

Siniscalchi, Marciano

Working Paper

Dynamic choice under ambiguity

Discussion Paper, No. 1430

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Siniscalchi, Marciano (2006) : Dynamic choice under ambiguity, Discussion Paper, No. 1430, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/31239>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Dynamic Choice under Ambiguity

Marciano Siniscalchi*

December 8, 2006

Abstract

This paper analyzes sophisticated dynamic choice for ambiguity-sensitive decision makers. It characterizes *Consistent Planning* via axioms on preferences over decision trees. Furthermore, it shows how to elicit conditional preferences from prior preferences. The key axiom is a weakening of Dynamic Consistency, deemed *Sophistication*. The analysis accommodates arbitrary decision models and updating rules. Hence, the results indicate that (i) ambiguity attitudes, (ii) updating rules, and (iii) sophisticated dynamic choice are mutually orthogonal aspects of preferences.

As an example, a characterization of prior-by-prior Bayesian updating and Consistent Planning for arbitrary maxmin-expected utility preferences is presented. The resulting sophisticated MEU preferences are then used to analyze the value of information under ambiguity; a basic *trade-off between information acquisition and commitment* is highlighted.

Web Appendix available at <http://faculty.econ.northwestern.edu/faculty/siniscalchi>

Note: for convenience, this file includes both the paper (pp 1–47) and the Web Appendix (pp. i–xiii).

1 Introduction

This paper provides robust behavioral foundations for sophisticated dynamic choice in the presence of ambiguity. It departs from the existing literature on dynamic choice under uncertainty by assuming that the objects of individuals' preferences are decision trees, rather than uncertain prospects (“acts”) or contingent consumption plans (“temporal acts”). This approach avoids a tension between ambiguity and coherent dynamic choice that arises if preferences over acts are taken as the only behavioral primitive.

To highlight this tension, recall that Ellsberg [6] demonstrated that ambiguity manifests itself through violations of the Sure-Thing principle—the central axiom in Savage’s axiomatic foundation for expected-utility theory. But, as Savage himself emphasizes, the Sure-Thing Principle also provides a foundation

*Economics Department, Northwestern University, Evanston, IL 60208; Email marciano@northwestern.edu. I have greatly benefited from many extensive conversations with Peter Klibanoff. I also thank Pierpaolo Battigalli, Eddie Dekel, Larry Epstein, Paolo Ghirardato, Alessandro Lizzeri and Fabio Maccheroni for insightful discussions. All errors are my own.

for Bayesian updating, and ensures that the resulting conditional preferences are *dynamically consistent* (cf. Savage [35], pp. 21–22 and 43–33; see also Ch. 10 in Kreps [27], and Ghirardato [11]). Models of ambiguity-sensitive preferences, such as Gilboa and Schmeidler’s [12] “maxmin-expected utility” (MEU) or Schmeidler’s [36] “Choquet-expected utility” (CEU) necessarily relax the Sure-Thing principle; as might be expected in light of the preceding observation, updating rules for these decision models (e.g. Gilboa and Schmeidler [13], Jaffray [20], Pires [33], Shafer [37] and Walley [42]) typically lead to violations of dynamic consistency. Indeed, Epstein and Le Breton [8] show that full dynamic consistency is generally incompatible with non-neutral attitudes towards ambiguity.¹

Dynamic consistency provides a simple rationale for backward induction. Specifically, it ensures that a sequentially optimal plan of action in a decision tree will also be optimal from the point of view of prior preferences, and conversely;² thus, backward induction can be viewed as an efficient way to compute a-priori optimal plans. But, when dynamic consistency fails, a-priori optimal and sequentially optimal plans may differ: an example is provided in §1.1 below. In such circumstances, the behavioral implications of a theory that relies *solely* on prior and conditional preferences over acts are not fully determined: will the individual follow her sequentially optimal plan, the a-priori optimal one, or a “compromise” plan? In fact, even testing whether an agent adopts a given dynamically-inconsistent updating rule may be problematic if the latter does not fully determine her dynamic-choice behavior.³

The approach proposed in this paper avoids these difficulties. Dynamic Consistency is a property of preferences over *acts*; on the other hand, the results in this paper rely primarily upon assumptions on preferences over non-trivial *decision trees*, and hence are unaffected by departures from Dynamic Consistency. As an added benefit, the proposed approach can accommodate a wide range of decision models and updating rules. To summarize the main results:

- Theorem 1 in Section 3.1 characterizes *Consistent Planning* for dynamic decision problems under uncertainty. This refinement of backward induction was introduced by R. H. Strotz [41] in the context of intertemporal choice with changing tastes, and has also been successfully employed in the setting of choice under risk: see e.g. Karni and Safra [21, 22], Caplin and Leahy [2]. These contributions adopt Consistent Planning as a solution concept for decision trees, but do not provide an axiomatic foundation.⁴ Theorem 1 provides such a foundation, formalizing the intuitive notion that a *sophisticated* decision-maker correctly anticipates her future preferences.
- The characterization of Consistent Planning requires that the individual’s conditional preferences

¹See also the discussion of Epstein and Schneider [9], Hanani and Klibanoff [18] and Klibanoff [24] in §1.4 below.

²Assuming all relevant conditioning events have positive prior probability.

³Note that conditional preferences cannot all be observed directly, especially in the kind of non-repeatable situations where ambiguity is typically of greatest interest: if an event occurs, one cannot observe preferences conditional on its complement.

⁴Gul and Pesendorfer [15] have recently axiomatized versions of Consistent Planning with changing tastes (under certainty).

over trees be specified. Section 3.2 shows that, under suitable conditions, these can be *elicited* from her prior preferences (Theorem 2). Together, Theorems 1 and 2 imply that, just like a coherent theory of dynamic choice can be rationalized in terms of prior preferences over *acts* if these conform to expected utility, a similarly coherent theory of dynamic choice can be founded upon prior preferences over *trees* for more general decision models.

- To exemplify the approach described in this paper, Section 4.1 characterizes consistent planning for arbitrary MEU preferences and prior-by-prior Bayesian updating (Theorem 3). In particular, no restriction need be imposed on the set of priors. Section 4.2 then employs the resulting model of “sophisticated MEU preferences” to analyze a simple value-of-information problem. Leveraging the framework and results of this paper, the analysis illustrates that a basic *trade-off between information acquisition and commitment* emerges under ambiguity.

As noted above, Theorems 1 and 2 require only relatively mild regularity conditions on preferences over acts. Hence, the results in this paper suggest that (i) ambiguity attitudes, (ii) updating rules, and (iii) sophisticated dynamic choice are essentially orthogonal aspects of behavior, which can be tested, modeled and axiomatized independently.

It is also worth emphasizing that, while the *characterization* of Consistent Planning involves preferences over trees, its *use as a solution concept* only requires prior and conditional preferences over Savage acts. Thus, in practice, when analyzing a dynamic choice problem, it is enough to specify the agent’s prior preferences over acts and her updating rule, then apply Consistent Planning. The results in this paper provide a rationale for this approach, regardless of ambiguity attitudes and choice of updating rules. From a conceptual point of view, as will be emphasized below, Consistent Planning can also be viewed as a way to “extend” the individual’s preferences from acts to arbitrary trees.

The remainder of this section documents the issues just discussed in specific examples; §1.4 discusses the related literature, including alternative approaches to dynamic choice under ambiguity. Section 2 describes the formal framework. Section 3 contains the main results on Consistent Planning (§3.1) and the elicitation of conditional preferences (§3.2). Section 4.1 considers the special case of MEU preferences and prior-by-prior updating; a value-of-information application is in §4.2; finally, §4.3 discusses Consistent Planning for infinite trees (the formal treatment is in §B.2 of the Web Appendix).

1.1 When Dynamic Consistency Fails

Consider the dynamic decision problem in Figure 1. A decision-maker (DM henceforth) is presented with an urn containing 90 balls, of which 30 are red and 60 green and blue, in unspecified proportions. There are two *prizes*, \$0 and \$10; x denotes an arbitrary prize. A ball will be drawn from the urn; the *state space* is denoted by $\Omega = \{r, g, b\}$, in obvious notation. If the ball drawn from the urn is blue, the DM

receives x . Otherwise, the DM is informed that the ball drawn is not blue, and can choose whether to bet on red or green. Finally, the DM learns the outcome of the draw, and receives the appropriate prize.




Figure 1: The tree f_x , representing deferred choice in the Ellsberg model; $x \in \{0, 10\}$.

Throughout this paper, “•” represents a decision node and “○” corresponds to a point in the tree where information is revealed to the DM (or, slightly improperly, a “chance node”). Trees are drawn from left to right, and from top to bottom; edges departing from decision nodes are decorated with action labels (A , R and G in Fig. 1), whereas edges departing from chance nodes are labelled with events ($\{r, g\}$, $\{r\}$ and $\{g\}$ in Fig. 1). The decision tree in Fig. 1 will be denoted f_x .

The tree f_x may be viewed as a menu of two *plans*: “choose A , then R ” (denoted AR_x) and “choose A , then G ” (denoted AG_x); see Fig. 2. Formally, a plan is a tree featuring a single action at every decision node; it is useful to think of a plan as representing a *commitment to carry out specific contingent choices*.




Figure 2: The plans AR_x and AG_x ; $x \in \{0, 10\}$.

The tree f_x contains two more subplans of interest: those beginning with the choice of R or G at the second decision point. These are “conditional” subtrees: they incorporate the information that the ball drawn was either red or green. For simplicity, in this section only, these subtrees (drawn in red and green respectively in Figs. 1 and 2) will also be denoted by R and G ; see §2 for a formal notation for subtrees.

Suppose that the DM has MEU preferences both ex-ante and after observing the event $E = \{r, g\}$; denote these preferences by \succsim and \succsim_E respectively. The DM’s utility function u satisfies $u(10) = 10$ and $u(0) = 0$; her sets of priors C and posteriors C_E on the state space $\Omega = \{r, g, b\}$ are given by

$$C = \left\{ \pi \in \Delta(\Omega) : \pi(\{r\}) = \frac{1}{3} \right\} \quad \text{and} \quad C_E = \left\{ \pi(\cdot | \{r, g\}) : \pi \in C \right\} = \left\{ \pi \in \Delta(\Omega) : \pi(\{r\}) \geq \frac{1}{3}, \pi(\{b\}) = 0 \right\}. \quad (1)$$

Thus, for all acts $f, g \in \{0, 10\}^\Omega$, $f \succsim g$ if and only if $\min_{q \in C} \int u(f(\omega)) dq \geq \min_{q \in C} \int u(g(\omega)) dq$, and similarly $f \succsim_E g$ if and only if $\min_{q \in C_E} \int u(f(\omega)) dq \geq \min_{q \in C_E} \int u(g(\omega)) dq$. The set C_E is obtained by updating every prior in C .

To rank *plans*, adopt the standard *Reduction* assumption. Any plan corresponds naturally to a Savage act; for instance, AR_x corresponds to the act $(10, 0, x)$ that yields the prizes 10, 0 and x in states r , g and b respectively. Reduction postulates that the DM ranks plans by comparing the corresponding acts: thus,

$$AR_0 \succ AG_0, \quad \text{and} \quad AR_{10} \prec AG_{10}. \quad (2)$$

Similarly, the subplans R and G correspond to restricted acts with domain E . The standard approach is to assume that their ranking is determined solely by the comparison of these restricted acts given E , regardless of whether one views R and G as subplans of f_0 or f_{10} . This assumption, which complements Reduction, is called *Consequentialism*. It is adopted in most (but not all) existing work on conditional preferences and dynamic choice under ambiguity: see §1.4 for details. In this example,

$$R \succ_E G. \quad (3)$$

Note that Reduction is assumed here for simplicity, but it is not required for the main results of this paper.

The backward-induction analysis of the trees f_0 and f_{10} is particularly simple. There is no choice to be made at the initial node, so only the second decision point must be considered. Since $R \succ_E G$, the inferior action G can be pruned; the Backward Induction algorithm then terminates.

Thus, for every $x \in \{0, 10\}$, the plan AR_x is the “backward-induction solution” of the tree f_x . Under expected utility and Bayesian updating, a backward-induction solution is also optimal from the perspective of prior preferences. However, in this example, this is not true for $x = 10$: the backward-induction solution of the tree f_{10} is AR_{10} , but Eq. (2) indicates that $AR_{10} \prec AG_{10}$. Because of this inconsistency, it is not clear exactly what kind of behavior should be expected in the tree f_{10} on the basis of the act preferences described above.

One way to avoid similar inconsistencies is to require that

$$\forall x \in \{0, 10\}, \quad R \succ_E G \iff AR_x \succ AG_x. \quad (4)$$

Eq. (4) is an implication of *Dynamic Consistency*; it is also clearly violated by the preferences considered here (cf. Eqs 2 and 3). The key observation is that *this violation is related to ambiguity*; to see this, note that Eq. (4) implies that $AR_0 \succ AG_0$ if and only if $AR_{10} \succ AG_{10}$. By Reduction, this corresponds to

$$(10, 0, 0) \succ (0, 10, 0) \iff (10, 0, 10) \succ (0, 10, 10). \quad (5)$$

Now recall that the modal preferences in Ellsberg’s original three-color-urn problem exhibit the pattern

$$(10, 0, 0) \succ (0, 10, 0) \quad \text{and} \quad (10, 0, 10) \prec (0, 10, 10); \quad (6)$$

this is also the pattern exhibited by preferences in the example under consideration here.

To summarize, in this simple example, in order to avoid inconsistencies between backward-induction and ex-ante analysis of the trees f_0 and f_{10} , it is necessary to *rule out Ellsberg-type preferences*. This is an instance of the tension between ambiguity and dynamic consistency referred to in the Introduction; indeed, a result emphasizing this tension was established in considerable generality by Epstein and Le Breton [8] (see also Ghirardato [11]).

Eq. (4) also provides a way to elicit the conditional ranking of R and G from prior preferences: this is a general property of Dynamic Consistency (cf. the observations at the end of §3.2.2). On the other hand, when Dynamic Consistency fails, even simply testing whether observed behavior confirms to or contradicts a given updating rule requires explicit additional assumptions on how the DM resolves conflicts between prior and conditional preferences.

To see this, consider a DM with MEU preferences and priors C as in Eq. (1); however, assume now that her set of posteriors is not known. An experimenter conjectures that she might be using the *maximum-likelihood* updating rule (Gilboa and Schmeidler [13]);⁵ in this case her set of posteriors would reduce to $C'_E = \{\pi\}$, where $\pi(\{r\}) = \frac{1}{3}$ and $\pi(\{g\}) = \frac{2}{3}$. For this posterior, G is strictly preferred to R given E . Now assume that this DM chooses R in the tree f_0 . This can be interpreted in (at least) two ways: (i) the DM's behavior conforms to backward-induction, and she does not use maximum-likelihood updating; or (ii) the DM's behavior is driven by the prior preference $AR_0 \succ AG_0$, so no conclusion can be drawn regarding her updating rule. Thus, even in this simple example, additional assumptions, such as Backward Induction, are required to relate observed behavior to specific updating rules.

1.2 Consistent Planning, Sophistication and Weak Commitment

The overall objective of this paper is to show that, despite the potential inconsistency between ex-ante and backward-induction analysis, it is still appropriate to employ backward induction to make specific predictions about dynamic choice behavior. In particular, to avoid indeterminacies due to ties, Strotz's concept of Consistent Planning will be the focus of the analysis.

In Strotz's own words, Consistent Planning captures the assumption that, when faced with a problem such as the one depicted in Fig. 1, the DM will "find the best plan among those that [s]he will actually follow." [41, p.173] Like backward induction, Consistent Planning is an algorithm that iteratively constructs a "solution" to the dynamic problem under consideration. However, when dynamic consistency fails, backward induction no longer has a clear rationale, despite its intuitive appeal: surely it can no longer be viewed as a way to construct ex-ante optimal plans. A fortiori, this is true for Consistent Planning, which refines backward induction by incorporating a specific tie-breaking rule.

Theorem 1 provides a simple behavioral characterization of Consistent Planning that addresses these

⁵This rule prescribes updating only those priors in C that maximize the probability of the conditioning event E .

issues; this subsection reviews the two main axioms. As noted above, the behavioral assumptions underlying Consistent Planning can only be formalized in terms of the individual’s preferences over general trees. This is in the spirit of the literature on menu choice originating from Kreps [26]: by carefully considering the DM’s preferences over trees, it is possible to identify “benchmark” situations in which she will, or will not, be able to carry out a-priori desirable future actions.

Consistent Planning can also be viewed as a way to *extend the DM’s preferences from acts to general trees*; the statement of Theorem 1 and of the main axioms themselves support this view. To clarify, in the example in Fig. 1, the fact that AR_x is the Consistent-Planning solution of f_x may be taken to imply that, for each $x \in \{0, 10\}$, the tree f_x is deemed equivalent to the plan AR_x : $f_x \sim AR_x$. A further consequence is that $f_0 \sim f_{10}$ —a statement about the DM’s preferences over trees.

The first main axiom is *Sophistication*: loosely speaking, it requires that the DM *hold correct expectations regarding her future choices*. However, this intuition is formalized solely in terms of prior and conditional preferences—there is no need to formally model the DM’s “beliefs” about her future preferences and/or choices. For instance, in the tree f_x of Fig. 1, Sophistication yields the following restrictions:

$$\forall x \in \{0, 10\}, \quad R \succ_E G \Rightarrow f_x \sim AR_x \quad \text{and} \quad G \succ_E R \Rightarrow f_x \sim AG_x \quad (7)$$

That is: the DM evaluates the trees f_{10} and f_0 *as if* they simply did not contain the continuation trees that she will surely not choose. This indirectly reflects the DM’s correct beliefs about her future preferences.

Notice that, unlike Eq. (4), Eq. (7) does not impose any restriction on the relative ranking of AR_x and AG_x ; thus, a DM can satisfy (Reduction, Consequentialism and) Sophistication while at the same time exhibiting the modal preferences in the Ellsberg paradox.

Also observe that Eq. (7) does not yield any restrictions in case $R \sim_E G$. This is immaterial for the preferences considered in §1.1, but may lead to indeterminacies in other settings; thus, a *tie-breaking* assumption is required. Consider for instance the decision tree in Fig. 3, denoted f .




Figure 3: Sophistication and Weak Commitment

The state space is $\Omega = \{r, g, b, w\}$; the DM has MEU preferences, with priors $C = \{\pi \in \Delta(\Omega) : \pi(\{r\}) = \pi(\{b\}), \pi(\{g\}) = \pi(\{w\})\}$. Also let $E = \{r, g\}$ and suppose that conditional preferences \succ_E are determined by prior-by-prior Bayesian updating, so the set of posteriors is $C_E = \{\pi \in \Delta(\Omega) : \pi(\{r\}) = 1 - \pi(\{g\})\}$.

Note that $R \sim_E G$; thus, Sophistication has no bite in this decision tree. However, this DM strictly prefers the act that yields 10 if g or b obtain, and 0 otherwise, to the act that yields 10 if r or b obtain, and 0 otherwise. Hence, a priori the DM would like to “commit” to choosing G at her second decision node.

Intuitively, allowing the DM to evaluate the tree f as if she could commit to G seems consistent with the logic of “finding the best plan among those she will actually follow:” the DM will have no incentive to deviate from the plan that prescribes G at her second decision node. The *Weak Commitment* axiom captures precisely this assumption. Denote by AR and AG the plans obtained from f by pruning the green and, respectively, the red subtrees; Weak Commitment then implies:

$$G \sim_E R, \quad AG \succ AR \quad \Rightarrow \quad f \sim AG$$

(and similarly if $AR \succ AG$). That is: if, tomorrow, the DM will be indifferent between G and R , but today she would like to be able to commit to G , then indeed she will be able to—and consequently she evaluates f as if the a priori inferior alternative R was not available. Again, observe that no restriction on prior preferences over acts or plans is required.

It is worth emphasizing that Sophistication and Weak Commitment do *not* provide a rationale for *recursion*, or replacing subtrees with the conditional certainty equivalent of the optimal continuation plans: they only rationalize pruning conditionally inferior actions.

To see this, consider the tree f_0 in Fig. 1 and the MEU preferences described in §1.1. To ensure the existence of certainty equivalents, assume that the prize set is $X = [0, 10]$ and, for simplicity, that utility is linear. Recall that $R \succ_E G$, so Consistent Planning implies that $f_0 \sim AR_0$. However, the conditional certainty equivalent of R is $\frac{10}{3}$; the prior MEU evaluation of the act $f' = (\frac{10}{3}, \frac{10}{3}, 0)$ is $\frac{10}{9}$, whereas (assuming Reduction) the prior MEU evaluation of AR_0 is $\frac{10}{3}$. Thus, the DM is *not* indifferent between f_0 and f' .

Recursion requires the full force of Dynamic Consistency: see Epstein and Schneider [9]. Informally, under Consistent Planning, the DM evaluates her future choices “from today’s perspective”; recursion, on the other hand, is only meaningful if today’s and tomorrow’s perspectives coincide.

1.3 Eliciting Conditional Preferences

Sophistication, together with basic structural assumptions that guarantee the existence of certainty equivalents, enables the elicitation of conditional preferences from prior preferences. Consider once again the tree f_x in Fig. 1 and the preferences in §1.1. To elicit the ranking of the subplans R and G conditional on $E = \{r, g\}$, consider the trees $AR_{x,y}$ and $AG_{x,y}$ and the plan $A_{x,y}$ in Fig. 4, where $y \in [0, 1]$.

To clarify, the choice Y in the trees of Fig. 4 yields a certain payoff of y . The conditional MEU evaluations of R and G are $\frac{10}{3}$ and 0 respectively. Now fix a prize $z \in (0, \frac{10}{3})$: then, for every $y \prec z$ (equivalently, $y < z$), the DM will prefer R to Y at the second decision node in the tree $AR_{x,y}$; similarly, for every $y \succ z$,




Figure 4: The trees $AR_{x,y}$ and $AG_{x,y}$, and the plan $A_{x,y}$; $x \in \{0, 10\}$, $y \in [0, 10]$.

she will prefer Y to G at the second node of $AG_{x,y}$. Intuitively, the prize z serves as a “wedge” between the plans R and G . Then Sophistication yields the following implications:

$$\forall y \in X: \quad y \prec z \Rightarrow AR_{x,y} \sim AR_x \quad \text{and} \quad y \succ z \Rightarrow AG_{x,y} \sim A_{x,y}; \quad (8)$$

the plan AR_x is as in Fig. 2. Observe that Eq. (8) only involves prior preferences: in other words, under Sophistication, Eq. (8) is a directly observable implication of the conditional preference ranking $R \succ_E G$.

This suggests the following approach: *stipulate that R is revealed weakly preferred to G given E if there exists a “wedge” prize $z \in X$ such that Eq. (8) holds.* In this example, any prize $z \in (0, \frac{10}{3})$ satisfies this condition; on the other hand, there is *no* prize $z \in X$ for which $y \prec z$ implies $AG_{x,y} \sim AG_x$ and $y \succ z$ implies $AR_{x,y} \sim A_{x,y}$: hence, G is *not* revealed weakly preferred to R given E . Taken together, these observations “reveal” that $R \succ_E G$.

Theorem 2 shows that, under suitable conditions, the notion of revealed conditional preference is well-defined and, loosely speaking, allows the analyst to elicit the DM’s “actual” conditional preferences. As noted above, the required conditions do not significantly restrict preferences over acts; thus, the approach proposed here applies to a variety of decision models and updating rules. Furthermore, Sophistication is only required to hold for trees such as $AR_{x,y}$ and $AG_{x,y}$: see §3.2 for details.

1.4 Related literature

A (small) sample of contributions on updating rules for MEU, CEU and other decision models are referenced at the beginning of this Introduction.

Myerson [30] characterizes EU preferences and Bayesian updating for “conditional probability systems” by considering axioms on a collection of conditional preferences; Dynamic Consistency plays a central role in his analysis. Skiadas [40] considers conditional preferences over pairs consisting of a state-contingent consumption plan and an information filtration, and axiomatizes a recursive expected-utility representation. A version of dynamic consistency is also central to his analysis.

Epstein and Schneider [9] characterize recursive MEU preferences over Savage-style acts. These authors retain Consequentialism and, implicitly, Reduction, and restrict Dynamic Consistency to a fixed, pre-specified collection of events. In particular, they analyze dynamic choice in decision trees generated

by a fixed filtration, or sequence of partitions; Dynamic Consistency is required to hold for all preferences conditional upon elements of these partitions. A related approach is investigated in Wang [43].

Consistently with the preceding discussion, this results in a restriction on prior preferences; in particular, it can be shown that, under the assumptions in [9], Savage’s Sure-Thing Principle and Eq. 4 will hold for all events the DM can condition upon. Loosely speaking, this rules out Ellsberg-type behavior with respect to “learnable” events; for instance, in the tree f_x of Fig. 1, the axioms in [9] rule out Ellsberg-type prior preferences (a fact that is also noted by Epstein and Schneider). By way of comparison, an objective of the present paper is precisely to avoid imposing any restriction on prior preferences.

