Dynamic Choice under Ambiguity

Marciano Siniscalchi*

September 10, 2004

Abstract

This paper analyzes sophisticated dynamic choice for ambiguity-sensitive decision makers. It characterizes *Backward Induction* and a single-person version of *Subgame Perfection* via axioms on preferences over decision trees. Furthermore, it indicates how to elicit conditional preferences from prior preferences. The key axiom is a weakening of Dynamic Consistency, deemed *Sophistication*.

The analysis is not restricted to specific decision models and/or updating rules. Hence, the results indicate that (i) ambiguity attitudes, (ii) updating rules, and (iii) sophisticated behavior in decision trees can essentially be modeled in a mutually orthogonal fashion. As an example, a characterization of prior-by-prior Bayesian updating and Backward Induction for *arbitrary* maxmin-expected utility preferences is presented.

1 Introduction

Models of ambiguity-sensitive preferences have received considerable attention in the decision-theoretic literature; a number of promising economic applications of these models have also appeared in recent years. This paper focuses on dynamic choice in the presence of ambiguity.

Most existing contributions[1] in this area analyze prior and conditional preferences over Savage acts (maps from states to outcomes) and focus on specific models of choice, such as the "maxmin expected utility" (MEU) model of Gilboa and Schmeidler [7] or the Choquet Expected Utility (CEU) model of Schmeidler [27]. Different updating rules for ambiguity-sensitive preferences have been characterized for such decision models: for instance, see Gilboa and Schmeidler [8], Jaffray

*Economics Department, Northwestern University, Evanston, IL 60208; Email `marciano@northwestern.edu`. Web: `faculty.econ.northwestern.edu/faculty/siniscalchi`. I have greatly benefited from conversations with Peter Klibanoff. I also thank Pierpaolo Battigalli, Larry Epstein, Paolo Ghirardato, and Fabio Maccheroni for extensive discussions of this paper and/or the related working paper [29]. All errors are my own.

[1]Notable exceptions, including Cubitt [1] and Klibanoff [17], are discussed in §1.2 below.

[15], Pires [24], Hanani and Klibanoff [13], as well as Walley [31] and Shafer [28]. These updating rules typically lead to failures of dynamic consistency in the presence of ambiguity. Indeed, Epstein and Le Breton [4] show that full dynamic consistency is generally incompatible with non-neutral attitudes towards ambiguity.[2] Thus, a consistent theory of dynamic choice *based solely upon preferences over Savage acts* must necessarily entail constraints on the patterns of ambiguity-sensitive behavior it can accommodate, and/or on the class of dynamic choice problems it can deal with.

This paper differs from the received literature in three respects. First, the DM is assumed to hold preferences over *decision trees*, rather than Savage acts. Second, no specific decision model (e.g. CEU, MEU) is assumed, and no particular updating rule is adopted. Third, dynamic consistency is weakened so as to accommodate arbitrary attitudes towards ambiguity; however, the DM is assumed to be *sophisticated* enough, so that arbitrary decision trees can be "solved" by Backward Induction, and conditional preferences can be elicited from prior preferences. Sophistication can be formulated as an assumption on preferences over decision trees that does *not* restrict the DM's preferences over acts. Hence, this approach yields a theory of sophisticated dynamic choice for general decision trees that allows for arbitrary patterns of ambiguity-sensitive behavior.

Indeed, the results in this paper indicate that (i) ambiguity attitudes, (ii) updating rules, and (iii) sophisticated choice in decision trees are essentially *orthogonal* aspects of behavior, and hence can be modeled independently. In particular, the analysis in this paper complements existing work on updating rules for CEU, MEU and related decision models. To summarize the main findings:

- Theorem 4 in Section 3.2 characterizes *Backward Induction* for decision trees; Theorem 5 in Section 3.3 characterizes a related notion of single-person *Subgame Perfection*. These are the main contributions of this paper. The key consistency axiom is deemed *Sophistication*; no structural property of preferences, besides completeness and transitivity, is required.

- Section 3.1 provides axioms on conditional preferences that make it possible to *elicit* the latter from prior preferences, and conversely shows how to *define* conditional preferences from prior preferences in such a way as to satisfy those axioms (Proposition 3).

- As an application of the approach described in this paper, Section 4 characterizes "full Bayesian updating" (Proposition 6) and Backward Induction (Theorem 7) for *arbitrary* MEU preferences. In particular, no restriction need be imposed on the set of priors.

The remainder of this Introduction expands upon these remarks, and provides examples; the related literature is further discussed in §1.2. Section 2 describes the formal framework, and in particular defines decision trees. Section 3 contains the main results, and Section 4 provides an application to MEU preferences. Finally, Section 5 discusses possible extensions.

[2] See also the discussion of Epstein and Schneider [5] and Klibanoff [17] in §1.2 below.

