that is not needed in other treatments of dynamic choice under uncertainty.

For simplicity, attention is restricted to *finite* trees associated with a single, *fixed* sequence of information partitions; see §5.3 for possible extensions.

⁴ Epstein [11] and Epstein, Noor and Sandroni [12] adopt a similar notation for decision trees, although they are not motivated by (and do not define) tree-surgery operations. In the context of risk, the notation in Sec. 3 of Kreps and Porteus [32] is similar, again except for tree surgery; see Sec. 6 for further details.

3.1 Actions, Trees and Histories

Fix a state space Ω , endowed with an algebra Σ , and a connected and separable space X of outcomes. Information is modeled as a sequence of progressively finer partitions $\mathcal{F}_0, \dots, \mathcal{F}_T$ of Ω , for some $0 \leq T < \infty$, such that $\mathcal{F}_0 = \{\Omega\}$ and $\mathcal{F}_t \subset \Sigma$ for all $t = 1, \dots, T$ (sometimes referred to as a *filtration*). For every $t = 0, \dots, T$, the cell of the partition \mathcal{F}_t containing the state $\omega \in \Omega$ is denoted by $\mathcal{F}_t(\omega)$; also, a pair (t, ω) , where $t \in \{0, \dots, T\}$ and $\omega \in \Omega$, will be referred to as a *node*.

Trees and actions can now be defined recursively, as “menus of contingent menus of contingent menus...”. A bit more rigorously, define first a “tree” beginning at the terminal date T in state ω simply as an outcome $x \in X$. Inductively, define an *action* available in node (t, ω) as a map associating with each state $\omega' \in \mathcal{F}_t(\omega)$ a continuation tree beginning at node $(t + 1, \omega')$; to complete the inductive step, define a *tree* beginning at node (t, ω) as a finite collection, or menus, of actions available at (t, ω) . The details are as follows:

Definition 1 Let $F_T(\omega) = F_T = X$ for all $\omega \in \Omega$. Inductively, for $t = T - 1, \dots, 0$ and $\omega \in \Omega$, let

1. $A_t(\omega)$ be the set of \mathcal{F}_{t+1} -measurable functions $a : \mathcal{F}_t(\omega) \rightarrow F_{t+1}$ such that, for all $\omega' \in \mathcal{F}_{t+1}(\omega)$, $a(\omega') \in F_{t+1}(\omega')$;
2. $F_t(\omega)$ be the collection of non-empty, finite subsets of $A_t(\omega)$; and
3. $F_t = \bigcup_{\omega \in \Omega} F_t(\omega)$.

The elements of $A_t(\omega)$ and F_t are called **actions** and **trees** respectively.

Observe that the maps $\omega \rightarrow A_t(\omega)$ and $\omega \rightarrow F_t(\omega)$ are \mathcal{F}_t -measurable.

A tree is interpreted throughout as an exhaustive description of the choices available in a given decision problem; in particular, if two or more actions are available at a node, the individual cannot also “randomize” among them. Of course, randomization can be explicitly modeled, by suitably extending the state space and the description of the tree.

A history describes a possible path connecting two nodes in a tree: specifically, it indicates the actions taken and events observed along the path. Given the filtration $\mathcal{F}_0, \dots, \mathcal{F}_T$, the sequence of events observed is fully determined by the prevailing state of nature; thus, formally, a history is identified by the initial time t , the prevailing state ω , and the (possibly empty) sequence of actions taken. The details, and some related notation and terminology, are as follows:

Definition 2 A **history** starting at a node (t, ω) is a tuple $h = [t, \omega, \mathbf{a}]$, where either

- $\mathbf{a} = (a_t, \dots, a_\tau)$, with $t \leq \tau \leq T - 1$, $a_t \in A_t(\omega)$ and, for all $\bar{t} = t + 1, \dots, \tau$, $a_{\bar{t}} \in A_{\bar{t}-1}(\omega)$; or
- $\mathbf{a} = \emptyset$ (an empty list).

The cardinality of \mathbf{a} is denoted $|\mathbf{a}|$. Furthermore:

1. If $h = [t, \omega, \mathbf{a}]$, $\mathbf{a} = \emptyset$ and $a_t \in A_t(\omega)$, then $\mathbf{a} \cup a_t = (a_t)$; and if $\mathbf{a} = (a_t, \dots, a_\tau)$, $\tau < T - 1$ and $a_{\tau+1} \in A_\tau(\omega)$, then $\mathbf{a} \cup a_{\tau+1} \equiv (a_t, \dots, a_\tau, a_{\tau+1})$.
2. A history $[t, \omega, \mathbf{a}]$ is **terminal** iff $t + |\mathbf{a}| = T$, and **initial** iff $\mathbf{a} = \emptyset$.
3. A history $h = [t', \omega', \mathbf{a}]$ is **consistent with a tree** $f \in F_t(\omega)$ if $t' = t$, $\omega' \in \mathcal{F}_t(\omega)$, and either $\mathbf{a} = \emptyset$ or the first action in \mathbf{a} is an element of f ; in this case, the **continuation tree of f starting at h** is $f(h) = f$ if $\mathbf{a} = \emptyset$, and $f(h) = a_\tau(\omega')$ if $\mathbf{a} = (a_t, \dots, a_\tau)$.

Certain special trees play an important role in the analysis. First, a **plan** is a tree where a single action is available at every decision point. Formally, a tree $f \in F_t$ is a plan if, for every history $h = [t, \omega, \mathbf{a}]$ consistent with f , $|f(h)| = 1$. The set of plans in F_t and $F_t(\omega)$ will be denoted by F_t^p and $F_t^p(\omega)$ respectively. Second, a **constant plan** yields the same outcome in every state of the world. Formally, $f_{t,\omega}^x \in F_t(\omega)$ is the unique plan such that, for every terminal history h consistent with $f_{t,\omega}^x$, $f_{t,\omega}^x(h) = x$. If the node (t, ω) can be understood from the context, the plan $f_{t,\omega}^x$ will be denoted simply by x .

As an example, the tree in Fig. 1, as well as its subtrees, can be formally defined as follows (recall that a simplified notation was used in the Introduction). Let $T = 2$, $\mathcal{F}_1 = \{\{\alpha, \beta\}, \{\gamma\}\}$, and $\mathcal{F}_2 = \{\{\alpha\}, \{\beta\}, \{\gamma\}\}$. The two choices available at the second decision node in Fig. 1 correspond to the time-1 actions $a, b \in A_1(\alpha) = A_1(\beta)$ defined by

$$a(\alpha) = 1, a(\beta) = 0 \quad \text{and} \quad b(\alpha) = 0, b(\beta) = 1. \quad (3)$$

Next, define the time-0 actions $c_x, s_x, ca_x, cb_x \in A_0(\alpha) = A_0(\beta) = A_0(\gamma)$ by

$$\text{for } \omega = \alpha, \beta, \quad c_x(\omega) = \{a, b\}, \quad s_x(\omega) = \frac{1}{2}, \quad ca_x(\omega) = \{a\}, \quad cb_x(\omega) = \{b\}; \quad (4)$$

$$c_x(\gamma) = s_x(\gamma) = ca_x(\gamma) = cb_x(\gamma) = x. \quad (5)$$

Here, x and $\frac{1}{2}$ denote the constant plans $f_{1,\gamma}^x$ and $f_{1,\gamma}^{\frac{1}{2}}$ respectively.

Now the full tree in Fig. 1 is formally defined as $f_x \equiv \{c_x, s_x\}$; the subtree beginning with the choice of c (respectively, s) is $\{c_x\}$ (respectively $\{s_x\}$); and the plans corresponding to the choice of c at the initial node, followed by a (respectively b) at the second decision node are $\{ca_x\}$ and $\{cb_x\}$. Finally, there are three non-terminal histories consistent with f_x : \emptyset , $[0, \alpha, c_x]$ and $[0, \beta, c_x]$.

3.2 Composite Trees

Fix $f \in F_t$, a history $h = [t, \omega, \mathbf{a}]$ consistent with f , and another tree $g \in F_{t+|\mathbf{a}|}(\omega)$. The **composite tree** $g_h f$ is, intuitively, a tree that coincides with f everywhere except at history h , where it coincides with g . Formalizing this notion is somewhat delicate, so I first provide some heuristics.

Since $h = [t, \omega, \mathbf{a}]$ is consistent with f and $\mathbf{a} = (a_t, \dots, a_\tau)$, with $\tau \geq t$, the last element a_τ of the action list \mathbf{a} satisfies $a_\tau(\omega') = f(h)$ for all $\omega' \in \mathcal{F}_{\tau+1}(\omega)$. To capture the idea that $f(h)$ is replaced with g , one would like to replace a_τ in the list \mathbf{a} with a new action \bar{a}_τ such that $\bar{a}_\tau(\omega') = g$ at such states, and $\bar{a}_\tau(\omega') = a_\tau(\omega')$ elsewhere. However, recall that, by definition, $a_{\tau-1}(\omega')$ must contain a_τ for all $\omega' \in \mathcal{F}_\tau(\omega)$; if a_τ is replaced with \bar{a}_τ , it is also necessary to “modify” $a_{\tau-1}$ so that it now contains \bar{a}_τ rather than a_τ in such states. These modifications must be carried out inductively for all actions $a_{\tau-1}, a_{\tau-2}, \dots, a_t$; this yields a new, well-defined action list $\bar{\mathbf{a}} = (\bar{a}_t, \dots, \bar{a}_\tau)$. Finally, recall that, by definition, the history $h = [t, \omega, \mathbf{a}]$ is consistent with f precisely when the first action a_t in the list \mathbf{a} is an element of f (trees are sets of actions). Then, the tree $g_h f$ differs from f precisely in that the action a_t is replaced with \bar{a}_t .

Now for the formal details. If $\mathbf{a} = \emptyset$, then let $g_h f \equiv g$. Otherwise, write $\mathbf{a} = (a_t, \dots, a_\tau)$, with $\tau \geq t$; let $\bar{a}_\tau(\omega') = g$ for all $\omega' \in \mathcal{F}_{\tau+1}(\omega)$, and $\bar{a}_\tau(\omega') = a_\tau(\omega')$ for $\omega' \in \mathcal{F}_\tau(\omega) \setminus \mathcal{F}_{\tau+1}(\omega)$. Inductively, for $\bar{t} = \tau - 1, \dots, t$, let $\bar{a}_{\bar{t}}(\omega') = \{\bar{a}_{\bar{t}+1}\} \cup (a_{\bar{t}}(\omega') \setminus \{a_{\bar{t}+1}\})$ for all $\omega' \in \mathcal{F}_{\bar{t}+1}(\omega)$, and $\bar{a}_{\bar{t}}(\omega') = a_{\bar{t}}(\omega')$ for $\omega' \in \mathcal{F}_{\bar{t}}(\omega) \setminus \mathcal{F}_{\bar{t}+1}(\omega)$. Finally, let $g_h f$ denote the set $\{\bar{a}_t\} \cup (f \setminus \{a_t\})$.

As a special case, consider a node (t, ω) and a plan $f \in F_0^p$. Since, by definition, a single action is available in f at any node, there is a unique history consistent with f that corresponds to the node (t, ω) ; it is then possible to define a tree that, informally, coincides with f everywhere except *at time t , in case event $\mathcal{F}_t(\omega)$ occurs*. Such tree will be denoted $g_{t,\omega} f$.

Formally, since f is a t -period plan, there is a unique action list $\mathbf{a} = (a_0, \dots, a_{t-1})$ such that

$h = [0, \omega, \mathbf{a}]$ is consistent with f . Then, for all $g \in F_t(\omega)$, let $g_{t,\omega}f \equiv g_h f$.⁵ The notation “ $g_{t,\omega}f$ ” is modeled after “ $g_E f$,” which is often used to indicate composite Savage acts.

3.3 Preferences, Reduction and Consequentialism

Definition 3 A **conditional preference system** (CPS) is a tuple $(\succ_{t,\omega})_{0 \leq t < T, \omega \in \Omega}$, such that, for every t and ω , $\succ_{t,\omega}$ is a binary relation on $F_t(\omega)$, and furthermore $\omega' \in \mathcal{F}_t(\omega)$ implies $\succ_{t,\omega} = \succ_{t,\omega'}$. The time-0 preference is also denoted simply by \succ .

Three aspects are worth emphasizing. First, preferences are assumed to be “adapted to \mathcal{F} ”: for each $t = 0, \dots, T-1$, $\succ_{t,\omega}$ is measurable with respect to \mathcal{F}_t . This reflects the assumption that $\succ_{t,\omega}$ is the DM’s ranking of trees conditional upon observing the event $\mathcal{F}_t(\omega)$ at time t .

Second, recall that, in the Introduction, preferences over *plans* were implicitly defined by first “reducing” plans to acts in the obvious way, and then invoking the DM’s preferences over acts, represented by the functional V . While this is the “textbook” approach to dynamic choice with EU preferences, there are compelling reasons to consider alternatives. For instance, the DM may display a preference for early or late resolution of the uncertainty, as in Kreps and Porteus [32], Epstein and Zin [10], Segal [41] and, in a fully subjective setting, Klibanoff and Ozdenoren [30]. To allow for such preference models, **reduction is not assumed** in the main results of this paper, Theorems 2 and 3. The proposed approach takes the DM’s preferences over plans as given, as part of her CPS, regardless of whether or not they are obtained from underlying preferences over acts by reduction.

Third, the assumption that the only “conditioning information” relevant to the preference relation $\succ_{t,\omega}$ is the event $\mathcal{F}_t(\omega)$ implies that our analysis is **consequentialist**: in particular, two actions $a, b \in A_t(\omega)$ are ranked in the same way, in any decision tree where they may be available. To elaborate, if the actions a and b are available at a history $h = [t, \omega, \mathbf{a}]$ consistent with a tree f , their ranking will of course, depend upon the realized event $\mathcal{F}_t(\omega)$; however, prior choices made and alternatives discarded on the path to h , choices the DM would have had to make at counterfactual histories in f , or events that might have obtained but didn’t, are irrele-

⁵ \mathbf{a} is also the unique list such that, for any $\omega' \in \mathcal{F}_t(\omega)$, the history $h = [0, \omega', \mathbf{a}]$ is consistent with f ; furthermore, the definition of composite trees implies that $g_h f = g_{h'} f$, consistently with the intended interpretation of $g_{t,\omega}f$.

vant. This is a standard property of EU preferences and Bayesian updating, and is preserved in most applications of non-EU and ambiguity-sensitive preferences. However, some alternative theoretical approaches to dynamic choice with non-EU preferences or under ambiguity relax consequentialism to salvage dynamic consistency; this important point is discussed in Sec. 6.

To conclude, recall that a binary relation is a **weak order** iff it is complete and transitive.

4 Main Results

This section presents the main results of this paper. Theorem 2 in §4.2 shows that sophistication provides a way to elicit conditional preferences *over acts and trees* from prior preferences *over trees*. §4.3 then takes as primitive a CPS and provides a definition (Def. 5) and characterization (Theorems 3 and 4) of CP in the context of choice under uncertainty. §4.4.1 considers CP for MEU preferences and prior-by-prior updating, and §4.4.2 analyzes the value of information under CP. All proofs are in the Appendix. Further motivation and discussion is provided in §5.

As a preliminary step, §4.1 formalizes the connection between dynamic consistency, Bayesian updating, and Savage’s Postulate P2 mentioned in the Introduction. This result constitutes a useful benchmark, and aids in the interpretation of Theorems 2–4.

4.1 Dynamic Consistency, Bayesian updating, and Postulate P2

The main result of this subsection should be considered a “folk theorem”: various versions of it exist in the literature, beginning with Savage’s own ([40, §2.7]). Its original statement concerns preferences over *acts*; I restate it in terms of preferences over *plans* merely to avoid introducing new notation.⁶ Also note that, while the definition of a CPS involves general trees, throughout this subsection axioms, definitions and results are explicitly restricted to preferences over plans.

For simplicity, assume that every event in the filtration $\mathcal{F}_0, \dots, \mathcal{F}_T$ is not *Savage-null*: for every node (t, ω) , it is *not* the case that $p \sim r_{t,\omega} p$ for all $p \in F_0^p$ and $r \in F_t^p$. I begin by formalizing DC, Savage’s Postulate P2, and Savage’s qualitative notion of Bayesian Updating; I follow Savage [40], §2.7 throughout, to which the reader is referred for interpretation (for DC, see also Epstein

⁶That is, to further clarify, the resulting additional generality is inessential for my purposes.

and Schneider [9]). Note however that Savage's postulates and definitions pertain to all possible conditioning events, whereas I restrict attention to the elements of the filtration $\mathcal{F}_0, \dots, \mathcal{F}_T$.

Axiom 4.1 (Dynamic Consistency — DC) *For all nodes (t, ω) with $t < T$, and all actions $a, b \in A_t(\omega)$ such that $\{a\}, \{b\} \in F_t^p$: if $a(\omega') \succ_{t+1, \omega'} b(\omega')$ for all $\omega' \in \mathcal{F}_t(\omega)$, then $\{a\} \succ_{t, \omega} \{b\}$; furthermore, if time- $(t + 1)$ preferences are strict for some $\omega^* \in \mathcal{F}_t(\omega)$, then $\{a\} \succ_{t, \omega} \{b\}$.⁷*

Axiom 4.2 (Postulate P2) *For all plans $p, q \in F_0^p$, all nodes (t, ω) , and all plans $r, s \in F_t^p(\omega)$,*

$$r_{t, \omega} p \succ_{t, \omega} s_{t, \omega} p \quad \Rightarrow \quad r_{t, \omega} q \succ_{t, \omega} s_{t, \omega} q.⁸$$

Note that DC relates preferences at different histories; on the other hand, P2 pertains to prior preferences alone. As asserted in Sec. 2, the MEU preferences specified in Sec. 2, jointly with the reduction assumption, yield a violation of Axiom 4.2: take $r = \{a\}$, $s = \{b\}$, $p = \{ca_0\}$ and $q = \{cb_0\}$. This is, of course, the main message conveyed by Ellsberg [7].