The discussion in the preceding subsections indicates that the possibility of Ellsberg-type behavior is precluded by the combination of Reduction, Consequentialism and Dynamic Consistency. Other authors have explored retaining Dynamic Consistency while dropping the other two axioms. In particular, to accommodate Dynamic Consistency, Klibanoff [24] drops Reduction, and also introduces a form of state-dependence of preferences over prizes. More recently, Hanani and Klibanoff [18] have proposed an updating rule for MEU preferences; their analysis drops Consequentialism (but maintains Reduction).

This paper should be viewed as complementary to these contributions. As noted above, the main results of this paper (i.e. Theorem 1 and 2) do *not* assume Reduction. Also, the present paper does not restrict attention to MEU preferences, or to specific updating rules. Indeed, as was mentioned above, the results in this paper indicate that the analysis of *updating* is essentially orthogonal to the characterization of *sophisticated behavior* in decision trees.

Assuming that preferences are defined over decision trees is much in the same spirit as Kreps’s seminal contribution on menu choice (Kreps [26]) and the temporal resolution of uncertainty (Kreps [28]). Gul and Pesendorfer [15] analyze the behavioral foundations of changing tastes in a model of temporal choice under certainty; as in the present paper, preferences are defined over intertemporal decision problems. Epstein [7] adopts a menu-choice framework to study “non-Bayesian” updating. Klibanoff and Ozdenoren [25] characterize a subjective version of Kreps’s [28] model of recursive expected utility; their decision setting is related to the one adopted here.

Versions of Consistent Planning have also been used in the literature of dynamic choice under risk: see for instance Karni and Safra [21, 22], who suggest the expression “behavioral consistency.” Cubitt, Starmer and Sugden [5] compare different non-expected utility models under risk in a dynamic-choice experiment, and find that “behavioral consistency” outperforms alternative specifications. Caplin and Leahy [2] analyze a class of problems where Consistent Planning admits a recursive formulation, and provide a general existence result. Machina [29] provides a critical discussion of Consistent Planning (“folding back”) under risk; some of his observations are arguably less applicable to dynamic choice under uncertainty, and others (such as issues related to the value of information) can be explicitly addressed once preferences are defined over decision trees, as this paper proposes. None of these contributions

provide an axiomatic analysis of backward induction or Consistent Planning.

Hammond [16, 17] and Cubitt [4] provide an analysis of Consequentialism and Dynamic Consistency in decision trees for EU preferences; see also Ghirardato [11] and, in a risk setting, Karni and Schmeidler [23]. Sarin and Wakker [34] consider non-EU behavior in decision trees, and in particular investigate the consequences of the assumption that prior and conditional preferences belong to the same class of models (e.g. MEU, CEU, etc.).

2 Decision Setting

This section introduces the basic notation for decision trees. The axioms of Section 3 involve modifying a tree at certain decision points in various ways. For this reason, a main objective of the notation adopted here is to allow a precise, yet relatively straightforward formalization of such “tree-surgery” operations.

2.1 Histories and Trees

Decision trees are described adapting the notation for “perfect-information game trees” in Osborne-Rubinstein [31].⁶ The basic building block in this formalism is the *history*: an ordered list of the DM’s actions and “chance moves” that describes a possible (partial or complete) unfolding of occurrences in the decision tree under consideration. Specifically, the DM’s actions are labels representing *choices* available to the DM in a given period; chance moves represent *information* that the DM may receive at the end of the period under consideration. Additionally, *terminal histories* indicate a prize, i.e. the ultimate outcome of the DM’s actions and the resolution of uncertainty in the history under consideration. A decision tree can then be represented simply as a set of partial and terminal histories.

Formally, fix a set Ω of states, an algebra Σ of events, and a set X of prizes; assume the latter is a connected separable topological space. Also fix a set A of *action labels* (e.g. letters or sequence of letters, numerals, etc.); assume that A is countably infinite.⁷

Definition 1 A (partial) history of length $T \geq 0$ starting at $E \in \Sigma$ is a sequence

$$h = [a_1, E_1, \dots, a_T, E_T],$$

such that, for all $t = 1, \dots, T$, $a_t \in A$ and $E_t \in \Sigma$, and $E \supset E_1 \supset \dots \supset E_T$.⁸ A **terminal history** of length

⁶Fudenberg and Tirole [10] use a similar notation for “multistage games with observable actions”.

⁷The topological assumptions on X are only required for the purposes of eliciting conditional preferences; the cardinality assumption on A are required in the proof of Theorem 1 (see §A.1.1 in the Appendix for details); it also ensures that Axiom 3.2 does not hold vacuously.

⁸Here and in the following, \subset denotes weak inclusion, and \subsetneq denotes strict inclusion; similarly for \supset and \supsetneq .

$T \geq 1$ starting at $E \in \Sigma$ is a sequence

$$h = [a_1, E_1, \dots, a_{T-1}, E_{T-1}, x],$$

where, for $t = 1, \dots, T-1$, a_t and E_t are as above and $x \in X$. The empty history (a partial history) is denoted by \emptyset . The set of all partial histories starting at $E \in \Sigma$ is denoted by \mathcal{H}_E ; the set of all terminal histories starting at $E \in \Sigma$ is denoted \mathcal{T}_E .

Table 1 defines additional terms and notation related to histories.

	Notation	Remarks and Examples
Length of history h	$\lambda(h)$	$\lambda(\emptyset) = 0$.
Last action taken at history $h \in \mathcal{H}_E \setminus \{\emptyset\}$	$a(h)$	$a([a_1, E_1, a_2, E_2]) = a_2$
Last realized event in history h	$E(h)$	$E([a_1, E_2, a_2, E_2]) = E_2$; $E([a_1, E_1, x]) = E_1$; if \emptyset is viewed as an element of \mathcal{H}_E , $E(\emptyset) = E$.
Prize at terminal history $h \in \mathcal{T}_E$	$\xi(h)$	$\xi([a_1, E_1, a_2, E_2, x]) = x$
Subhistory of $h = [a_1, E_1, \dots, a_T, E_T]$ or $h = [a_1, E_1, \dots, a_T, E_T, x]$	$h_t = [a_1, E_1, \dots, a_t, E_t]$	Note: requires $t \leq T$. $h_0 = \emptyset$.
Subhistory relation	$h' \leq h$ $h' < h$	Means $\exists t : h' = h_t$. Means $h' \leq h$ and $h' \neq h$.
Composition of histories	$[h, h']$ $[h, a, E]$, $[h, x]$.	Needs $h' \in \mathcal{H}_{E(h)}$ and $E \subset E(h)$; also $h \in \mathcal{H}_E$ if $h' \neq \emptyset$; $[\emptyset, h] = [h, \emptyset] = h$.

Table 1: Terms and notation for histories

Definition 2 Let $E \in \Sigma \setminus \{\emptyset\}$. A **decision tree starting at E** is a subset f of $\mathcal{H}_E \cup \mathcal{T}_E$ such that

1. if $h \in f$ and $\lambda(h) > 0$, then $h_{\lambda(h)-1} \in f$;
2. for every $h \in f$ and every $a \in A$ such that $[h, a, F] \in f$ for some $F \in \Sigma$, the collection $\{F : [h, a, F] \in f\}$ is a (possibly trivial) partition of $E(h)$;
3. $h \in f \cap \mathcal{T}_E$ if and only if there is no $h' \in f$ such that $h < h'$.

A tree f is **finite** if it is a finite set. The sets of all finite trees starting at E is denoted by F_E .

By Condition 1, if a history h can occur in a tree, then all truncations h' such that $h' \leq h$ can also occur; in particular, the empty history can occur ($\emptyset \in f$). In Condition 2, the assumption that $\{F : [h, (a, F)] \in f\}$ is a partition of $E(h)$ ensures that continuation histories are specified for every state $\omega \in E(h)$. Finally, Condition 3 states that a history is terminal precisely when it is not followed by other actions (or immediate prizes) in f .

Section 3 focuses on finite trees; §4.3 discusses extensions of the main results to a class of infinite trees. Table 2 below defines additional useful notation.

	Notation and Definition
Choices (actions and prizes) available at $h \in f \cap \mathcal{H}_E$	$C_f(h) = \{a \in A : \exists F \in \Sigma, [h, a, F] \in H\} \cup \{x \in X : [h, x] \in f\}$
Information partition following h and $c \in C_f(h)$	$\mathcal{F}_f(h, c) = \{F \in \Sigma : [h, c, F] \in f\}$. Note: implies $\mathcal{F}_f(h, c) = \emptyset$ if $c \in C_f(h) \cap X$.

Table 2: Additional notation for a decision tree $f \in F_E$.

A number of special types of trees will now be described, both to clarify the notation and because they will play a role in the analysis

Constant trees. The simplest possible tree corresponds to a prize to be received immediately. Formally, if $x \in X$, the constant tree corresponding to x is $\{\emptyset, [x]\}$. In the the following, we shall abuse notation slightly and refer to such a tree simply as “ x ”.

(Simple) Savage acts. In the standard Savage [35] setting, an act is a Σ -measurable map $\varphi : \Omega \rightarrow X$; an act φ is simple if $\varphi(\Omega)$ is finite. For every simple Savage act φ , one can construct a corresponding tree $f \in F_\Omega$ by choosing an arbitrary action label $a \in A$ and letting $f = \{\emptyset\} \cup \{[a, \varphi^{-1}(x)], [a, \varphi^{-1}(x), x] : x \in \varphi(\Omega)\}$

Plans. A tree $f \in F_E$ is a plan if the set $C_f(h)$ is a singleton for every $h \in f \cap \mathcal{H}_E$. That is, there is a unique choice at every decision point in f .

Partitional Trees. A tree $f \in F_\Omega$ is partitional if there is a sequence $\mathcal{F}_1, \dots, \mathcal{F}_T$ of progressively finer Σ -measurable partitions of Ω such that all terminal histories have length $T+1$, and furthermore $\mathcal{F}_f(h, a) = \{F \in \mathcal{F}_{\lambda(h)+1} : F \subset E(h)\}$ for every non-terminal history h and action $a \in C_f(h) \cap A$. Thus, in a partitional tree, Nature’s moves are independent of the DM’s choices. Definition 2 allows for greater flexibility; for instance, it can describe a situation in which the DM can acquire different signals about the prevailing state of the world.

Notation. For every $E \in \Sigma$, the set of plans in F_E will be denoted by F_E^p .

2.2 Preferences: Consequentialism and Relabeling

The main object of interest in this paper is a collection $\{\succsim_E\}_{\emptyset \neq E \in \Sigma}$ of binary relations; in particular, for all non-empty $E \in \Sigma$, \succsim_E is a preference relation on F_E . For notational simplicity, \succsim_Ω will be denoted by \succsim .

Before turning to the actual axiomatic analysis in Section 3, it is worth commenting on two important aspects of conditional preferences in the present framework. First, since each conditional preference \succsim_E is defined over the set of decision trees starting at E , **Consequentialism** is implicitly assumed throughout.

Second, two trees may, loosely speaking, share a common overall structure and payoffs, and differ only in the choice of action labels at decision points; for instance, by replacing the action labels A, R, G in the tree in Fig. 1 with, say, X, Y, Z , one obtains a new tree that has the same structure and yields the same payoffs as the original one. Such trees are distinct objects in the formalism adopted here, and preferences may in principle treat them differently. The following definition and assumption rule out this possibility: it will be assumed that preferences over trees are invariant to relabeling. This is necessary for a correct interpretation of the tie-breaking assumption (Axiom 3.2 in Sec. 3); moreover, it seems intuitively consistent with the logic of backward induction and consistent planning.

Definition 3 Consider two trees $f, g \in F_E$. Then g is a **relabeling** of f , written “ $f \approx g$ ”, if there exists a bijection $\varphi : f \rightarrow g$ such that

- (i) for all $h, \bar{h} \in f$, $h \leq \bar{h}$ iff $\varphi(h) \leq \varphi(\bar{h})$;
- (ii) for all $h \in f \cap \mathcal{H}_E$, $E(h) = E(\varphi(h))$;
- (iii) for all $h \in f \cap \mathcal{H}_E$, $a \in C_f(h)$, and $D, D' \in \mathcal{F}_f(h, a)$, $a(\varphi([h, a, D])) = a(\varphi([h, a, D']))$;
- (iv) for all $h \in f \cap \mathcal{T}_E$, $\xi(h) = \xi(\varphi(h))$.⁹

Condition (i) requires that the relabeling φ preserve the ordering of histories: if h precedes \bar{h} in f , then the same must be true of the histories corresponding to h and \bar{h} in g . Condition (ii) requires that corresponding histories in f and g terminate with the same event; similarly, Condition (iv) ensures that the same payoffs are assigned at corresponding terminal histories. Finally, Condition (iii) ensures that the same actions are available at corresponding terminal histories. Specifically, if there is an action a at a history h of f that can be followed by two events D and D' , then there must be an action a' at the history h' of g corresponding to h that can be followed by the same two events D and D' .

Relabeling is an equivalence relation; for this and other properties of relabeling, see Section A.1.1 in the Appendix. It is now possible to formalize the assumption that “action labels don’t matter”:

Assumption 2.1 (Irrelevance of action labels) For all $E \in \Sigma$ and $f, g \in F_E$: $f \approx g$ implies $f \sim_E g$.

2.3 Tree Surgery

Finally, the key notions of *subtree*, *continuation tree* and *replacement tree* will be formally introduced. Since a decision tree is a set of histories, it is possible to formalize these relations and operations using straightforward set-theoretic notions. Some of the axioms in Sec. 3 require modifying trees in specific ways; the definitions in this subsection provide the basic formal language to describe such operations.

⁹Notice that (i)–(iv) imply further intuitive restrictions: in particular, by (i), $h \in f \cap \mathcal{T}_E$ iff $\varphi(h) \in g \cap \mathcal{T}_E$, so (iv) is well-posed. Furthermore, $\varphi(\emptyset) = \emptyset$ and $\lambda(h) = \lambda(\varphi(h))$ for all $h \in f$.

Definition 4 A tree $g \in F_E$ is a **subtree** of $f \in F_E$ if $g \subset f$.

Observe that an *arbitrary* subset of f need *not* be a subtree of f —the above definition explicitly requires that a subtree be a tree in its own right. Also notice that the definition implies that f and g can only differ in that some actions available at certain histories of f are removed; however, no new terminal histories are introduced (i.e. not all actions available at a history can be removed).

Definition 5 Consider a tree $f \in F_E$ and a $h \in f$. The **continuation tree** beginning at h , denoted $f(h)$, is the set $\{h' \in \mathcal{H}_{E(h)} \cup \mathcal{T}_{E(h)} : [h, h'] \in f\}$ if $h \in \mathcal{H}_E$, and the set $\{\emptyset, [\xi(h)]\}$ if $h \in \mathcal{T}_E$.

Thus, if h is a partial history, the continuation tree $f(h)$ contains all the histories h' such that $[h, h']$ is a history of f ; if instead h is terminal and in particular the tree f yields the prize x at h , then the only possible “continuation” for f at h is the (degenerate) tree that yields x with certainty and immediately. Notice that, according to this definition, for $h = \emptyset$, $f(h) = f$.

Additional notation: The preceding two notions can be usefully combined to characterize “partial” continuation trees. Formally, consider $f \in F_E$, a non-terminal history $h \in f \cap \mathcal{H}_E$ and a set $B \subset C_f(h)$ of actions and prizes available at h in the tree f . Then $f(h, B)$ denotes the (unique) \subset -maximal element of $F_{E(h)}$ such that $f(h, B) \subset f(h)$ and $C_{f(h, B)}(\emptyset) = B$.

For simplicity, if $B = \{c\}$ for $c \in C_f(h)$, I shall write $f(h, c)$ in lieu of $f(h, \{c\})$.

Definition 6 Consider a tree $f \in F_E$, a non-terminal history $h \in f \cap \mathcal{H}_E$, and another tree $g \in F_{E(h)}$. The **replacement tree** $g_h f \in F_E$ is the collection $\{h' \in f : h \not\preceq h'\} \cup \{[h, h'] : h' \in g\}$.

That is: $g_h f$ comprises all histories in f that do not weakly follow h , plus all histories obtained by concatenating h with histories in g . If $h = \emptyset$, then $g_h f = g$.

3 Axioms and Results

It is now possible to present the main results of this paper. Subsection 3.1 provides a characterization of Consistent Planning (Theorem 1); Subsection 3.2 shows that, under suitable assumptions, conditional preferences can be derived from prior preferences (Theorem 2).

From a formal standpoint, the material in Subsections 3.1 and 3.2 is essentially orthogonal. For instance, the assumptions required for the elicitation result of Theorem 2 allow for considerable departures from sophistication (and weak commitment); for another example, see Footnote 14 below.

All proofs are in the Appendix.

3.1 A decision-theoretic analysis of Consistent Planning

Subsection 3.1.1 formalizes the Sophistication and Weak Commitment axioms discussed in the Introduction, as well as an additional required axiom; the definition of Consistent Planning and the characterization result are provided in §3.1.2.

3.1.1 Sophistication, Weak Commitment and Simplification

As discussed in the Introduction, *Sophistication* reflects the assumption that the DM correctly anticipates her future preferences, and recognizes that she will not be able to carry out plans that involve conditionally dominated actions at future histories. To capture its implications solely in terms of the individual’s preferences, without explicitly modeling her introspective beliefs, it is enough to assume that *pruning conditionally dominated actions leaves the DM indifferent*.

The notation introduced in §2.3 makes it easy to formalize this assumption. Recall that any choice c available at a non-terminal history h of a tree f corresponds to a continuation tree, denoted $f(h, c)$.¹⁰ Thus, choice b (strictly) dominates choice w at history h if $f(h, b) \succ_{E(h)} f(h, c)$. If B is a collection of (strictly undominated) actions available at h , then $f(h, B)$ is the subtree of $f(h)$ corresponding to these undominated actions, and $f(h, B)_h f$ is the subtree of f obtained by pruning actions available at h but not in the set B .¹¹ The interpretation of the following axiom should now be straightforward:

Axiom 3.1 (Sophistication) *For all $f \in F_E$, all $h \in f \cap \mathcal{H}_E$, and all $B \subset C_f(h)$: if, for all $b \in B$ and $w \in C_f(h) \setminus B$, $f(h, b) \succ_{E(h)} f(h, w)$, then $f \sim_E f(h, B)_h f$.*

It is worth emphasizing that Axiom 3.1 also applies to $h = \emptyset$, the initial history. For this special case, the axiom is a version of the usual “strategic rationality” property of standard menu preferences (Kreps [27]).¹² 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). However, the availability of choices at the *initial* history of f that are deemed inferior given the same *initial* preference relation \succ_E 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 assumption makes it possible to focus on deviations from standard behavior due solely to differences in information and perceived ambiguity at distinct points in time; it abstracts away from differently motivated deviations.

It should also be noted that Axiom 3.1 is meaningful even in case the history h is “irrelevant” in the tree f . More precisely, two cases must be considered. First, suppose that the history h is not reached in

¹⁰Recall that the choice c may be a prize in X ; in this case $f(h, c)$ is a degenerate tree, as noted in the preceding subsection. This does not affect the discussion in the text.

¹¹Thus, not all dominated actions need be pruned.

¹²Recall that this section only considers finite trees.

the tree f because of choices the DM makes prior to h . In this case, a sophisticated DM should anticipate that h will not be reached; consequently, removing alternatives at h that the DM would not have chosen should not affect her evaluation of the tree f .¹³ This is precisely what Axiom 3.1 requires.

Second, if the DM does not expect the event $E(h)$ to occur, then she will be indifferent between trees that only differ in case this event occurs (indeed, this is what is usually meant by “null event”). The behavioral restriction imposed by Axiom 3.1 takes the form of an indifference, and so it is clearly consistent with this, regardless of how preferences conditional on $E(h)$ are defined.¹⁴

Analogous observations hold for the remaining two axioms, Weak Commitment and Simplification.

As noted in the Introduction, a tie-breaking rule may be required in addition to Sophistication. For definiteness, the assumption adopted here reflects a notion of *one-period-ahead commitment*. Refer back to the tree in Fig. 3, henceforth denoted f . Recall that, conditional upon learning that the ball drawn is red or green, the DM is indifferent between the actions R and G ; however, ex-ante, she would like to commit to R . In this case, it will be assumed that the DM “can” in fact carry out her plan and choose R upon learning that the event $\{r, g\}$ has occurred. More precisely, the DM evaluates the tree in Fig. 3 *as if* she could indeed commit to choosing R .

As was the case for Sophistication, this interpretation implicitly involves the DM’s beliefs about her own future preferences. However, it is possible to formalize one-period-ahead commitment solely in terms of the DM’s preferences. To do so, the notion of a *one-period-commitment version* of a tree is required. Again, refer to the tree in Fig. 3, and consider a modified tree where the action A at the initial history \emptyset is replaced by *two* actions, labelled AR and AG . If the event $\{r, g\}$ occurs and the DM has chosen AR , then only the action R is available to the DM: more precisely, the continuation tree following $[AR, \{r, g\}]$ equals $f([A, \{r, g\}], R)$. Similarly, the continuation tree following $[AG, \{r, g\}]$ equals $f([A, \{r, g\}], G)$. The resulting tree, denoted by g henceforth, is depicted in Fig. 5.

Relative to the original tree f , the modified tree g allows the DM to commit at the initial history to a specific choice of action in the following period. For this reason, g is deemed a one-period commitment version of the tree f .¹⁵ Axiom 3.2 below requires that the DM be *indifferent between f and its one-period-commitment version g at the initial node*. That is: the structure of the original tree f does not offer the possibility of commitment at the initial history; yet, the DM evaluates it f “as if she could actually

¹³On the other hand, removing alternatives at h that the DM would choose *might* lead the DM to reoptimize and take actions that no longer prevent history h from being reached. As can be seen, Axiom 3.1 is silent in these cases.

¹⁴The standard approach is to assume that $g \sim_{E(h)} g'$ for all $g, g' \in F_{E(h)}$ (which, for instance, follows from the characterization of conditional preferences in Sec. 3.2). Axiom 3.1 holds trivially in this case; indeed, notice that the Axiom would also be satisfied if \succsim_E itself was trivial (i.e. if E was null for the ex-ante preference). However, for instance, Myerson [30] assumes that non-trivial preferences conditional on every non-empty event are given. Axiom 3.1 accommodates this possibility, too.

¹⁵As will be clear momentarily, f admits many one-period commitment versions, but these only differ in the action labels they use. That is, one-period-commitment versions are unique up to relabeling.




Figure 5: One-period commitment version of Fig. 3

commit” to R (or, for that matter, G). Jointly with Sophistication, this axiom captures the intuition that the DM expects to be able to carry out her preferred one-period-ahead plan.

The tree in Fig. 3 has three simplifying characteristics. First, there are only two decision epochs—“time 0”, corresponding to the initial history \emptyset , and “time 1”, corresponding to history $[A, \{r, g\}]$. Thus, there is no distinction between one-period and full commitment. Second, there is only one time-1 history where the DM has to make a choice. Third, there is only one action available to the DM at the initial history. A general notion of one-period commitment must allow for more than two decision epochs, for multiple time-1 decision points, and for multiple actions available at \emptyset .

As a first step, the following definition identifies subtrees of a general tree f that feature exactly one action at every (non-terminal) time-0 and time-1 history, and otherwise coincide with f .

Definition 7 Consider an event $E \in \Sigma \setminus \{\emptyset\}$ and a tree $f \in F_E$. A tree $\bar{g} \in F_E$ **allows one-period commitment in f** if $\bar{g} \subset f$ and, for all $h \in \bar{g} \cap \mathcal{H}_E$, $\lambda(h) \leq 1$ implies $|C_{\bar{g}}(h)| = 1$, and $\lambda(h) = 2$ implies $\bar{g}(h) = f(h)$.

That is, a tree \bar{g} allows one-period commitment in f if (i) it is a subtree of f ; (ii) it features only one choice at the initial history, and at any non-terminal history of length one; and (iii) agrees with f otherwise.

The one-period commitment version of a tree f can now be defined as a new tree g that, loosely speaking, “contains” *all* subtrees \bar{g} that allow one-period commitment in f (and no other subtree). Notice that one-period commitment versions are only identified up to relabeling.

Definition 8 A tree $g \in F_E$ is a **one-period commitment version of $f \in F_E$** iff

- (i) for every tree \bar{g} that allows one-period commitment in f there is a unique $a \in C_g(\emptyset)$ such that $g(\emptyset, a) \approx \bar{g}$; and
- (ii) for every $a \in C_g(\emptyset)$ there is a tree \bar{g} that allows one-period commitment in f such that $g(\emptyset, a) \approx \bar{g}$.

Notice that, since A is assumed to be countably infinite, every (finite) tree $g \in F_E$ admits at least one (indeed, infinitely many) trees that satisfy Def. 8.