1.1 Motivating Examples

1.1.1 Background: A "folk theorem"

It is useful to begin by briefly recalling a well-known result that relates properties of conditional and unconditional preferences. The decision setting is also standard, so a brief presentation will suffice. Denote the state and prize spaces by Ω and X respectively. *Savage acts* are maps from Ω to X; for any pair of Savage acts f, g, and any event E, let fEg denote the Savage act defined by $fEg(\omega) = f(\omega)$ for $\omega \in E$, and $fEg(\omega) = g(\omega)$ for $\omega \notin E$.

Consider two binary relations $\succcurlyeq, \succcurlyeq_E$ on the set of Savage acts; $\succcurlyeq$ is the DM's *prior* preferences, whereas $\succcurlyeq_E$ represents her preferences *conditional upon* E. Say that $\succcurlyeq_E$ satisfies *Consequentialism* if $f \sim_E fEg$ for all Savage acts f, g. Say that $\succcurlyeq, \succcurlyeq_E$ jointly satisfy *Dynamic Consistency* if $f \succcurlyeq_E g \Leftrightarrow fEg \succcurlyeq g$ (cf. Ghirardato [6]). An event E is *null* if $fEg \sim g$ for all Savage acts f, g. Finally, $\succcurlyeq$ satisfies Savage's *Postulate P2* on E if $fEh \succcurlyeq gEh \Leftrightarrow fEh' \succcurlyeq gEh'$ for all Savage acts f, g, h, h'. The interpretation of these definitions and axioms is well-known; in particular, the version of Dynamic Consistency adopted here states that f is a (weakly) "profitable deviation" from g conditional upon E, if and only if, a priori, the DM would also (weakly) prefer to modify her "plan" g so as to obtain the same outcomes as f at states in E.

Theorem 1 (The "folk theorem" of dynamic choice under uncertainty) [3] *Assume that $\succcurlyeq$ is complete and transitive and E is not null. Then $\succcurlyeq_E$ is complete and transitive, it satisfies Consequentialism, and $\succcurlyeq, \succcurlyeq_E$ jointly satisfy Dynamic Consistency* if and only if *$\succcurlyeq$ satisfies P2, and furthermore, for all Savage acts f, g, h,*

$$f \succcurlyeq_E g \quad \Leftrightarrow \quad fEh \succcurlyeq gEh. \tag{1}$$

It is useful to emphasize two notable and well-understood consequences of this "folk theorem". First, given the DM's prior preferences $\succcurlyeq$, the "Bayesian updating rule" in Eq. 1 is the only possible *definition* of conditional preferences $\succcurlyeq_E$ that satisfies Consequentialism and Dynamic Consistency; conversely, if the latter axioms hold, then $\succcurlyeq_E$ can be *elicited* from $\succcurlyeq$ via Eq. 1.

Second, Eq. 1, and hence Dynamic Consistency, is the basis for the *Backward Induction* analysis of decision trees; the following subsection elaborates this point.

---

[3] One exposition of this result can be found in the appendix of Ghirardato [6] (who also refers to it as a "folk theorem"). Epstein and Le Breton [4] show that a similar result holds even if Consequentialism is dropped. For a related result in the setting of dynamic choice under risk, see Karni and Schmeidler [16]. It should also be noted that this equivalence is implicit in Savage's justification of his Postulate P2 (cf. [26], pp. 21–23).

1.1.2 A dynamic version of the Ellsberg Paradox

Consider the dynamic decision problem in Figure 1. As in the celebrated three-color urn example by Daniel Ellsberg [2], a DM is presented with an urn containing 90 balls, of which 30 are red and 60 are either green or blue, in unspecified proportions. There are two prizes available, \$0 and \$10; x denotes an arbitrary prize, so $x \in \{0, 10\}$. The state space is $\Omega = \{r, g, b\}$, in obvious notation. For the purposes of this Introduction, it is not necessary to specify a representation for (or otherwise restrict) the DM's prior preferences.

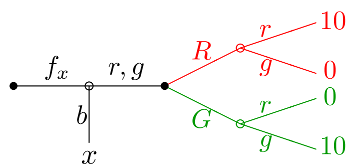

Figure 1: Deferred Choice in the Ellsberg model; $x \in \{0, 10\}$.

The DM is first informed whether the ball drawn from the urn is blue, in which case she receives x. If the ball drawn is not blue, the she can choose whether to bet on red or green. The DM is then informed of the outcome of the draw, and receives the appropriate prize.

A few remarks on graphical conventions are in order. Filled dots ($\bullet$) represent decision nodes, whereas empty circles ($\circ$) correspond to points in the tree where information is revealed to the DM; it is often convenient, if imprecise, to refer to these as "chance nodes". Trees are drawn from left to right, and from top to bottom; edges departing from decision nodes are decorated with action labels (f_x, R and G in Fig. 1), whereas edges departing from chance nodes are labelled with events ($\{r, g\}$, $\{b\}$, $\{r\}$ and $\{g\}$ in Fig. 1).