Finally, say that the restriction of $\succ_{t, \omega}$ to F_t^p is derived from \succ via *Bayesian updating* (cf. Savage [40], p. 22) if, for all plans $r, s \in F_t^p$,

$$r \succ_{t, \omega} s \quad \Leftrightarrow \quad r_{t, \omega} p \succ_{t, \omega} s_{t, \omega} p \quad \text{for some plan } p \in F_0.$$

For ex-ante EU preferences the above condition indeed characterizes Bayesian updating of the DM's prior. The following result is then straightforward:⁹

Proposition 1 *Consider a CPS $(\succ_{t, \omega})_{t, \omega}$. The following statements are equivalent:*

(1) *\succ is a weak order on F_0^p , Axiom 4.2 (Postulate P2) holds, and for every node (t, ω) , the restriction of $\succ_{t, \omega}$ to F_t^p is derived from \succ via Bayesian updating;*

(2) *every $\succ_{t, \omega}$ is a weak order on F_t^p , and Axiom 4.1 (DC) holds.*

This result depends upon the assumption of consequentialism implicit in the framework: according to Def. 5, each preference $\succ_{t, \omega}$ is defined on (sub)trees with initial event $\mathcal{F}_t(\omega)$ (cf. §6). For a version of this result that relaxes consequentialism, see Epstein and Le Breton [8].

⁷ a and b are actions, whereas the singleton sets $\{a\}$ and $\{b\}$ are trees; on the other hand, $a(\omega')$ and $b(\omega')$ are trees in $F_{t+1}(\omega)$. Finally, F_t^p is a set of plans, i.e. special types of trees, and $\succ_{t, \omega}$ and $\succ_{t+1, \omega'}$ are defined over trees.

⁸A plan is, a fortiori, a t -period plan, so $r_{t, \omega} p$, etc. are well-defined: cf. §3.2.

⁹The proof is similar to that of analogous results (e.g. Ghirardato [13]); hence, it is available upon request.

Prop. 1 highlights the tension between dynamic consistency and ambiguity that was anticipated in the Introduction. However, I now wish to emphasize the implications of this result for Bayesian updating. If one assumes that prior preferences satisfy P2, then one can *define* conditional preferences via Bayesian updating, and in this case Prop. 1 implies that DC will hold. Conversely, if one assumes that DC holds, Prop. 1 implies that Bayesian updating provides a way to *elicit* conditional preferences; furthermore, ex-ante preferences necessarily satisfy P2.

4.2 Eliciting Conditional Preferences

Turn now to the main results of the paper, beginning with the elicitation of conditional preferences. First, we adopt a standard requirement: the conditioning event should “matter.”

Assumption 4.1 (Non-null conditioning events) For every node (t, ω) and prizes x, y such that $x \succ y$, there exists a plan $g \in F_0^p$ such that $x_{t,\omega} g \succ y_{t,\omega} g$.

For general preferences over acts or plans, Assumption 4.1 is stronger than the requirement that every set $\mathcal{F}_t(\omega)$ not be Savage-null (cf. §4.1); however, the two notions coincide, for instance, for MEU (and of course EU) preferences. Assumption 4.1 is weaker than analogous conditions in the literature—e.g. the notions of “non-null” event in Ghirardato and Marinacci [14], which requires that $g = y$.

4.2.1 Beliefs about Conditional Preferences

I begin by proposing a procedure that elicits the DM’s *beliefs about her own future preferences*; the details are in Def. 4. To motivate it, refer to the decision tree in Fig. 1 with $x = 1$; adopt the notation in Eqs. (3)–(5). Since $\{cb_1\} \succ \{ca_1\} \sim \{c_1\}$ ex-ante, it was argued in §2 that $\{a\} \succ_{1,\alpha} \{b\}$ [equivalently, $\{a\} \succ_{1,\beta} \{b\}$]: if the DM would like to commit to b at the second decision node, but deems the tree $\{c_1\}$ just as good as committing to a , it must be the case that the DM *expects* to choose a at the second decision node, if both a and b are available.

However, this argument fails if $\{ca_1\} \sim \{cb_1\}$: in this case, the indifference $\{c_1\} \sim \{ca_1\}$ is not sufficiently informative as to the relative conditional ranking of a vs. b . Detecting conditional indifferences is even more delicate. Thus, a different, but related approach must be adopted:

Definition 4 For all nodes (t, ω) and trees $f, f' \in F_t(\omega)$, f is **conjecturally weakly preferred to f' given (t, ω)** , written $f \succsim_{t,\omega}^0 f'$, iff there exists a prize $z \in X$ such that, for all plans $g \in F_0^p$,

$$\forall y \in X, \quad y \succ z \Rightarrow (f' \cup y)_{t,\omega} g \sim y_{t,\omega} g \quad \text{and} \quad z \succ y \Rightarrow (f \cup y)_{t,\omega} g \sim f_{t,\omega} g. \quad (6)$$

The superscript 0 in the notation $\succsim_{t,\omega}^0$ emphasizes that this conjectural conditional preference relation is defined *solely* in terms of the DM's time-0, i.e. prior preferences.

The logic behind Def. 4 is as follows. Suppose that the DM *believes* that $f \succsim_{t,\omega} f'$. Under suitable regularity (in particular, solvability) assumptions that are captured by the axioms in the next subsection, there exists a prize $z \in X$ such that (the DM will also believe that) $f \succ_{t,\omega} z \succ_{t,\omega} f'$. Now consider another prize $y \in X$ such that a priori $y \succ z$; if the DM does not expect her preferences over prizes to change, then (she will believe that) $y \succ_{t,\omega} z$ as well, and hence that $y \succ_{t,\omega} f'$. But this implies that she will expect y to be chosen rather than f' in the tree $(f' \cup y)_{t,\omega} g$ at node (t, ω) .¹⁰ As in the example of Sec. 2, the ex-ante indifference between $(f' \cup y)_{t,\omega} g$ and $y_{t,\omega} g$ now reflects this belief. The argument for the case $z \succ y$ is similar.

Note that, for every tree considered in Eq. (6), there is a unique path from the initial history to the node (t, ω) , because g is a plan; furthermore, the event $\mathcal{F}_t(\omega)$ is not Savage-null. Hence, the DM “cannot avoid” contemplating her choices at that node.

4.2.2 Axioms and Characterization

The axioms I consider are divided into two groups. Axioms 4.3–4.6 relate the DM's *actual* conditional preferences with her prior preferences; Axioms 4.7–4.9 instead concern the DM's prior preferences only, and ensure that the definition of *conjectural* conditional preferences (in Def. 4) is well-posed (that is, non-contradictory).

Axiom 4.3 (Stable Tastes) For all $x, x' \in X$, and all nodes (t, ω) : $x \succ_{t,\omega} x'$ if and only if $x \succ x'$.

Axiom 4.4 (Conditional Dominance) For all nodes (t, ω) , all $f \in F_t(\omega)$, and all $x', x'' \in X$: if $x' \succ f(h) \succ x''$ for all terminal histories h of f , then $x' \succ_{t,\omega} f \succ_{t,\omega} x''$.

¹⁰ $f' \cup y$ denotes the time- t tree that contains all actions in f' , plus the unique initial action in the plan $f'_{t,\omega}$ that leads to the prize y in every state of nature. In other words, the notation exploits (a) the simplified notation for prizes, and (b) the fact that trees are just sets of acts, and therefore unions of trees are also well-defined trees.

Axiom 4.5 (Conditional Prize-Tree Continuity) For all nodes (t, ω) and all $f \in F_t(\omega)$, the sets $\{x \in X : x \succ_{t,\omega} f\}$ and $\{x \in X : x \preceq_{t,\omega} f\}$ are closed in X .

Axiom 4.6 (Weak Sophistication) For all nodes (t, ω) , plans $g \in F_0^p$, trees $f \in F_t(\omega)$, and prizes $x \in X$:

$$x \succ_{t,\omega} f \Rightarrow (f \cup x)_{t,\omega} g \sim x_{t,\omega} g \quad \text{and} \quad x \preceq_{t,\omega} f \Rightarrow (f \cup x)_{t,\omega} g \sim f_{t,\omega} g.$$

Axiom 4.3 states that tastes, i.e. preferences over prizes, are unaffected by conditioning.¹¹ Axioms 4.4 and 4.5 are standard, and ensure that conditional certainty equivalents exist (recall that X is assumed to be a connected and separable topological space).

Axiom 4.6 assumes just enough sophistication to ensure that *conjectural* and *actual* conditional preferences coincide: in particular, the logic of sophistication is applied only to comparisons between a tree and a constant prize, and then only if the DM has no other choice available on the path to the node (t, ω) . Preferences at times $t > 0$ are *not* required to be sophisticated.

Turn now to the second group of axioms.

Axiom 4.7 (Prize Continuity) For all $\bar{x} \in X$, the sets $\{x \in X : x \succ \bar{x}\}$ and $\{x \in X : x \preceq \bar{x}\}$ are closed in X .

Axiom 4.8 (Dominance) Fix a node (t, ω) , a tree $f \in F_t(\omega)$, a plan $g \in F_0^p$ and a prize $x \in X$.

(i) If $f(h) \succ x$ for all terminal histories h of f , then $(f \cup x)_{t,\omega} g \sim f_{t,\omega} g$.

(ii) If $f(h) \preceq x$ for all terminal histories h of f , then $(f \cup x)_{t,\omega} g \sim x_{t,\omega} g$.

Axiom 4.8 reflects stability of preferences over outcomes. If the individual's preferences over X do not change when conditioning on $\mathcal{F}_t(\omega)$, then in (i) she will expect *not* to choose x at node (t, ω) , because f yields strictly better outcomes at every terminal history; similarly for (ii). As in Sec. 2, the indifferences in (i) and (ii) capture the DM's expectations.

The next axiom is a “beliefs-based” counterpart to Weak Sophistication:

¹¹For the present purposes, it would be sufficient to impose this requirement on a suitably rich subset of prizes. For instance, if X consists of consumption streams, it would be enough to restrict Axiom 4.3 to constant streams.

Axiom 4.9 (Separability) Consider a node (t, ω) , $f \in F_t(\omega)$, plans $g, g' \in F_0^p$ and $x, y \in X$. Then:

(i) $(f \cup y)_{t, \omega} g \not\prec f_{t, \omega} g$ and $x \succ y$ imply $(f \cup x)_{t, \omega} g' \sim x_{t, \omega} g'$;

(ii) $(f \cup y)_{t, \omega} g \not\prec y_{t, \omega} g$ and $x \prec y$ imply $(f \cup x)_{t, \omega} g' \sim f_{t, \omega} g'$.

To interpret, consider first the case $g = g'$ and fix a prize y . According to the by-now familiar logic of belief elicitation, $(f \cup y)_{t, \omega} g \not\prec f_{t, \omega} g$ indicates that the DM believes that she will *not* strictly prefer f to y given $\mathcal{F}_t(\omega)$ —otherwise indifference would have to obtain. Thus, if $x \succ y$ and the DM's preferences over X are stable, she will also expect to strictly prefer x to f given $\mathcal{F}_t(\omega)$; again, the elicitation logic yields $(f \cup x)_{t, \omega} g \sim x_{t, \omega} g$. The interpretation of (ii) is similar.

Additionally, Axiom 4.9 implies that these conclusions are independent of the particular t -period plan under consideration, and hence of what the decision problem looks like if the event $\mathcal{F}_t(\omega)$ does not obtain. In this respect, Axiom 4.9 reflects a form of “separability.” More generally, Axiom 4.9 essentially requires that Eq. (6) in Def. 4 hold for all plans g , or for none. There is a close analogy with the role of Savage's Postulate P2: see §4.1 for details.

The main result of this section can now be stated.

Theorem 2 Suppose that Assumption 4.1 holds. Consider the CPS $(\succ_{t, \omega})$, and assume that \succ is a weak order on F_0 . Then the following statements are equivalent.

1. \succ satisfies Axioms 4.7–4.9; furthermore, for all nodes (t, ω) , $\succ_{t, \omega} = \succ_{t, \omega}^0$.
2. For every node (t, ω) , $\succ_{t, \omega}$ is a weak order, and satisfies Axioms 4.3–4.6.

Theorem 2 and Proposition 1 in §4.1 are structurally similar: Axioms 4.7–4.9 play the role of Postulate P2 (but add solvability requirements), the definition of conjectural conditional preferences corresponds to Bayesian updating, and Axioms 4.3–4.6 correspond to DC (but again add solvability requirements). The interpretation is also similar: under Axioms 4.7–4.9, Def. 4 yields well-behaved conditional preferences, and hence can be taken as the *definition* of conditional preferences; in this case, Axioms 4.3–4.6 will hold. Conversely, if Axioms 4.3–4.6 hold, the beliefs derived via Def. 4 from prior preferences are actually correct, so that Def. 4 can be seen as a way to *elicit* actual conditional preferences. The main differences are that, of course, Theorem 2 does *not* rely on P2 or DC, and concerns preferences over non-degenerate trees.

4.3 A decision-theoretic analysis of Consistent Planning

4.3.1 Consistent Planning under Uncertainty

As noted in the Introduction, *Consistent Planning* (CP) is a refinement of backward induction. If there are unique optimal actions at any point in the tree, the two concepts coincide. Otherwise, CP complements backward induction with a specific *tie-breaking* rule: indifferences at a history h are resolved by considering preferences at the history that immediately precedes h .

To illustrate, consider the tree in Fig. 1 with $x = 1$, but assume MEU preferences with priors $C = \{q \in \Delta(\Omega) : \frac{1}{90} \leq q(\alpha) \leq \frac{30}{90}, \frac{2}{90} \leq q(\beta) \leq \frac{15}{90}\}$. Continue to assume prior-by-prior updating and reduction, and again adopt the notation in Eqs. (3)–(5). It can then be verified that $\{a\} \sim_{1,\alpha} \{b\}$; however, $\{ca_1\} \succ \{cb_1\}$, so CP prescribes that the DM will follow c with a . The corresponding plan $\{ca_1\}$ is strictly preferred to $\{s_1\}$, so the unique CP “solution” of this tree is the plan $\{ca_1\}$.

Algorithmically, CP operates as follows. For each history $h = [t, \omega, \mathbf{a}]$ in a tree f , consider first the set $\text{CP}_f^0(h)$ of actions $b \in A_{t+|\mathbf{a}|}(\omega)$ that, for every realization $\omega' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega)$, prescribe a continuation action $a_{t+|\mathbf{a}|+1, \omega'}$ that has survived prior iterations of the procedure. Intuitively, such actions b correspond to plans that the DM “will actually follow.” Then, out of these actions, select the conditionally optimal ones: this completes the induction step and defines the set $\text{CP}_f(h)$. Def. 5 is modeled after analogous definitions in Strotz [46] and Gul and Pesendorfer [20], except that it is phrased in terms of preferences, rather than numerical representations.

Definition 5 (Consistent Planning) Consider a tree $f \in F_t(\omega)$. For every terminal history $h = [t, \omega, \mathbf{a}]$ consistent with f , let $\text{CP}_f(h) = \{f(h)\}$. Inductively, if $h = [t, \omega, \mathbf{a}]$ is consistent with f and $\text{CP}_f([t, \omega', \mathbf{a} \cup a])$ has been defined for every $\omega' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega)$ and $a \in f(h)$, let

$$\begin{aligned} \text{CP}_f^0(h) &= \left\{ b \in A_{t+|\mathbf{a}|}(\omega) : \exists a \in f(h) \text{ s.t. } \forall \omega' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega), \right. \\ &\quad \left. b(\omega') = \{a_{+1, \omega'}\} \text{ for some } a_{+1, \omega'} \in \text{CP}_f([t, \omega', \mathbf{a} \cup a]) \right\} \text{ and} \\ \text{CP}_f(h) &= \left\{ b \in \text{CP}_f^0(h) : \forall a \in \text{CP}_f^0(h), \{b\} \succ_{t+|\mathbf{a}|, \omega} \{a\} \right\}. \end{aligned}$$

A plan $\{a\} \in F_t(\omega)$ is a **consistent-planning solution** of f if $a \in \text{CP}_f([t, \omega, \emptyset])$.¹²

¹²To help parse notation, a , $a_{+1, \omega'}$ and b in this definition are acts; $b(\omega')$ must therefore be a tree, and in particular the definition requires that it be the tree $\{a_{+1, \omega'}\}$ having a single initial action $a_{+1, \omega'}$ taken from the set $\text{CP}_f([t, \omega', \mathbf{a} \cup a])$. Finally, braces in $\{b\} \succ_{t+|\mathbf{a}|, \omega} \{a\}$ are required because $\succ_{t+|\mathbf{a}|, \omega}$ is defined over trees, not actions.

Note that, in order to carry out the CP procedure, it is only necessary to specify the DM's preferences over *plans*. The output of the CP algorithm is also a set of plans.¹³ Moreover, it is straightforward to verify that, if preferences over plans are complete and transitive, then Def. 5 is well-posed: it always delivers a non-empty set of solutions that the DM deems equally good.

4.3.2 Behavioral Characterization of Consistent Planning

The behavioral analysis of CP takes as input the DM's CPS $(\succsim_{t,\omega})$. The key assumption of *Sophistication* was introduced in Sec. 2; Axiom 4.6 applies the same principle to a small set of trees, with unique features. To capture the implications of *Sophistication* in general trees, it will be assumed that *pruning conditionally dominated actions leaves the DM indifferent*. Formally, if g is a subset of actions available in the tree f at the history h , and every action $b \in g$ is *strictly preferred* to every action w that lies in $f(h)$ but not in g , then ex-ante the DM must be indifferent between f and the tree $g_h f$ in which the inferior actions have been pruned:

Axiom 4.10 (Sophistication) *For all $f \in F_t$, all histories $h = [t, \omega, \mathbf{a}]$ consistent with f and such that $\mathbf{a} \neq \emptyset$, and all $g \subset f(h)$: if, for all $b \in g$ and $w \in f(h) \setminus g$, $\{b\} \succ_{t+|\mathbf{a}|,\omega} \{w\}$, then $f \sim_{t,\omega} g_h f$.*

Observe that Axiom 4.10 is silent as far as *indifferences* at node $(t + |\mathbf{a}|, \omega)$ are concerned. For instance, if $f(h) = \{a, b\}$ and $\{a\} \sim_{t+|\mathbf{a}|,\omega} \{b\}$, the axiom does *not* require that $f \sim_{t,\omega} \{a\}_h f \sim_{t,\omega} \{b\}_h f$. This allows for the possibility that, ex-ante, the DM have a strict preference for *commitment* to a or b ; Axiom 4.11 deals with these situations. Axiom 4.10 is also silent if h is the initial history of f : Axiom 4.12 below encodes the assumptions required in this case. This “division of labor” is *crucial* so as to avoid unduly restricting ambiguity attitudes: see §5.2.