The notion of one-period-ahead commitment discussed above can now be formalized. Consider a tree f and a history $h \in f$, and suppose that, at every history h' that immediately follows h , the DM is conditionally indifferent among all available actions; then replacing the continuation tree $f(h)$ with one of its one-period-commitment versions g must leave the DM indifferent ex-ante:

Axiom 3.2 (Weak Commitment) *For all $f \in F_E$ and all histories $h \in f$: if, for all $h' \in f$ with $h < h'$ and $\lambda(h') = \lambda(h) + 1$, and all $c, c' \in C_f(h')$, $f(h', c) \sim_{E(h')} f(h', c')$, then $f \sim_E g_h f$ for all one-period commitment versions $g \in F_{E(h)}$ of $f(h)$.*

One last simple axiom is required. Consider a tree $f \in F_E$, and suppose that the DM is indifferent among all actions available at the initial history. Then, again *at the initial history*, the tree f should be deemed equivalent to a subtree f' where one or more (but not all!) initial actions have been removed.

Axiom 3.3 (Simplification) *For all $f \in F_E$: if, for all $c, c' \in C_f(\emptyset)$, $f(\emptyset, c) \sim_E f(\emptyset, c')$, then for all non-empty $B \subset C_f(\emptyset)$, $f \sim_E f(\emptyset, B)$.*

Two observations are in order. First, Simplification is not implied by the previous axioms. Second, consider a tree $f \in F_E$ and a non-terminal history h of f . If the DM is indifferent among all actions available at h , then Simplification requires that $f(h) \sim_{E(h)} f(h, A)$ for all non-empty subsets A of actions at h . However, ex-ante, it may well be the case that the DM would strictly prefer, to have certain actions removed or not removed at h . Refer to the tree in Fig. 3: conditional upon observing $\{r, g\}$, the DM is indifferent between R and G ; however, ex ante, she strictly prefers that action R be available at the second decision point. Formally, if $h = [A, \{r, g\}]$, then, consistently with Simplification, $f(h) \sim_{\{r, g\}} f(h, G)$: however, $f \succ f(h, G)_h f$, i.e. a priori the DM dislikes being forced to choose G at h .

3.1.2 Formulation and Characterization of Consistent Planning

As noted in the Introduction, Strotz's notion of "consistent planning" corresponds to backward induction with a particular tie-breaking rule. The following definition provides the details.

Definition 9 (Consistent Planning) Consider a tree $f \in F_E$. For every terminal history $h \in f \cap \mathcal{T}_E$, let $CP_f(h) = \{f(h)\}$. Inductively, if $h \in f \cap \mathcal{H}_E$ and $CP_f(h')$ has been defined for all $h' \in f$ with $h < h'$, let

$$\begin{aligned} CP_f^0(h) &= \left\{ p \in F_{E(h)}^p : \exists c \in C_f(h) \text{ s.t. } C_p(\emptyset) = \{c\}, \mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h, c) \text{ and} \right. \\ &\quad \left. \forall D \in \mathcal{F}_p(\emptyset, c), p([c, D]) \in CP_f([h, c, D]) \right\} \text{ and} \\ CP_f(h) &= \left\{ p \in CP_f^0(h) : \forall p' \in CP_f^0(h), p \succ_{E(h)} p' \right\}. \end{aligned}$$

A plan $p \in F_E$ is a **consistent-planning solution** of f if $p \in \text{CP}_f(\emptyset)$.

Consistent Planning inductively associates a set of continuation *plans* to each history in a decision tree f . Def. 9 is modeled after analogous definitions in Strotz [41] and Gul and Pesendorfer [15], except that it is phrased in terms of preferences, rather than via their numerical representation. To further clarify how Consistent Planning operates, it is useful to rephrase Def. 9 as an algorithm, as follows.

1. First, define $\text{CP}_f(\cdot)$ for each terminal history h of f as the singleton set consisting solely of the (degenerate) plan $f(h)$, which corresponds to the outcome that f delivers at h (cf. §2.1).
2. Next, suppose that CP_f has been defined for a collection of histories $f^{\text{done}} \subset f$.¹⁶ Now define $\text{CP}_f(\cdot)$ for histories h that are immediately followed only by histories in f^{done} , as follows:
 - (a) Define $\text{CP}_f^0(h)$ as the set of plans that choose one of the alternatives available in tree f at history h , and then continue in a manner consistent with previous iterations of the algorithm;
 - (b) Next, define $\text{CP}_f(h)$ as the set of $\succ_{E(h)}$ -best elements in $\text{CP}_f^0(h)$.
3. Repeat Step 2 until $\text{CP}_f(\cdot)$ has been defined for all histories in f .

Four features and immediate consequences of the preceding definition need to be emphasized:

1. *Consistent Planning only requires that complete and transitive preferences over plans be specified:* to see this, observe that the definition of $\text{CP}_f(h)$ only requires comparisons of plans, not general trees. This feature of Def. 9 is essential, as a main objective of the present approach is to employ Consistent Planning to extend preferences from plans to arbitrary trees.

2. *For all $f \in F_E$, the set $\text{CP}_f(\emptyset)$ of consistent-planning solutions of f is non-empty,* provided preferences over plans are complete and transitive. This follows immediately by noting that each set $\text{CP}_f^0(h)$ is finite. See §4.3 for a discussion of infinite trees.

3. *For every history h , the plans in the set $\text{CP}_f(h)$ are mutually indifferent conditional on $E(h)$. In particular, all consistent-planning solutions of a tree are mutually indifferent.* Provided the restriction of each conditional preference \succ_E to plans is complete and transitive, this follows directly from Def. 9, regardless of whether or not the system of conditional preferences under consideration satisfies the axioms of this section. Again, if this was not the case, it would not be possible to employ Consistent Planning to extend preferences over plans to preferences over trees.

¹⁶To clarify, f^{done} is typically *not* a tree itself. For instance, immediately after completing the first step in the algorithm, f^{done} is defined as the set of all terminal histories in f .

4. *Consistent Planning* does not replace optimal continuation plans with a “continuation value”, as was noted in Sec. 1.2 of the Introduction. Rather, the Consistent Planning procedure iteratively deletes inferior continuation plans.¹⁷

The main result of this paper may now be stated. It states that preferences over plans can be extended to preferences over trees via Def. 9 *if and only if* the system of conditional preferences under consideration satisfies Sophistication, Weak Commitment and Simplification.

Theorem 1 *Consider a system of preferences $\{\succ_E\}_{\emptyset \neq E \in \Sigma}$ that satisfy Assumption 2.1 and such that, for every non-empty $E \in \Sigma$, \succ_E is a complete and transitive binary relation on F_E^p . Then the following statements are equivalent.*

1. *For every non-empty $E \in \Sigma$, \succ_E is complete and transitive on all of F_E ; furthermore, Axioms 3.1, 3.2 and 3.3 hold;*
2. *for any $E \in \Sigma$, and every pair of trees $f, g \in F_E$: $f \succ_E g$ if and only if $p \succ q$ for some (hence all) $p \in CP_f(\emptyset)$ and $q \in CP_g(\emptyset)$.*

3.2 Eliciting Conditional Preferences

As can be seen from Eq. (4) in the Introduction, Dynamic Consistency also provides a way to elicit conditional preferences *over acts* from prior preferences *over acts*; indeed, it characterizes Bayesian updating for expected-utility preferences. This subsection establishes a similar result in the setting of choice among trees: Sophistication provides a way to elicit conditional preferences *over acts and trees* from prior preferences *over trees*. A weak form of Sophistication actually suffices.

3.2.1 Revealed Conditional Preferences over Trees

To this end, the notion of *test tree* is useful. Consider an event $E \in \Sigma$. In order to elicit preferences conditional on E from prior preferences, one clearly needs to consider trees in which the DM *can* potentially receive the information that E has occurred, and has to make a choice conditional upon E : formally, one is led to focus on trees $f \in F_\Omega$ such that $E(h) = E$ for some $h \in f \cap \mathcal{H}$. But even if a tree f contains one such history h , it may be the case that the DM wishes to avoid reaching them, and is able to do so by choosing suitable actions at histories preceding h . An E -test tree is, intuitively, the simplest tree in which the DM “has” to face a decision conditional on E (provided of course E obtains): there is a single initial action, and E is one of the events that may follow it. Formally:

¹⁷Indeed, since no solvability assumption is made in this subsection, conditional certainty equivalents may fail to exist.

Definition 10 Let $E \in \Sigma$. A tree $g \in F_\Omega$ is an E -test tree iff $C_g(\emptyset) = \{a\}$ for some $a \in A$ and $E \in \mathcal{F}_g(\emptyset, a)$. The set of E -test trees is denoted by T_E ; furthermore, for all $f \in F_E$ and $g \in T_E$, “ $f_E g$ ” denotes the tree $f_{[a,E]}g$, where $A_g(\emptyset) = \{a\}$.

Also, as in the case of expected-utility preferences, conditioning events must “matter” to the DM in order for any updating or elicitation rule to be meaningful; the relevant notion is formalized below.

Definition 11 An event $E \in \Sigma$ is **non-null** (for \succ) iff, for all trees $g \in T_E$, and for all $x, x' \in X$ with $x \succ x'$, $x_E g \succ x'_E g$.

The following notation simplifies the statement of the axioms and the proposed updating rule.

Definition 12 Let $E \in \Sigma$ be non-null; fix $f \in F_E$ and $x \in X$. Then $\{f, x\}$ denotes the tree $f \cup \{[x]\} \in F_E$.

Definitions 10 and 12 will often be used together, as in “ $\{f, x\}_E g$ ” for a test tree $g \in T_E$.

As an example, consider the tree f_x in Fig. 1, its subplans R and G starting at $E = \{r, g\}$, and the trees $AR_{x,y}$ and $AG_{x,y}$ in Fig. 4. Then f_x , $AR_{x,y}$ and $AG_{x,y}$ are all E -test trees. Furthermore, $AR_{x,y} = \{R, y\}_E f_x$ and $AG_{x,y} = \{G, y\}_E f_x$. Also, according to the preferences in §1.1, the event E is not null.

It is now possible to formalize the notion of revealed conditional preferences discussed in §1.3.

Definition 13 Let $E \in \Sigma$ be non-null and consider $f, f' \in F_E$ and $g \in T_E$. Then f is **revealed weakly preferred to f' given E and g** , written $f \succ_{E,g}^* f'$, iff there exists $z \in X$ such that

$$\forall y \in X, \quad y \succ z \Rightarrow \{f', y\}_E g \sim y_E g \quad \text{and} \quad z \succ y \Rightarrow \{f, y\}_E g \sim f_E g.$$

Again, refer to §1.3; note that the plans AR_x and AG_x in Fig. 2 can be written as $R_E f_x$ and $G_E f_x$ respectively. Then, taking $g = f_x$, $f = R$ and $f' = G$, the condition in Def. 13 corresponds to Eq. (8).

Two observations are in order. First of all, the above is simply a definition: without further assumptions, there is no guarantee that the relation $\succ_{E,g}^*$ will be well-defined, that it will not depend upon the choice of test tree g , and in particular that it will coincide with the DM’s “actual” conditional preference \succ_E , if one is given. However, as Theorem 2 below states, revealed conditional preferences *do* enjoy these properties under suitable (and relatively weak) axioms.

Second, Def. 13 and the discussion that motivates it purportedly focus on the *strict* preferences $y \succ z$, $f \succ_E y$ and $y \succ_E f'$, $z \succ y$. This avoids introducing specific assumptions about tie-breaking: in particular, it is *not* necessary to assume even a restricted version of Weak Commitment (Axiom 3.2 in Sec. 3.1.1). In addition to basic solvability and “taste consistency” requirements, only a limited form of Sophistication is required to elicit conditional preferences. Incidentally, avoiding Weak Commitment makes the extension of Theorem 2 to infinite trees straightforward.¹⁸

¹⁸Also, assuming Weak Commitment does not substantially simplify the definition of revealed conditional preference.

3.2.2 Axioms and Characterization

Turn now to the formal characterization of conditional preferences. First, I formulate four “linkage” axioms, relating prior and conditional preferences. The first axiom states that preferences over prizes (constant trees) are unaffected by conditioning.¹⁹

Axiom 3.4 (Stable Tastes) For all $x, x' \in X$, and all non-null $E \in \Sigma$: $x \succ_E x'$ if and only if $x \succ x'$.

The following two standard axioms ensure that conditional certainty equivalents exist (recall that X is assumed to be a connected and separable topological space).

Axiom 3.5 (Conditional Dominance) For all non-null $E \in \Sigma$, $f = (E, H, \xi) \in F_E$, and all $x', x'' \in X$: if $x' \succ \xi(z) \succ x''$ for all terminal histories z of f , then $x' \succ_E f \succ_E x''$.

Axiom 3.6 (Conditional Prize-Tree Continuity) For all non-null $E \in \Sigma$ and all $f \in F_E$, the sets $\{x \in X : x \succ_E f\}$ and $\{x \in X : x \preceq_E f\}$ are closed in X .

Finally, I impose a relatively mild, but essential *sophistication* requirement. Informally, it assumes “just enough” sophistication to ensure that the argument motivating Def. 13 is actually correct.

Axiom 3.7 (Weak Sophistication) For all non-null $E \in \Sigma$, $g \in T_E$, $f \in F_E$, and $x \in X$:

$$x \succ_E f \Rightarrow \{f, x\}_E g \sim x_E g \quad \text{and} \quad x \prec_E f \Rightarrow \{f, x\}_E g \sim f_E g.$$

Axiom 3.7 is considerably weaker than the Sophistication axiom considered in Section 3.1. For instance, consider a tree g that features two actions a, b at the initial history, such that a may be followed by E (formally, $[a, E] \in g$). Then, even if one assumes that Axiom 3.7 holds, it is still possible that $x \succ_E f$ and $\{f, x\}_{[a, E]} g \sim f_{[a, E]} g \succ x_{[a, E]} g$: in words, the DM *naively* expects to be able to stick to her ex-ante preferred choice of f . Thus, the present approach makes it possible to address the distinct issues of elicitation and sophistication in a relatively independent way.

Next, I formalize three “structural” axioms on prior preferences that ensure that the proposed definition of revealed conditional preferences is well-posed. It may be helpful to recall that Savage’s Sure-Thing Principle (Postulate P2) plays a similar role for his definition of conditional preferences.²⁰

The first two axioms ensure the existence of “conjectural” certainty equivalents; they may be viewed as the prior-preference counterpart to Axioms 3.5 and 3.6 above:

¹⁹For 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 3.4 to constant streams.

²⁰Actually, Savage himself restates P2 as follows: “for any [pair of acts] f and g and for every [event] B , $f \leq g$ given B or $g \leq f$ given B ” [35, inside back cover].

Axiom 3.8 (Prize Continuity) For all $\bar{x} \in X$, the sets $\{x \in X : x \succ \bar{x}\}$ and $\{x \in X : x \preccurlyeq \bar{x}\}$ are closed in X .

Axiom 3.9 (Conjectural Dominance) Consider a non-null $E \in \Sigma$, $f \in F_E$, $g \in T_E$ and $x \in X$. Then:

- (i) if $\xi(z) \succ x$ for all terminal histories z of f , then $\{f, x\}_E g \sim f_E g$;
- (ii) if $\xi(z) \prec x$ for all terminal histories z of f , then $\{f, x\}_E g \sim x_E g$.

Notice that Axiom 3.9 reflects considerations of sophistication and stability of preferences over outcomes. To elaborate, if the individual's preferences over X do not change when conditioning on E , then in (i) she will not choose x after observing that E has occurred, because f yields strictly better outcomes at every terminal history;²¹ similarly, in (ii), she will never choose f conditional on E . The indifferences in (i) and (ii) thus reflect the assumption that the individual correctly anticipates her future choices.

The key assumption on unconditional preferences is a counterpart to Weak Sophistication:

Axiom 3.10 (Conjectural Separability) Consider a non-null $E \in \Sigma$, $f \in F_E$, $g, g' \in T_E$ and $x, y \in X$. Then:

- (i) $\{f, y\}_E g \not\sim f_E g$ and $x \succ y$ imply $\{f, x\}_E g' \sim x_E g'$;
- (ii) $\{f, y\}_E g \not\sim y_E g$ and $x \prec y$ imply $\{f, x\}_E g' \sim f_E g'$.

To interpret this axiom, consider first the case $g = g'$ and fix a prize y . According to the logic of sophistication, the hypothesis that $\{f, y\}_E g \not\sim f_E g$ indicates that the DM believes that she will *not* strictly prefer f to y given E —otherwise indifference would have to obtain. Thus, if $x \succ y$ and the DM's preferences over X are stable, she will also strictly prefer x to f given E ; now sophistication yields the conclusion that $\{f, x\}_E g \sim x_E g$. The interpretation of (ii) is similar.

Additionally, Axiom 3.10 implies that these conclusions are independent of the particular test tree under consideration, and hence of what the decision problem looks like if the event E does not obtain. In this respect, Axiom 3.10 reflects a form of “separability,” much in the spirit of Savage's Postulate P2. More generally, just like P2 (cf. Footnote 20), Axiom 3.10 essentially requires that the revealed conditional preference relation $\succ_{E,g}^*$ be well-defined and independent of the test-tree g .

The main characterization result can finally be stated.

Theorem 2 Consider the conditional preference system $\{\succ_E\}_{0 \neq E \in \Sigma}$. Assume that \succ is a complete and transitive relation on F_Ω . Then the following statements are equivalent.

²¹ Recall that the trees under considerations are finite. For infinite trees, Axiom 3.9 needs to be modified slightly [e.g. in (i) one needs to assume that $\xi(z) \succ x'$ for some $x' \succ x$, and similarly in (ii)]. However, it should be emphasized that this is the *only* axiom that requires adjusting to extend the results of this section to infinite trees.

1. \succsim satisfies Axioms 3.8, 3.10 and 3.9; furthermore, for all non-null $E \in \Sigma$ and all $f, f' \in F_E$, $f \succsim_E f'$ if and only if $f \succsim_{E,g}^* f'$ for some (hence all) $g \in T_E$.
2. For every non-null $E \in \Sigma$, \succsim_E is complete and transitive, and satisfies Axioms 3.4, 3.5, 3.6 and 3.7.

This result can be interpreted as follows. If a given unconditional preference satisfies the above “structure” axioms, and conditional preferences are *defined* from it as in Def. 13, then the resulting system of conditional preferences satisfies the “linkage” axioms. Conversely, if one assumes that the DM is characterized by a system of conditional preferences that satisfy the above “linkage” axioms, one can *elicit* these preferences via Def. 13; furthermore, it will also be the case that the DM’s unconditional preference satisfies the “structure” axioms.

Theorem 2 mirrors a similar statement that applies to Savage’s updating rule for preferences over acts. In that setting, if prior preferences satisfy the Sure-Thing Principle (a “structure” axiom), then conditional preferences defined via Savage’s updating rule satisfy Dynamic Consistency and Relevance (the “linkage” axioms; Relevance requires that the DM be conditionally indifferent between any two acts that agree on the conditioning event). Conversely, if the “linkage” axioms of Dynamic Consistency and Relevance hold, conditional preferences can be elicited from prior preferences via Savage’s rule; furthermore, prior preferences satisfy the “structural” Sure-Thing axiom. See Ghirardato [11].

4 Examples and Extensions

4.1 Consistent Planning for MEU preferences and Full Bayesian Updating

As noted in the Introduction, a main objective of the present paper is to show that sophisticated dynamic choice can be guaranteed for any choice of preference model and updating rule, regardless of whether or not the latter yields dynamically consistent preferences over Savage acts.

This subsection shows how to apply the proposed approach to the MEU decision model (Gilboa and Schmeidler [12]) and prior-by-prior, or “full” Bayesian updating. While this is not the only possible updating rule for MEU preferences (see the next subsection for examples and references), it has received considerable attention in the literature. As will be clear, it is straightforward to adapt the analysis in this subsection to different representations of preferences (e.g. Choquet-expected utility) and different updating rules (e.g. the Dempster-Shafer rule).

I begin by formalizing the assumption that preferences are consistent with MEU. As a preliminary step, recall that a tree $f \in F_E$ is a *plan* iff, for every $h \in F_E$, $C_f(h)$ is a singleton. This immediately implies that every state ω determines a unique path through the tree f : formally, for every $\omega \in E$, there is a unique $h \in f \cap \mathcal{T}_E$ such that $\omega \in E(h)$. Throughout this subsection, for every plan $f \in F_E$ and $\omega \in E$, the notation $f(\omega)$ indicates the prize $\xi(h)$, where h is the unique terminal history of f with $\omega \in E(h)$.

The required assumption on preferences can now be stated.

Assumption 4.1 (Non-trivial MEU with reduction) 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 $f, g \in F_\Omega$,

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

Moreover, there exist plans $f, g \in F_\Omega$ such that $f \succ g$.

If prior preferences satisfy Assumption 4.1 and the event E is not null (cf. Def. 11), then, according to prior-by-prior Bayesian updating, conditional preferences \succ_E over plans are defined stipulating that, for all plans $f, g \in F_E$,

$$f \succ_E g \iff \min_{q \in C} \int_E u(f(\omega))q(d\omega|E) \geq \min_{q \in C} \int_E u(g(\omega))q(d\omega|E).$$

Thus, one can adapt Def. 9 to obtain a version of Consistent Planning specific to MEU preferences and prior-by-prior updating; the details are as follows. Consider a tree $f \in F_E$. For every terminal history $h \in f \cap \mathcal{T}_E$, let $\text{CPMEU}_f(h) = \{f(h)\}$. Inductively, if $h \in f \cap \mathcal{H}_E$ and $\text{CPMEU}_f(h')$ has been defined for all $h' \in f$ with $h < h'$, let

$$\begin{aligned} \text{CPMEU}_f^0(h) &= \left\{ p \in F_{E(h)} : \exists c \in C_f(h) \text{ s.t. } C_p(\emptyset) = \{c\}, \mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h, c) \text{ and} \right. \\ &\quad \left. \forall D \in \mathcal{F}_p(\emptyset, c), p([c, D]) \in \text{CP}_f([h, c, D]) \right\} \text{ and} \\ \text{CPMEU}_f(h) &= \left\{ p \in \text{CPMEU}_f^0(h) : \forall p' \in \text{CPMEU}_f^0(h), \right. \\ &\quad \left. \min_{q \in C} \int_E u(p(\omega))q(d\omega|E) \geq \min_{q \in C} \int_E u(p'(\omega))q(d\omega|E) \right\}. \end{aligned}$$

where u and C are as in Assumption 4.1. In the next subsection, this algorithm will be used to analyze a model of information acquisition.

Prior-by-prior Bayesian updating can be characterized via axioms on conditional preferences over acts; see e.g. Jaffray [20] and Pires [33].²² It will now be shown that a version of the main axiom in Pires [33] and Siniscalchi [38], augmented with the axioms in Sec. 3.1.1, yields a characterization of consistent planning for MEU preferences and full Bayesian updating. This may be viewed as a specialization of Theorem 1, just like the definition of $\text{CPMEU}_f(\cdot)$ specializes Def. 9.

To fix ideas, it is useful to consider the standard Dynamic Consistency axiom as a starting point:

²²Siniscalchi [39] characterizes prior-by-prior updating for a broad class of preferences. Also, although it is not explicitly decision-theoretic, the earliest characterization is probably due to Walley [42].

Axiom 4.1 (Dynamic Consistency) For all non-null $E \in \Sigma$, all $g \in T_E$, and all plans $f, f' \in F_E$:

$$f \succ_E f' \iff f_E g \succ f'_E g.$$

The key axiom characterizing prior-by-prior updating has a much narrower scope:

Axiom 4.2 (Constant-Act Dynamic Consistency) For all non-null $E \in \Sigma$, all plans $f \in F_E$ and $x \in X$: if $g \in T_E$, $C_g(\emptyset) = \{a\}$, $\mathcal{F}_g(\emptyset, a) = \{E, \Omega \setminus E\}$ and $[a, \Omega \setminus E, x] \in g$, then

$$f \succ_E x \iff f_E g \succ x_E g.$$

Axiom 4.2 differs from Axiom 4.1 in two important respects. First, Axiom 4.2 only considers conditional comparisons between a tree f and a prize x , whereas Axiom 4.1 has implications whenever two arbitrary trees f and f' are compared conditional on an essential event E . Second, the E -test trees g considered in Axiom 4.2 are rather special: following the (unique) initial action a , the individual can observe either the event E or the event $\Omega \setminus E$; in the latter case, she receives the prize x .²³ Thus, in particular, the tree $x_E g$ is a plan, and by Assumption 4.1 the DM deems it equivalent to the prize x itself.²⁴ No such restriction is imposed on test trees in Axiom 4.1. The motivations for these restrictions are discussed in the sources cited above (see especially [33] and [38]).

The counterpart to Theorem 1 can then be stated.

Theorem 3 Consider a system of preferences $\{\succ_E\}_{\emptyset \neq E \in \Sigma}$ that satisfies Assumption 2.1. Suppose that Assumption 4.1 holds, and that every event $E \in \Sigma \setminus \{\emptyset\}$ is non-null. Then the following statements are equivalent.