The next axiom formalizes the tie-breaking assumption that characterizes CP within the class of backward-induction solutions: if the DM is indifferent among two or more actions at a history h , then she can *precommit* (more precisely, expects to be able to precommit) to any of them at the history that immediately precedes h . It is important to emphasize that *no* such precommitment is possible in case the individual has strict preferences over actions at h : in such cases, the full force of the Sophistication axiom applies.

¹³Formally, $CP_f([t, \omega, \emptyset])$ is a set of *actions*, not plans; however, if $a \in CP_f([t, \omega, \emptyset])$, then $\{a\}$ is a plan.

To formalize this assumption, the notion of a *next-period commitment version* of a tree is required. Again, refer to the tree f_x in Fig. 1; as it turns out, the notation in Eqs. (3)–(5) greatly simplifies the exposition. Consider a modified version of the tree $f_x = \{c_x, s_x\}$ where the action c_x at the initial history \emptyset is replaced with the actions ca_x and cb_x . Recall that, while c_x allows a choice between a and b at the second decision node, ca_x and cb_x enforce a commitment to a and, respectively, b : cf. Eq. (4). The resulting tree $\{ca_x, cb_x, s_x\}$, referred to as the “next-period commitment version” of f_x , is depicted in Fig. 3.




Figure 3: Next-period commitment version of Fig. 1

To reflect the DM’s ability to precommit in case of future indifferences, it will be assumed that, if $\{a\} \sim_{1,\alpha} \{b\}$, the DM is *indifferent ex-ante* between $f_x = \{c_x, s_x\}$ and its *next-period commitment version* $\{ca_x, cb_x, s_x\}$. Intuitively, if $\{a\} \sim_{1,\alpha} \{b\}$, the DM regards the original tree as if it afforded the same “physical” ability to commit as its next-period commitment version.

In the tree f_x , non-trivial future choices must be made only following c_x , and only if $\omega \in \{\alpha, \beta\}$; this simplifies the construction of its next-period commitment version. For a general tree, one proceeds as follows. Given a tree f at a node (t, ω) , one fixes an initial action a in the tree f ; in every state $\omega' \in \mathcal{F}_t(\omega)$, a leads to a continuation tree $a(\omega')$, which by definition is a set of time- $(t + 1)$ actions (in the intended application of this definition, i.e. Axiom 4.11, such actions are mutually indifferent, but the following definition does not require this). Out of the time- $(t + 1)$ actions in $a(\omega')$, one picks a distinguished one $a_{+1,\omega'}$. Finally, one constructs a new action b available at time t that, for any state $\omega' \in \mathcal{F}_t(\omega)$, leads to the time- $(t + 1)$ tree containing the single initial action $a_{+1,\omega}$. Each possible choice of initial action a and subsequent actions $a_{+1,\omega'}$ leads to a different initial action b in the next-period commitment version of f . Formally:

Definition 6 Fix a tree $f \in F_t(\omega)$. The **next-period commitment version** of f is the tree

$$g = \left\{ b \in A_t(\omega) : \exists a \in f \text{ s.t. } \forall \omega' \in \mathcal{F}_t(\omega) \exists a_{+1, \omega'} \in a(\omega') \text{ s.t. } b(\omega') = \{a_{+1, \omega'}\} \right\}.$$

Now consider a tree f at a node (t, ω) and a history h consistent with f ; suppose that every action $a \in f(h)$, and every realization of the uncertainty $\omega' \in \mathcal{F}_t(\omega)$, leads to a new history where the DM is indifferent among all available actions. Then replacing the continuation tree $f(h)$ with its next-period-commitment version g must leave the DM indifferent at (t, ω) :

Axiom 4.11 (Weak Commitment) *For all $f \in F_t$ and all histories $h = [t, \omega, \mathbf{a}]$ consistent with f : if, for all $a \in f(h)$, all $\omega' \in \mathcal{F}_{t+|a|}(\omega)$, and all $a_{+1}, b_{+1} \in a(\omega')$, it is the case that $\{a_{+1}\} \sim_{t+|a|+1, \omega'} \{b_{+1}\}$, then $f \sim_{t, \omega} g_h f$, where g is the next-period commitment version of $f(h)$.*

Finally, Sophistication allows for the possibility that actions at *future* histories might be tempting for *future* preferences, even though they are unappealing for *initial* preferences (or vice versa). The following, standard axiom ensures that, by way of contrast, the availability of choices at the *initial* history of f that are deemed inferior given the same *initial* preference relation $\succsim_{t, \omega}$ is considered neither harmful (as might be the case if the DM were subject to temptation) nor beneficial (as it would be for a DM who has a preference for flexibility). This rules out deviations from standard behavior that are *not* due to differences in information and perceived ambiguity at distinct points in time:

Axiom 4.12 (Strategic Rationality) *For all $f, g \in F_t(\omega)$ such that $f \subset g$: if, for all $b \in f$ and $w \in g$, $\{b\} \succsim_{t, \omega} \{w\}$, then $f \sim_{t, \omega} g$.*

It is now possible to state the main result of this section. Two related characterizations of CP will be provided. The first is better suited to the analysis of specific preference models and updating rules (as in §4.4.1) and applications (as in §4.4.2). The second emphasizes that all behavioral implications of CP can be identified on the basis of prior preferences alone (as noted in the Introduction), and also has implications for policy evaluation (cf. Sec. 5.3).

Begin by specifying the DM's prior and conditional preferences over *plans*. Next, assume that this DM employs CP to determine her course of action in any given tree. Then, the DM's CPS should indicate *indifference between a tree f and any one of its CP solution(s)*. The following theorem shows that this is the case precisely when Axioms 4.10–4.12 hold.

Theorem 3 Consider a CPS $(\succ_{t,\omega})_{0 \leq t < T, \omega \in \Omega}$ such that, for every $t = 0, \dots, T$ and $\omega \in \Omega$, $\succ_{t,\omega}$ is a weak order on $F_t^p(\omega)$. The following statements are equivalent.

1. Every $\succ_{t,\omega}$ is a weak order on all of $F_t(\omega)$; furthermore, Axioms 4.10–4.12 hold.
2. for every node (t, ω) , every tree $f \in F_t(\omega)$, and every action $a \in \text{CP}_f([t, \omega, \emptyset])$: $f \sim_{t,\omega} \{a\}$.

Suppose instead that the axioms of this section are applied to the CPS $(\succ_{t,\omega}^0)_{0 \leq t < T, \omega \in \Omega}$ derived from the DM's prior preference \succ via Def. 4. In this case, Axioms 4.10–4.12 are effectively assumptions on the DM's *prior* preferences; formulating them in terms of the revealed conditional preferences $\succ_{t,\omega}^0$ is merely a matter of notational convenience. Leveraging Theorems 2 and 3, one then obtains

Theorem 4 Consider a weak order \succ on F_0 that satisfies Assumption 4.1 and Axioms 4.7–4.9, and the CPS $(\succ_{t,\omega}^0)_{0 \leq t < T, \omega \in \Omega}$ obtained from \succ via Def. 4. The following statements are equivalent.

1. Axioms 4.10–4.12 hold.
2. For every tree $f \in F_0$ and action $a \in \text{CP}_f([0, \omega, \emptyset])$: $f \sim \{a\}$.

4.4 Applications

4.4.1 Consistent Planning for MEU preferences and Prior-by-Prior Updating

To illustrate the results of Sec. 4.3.2, this subsection specializes Theorem 3 to the MEU decision model and prior-by-prior Bayesian updating, assuming reduction of plans to acts. It is straightforward to adapt the analysis to different representations of preferences and different updating rules (cf. e.g. Gilboa and Schmeidler [16], or Eichberger, Grant and Kelsey [6] and Horie [25]).

Begin by noting that, if $f \in F_t$ is a plan, every state ω determines a unique path through f : formally, for every $\omega' \in \mathcal{F}_t(\omega)$, there is a unique list of actions \mathbf{a} such that $[t, \omega', \mathbf{a}]$ is terminal and consistent with f . Throughout this subsection, for every node (t, ω) and plan $f \in F_t(\omega)$, the notation $f(\omega)$ indicates the prize $f(h)$, where $h = [t, \omega, \mathbf{a}]$ is the unique terminal history consistent with f . The required assumption on preferences can now be stated.

Assumption 4.2 (MEU) There exists a weak*-closed, convex set C of finitely-additive probabilities on (Ω, Σ) and a continuous function $u : X \rightarrow \mathbb{R}$ such that, for all plans $f, g \in F_0$,

$$f \succcurlyeq g \iff \min_{q \in C} \int_{\Omega} u(f(\omega))q(d\omega) \geq \min_{q \in C} \int_{\Omega} u(g(\omega))q(d\omega).$$

Moreover, (i) there exist plans $f, g \in F_0$ such that $f \succ g$; and (ii) for every node (t, ω) and all $q \in C$, $q(\mathcal{F}_t(\omega)) > 0$.

Note that the MEU decision rule is often seen as embodying “pessimistic” expectations; by contrast, Axiom 4.11 in §4.3.2 is “optimistic” about one’s ability to carry out ex-ante preferred courses of action (provided one does not have opposite strict preferences in the future).¹⁴

Part (i) of Assumption 4.2 states that ex-ante preferences over acts are not trivial. Part (ii) is a strengthening of the assumption that every conditioning event $\mathcal{F}_t(\omega)$ is not Savage-null; it ensures that prior-by-prior Bayesian updating is well-defined (cf. Pires [38], Prop. 1 and p. 150).

Assumption 4.2 pertains solely to *prior* preferences (over plans); Axiom 4.13 below provides a link with conditional preferences over plans, and in particular will be shown to characterize prior-by-prior updating. This axiom (see Siniscalchi [42]) recasts the main axiom in Pires [38] and Jaffray [26] in a form that is more easily compared with Axiom 4.1 (DC) of Sec. 4.1.¹⁵

Axiom 4.13 (Constant-act dynamic consistency) For all plans $p \in F_0^p$, prizes $x \in X$, and non-terminal histories $h = [0, \omega, \mathbf{a}]$ consistent with p :¹⁶

$$\begin{aligned} (p(h) \succcurlyeq_{|\mathbf{a}|, \omega} x) \wedge (\forall \omega' \notin \mathcal{F}_{|\mathbf{a}|}(\omega), p(\omega') \succcurlyeq x) &\implies p \succcurlyeq x, \\ (p(h) \succ_{|\mathbf{a}|, \omega} x) \wedge (\forall \omega' \notin \mathcal{F}_{|\mathbf{a}|}(\omega), p(\omega') \succcurlyeq x) &\implies p \succ x; \quad \text{and} \\ (p(h) \preccurlyeq_{|\mathbf{a}|, \omega} x) \wedge (\forall \omega' \notin \mathcal{F}_{|\mathbf{a}|}(\omega), p(\omega') \preccurlyeq x) &\implies p \preccurlyeq x, \\ (p(h) \prec_{|\mathbf{a}|, \omega} x) \wedge (\forall \omega' \notin \mathcal{F}_{|\mathbf{a}|}(\omega), p(\omega') \preccurlyeq x) &\implies p \prec x. \end{aligned}$$

Axiom 4.13 differs from Axiom DC in two respects. First, Axiom 4.13 only considers conditional comparisons between a plan p and a prize x ¹⁷, whereas DC has implications whenever two arbitrary plans are compared conditional on $\mathcal{F}_{|\mathbf{a}|}(\omega)$. Second, dominance, rather than conditional

¹⁴I thank a referee for this observation.

¹⁵For a non-decision-theoretic analysis, see Walley [47].

¹⁶Note that the history h reaches node $(|\mathbf{a}|, \omega)$; hence the notation $\mathcal{F}_{|\mathbf{a}|}(\omega)$, $\succ_{|\mathbf{a}|, \omega}$, etc.

¹⁷This is why the four cases $p(h) \succ_{t, \omega} x$, $p(h) \succ_{t, \omega} x$, $p(h) \preccurlyeq_{t, \omega} x$ and $p(h) \prec_{t, \omega} x$ must all be explicitly considered.

preference, is required outside of the conditioning event $\mathcal{F}_{|\mathbf{a}}(\omega)$. The motivations for these restrictions are discussed in the sources cited above (esp. [38] and [42]).

I now specialize the definition of CP to reflect the assumption that preferences over plans at a node (t, ω) are derived from an ex-ante MEU preference via prior-by-prior updating. Let u and C be as in Assumption 4.2; consider a tree $f \in F_t(\omega)$. For every terminal history $h = [t, \omega, \mathbf{a}]$ consistent with f , let $\text{CPMEU}_f(h) = \{f(h)\}$. Inductively, if $h = [t, \omega, \mathbf{a}]$ is consistent with f and $\text{CPMEU}_f([t, \omega', \mathbf{a} \cup a])$ has been defined for every $\omega' \in \mathcal{F}_{t+|\mathbf{a}}(\omega)$ and $a \in f(h)$, let

$$\begin{aligned} \text{CPMEU}_f^0(h) &= \left\{ \{b\} \subset A_{t+|\mathbf{a}}(\omega) : \exists a \in f(h) \text{ s.t. } \forall \omega' \in \mathcal{F}_{t+|\mathbf{a}}(\omega), \right. \\ &\quad \left. b(\omega') = p_{\omega'} \text{ for some } p_{\omega'} \in \text{CPMEU}_f([t, \omega', \mathbf{a} \cup a]) \right\} \text{ and} \\ \text{CPMEU}_f(h) &= \left\{ p \in \text{CPMEU}_f^0(h) : \forall p' \in \text{CPMEU}_f^0(h), \right. \\ &\quad \left. \min_{q \in C} \int_{\mathcal{F}_{t+|\mathbf{a}}(\omega)} u(p(\omega')) q(d\omega' | \mathcal{F}_{t+|\mathbf{a}}(\omega)) \geq \min_{q \in C} \int_{\mathcal{F}_{t+|\mathbf{a}}(\omega)} u(p'(\omega')) q(d\omega' | \mathcal{F}_{t+|\mathbf{a}}(\omega)) \right\}, \end{aligned}$$

Note that the assumption of prior-by-prior updating is embodied in the second line in the definition of $\text{CPMEU}_f(h)$. The counterpart to Theorem 3 can then be stated.

Theorem 5 *Consider a CPS $(\succsim_{t,\omega})_{0 \leq t < T, \omega \in \Omega}$. Suppose that Assumption 4.2 holds, and that every event $E \in \bigcup_{t=0}^T \mathcal{F}_t$ is non-null. Then the following statements are equivalent.*

1. *For every node (t, ω) , $\succsim_{t,\omega}$ is a weak order on F_t ; also, Axioms 4.10, 4.11, 4.12 and 4.13 hold;*
2. *for every node (t, ω) , tree $f \in F_t(\omega)$, and action $a \in \text{CPMEU}_f([t, \omega, \emptyset])$, $f \sim_{t,\omega} \{a\}$.*

Unlike Theorem 3, the above result (see statement 2 and the definition of CPMEU_f) has a specific implication for the way *plans* are evaluated at a node (t, ω) : prior-by-prior updating is employed. Again, this follows from Axiom 4.13, which appears in statement 1 of Theorem 5.

4.4.2 Sophistication and the Value of Information

This subsection analyzes a simple model of information acquisition and addresses the concern noted in the Introduction regarding the implications of CP: a sophisticated DM may reject

freely available information. I shall argue that this behavior reflects a basic *trade-off between information acquisition and commitment*; this trade-off is difficult to uncover when preferences over *acts* only are considered, but becomes transparent in the richer setting of this paper.¹⁸

Consider an individual facing a choice between two alternative actions, a and b (the term “action” is used informally here). Uncertainty is represented by a state space $\Omega = \Omega_1 \times \Omega_2$, where $\Omega_1 = \Omega_2 = \{\alpha, \beta\}$. The individual receives H dollars if she chooses action a and the second coordinate of the prevailing state is α , or if she chooses action b and the second coordinate of the prevailing state is β ; otherwise, she receives $L < H$ dollars. Finally, prior to choosing an action, the DM can observe the first coordinate of the prevailing state.

The DM has risk-neutral MEU preferences over acts, and evaluates plans by reduction. Her set of priors is $C = \{\lambda P + (1 - \lambda)Q : \lambda \in [0, 1]\}$, where $P, Q \in \Delta(\Omega)$ are defined by

$$P(\alpha, \alpha) = Q(\beta, \beta) = 1 - 2\epsilon, \quad P(\alpha, \beta) = P(\beta, \alpha) = \epsilon = Q(\alpha, \beta) = Q(\beta, \alpha), \quad P(\beta, \beta) = Q(\alpha, \alpha) = 0.$$

The parameter ϵ lies in the interval $(0, \frac{1}{4})$, and should be thought of as being “small”. In other words, this individual believes that the signal (ω_1) is most likely equal to the payoff-relevant component of the state (ω_2), but the relative likelihood of $\omega_2 = \alpha$ vs. $\omega_2 = \beta$ is ambiguous; furthermore, she assigns a (small and unambiguous) probability ϵ to each state where the signal is “wrong” (i.e. different from the payoff-relevant component). Finally, assume prior-by-prior updating. Note that the resulting conditional preferences over acts are dynamically inconsistent (they violate Axiom 4.1 in Sec. 4.1).