1. For every $E \in \Sigma \setminus \{\emptyset\}$, \succ_E is complete and transitive on all of F_E ; furthermore, Axioms 3.1, 3.2, 3.3 and 4.2 hold;
2. for any $E \in \Sigma \setminus \{\emptyset\}$ and every pair of trees $f, g \in F_E$: $f \succ_E g$ if and only if $p \succ_E q$ for some $p \in \text{CPMEU}_f(\emptyset)$ and $q \in \text{CPMEU}_g(\emptyset)$.

4.2 Sophistication and the Value of Information

This subsection analyzes a simple model of information acquisition. This example provides an application of the analytic framework proposed in this paper. From a substantive point of view, it highlights a basic trade-off between information acquisition and commitment.

²³Formally, there are many such test trees; these only differ in the label attached to the initial action, and in the (irrelevant) continuation trees specified after E occurs.

²⁴The reader may wonder why the axiom was not stated in the simpler form: for all $f \in F_E$ and $x \in X$, " $f \succ_E x$ if and only if $f_E x \succ x$ ". The problem is that " $f_E x$ " would need a separate definition: recall that the prize x is identified with the tree $\{\emptyset, [x]\}$, which is, formally, not an E -test tree, and therefore does not allow one to invoke the notation in Def. 10.

4.2.1 A Parametric Information-Acquisition Model

Consider an individual facing a choice between two alternative actions, a and b . 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 individual observes the first coordinate of the prevailing state.

Thus, at any state (ω_1, ω_2) , payoffs are determined by ω_2 , which the individual does not observe; however, she does observe the signal ω_1 .

Assume that the individual reduces plans to acts, has MEU preferences over acts, and linear utility over monetary prizes. Beliefs are represented by the set of priors $C = \{\alpha P + (1 - \alpha)Q : \alpha \in [0, 1]\}$, where $P, Q \in \Delta(\Omega)$ are defined in Table 3.

Prior	(α, α)	(α, β)	(β, α)	(β, β)
P	$1 - 2\varepsilon$	ε	ε	0
Q	0	ε	ε	$1 - 2\varepsilon$

Table 3: The priors P and Q ; $\varepsilon \in (0, \frac{1}{4})$.

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 that the individual updates her beliefs prior-by-prior. The set C does not satisfy the “rectangularity” condition of Epstein and Schneider [9] with respect to the filtration corresponding to the information-acquisition problem under consideration; thus, dynamic consistency does not hold.²⁵

The objective is to determine the value of the information conveyed by the signal ω_1 . However, since preferences are dynamically inconsistent, *it is crucial to specify whether or not the individual can commit to specific actions contingent upon the realization of the signal ω_1 .*

If the individual can commit, then she effectively faces a choice among four *plans*, denoted p_{aa}, p_{ab}, p_{ba} and p_{bb} : for instance, p_{ab} is the plan that prescribes the choice a after seeing $\omega_1 = \alpha$ and the choice b after observing $\omega_1 = \beta$. Table 4 indicates the prize delivered by each of these plans in every state, as well as their MEU evaluation (assuming Reduction).

²⁵For completeness, the relevant filtration is $\mathcal{F}_0, \mathcal{F}_1$, where $\mathcal{F}_0 = \{\Omega\}$ and $\mathcal{F}_1 = \{(\alpha, \alpha), (\alpha, \beta)\}, \{(\beta, \alpha), (\beta, \beta)\}$.

Plan	(α, α)	(α, β)	(β, α)	(β, β)	MEU
p_{aa}	H	L	H	L	$(1 - \varepsilon)L + \varepsilon H$
p_{bb}	L	H	L	H	$(1 - \varepsilon)L + \varepsilon H$
p_{ab}	H	L	L	H	$2\varepsilon L + (1 - 2\varepsilon)H$
p_{ba}	L	H	H	L	$(1 - 2\varepsilon)L + 2\varepsilon H$

Table 4: Plans under full commitment.

Since $0 < \varepsilon < \frac{1}{4}$, the individual's ranking of the four plans in Table 4 is

$$p_{ab} \succ p_{ba} \succ p_{aa} \sim p_{bb}.$$

Furthermore, note that, if the individual does not acquire information, her feasible choices are the two acts corresponding to p_{aa} and p_{bb} respectively: thus, the value of the signal ω_1 under commitment is $(1 - 3\varepsilon)(H - L)$.

If the individual cannot commit, then the individual's dynamic-choice problem is represented by the tree f depicted in Fig. 6.²⁶




Figure 6: The tree f (value of information without commitment).

If the individual is sophisticated,²⁷ Theorem 1 implies that she deems f equivalent to its Consistent-Planning solution—which, as will be seen momentarily, is unique in this problem. To compute it, note

²⁶As a matter of notation, one must label the sole (dummy) initial action; the symbol “*” is used here.

²⁷Weak Commitment is not required in this particular problem.

that the set of posteriors given each realization of the signal ω_1 is obtained by updating P and Q , then taking the convex hull of the resulting probabilities. These posteriors are given by Table 5.

	$\omega_2 = \alpha$	$\omega_2 = \beta$		$\omega_2 = \alpha$	$\omega_2 = \beta$
$P(\cdot \omega_1 = \alpha)$	$\frac{1-2\varepsilon}{1-\varepsilon}$	$\frac{\varepsilon}{1-\varepsilon}$	$P(\cdot \omega_1 = \beta)$	1	0
$Q(\cdot \omega_1 = \alpha)$	0	1	$Q(\cdot \omega_1 = \beta)$	$\frac{\varepsilon}{1-\varepsilon}$	$\frac{1-2\varepsilon}{1-\varepsilon}$

Table 5: The individual's posterior beliefs given $\omega_1 = \alpha$ (left) and $\omega_1 = \beta$ (right).

It is immediate to see that *the DM will choose b after observing $\omega_1 = \alpha$ and a after observing $\omega_1 = \beta$* . To understand the motivation for these choices, consider for definiteness the case $\omega_1 = \alpha$. A priori, the individual believes that ω_1 equals ω_2 with high probability, *but* she also feels that the relative likelihood of $\omega_2 = \alpha$ vs. $\omega_2 = \beta$ is highly ambiguous. On the other hand, she assigns a small, but positive and unambiguous prior probability ε to the event that the signal ω_1 is actually wrong. When contemplating the choice of a after seeing $\omega_1 = \alpha$, the latter consideration looms large: the individual is convinced that the signal might be wrong, whereas, for all she knows, the event $\omega_2 = \alpha$ may actually be extremely unlikely. Thus, this ambiguity-averse individual perceives a as a more “dangerous” choice than b . A similar argument applies to the case $\omega_1 = \beta$.

Summing up, p_{ba} is the unique Consistent-Planning solution of the tree f . From Table 4, the value of the signal ω_1 when the DM cannot commit is then $\varepsilon(H - L)$. Since $\varepsilon \in (0, \frac{1}{4})$, it follows that the value of information is positive, but smaller than in the commitment case.

4.2.2 Value of Information and Value of Commitment

The above example suggests that acquiring information has two consequences for sophisticated but (possibly) dynamically inconsistent decision-makers. On one hand, information *potentially* leads to more fine-grained decisions; in this sense, it is “intrinsically” beneficial, regardless of ambiguity attitudes or other features of preferences. On the other hand, the *actual choices* made contingent upon the realization of a signal may fail to be optimal from the ex-ante point of view. Conversely, not acquiring information entails a cost (coarser decisions), but also yields a benefit—partial commitment.

In the example, without acquiring any information, the feasible set of plans available to the individual is (up to Reduction) $\{p_{aa}, p_{bb}\}$. If the DM acquires information and can commit, then the set of feasible plans is $\{p_{aa}, p_{ab}, p_{ba}, p_{bb}\}$. Finally, if she acquires information but cannot commit, then the feasible set only consists of the plans she will actually follow, i.e. $CP_f(\emptyset) = \{p_{ba}\}$. The fact that both information and commitment are intrinsically valuable is simply a consequence of the fact that the $\{p_{aa}, p_{bb}\}$ and $CP_f(\emptyset)$ are both subsets of $\{p_{aa}, p_{ab}, p_{ba}, p_{bb}\}$. Similarly, the fact that $CP_f(\emptyset)$ and $\{p_{aa}, p_{bb}\}$ are not ordered by inclusion indicates that, in general, it may or may not be beneficial for the individual to ac-

quire information. For the preferences considered here, information is valuable; however, the example can be modified so that not acquiring information is optimal; see also Epstein and Le Breton [8].

The decision whether or not to acquire a signal (and what price to pay for it) is thus driven by a trade-off between the intrinsic value of information and the value of commitment; this trade-off is captured by Consistent Planning. In particular, if the commitment problem is severe enough, the DM may rationally choose to pay a price so as to *avoid* making choices after observing the realization of a particular signal, despite the fact that information is intrinsically valuable. Similar patterns of behavior emerges in other related contexts featuring time-inconsistent but sophisticated decision-makers: see e.g. Carrillo and Mariotti [3] and references therein.

The framework proposed in this paper allows a straightforward formalization of these observations in a general setting: see Section B.1 in the Online Appendix for details.

4.3 Consistent Planning for Simple Infinite Trees

In the context of time-inconsistent choice under conditions of certainty, it is well-known (e.g. Peleg and Yaari [32]; Gul and Pesendorfer [15]) that consistent-planning solutions may fail to exist for certain infinite decision trees. Similar situations can arise in the present setting as well. This subsection provides an example, and uses it to motivate an extension of the axioms in Sec. 3.1.1 that addresses this issue for a class of “simple infinite games” in a fully decision-theoretic framework.




Figure 7: The subtree $f(\theta, a)$ corresponding to action $a \in [0, 1]$.

Let $\Omega = \{\alpha, \beta, \gamma, \delta\}$ and $X = \mathbb{R}$; consider a tree f such that $C_f(\emptyset) = [0, 1]$ and, for every $a \in [0, 1]$, the continuation tree $f(\theta, a)$ corresponding to the choice of a is as depicted in Fig. 7. Assume that the DM

has MEU preferences with linear utility and priors C given by the convex hull of $P, Q \in \Delta(\Omega)$, where

$$P(\{\beta\}) = P(\{\gamma\}) = \frac{1}{2} \quad \text{and} \quad Q(\{\alpha\}) = Q(\{\delta\}) = \frac{1}{2}.$$

Finally, assume that the DM adopts the prior-by-prior Bayesian updating rule.

Consider the last (i.e. rightmost) decision history. With these assumptions, the minimum expected payoffs conditional on $\{\alpha, \beta\}$ are $-a$ for t and 0 for b , the DM will strictly prefer b if $a > 0$, and will be indifferent between t and b if $a = 0$.

Next, consider the middle decision node. Action s is expected to yield 4 regardless of the value of a , whereas c is expected to yield 0 if $a > 0$ (because c will be followed with b) and, by Weak Commitment and Sophistication, 5 if $a = 0$ (because she can commit to t at the last decision history). Therefore, at the middle decision node, the individual will choose s if $a > 0$ and c if $a = 0$.

Now consider the point of view of the DM when she is at initial node and chooses $a \in [0, 1]$. If $a > 0$, then she expects to continue with s at the middle node, so the tree $f(\emptyset, a)$ evaluates to $10 - a \geq 9$. If instead $a = 0$, then she expects to continue with c and t , which she evaluates at 5. It is then clear that there is no $a \in [0, 1]$ that maximizes the DM's ex-ante minimum expected payoffs. Consequently, applying Def. 9 yields an empty set of “consistent-planning solutions” for the tree f .

Observe that the DM effectively faces a “non-compact” set of alternative continuation plans at the initial history.²⁸ This is mainly due to the specific tie-breaking rule implicit in consistent planning; introducing some flexibility in the tie-breaking rule resolves the issue. The “multi-selves” approach, which is common in the literature on time-inconsistent preferences under certainty (see e.g. Peleg and Yaari [32]), provides a way to formalize this notion of flexibility.²⁹

However, there is an alternative resolution to the difficulties highlighted here: namely, maintain the tie-breaking assumption embodied in the Weak Commitment axiom, but *strengthen Sophistication so as to specify how the DM evaluates non-compact sets of continuation plans*.

In the example, the DM cannot achieve an expected payoff of exactly 10; however, she can approximate it by choosing a positive but small. In light of this observation, it does not seem unreasonable to assume that *the DM will evaluate the tree f to be worth 10 payoff units*. For instance, she will strictly prefer the tree f to any constant prize smaller than 10. More generally, one can *assume that a collection of continuation plans is evaluated according to the “supremum”, not the “maximum” of expected payoffs*.

It is easy to modify the definition of consistent planning to reflect this assumption. Refer to Def. 9;

²⁸The relevant topological notions are made precise in §B.2 of the Web Appendix. In the example, it is sufficient to identify plans with acts, i.e. payoff vectors in \mathbb{R}^4 ; in this case, the preceding analysis makes it clear that the set of feasible payoff vectors is not compact.

²⁹Harris [19] shows that subgame-perfect equilibria exist for a broad class of “regular” infinite games; his analysis can be adapted to the present setting as well.

once a suitable topology on the set of plans has been introduced, the set $CP_f(h)$ can be redefined to be the collection of $\succ_{E(h)}$ -best elements in the *closure* of $CP_f^0(h)$.

A characterization of this modified definition analogous to Theorem 1 can be obtained by suitably strengthening the axioms in Sec. 3.1.1. This requires some care, but the basic idea is simple. Recall that Sophistication requires that, whenever a set B of actions at a history h of a tree f dominates every other action at h , then f is deemed equivalent to a tree where the actions not in B are pruned. The strengthened Sophistication axiom instead requires that f be equivalent to a modified tree f' where the continuation plans available at h are, loosely speaking, the *closure* of the set of continuation plans corresponding to actions in B .³⁰

The precise statement of the axioms, algorithmic procedure, and characterization result, can be found in §B.2 of the Web Appendix. The remainder of this subsection focuses on the interpretation of the assumptions implicit in this approach.

As noted in the Introduction, Consistent Planning can be seen as a way to extend a collection of preferences over *plans* to preferences over more general *trees*. It should be clear from the above, informal discussion that the proposed approach achieves this objective for simple infinite trees.

It is also worth emphasizing that this approach does not entail any departure from the fully decision-theoretic point of view maintained throughout this paper. There are two aspects to this claim. First, to the extent that one can construct actual choice experiments featuring infinitely many alternatives, the strengthened axioms are *testable*; they reflect aspects of the DM's behavior. Second, the modified consistent-planning algorithm is simply a description of a specific individual's preferences, so there is no ambiguity as to the interpretation of *welfare* statements. By way of contrast, welfare analysis presents methodological difficulties in the multi-selves approach: see especially Gul and Pesendorfer [14, §6.4].

One final issue warrants further discussion. Again refer to the above example. The proposed assumption is that the DM deem the tree f to be worth 10 units, the “supremum” of the expected continuation payoffs she can actually secure in f . *This assumption pertains to the DM's evaluation of a (non-compact) set of feasible continuation plans, not to the DM's beliefs about her future choices.* In particular, the DM does *not* expect to be able to commit to s at the second decision point if she chooses $a = 0$.

This can be seen by applying the modified axioms to the tree f . Following the arguments given above, the tree is first pruned so that each initial action $a > 0$ leads to a choice of s , and the initial action $a = 0$ leads to c followed by t . The modified Sophistication axiom now implies that the resulting pruned tree is deemed equivalent to an *expanded* tree that includes a new initial action, say “0*”, corresponding to the plan in f characterized by the choices $a = 0$ and s . Thus, if the DM chooses the “real” action $a = 0$, she still expects this to be followed by c and t . The assumption adopted here is simply that she evaluates the

³⁰Other changes to the axioms are necessary to handle infinite trees; however, these changes are not related to consistent planning per se, and are largely of a notational and technical nature.

tree f “as if” the fictitious action $a = 0^*$ was also available.

A Appendix

A.1 Proof of Theorem 1 (Consistent Planning)

A.1.1 Preliminaries

First, for completeness, I show that the relabeling relation \approx is an equivalence. Reflexivity and transitivity of \approx are immediate. For symmetry, suppose that $f \approx g$ and let φ be as in Def. 3. To show that also $g \approx f$, note first that φ^{-1} satisfies properties (i), (ii) and (iv) in the Definition. For (iii), I prove the following remark, which will also be used in the proof of Lemma 4 below.

Remark A.1 Consider $f, g \in F_E$ such that $f \approx g$, and let φ be as in Def. 3. Then, for all $h' \in g$, $a' \in C_g(h')$, and $D, D' \in \mathcal{F}_g(h', a')$, $a(\varphi^{-1}([h', a', D])) = a(\varphi^{-1}([h', a', D'])).$

Proof: If $C_g(h') = \emptyset$, so $h' \in \mathcal{T}_E$, then by (i) in Def. 3 also $\varphi^{-1}(h') \in \mathcal{T}_E$; similarly, if $x \in C_g(h') \cap X$, then $\mathcal{F}_g(h', x) = \emptyset$, so $[h', x] \in \mathcal{T}_E$ and thus $\varphi^{-1}([h', x]) \in \mathcal{T}_E$ as well: in turn, this implies that $\mathcal{F}_f([h, x]) = \emptyset$. Thus, the claim is trivially true in these cases.

Assume now that $a' \in C_g(h') \cap A$, and consider $D, D' \in \mathcal{F}_g(h', a')$. Suppose by contradiction that $a = a(\varphi^{-1}([h', a', D])) \neq b = a(\varphi^{-1}([h', a', D'])),$ so in particular $D \neq D'$, and let $h = \varphi^{-1}(h')$. Then $a, b \in C_f(h)$; since φ satisfies (iii) by assumption, for any $G \in \mathcal{F}_f(h, a)$, $a(\varphi([h, a, G])) = a(\varphi([h, a, D])) = a'$, and similarly, for any $G \in \mathcal{F}_f(h, b)$, $a(\varphi([h, b, G])) = a(\varphi([h, b, D'])) = a'$. This implies that $\mathcal{F}_g(h', a') \supset \mathcal{F}_f(h, a) \cup \mathcal{F}_f(h, b)$. Since $\mathcal{F}_g(h', a'), \mathcal{F}_f(h, a), \mathcal{F}_f(h, b)$ are all partitions of $E(h') = E(h)$, this implies that they are all equal. In particular, $D' \in \mathcal{F}_f(h, a)$: but then $\varphi([h, a, D']) = [h', a', D'] = \varphi([h, b, D']),$ which contradicts the fact that φ is one-to-one. ■

Next, the following straightforward result on relabelings that will be used several times below. (This Lemma also suggests an alternative, inductive characterization of relabeling based on bijections between sets of choices available at different histories; such a characterization would be possibly more explicit, but also more involved, than Def. 3).

Lemma 4 Consider $E \in \Sigma \setminus \emptyset$ and $f, g \in F_E$. Then $f \approx g$ if and only if there is a bijection $\gamma : C_f(\emptyset) \rightarrow C_g(\emptyset)$ such that, for all $c \in C_f(\emptyset)$,

(i) $c \in X$ implies $\gamma(c) = c$, and

(ii) $c \in A$ implies $\mathcal{F}_f(\emptyset, c) = \mathcal{F}_g(\emptyset, \gamma(c))$ and, for all $D \in \mathcal{F}_f(\emptyset, c)$, $f([c, D]) \approx g([\gamma(c), D]).$

Proof: (If): for each $a \in C_f(\emptyset) \cap A$ and $D \in \mathcal{F}_f(\emptyset, a)$, let $\varphi_{[a,D]}$ be the bijection satisfying (i)–(iv) in Def. 3 for $f([a,D])$ and $g([\alpha(a), D])$. Define $\varphi : f \rightarrow g$ as follows: first, let $\varphi(\emptyset) = \emptyset$ and $\varphi([x]) = [x]$ for every $x \in C_f(\emptyset) \cap X$; then, for any history $h = (a_1, E_1, \dots, a_T, E_T) \in f \cap \mathcal{A}_E$ with $T \geq 1$, let $\varphi(h) = [\gamma(a_1), E_1, \varphi_{[a_1, E_1]}([a_2, E_2, \dots, a_T, E_T])]$;³¹ similarly, for any history $h = (a_1, E_1, \dots, a_{T-1}, E_{T-1}, x) \in f \cap \mathcal{T}_E$ with $T \geq 2$, let $\varphi(h) = [\gamma(a_1), E_1, \varphi_{[a_1, E_1]}([a_2, E_2, \dots, a_{T-1}, E_{T-1}, x])]$. It is now routine to verify that, since the bijections $\varphi_{[a,D]}$ satisfy (i)–(iv) in Def. 3, φ is a bijection that also satisfies the same properties.

(Only if): let $\varphi : f \rightarrow g$ be the bijection in Def. 3 for f and g , and let $\gamma : C_f(\emptyset) \rightarrow C_g(\emptyset)$ be such that, for all $c \in C_f(\emptyset)$, if $c \in A$ and $D \in \mathcal{F}_f(\emptyset, a)$, then $\varphi([c, D]) = [\gamma(c), D]$ (the fact that $E(\varphi([a, D])) = D$ follows from (ii) in Def. 3); and if $c \in X$, then $\varphi([c]) = [\gamma(c)]$, so $\gamma(c) = c$ (recall that $\varphi([c])$ must have length 1, be terminal, and satisfy $\xi(\varphi([c])) = \xi([c]) = c$). Then γ is well-defined on $C_f(\emptyset) \cap A$: if $[a, D], [a, D'] \in f$, then $D, D' \in \mathcal{F}_f(\emptyset, a)$, and (iii) in Def. 3 ensures that $a(\varphi([a, D])) = a(\varphi([a, D']))$; it is also trivially well-defined on $C_f(\emptyset) \cap X$. Also, consider $c' \in C_g(\emptyset)$: if $c' \in X$, then $[c'] \in g \cap \mathcal{T}_E$, so $\varphi^{-1}([c'])$ is terminal in f and has length 1, which implies that $c' \in C_f(\emptyset)$. If instead $c' \in A$, then there is $D \in \Sigma$ such that $[a', D] \in g$, and since φ is onto, there exists $a \in C_f(\emptyset)$ such that $\varphi([a, D]) = [a', D]$; this shows that γ is onto. Similarly, consider $c, c' \in C_f(\emptyset)$. If $c \in A$ and $c' \in X$, clearly $\gamma(c) \neq \gamma(c')$. If $c, c' \in X$ and $c \neq c'$, then $\gamma(c) = c \neq c' = \gamma(c')$. Finally, consider the case $c, c' \in A$, so there are $D, D' \in \Sigma$ such that $[c, D], [c', D'] \in f$. Suppose $\gamma(c) = \gamma(c') = a'$: then $a(\varphi([c, D])) = a(\varphi([c', D'])) = a'$, which implies that $D, D' \in \mathcal{F}_g(\emptyset, a')$. But then Remark A.1 implies that $c = a(\varphi^{-1}([a', D])) = a(\varphi^{-1}([a', D'])) = c'$; thus, γ is one-to-one. It is now routine to verify that the bijection γ satisfies properties (i) and (ii) in the above claim. ■

Corollary 5 *Assume further that $f, g \in F_E$ are such that $C_f(\emptyset) = \{c\}$ and $C_g(\emptyset) = \{c'\}$. Then:*

- (i) *If $c \in X$, then $f \approx g$ iff $f = g$;*
- (ii) *if $c \in A$, then $f \approx g$ iff $\mathcal{F}_f(\emptyset, c) = \mathcal{F}_g(\emptyset, c')$ and, for all $D \in \mathcal{F}_f(\emptyset, c)$, $f([c, D]) \approx g([c', D])$.*

It is also useful to extend the notion of “equivalence up to relabeling” from individual trees to sets of trees. Let $E \in \Sigma$ and consider $G, G' \subset F_E$. Then $G \approx G'$ means that there is a bijection $\psi : G \rightarrow G'$ such that, for every $g \in F$, $g \approx \psi(g)$. The following corollary to Lemma 4 then follows easily:

Corollary 6 *Consider $E \in \Sigma \setminus \{\emptyset\}$ and $f, g \in F_E$. Then $f \approx g$ iff the set $\{f(\emptyset, c) : c \in C_f(\emptyset)\}$ is a relabeling of $\{g(\emptyset, c') : c' \in C_g(\emptyset)\}$.*

Proof: (If) : let ψ the bijection between the two sets of continuation trees. Define $\gamma : C_f(\emptyset) \rightarrow C_g(\emptyset)$ by $\gamma(c) = c'$ iff $\psi(f(\emptyset, c)) = g(\emptyset, c')$. The claim now follows from Corollary 5.

³¹In particular, $\varphi([a, D]) = [\gamma(a), D, \varphi_{[a,D]}(\emptyset) = [\gamma(a), D]]$ for suitable a and D .

(Only If): let γ be the bijection in Lemma 4 and define $\psi(f(\emptyset, c)) = g(\emptyset, \gamma(c))$. It is now easy to verify that ψ has the required properties. ■

Finally, the definition of Consistent Planning (Def. 9) can be augmented with an explicit iterative structure. Fix a tree $f \in F_E$. A **feasible sequence** for the Consistent-Planning algorithm is an ordering h^1, \dots, h^N of the non-terminal histories of f such that, for all $n, m \in \{1, \dots, N\}$, $n < m$ implies not $h^n \leq h^m$; that is, either h^n and h^m are not ordered (neither one follows the other), or h^n follows h^m . Clearly, in every feasible sequence, $h^N = \emptyset$.

Corresponding to h^1, \dots, h^N , one can construct a sequence of trees $f^1, \dots, f^{N+1} \in F_E$ as follows. First, let $f^1 = f$; then, for all $n = 1, \dots, N$, let $f^{n+1} = (g^n)_{h^n} f^n$, where the tree g^n is such that $\{g^n(\emptyset, c) : c \in C_{g^n}(\emptyset)\} \approx \text{CP}_f(h^n)$. Observe that, in this construction, $C_{f^n}(h^n) = C_f(h^n)$ for every $n = 1, \dots, N$.

The assumption that A is countably infinite ensures that this construction is always possible: in particular, in the definition of the tree g^n at step n , one can always find as many distinct action labels as there are plans in $\text{CP}_f(h^n)$.

A.1.2 Sufficiency

Assume first that (i) in Theorem 1 holds, and let f^1, \dots, f^{N+1} be the sequence of trees corresponding to the feasible sequence h^1, \dots, h^N . It will be shown that, for all $n = 1, \dots, N$, $f^n \sim_E f^{n+1}$.

Claim 1: Fix $n \in \{1, \dots, N\}$ and consider the history h^n . Then:

- (a) for every action $a \in C_{f^n}(h^n) \cap A$ and event $D \in \mathcal{F}_{f^n}(h^n, a)$,

$$\{f^n([h^n, a, D], c) : c \in C_{f^n}([h^n, a, D])\} \approx \text{CP}_f([h^n, a, D]);$$

- (b) for every prize $x \in C_{f^n}(h^n) \cap X$, $\text{CP}_f([h^n, x]) = \{f^n(h^n, x)\}$ and $f^n(h^n, x) = f(h^n, x) = \{\emptyset, [x]\}$.

Proof: Note that, for every action a and event D as in (a), $[h^n, a, D] = h^m$ for some index $m < n$; now, for all $\ell \in \{m+1, \dots, n-1\}$, $h^m \not\leq h^\ell$, and clearly it is also note the case that $h^m > h^\ell > h^n$; this implies that $f^n(h^m) = f^{m+1}(h^m) = g^m$ [that is, no further modification of the tree at histories weakly following $h^m = [h^n, a, D]$ has occurred after step m]. Thus, by construction, $\{f^n([h^n, a, D], c) : c \in C_{f^n}([h^n, a, D])\} = \{g^m(\emptyset, c) : c \in C_{g^m}(\emptyset)\} \approx \text{CP}_f([h^n, a, D])$ by the definition of g^m .

If instead x is as in (b), then $x \in C_f(h^n) = C_{f^n}(h^n)$, so by the relevant definitions $\text{CP}_f([h^n, c]) = \{f(h^n, x)\}$ and $f(h^n, x) = \{\emptyset, [x]\}$; furthermore, clearly, for $m = 1, \dots, n-1$, $[h^n, x] \not\leq h^m$, which implies that $f^n([h^n, x]) = f([h^n, x])$, as needed. ■

Claim 2. Consider a tree $g_0^n \in F_{E(h^n)}$ such that $\{g_0^n(\emptyset, c) : c \in C_{g_0^n}(\emptyset)\} \approx \text{CP}_f^0(h^n)$. Then g_0^n is a one-period commitment version of $f^n(h^n)$.

Proof: Let \bar{g} be a tree that allows one-period commitment in $f^n(h^n)$. Note first that, by Def. 7, \bar{g} must be a plan, because the tree f^n is such that every choice at histories following h^n corresponds to a plan. Denote the unique choice in $C_{\bar{g}}(\emptyset)$ by c . If $c \in X$, so $\bar{g} = \{\emptyset, [x]\}$, then $c \in C_{f^n}(h^n) = C_f(h^n)$, and so by definition $\bar{g} \in \text{CP}_f^0(h^n)$: this implies that also $c \in C_{g_0^n}(\emptyset)$, i.e. $g_0^n(\emptyset, c) = \bar{g}$. If instead $c \in A$, consider an arbitrary $D \in \mathcal{F}_{\bar{g}}(\emptyset, c)$. Again by Def. 7, $\bar{g}([c, D]) = f^n(h^n)([c, D], b_{[c, D]}) = f^n([h^n, c, D], b_{[c, D]})$ for some $b_{[c, D]} \in C_{f^n}([h^n, c, D])$; therefore, Claim 1 implies that there exists a unique $p_{[c, D]} \in \text{CP}_f([h^n, c, D])$ such that $\bar{g}([c, D]) \approx p_{[c, D]}$. Now consider the (unique) plan $p \in F_{E(h^n)}$ such that $C_p(\emptyset) = \{c\}$, $\mathcal{F}_p(\emptyset, c) = \mathcal{F}_{\bar{g}}(\emptyset, c) = \mathcal{F}_f(h^n, c)$, and for all $D \in \mathcal{F}_p(\emptyset, c)$, $p([c, D]) = p_{[c, D]}$. Clearly, $p \in \text{CP}_f^0(h^n)$; since $\{g_0^n(\emptyset, c) : c \in C_{g_0^n}(\emptyset)\} \approx \text{CP}_f^0(h^n)$, there exists a unique $\bar{c} \in C_{g_0^n}(\emptyset)$ such that $g_0^n(\emptyset, \bar{c}) \approx p$. Furthermore, by Lemma 4, $p \approx \bar{g}$; hence, as required, there is a unique action $\bar{c} \in C_{g_0^n}(\emptyset)$ such that $g_0^n(\emptyset, \bar{c}) \approx \bar{g}$.

Conversely, fix $\bar{c} \in C_{g_0^n}(\emptyset)$. By assumption there is a (unique) plan $p \in \text{CP}_f^0(h^n)$ such that $g_0^n(\emptyset, \bar{c}) \approx p$. If $\bar{c} \in X$, then $g_0^n(\emptyset, \bar{c}) = p = \{\emptyset, [\bar{c}]\}$ and, by Def. 9, $\bar{c} \in C_f(h^n) = C_{f^n}(h^n)$: this implies that, trivially, $g_0^n(\emptyset, \bar{c})$ allows one-period commitment in $f^n(h^n)$. If instead $\bar{c} \in A$, then by Def. 9, there is $c \in C_f(h^n) = C_{f^n}(h^n)$ such that $C_p(\emptyset) = \{c\}$, $\mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h^n, c) = \mathcal{F}_{f^n}(h^n, c)$, and furthermore, for every $D \in \mathcal{F}_p(\emptyset, c)$, $p([c, D]) \in \text{CP}_f([h^n, c, D])$. Claim 1 implies that there exists a (unique) $b_{[c, D]} \in C_{f^n}([h^n, c, D])$ such that $p([c, D]) \approx f^n([h^n, c, D], b_{[c, D]})$. Thus, consider the plan $\bar{g} \subset f^n(h^n)$ such that $C_{\bar{g}}(\emptyset) = \{c\}$, $\mathcal{F}_{\bar{g}}(\emptyset, c) = \mathcal{F}_{f^n}(h^n, c)$, and for all $D \in \mathcal{F}_{\bar{g}}(\emptyset, c)$, $\bar{g}([c, D]) = f^n([h^n, c, D], b_{[c, D]})$. By Lemma 4, $p \approx \bar{g}$, hence $g_0^n(\emptyset, \bar{c}) \approx \bar{g}$, and furthermore \bar{g} allows one-period commitment in $f^n(h^n)$, as needed. ■

Claim 3. For all $n = 1, \dots, N$, $f^n \sim_E f^{n+1}$.

Proof: Consider g_0^n as in Claim 2. The latter asserts that g_0^n is a one-period commitment version of $f^n(h^n)$. Furthermore, if h^m immediately follows h^n , then $m < n$ and $f^n(h^m) = g^m$; consider $c, c' \in C_{f^n}(h^m)$; by construction, there exist $p, p' \in \text{CP}_f(h^m)$ such that $f^n(h^m, c) = g^m(\emptyset, c) \approx p$ and $f^n(h^m, c') = g^m(\emptyset, c') \approx p'$, which implies that $f^n(h^m, c) \sim_{E(h^m)} f^n(h^m, c')$. Therefore Weak Commitment implies that $(g_0^n)_{h^n} f^n \sim_E f^n$. Now consider the set of $\succ_{E(h^n)}$ -best actions at the initial history of g_0^n , i.e. $C^n = \{c \in C_{g_0^n}(\emptyset) : \forall c' \in C_{g_0^n}(\emptyset), g_0^n(\emptyset, c) \succ_{E(h^n)} g_0^n(\emptyset, c')\}$. Then Sophistication implies that $(g_0^n(\emptyset, C^n))_{h^n} f^n \sim_E (g_0^n)_{h^n} f^n$.

I claim that $g_0^n(\emptyset, C^n)$ is a relabeling of g^n . To prove this, by Corollary 6 it is enough to show that $\{g_0^n(\emptyset, c) : c \in C^n\} \approx \text{CP}_f(h^n)$, as $\text{CP}_f(h^n) \approx \{g^n(\emptyset, c) : c \in C_{g^n}(\emptyset)\}$ by construction. Thus, consider $c \in C^n$, and let $p \in \text{CP}_f^0(h^n)$ be the unique plan in the latter set such that $g_0^n(\emptyset, c) \approx p$. Now consider an arbitrary $p' \in \text{CP}_f^0(h^n)$ and let $c' \in C_{g_0^n}(\emptyset)$ be such that $g_0^n(\emptyset, c') \approx p'$. Then $g_0^n(\emptyset, c) \succ_{E(h^n)} g_0^n(\emptyset, c')$; by Relabeling Invariance, this implies that $p \succ_{E(h^n)} p'$. Thus, $p \in \text{CP}_f(h^n)$. Conversely, pick $p \in \text{CP}_f(h^n) \subset \text{CP}_f^0(h^n)$, and let $c \in C_{g_0^n}(\emptyset)$ be such that $p \approx g_0^n(\emptyset, c)$. Consider an arbitrary $c' \in C_{g_0^n}(\emptyset)$, and let $p' \in \text{CP}_f^0(h^n)$ be such that $g_0^n(\emptyset, c') \approx p'$. Then, by the definition of $\text{CP}_f(h^n)$, $p \succ_{E(h^n)} p'$, so by Relabeling Invariance also $g_0^n(\emptyset, c) \succ_{E(h^n)} g_0^n(\emptyset, c')$. Thus, $c \in C^n$.

Finally, $g_0^n(\emptyset, C^n) \approx g^n$ readily implies³² that $(g_0^n(\emptyset, C^n))_{h^n} f^n \approx (g^n)_{h^n} f^n = f^{n+1}$. Thus, by Relabeling Invariance, $f^n \sim_E (g_0^n(\emptyset, C^n))_{h^n} f^n \sim_E (g^n)_{h^n} f^n = f^{n+1}$, as required. ■

The proof of sufficiency can now be completed. Observe that $f = f^1 \sim_E f^{N+1}$ by transitivity; furthermore, $\{f^{N+1}(\emptyset, c) : c \in C_{f^{N+1}}(\emptyset)\} \approx \text{CP}_f(\emptyset)$, so for all $c, c' \in C_{f^{N+1}}(\emptyset)$, $f^{N+1}(\emptyset, c) \sim_E f^{N+1}(\emptyset, c')$: thus, Simplification implies $f^{N+1}(\emptyset, c) \sim_E f$ for all $c \in C_{f^{N+1}}(\emptyset)$. Now consider $g \in F_E$ such that $g \approx p$ for some $p \in \text{CP}_f(\emptyset)$: by construction, there is some $c_p \in C_{f^{N+1}}(\emptyset)$ such that $p \approx f^{N+1}(\emptyset, c_p)$; by transitivity of the relabeling relation, $g \approx f^{N+1}(\emptyset, c_p)$. Now, by Relabeling Invariance, $g \sim_E f^{N+1}(\emptyset, c_p) \sim_E f$, as required.

A.1.3 Necessity

First of all, observe that Def. 9 immediately implies that, for every tree f and every history h of f , $\text{CP}_f(h) = \text{CP}_{f(h)}(\emptyset)$. The following property of Consistent Planning is also easy to establish.

Lemma 7 *Let $f, g \in F_E$ be such that $f \approx g$. Then $\text{CP}_f^0(\emptyset) \approx \text{CP}_g^0(\emptyset)$ and $\text{CP}_f(\emptyset) \approx \text{CP}_g(\emptyset)$.*

Proof: The argument is by induction on the maximum length of histories in f (hence, also in g). If f and g are constant trees (so the maximum length is one), the claim is immediate from Corollary 5. Thus, suppose that the claim is true for trees of maximum history length n , and assume that f and g have maximum history length equal to $n + 1$. Let $\varphi : f \rightarrow g$ be the bijection in Def. 3. Lemma 4 implies that, for every $a \in C_f(\emptyset) \cap A$ and $D \in \mathcal{F}_f(\emptyset, a)$, $f([a, D]) \approx g(\varphi([a, D]))$; by the induction hypothesis, the claim is true for these two (sub)trees, so $\text{CP}_f([a, D]) = \text{CP}_{f([a, D])}(\emptyset) \approx \text{CP}_{g([a, D])}(\emptyset) = \text{CP}_g([a, D])$. Moreover, for every $x \in C_f(\emptyset) \cap X$, $\varphi([x]) = [x]$.

It is clear that, if $\{\emptyset, [x]\} \in \text{CP}_f^0(\emptyset)$, then $x \in C_f(\emptyset)$, and so also $\{\emptyset, [x]\} \in \text{CP}_g^0(\emptyset)$. Next, suppose $p \in \text{CP}_f^0(\emptyset)$ is such that $C_p(\emptyset) = \{a\}$ for some $a \in A$. Let $\gamma : C_f(\emptyset) \rightarrow C_g(\emptyset)$ be the bijection in Lemma 4. For every $D \in \mathcal{F}_p(\emptyset, a) = \mathcal{F}_f(\emptyset, a)$, $p([a, D]) \in \text{CP}_f([a, D])$, and by the preceding argument there is $q_{[a, D]} \in \text{CP}_g(\varphi([a, D]))$ such that $p([a, D]) \approx q_{[a, D]}$. Hence, there is $q \in \text{CP}_g^0(\emptyset)$ such that $C_q(\emptyset) = \{\gamma(a)\}$ and, for all $D \in \mathcal{F}_q(\emptyset, \gamma(a)) = \mathcal{F}_g(\emptyset, \gamma(a))$, $q([\gamma(a), D]) = q_{[a, D]}$. By Lemma 4, $p \approx q$.

Furthermore, suppose $p, p' \in \text{CP}_f^0(\emptyset)$, with $p \neq p'$. If $C_p(\emptyset) \neq C_{p'}(\emptyset)$, or if $C_p(\emptyset) = C_{p'}(\emptyset) = \{a\}$ and $\mathcal{F}_p(\emptyset, a) \neq \mathcal{F}_{p'}(\emptyset, a)$, the preceding construction associates distinct elements of $\text{CP}_g^0(\emptyset)$ to p and p' . Otherwise, there must be $a \in A$ and $D \in \mathcal{F}_p(\emptyset, a) = \mathcal{F}_{p'}(\emptyset, a)$ such that $p([a, D]) \neq p'([a, D])$. The induction hypothesis implies that $p([a, D])$ and $p'([a, D])$ map to different elements of $\text{CP}_g([\gamma(a), D])$, so it follows that p and p' are mapped to different elements of $\text{CP}_g^0(\emptyset)$.

³²Let ψ be a bijection from histories of $g_0^n(\emptyset, C^n)$ to histories of g^n with the properties in Def. 3. Then a suitable bijection φ from histories in $[g_0^n(\emptyset, C^n)]_{h^n} f^n$ to histories in $[g^n]_{h^n} f^n$ can be defined by letting $\varphi(h) = h$ if h is a history of f^n that does not weakly follow h^n , and $\varphi([h^n, \bar{h}]) = [h^n, \psi(\bar{h})]$ for every history \bar{h} of $g_0^n(\emptyset, C^n)$.

Thus, the proposed construction yields a one-to-one map of $\text{CP}_f^0(\emptyset)$ to $\text{CP}_g^0(\emptyset)$. A symmetric argument shows that this map is also onto.

Thus, $\text{CP}_f^0(\emptyset) \approx \text{CP}_g^0(\emptyset)$. Now consider $p \in \text{CP}_f(\emptyset)$, so for all $p' \in \text{CP}_f^0(\emptyset)$, $p \succ_E p'$. Then there exists $q \in \text{CP}_g^0(\emptyset)$ such that $p \approx q$; furthermore, for every $q' \in \text{CP}_g^0(\emptyset)$, there is $p' \in \text{CP}_f^0(\emptyset)$ with $q' \approx p'$, which, by Relabeling Invariance implies that $q \succ_E q'$. Hence, $q \in \text{CP}_g(\emptyset)$. By a symmetric argument, if $q \in \text{CP}_g(\emptyset)$, there exists $p \in \text{CP}_f(\emptyset)$ such that $p \approx q$. This completes the proof of the inductive step. ■

Turn now to the actual proof of necessity. Assume that (ii) holds in Theorem 1. Since each \succ_E is complete and transitive on F_E^p , and $\text{CP}_f(\emptyset) \neq \emptyset$ for all $f \in F_E$, it is clear that \succ_E must also be complete and transitive on all of F_E . Next, the three axioms in (i) will be considered in turn. As will be clear, only Weak Commitment requires somewhat special care.

Axiom 3.1, Sophistication. Let $f \in F_E$ and fix a history $h \in f \cap \mathcal{H}_E$ and a set $B \subset C_f(h)$ as in the Axiom. Clearly, Def. 9 implies that $\text{CP}_f(h)^0 = \bigcup_{c \in C_f(h)} \text{CP}_{f(h,c)}^0(\emptyset)$ and similarly $\text{CP}_{f(h,B)_h f}^0(h) = \bigcup_{b \in B} \text{CP}_{f(h,b)}^0(\emptyset)$. Fix $p \in \text{CP}_f(h)$ and $q \in \text{CP}_{f(h,B)_h f}(h)$, and assume for definiteness that $p \in \text{CP}_{f(h,c)}^0(\emptyset)$ and $q \in \text{CP}_{f(h,b)}^0(\emptyset)$; then in particular $p \sim_{E(h)} f(h,c) \succ_{E(h)} f(h,b) \sim q$. Since $f(h,b) \succ_{E(h)} f(h,c')$ for all $c' \in B \setminus C_f(h)$, it must be the case that $c \in B$, and hence $p \in \text{CP}_{f(h,B)_h f}(h)$. Furthermore, it must be the case that $f(h,c) \sim_{E(h)} f(h,b)$, and so $q \in \text{CP}_f(h)$. Thus, $\text{CP}_f(h) = \text{CP}_{f(h,B)_h f}(h)$.

It is clear that, if $\bar{h} \in f$ neither strictly follows nor strictly precedes h , then $\text{CP}_f(\bar{h}) = \text{CP}_{f(h,B)_h f}(\bar{h})$. Since $\text{CP}_f(h) = \text{CP}_{f(h,B)_h f}(h)$, a simple induction argument shows that actually $\text{CP}_f(\bar{h}) = \text{CP}_{f(h,B)_h f}(\bar{h})$ for all histories \bar{h} that weakly precede h , including $\bar{h} = \emptyset$. By (ii) in Theorem 1, $f \sim_E f(h,B)_h f$.

Axiom 3.2, Weak Commitment. Let $f \in F_E$ and fix a history h with the properties indicated in the Axiom. In particular, note that, if $a \in C_f(h) \cap A$, $D \in \mathcal{F}_f(h,a)$ and $b, b' \in C_f([h,a,D])$, then $f([h,a,D], b) \sim f([h,a,D], b')$; this implies that $\text{CP}_f([h,a,D]) = \bigcup_{b \in C_f([h,a,D])} \text{CP}_{f([h,a,D],b)}(\emptyset)$. Now let g be a one-period commitment version of $f(h)$; I claim that $\text{CP}_f^0(h) \approx \text{CP}_{g_h f}^0(h)$. To prove this claim, a suitable map $\psi : \text{CP}_f^0(h) \rightarrow \text{CP}_{g_h f}^0(h)$ will be constructed

Fix $p \in \text{CP}_f^0(h)$. If $p = \{\emptyset, [x]\}$, let $\psi(p) = p$. Otherwise, there is $a \in A$ such that $C_p(\emptyset) = \{a\} \subset C_f(h)$, and furthermore, for all $D \in \mathcal{F}_p(\emptyset, a) = \mathcal{F}_f(h,a)$, $p([a,D]) \in \text{CP}_f([h,a,D])$. By the argument in the previous paragraph, there is $b_{[a,D]} \in C_f([h,a,D])$ such that $p([a,D]) \in \text{CP}_{f([h,a,D],b_{[a,D]})}(\emptyset)$. Now consider the tree \bar{g} that allows one-period commitment in $f(h)$ and such that $C_{\bar{g}}(\emptyset) = \{a\}$ and, for all $D \in \mathcal{F}_{\bar{g}}(\emptyset, a) = \mathcal{F}_f(h,a)$, $C_{\bar{g}}([a,D]) = \{b_{[a,D]}\}$. Then for all D , $\bar{g}([a,D]) = f([h,a,D], b_{[a,D]})$, and thus $p([a,D]) \in \text{CP}_{\bar{g}}([a,D])$. Therefore, $p \in \text{CP}_{\bar{g}}^0(\emptyset)$. Since g is a one-period commitment version of $f(h)$, there is $c \in C_g(\emptyset)$ such that $g(\emptyset, c) \approx \bar{g}$, and Lemma 7 implies that there exists $q \in \text{CP}_{g(\emptyset,c)}^0(\emptyset)$ such that $p \approx q$. Finally, it is clear from Def. 9 that $\text{CP}_g^0(\emptyset) = \bigcup_{c \in C_g(\emptyset)} \text{CP}_{g(\emptyset,c)}^0(\emptyset)$, so $q \in \text{CP}_g^0(\emptyset) = \text{CP}_{g_h f}^0(h)$. Thus, let $\psi(p) = q$.

It is clear from the above construction that ψ is one-to-one. To see that it is onto, let $q \in \text{CP}_{g_h f}^0(h) = \text{CP}_{\bar{g}}^0(\emptyset)$. Then there is $c \in C_g(\emptyset)$ such that $q \in \text{CP}_{g(\emptyset, c)}^0(\emptyset)$. By Def. 8, there is a tree \bar{g} that allows one-period commitment in $f(h)$ and such that $\bar{g} \approx g(\emptyset, c)$; by Lemma 7, this implies that there exists a unique $p \in \text{CP}_{\bar{g}}^0(\emptyset)$ such that $p \approx q$. I claim that $p \in \text{CP}_f^0(h)$ and $\psi(p) = q$. Note first that, if $p = \{\emptyset, [x]\}$, then $p = q$ and indeed $\psi(p) = p = q$; furthermore, also $\bar{g} = p$, and so $p \in \text{CP}_f^0(h)$. Otherwise, let $C_p(\emptyset) = \{a\} \subset C_f(h)$, and note that $\mathcal{F}_p(\emptyset, a) = \mathcal{F}_{\bar{g}}(\emptyset, c) = \mathcal{F}_f(h, a)$. By Def. 9, for all $D \in \mathcal{F}_p(\emptyset)$, $p([a, D]) \in \text{CP}_{\bar{g}}^0([a, D])$; furthermore, $\bar{g}([a, D]) = f([h, a, D], b_{[a, D]})$ for some $b_{[a, D]} \in C_f([h, a, D])$. Thus, in the latter case, $p([a, D]) \in \text{CP}_{f([h, a, D], b_{[a, D]})}^0(\emptyset)$; since all actions at $[h, a, D]$ are equivalent, as argued above this implies that $p([a, D]) \in \text{CP}_f^0([h, a, D])$. Thus, $p \in \text{CP}_f^0(h)$. Furthermore, it is clear that \bar{g} is the tree that allows one-period commitment constructed in the definition of $\psi(p)$, and it follows that $\psi(p) = q$.

Hence, $\text{CP}_f^0(h) \approx \text{CP}_{g_h f}^0(h)$. It is easy to verify³³ that then also $\text{CP}_f(h) \approx \text{CP}_{g_h f}(h)$. Clearly, if a history \bar{h} of f neither precedes nor follows h , then \bar{h} is also a history of $g_h f$; furthermore, $[g_h f](\bar{h}) = f(\bar{h})$, and therefore $\text{CP}_f(h') = \text{CP}_{g_h f}(h')$. If instead \bar{h} weakly precedes h , then \bar{h} is also a history of $g_h f$, but the corresponding continuation trees are different; to complete the proof, in light of (ii) of Theorem 1 it is sufficient to show that, if the history \bar{h} of f (weakly) precedes h , then $\text{CP}_f(\bar{h}) \approx \text{CP}_{g_h f}(\bar{h})$.