The objective is to determine the value of the information conveyed by the signal ω_1 ; this value turns out to depend upon whether or not the DM has the opportunity to *commit* to subsequent, ω_1 -contingent choices (e.g. by writing a binding contract, or delegating choices to an agent or machine). To adopt the formal framework of Sec. 3, it is useful to consider four *plans*, denoted p_{aa}, p_{ab}, p_{ba} and p_{bb} (the formal definitions are omitted for brevity). For instance, p_{ab} is the plan that prescribes the choice a after seeing $\omega_1 = \alpha$ and the choice b after observing $\omega_1 = \beta$; the DM evaluates it by “reducing” it to the act that yields H if $\omega \in \{(\alpha, \beta), (\beta, \alpha)\}$ and L

¹⁸Footnote 35 in Machina [34] attributes a similar observation, albeit expressed in the language of multiple selves, to Edi Karni.

elsewhere. Under the assumed preferences,

$$p_{ab} \succ p_{ba} \succ p_{aa} \sim p_{bb}. \quad (7)$$

If the individual acquires information and can commit, then she effectively faces the tree $f^{I,C} \equiv p_{aa} \cup p_{ab} \cup p_{ba} \cup p_{bb}$. Her preferred ex-ante choice is p_{ab} , so $f^{I,C} \sim p_{ab}$.

If the DM does not acquire information, her feasible choices are the plans p_{aa} and p_{bb} : thus, she can trivially “commit” to either a or b regardless of the realization of ω_1 , which she does not observe. Formally, she faces the tree $f^{NI} = p_{aa} \cup p_{bb}$. By Eq. (7), CP implies that $f^{NI} \sim p_{aa} \sim p_{bb}$. The value of the signal ω_1 under commitment is then the difference between the MEU evaluation of p_{ab} and that of p_{aa} , namely $(1 - 3\varepsilon)(H - L)$.

If the individual acquires information but cannot commit, then she faces a tree $f^{I,NC}$ wherein the choice of a vs. b is made *after* observing ω_1 .¹⁹ If the individual is *sophisticated* (as well as strategically rational), she will determine her willingness to pay for the information by taking into account the choices she will *actually* make after observing ω_1 : in other words, she will evaluate the tree $f^{I,NC}$ according to its CP solution.

Under prior-by-prior updating, one can verify that *the DM will strictly prefer b after observing $\omega_1 = \alpha$ and a after observing $\omega_1 = \beta$* ; therefore, by CP, $f^{I,NC} \sim p_{ba}$. The value of the signal ω_1 is then the difference between the MEU evaluations of p_{ba} and p_{aa} , namely $\varepsilon(H - L)$; since $\varepsilon \in (0, \frac{1}{4})$, this is positive, but smaller than in the commitment case.

To summarize, if the DM can exogenously commit, information is valuable, as usual: the DM has more options in the tree $f^{I,C}$ than in the tree f^{NI} (formally, $f^{NI} \subset f^{I,C}$) and this is of course beneficial. Furthermore, and symmetrically, if the DM “exogenously” receives information, then *commitment is also valuable*: it expands the *effective* choice set from just p_{ba} , the CP solution of $f^{I,NC}$, to $f^{I,C}$. Finally, *there is a trade-off between information and commitment*: the CP solution p_{ba} of $f^{I,NC}$ is not a subset or superset of f^{NI} , so one cannot say a priori whether this sophisticated but dynamically inconsistent DM should acquire information. For the preferences considered here, information *is* valuable; however, in other settings, the commitment problem may be so severe that the DM may rationally choose to pay a price so as to *avoid* in-

¹⁹Formally, $f^{I,NC} = \{c\}$, where the action c satisfies, for instance, $c(\alpha) = \{a_\alpha, b_\alpha\}$, with $a_\alpha, b_\alpha : \{\omega' : \omega'_1 = \alpha\} \rightarrow \{0, 1\}$, $a_\alpha(\omega') = 1$ if $\omega'_2 = \alpha$ and $a_\alpha(\omega') = 0$ otherwise, and similarly for b_α .

formation: for an interesting example, see Eichberger, Grant and Kelsey [6, p. 892]. Similar patterns of behavior also emerges in related contexts featuring time-inconsistent but sophisticated decision-makers: see e.g. Carrillo and Mariotti [3] and references therein.

5 Discussion of Theorems 2–4

5.1 Counterfactuals and conditional preferences²⁰

Any treatment of dynamic choice involves statements about preferences at potentially counterfactual decision points. In the tree of Fig. 1, the second node is not reached if the ball drawn is green; in such case, one cannot directly observe the DM’s preferences at that node. Consequently, substantive assumptions about conditional preferences are required.

This issue arises even with dynamically-consistent preferences (e.g. under EU). As noted in §4.1, one may employ Bayesian updating to *define* conditional preferences based on prior ones, thereby ensuring that DC holds per Proposition 1; however, the preferences thus defined need not be the DM’s “actual” conditional preferences. As noted in §4.1, one may equivalently assume that DC holds, and employ Bayesian updating to *elicit* conditional preferences; however, the DM’s “actual” conditional preferences may be dynamically inconsistent, in which case Bayesian updating elicits a spurious object. In other words, the Bayesian updating and DC assumptions may well be incorrect from a descriptive point of view.

Theorem 2 is subject to the same qualifications. Whether one views Def. 4 as a way to *define* or *elicit* conditional preferences, a substantive assumption about conditional preferences must be made: one either stipulates directly that $\succ_{t,\omega} \equiv \succ_{t,\omega}^0$ “by fiat,” or else stipulates that Axioms 4.3–4.6 holds, so beliefs are correct and hence $\succ_{t,\omega} \equiv \succ_{t,\omega}^0$, as Theorem 2 shows. Again, either of these substantive assumptions may be incorrect from a descriptive standpoint.

On the other hand, Theorem 4 can be “safely” interpreted as a behavioral characterization of CP in terms of the DM’s *prior* preferences over trees alone. The conjectural preferences $\succ_{t,\omega}^0$ can be interpreted as reflecting the DM’s prior beliefs about her future behavior, and Axioms 4.10–4.12 then ensure that such beliefs “support” or “explain” her ex-ante choices.

²⁰I thank the Coeditor and referees for several observations that guided and motivated this discussion.

5.2 An important caveat: Strong Sophistication

Recall that the Sophistication axiom has no implications in the case of indifferences at future nodes. *This is crucial to avoid unduly restricting preferences over plans.* If Axiom 4.10 is strengthened by replacing strict preferences at future nodes with weak preferences, one obtains

Axiom 5.1 (Strong Sophistication) *For all $f \in F_t$, all histories $h = [t, \omega, \mathbf{a}]$ with $\mathbf{a} \neq \emptyset$ consistent with f , and all $g \subset f(h)$: if, for all $b \in g$ and $w \in f(h) \setminus g$, $\{b\} \succsim_{t+|\mathbf{a}|, \omega} \{w\}$, then $f \sim_{t, \omega} g_h f$.*

Refer to the tree in Fig. 1, with $x = 1$ and notation as per Eqs. (3)–(5); consider the MEU preferences described in §4.3.1, so in particular $\{a\} \sim_{1, \alpha} \{b\}$ and $\{ca_1\} \succ \{cb_1\}$, and the history $h = [1, \alpha, c_1]$, i.e. the second decision point: under Strong Sophistication, $\{a\} \sim_{1, \alpha} \{b\}$ would imply that $\{ca_1\} = \{a\}_h \{c_1\} \sim \{c_1\} \sim \{b\}_h \{c_1\} = \{cb_1\}$, which is inconsistent with the preferences over acts (and plans) specified at the beginning of §4.3.1.

The example points out the key problematic implication of Strong Sophistication: it implies that, loosely speaking, if the DM is indifferent between two actions at a given history, she must also be indifferent between them at any earlier history. Furthermore, unlike the Sophistication axiom adopted here to characterize CP (i.e. Axiom 4.10), Strong Sophistication *does* impose restrictions on preferences over acts (or, more generally, plans).

Indeed, it turns out that these restrictions are overly strong for a very broad class of preferences. Say that preferences *admit certainty equivalents* if, for every (t, ω) and $p \in F_t^p(\omega)$ there is $x \in X$ with $p \sim_{t, \omega} x$; this includes virtually all parametric models of preferences over acts, assuming reduction. Also say that $\mathcal{F}_t(\omega)$ is *strongly non-null* (for \succsim) if, for all $x, y \in X$ and $p \in F_t^p(\omega)$, $x \succ y$ implies $x_{t, \omega} p \succ y_{t, \omega} p$.²¹

Proposition 6 *If each $\succsim_{t, \omega}$ is a weak order that admits certainty equivalents, each $\mathcal{F}_t(\omega)$ is strongly non-null, and Axioms 4.3 and 5.1 hold, then \succsim satisfies Axiom 4.2 (postulate P2), the restriction of each $\succsim_{t, \omega}$ to F_t^p is derived from \succsim via Bayesian updating, and Axiom 4.1 (DC) holds.*

The basic intuition behind this result is fairly straightforward (although the actual proof takes a different approach). Given a plan p and a history h consistent with p , if $x \in X$ is the certainty equivalent of $p(h)$ at h , then Strong Sophistication implies that, loosely speaking, x is also the

²¹For MEU preferences with priors C , this corresponds to $\min_{q \in C} q(\mathcal{F}_t(\omega)) > 0$.

“value” that the DM attaches to $p(h)$ at any history immediately preceding h . In other words, the DM can evaluate the plan p by *recursion*, as with standard EU preferences. This recursive structure implies dynamic consistency.

Proposition 6 thus implies that, in particular, *Axiom 5.1 can preclude ex-ante Ellsberg preferences* in the tree of Fig. 1.²² More broadly, Axiom 5.1 rules out precisely the kind of behavior that is the focus of the present paper. Axioms 4.10 and 4.11 are formulated so as to avoid this.

5.3 Miscellaneous

Policy evaluation. Welfare analysis is problematic in the presence of dynamic inconsistency. Refer to the tree in Fig. 1, with $x = 1$ and preferences as in the Introduction; consider a “policy” that removes action a . Suppose that an irreversible decision to implement the policy must be made at time 0; can a definite recommendation be made, despite the noted inconsistency of the DM’s preferences over acts? If the DM is sophisticated, then at the initial node she strictly prefers that a be removed, even though she anticipates being unhappy at the second node if a is indeed deleted. Gul and Pesendorfer [21, p.31] observe that “standard economic models identify choice with welfare”; from this point of view, the argument just given supports a policy that removes a (again, assuming an irreversible decision must be made at time 0).

The crux of the argument is that, in view of the results of this paper, there is no ambiguity as to whose “choice” and “welfare” one is concerned with: there is a single individual, characterized by her time-0 preferences over decision trees, who in particular strictly prefers the subtree with a removed to the tree f_1 . By way of contrast, if the decision problem is interpreted as a game played by multiple selves, it is no longer clear whose choices and welfare one should focus on; clear-cut policy prescriptions thus necessitate the introduction of an *exogenous* welfare criterion—perhaps one that trades off the well-being of the different selves.

Extensions. In the decision framework adopted in this paper, all trees are defined with respect to a fixed filtration $\mathcal{F}_0, \dots, \mathcal{F}_T$. However, if the DM holds well-defined preferences over *plans* conditional upon arbitrary (non-null) events, Theorem 3 suggests that she can also com-

²²In Sec. 2, modify the ex-ante MEU preferences so $C = \{q(\{\alpha\}) = \frac{1}{3}, q(\{\beta\}) \in [\epsilon, \frac{2}{3} - \epsilon]\}$ for some small $\epsilon > 0$: then $\{\alpha, \beta\}$ and $\{\gamma\}$ are strongly non-null and Prop. 6 applies. Of course, the analysis in Sec. 2 does not change.

pare two trees f and g that are each adapted to a *different* filtration. This is required, for instance, in order to model a DM who, at time 0, faces a choice between two different information structures, of which neither is finer than the other. Intuitively, the DM can apply CP to f and g separately, and then rank these trees according to her preferences over their respective CP solutions. By Theorem 3, this is equivalent to the assumption that the axioms in Sec. 4.3.2 apply “filtration-by-filtration”; the straightforward details are omitted.

The previous version of this paper, [43], provides results analogous to those in Sec. 4 for a more general class of decision trees that allows the agent’s actions at any point in time to determine the information she can receive at subsequent nodes. It also discusses the extension to a class of infinite decision trees.

6 Related Literature and Alternative Approaches

Kreps and Porteus [32]. As was noted above, the notation adopted here for decision trees is closely related to the formalization of a “decision problem” under risk in Sec. 3 of Kreps and Porteus [32, KP henceforth]. Specifically, a time- t tree in the sense of the present paper is a (finite) set of state-contingent menus of time- $(t + 1)$ trees, whereas in KP, a time- t decision problem is a (closed) set of lotteries over (contemporaneous payoffs and) time- $(t + 1)$ continuation problems. The key difference with this paper is that KP propose a model of *recursive* preferences, which satisfies dynamic consistency: see their central Axiom 3.1. By way of contrast, the present paper is concerned precisely with violations of dynamic consistency.

Segal [41]. Again in the setting of risky choice, Segal [41] studies preferences over two-stage lotteries that do not satisfy the Reduction axiom; he also allows for non-EU risk attitudes. As Segal states explicitly [41, p. 353], decisions are made only ex-ante in his framework (i.e. before first-stage lotteries are resolved); therefore, the issue of dynamic (in)consistency simply does not arise. By way of contrast, the present paper focuses on non-degenerate decision situations in which choices are made at two or more histories.

Karni and Safra [27, 28]. As noted in Sec. 2, these authors study economic applications of CP (which they call “behavioral consistency”) to choice under risk with non-EU preferences.

These papers employ CP as solution concept of a game played by agents, or selves, of the decision maker; by way of contrast, the present paper is concerned with the decision-theoretic foundations for CP.

The Menu Choice Literature and Gul and Pesendorfer [20]. The approach in the present paper is influenced by the menu-choice literature initiated by Kreps [31] and further developed by Dekel, Lipman and Rustichini [5] and Gul and Pesendorfer [19] in the context of certainty. The work of Epstein [11] and Epstein, Noor and Sandroni [12], which deals with non-Bayesian updating for EU preferences but does *not* allow for broader risk or ambiguity attitudes, and does not focus on CP, have already been mentioned (cf. footnote 4).

Gul and Pesendorfer [20] axiomatize a version of CP in the setting of dynamic choice under certainty. Their Theorems 1 and 2 are loosely related to Theorem 2 in §4.2 of this paper; they axiomatize ex-ante preferences that admit a “weak Strotz representation,” i.e., roughly speaking, a system of conditional preferences that generates it via CP. However, multiple conditional preference systems can generate the same ex-ante preferences (p. 437); by way of contrast, Theorems 2 and 4 relate ex-ante preferences to a *unique* CPS. Gul and Pesendorfer’s Theorem 3 ([20], p. 439) is the closest counterpart to Theorem 3 in this paper: it characterizes CP for a *given* time-0 preference on decision problems *and* a given collection of conditional choice correspondences.²³ However, their key axiom IRA corresponds to the Strong Sophistication axiom discussed in §5.2. As noted in §5.2, this is too strong an assumption for the purposes of this paper, as it implies a strong form of dynamic consistency.

Hammond [22] and related contributions. Hammond [22] also takes the DM’s behavior in decision trees as given. He proposes a notion of “consequentialism” that differs significantly from the one discussed in Sec. 3.3 (cf. [34], p. 1641); call this property “H-consequentialism.” Hammond’s main result shows that a behavioral rule satisfies H-consequentialism, consequentialism in the sense of Sec. 3.3, and continuity if and only if it is consistent with EU. In other words, Hammond emphasizes that assumptions about dynamic choice behavior can provide a *foundation* for (atemporal) EU preferences over acts. By way of contrast, the assumptions

²³In the statement of Theorem 3 in [20], “weak Strotz representation” should actually read “canonical Strotz representation,” as I have confirmed with the authors (the term “canonical” is formally defined in their paper).

on dynamic choice behavior considered in this paper are specifically designed *not* to restrict preferences *over acts* in any way. Rather, Hammond's result is related to Propositions 1 and 6 (loosely speaking, with H-consequentialism playing the role of FDC or Strong Sophistication).

For two-period decision problems under risk, Gul and Lantto [18] propose weakenings of H-consequentialism and dynamic consistency, and show that, under reduction of compound lotteries, the properties they propose are equivalent to the assumption that preferences over lotteries satisfy Dekel's betweenness axiom (Dekel [4]). Grant, Kaji and Polak [17] focus on value-of-information problems; they do not require reduction of two-stage lotteries, and assume suitable versions of DC and strong sophistication (cf. Sec. 5.2). They identify conditions on ex-ante preferences over two-stage lotteries that are necessary and sufficient for the DM to prefer more information to less. Thus, like Hammond [22], both papers relate dynamic choice behavior to properties of ex-ante preferences over lotteries; by way of contrast, the present paper does not impose or derive restrictions on ex-ante preferences over acts or plans.

Recursive preferences under ambiguity. In an influential paper, Epstein and Schneider [9] characterize the class of MEU preferences over acts that are *recursive*, and hence dynamically consistent, in all decision trees consistent with a given filtration. This permits the application of standard dynamic-programming techniques even in the presence of ambiguity. Maccheroni, Marinacci and Rustichini [33], and Klibanoff, Marinacci and Mukerji [29] adapt this approach to different preference models.

Epstein and Schneider's dynamic consistency requirement corresponds to Axiom 4.1 (DC) in §4.1; while this assumption does allow for non-trivial manifestations of ambiguity aversion at each decision node, by Prop. 1 it implies that prior preferences must satisfy Savage's Postulate P2 relative to every conditioning event in the filtration under consideration. In particular, as noted in §4.1, in the tree of Fig. 1, their requirement rules out the modal (ambiguity-averse) preferences at the initial node. By way of contrast, the approach in this paper does not restrict preferences over acts.

Non-consequentialist choice. An alternative approach to dynamic choice with non-EU preferences, advocated in the context of risky choice by Machina [34] and McClennen [35]

among others, essentially²⁴ allows conditional preferences at a history h to depend upon the “context” of the entire decision tree, so as to preserve optimality of the ex-ante optimal plan. Thus, this approach espouses dynamic consistency, but drops consequentialism. Hanany and Klibanoff [23, 24] implement this approach for a broad class of ambiguity-averse preferences.

When consequentialism is relaxed *in the presence of ambiguity*, interpreting the effect of *information* on preferences can be problematic. In particular, the *same* information may appear to eliminate ambiguity (or perception thereof) in one decision tree, and preserve it in another: I provide an example in Siniscalchi [44]. This conclusion stands in sharp contrast with the prevailing view of ambiguity as an informational phenomenon.

Al-Najjar and Weinstein [1] discuss further difficulties with violations of consequentialism.

A Appendix

A.1 Proof of Theorem 2 (Eliciting Conditional Preferences)

Remark A.1 Fix a node (t, ω) and let \geq be a weak order on $F_t(\omega)$ such that (i) for all $x, y \in X$, $x \geq y$ iff $x \succcurlyeq y$; (ii) if $x \in X$ satisfies $x \succcurlyeq f(h)$ [resp. $x \preccurlyeq f(h)$] for all terminal histories h consistent with $f \in F_t(\omega)$, then $x \geq f$ [resp. $f \geq x$]; and (iii) the sets $U = \{x \in X : x \geq f\}$ and $\{x \in X : f \geq x\}$ are closed in X . Then:

- (a) for every $f \in F_{t,\omega}$, there exists $x \in X$ such that $x \geq f$ and $f \geq x$ (abbreviated $x = f$);
- (b) if $f > g$, there is $x \in X$ such that $f > x > g$.

The proof of this remark is routine, hence omitted.

Turn to Theorem 2. For $f, f' \in F_t(\omega)$ and $g \in F_0^t$, write $f \succcurlyeq_{t,\omega|g}^0 f'$ to denote that Eq. (6) holds for g and for a suitable $z \in X$. Thus, $f \succcurlyeq_{t,\omega}^0 f'$ iff $f \succcurlyeq_{t,\omega|g}^0 f'$ for all $g \in F_0^t$.

Assume that (2) holds, and consider a node (t, ω) . Suppose that $f \succ_{t,\omega} f'$ and let $z \in X$ be such that $z \sim_{t,\omega} f'$: such a prize exists by Remark A.1. Fix $g \in F_0$ arbitrarily: I claim that $f \succcurlyeq_{t,\omega|g}^0 f'$, so $f \succcurlyeq_{t,\omega}^0 f'$. To see this, suppose first $y \succ z$, so $y \succ_{t,\omega} z$ by Axiom 4.3; then $y \succ_{t,\omega} f'$ by transitivity, and Axiom 4.6 implies that $(f' \cup y)_{t,\omega} g \sim y_{t,\omega} g$. Next, suppose that $y \prec z$: again invoking Axiom 4.3 and transitivity we get $y \prec_{t,\omega} f$, and Axiom 4.6 implies $(f \cup y)_{t,\omega} g \sim f_{t,\omega} g$.

²⁴This necessarily brief summary overlooks nuances among different proponents of this approach.

This proves the claim. In the opposite direction, consider $f, f' \in F_t(\omega)$ and suppose that $f \succ_{t,\omega}^0 f'$; let $z \in X$ be such that Eq. (6) holds for all $g \in F_0^p$. Suppose by contradiction that $z \succ_{t,\omega} f$, so there exist $y', y'' \in X$ such that $z \succ_{t,\omega} y' \succ_{t,\omega} y'' \succ_{t,\omega} f$ (by Remark A.1). Now Def. 4 implies $(f \cup y')_{t,\omega} g \sim f_{t,\omega} g \sim (f \cup y'')_{t,\omega} g$ for all $g \in F_0^p$, but Axiom 4.6 and the assumption that $\mathcal{F}_t(\omega)$ is not null imply $(f \cup y')_{t,\omega} g \sim y'_{t,\omega} g \succ y''_{t,\omega} g \sim (f \cup y'')_{t,\omega} g$ for some such g : contradiction. Hence, $f \succ_{t,\omega} z$; similarly, $z \succ_{t,\omega} f'$, and it follows that $f \succ_{t,\omega} f'$.

It remains to be shown that \succ satisfies Axioms 4.9 and 4.8 (Axiom 4.7 is immediately implied by Axioms 4.3 and 4.5). Consider first Axiom 4.9: fix a node (t, ω) , $f \in F_{t,\omega}$, $g, g' \in F_0^p$ and $x, y \in X$. For (i), suppose that $(f \cup y)_{t,\omega} g \not\sim f_{t,\omega} g$ and $x \succ y$. By Axiom 4.6, the first relation implies that $f \preceq_{t,\omega} y$, so by transitivity $f \preceq_{t,\omega} x$, and Axiom 4.6 implies that $(f \cup x)_{t,\omega} g' \sim x_{t,\omega} g'$. The argument for (ii) is similar. Finally, consider Axiom 4.8. If $f(h) \succ x$ for all terminal histories h consistent with f , then, since f is finite, there is $y \in X$ such that $y \succ x$ and $f(h) \succcurlyeq y$ for all terminal h . Now Axiom 4.4 implies $f \succ_{t,\omega} y$, and hence $f \succ_{t,\omega} x$; then Axiom 4.6 implies $(f \cup x)_{t,\omega} g \sim f_{t,\omega} g$, as required. The argument for (ii) is similar.

Now assume that (1) holds. To streamline the exposition, for any node (t, ω) and $f, f' \in F_{t,\omega}$ call any $z \in X$ with the properties in Def. 4 for all $g \in F_0^p$ a *cutoff* for $f \succ_{t,\omega}^0 f'$.

Claim 1: For every node (t, ω) , $\succ_{t,\omega}^0$ is transitive.

Consider $f, f', f'' \in F_{t,\omega}$ such that $f \succ_{t,\omega}^0 f'$ and $f' \succ_{t,\omega}^0 f''$, and let $x, x' \in X$ be the respective cutoffs (which, remember, must apply for all $g \in F_0^p$). Then it must be the case that $x \succ x'$; otherwise, consider $y', y'' \in X$ such that $x' \succ y' \succ y'' \succ x$ (which exist by Remark A.1): by Assumption 4.1, for some $g \in F_0^p$, $y'_{t,\omega} g \succ y''_{t,\omega} g$, and since $f \succ_{t,\omega|g}^0 f'$ must hold, we conclude that $(f' \cup y')_{t,\omega} g \sim y'_{t,\omega} g \succ y''_{t,\omega} g \sim (f' \cup y'')_{t,\omega} g$; but $f' \succ_{t,\omega|g}^0 f''$ must also hold, and it implies $(f' \cup y')_{t,\omega} g \sim f'_{t,\omega} g \sim (f' \cup y'')_{t,\omega} g$, so a contradiction results.

Now consider $y \in X$ and fix an arbitrary $g \in F_0^p$. If $y \succ x'$, then $f' \succ_{t,\omega|g}^0 f''$ implies $(y \cup f'')_{t,\omega} g \sim y_{t,\omega} g$; if instead $y \preceq x'$, then $y \preceq x$, and $f \succ_{t,\omega|g}^0 f'$ implies $(f \cup y)_{t,\omega} g \sim f_{t,\omega} g$. Hence, x' is a cutoff for $f \succ_{t,\omega}^0 f''$.

Claim 2: Fix a node (t, ω) and $x, y \in X$. Then $x \succcurlyeq y$ iff $x \succ_{t,\omega}^0 y$. In particular, $x \succ y$ implies $(x \cup y)_{t,\omega} g \sim x_{t,\omega} g$ for all $g \in F_0^p$.

Suppose $x \succcurlyeq y$ and fix an arbitrary $g \in F_0^t$. For all $x' \succ y$, Axiom 4.8 implies that $(x' \cup y)_{t,\omega} g \sim$

$x'_{t,\omega}g$; similarly, for all $x' \prec y$, also $x' \prec x$, and Axiom 4.8 implies $(x \cup x')_{t,\omega}g \sim x_{t,\omega}g$. Hence, y is a cutoff for $x \succcurlyeq_{t,\omega}^0 y$.

Conversely, suppose $x \succcurlyeq_{t,\omega}^0 y$ and let y' be a cutoff. If $y' \prec z \prec y$, then for any $g \in F_0^p$, Axiom 4.8 implies $(z \cup y)_{t,\omega}g \sim y_{t,\omega}g$, but Def. 4 requires $(z \cup y)_{t,\omega}g \sim z_{t,\omega}g$: since $\mathcal{F}_t(\omega)$ is non-null, this is a contradiction. Hence, $y' \succcurlyeq y$, and similarly $x \succcurlyeq y'$. By transitivity, $x \succcurlyeq y$.

Claim 3: fix a node (t, ω) , $f \in F_{t,\omega}$ and $x \in X$. Then either $f \succcurlyeq_{t,\omega}^0 x$ or $x \succcurlyeq_{t,\omega}^0 f$ (or both). In particular, if $x, x' \in X$ satisfy $x \succcurlyeq f(h) \succcurlyeq x'$ for all terminal histories h consistent with f , then $x \succcurlyeq_{t,\omega}^0 f$ and $f \succcurlyeq_{t,\omega}^0 x'$.

Suppose that it is not the case that $f \succcurlyeq_{t,\omega}^0 x$. Then in particular x is not a cutoff; by Claim 2, for all $y \succ x$, $(y \cup x)_{t,\omega}g \sim y_{t,\omega}g$ for all $g \in F_0^p$, so there must be $y \prec x$ and $g^* \in F_0^p$ such that $(f \cup y)_{t,\omega}g^* \not\sim f_{t,\omega}g^*$. Then Axiom 4.9 implies that, for all $y' \succ y$ and all $g \in F_0^p$, $(f \cup y')_{t,\omega}g \sim y'_{t,\omega}g$. On the other hand, for all $y' \prec y$, also $y' \prec x$, so Claim 2 implies $(x \cup y')_{t,\omega}g \sim x_{t,\omega}g$ for all $g \in F_0^p$. Hence, y' is a cutoff for $x \succcurlyeq_{t,\omega}^0 f$.

If x, x' are as above, then Axiom 4.8 implies that, for every $y \prec x'$ and $g \in F_0^p$, $(f \cup y)_{t,\omega}g \sim f_{t,\omega}g$, and Claim 2 implies that, for every $y \succ x'$ and $g \in F_0^p$, $(y \cup x')_{t,\omega}g \sim y_{t,\omega}g$. Thus, $f \succcurlyeq_{t,\omega}^0 x'$, and the other relation follows similarly.

Claim 4: fix a node (t, ω) and $f \in F_t(\omega)$. Then there exists $x \in X$ such that $x \sim_{t,\omega}^0 f$ (i.e. $x \succcurlyeq_{t,\omega}^0 f$ and $f \succcurlyeq_{t,\omega}^0 x$ both hold). Hence, $\succcurlyeq_{t,\omega}^0$ is complete on $F_t(\omega)$.

Let $L = \bigcap_{x: x \succcurlyeq_{t,\omega}^0 f} \{y : x \succcurlyeq y\}$. Notice that L is an intersection of closed sets by Axiom 4.7, and hence is closed. Also, the last part of Claim 3 shows that there always exists $x \in X$ such that $x \succcurlyeq_{t,\omega}^0 f$. Since $\succcurlyeq_{t,\omega}^0$ is transitive by Claim 1, if $f \succcurlyeq_{t,\omega}^0 y$, then $x \succcurlyeq_{t,\omega}^0 y$ (and hence $x \succcurlyeq y$) for every $x \in X$ such that $x \succcurlyeq_{t,\omega}^0 f$: thus, $f \succcurlyeq_{t,\omega}^0 y$ implies $y \in L$. On the other hand, suppose $f \not\succcurlyeq_{t,\omega}^0 y$: then in particular y cannot be a cutoff and, as in the proof of Claim 3, Claim 2 implies that there must exist $x \prec y$ and $g^* \in F_0^p$ such that $(f \cup x)_{t,\omega}g^* \not\sim f_{t,\omega}g^*$. Then Axiom 4.9 implies that, for all $x' \succ x$, and $g \in F_0^p$, $(f \cup x')_{t,\omega}g \sim x'_{t,\omega}g$; also, by Claim 2, for all $x' \prec x$ and g , $(x \cup x')_{t,\omega}g \sim x_{t,\omega}g$. Thus, x is a cutoff for $x \succcurlyeq_{t,\omega}^0 f$, and since $y \notin \{y' : x \succcurlyeq y'\}$, $y \notin L$. Thus, $L = \{y : f \succcurlyeq_{t,\omega}^0 y\}$; as noted above, this set is non-empty. Similarly, the set $U = \{y : y \succcurlyeq_{t,\omega}^0 f\}$ is non-empty and closed.

By Claim 3, $U \cup L = X$, so there exists $x \in U \cap L$, which by definition satisfies $x \sim_{t,\omega}^0 f$.

To complete the proof of Theorem 2, note first that $\succcurlyeq_{t,\omega}^0$ is complete and transitive on $F_t(\omega)$

by Claims 4 and 1 respectively; by Claim 2, it satisfies Axiom 4.3; by Claim 3, it satisfies Axiom 4.4; by Claim 4 and Axiom 4.7, it satisfies Axiom 4.5.

Finally, we verify that it also satisfies Axiom 4.6. Fix a node (t, ω) , $f \in F_t(\omega)$, and $x \in X$. Suppose $f \succ_{t, \omega}^0 x$; if $(f \cup x)_{t, \omega} g^* \not\sim f_{t, \omega} g^*$ for some $g^* \in F_0^p$, then Axiom 4.9 and Claim 2 imply that, for all $g \in F_0^p$, $y \succ x$ implies $(f \cup y)_{t, \omega} g \sim y_{t, \omega} g$ and $y \prec x$ implies $(x \cup y)_{t, \omega} g \sim x_{t, \omega} g$. Thus, by definition $x \succ_{t, \omega}^0 f$: contradiction. Similarly, suppose $x \succ_{t, \omega}^0 f$; if $(f \cup x)_{t, \omega} g^* \not\sim x_{t, \omega} g^*$ for some g^* , then for all $g \in F_0^p$, $y \prec x$ implies $(f \cup y)_{t, \omega} g \sim f_{t, \omega} g$ and $y \succ x$ implies $(x \cup y)_{t, \omega} g \sim x_{t, \omega} g$, i.e. x is a cutoff for $f \succ_{t, \omega}^0 x$: contradiction.

A.2 Proof of Theorems 3 and 4 (Consistent Planning)

Say that history $h = [t, \omega, \mathbf{a}]$ **precedes** history $h' = [t', \omega', \mathbf{a}']$ iff $t = t'$, $\mathcal{F}_{t+|\mathbf{a}|}(\omega) = \mathcal{F}_{t+|\mathbf{a}'|}(\omega')$, and either $\mathbf{a} = \emptyset$, or else $\mathbf{a} = (a_t, \dots, a_\tau)$ and $\mathbf{a}' = (a_t, \dots, a_\tau, a'_{\tau+1}, \dots, a'_{\tau+\tau'})$ for some $\tau' \geq 0$. In this case, write $h \leq_H h'$. The notation $h <_H h'$ means $h \leq_H h'$ and not $h' \leq_H h$; $h =_H h'$ instead means that $h \leq_H h'$ and $h' \leq_H h$. Observe that $h =_H h'$ iff $h = [t, \omega, \mathbf{a}]$ and $h' = [t, \omega', \mathbf{a}]$ for some $\omega' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega)$. Finally, if h, h' are consistent with f , then h **immediately precedes** h' , written $h <_H^* h'$, if $h <_H h'$ and there is no history h'' such that $h <_H h'' <_H h'$.

Begin with two preliminary remarks. First, the set of actions that CP associates with a given history h in a tree f only depends upon the continuation tree $f(h)$.

Remark A.2 For every node (t, ω) with $t < T$, tree $f \in F_t(\omega)$, and non-terminal history $h = [t, \omega', \mathbf{a}]$ consistent with f : $\text{CP}_f^0(h) = \text{CP}_{f(h)}^0([t + |\mathbf{a}|, \omega', \emptyset])$ and $\text{CP}_f(h) = \text{CP}_{f(h)}([t + |\mathbf{a}|, \omega', \emptyset])$.

Proof: Suppose first $t + |\mathbf{a}| = T - 1$. Note that, for every $a \in f(h)$ and $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, $[t, \omega'', \mathbf{a} \cup a]$ is consistent with f and $[t + |\mathbf{a}|, \omega'', \emptyset \cup a]$ is consistent with $f(h)$; furthermore, $f([t, \omega'', \mathbf{a} \cup a]) = a(\omega'') = (f(h))([t + |\mathbf{a}|, \omega'', \emptyset \cup a])$. Therefore, for every $a \in f(h)$ and $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, $\text{CP}_f([t, \omega'', \mathbf{a} \cup a]) = \{a(\omega'')\} = \text{CP}_{f(h)}([t + |\mathbf{a}|, \omega'', \emptyset \cup a])$. This immediately implies that $\text{CP}_f^0(h) = \text{CP}_{f(h)}^0([t + |\mathbf{a}|, \omega', \emptyset])$, and thus also $\text{CP}_f(h) = \text{CP}_{f(h)}([t + |\mathbf{a}|, \omega', \emptyset])$.

The induction step is immediate: if, for every $a \in f(h)$ and $\omega'' \in \mathcal{F}_t(\omega')$, it is the case that $\text{CP}_f([t, \omega'', \mathbf{a} \cup a]) = \text{CP}_{f(h)}([t + |\mathbf{a}|, \omega'', \emptyset \cup a])$, then Def. 5 readily implies that $\text{CP}_f^0(h) = \text{CP}_{f(h)}^0([t + |\mathbf{a}|, \omega', \emptyset])$, and thus also $\text{CP}_f(h) = \text{CP}_{f(h)}([t + |\mathbf{a}|, \omega', \emptyset])$. ■

Second, CP solutions are measurable with respect to the partitions $\mathcal{F}_0, \dots, \mathcal{F}_{T-1}$. Hence, only finitely many histories of any given tree need be considered in order to evaluate it.

Remark A.3 For every node (t, ω) with $t < T$, tree $f \in F_t(\omega)$, and history h consistent with f : if $h' =_H h$, then h' is consistent with f , and $f(h') = f(h)$. Consequently, $\text{CP}_f(h) = \text{CP}_f(h')$.

Proof: Since $h =_H h'$, one can write $h = [t, \omega', \mathbf{a}]$ and $h' = [t, \omega'', \mathbf{a}]$, with $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$ and $\omega' \in \mathcal{F}_t(\omega)$. Since $\mathcal{F}_{t+|\mathbf{a}|}(\omega') \subset \mathcal{F}_{\bar{t}}(\omega')$ for $\bar{t} = t, \dots, t + |\mathbf{a}| - 1$, it is the case that $\omega'' \in \mathcal{F}_{\bar{t}}(\omega')$ for such \bar{t} ; in particular, $\omega'' \in \mathcal{F}_t(\omega') = \mathcal{F}_t(\omega)$, so h' is consistent with f .