The claim is true if $\bar{h} = h$, so argue by induction and assume it is true for all histories $h' > \bar{h}$ that (weakly) precede h and strictly follow \bar{h} . Again, it is sufficient to show that $\text{CP}_f^0(\bar{h}) \approx \text{CP}_{g_h f}^0(\bar{h})$. Note first that $C_f(\bar{h}) = C_{g_h f}(\bar{h})$ and, for all $a \in C_f(\bar{h}) \cap A$, $\mathcal{F}_f(\bar{h}, a) = \mathcal{F}_{g_h f}(\bar{h}, a)$. Now pick $p \in \text{CP}_f^0(\bar{h})$; if $p = \{\emptyset, [x]\}$ for $x \in X$, then it is easy to see that also $p \in \text{CP}_{g_h f}^0(\bar{h})$. Otherwise, let $\{a\} = C_p(\emptyset)$ and, for all $D \in \mathcal{F}_p(\emptyset, a) = \mathcal{F}_f(\bar{h}, a)$, note that $p([a, D]) \in \text{CP}_f^0([a, D])$. Hence, by the inductive hypothesis, there is $q_{[a, D]} \in \text{CP}_{g_h f}^0([a, D])$ such that $p([a, D]) \approx q_{[a, D]}$. Thus, let $q \in F_{E(\bar{h})}$ be such that $C_q(\emptyset) = \{a\}$, $\mathcal{F}_q(\emptyset, a) = \mathcal{F}_{g_h f}(\bar{h}, a)$, and for all $D \in \mathcal{F}_q(\emptyset, a)$, $q([a, D]) = q_{[a, D]}$. Then $q \in \text{CP}_{g_h f}^0(\bar{h})$, and by Lemma 4, $p \approx q$. By-now usual arguments imply that this construction yields a one-to-one, onto map of $\text{CP}_f^0(\bar{h})$ onto $\text{CP}_{g_h f}^0(\bar{h})$, and again this implies that also $\text{CP}_f(\bar{h}) \approx \text{CP}_{g_h f}(\bar{h})$. This completes the proof of the claim, and establishes that Weak Commitment holds.

Axiom 3.3, Simplification. Let f and B be as in the Axiom. Clearly, $\text{CP}_f^0(\emptyset) = \bigcup_{c \in C_f(\emptyset)} \text{CP}_{f(\emptyset, c)}^0(\emptyset) \supset \bigcup_{c \in B} \text{CP}_{f(\emptyset, c)}^0(\emptyset) = \text{CP}_{f(\emptyset, B)}^0(\emptyset)$. Pick $p \in \text{CP}_f(\emptyset)$ and $q \in \text{CP}_{f(\emptyset, B)}(\emptyset)$, and for definiteness assume that $p \in \text{CP}_{f(\emptyset, c')}^0(\emptyset)$ and $q \in \text{CP}_{f(\emptyset, c'')}^0(\emptyset)$. Then $f \sim_E p \sim_E f(\emptyset, c') \sim_E f(\emptyset, c'') \sim_E q \sim_E f(\emptyset, B)$, as required.

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

To simplify the exposition, let Σ_1 be the collection of non-null events in Σ .

Remark A.2 Fix $E \in \Sigma_1$ and let \geq be a complete and transitive binary relation on F_E such that

- (i) for all $x, y \in X$, $x \geq y$ iff $x \succ y$;

³³See e.g. the last step in the proof of Lemma 7

- (ii) if $x \in X$ satisfies $x \succ \xi(h)$ [resp. $x \prec \xi(h)$] for all $h \in f \cap \mathcal{T}_E$, 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 $f \in F_E$, 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 $f > x > g$.

Proof: Since f is finite, there exist x', x'' such that $x' \succ x(z) \succ x''$ for all terminal histories z of f . By (i), (ii) and transitivity, $x' \in U$ and $x'' \in L$. By completeness, $U \cup L = X$. Since X is separable, the non-empty, closed [by (iii)] sets U and L satisfy $U \cap L \neq \emptyset$; any $x \in U \cap L$ satisfies $x \sim_E f$, which proves (a). For (b), take $x, y \in X$ such that $x = f$ and $y = g$. The claim follows from separability of X and (iii). ■

Turn now to the proof of Theorem 2. Assume that (ii) holds, and consider $E \in \Sigma_1$. Suppose that $f \succ_E f'$ and let $x \in X$ be such that $x \sim_E f'$: such a prize exists by Remark A.2. Fix $g \in T_E$: I claim that $f \succ_{E,g}^* f'$. To see this, suppose first $y \succ x$, so $y \succ_E x$ by Axiom 3.4; then $y \succ_E f'$ by transitivity, and Axiom 3.7 implies that $\{f', y\}_E g \sim y_E g$. Next, suppose that $y \prec x$: again invoking Axiom 3.4 and transitivity we get $y \prec_E f$, and Axiom 3.7 implies $\{f, y\}_E g \sim f_E g$. This proves the claim. In the opposite direction, consider $f, f' \in F_E$ and $g \in T_E$, and suppose that $f \succ_{E,g}^* f'$; let $x \in X$ be as in Def. 13. Suppose by contradiction that $x \succ_E f$, so there exist $y', y'' \in X$ such that $x \succ_E y' \succ_E y'' \succ_E f$ (by Remark A.2). Now Def. 13 implies $\{f, y'\}_E g \sim f_E g \sim \{f, y''\}_E g$, but Axiom 3.7 and the assumption that E is not null imply $\{f, y'\}_E g \sim y'_E g \succ y''_E g \sim \{f, y''\}_E g$: contradiction. Hence, $f \succ_E x$; similarly, $x \succ_E f'$, and it follows that $f \succ_E f'$. Thus, $f \succ_E f'$ iff $f \succ_{E,g}^* f'$ for all $g \in T_E$.

It remains to be shown that \succ satisfies Axioms 3.10 and 3.9 (Axiom 3.8 is immediately implied by Axioms 3.4 and 3.6). Consider first Conjectural Separability (Axiom 3.10): let $E \in \Sigma_1$, $f \in F_E$, $g, g' \in T_E$ and $x, y \in X$. For (i), suppose that $\{f, y\}_E g \not\sim f_E g$ and $x \succ y$. If $f \prec_E y$, then $f \prec_E x$ as well, and Axiom 3.7 implies that $\{f, x\}_E g' \sim x_E g'$. By Axiom 3.7, it cannot be the case that $f \succ_E y$, so the only remaining case is $f \sim_E y$. Then also $x \succ_E f$, and again, Axiom 3.7 yields the required indifference. The argument for (ii) is similar. Finally, consider Axiom 3.8. If $\xi(h) \succ x$ for all $h \in f \cap \mathcal{T}_E$, then, since f is finite, there is $y \in X$ such that $\xi(h) \succ y$ for all $h \in f \cap \mathcal{T}_E$ and $y \succ x$. Now Axiom 3.5 implies $f \succ_E y$, and hence $f \succ_E x$; now Axiom 3.7 implies $\{f, x\}_E g \sim_E f_E g$, as required. The argument for (ii) is similar.

Now assume that (i) holds. To streamline the exposition, given $E \in \Sigma_1$, $g \in T_E$ and $f, f' \in F_E$ call any $x \in X$ with the properties in Def. 13 a *cutoff* for $f \succ_{E,g}^* f'$.

Claim 1: For every $E \in \Sigma_1$ and $g \in T_E$, $\succ_{E,g}^*$ is transitive.

Consider $f, f', f'' \in F_E$ such that $f \succ_{E,g}^* f'$ and $f' \succ_{E,g}^* f''$, and let $x, x' \in X$ be the respective cutoffs. 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.2): then $f \succ_{E,g}^* f'$ and $E \in \Sigma_1$ imply $\{f', y'\}_{Eg} \sim y'_E g \succ y''_E g \sim \{f', y''\}_{Eg}$, but $f' \succ_{E,g}^* f''$ implies $\{f', y'\}_{Eg} \sim f'_E g \sim \{f', y''\}_{Eg}$, a contradiction.

Now consider $y \in X$. If $y \succ x'$, then $f' \succ_{E,g}^* f''$ implies $\{y, f''\}_{Eg} \sim y_E g$; if instead $y \prec x'$, then $y \prec x$, and $f \succ_{E,g}^* f'$ implies $\{f, y\}_{Eg} \sim f_E g$. But this means precisely that x' is a cutoff for $f \succ_{E,g}^* f''$.

Claim 2: For all $E \in \Sigma_1$, $g \in T_E$ and $x, y \in X$: $x \succ y$ iff $x \succ_{E,g}^* y$. In particular, $x \succ y$ implies $\{x, y\}_{Eg} \sim x_E g$.

Suppose $x \succ y$. For all $x' \succ y$, Axiom 3.9 implies that $\{x', y\}_{Eg} \sim x'_E g$; similarly, for all $x' \prec y$, also $x' \prec x$, and the same axiom implies $\{x, x'\}_{Eg} \sim x_E g$. Hence, y is a cutoff for $x \succ_{E,g}^* y$.

Conversely, suppose $x \succ_{E,g}^* y$ and let y' be a cutoff. If $y' \prec z \prec y$, then Axiom 3.9 implies $\{z, y\}_{Eg} \sim y_E g$, but Def. 13 requires $\{z, y\}_{Eg} \sim z_E g$: since E is non-null, this is a contradiction. Hence, $y' \succ y$, and similarly $x \succ y'$. By transitivity, $x \succ y$.

Claim 3: For every $E \in \Sigma_1$, $g \in T_E$, $f \in F_E$ and $x \in X$, either $f \succ_{E,g}^* x$ or $x \succ_{E,g}^* f$ (or both). In particular, if $x, x' \in X$ satisfy $x \succ \xi(h) \succ x'$ for all $h \in f \cap \mathcal{T}_E$, then $x \succ_{E,g}^* f$ and $f \succ_{E,g}^* x'$.

Suppose that it is not the case that $f \succ_{E,g}^* x$. Then in particular x is not a cutoff; by Claim 2, for all $y \succ x$, $\{y, x\}_{Eg} \sim y_E g$, so there must be $y \prec x$ such that $\{f, y\}_{Eg} \not\sim f_E g$. Then Axiom 3.10 implies that, for all $y' \succ y$, $\{f, y'\}_{Eg} \sim y'_E g$. On the other hand, for all $y' \prec y$, also $y' \prec x$, so Claim 2 implies $\{x, y'\}_{Eg} \sim x_E g$. Hence, y' is a cutoff for $x \succ_{E,g}^* f$.

If x, x' are as above, then Axiom 3.9 implies that, for every $y \prec x'$, $\{f, y\}_{Eg} \sim f_E g$, and Claim 2 implies that, for every $y \succ x'$, $\{y, x'\}_{Eg} \sim y_E g$. Thus, $f \succ_{E,g}^* x'$, and the other relation follows similarly.

Claim 4: For every $E \in \Sigma_1$, $g \in T_E$ and $f \in F_E$, there exists $x \in X$ such that $x \sim_{E,g}^* f$ (i.e. $x \succ_{E,g}^* f$ and $f \succ_{E,g}^* x$ both hold). Hence, $\succ_{E,g}^*$ is complete on F_E .

Let $L = \bigcap_{x: x \succ_{E,g}^* f} \{y : x \succ y\}$. Notice that L is an intersection of closed sets by Axiom 3.8, and hence is closed. Also, the last part of Claim 3 shows that there always exists $x \in X$ such that $x \succ_{E,g}^* f$. Since $\succ_{E,g}^*$ is transitive, if $f \succ_{E,g}^* y$, then $x \succ_{E,g}^* y$ (and hence $x \succ y$) for every $x \in X$ such that $x \succ_{E,g}^* f$: thus, $f \succ_{E,g}^* y$ implies $y \in L$. On the other hand, suppose $f \not\succeq_{E,g}^* y$: then in particular y cannot be a cutoff and, as in the proof of Claim 3, there must exist $x \prec y$ such that $\{f, x\}_{Eg} \not\sim f_E g$. Then Axiom 3.10 implies that, for all $x' \succ x$, $\{f, x'\}_{Eg} \sim x'_E g$; also, by Claim 2, for all $x' \prec x$, $\{x, x'\}_{Eg} \sim x_E g$. Thus, x is a cutoff for $x \succ_{E,g}^* f$, and since $y \notin \{y' : x \succ y'\}$, $y \notin L$. Thus, $L = \{y : f \succ_{E,g}^* y\}$; by the last part of the preceding claim, this set is non-empty. Similarly, $U = \{y : y \succ_{E,g}^* f\}$, and this set is non-empty and closed.

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

Claim 5 For every $E \in \Sigma_1$ and $g, g' \in T_E$, $\succ_{E,g}^* = \succ_{E,g'}^*$.

Observe first that, if $f \succ_{E,g}^* x$, then one can take x as cutoff. Suppose x' is another cutoff: then $x' \succ x$, because otherwise any $y \in X$ such that $x \succ y \succ x'$ would satisfy $\{x, y\}_{Eg} \sim_E x_E g \succ y_E g$, which

contradicts the assumption that x' is a cutoff for $f \succ_{E,g}^* x$. Now take $y \prec x$: then also $y \prec x'$, and by definition $\{f, y\}_E g \sim f_E g$. Since $y \succ x$ immediately implies $\{x, y\}_E g \sim y_E g$, it follows that indeed x is a cutoff for $f \succ_{E,g}^* x$. Similarly, if $x \succ_{E,g}^* f$, then any cutoff x' must satisfy $x \succ x'$, and in particular one can take x as cutoff.

Also, it is enough to prove that, for $f \in F_E$ and $x \in X$, $f \succ_{E,g}^* x$ iff $f \succ_{E,g'}^* x$ and $x \succ_{E,g}^* f$ iff $x \succ_{E,g'}^* f$.

Suppose $f \succ_{E,g}^* x$; take x as the cutoff. To verify that $f \succ_{E,g'}^* x$, we also take x as cutoff, and by the by-now usual argument, only check that $y \prec x$ implies $\{f, y\}_E g' \sim f_E g'$. If instead $\{f, y\}_E g' \not\sim f_E g'$ for some $y \prec x$, then Axiom 3.10 implies that, for all $x' \succ y$, $\{f, x'\}_E g \sim_E x'_E g$ [observe that g is used instead of g']. In particular, this holds for y', y'' such that $x \succ y' \succ y'' \succ y$, and since $E \in \Sigma_1$, $\{f, y'\}_E g \sim y'_E g \succ y''_E g \sim \{f, y''\}_E g$; but since x is a cutoff for $f \succ_{E,g}^* x$, it must be the case that $\{f, y'\}_E g \sim f_E g \sim \{f, y''\}_E g$: contradiction. Reversing the roles of g and g' yields the first part of the claim.

Similarly, suppose $x \succ_{E,g}^* f$; again, take x as cutoff. Then $x \succ_{E,g'}^* f$ can fail only if $\{y, f\}_E g' \not\sim y_E g'$ for some $y \succ x$. But, in this case, Axiom 3.10 implies that $\{y', f\}_E g \sim f_E g \sim \{y'', f\}_E g$ for $y \succ y' \succ y'' \succ x$, and this again yields a contradiction.

The proof of Theorem 2 can now be completed. Letting $f \succ_E^* f'$ iff $f \succ_{E,g}^* f'$ for some $g \in T_E$ yields a well-defined, complete and transitive binary relation on F_E . By Claim 2, this relation satisfies Axiom 3.4; by Claim 3, it satisfies Axiom 3.5; by the arguments in Claim 4, it satisfies Axiom 3.6.

Finally, we verify that it also satisfies Axiom 3.7. Let $E \in \Sigma_1$, $f \in F_E$, $x \in X$ and $g \in T_E$. Suppose $f \succ_E^* x$; if $\{f, x\}_E g \not\sim f_E g$, then Axiom 3.10 implies that, for all $y \succ x$, $\{f, y\}_E g \sim y_E g$; furthermore, for all $y \prec x$, $\{x, y\}_E g \sim x_E g$. Thus, by definition $x \succ_{E,g}^* f$, which is a contradiction.

Similarly, suppose $x \succ_E^* f$; if $\{f, x\}_E g \not\sim x_E g$, then for all $y \prec x$, $\{f, y\}_E g \sim f_E g$: again, for all $y \succ x$, $\{x, y\}_E g \sim x_E g$, which by definition means that x is a cutoff for $f \succ_{E,g}^* x$. Again, a contradiction is found.

A.3 Proof of Theorem 3 (Consistent Planning for MEU Preferences)

Note first that, if E is non-null according to Def. 11, then $\min_{q \in C} q(E) > 0$. Thus, all relevant conditional preferences are well-defined.

Now suppose (1) holds. Notice that Axiom 4.2 implies that, for all non-empty $E \in \Sigma$ plans $f \in F_E$, $f \sim_E x$ if and only if $f_E g \sim x_E g$, where g has the properties in Axiom 4.2. This is Axiom A9 in Pires [33], and the results in that paper imply that \succ_E is derived from \succ via prior-by-prior Bayesian updating. Hence, CPMEU and CP coincide, and (2) follows from Theorem 1.

Conversely, assume that (2) holds. Consider a plan $f \in F_E$ and a prize $x \in X$ such that $u(x) = \min_{q \in C} \int_E u(f(\omega)) q(d\omega|E)$; one such prize must exist because X is connected and u is continuous. Now consider the tree $\{f, x\} \in F_E$ (cf. Def. 12); clearly, $\text{CPMEU}_{\{f, x\}}(\emptyset)$ is precisely the set containing f and x ; by (2), we have $f \sim_E \{f, x\} \sim_E x$, i.e. $f \sim_E x$.

Thus, \succ_E is consistent with MEU and prior-by-prior Bayesian updating of C . This implies that CPMEU and CP coincide, so Theorem (1) ensures that each preference is complete and transitive, and that Axioms 3.1, 3.2 and 3.3 hold.

Finally, it follows from the properties of the MEU functional (see also [33] and [38]) that, if f is a plan and g has the properties in Axiom 4.2, then $f \succ_E x$ if and only if $f_E g \succ_{x_E} g$, i.e. Axiom 4.2 holds.

References

- [1] Charalambos Aliprantis and Kim Border. *Infinite Dimensional Analysis*. Springer Verlag, Berlin, 1994.
- [2] Andrew Caplin and John Leahy. The recursive approach to time inconsistency. *Journal of Economic Theory*, 131:134–156, 2006.
- [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] Robin Cubitt. Rational dynamic choice and expected utility theory. *Oxford Economic Papers*, 48(1):1–19, January 1996. New Series.
- [5] Robin Cubitt, Chris Starmer, and Robert Sugden. Dynamic choice and the common ratio effect: An experimental investigation. *Economic Journal*, pages 1362–1380, 1998.
- [6] Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.
- [7] Larry Epstein. An axiomatic model of non-bayesian updating. University of Rochester, March 2004.
- [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, 2001.
- [10] Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, Cambridge, Massachusetts, 1991.
- [11] Paolo Ghirardato. Revisiting Savage in a conditional world. *Economic Theory*, 20:83–92, March 2002.
- [12] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

- [13] Itzhak Gilboa and David Schmeidler. Updating ambiguous beliefs. *Journal of Economic Theory*, 59:33–49, 1993.
- [14] Faruk Gul and Wolfgang Pesendorfer. The case for mindless economics. November 2005. mimeo, Princeton University.
- [15] Faruk Gul and Wolfgang Pesendorfer. The revealed preference theory of changing tastes. *Review of Economic Studies*, forthcoming.
- [16] Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.
- [17] Peter J. Hammond. Consistent plans, consequentialism, and expected utility. *Econometrica*, 57:1445–1449, 1989.
- [18] Eran Hanani and Peter Klibanoff. Updating multiple-prior preferences. Kellogg Graduate School of Management, July 2004.
- [19] Christopher Harris. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica*, 53(3):613–628, May 1985.
- [20] Jean-Yves Jaffray. Dynamic decision making with belief functions. In Yaeger et al., editor, *Advances in the Dempster-Shafer Theory of Evidence*. Wiley, 1994.
- [21] Edi Karni and Zvi Safra. Ascending bid auctions with behaviorally consistent bidders. *Annals of Operations Research*, pages 435–446, 1989.
- [22] Edi Karni and Zvi Safra. Behaviorally consistent optimal stopping rules. *Journal of Economic Theory*, 51:391–402, 1990.
- [23] Edi Karni and David Schmeidler. Atemporal dynamic consistency and expected utility theory. *Journal of Economic Theory*, 54:401–408, 1991.
- [24] Peter Klibanoff. Dynamic choice with uncertainty aversion. mimeo, Kellogg School of Management, March 1995.
- [25] Peter Klibanoff and Emre Ozdenoren. Subjective recursive expected utility. mimeo, Kellogg School of Management, 2004.
- [26] David M. Kreps. A representation theorem for “preference for flexibility”. *Econometrica*, 47(3):565–577, May 1979.

- [27] David M. Kreps. *Notes on the Theory of Choice*. Westview Press, Boulder and London, 1988.
- [28] David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46:185–200, 1978.
- [29] Mark J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27:1622–1668, 1989.
- [30] Roger B. Myerson. Axiomatic foundations of bayesian decision theory. Discussion Paper 671, The Center for Mathematical Studies in Economics and Management Science, Northwestern University, January 1986.
- [31] Martin J. Osborne and A. Rubinstein. *A Course on Game Theory*. MIT Press, Cambridge, MA, 1994.
- [32] Bezalel Peleg and Menahem Yaari. On the existence of a consistent course of action when tastes are changing. *Review of Economic Studies*, 40(3):391–401, July 1973.
- [33] C. P. Pires. A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152, 2002.
- [34] Rakesh Sarin and Peter Wakker. Dynamic choice and non-expected utility. *Journal of Risk and Uncertainty*, 17:87–119, 1998.
- [35] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [36] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.
- [37] Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, New Jersey, 1976.
- [38] Marciano Siniscalchi. Bayesian updating for general maxmin expected utility preferences. Mimeo, Princeton University, September 2001.
- [39] Marciano Siniscalchi. A behavioral characterization of plausible priors. *Journal of Economic Theory*, forthcoming.
- [40] Costis Skiadas. Recursive utility and preferences for information. *Economic Theory*, (12):293–312, 1998.
- [41] R.H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1955-1956.

- [42] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London and New York, 1991.
- [43] Tan Wang. Conditional preferences and updating. *Journal of Economic Theory*, 108:286–321, 2003.

B Web Appendix for “Dynamic Choice Under Ambiguity”

B.1 Sophistication and the Value of Information

For simplicity, we assume that the individual’s preferences over plans admit a numerical representation, denoted $V(\cdot)$, and that the individual is risk-neutral (so the restriction of V to X is affine).

An *information-acquisition problem* is described by a (finite, for simplicity) state space Ω , an information partition $\mathcal{I} = \{E_1, \dots, E_n\}$, a set A of actions, and a payoff function $\xi : \Omega \times A \rightarrow \mathbb{R}$. Let g be the partitional tree defined by $C_g(\emptyset) = \{*\}$, $\mathcal{F}_g(\emptyset, *) = \mathcal{I}$, $C_g([*, E_i]) = A$ for all $i = 1, \dots, n$ and $C_g([*, E_i, a, \{\omega\}]) = \{\xi(\omega, a)\}$ for all $i = 1, \dots, n$, $a \in A$ and $\omega \in E_i$. Clearly, the tree f in Fig. 6 exhibits this structure.

Now consider the following sets of plans consistent with g :

- $N_g = \{p \in g : \exists a \in \bar{A} \text{ such that } \forall i = 1, \dots, n, p([*, E_i]) = a\}$ contains plans that correspond to the choice of a single action without taking the information \mathcal{I} into account;
- $C_g = \{p \in F_\Omega^p : p \in g\}$ corresponds to the contingent plans that could be implemented by the individual *if* she could commit; and
- $CP_g(\emptyset)$ represents of course the plans that the individual can actually carry out.

Then:

- $\max_{p \in C_g} V(p) - \max_{p \in N_g} V(p)$ is the value of information with commitment; it represents the “intrinsic” value of information;
- $\max_{p \in CP_g(\emptyset)} V(p) - \max_{p \in N_g} V(p)$ is the value of information without commitment (assuming sophistication); it reflects the “realized” value of information; and
- $\max_{p \in C_g} V(p) - \max_{p \in CP_g(\emptyset)} V(p)$ is the value of commitment.

By a simple set-inclusion argument, the first and third quantities are positive; the second, on the other hand, may be positive or negative.

Thus, as noted above, the present framework provides a straightforward way to disentangle the intrinsic value of information from the value of commitment, and separately assess their impact on the realized value of information.

B.2 Consistent Planning for Infinite Trees

This section adapts the Sophistication and Weak Commitment axioms of Sec. 3.1.1 and the definition of Consistent Planning, Def. 9 to a class of infinite decision trees. This makes it possible to state an

extension of the main characterization result of Sec. 3.1.2 to this class. The main ideas behind this extension have been discussed above in §4.3.

Continue to assume that Ω is an arbitrary state space and Σ an algebra of events. The following topological assumptions on the set of labels A and the set of prizes X will be maintained.

Assumption B.1 X is a connected, separable metric space, and A is an uncountable, separable metric space. The metrics on X and A will be denoted by d_X and d_A respectively.

The role of the assumptions about A is to ensure that “enough labels” exist to represent any collection of continuation plans (recall that a cardinality assumption of A plays a similar role for finite trees): see the discussion preceding Lemma 9 for details.

The metrizable assumption on X simplifies the definition of a (pseudo)metric on the set of plans (see below). Connectedness and separability ensure that (with the sole modification described in Footnote 21), *the elicitation result of Section 3.2 will hold for infinite trees as well.*

The class of trees considered here allows for uncountably many actions at each history. For simplicity, the horizon is assumed finite; also, only a finite number of different events may be observed. For instance, this class includes, but is *not* limited to, all decision trees generated by a finite filtration. Finally, all prizes assigned at terminal histories must lie in some compact subset of X . Jointly with suitable assumptions on preferences, this implies that such trees are “bounded in preference”.

Definition 14 Consider an event $E \in \Sigma \setminus \{\emptyset\}$. A tree $f \subset \mathcal{H}_E \cup \mathcal{T}_E$ is **simple** if $\max_{h \in H} \lambda(h)$ is finite, the set $\{\xi(h) : h \in f \cap \mathcal{T}_E\}$ is contained in a compact subset K of X , and furthermore the collection $\{E(h) : h \in f \cap \mathcal{H}_E\}$ is finite. The set of simple trees starting at E is denoted F_E^* .

Clearly, $F_E \subset F_E^*$; furthermore, observe that a simple plan (i.e. a tree $p \in F_E^* \cap F_E^p$) is finite.

The relabeling relation \approx was defined above (cf. Def. 3) for trees in F_E ; however, the same definition applies verbatim to trees in F_E^* . Furthermore, it is convenient to introduce notation to indicate the collection of all relabelings of a tree f (finite or otherwise).

Definition 15 For every $f \in F_E^*$, denote the set of \approx -equivalence classes of f by $[f] = \{g \in F_E : g \approx f\}$.

A semimetric on the set of plans will now be introduced.³⁴ The following definition is motivated by two intuitive considerations. First, trees that exhibit a different structure must be “far” from one another, regardless of their payoffs. Second, the distance between two trees that have the same structure is a function of the distance between payoffs assigned at corresponding histories. Thus, action labels do not

³⁴While one could extend the definition provided below to arbitrary trees, this is not required for our purposes.

affect the proposed notion of distance between trees. The similarity of the following definition with that of relabeling (Def. 3) should then not be surprising:

Definition 16 (Plan distance) Consider an event $E \in \Sigma \setminus \{\emptyset\}$ and plans $p, p' \in F_E^*$. The distance between p and p' , denoted $d(p, p')$, is defined as follows. If there exists a bijection $\varphi : p \rightarrow p'$ such that

- (i) for all $h, \bar{h} \in p$, $h \leq \bar{h}$ iff $\varphi(h) \leq \varphi(\bar{h})$; and
- (ii) for all $h \in p \cap \mathcal{H}_E$, $E(h) = E(\varphi(h))$;

then

$$d(p, p') = \max_{h \in p \cap \mathcal{H}_E} \frac{d_X(\xi(h), \xi(\varphi(h)))}{1 + d_X(\xi(h), \xi(\varphi(h)))};$$

otherwise, $d(p, p') = 1$.

Notice that conditions (i) and (ii) are exactly as in Def. 3; there is no need for a counterpart to condition (iii) in the definition of relabeling, because attention is restricted to plans. In other words, conditions (i) and (ii) in the definition of relabeling fully characterize what it means for two plans to “have the same structure.” Therefore, according to Def. 16, the plans p and p' have the same structure if and only if $d(p, p') < 1$; otherwise, $d(p, p') = 1$.

The maximum in Def. 16 is achieved because a plan has only finitely many terminal histories.

Observe that d is *not* a metric on the set of plans in F_E^* . In particular, if $p \approx q$, then $d(p, q) = 0$. It is, however, a metric on the set of \approx -equivalence classes of plans, F_E / \approx .

Additional Notation: Closure of a set G of plans with respect to d will be denoted by $\text{cl } G$, and convergence of a sequence $\{p_n\}$ of plans to a plan p will be denoted by $p_n \rightarrow p$. Abusing notation, if G is a set of *equivalence classes* of plans in F_E , then $\text{cl } G$ will also be used to denote the d -closure of G in F_E / \approx .

Turn now to the axioms characterizing Consistent Planning. As should be expected, a notion of continuity of preferences is necessary. It is enough to require continuity for preferences over plans:

Axiom B.1 (Plan Continuity) For every non-empty $E \in \Sigma$ and every collection of plans $p, q, p_1, p_2, \dots \in F_E^p$: if $p_n \succ_E q$ (resp. $p_n \preccurlyeq q$) for all $n \geq 1$ and $p_n \rightarrow p$, then $p \succ_E q$ (resp. $p \preccurlyeq q$).

Next, consider Sophistication and Weak Commitment, the main behavioral axioms in Sec. 3.1.1. As noted in §4.3, Sophistication must be modified by adding a “closure” condition, which will be described shortly. However, both Sophistication and Weak Commitment require further modifications, arising from entirely different (and more mundane) considerations.

To elaborate, note that the finite-tree versions of Sophistication and Weak Commitment (Axioms 3.1 and 3.2) allow modifying a tree f at one history only. However, since only finite trees are considered, it is intuitively clear that, by repeatedly applying one or the other axiom, one eventually obtains a tree f' that

does not permit further applications of either axiom—an “irreducible” tree. This fact is used in the proof of sufficiency in Theorem 1.

It is clear that this will *not* be the case for an infinite tree: one cannot reach an irreducible tree in finitely many steps. To address this issue, both Sophistication and Weak Commitment must be strengthened so as to allow for the simultaneous modification of the original tree f at infinitely many histories. In particular, the formulation chosen here allows for the *modification of continuation plans at all non-terminal histories of a certain length t* . This is notationally intensive, but conceptually straightforward.

Preliminary notation is required to indicate that two trees share all histories of length t :

Definition 17 Consider two trees $f, f' \in F_E$. Then f and f' **agree up to time t** , denoted $f =_t f'$, if $\{h \in f : \lambda(h) \leq t\} = \{h \in f' : \lambda(h) \leq t\}$.

The appropriate version of Sophistication for simple infinite trees can now be stated.

Axiom B.2 (Sophistication for Simple Infinite Trees) Consider $E \in \Sigma \setminus \{\emptyset\}$, a tree $f \in F_E^*$, and an integer $t \geq 0$. If, for every history $h \in f$ of length t :

(i) for every $c \in C_f(h)$, $f(h, c)$ is a plan, and

(ii) there is a non-empty $B(h) \subset C_f(h)$ such that $p \succ_{E(h)} q$ for all $p \in \text{cl} \{f(h, b) : b \in B(h)\}$ and $q \in \text{cl} \{f(h, c) : c \in C_f(h) \setminus \text{cl} \{f(h, b) : b \in B(h)\}\}$,
then $f \sim g$ for every tree $g \in F_E^*$ that satisfies $f =_t g$ and $\{[g(h, c)] : c \in C_g(h)\} = \text{cl} \{[f(h, b)] : b \in B(h)\}$ for every $h \in f$ of length t .

As noted above, this axiom essentially entails modifying a tree f at all histories of a specified length t . Condition (i) in the present axiom restricts attention to trees wherein all time- t continuation subtrees are plans. This is for simplicity: a notion of distance, and an attendant notion of continuity, was only defined for plans, not arbitrary trees; hence, to apply the “closure” condition discussed in §4.3 to all length- t histories, it is necessary to ensure that each action at such histories correspond to a plan.³⁵

Condition (ii) is the key “closure” requirement. To ease interpretation, informally identify each choice available at h with the corresponding continuation plan; then condition (ii) requires that every choice in the “closure” of $B(h)$ be strictly preferred to any choice that lies in the “closure” of $C_f(h)$, but not in the closure of $B(h)$.

If Conditions (i) and (ii) hold, then the axiom asserts the indifference of f with a tree g that differs from f in that, at every history h of length t , the continuation tree $f(h)$ is, again loosely speaking, replaced with the “closure” of $B(h)$.

³⁵An extension that drops this restriction is feasible, but complicates the notation further.

Temporarily defer the discussion of technical aspects; to fix ideas, suppose that $X = \mathbb{R}$, and consider a very simple tree f where the set of choices available at the initial history is the half-open interval $[0, 1]$.³⁶ Then, with the obvious ordering of prizes in $X = \mathbb{R}$, the set $B(\emptyset) = (\frac{1}{2}, 1)$ satisfies Condition (ii), and the axiom implies that f is indifferent to a tree g that offers the choices $[\frac{1}{2}, 1]$ at the initial history. If one now applies the same axiom to the tree g , taking $B(\emptyset) = \{1\}$, one concludes that f and g are both indifferent to the constant tree corresponding to the prize 1. As can be seen, the axiom encodes the key assumption discussed in Sec. 4.3: a set of alternatives is just as good as its “supremum”, regardless of whether or not this can be achieved.

Turn now to technical aspects. The equality $\{[g(h, c)] : c \in C_g(h)\} = \text{cl} \{[f(h, b)] : b \in B(h)\}$, where h is a history of length t in f , is easiest to interpret if one again identifies choices in $B(h)$ with the corresponding continuation plans in f . With this informal convention, the above equality requires that every choice c available at the history h in the tree g correspond (up to relabeling) to a choice in the “closure” of $B(h)$, and conversely.³⁷

Notice that, in particular, taking $B(h) = C_f(h)$ for all histories h of length t trivially satisfies condition (ii); thus, Axiom B.2 requires that the individual be indifferent to taking the “closure” of all choice sets at length- t histories. As a further special case, if h is a terminal history of length t , so $C_f(h)$ is a singleton set $\{x\}$, then conditions (i) and (ii) both hold at h , and furthermore the tree g in Axiom B.2 must coincide with f at history h . In other words, terminal histories of length t are left unchanged.

Finally, observe that, since all prizes in f lie in a compact subset K of the metric space X , the same is true of all prizes in any tree g as in the Axiom. Thus, g is indeed a simple tree according to Def. 14.

The Weak Commitment axiom only requires mechanical modifications of the type described immediately before Axiom B.2.

Axiom B.3 (Weak Commitment for Simple Infinite Trees) Consider $E \in \Sigma \setminus \{\emptyset\}$, a tree $f \in F_E$, and an integer $t \geq 0$. If, for every history $h' \in f$ with $\lambda(h') = t + 1$,

(i) for every $c \in C_f(h')$, $f(h', c)$ is a plan, and

(ii) for all $c, c' \in C_f(h')$, $f(h', c) \sim_{E(h')} f(h', c')$,

then $f \sim_E g$ for all $g \in F_E$ such that $f =_t g$ and $\{[g(h, c)] : c \in C_g(h)\} = \{[g'] : g' \in F_{E(h)} \text{ allows one-period commitment in } f(h)\}$ for every history $h \in f$ of length t .

³⁶Formally, $f = \{\emptyset\} \cup \{[x] : x \in [0, 1]\}$.

³⁷To see this, pick $c \in C_g(h)$. Then the axiom requires that the \approx -equivalence class $[g(h, c)]$ lie in the closure of $\{[f(h, b)] : b \in B(h)\}$, which is the case precisely if there is a sequence $\{b^n\} \subset B(h)$ such that $f(h, b^n) \rightarrow g(h, c)$. Conversely, if $f(h, b^n) \rightarrow p$ for some $\{b^n\} \subset B(h)$, then $[p] \in \{[g(h, c)] : c \in C_g(h)\}$, which is the case precisely if $p \approx g(h, c)$ for some $c \in C_g(h)$. Notice that this formulation allows for the possibility that two or more choices in $C_g(h)$ correspond to plans that are equivalent up to relabeling: this is not the case in the finite-tree version of Sophistication. However, the Simplification axiom renders this difference irrelevant.

As was the case for Sophistication, the above version of Weak Commitment restricts attention to trees wherein continuation plans at histories of length $t + 1$ are plans: see condition (i). Observe that, consequently, each tree allowing one-period commitment in $f(h)$ (cf. Def. 7) is a plan.

The axiom requires that, if conditions (i) and (ii) hold, then, for every history h of length t , $f(h)$ can be replaced with one of its one-period-commitment versions (cf. Def. 8).³⁸

Finally, the definition of consistent planning requires only minimal modifications.

Definition 18 (Consistent Planning for Simple Trees) Consider a tree $f \in F_E^*$. For every terminal history $h \in f \cap \mathcal{T}_E$, let $\text{CPS}_f(h) = \{f(h)\}$. Inductively, if $h \in f \cap \mathcal{H}_E$ and $\text{CPS}_f(h')$ has been defined for all $h' \in f$ with $h < h'$, let

$$\begin{aligned} \text{CPS}_f^0(h) &= \left\{ p \in F_{E(h)} : \exists c \in C_f(h) \text{ s.t. } C_p(\emptyset) = \{c\}, \mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h, c) \text{ and} \right. \\ &\quad \left. \forall D \in \mathcal{F}_p(\emptyset, c), p([c, D]) \in \text{CPS}_f([h, c, D]) \right\} \\ \text{CPS}_f(h) &= \left\{ p \in \text{cl } \text{CPS}_f^0(h) : \forall p' \in \text{CPS}_f^0(h), p \succ_{E(h)} p' \right\} \end{aligned}$$

A tree $g \in F_E$ is a **consistent-planning solution** of f if $g \in \text{CPS}_f(\emptyset)$.

Due to the definition of distance on F_E^p , if $p \in \text{CPS}_f^0(h)$, then any $q \approx p$ is an element of $\text{cl } \text{CPS}_f^0(h)$ as well. Consequently, if p is a consistent-planning solution of f , then so is any relabeling of p .³⁹

It is formally necessary to restate the relabeling-invariance assumption for simple infinite trees.

Assumption B.2 For all non-empty $E \in \Sigma$ and all $f, g \in F_E^*$: $f \approx g$ implies $f \sim_E g$.

The counterpart to Theorem 1 for simple infinite trees can finally be stated.

Theorem 8 Consider a system of preferences $\{\succ_E\}_{\emptyset \neq E \in \Sigma}$ that satisfies Assumption B.2 and such that, for every non-empty $E \in \Sigma$, the restriction of \succ_E to F_E^p is complete and transitive, and satisfies Axiom B.1. Then the following statements are equivalent.

1. For every non-empty $E \in \Sigma$, \succ_E is complete and transitive on all of F_E^* ; furthermore, Axioms B.2, B.3 and 3.3 hold;
2. for any $E \in \Sigma$ and every pair of acts $f, g \in F_E^*$: $f \succ_E g$ if and only if $p \succ_E q$ for some (hence all) $p \in \text{CPS}_f(\emptyset)$ and $q \in \text{CPS}_g(\emptyset)$.

³⁸There is a slight difference with Def. 8: Axiom B.3 allows for the possibility that two or more choices in $C_h(h)$ lead to plans that are equivalent up to relabeling (as was the case in Axiom B.2). This slightly simplifies notation in the proof of Theorem 8, and is irrelevant under the Simplification axiom.

³⁹In the proposed definition for finite trees, a relabeling of a consistent-planning solution is not itself a consistent-planning solution; however, it is indifferent to it. Therefore, under relabeling indifference, Theorems 1 and 8 have exactly the same behavioral content.

B.3 Proof of Theorem 8 (Consistent Planning for Simple Infinite Trees)

B.3.1 Preliminaries

Begin with two simple results related to the notion of plan convergence. The first follows immediately from Def. 16.

Remark B.1 Consider a sequence of plans $\{p^n\} \subset F_E^*$ and a plan $p \in F_E^*$. Let $\{a^n\} = C_{p^n}(\emptyset)$ and $\{a\} = C_p(\emptyset)$. Assume that $a \in A$. Then $p^n \rightarrow p$ iff for every n , $\mathcal{F}_{p^n}(\emptyset, a^n) = \mathcal{F}_p(\emptyset, a)$ and $p^n([a^n, D]) \rightarrow p([a, D])$ for every $D \in \mathcal{F}_p([\emptyset, a])$.

Of course, it is also the case that, if $a \in X$, i.e. $p = \{\emptyset, [a]\}$, then $p^n = \{\emptyset, [a^n]\}$ for some $\{a^n\} \subset X$ such that $a^n \rightarrow a$.

The second result relates the closure of sets of plans to the closure of *equivalence classes* of plans. Recall that the semimetric d in Def. 16 can also be viewed as a metric on each F_E^p / \approx , by letting $d([p], [q]) = d(p, q)$ for every $p, q \in F_E^p$. Thus, $[p^n] \rightarrow [p]$ means that $d([p^n], [p]) = d(p^n, p) \rightarrow 0$; furthermore, if $G \subset F_E^p / \approx$, then $\text{cl } G = \{[p] \in F_E^p / \approx : \exists \{p^n\} \subset G, [p^n] \rightarrow [p]\}$.

Remark B.2 Let $E \in \Sigma \setminus \emptyset$ and $G \subset F_E^p$. Then $\text{cl } G = \bigcup \text{cl } \{[g] : g \in G\} = \text{cl } \bigcup \{[g] : g \in G\}$.

Proof: Let $[G] = \{[g] : g \in G\}$. Suppose $p \in \text{cl } G$, so there is $\{p^n\} \subset G$ such that $p^n \rightarrow p$. Then $[p^n] \in [G]$, and furthermore $d([p], [p^n]) \rightarrow 0$. Thus, $[p] \in \text{cl } [G]$. Conversely, suppose that $[p^n] \rightarrow [p]$ and $[p^n] \in [G]$. Then by def. $p^n \in G$ and $d(p^n, p) \rightarrow 0$. This proves the first equality.

For the second equality, suppose $p \in \bigcup \text{cl } [G]$, so $[p] \in \text{cl } [G]$. Hence there is some sequence $\{p^n\} \subset G$ such that $d([p^n], [p]) \rightarrow 0$: but this by definition means that $p^n \rightarrow p$. Furthermore, $p^n \in \bigcup [G]$, so $p \in \text{cl } \bigcup [G]$. Conversely, suppose $p \in \text{cl } \bigcup [G]$, so $p^n \rightarrow p$ for some sequence $\{p^n\} \subset \bigcup [G]$. Then, for every n , $[p^n] \in [G]$, and by definition $d([p^n], [p]) \rightarrow 0$. Thus, $[p] \in \text{cl } [G]$, and therefore $p \in \bigcup \text{cl } [G]$. ■

Notation. It is useful to let $F^X = \{\emptyset, [x] : x \in X\}$ be the set of plans corresponding to an immediate payoff. Observe that, for every $p \in F^X$, $[p] = \{p\}$ and similarly $\text{cl } \{p\} = \{p\}$.

As noted above, A is assumed to be an uncountable, separable metric space; thus, it has the cardinality of the continuum. Together with the assumption that X is a metric space, this leads to the following simple result:

Lemma 9 Fix $E \in \Sigma \setminus \{\emptyset\}$ and $f \in F_E^*$. If $G \subset \{p \in F_E^p : p \subset f\}$, then $\text{cl } \{[p] : p \in G\}$ is a compact subset of F_E^p / \approx , and there exists $g \in F_E$ such that $\{[g(\emptyset, c)] : c \in C_g(\emptyset)\} = \text{cl } \{[p] : p \in G\}$.

In particular, *this ensures that Axioms B.2 and B.3 do not hold vacuously*. It will also be invoked several times in the next section.

Proof: It is enough to prove the result for the case $C_f(\emptyset) \cap X = \emptyset$. To see this, let $G^X = G \cap F^X$; by assumption, $C_f(\emptyset) \cap X$ lies in a compact subset of X , and it is immediate from the definition of d that G^X lies in a compact subset of F^X , and hence of F_E^p / \approx . Since $\text{cl} \{[p] : p \in G\} = \text{cl} \{[p] : p \in G^X\} \cup \text{cl} \{[p] : p \in G \setminus G^X\}$ by our definition of distance, and $\text{cl} \{[p] : p \in G^X\}$ is compact, then $\text{cl} \{[p] : p \in G\}$ is closed if $\text{cl} \{[p] : p \in G \setminus G^X\}$ is. Furthermore, if $\{[g(\emptyset, a)] : a \in C_g(\emptyset)\} = \text{cl} \{[p] : p \in G \setminus G^X\}$, then the tree $g' = g \cup \text{cl} G^X$ has the required properties.

Thus, assume $C_f(\emptyset) \cap X = \emptyset$. As above, for any set C of plans, let $[C] = \{[p] : p \in C\}$. Since $\{E(h) : h \in H\}$ is finite, there is a finite partition G_1, \dots, G_N of G such that $p \in G_n$ and $p' \in G_m$ have the same “structure”, i.e. admit a bijection φ with the properties in Def. 16, if and only if $n = m$. Also, clearly $\text{cl} [G] = \text{cl} [G_1] \cup \dots \cup \text{cl} [G_N]$, and $\text{cl} [G_n] \cap \text{cl} [G_m] = \emptyset$ whenever $n \neq m$.

Since f is simple, there exists a compact set $K \subset X$ such that $\xi(z) \in K$ for all $z \in f \cap \mathcal{T}_E$. Then, for every n , $\text{cl} [G_n]$ is (homeomorphic to) a subset of K^{M_n} , where M_n is the number of terminal histories in any $p \in G_n$; therefore, this set is compact, and hence so is $\text{cl} [G]$. This implies that $\text{cl} [G]$ has at most the cardinality of the continuum.

It follows that there exists a subset \bar{A} of A such that \bar{A} and $\text{cl} [G]$ have the same cardinality, so there exists a bijection $\alpha : \text{cl} [G] \rightarrow \bar{A}$. Now define $g \in F_{E(h)}^*$ as follows. First, let $C_g(\emptyset) = \{\alpha([p]) : [p] \in \text{cl} [G]\}$. Then, for every $[p] \in \text{cl} [G]$, let $\mathcal{F}_g(\emptyset, \alpha([p])) = \mathcal{F}_p(\emptyset, a)$, where p is a member of $[p]$ and $C_p(\emptyset) = \{a\}$; and for every $D \in \mathcal{F}_g(\emptyset, \alpha([p]))$, let $g([\alpha([p]), D]) = p([a, D])$. Clearly, $g(\emptyset, \alpha([p])) \approx p$ by Lemma 7. Thus, for every $a \in C_g(\emptyset)$, $g(\emptyset, a) \approx p$ for some $p \in [p] \in \text{cl} [G]$, and conversely for every $[p] \in \text{cl} [G]$, $\alpha([p])$ satisfies $g(\emptyset, \alpha([p])) \approx p$. This proves the claim. ■

Clearly, the Lemma also implies the existence of $g \in F_E$ such that $\{[g(\emptyset, c)] : c \in C_g(\emptyset)\} = \{[p] : p \in G\}$. Furthermore:

Lemma 10 *Fix $E \in \Sigma \setminus \{\emptyset\}$ and $f \in F_E^*$. If $G \subset \{p \in F_E : p \text{ is a plan and } p \subset f\}$ and \succ_E satisfies Axiom B.1, then the set $\{p \in \text{cl} G : \forall p' \in \text{cl} G, p \succ_E p'\}$ is non-empty and closed.*

Proof: Lemma 9 shows that $\text{cl} [G] = \text{cl} \{[p] : p \in G\}$ is a compact subset of F_E^p / \approx . Now define a binary relation \succ_E^\approx on F_E^p / \approx by letting $[p] \succ_E^\approx [q]$ iff $p \succ_E q$ for all $p, q \in F_E^p$. By Relabeling Invariance, \succ_E^\approx is well-defined, and it satisfies the appropriate version of Axiom B.1. The existence of a \succ_E^\approx -maximal tree in $\text{cl} [G]$ now follows e.g. from [1, Theorem 2.41].

By Remark B.2, $\text{cl} G = \bigcup \text{cl} [G]$; it is then clear that, if $[p] \in \text{cl} [G]$ satisfies $[p] \succ_E^\approx [q]$ for all $[q] \in \text{cl} [G]$, then $p \in \text{cl} G$ and $p \succ_E q$ for all $q \in \text{cl} G$: thus, the set $G^* = \{p \in \text{cl} G : \forall p' \in \text{cl} G, p \succ_E p'\}$ is non-empty. Conversely, if $p \in G^*$, then $[p] \in \text{cl} [G]$ and, for every $[q] \in \text{cl} [G]$, $[p] \succ_E^\approx [q]$.

Furthermore, suppose $p^n \rightarrow p$ for some $\{p^n\} \subset G^*$ and $p \in F_E^p$. Then $[p^n] \rightarrow [p]$, and for every $[q] \in \text{cl}[G]$, $[p^n] \succ_E^{\approx} [q]$. By Plan continuity for \succ_E^{\approx} , $[p]$ is also \succ_E^{\approx} -maximal, so $p \in G^*$. Thus, G^* is closed. ■

Notice that, in particular, this ensures that the sets $\text{CPS}_f(h)$ in Def. 18 are non-empty and closed.