If $\mathbf{a} = \emptyset$, then $f(h') = f = f(h)$. Otherwise, let $\mathbf{a} = (a_t, \dots, a_\tau)$, with $a_\tau \in A_\tau(\omega')$; since $t + |\mathbf{a}| = t + (\tau - t + 1) = \tau + 1$, by assumption $\mathcal{F}_{\tau+1}(\omega'') = \mathcal{F}_{\tau+1}(\omega')$, and a_τ is $\mathcal{F}_{\tau+1}$ -measurable; hence, $f(h) = a_\tau(\omega') = a_\tau(\omega'') = f(h')$. The last implication follows from Remark A.2. ■

Next, a technical issue must be addressed. Recall from Sec. 3.2 that, if $h = [t, \omega, \mathbf{a}]$ is a history consistent with the tree f , with $\mathbf{a} = [a_t, \dots, a_\tau]$, the *composite tree* $g_h f$ is obtained by iteratively constructing actions $\bar{a}_t, \dots, \bar{a}_\tau$, and replacing the initial action a_t in f with \bar{a}_t . One consequence is that the history h itself is no longer consistent with $g_h f$: rather, in a natural sense, it corresponds to the history $[t, \omega, (\bar{a}_t, \dots, \bar{a}_\tau)]$. Indeed, *any* history $h' = [t, \omega', \mathbf{a}']$ consistent with f , with $\mathbf{a}' = (a'_t, \dots, a'_\tau) \neq \emptyset$ and $a'_t = a_t$, is no longer consistent with $g_h f$. Thus, it is necessary to construct the history corresponding to such h' in $g_h f$, and define notation for it.

To this end, continue to denote by $(\bar{a}_t, \dots, \bar{a}_\tau)$ the action sequence constructed in the definition of the composite act $g_h f$. Assume first that h' is not initial and $a'_t = a_t$. Let σ be the largest number in $\{t, \dots, \tau\}$ such that $a'_i = a_i$ and $\omega' \in \mathcal{F}_{\bar{i}}(\omega)$ for $\bar{i} = t, \dots, \sigma$. Then define the **image of h' in $g_h f$** , denoted $h'|g_h f$, as $[t, \omega', (\bar{a}_t, \dots, \bar{a}_\sigma, a'_{\sigma+1}, \dots, a'_\tau)]$. If instead h' is initial, or $a'_t \neq a_t$, then simply let $h'|g_h f = h'$.

While $h'|g_h f$ is formally defined for *all* histories h' consistent with f , it will *not* be a history (let alone one consistent with $g_h f$) in case h strictly precedes h' (except for $g = \{a'_{\tau+1}\}$). However, $h'|g_h f$ is a history consistent with $g_h f$ if h does not strictly precede h' (the case of interest in the proof of Theorem 3). The following remark establishes this and other useful facts:

Lemma 7 *Let h, h' be histories consistent with $f \in F_t$, with $h \not\prec_H h'$. Then:*

1. $h'|g_h f$ is a history consistent with $g_h f$;

2. if h'' is another history consistent with f and such that $h \not\leq_H h''$, then $h' \leq_H h''$ if and only if $h'|g_h f \leq_H h''|g_h f$;
3. if also $h' \neq_H h$, then there is a surjection $\alpha : f(h') \rightarrow (g_h f)(h'|g_h f)$ such that, if $a' \in f(h')$, $h' = [t, \omega', \mathbf{a}']$, and $h'|g_h f = [t, \omega', \mathbf{b}']$, then $[t, \omega'', \mathbf{a}' \cup \alpha(a')]|g_h f = [t, \omega'', \mathbf{b}' \cup \alpha(a')]$ for all $\omega'' \in \mathcal{F}_{t+|\mathbf{a}'|}(\omega')$.
4. if neither $h \leq_H h'$ nor $h' \leq_H h$, then $(g_h f)(h'|g_h f) = f(h')$.

Proof: Write $h = [t, \omega, \mathbf{a}]$ and let notation be as in the construction of $h'|g_h f$. The assumption that $h \not\leq h'$ rules out the possibility that $\sigma = \tau < \tau'$. Also, the first claim is immediate if $\sigma = \tau' \leq \tau$. Thus, assume that $\sigma < \min(\tau, \tau')$. It must be shown that $a'_{\sigma+1} \in \bar{a}_\sigma(\omega')$; there are two cases.

If $\omega' \notin \mathcal{F}_{\sigma+1}(\omega)$, then, according to the definition of $(\bar{a}_t, \dots, \bar{a}_\tau)$, $\bar{a}_\sigma(\omega') = a_\sigma(\omega')$, and by assumption $a'_{\sigma+1} \in a'_\sigma(\omega')$; furthermore, by definition $a'_\sigma = a_\sigma$. Thus, $a'_{\sigma+1} \in \bar{a}_\sigma(\omega')$.

If instead $\omega' \in \mathcal{F}_{\sigma+1}(\omega)$, then $\bar{a}_\sigma(\omega') = \bar{a}_\sigma(\omega) = \{\bar{a}_{\sigma+1}\} \cup (a_\sigma(\omega) \setminus \{a_{\sigma+1}\})$; furthermore, by definition $a'_{\sigma+1} \neq a_{\sigma+1}$ and $a'_\sigma = a_\sigma$. Because h' is a history, $a'_{\sigma+1} \in a'_\sigma(\omega') = a_\sigma(\omega') = a_\sigma(\omega)$; but since $a'_{\sigma+1} \neq a_{\sigma+1}$, it follows that $a'_{\sigma+1} \in a_\sigma(\omega) \setminus \{a_{\sigma+1}\}$, and therefore $a'_{\sigma+1} \in \bar{a}_\sigma(\omega) = \bar{a}_\sigma(\omega')$.

The second claim is immediate if h' is initial; otherwise, write $h' = [t, \omega', (a'_t, \dots, a'_{\tau'})]$ and $h'' = [t, \omega'', (a''_t, \dots, a''_{\tau''})]$, let σ' be the largest index in $\{t, \dots, \tau\}$ such that $a_{\bar{t}} = a'_{\bar{t}}$ and $\mathcal{F}_{\bar{t}}(\omega) = \mathcal{F}_{\bar{t}}(\omega')$ for $\bar{t} = t, \dots, \sigma$, and define σ'' analogously for h'' . Since $a'_{\bar{t}} = a''_{\bar{t}}$ for $\bar{t} = t, \dots, \tau'$, and $\mathcal{F}_{\tau'+1}(\omega') = \mathcal{F}_{\tau'+1}(\omega'')$, it must be the case that $\sigma' \leq \sigma''$; furthermore, the argument for the first claim indicates that $\sigma' \leq \tau'$. The second claim then follows immediately if $\sigma' = \tau'$; moreover, if $\sigma' < \tau'$, then we must also have $\sigma' = \sigma''$, because $a'_{\sigma'+1} = a''_{\sigma'+1}$ and $\mathcal{F}_{\sigma'+1}(\omega') = \mathcal{F}_{\sigma'+1}(\omega'')$; but then $h'|g_h f = [t, \omega', (\bar{a}_t, \dots, \bar{a}_\sigma, a'_{\sigma'+1}, \dots, a'_{\tau'})] \leq_H [t, \omega'', (\bar{a}_t, \dots, \bar{a}_\sigma, a'_{\sigma'+1}, \dots, a'_{\tau'}, a''_{\tau'+1}, \dots, a''_{\tau''})] = h''|g_h f$, as required. The converse implication follows by reversing the roles of f and $g_h f$, and correspondingly those of h', h'' and $h'|g_h f$ and $h''|g_h f$, because $f = [f(h)]_{h|g_h f}[g_h f]$.

For the third claim, again refer to the notation in the construction of $h'|g_h f$. If h' is not initial and $a'_t \neq a_t$, then $h'|g_h f = h'$, and furthermore $(g_h f)(h') = a'_{\tau'}(\omega') = f(h')$, so α can be taken to be the identity map. Next, if h' is initial, a surjection α with the required properties can be obtained by letting $\alpha(a) = a$ for every $a \in f(h') \setminus \{a_t\}$, and $\alpha(a_t) = \bar{a}_t$. Finally, if h' is not initial, let σ be as in the construction of $h'|g_h f$. As in the proof of the first claim, it cannot be the case that $\sigma = \tau < \tau'$, because $h \not\leq h'$, so either $\sigma < \min(\tau, \tau')$ or $\sigma = \tau' \leq \tau$. In the first case, the

last action in both h' and $h'|g_h f$ is $a'_{\tau'}$, so $(g_h f)(h'|g_h f) = a'_{\tau'}(\omega') = f'(h')$ and α can be taken to be the identity map. In the second case, the last action in $h'|g_h f$ is $\bar{a}_{\tau'}$, whereas the last action in h' is $a'_{\tau'} = a_{\tau'}$. There are two sub-cases: if $\omega' \in \mathcal{F}_{\tau'+1}(\omega)$, then we must have $\tau' < \tau$ because $h \neq_H h'$: then $(g_h f)(h'|g_h f) = \bar{a}_{\tau'}(\omega') = \{\bar{a}_{\tau'+1}\} \cup (a_{\tau'}(\omega') \setminus a_{\tau'+1})$, $f'(h') = a'_{\tau'}(\omega') = a_{\tau'}(\omega')$, and a suitable α is given by $\alpha(a) = a$ for $a \neq a_{\tau'+1}$, and $\alpha(a_{\tau'+1}) = \bar{a}_{\tau'+1}$. If instead $\omega' \notin \mathcal{F}_{\tau'+1}(\omega)$, then $(g_h f)(h'|g_h f) = \bar{a}_{\tau'}(\omega') = a_{\tau'}(\omega') = a'_{\tau'}(\omega') = f'(h')$, and α can again be taken to be the identity.

Finally, let h, h' be as in the fourth claim. Since h, h' are unranked by \leq_H , neither can be initial, so $\mathbf{a}' = [a'_{t'}, \dots, a'_{\tau'}]$; also, it also cannot be the case that $\omega' \in \mathcal{F}_{\tau'+1}(\omega)$, because otherwise $h' \leq_H h$. Then the proof of the third claim shows that $\alpha : f(h') \rightarrow (g_h f)(h'|g_h f)$ can be taken to be the identity map, and the result follows. ■

For **sufficiency**, assume (1) in Theorem 3. I will show that, for all (t, ω) and $f \in F_t(\omega)$,

$$f \sim_{t, \omega} \text{CP}_f([t, \omega, \emptyset]); \quad (8)$$

to interpret, recall that, for every history h consistent with f , $\text{CP}_f(h)$ is a set of acts, and hence can itself be viewed as a tree in $F_t(\omega)$. By Axiom 4.12 and the definition of $\text{CP}_f(h)$, for every $a \in \text{CP}_f([t, \omega, \emptyset])$, $\{a\} \sim_{t, \omega} \text{CP}_f([t, \omega, \emptyset])$; transitivity then implies that (ii) in Theorem 3 holds.

Fix a node (t, ω) and a tree $f \in F_t(\omega)$. Now construct a sequence f^0, \dots, f^N of trees by iteratively replacing continuation trees $f(h)$ with the corresponding CP solutions $\text{CP}_f(h)$. To do so, two issues must be addressed. First, if Ω is infinite, there are infinitely many histories consistent with f ; however, by Remark A.3, these can be partitioned into equivalence classes, each element of which yields the same continuation tree and set of CP solutions. Second, as the tree f is iteratively modified, one must keep track of the *image* of its histories in the modified trees; however, the notation developed above and in Lemma 7 makes this relatively straightforward.

To address the first issue, for every $\tau = t, \dots, T$, fix a collection $\Omega_\tau \subset \Omega$ such that, for every $E_\tau \in \mathcal{F}_\tau$, there is a unique $\omega(E_\tau) \in \Omega_\tau$ such that $\omega(E_\tau) \in E_\tau$. Then let H^0 be the collection of all non-terminal histories $h = [t, \bar{\omega}, \mathbf{a}]$ consistent with f and such that $\bar{\omega} \in \Omega_{t+|\mathbf{a}|}$ (the reason for the superscript 0 will be clear momentarily). Since $f(h)$ is finite for every history h consistent with f , and every Ω_τ is finite, the set H^0 is also finite.

Next, to address the second issue, define a sequence of iteratively modified trees as follows. First, let $f^0 = f$. Then, enumerate the elements of H^0 as $h^{0,1}, \dots, h^{0,N}$ in such a way that, for all $n, m \in \{0, \dots, N\}$, $n < m$ implies $h^{0,n} \not\prec_H h^{0,m}$: that is, since by construction $h^{0,n} \neq_H h^{0,m}$, either $h^{0,m}$ strictly precedes $h^{0,n}$, or the two histories are not ordered by the precedence relation.

The induction step consists of the following two sub-steps. Let $n > 0$ and assume that the tree f^{n-1} has already been defined, along with the collection $\{h^{n-1,1}, \dots, h^{n-1,N}\}$. Then:

- let $f^n = \text{CP}_f(h^{0,n})_{h^{n-1,n}} f^{n-1}$;
- for $m = 1, \dots, N$, let $h^{n,m} = h^{n-1,m} |_{g_{h^{n-1,n}}} f^{n-1}$ if $h^{n-1,n} \not\prec_H h^{n-1,m}$; else, let $h^{n,m} = h^{n-1,m}$.

To elaborate, the tree f^n is obtained from f^{n-1} , by replacing the current continuation at the history $h^{n-1,n}$, which intuitively corresponds to $h^{0,n}$, with the set of consistent-planning solutions of f at $h^{0,n}$. Then, for each history in f^{n-1} that does not strictly follow $h^{n-1,n}$, the image in f^n is constructed (formally, $h^{n-1,m}$ is defined for all $m = 1, \dots, N$, but the particular assignment chosen is irrelevant if $h^{n-1,n} < h^{n-1,m}$). Inductively, this ensures that the structure of actions and histories in f^n reflects that of the corresponding actions and histories in $f = f^0$:

Lemma 8 For all $\ell = 0, \dots, N$:

1. for all $n, m \in \{\ell, \dots, N\}$, $h^{\ell,n} \leq_H h^{\ell,m}$ iff $h^{0,n} \leq_H h^{0,m}$;
2. for all $n \in \{\ell + 1, \dots, N\}$, there is a surjection $\alpha^{\ell,n} : f(h^{0,n}) \rightarrow f^\ell(h^{\ell,n})$ such that, for all $a \in f(h^{0,n})$, $h^{0,n} = [t, \omega^n, \mathbf{a}^{0,n}]$, $h^{\ell,n} = [t, \omega^n, \mathbf{a}^{\ell,n}]$ and $h^{0,m} = [t, \omega^m, \mathbf{a}^{0,n} \cup a]$ for some $m \in \{1, \dots, n\}$ imply $h^{\ell,m} = [t, \omega^m, \mathbf{a}^{\ell,n} \cup \alpha^{\ell,n}(a)]$.

Proof: The first statement is obviously true for $\ell = 0$. Inductively, suppose it holds for some $\ell < N$ and consider $n, m \geq \ell + 1$. By the induction hypothesis, $h^{\ell,n} \leq_H h^{\ell,m}$ iff $h^{0,n} \leq_H h^{0,m}$; furthermore, by construction $h^{0,\ell+1} \not\prec_H h^{0,n}$ and $h^{0,\ell+1} \not\prec_H h^{0,m}$, so again by the induction hypothesis $h^{\ell,\ell+1} \not\prec_H h^{\ell,n}$ and $h^{\ell,\ell+1} \not\prec_H h^{\ell,m}$. Apply Lemma 7 Part 2 to conclude that $h^{\ell,n} \leq_H h^{\ell,m}$ iff $h^{\ell+1,n} \leq_H h^{\ell+1,m}$; the assertion then follows.

The second claim is trivially true for $\ell = 0$. Inductively, suppose it holds for some $\ell < N$ and consider $n > \ell + 1$ (if $\ell = N - 1$, there is nothing to show). Since $n > \ell$, the induction hypothesis yields a surjection $\alpha^{\ell,n} : f(h^{0,n}) \rightarrow f^\ell(h^{\ell,n})$ with the properties stated in the Lemma. By the

choice of ordering on H^0 , $h^{0,\ell+1} \not\prec_H h^{0,n}$; thus, by the first claim, $h^{\ell,\ell+1} \not\prec_H h^{\ell,n}$. Moreover, by assumption $n \neq \ell+1$, so also $h^{\ell,\ell+1} \neq_H h^{\ell,n}$. Part 3 of Lemma 7 then yields a surjection $\alpha : f^\ell(h^{\ell,n}) \rightarrow f^{\ell+1}(h^{\ell+1,n})$ such that $h^{\ell,m} = [t, \omega^m, \mathbf{a}^{\ell,n} \cup a]$ implies $h^{\ell+1,m} = [t, \omega^m, \mathbf{a}^{\ell+1,n} \cup \alpha(a)]$. Thus, fix $a \in f(h^{0,n})$, $h^{0,n} = [t, \omega^n, \mathbf{a}^{0,n}]$, $h^{\ell,n} = [t, \omega^n, \mathbf{a}^{\ell,n}]$, $h^{\ell+1,n} = [t, \omega^n, \mathbf{a}^{\ell+1,n}]$ and $h^{0,m} = [t, \omega^m, \mathbf{a}^{0,n} \cup a]$. Then $h^{\ell,m} = [t, \omega^m, \mathbf{a}^{\ell,n} \cup \alpha^{\ell,n}(a)]$ and therefore $h^{\ell+1,m} = [t, \omega^m, \mathbf{a}^{\ell+1,n} \cup \alpha(\alpha^{\ell,n}(a))]$. Thus, $\alpha^{\ell+1,n} = \alpha \circ \alpha^{\ell,n}$ has the required properties; furthermore, it is onto, as both α and $\alpha^{\ell,n}$ are. ■

Next, there is a unique $\omega^* \in \Omega_t \cap \mathcal{F}_t(\omega)$, and hence a unique initial history in H^0 ; since this history necessarily precedes every other history consistent with f , it must be $h^{0,N}$. Then, by construction, $h^{n,N} = h^{0,N} = [t, \omega^*, \emptyset]$ is the only initial history in f^n for all n : in particular, this is true for $n = N - 1$, so $\text{CP}_f(h^{0,N})_{h^{N-1,N}} f^{N-1} = \text{CP}_f(h^{0,N})$, i.e. $f^N = \text{CP}_f([t, \omega, \emptyset])$, the r.h.s. of Eq. (8). Thus, to prove sufficiency, it is enough to show that $f^{n-1} \sim_{t,\omega} f^n$ for all $n = 1, \dots, N$.

Thus, consider $n \in \{1, \dots, N\}$. By construction, the history $h^{n-1,n}$ is consistent with f^{n-1} and intuitively corresponds to $h^{0,n}$. It will now be shown that, at every history $h^{n-1,m}$ that immediately follows $h^{n-1,n}$ in f^{n-1} , the continuation tree $f^{n-1}(h^{n-1,m})$ is $\text{CP}_f(h^{0,m})$.

Lemma 9 *For every $m \in \{1, \dots, n\}$ such that $h^{0,n} <^*_H h^{0,m}$, $f^{n-1}(h^{n-1,m}) = \text{CP}_f(h^{0,m})$.*

Proof: It must be the case that $m < n$. I claim that, for $\ell = m+1, \dots, n-1$, and for $k = 0, \dots, \ell-1$, $h^{k,\ell}$ and $h^{k,m}$ are unordered, and furthermore $h^{k,k+1} \not\prec_H h^{k,\ell}$ for $\ell = m, \dots, n-1$.

To see this, consider first $k = 0$: for $\ell = m+1, \dots, n-1$, $h^{0,\ell} \not\prec_H h^{0,m}$, because $h^{0,n} <^*_H h^{0,m}$ by assumption. Furthermore, by construction $h^{0,m} \not\prec_H h^{0,\ell}$, so $h^{0,\ell}$ and $h^{0,m}$ are unordered. Finally, by construction $h^{0,1} \not\prec_H h^{0,\ell}$ for $\ell = m, \dots, n-1$.

Inductively, assume the claim is true for $k-1 < \ell-1$. By the inductive hypothesis, $h^{k-1,k} \not\prec_H h^{k-1,\ell}$ for $\ell = m, \dots, n-1$, and $h^{k-1,\ell}$ and $h^{k-1,m}$ are unordered for $\ell = m+1, \dots, n-1$. By Part 2 in Lemma 7, $h^{k,\ell}$ and $h^{k,m}$ are also unordered for $\ell = m+1, \dots, n-1$. Moreover, since $k < \ell$, $k+1 \leq \ell$, so $h^{0,k+1} \not\prec_H h^{0,\ell}$; the last part of the claim then follows from Lemma 8.

The claim implies in particular that, for $\ell = m+1, \dots, n-1$, $h^{\ell-1,\ell}$ and $h^{\ell-1,m}$ are unordered; Part 3 of Lemma 7 now implies that $f^\ell(h^{\ell,m}) = f^{\ell-1}(h^{\ell-1,m})$ for such ℓ . Therefore, $f^{n-1}(h^{n-1,m}) = f^m(h^{m,m})$, and the result follows from the construction of f^m . ■

Lemma 10 $CP_f^0(h^{0,n})$ is the next-period commitment version of $f^{n-1}(h^{n-1,n})$.

Proof: Write $h^{0,n} = [t, \omega^n, \mathbf{a}^{0,n}]$. Consider $b \in CP_f^0(h^{0,n})$; by definition, there exists $a \in f(h^{0,n})$ such that, for every $\omega' \in \mathcal{F}_{t+|\mathbf{a}^{0,n}|}(\omega^n)$, $b(\omega') = \{a_{+1,\omega'}\}$ for some $a_{+1,\omega'} \in CP_f([t, \omega', \mathbf{a}^{0,n} \cup a])$.

Now, for every $\omega' \in \mathcal{F}_{t+|\mathbf{a}^{0,n}|}(\omega^n)$, let $m(\omega') \in \{1, \dots, n\}$ be such that $h^{0,m(\omega')} =_H [t, \omega', \mathbf{a}^{0,n} \cup a]$. By Part 2 of Lemma 8, there is $\bar{a} \in f^{n-1}(h^{n-1,n})$ such that $h^{n-1,m(\omega')} = [t, \omega^{m(\omega')}, \mathbf{a}^{n-1,n} \cup \bar{a}]$ for all $\omega' \in \mathcal{F}_{t+|\mathbf{a}^{0,n}|}(\omega^n) = \mathcal{F}_{t+|\mathbf{a}^{n-1,n}|}(\omega^n)$. By Lemma 9 and Remark A.3, $f^{n-1}(h^{n-1,m(\omega')}) = CP_f(h^{0,m(\omega')}) = CP_f([t, \omega', \mathbf{a}^{0,n} \cup a])$. Conclude that, for every $b \in CP_f^0(h^{0,n})$, there exists $\bar{a} \in f^{n-1}(h^{n-1,n})$ such that, for all $\omega' \in \mathcal{F}_{t+|\mathbf{a}^{0,n}|}(\omega^n)$, there exists $\{a_{+1,\omega'}\} \in CP_f([t, \omega', \mathbf{a}^{0,n} \cup a]) = CP_f(h^{0,m}) = f^{n-1}(h^{n-1,m}) = \bar{a}(\omega')$ such that $b(\omega') = \{a_{+1,\omega'}\}$: that is, b is an element of the next-period commitment version of $f^{n-1}(h^{n-1,n})$.

In the opposite direction, let b be an element of the next-period commitment version of $f^{n-1,n}$, so there is $\bar{a} \in f^{n-1}(h^{n-1,n})$ such that, for all $\omega' \in \mathcal{F}_{t+|\mathbf{a}^{n-1,n}|}(\omega^n)$, there is $a_{+1,\omega'} \in \bar{a}(\omega')$ such that $b(\omega') = \{a_{+1,\omega'}\}$. As above, for every such ω' let $m(\omega')$ be such that $h^{n-1,m} =_H [t, \omega', \mathbf{a}^{n-1,n} \cup \bar{a}]$. By measurability, $\bar{a}(\omega') = \bar{a}(\omega^{m(\omega')})$, so by Lemma 9, $a_{+1,\omega'} \in CP_f(h^{0,m(\omega')})$; furthermore, by Lemma 8, there is $a \in f(h^{0,n})$ such that $a^{n-1,n}(a) = \bar{a}$ and $h^{0,m(\omega')} = [t, \omega^{m(\omega')}, \mathbf{a}^{0,n} \cup a] =_H [t, \omega', \mathbf{a}^{0,n} \cup a]$. Hence, for every $\omega' \in \mathcal{F}_{t+|\mathbf{a}^{0,n}|}(\omega^n)$, $b(\omega') = \{a_{+1,\omega'}\}$ for some $\{a_{+1,\omega'}\} \in CP_f([t, \omega', \mathbf{a}^{0,n} \cup a]) = CP_f(h^{0,m})$, where the equality follows from Remark A.3: thus, $b \in CP_f^0(h^{0,n})$. ■

The proof of sufficiency can now be completed. By Lemma 10, the set $CP_f^0(h^{0,n})$, viewed as a tree, is the next-period commitment version of $f^{n-1}(h^{n-1,n})$, so by Axiom 4.11, $f^{n-1} \sim_{t,\omega} CP_f^0(h^{0,n})_{h^{n-1,n}} f^{n-1}$. If now $n = N$, then $h^{N-1,N}$ is initial, so actually $f^{N-1} \sim_{t,\omega} CP_f^0(h^{0,n})$; furthermore, Axiom 4.12 implies that $CP_f^0(h^{0,n}) \sim_{t,\omega} \{b \in CP_f^0(h^{0,n}) : \forall a \in g, \{b\} \succ_{t,\omega} \{a\}\} = CP_f(h^{0,N}) = f^N$. If instead $n < N$, then Axiom 4.10 implies that $CP_f^0(h^{0,n})_{h^{n-1,n}} f^{n-1} \sim_{t,\omega} CP_f(h^{0,N})_{h^{n-1,n}} f^{n-1} = f^n$. Thus, in either case, $f^{n-1} \sim_{t,\omega} f^n$, as required.

For **necessity**, begin with a preliminary

Lemma 11 Consider a node (t, ω^*) with $t < T$, a tree $f \in F_t(\omega^*)$, and a non-terminal history $h = [t, \omega, \mathbf{a}]$ consistent with f . Then:

1. $CP_f^0(h) = \bigcup_{a \in f(h)} CP_{\{a\}}^0([t + |\mathbf{a}|, \omega, \emptyset])$.

2. If $g \in F_{t+|\mathbf{a}|}(\omega)$ is such that $\text{CP}_g([t+|\mathbf{a}|, \omega, \emptyset]) = \text{CP}_f(h)$, then $\text{CP}_{g_{hf}}([t, \omega, \emptyset]) = \text{CP}_f([t, \omega, \emptyset])$.

Proof: Claim 1 holds because $\text{CP}_f^0(h) = \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega, \emptyset])$ by Remark A.2 and Def. 5.

To prove the second claim, it will be shown that, for any history $h' = [t, \omega', \mathbf{a}']$ consistent with f and such that $t+|\mathbf{a}'| \in \{t, \dots, t+|\mathbf{a}|\}$, $\text{CP}_f(h') = \text{CP}_{g_{hf}}(h'|g_{hf})$. The claim is obviously true for $h' = h$; also, if $t+|\mathbf{a}'| = t+|\mathbf{a}|$ and $h \neq h'$, then neither $h \leq_H h'$ nor $h' \leq_H h$, so Part 4 of Lemma 7 implies that $f(h') = (g_{hf})(h'|g_{hf})$, and Remark A.2 then implies that $\text{CP}_f(h') = \text{CP}_{g_{hf}}(h'|g_{hf})$.

Now consider $h' = [t, \omega', \mathbf{a}']$ such that $t+|\mathbf{a}'| < t+|\mathbf{a}|$, and assume that the claim has been proved for all histories that immediately follow h' . Pick $b \in \text{CP}_f(h')$, and let $a \in f(h')$ be such that, for all $\omega'' \in \mathcal{F}_{t+|\mathbf{a}'|}(\omega')$, $b(\omega'') = \{a_{+1, \omega''}\} \subset \text{CP}_f([t, \omega'', \mathbf{a}' \cup a])$. Write $h'|g_{hf} = [t, \omega', \mathbf{b}']$. Since $h \not\leq_H h'$, Lemma 7 yields a surjection $\alpha : f(h') \rightarrow (g_{hf})(h'|g_{hf})$ such that, for all $\omega'' \in \mathcal{F}_{t+|\mathbf{a}'|}(\omega')$, $[t, \omega'', \mathbf{a}' \cup a]|g_{hf} = [t, \omega'', \mathbf{b}' \cup \alpha(a)]$. The induction hypothesis implies that, for all such ω'' , $\text{CP}_f([t, \omega'', \mathbf{a}' \cup a]) = \text{CP}_{g_{hf}}([t, \omega'', \mathbf{b}' \cup \alpha(a)])$. Therefore, $b \in \text{CP}_{g_{hf}}^0(h'|g_{hf})$.

Conversely, if $b \in \text{CP}_{g_{hf}}^0(h'|g_{hf})$, there is $\bar{a} \in (g_{hf})(h'|g_{hf})$ such that, for all $\omega'' \in \mathcal{F}_{t+|\mathbf{a}'|}(\omega')$, $b(\omega'') = \{a_{+1, \omega''}\} \subset \text{CP}_{g_{hf}}([t, \omega'', \mathbf{b}' \cup \bar{a}])$. Again by Lemma 7, there is $a \in f(h')$ such that $[t, \omega'', \mathbf{b}' \cup \bar{a}] = [t, \omega'', \mathbf{a}' \cup a]|g_{hf}$ for all such ω'' , and therefore, by the induction hypothesis, $\text{CP}_{g_{hf}}([t, \omega'', \mathbf{b}' \cup \bar{a}]) = \text{CP}_f([t, \omega'', \mathbf{a}' \cup a])$. Therefore, $b \in \text{CP}_f^0(h')$.

Thus, $\text{CP}_f^0(h') = \text{CP}_{g_{hf}}^0(h'|g_{hf})$, which implies that $\text{CP}_f(h') = \text{CP}_{g_{hf}}(h'|g_{hf})$, as required. ■

Now assume that (2) holds in Theorem 3. Since each $\succ_{t, \omega}$ is complete and transitive on $F_t^p(\omega)$, and $\text{CP}_f([t, \omega, \emptyset]) \neq \emptyset$ for all $f \in F_t(\omega)$, $\succ_{t, \omega}$ is also complete and transitive on all of $F_t(\omega)$. Next, the three axioms in (1) will be considered in turn.

Axiom 4.10, Sophistication. Let $f \in F_t(\omega)$ and fix a history $h = [t, \omega', \mathbf{a}]$ consistent with f that is neither initial nor terminal, and finally $g \subset f(h)$ as in the Axiom. By the first claim in Lemma 11, $\text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset]) \subset \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset])$. Now fix $b^{\text{CP}} \in \text{CP}_g([t+|\mathbf{a}|, \omega', \emptyset])$ and $w^{\text{CP}} \in \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega', \emptyset])$. By the definition of Consistent Planning and Lemma 11, there are $b, w \in f(h)$ such that $b \in g$, $b^{\text{CP}} \in \text{CP}_{\{b\}}^0([t+|\mathbf{a}|, \omega', \emptyset])$, and $w^{\text{CP}} \in \text{CP}_{\{w\}}^0([t+|\mathbf{a}|, \omega', \emptyset])$.

Suppose that $w \in f(h) \setminus g$: then, by the assumption in the Axiom, $\{b\} \succ_{t+|\mathbf{a}|, \omega'} \{w\}$, and by (2) in Theorem 3, $\{b^{\text{CP}}\} \succ_{t+|\mathbf{a}|, \omega'} \{w^{\text{CP}}\}$. But since $b^{\text{CP}} \in \text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset]) \subset \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset])$ and $w^{\text{CP}} \in \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega', \emptyset])$, $\{w^{\text{CP}}\} \succ_{t+|\mathbf{a}|, \omega'} \{b^{\text{CP}}\}$: contradiction.

It follows that $w \in g$. That is: for every $w^{\text{CP}} \in \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega', \emptyset])$, $w^{\text{CP}} \in \text{CP}_{\{w\}}^0([t+|\mathbf{a}|, \omega', \emptyset]) \subset \text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset])$. Since, furthermore, $w^{\text{CP}} \succ_{t+|\mathbf{a}|, \omega'} a^{\text{CP}}$ for all $a^{\text{CP}} \in \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset])$, hence a fortiori for all $a^{\text{CP}} \in \text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset])$, conclude that $w^{\text{CP}} \in \text{CP}_g([t+|\mathbf{a}|, \omega', \emptyset])$: that is, $\text{CP}_{f(h)}([t+|\mathbf{a}|, \omega', \emptyset]) \subset \text{CP}_g([t+|\mathbf{a}|, \omega', \emptyset])$.

Conversely, if $b^{\text{CP}} \in \text{CP}_g([t+|\mathbf{a}|, \omega', \emptyset])$, then $b^{\text{CP}} \in \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset])$, and the argument just given implies that also $b^{\text{CP}} \sim_{t+|\mathbf{a}|, \omega'} w^{\text{CP}}$ for any $w^{\text{CP}} \in \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega', \emptyset])$: thus, it is also the case that $\text{CP}_g([t+|\mathbf{a}|, \omega', \emptyset]) \subset \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega', \emptyset])$.

Part 2 of Lemma 11 then implies that $\text{CP}_f([t, \omega, \emptyset]) = \text{CP}_{g_h f}([t, \omega, \emptyset])$, and the definition of $\succ_{t, \omega}$ in (ii) of the Theorem then implies that $f \sim_{t, \omega} g_h f$, as required.

Axiom 4.11, Weak Commitment. Let $f \in F_t(\omega)$ and fix a history $h = [t, \omega', \mathbf{a}]$ with the properties indicated in the Axiom. In particular, if $a \in f(h)$, $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, and $b_{+1}, b'_{+1} \in a(\omega'') = f([t, \omega'', \mathbf{a} \cup a])$, then $\{b_{+1}\} \sim_{t+|\mathbf{a}|+1, \omega''} \{b'_{+1}\}$; by (ii) in Theorem 3, for all $a_{+1} \in \text{CP}_{\{b_{+1}\}}([t+|\mathbf{a}|+1, \omega'', \emptyset])$ and $a'_{+1} \in \text{CP}_{\{b'_{+1}\}}([t+|\mathbf{a}|+1, \omega'', \emptyset])$, $\{a_{+1}\} \sim_{t+|\mathbf{a}|+1, \omega''} \{a'_{+1}\}$. By Part 1 of Lemma 11,

$$\text{CP}_{a(\omega'')}([t+|\mathbf{a}|+1, \omega'', \emptyset]) = \bigcup_{b_{+1} \in a(\omega'')} \text{CP}_{\{b_{+1}\}}([t+|\mathbf{a}|+1, \omega'', \emptyset]).$$

Now let g be the next-period commitment version of $f(h)$; I claim that $\text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset]) = \text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset])$: by Part 2 of Lemma 11, this implies that $\text{CP}_f([t, \omega, \emptyset]) = \text{CP}_{g_h f}([t, \omega, \emptyset])$ and hence, by (ii) of Theorem 3, $f \sim_{t, \omega} g_h f$, as required.