B.3.2 Sufficiency

Assume first that (i) in Theorem 8 holds. Fix a tree $f \in F_E^*$, let $T = \max_{h \in H} \lambda(h)$ and let f^0, \dots, f^T be the sequence of trees constructed as follows. First, let $f^T = f$. Then, for $t = T - 1, \dots, 0$, construct a tree f^t such that $f^t =_{t-} f$ and, for every length- t history h , $\{[f^t(h, c)] : c \in C_{f^t}(h)\} = \{[p] : p \in \text{CPS}_f(h)\}$; the existence of such trees follows from Lemma 9.

It will be shown that, for all $t = 0, \dots, T - 1$, $f^t \sim_E f^{t+1}$. To streamline the presentation, for all $h \in f \cap \mathcal{T}_E$, let

$$\text{Comm}_f^t(h) = \{\bar{g} \in F_{E(h)} : \bar{g} \text{ allows one-period commitment in } f^t(h)\}.$$

Claim 1. Fix $t \in \{0, \dots, T - 1\}$ and consider a non-terminal history $h \in H$ with $\lambda(h) = t$. Then $\{[p] : p \in \text{CPS}_f^0(h)\} = \{[p] : p \in \text{Comm}_f^{t+1}(h)\}$.

Proof: Clearly, any $p \in F^X$ allows one-period commitment in any tree; conversely, if some $p = \{\emptyset, [x]\} \in F^X$ allows one-period commitment in $f^{t+1}(h)$, then $x \in C_{f^{t+1}}(h)$; since $f^{t+1} =_{t+1-} f$, $x \in C_f(h)$ and so obviously $p \in \text{CPS}_f^0(h)$. So, it is enough to restrict attention to equivalence classes of plans $p \notin F^X$.

Consider $p \in \text{CPS}_f^0(h) \setminus F^X$; let $\{a\} = C_p(\emptyset)$, so $a \in A$, and note that, by definition, for every $D \in \mathcal{F}_p(\emptyset, a) = \mathcal{F}_f(h, a)$, $p([a, D]) \in \text{CPS}_f([h, a, D])$. By construction, $\{[f^{t+1}([h, a, D], b)] : b \in C_{f^{t+1}}([h, a, D])\} = \{[q] : q \in \text{CPS}_f([h, a, D])\}$; thus, there exists $b_D \in C_{f^{t+1}}([h, a, D])$ such that $p([a, D]) \approx f^{t+1}([h, a, D], b_D)$. Finally, let $p' \in F_{E(h)}$ be a plan such that $p' =_{t+1-} p$ and, for all $D \in \mathcal{F}_p(\emptyset, a)$, $p'([a, D]) = f^{t+1}([h, a, D], b_D)$. Clearly, $p' \in \text{Comm}_f^{t+1}(h)$ and $p \approx p'$ by Lemma 4. Thus, $\{[p] : p \in \text{CPS}_f^0(h)\} \subset \{[p] : p \in \text{Comm}_f^{t+1}(h)\}$.

In the opposite direction, consider $p \in \text{Comm}_f^{t+1}(h) \setminus F^X$; thus, there is $a \in C_f(h) = C_{f^{t+1}}(h)$ with $a \in A$ and, for every $D \in \mathcal{F}_f(h, a) = \mathcal{F}_{f^{t+1}}(h, a)$, an action $b_D \in C_{f^{t+1}}([h, a, D])$ such that $C_p(\emptyset) = \{a\}$, $\mathcal{F}_p(\emptyset, a) = \mathcal{F}_f(h, a)$ and for every D as above $p([a, D]) = f^{t+1}([h, a, D], b_D)$. But since $\{[f^{t+1}([h, a, D], b)] : b \in C_{f^{t+1}}([h, a, D])\} = \{[q] : q \in \text{CPS}_f([h, a, D])\}$, there exists $q_D \in \text{CPS}_f([h, a, D])$ such that $p([a, D]) = f^{t+1}([h, a, D], b_D) \approx q_D$. Finally, let $p' \in F_{E(h)}$ be a plan such that $p' =_{t+1-} p$ and, for all $D \in \mathcal{F}_p(\emptyset, a)$. Then clearly $p' \in \text{CPS}_f^0(h)$ and $p' \approx p$. The claim follows. ■

Claim 2. For all $t = 0, \dots, T - 1$, $f^t \sim_E f^{t+1}$.

Proof: Fix one such t and let $f_0^t \in F_E$ be a tree such that (i) $f_0^t =_t f$, and (ii) for all $h \in f$ with $\lambda(h) = t$, $\{[f_0^t(h, c)] : c \in C_{f_0^t}(h)\} = \{[p] : p \in \text{CPS}_f^0(h)\}$; such a tree exists by Lemma 9.⁴⁰ By (i), $f_0^t =_t f^{t+1}$, and by Claim 1, for all $h \in f$ with $\lambda(h) = t$, $\{[f_0^t(h, c)] : c \in C_{f_0^t}(h)\} = \{[p] : p \in \text{Comm}_f^{t+1}(h)\}$.

Furthermore, consider a, D such that $[h, a, D] \in f$; fix two choices $c, c' \in C_{f^{t+1}}([h, a, D])$. By the definition of f^{t+1} , there exist $p, p' \in \text{CPS}_f([h, a, D])$ such that $f^{t+1}([h, a, D], c) \approx p$ and $f^{t+1}([h, a, D], c') \approx p'$; since $p \sim_{[h, a, D]} p'$, it follows that $f^{t+1}([h, a, D], c) \sim_{[h, a, D]} f^{t+1}([h, a, D], c')$ as well.

Weak Commitment (Axiom B.3) can now be invoked (letting $f = f^{t+1}$ and $g = f_0^t$ in the statement of the axiom) to conclude that $f^{t+1} \sim_E f_0^t$.

Next, let $f_1^t \in F_E$ be a tree such that (i) $f_1^t =_t f$ and (ii) for every $h \in f$ with $\lambda(h) = t$, $\{[f_1^t(h, a)] : a \in C_{f_1^t}(h)\} = \text{cl} \{[p] : p \in \text{CPS}_f^0(h)\}$. Then also $f_1^t =_t f_0^t$, and one can invoke Axiom B.2 [with $f = f_0^t$, $B(h) = C_{f_0^t}(h)$ and $g = f_1^t$] to conclude that $f_1^t \sim_E f_0^t$.

Now, for every $h \in f$ with $\lambda(h) = t$, let $C(h) = \{c \in C_{f_1^t}(h) : \forall c' \in C_{f_1^t}(h), f_1^t(h, c) \succ_{E(h)} f_1^t(h, c')\}$. Let $G = \{f_1^t(h, c) : c \in C_{f_1^t}(h)\}$ for simplicity; then, by Lemma 10, the set $G^* = \{p \in \text{cl } G : \forall p' \in G, p \succ_E p'\}$ is non-empty and closed.

Observe that $[G] \equiv \{[p] : p \in G\} = \{[f_1^t(h, c)] : c \in C_{f_1^t}(h)\} = \text{cl} \{[p] : p \in \text{CPS}_f^0(h)\}$, i.e. $[G] = \text{cl } [G]$. Remark B.2 then implies that $\text{cl } G = \bigcup \text{cl } [G] = \bigcup [G] = \{q \in F_{E(h)}^p : \exists c \in C_{f_1^t}(h), q \approx f_1^t(h, c)\}$. Hence, every $p \in G^*$ is such that $p \approx f_1^t(h, c)$ for some $c \in C_{f_1^t}(h)$; indeed, it is clear that $c \in C(h)$, because $G \subset \text{cl } G$. Since $d(p, f_1^t(h, c)) = 0$, $G^* \subset \text{cl} \{f_1^t(h, c) : c \in C(h)\}$; furthermore, for every $c \in C(h)$ and $q \in \text{cl } G$, we have $q \approx f_1^t(h, c')$ for some $c' \in C_{f_1^t}(h)$, so $f_1^t(h, c) \succ_{E(h)} q$: thus, $\{f_1^t(h, c) : c \in C(h)\} \subset G^*$. Hence, $G^* = \text{cl} \{f_1^t(h, c) : c \in C(h)\}$. In particular, $C(h)$ is non-empty.

It is clear that the tree f_1^t satisfies Condition (i) in Axiom B.2. To verify that it satisfies Condition (ii) as well for $B(h) = C(h)$, let $p \in \text{cl} \{f_1^t(h, c) : c \in C(h)\} = G^*$ and $q \in \text{cl} \{f_1^t(h, c') : c' \in C_{f_1^t}(h)\} \setminus \text{cl} \{f_1^t(h, c) : c \in C(h)\} = \{q \in F_{E(h)}^p : \exists c' \in C_{f_1^t}(h), q \approx f_1^t(h, c')\} \setminus G^*$. Then it is clear that $p \succ_{E(h)} q$, as required.

It will now be shown that $\{[f^t([h, c])] : c \in C_{f^t}(h)\} = \text{cl} \{[f_1^t(h, c)] : c \in C(h)\}$ for any history h as above. Note first that the set on the r.h.s. is simply $\{[f_1^t(h, c)] : c \in C(h)\}$, because Lemma 10 implies that the latter is closed. Consider first $c \in C_{f^t}(h)$; thus, by construction, $f^t(h, c) \in \text{CPS}_f(h)$. [Recall that $\text{CPS}_f(h)$ is closed under relabeling, unlike $\text{CP}_f(h)$.] Hence, $f^t(h, c) \succ_{E(h)} q$ for all $q \in \text{cl } \text{CPS}_f^0(h)$: in particular, $f^t(h, c) \succ_{E(h)} f_1^t(h, c')$ for all $c' \in C_{f_1^t}(h)$ [again using the fact that $\text{cl } \text{CPS}_f^0(h)$ is closed under relabeling]. Moreover, since $f^t(h, c) \in \text{cl } \text{CPS}_f^0(h)$, there is $c_1 \in C_{f_1^t}(h)$ such that $f^t(h, c) \approx f_1^t(h, c_1)$. It follows that $c_1 \in C(h)$. Thus, $\{[f^t([h, c])] : c \in C_{f^t}(h)\} \subset \text{cl} \{[f_1^t(h, c)] : c \in C(h)\}$.

Conversely, consider $c \in C(h)$, so $f_1^t(h, c) \in \text{cl } \text{CPS}_f^0(h)$. By assumption, for every $q \in \text{cl } \text{CPS}_f^0(h)$, there is $c' \in C_{f_1^t}(h)$ such that $f_1^t(h, c') \approx q$, and by construction $f_1^t(h, c) \succ_{E(h)} q$. Thus, $f_1^t(h, c) \in \text{CPS}_f(h)$. This implies that there exists $c' \in C_{f^t}(h)$ such that $f_1^t(h, c) \approx f^t(h, c')$. Thus, $\{[f^t([h, c])] : c \in C_{f^t}(h)\} \supset \text{cl} \{[f_1^t(h, c)] : c \in C(h)\}$, and the proof of the subclaim is complete.

⁴⁰Also note that this is automatically true for $f \in f \cap \mathcal{T}_E$.

Finally, invoke Sophistication (Axiom B.2 [with $f = f_1^t$, $B(h) = C(h)$ and $g = f^t$] to conclude that $f_1^t \sim_E f^t$. By transitivity, $f^{t+1} \sim_E f_0^t \sim_E f_1^t \sim_E f^t$, and the proof of the Claim is complete. ■

The proof of sufficiency can now be completed, essentially as for finite trees. Observe that $f = f^T \sim_E f^0$ by transitivity, and for all $c, c' \in C_{f^0}(\emptyset)$, $f^0(\emptyset, c) \sim_E f^0(\emptyset, c')$: thus, Simplification implies $f^0(\emptyset, c) \sim_E f$ for every $c \in C_f(\emptyset)$. Now consider $g \in \text{CPS}_f(\emptyset)$: by construction, there is some $c_p \in C_{f^0}(\emptyset)$ such that $p \approx f^0(\emptyset, c_p)$; by Relabeling Invariance and transitivity, $g \sim_E f^0(\emptyset, c_p) \sim_E f$, as required.

B.3.3 Necessity

As is the case for finite trees, $\text{CPS}_f(h) = \text{CPS}_{f(h)}(\emptyset)$. Furthermore, we have the following simple facts.

Lemma 11 *Let $f \in F_E^*$ and $h \in f$. If $f(h, c) \in F_{E(h)}^p$ for every $c \in C_f(h)$, then:*

- (i) $\text{CPS}_f^0(h) \subset \bigcup \{[f(h, c)] : c \in C_f(h)\}$;
- (ii) $\bigcup \{[f(h, c)] : c \in C_f(h)\} \subset \text{cl } \text{CPS}_f^0(h)$;
- (iii) $\text{cl } \text{CPS}_f^0(h) = \bigcup \text{cl } \{[f(h, c)] : c \in C_f(h)\} = \text{cl } \{f(h, c) : c \in C_f(h)\}$;
- (iv) if $C_f(h) = \{c\}$, then $\text{CPS}_f(h) = [f(h, c)]$.

Proof: Begin with (iv). Argue by induction on the maximum length of $f(h, c)$. If $c \in X$, the claim is trivially true. Thus, suppose the claim is true for $h \in f$ and $c \in C_f(h)$ with $\max\{\lambda(h') : h' \in f(h, c)\} \leq L$, and consider h such that $\max\{\lambda(h') : h' \in f(h, c)\} = L + 1$. Clearly, $c \in A$; by the induction hypothesis, the claim is true at all histories $[h, c, D]$, for every $D \in \mathcal{F}_f(h, c)$. Then $\text{CPS}_f^0(h) = \{p : C_p(\emptyset) = \{c\} \text{ and } \forall D \in \mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h, c), p([c, D]) \approx f([h, c, D])\}$, which implies that the claim is again true.

Now consider (i). Since, for every $a \in C_f(h) \cap A$ and $D \in \mathcal{F}_f(h, a)$, $f([h, a, D]) \in F_D^p$, (iv) implies that $\text{CPS}_f([h, a, D]) = [f([h, a, D])]$. Thus, consider $p \in \text{CPS}_f^0(h)$, and let $\{c\} = C_p(\emptyset) \subset C_f(h)$. Any $c \in X$, trivially belongs to the union in the r.h.s. of (i). If $c \in A$, then $\mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h, c)$ and for every D in this set, $p([c, D]) \approx f([h, c, D])$; Lemma 4 then implies that $p \approx f(h, c)$, which implies the claim.

For (ii), it is clear that any $f(h, c)$ with $c \in C_f(h) \cap X$ also belongs to $\text{CPS}_f^0(h)$. Thus, consider $p' \in \bigcup \{[f(h, a)] : a \in C_f(h) \cap A\}$, so $p' \approx f(h, a)$ for some $a \in C_f(h) \cap A$. Let $\{a'\} = C_{p'}(\emptyset)$. Then $\mathcal{F}_{p'}(\emptyset, a') = \mathcal{F}_f(h, a)$ and for every D in this collection $p'([a', D]) \approx f([h, a, D])$. Now consider the plan $p \in F_{E(h)}^p$ such that $C_p(\emptyset) = \{a\}$, $\mathcal{F}_p(\emptyset, a) = \mathcal{F}_{p'}(\emptyset, a')$ and for every D in this set $p([a, D]) = p'([a', D])$. Then clearly $p \approx p'$, and furthermore, arguing as above, $p \in \text{CPS}_f^0(h)$. Since $d(p, p') = 0$, the claim follows.

Finally, (iii) follows from (i), (ii) and Remark B.2. ■

Turn now to the actual proof of necessity. Assume that (ii) in Theorem 8 holds. Completeness and transitivity is immediate. Thus, the three axioms in (i) will be considered in turn.

Axiom B.2, Sophistication. Consider $t, f \in F_E^*, H_t, g$ and $B(h), h \in H_t$ as in the Axiom. Consider $p \in \text{CPS}_f(h) \subset \text{cl CPS}_f^0(h) = \text{cl} \{f(h, c) : c \in C_f(h)\}$, where the equality follows from Lemma 11; by construction, $p \succ_{E(h)} f(h, b)$ for all $b \in B(h)$, so condition (ii) in the axiom implies that $p \in \text{cl} \{f(h, b) : b \in B(h)\}$.⁴¹ By assumption $\text{cl} \{f(h, b) : b \in B(h)\} = \{[g(h, c)] : c \in C_g(h)\}$, so by Remark B.2 and (iii) in Lemma 11, $p \in \text{cl} \{g(h, c) : c \in C_g(h)\} = \text{cl CPS}_g^0(h)$. Furthermore, consider $q \in \text{cl CPS}_g^0(h) = \text{cl} \{g(h, c) : c \in C_g(h)\}$, where again the equality follows from (iii) in Lemma 11. By Remark B.2, $q \in \bigcup \text{cl} \{[g(h, c)] : a \in C_g(h)\} = \bigcup \{[g(h, c)] : c \in C_g(h)\} = \bigcup \text{cl} \{f(h, b) : b \in B(h)\} = \text{cl} \{f(h, b) : b \in B(h)\}$; hence, there is a sequence $\{b^n\} \subset B(h)$ such that $f(h, b^n) \rightarrow q$, and the assumptions imply that $p \succ_{E(h)} f(h, b^n)$ for each n . Plan Continuity now ensures that $p \succ_{E(h)} q$. Thus, $p \in \text{CPS}_g(h)$, i.e. $\text{CPS}_f(h) \subset \text{CPS}_g(h)$.

Now consider $q \in \text{CPS}_g(h)$. Fix $p \in \text{CPS}_f(h)$, which must exist by Lemma 10; the result just established implies that $p \sim_{E(h)} q$. Then, for every $p' \in \text{cl CPS}_f^0(h)$, we have $q \sim_{E(h)} p \succ_{E(h)} p'$. Thus, $q \in \text{CPS}_f(h)$.

Therefore, for every $h \in H_t$, $\text{CPS}_f(h) = \text{CPS}_g(h)$. Inductively, this implies that $\text{CPS}_f(\emptyset) = \text{CPS}_g(\emptyset)$, and hence, by (ii) in Theorem 8, $f \sim_E g$.

Axiom B.3, Weak Commitment. Consider $t, f \in F_E, H_t$ and g as in the Axiom. Analogously to the notation in the sufficiency part of the proof, let $\text{Comm}_f(h)$ be the set of plans that allow 1-period commitment in $f(h)$, where $h \in H_t$; thus, by assumption, for every $h \in H_t$, $\{[g(h, c)] : c \in C_g(h)\} = \{[p] : p \in \text{Comm}_f(h)\}$.

Consider $h \in H_t$ and an arbitrary history of the form $[h, a, D] \in f$. Since every choice $c \in C_f([h, a, D])$ corresponds to a plan, Lemma 11 implies that $\text{CPS}_f^0([h, a, D]) \subset \bigcup \{[f([h, a, D], c)] : c \in C_f([h, a, D])\}$, $\text{cl CPS}_f^0([h, a, D]) \supset \{[f([h, a, D], c)] : c \in C_f([h, a, D])\}$ and $\text{cl CPS}_f^0([h, a, D]) = \text{cl} \{[f([h, a, D], c)] : c \in C_f([h, a, D])\} = \bigcup \text{cl} \{[f([h, a, D], c)] : c \in C_f([h, a, D])\}$.

Furthermore, for any such continuation history $[h, a, D]$, the assumptions in the Axiom imply that every $p, p' \in \text{CPS}_f^0([h, a, D])$ satisfies $p \sim_D p'$; hence, by Plan Continuity, the same holds for $p, p' \in \text{cl CPS}_f^0([h, a, D])$. Thus, $\text{CPS}_f([h, a, D]) = \text{cl CPS}_f^0([h, a, D])$.

Now fix $h \in H_t$: it will be shown that $\text{cl CPS}_f^0(h) = \text{cl Comm}_f(h)$. Consider first $p \in \text{CPS}_f^0(h)$, and let c be its initial choice. Clearly, if $c \in X$, then $c \in C_f(h)$ and so $p \in \text{Comm}_f(h)$. Thus, assume instead that $c \in A$; then, for every $D \in \mathcal{F}_p(h, c)$, $p([c, D]) \in \text{CPS}_f([h, c, D]) = \text{cl CPS}_f^0([h, c, D])$; hence, there is $\{p_{c,D}^n\} \subset \text{CPS}_f^0([h, c, D])$ such that $p_{c,D}^n \rightarrow p([c, D])$. Now define $\{p^n\} \subset F_{E(h)}$ by letting $C_{p^n}(\emptyset) = \{c\}$, $\mathcal{F}_{p^n}(\emptyset, c) = \mathcal{F}_p(\emptyset, c)$ and, for all $D \in \mathcal{F}_p(\emptyset, c)$, $p^n([c, D]) = p_{c,D}^n$. Recall that $c \in C_f(h)$ and $\mathcal{F}_p(\emptyset, c) = \mathcal{F}_f(h, c)$; also, since it was shown above that, for all $[h, c, D] \in f$, $\text{CPS}_f^0([h, c, D]) \subset \bigcup \{[f([h, c, D], b)] : b \in C_f([h, c, D])\}$, $p_{c,D}^n \approx f([h, c, D], b_{c,D})$ for some $b_{c,D} \in C_f([h, c, D])$. Therefore, Lemma 4 implies that, for every n , $p^n \approx q^n$ for some $q^n \in \text{Comm}_f(h)$; hence, $q^n \rightarrow p$, so $p \in \text{cl Comm}_f(h)$. It follows that $\text{cl CPS}_f^0(h) \subset \text{cl Comm}_f(h)$.

⁴¹Suppose not: then Condition (ii) in the axiom implies that, in particular, $f(h, b) \succ_{E(h)} p$ for any $b \in B(h)$.

Conversely, fix $p \in \text{cl Comm}_f(h)$, and let $\{p^n\} \subset \text{Comm}_f(h)$ be such that $p^n \rightarrow p$. Again let $\{c\} = C_p(\emptyset)$ and $\{c^n\} = C_{p^n}(\emptyset)$. If $c, c^n \in X$, then $c^n \in C_f(h)$, which implies that $p \in \text{cl CPS}_f^0(h)$, as required. Thus, assume that $c, c^n \in A$. Then, for all $D \in \mathcal{F}_p(\emptyset, c) = \mathcal{F}_{p^n}(\emptyset, c^n)$, $p^n([c^n, D]) \rightarrow p([c, D])$. Furthermore, for every n , $p^n([c^n, D]) = f([h, c^n, D], b^n)$ for some $b^n \in C_f([h, c^n, D])$; but it was shown above that $f([h, c^n, D], b^n) \in \text{CPS}_f([h, c^n, D])$. Therefore $p^n \in \text{CPS}_f^0(h)$, and so $p \in \text{cl CPS}_f^0(h)$.

Therefore, as claimed, for every $h \in H_t$, $\text{cl Comm}_f(h) = \text{cl CPS}_f^0(h)$. On the other hand, by assumption $\{[g(h, c)] : c \in C_g(h)\} = \{[p] : p \in \text{Comm}_f(h)\}$. In particular, this implies that $g(h, c)$ is a plan for every $c \in C_g(h)$, and so $\text{cl CPS}_g^0(h) = \text{cl } \{g(h, c) : c \in C_g(h)\}$ by (iii) in Lemma 11. But $\text{cl } \{g(h, c) : c \in C_g(h)\} = \bigcup \text{cl } \{[g(h, c)] : c \in C_g(h)\} = \bigcup \text{cl } \{[p] : p \in \text{Comm}_f(h)\} = \text{cl Comm}_f(h)$ by Remark B.2. Hence, $\text{cl CPS}_g^0(h) = \text{cl CPS}_f^0(h)$, so $\text{CPS}_g(h) = \text{CPS}_f(h)$ and hence, inductively, $\text{CPS}_g(\emptyset) = \text{CPS}_f(\emptyset)$, which implies the claim.

Axiom 3.3, Simplification. Let $f \in F_E^*$ be as in the Axiom; consider first the case $B = \{c\}$. Consider $p \in \text{CPS}_{f(\emptyset, c)}(\emptyset)$ and $q \in \text{CPS}_f(\emptyset)$. Since $\text{CPS}_f^0(\emptyset) \supset \text{CPS}_{f(\emptyset, c)}^0(\emptyset)$, the same holds for the closures of these sets, and so surely $q \succ_E p$. If the preference is strict, consider a sequence $\{q^n\} \subset \text{CPS}_f^0(\emptyset)$ such that $q^n \rightarrow q$; then, for n large, $q^n \succ_E p$ by Plan Continuity. Fix one such n and suppose for definiteness that $q^n \in \text{CPS}_{f(\emptyset, b^n)}^0(\emptyset)$. Then, a fortiori, for any $\bar{q}^n \in \text{CPS}_{f(\emptyset, b^n)}(\emptyset)$, $\bar{q}^n \succ p$. But by (ii) in Theorem 8, $f(\emptyset, b^n) \sim_E \bar{q}^n$ and $f(\emptyset, a) \sim_E p$, so $f(\emptyset, b^n) \succ_E f(\emptyset, c)$, which contradicts the assumptions of the Axiom. Thus, $p \sim_E q$, which implies that $f \sim_E f(\emptyset, c)$ as required.

For general $B \subset C_f(\emptyset)$, pick any $c \in B$: then, by the above argument, $f \sim_E f(\emptyset, c) \sim_E f(\emptyset, B)$.