Fix $a^0 \in \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset])$ and let $a \in f(h)$ be such that, for every $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, $a^0(\omega'') = \{a_{+1, \omega''}\}$ for some $a_{+1, \omega''} \in \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega'', \emptyset \cup a]) = \text{CP}_{a(\omega'')}([t+|\mathbf{a}|+1, \omega'', \emptyset])$. Then, by the above argument, $a_{+1, \omega''} \in \text{CP}_{\{b_{+1, \omega''}\}}([t+|\mathbf{a}|+1, \omega'', \emptyset])$ for some $b_{+1, \omega''} \in a(\omega'')$. Now let $b \in A_{t+|\mathbf{a}|}(\omega')$ be such that, for all $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, $b(\omega'') = \{b_{+1, \omega''}\}$; then $b \in g$, and furthermore a^0 satisfies the following property: for every $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, $a^0(\omega'') = \{a_{+1, \omega''}\}$ for some $a_{+1, \omega''} \in \text{CP}_g([t+|\mathbf{a}|, \omega'', \emptyset \cup b]) = \text{CP}_{\{b_{+1, \omega''}\}}([t+|\mathbf{a}|+1, \omega'', \emptyset])$, where the equality follows from Remark A.2 and the fact that $g([t+|\mathbf{a}|, \omega'', \emptyset \cup b]) = b(\omega'') = \{b_{+1, \omega''}\}$. Therefore, $a^0 \in \text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset])$.

Conversely, suppose $a^0 \in \text{CP}_g^0([t+|\mathbf{a}|, \omega', \emptyset])$, so there is $b \in g$ such that, for every $\omega'' \in \mathcal{F}_{t+|\mathbf{a}|}(\omega')$, $a^0(\omega'') = \{a_{+1, \omega''}\}$ for some $a_{+1, \omega''} \in \text{CP}_g([t+|\mathbf{a}|, \omega'', \emptyset \cup b])$. But by the definition of g , there is $a \in f(h)$ such that, for all such ω'' , $b(\omega'') = \{b_{+1, \omega''}\}$ for some $b_{+1, \omega''} \in a(\omega'')$. Hence, $\text{CP}_g([t+|\mathbf{a}|, \omega'', \emptyset \cup b]) = \text{CP}_{\{b_{+1, \omega''}\}}([t+|\mathbf{a}|+1, \omega'', \emptyset]) \subset \text{CP}_{a(\omega'')}([t+|\mathbf{a}|+1, \omega'', \emptyset]) = \text{CP}_{f(h)}([t+|\mathbf{a}|, \omega'', \emptyset \cup a])$. It follows that $a^0 \in \text{CP}_{f(h)}^0([t+|\mathbf{a}|, \omega', \emptyset])$, as claimed.

Axiom 4.12, Strategic Rationality. Let f and g be as in the Axiom. Arguing as for Sophistication, Lemma 11 implies that $\text{CP}_g^0([t, \omega, \emptyset]) \subset \text{CP}_f^0([t, \omega, \emptyset])$. Fix $b^{\text{CP}} \in \text{CP}_g([t, \omega, \emptyset])$ and $w^{\text{CP}} \in \text{CP}_f([t, \omega, \emptyset])$. By the definition of Consistent Planning and Lemma 11, there are $b, w \in f$ such that $b \in g$, $b^{\text{CP}} \in \text{CP}_{\{b\}}^0([t, \omega, \emptyset])$, and $w^{\text{CP}} \in \text{CP}_{\{w\}}^0([t, \omega, \emptyset])$.

Suppose that $w \in f \setminus g$: then, by the assumption in the Axiom, $\{b\} \succ_{t, \omega} \{w\}$, and by the way $\succ_{t, \omega}$ is defined in (ii) of Theorem 3, $\{b^{\text{CP}}\} \succ_{t, \omega} \{w^{\text{CP}}\}$. Since $b^{\text{CP}} \in \text{CP}_g^0([t, \omega, \emptyset]) \subset \text{CP}_f^0([t, \omega, \emptyset])$ and $w^{\text{CP}} \in \text{CP}_f([t, \omega, \emptyset])$, $\{w^{\text{CP}}\} \succ_{t, \omega} \{b^{\text{CP}}\}$: thus, $\{w^{\text{CP}}\} \sim_{t, \omega} \{b^{\text{CP}}\}$. If instead $w \in g$, then $w^{\text{CP}} \in \text{CP}_g^0([t, \omega, \emptyset])$, so $\{b^{\text{CP}}\} \succ_{t, \omega} \{w^{\text{CP}}\}$; but since $\{w^{\text{CP}}\} \succ_{t, \omega} \{b^{\text{CP}}\}$ as well, again $\{w^{\text{CP}}\} \sim_{t, \omega} \{b^{\text{CP}}\}$. The definition of $\succ_{t, \omega}$ in (2) of Theorem 3 now implies that $f \sim_{t, \omega} g$, as required.

A.2.1 Theorem 4

The result follows readily from Theorems 2 and 3, except that, to show that (2) implies (1), we must show that (2) in the present Theorem implies (2) in Theorem 3. In doing so, since \succ satisfies Axioms 4.7–4.9, Theorem 2 implies that $(\succ_{t, \omega}^0)$ satisfies Axioms 4.3–4.6 and each relation is a weak order; we will make use of this fact. Thus, assume that (2) holds, fix a node (t, ω) , $f \in F_t(\omega)$ and $a \in \text{CP}_f([t, \omega, \emptyset])$; we apply Def. 4 to show that $f \sim_{t, \omega}^0 \{a\}$.

As in the proof of Theorem 2, there exists $z \in X$ such that $z \sim_{t, \omega}^0 \{a\}$. Fix $g \in F_0^p$ and choose the unique sequence of t actions \mathbf{a} in g such that $h = [0, \omega, \mathbf{a}]$ is a history of g . Note that, for any $y \in X$, part 1 of Lemma 11 implies that $\text{CP}_{(f \cup y)_{t, \omega} g}^0(h) = \text{CP}_{f \cup y}^0([t, \omega, \emptyset]) = \{a_{t, \omega}^y\} \cup \text{CP}_f^0([t, \omega, \emptyset])$, where we explicitly represent y as the plan $f_{t, \omega}^y = \{a_{t, \omega}^y\} \in F_t^p(\omega)$. Similarly, $\text{CP}_{(\{a\} \cup y)_{t, \omega} g}^0(h) = \text{CP}_{\{a\} \cup y}^0([t, \omega, \emptyset]) = \{a, a_{t, \omega}^y\}$, because $\{a\}$ is a plan. It is also clear that $\text{CP}_y([t, \omega, \emptyset]) = \text{CP}_{f_{t, \omega}^y}([t, \omega, \emptyset]) = \{a_{t, \omega}^y\}$ and $\text{CP}_{\{a\}}([t, \omega, \emptyset]) = \{a\}$. I shall use these calculations below without explicit notice.

First, we claim that $\{a\} \succ_{t, \omega | g}^0 f$, using z as cutoff. If $y \succ z$, so also $y \succ_{t, \omega}^0 z$ by Axiom 4.3, by transitivity $y = \{a_{t, \omega}^y\} \succ_{t, \omega}^0 \{a\} \sim_{t, \omega}^0 \{a'\}$ for every $a' \in \text{CP}_f([t, \omega, \emptyset])$, and hence $y \succ_{t, \omega}^0 \{a''\}$ for every $a'' \in \text{CP}_f([t, \omega, \emptyset])$: therefore, $\text{CP}_{(f \cup y)_{t, \omega} g}(h) = \{a_{t, \omega}^y\} = \text{CP}_{f_{t, \omega}^y}([t, \omega, \emptyset])$. Then, part 2 of Lemma 11 implies that $\text{CP}_{(f \cup y)_{t, \omega} g}([0, \omega, \emptyset]) = \text{CP}_{y_{t, \omega} g}([0, \omega, \emptyset])$, and so by (2) in Theorem 4 also $(f \cup y)_{t, \omega} g \sim y_{t, \omega} g$. If instead $y \prec z$, then $\{a\} \succ_{t, \omega}^0 y$; then $\text{CP}_{(\{a\} \cup y)}([t, \omega, \emptyset]) = \{a\} = \text{CP}_{\{a\}}([t, \omega, \emptyset])$, and therefore $\text{CP}_{(\{a\} \cup y)_{t, \omega} g}([0, \omega, \emptyset]) = \text{CP}_{\{a\}_{t, \omega} g}([0, \omega, \emptyset])$ by part 2 of Lemma 11: thus, $(\{a\} \cup y)_{t, \omega} g \sim \{a\}_{t, \omega} g$. This proves the claim.

Next, we claim that $f \succ_{t, \omega | g}^0 \{a\}$. If $y \succ z$, then $y \succ_{t, \omega}^0 \{a\}$. Then $\text{CP}_{\{a\} \cup y}([t, \omega, \emptyset]) = \{a_{t, \omega}^y\} =$

$CP_y([t, \omega, \emptyset])$, and so $(\{a\} \cup y)_{t, \omega} g \sim y_{t, \omega} g$ by part 2 of Lemma 11 and (2) of Theorem 4. If instead $y \prec z$, then also $\{a\} \succ_{t, \omega}^0 y$, and hence $\{a'\} \succ_{t, \omega}^0 y$ for all $a' \in CP_f([t, \omega, \emptyset])$. Then $CP_{f \cup y}([t, \omega, \emptyset]) = CP_f([t, \omega, \emptyset])$. Part 2 of Lemma 11 now implies that $CP_{(f \cup y)_{t, \omega} g}([0, \omega, \emptyset]) = CP_{f_{t, \omega} g}([0, \omega, \emptyset])$, and so $(f \cup y)_{t, \omega} g \sim f_{t, \omega} g$. Thus, $f \sim_{t, \omega | g}^0 \{a\}$ for all $g \in F_0^p$, which concludes the proof.

A.3 Other results

Proof of Theorem 5 Recall that all relevant conditional preferences are well-defined. Moreover, Axiom 4.13 holds *iff* for all nodes (t, ω) , prizes $x \in X$, and plans $p \in F_t^p(\omega)$, $p \succ_{t, \omega} x$ if and only if $p_{t, \omega} x \succ x$. The ‘only if’ direction is immediate; for the converse, let p, x, h be as in the Axiom; by assumption $p \succ_{t, \omega} x$ implies $p_{t, \omega} x \succ x$, and for all $\omega' \notin \mathcal{F}_t(\omega)$, $p(\omega') \succ x$: thus, by monotonicity of MEU preferences, $p \succ p_{t, \omega} x \succ x$. The other cases are similar.

Now suppose (1) holds. This implies that Axiom A9 in Pires [38] holds, and the results therein imply that $\succ_{t, \omega}$ is derived from \succ via prior-by-prior Bayesian updating. Hence, CPMEU and CP coincide, and (2) follows from Theorem 3. Conversely, assume that (2) holds. Consider a plan $p \in F_t^0(\omega)$ and a prize $x \in X$ such that $u(x) = \min_{q \in C} \int_{t, \omega} u(p(\omega)) q(d\omega | E)$; one such prize must exist because X is connected and u is continuous. Now consider the tree $(p \cup x) \in F_t(\omega)$; clearly, $CPMEU_{(p \cup x)}(\emptyset)$ is precisely the set containing p and x ; by (2), we have $p \sim_{t, \omega} (p \cup x) \sim_{t, \omega} x$, i.e. $p \sim_{t, \omega} x$. Thus, $\succ_{t, \omega}$ is consistent with MEU and prior-by-prior Bayesian updating of C . This implies that CPMEU and CP coincide, so Theorem 3 ensures that each preference is complete and transitive, and that Axioms 4.10, 4.11 and 4.12 hold. Finally, prior-by-prior updating implies that the above restatement of Axiom 4.13 holds. ■

Proof of Proposition 6: By Axiom 5.1, for all (t, ω) , all $r, s \in F_t^p(\omega)$, and all $p \in F_0^p$, $r \sim_{t, \omega} s$ implies $r_{t, \omega} p \sim (r \cup s)_{t, \omega} p \sim s_{t, \omega} p$. Thus, by Axiom 4.3, for all $x, y \in X$ and $p \in F_0^p$, $x \sim y$ implies $x_{t, \omega} p \sim y_{t, \omega} p$ (for $x = y$, the claim is true by reflexivity), and since $\mathcal{F}_t(\omega)$ is strongly non-null, $x \succ y$ implies $x_{t, \omega} p \succ y_{t, \omega} p$. By assumption, for all $r \in F_t^p(\omega)$ there is $x \in X$ such that $x \sim_{t, \omega} r$, and so also $x_{t, \omega} p \sim r_{t, \omega} p$ for all $p \in F_0^p$.

Now suppose that, for r, s, p, q as in Postulate P2, it is the case that $r_{t, \omega} p \succ s_{t, \omega} p$. Let $x, y \in X$ be such that $x \sim_{t, \omega} r$ and $y \sim_{t, \omega} s$. If $r \prec_{t, \omega} s$, then $x \prec y$ and so $r_{t, \omega} p \sim x_{t, \omega} p \prec y_{t, \omega} p \sim s_{t, \omega} p$, a

contradiction: thus, $r \succ_{t,\omega} s$. Therefore, $x \succ y$, and so also $r_{t,\omega}q \sim x_{t,\omega}q \succ y_{t,\omega}q \sim s_{t,\omega}q$. Thus, P2 holds; furthermore, this argument also shows that the restriction of $\succ_{t,\omega}$ to $F_t^p(\omega)$ is derived from \succ via Bayesian updating. ■

References

- [1] N.I. Al-Najjar and J. Weinstein. The ambiguity aversion literature: a critical assessment. *Economics and Philosophy*, 25(03):249–284, 2009.
- [2] Maurice Allais. Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine. *Econometrica*, 21:503–546, 1953.
- [3] Juan D. Carrillo and T. Mariotti. Strategic ignorance as a self-disciplining device. *Review of Economic Studies*, 67(3):529–544, July 2002.
- [4] Eddie Dekel. An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom. *Journal of Economic Theory*, 40:304–318, 1986.
- [5] Eddie Dekel, Barton L. Lipman, and Aldo Rustichini. Representing preferences with a unique subjective state space. *Econometrica*, 69(4):891–934, July 2001.
- [6] J. Eichberger, S. Grant, and D. Kelsey. Updating Choquet beliefs. *Journal of Mathematical Economics*, 43(7-8):888–899, 2007.
- [7] Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.
- [8] Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, 61:1–22, 1993.
- [9] Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113:1–31, 2003.

- [10] Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57:937–969, 1989.
- [11] L.G. Epstein. An Axiomatic Model of Non-Bayesian Updating. *Review of Economic Studies*, 73(2):413–436, 2006.
- [12] L.G. Epstein, J. Noor, and A. Sandroni. Non-Bayesian updating: A theoretical framework. *Theoretical Economics*, 3(2):193–229, 2008.
- [13] Paolo Ghirardato. Revisiting Savage in a conditional world. *Economic Theory*, 20:83–92, March 2002.
- [14] Paolo Ghirardato and Massimo Marinacci. Risk, ambiguity, and the separation of utility and beliefs. *Mathematics of Operations Research*, 26:864–890, November 2001.
- [15] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.
- [16] Itzhak Gilboa and David Schmeidler. Updating ambiguous beliefs. *Journal of Economic Theory*, 59:33–49, 1993.
- [17] S. Grant, A. Kajii, and B. Polak. Preference for information and dynamic consistency. *Theory and decision*, 48(3):263–286, 2000.
- [18] F. Gul and O. Lantto. Betweenness satisfying preferences and dynamic choice* 1. *Journal of Economic Theory*, 52(1):162–177, 1990.
- [19] F. Gul and W. Pesendorfer. Temptation and Self-Control. *Econometrica*, 69(6):1403–1435, 2001.
- [20] F. Gul and W. Pesendorfer. The Revealed Preference Theory of Changing Tastes. *Review of Economic Studies*, 72(2):429–448, 2005.
- [21] F. Gul and W. Pesendorfer. The case for mindless economics. *The foundations of Positive and normative Economics: A handbook*, pages 3–42, 2008.

- [22] Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.
- [23] Eran Hanany and Peter Klibanoff. Updating preferences with multiple priors. *Theoretical Economics*, 2(3), September 2007.
- [24] Eran Hanany and Peter Klibanoff. Updating uncertainty averse preferences. mimeo, Kellogg Graduate School, 2007.
- [25] M. Horie. Reexamination on Updating Choquet Beliefs. *KIER Working Papers*, 2007.
- [26] Jean-Yves Jaffray. Dynamic decision making with belief functions. In Yaeger et al., editor, *Advances in the Dempster-Shafer Theory of Evidence*. Wiley, 1994.
- [27] Edi Karni and Zvi Safra. Ascending bid auctions with behaviorally consistent bidders. *Annals of Operations Research*, pages 435–446, 1989.
- [28] Edi Karni and Zvi Safra. Behaviorally consistent optimal stopping rules. *Journal of Economic Theory*, 51:391–402, 1990.
- [29] Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930 – 976, 2009.
- [30] Peter Klibanoff and Emre Ozdenoren. Subjective recursive expected utility. mimeo, Kellogg School of Management, 2004.
- [31] David M. Kreps. A representation theorem for “preference for flexibility”. *Econometrica*, 47(3):565–577, May 1979.
- [32] David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46:185–200, 1978.
- [33] F. Maccheroni, M. Marinacci, and A. Rustichini. Dynamic variational preferences. *Journal of Economic Theory*, 128(1):4–44, 2006.
- [34] Mark J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27:1622–1668, 1989.

- [35] E.F. McClennen. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press, 1990.
- [36] T. O'Donoghue and M. Rabin. Doing It Now or Later. *The American Economic Review*, 89(1):103–124, 1999.
- [37] M. Piccione and A. Rubinstein. On the Interpretation of Decision Problems with Imperfect Recall* 1. *Games and Economic Behavior*, 20(1):3–24, 1997.
- [38] C. P. Pires. A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152, 2002.
- [39] John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3:94–101, 1982.
- [40] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [41] Uzi Segal. Two-stage lotteries without the reduction axiom. *Econometrica*, 58:349–377, 1990.
- [42] Marciano Siniscalchi. Bayesian updating for general maxmin expected utility preferences. Mimeo, Princeton University, September 2001.
- [43] Marciano Siniscalchi. Dynamic choice under ambiguity. Technical Report 1430, Northwestern University CMS-EMS Discussion Paper, February 2009.
- [44] Marciano Siniscalchi. Two out of three ain't bad: a comment on "The ambiguity aversion literature: a critical assessment". *Economics and Philosophy*, 25(03):335–356, 2009.
- [45] C. Starmer. Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. *Journal of Economic Literature*, 38:332–382, 2000.
- [46] R.H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1955-1956.
- [47] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London and New York, 1991.