The decision tree in Fig. 1 will be denoted f_x. Observe that each of the subtrees beginning with the second decision point and continuing with either R or G can be viewed as a *conditional* decision tree in its own right. For simplicity, in this section only, these subtrees (drawn in red and green colors respectively in Fig. 1) will also be denoted by R and G respectively; a formal notation for subtrees will be introduced in Section 2.

Now let $E = \{r, g\}$ and consider a DM endowed with prior and conditional preferences $\succcurlyeq, \succcurlyeq_E$ over Savage acts. If Consequentialism and Dynamic Consistency hold, this DM can "evaluate" the tree f_x in Fig. 1 by Backward Induction. It is useful to spell out the (admittedly tedious) details to highlight the crucial role of Dynamic Consistency and related axioms in the development of a coherent theory of dynamic choice built solely upon the DM's "static" preferences over Savage acts. The analysis will also clarify why Dynamic Consistency necessarily conflicts with ambiguity-sensitive preferences, and thus provide a rationale for the central notion of Sophistication.

By Consequentialism, the DM's conditional preferences $\succcurlyeq_E$ over Savage acts have a straightforward extension to the (trivial) subtrees R and G. Moreover, a priori, the decision tree f_x may be seen as providing a menu of two *plans*: "choose R at the second decision node" (denoted f_x^R) and "choose G at the second decision node" (denoted f_x^G). These plans are depicted in Fig. 2; formally, a plan is a decision tree wherein a single action is available at every decision node.

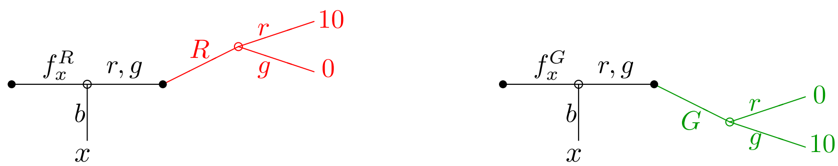

Figure 2: The plans f_x^R and f_x^G; $x \in \{0, 10\}$.

Recall that the DM's preferences $\succcurlyeq, \succcurlyeq_E$ are formally defined over Savage acts. However, under the additional, standard assumption (hereinafter "Reduction") that the DM is indifferent between a plan and the Savage act that can be derived from it in the obvious way,[4] her preferences can be extended to plans, including f_x^R and f_x^G.

The key observation is now that Eq. 1 in the "folk theorem" of §1.1.1 relates the DM's prior preferences over f_x^R and f_x^G with her conditional preferences over R and G:

$$\forall x \in \{0, 10\}, \qquad R \succcurlyeq_E G \quad \Leftrightarrow \quad f_x^R \succcurlyeq f_x^G. \tag{2}$$

As noted above, f_x can be viewed as a menu consisting of the plans f_x^R and f_x^G. Eq. 2 then implies that the search for an a priori optimal plan in this menu can be restricted to plans that are conditionally optimal: this justifies "solving" the tree f_x by *Backward Induction*. Moreover, Eq. 2 also rules out *commitment* problems: if the DM's a priori optimal plan is f_x^R, she will actually wish to carry out this plan at her second decision node.

Thus, Consequentialism and Dynamic Consistency, jointly with Reduction, provide a complete and coherent theory of dynamic choice. However, the preceding analysis also readily implies that, under these assumptions, by Eq. 2,

$$f_0^R \succcurlyeq f_0^G \quad \Leftrightarrow \quad f_{10}^R \succcurlyeq f_{10}^G:$$

Again assuming Reduction, this is easily seen to rule out the modal preferences in the Ellsberg paradox: f_0^R and f_0^G are (compound) bets on red and green, whereas f_{10}^R and f_{10}^G are bets on "red or blue" and "green or blue"—the four bets in the static version of the paradox.

[4]For instance, f_x^R corresponds to the Savage act that yields x if b obtains, 10 if r obtains, and 0 if g obtains.

Thus, a DM cannot satisfy Reduction, Consequentialism and Dynamic Consistency, *and* at the same time exhibit Ellsberg-type behavior.

1.1.3 Sophistication

The approach proposed in this paper involves relaxing Dynamic Consistency in a way that still *supports backward-induction reasoning in decision trees, but does not constrain prior preferences over Savage acts.* This makes it possible to accommodate Ellsberg-type behavior, while retaining Reduction and Consequentialism.[5]

To formalize the appropriate generalization of Dynamic Consistency, the DM is assumed to hold well-defined preferences over decision trees, not just Savage acts. This is in the spirit of the literature on menu choice originating from Kreps [19].

The central axiom of this paper is deemed *Sophistication*: loosely speaking, it requires that the DM *hold correct expectations regarding her future choices.* In the decision tree f_x under consideration, Sophistication yields the following restriction:

$$\forall x \in \{0, 10\}, \qquad R \succ_E G \quad \Rightarrow \quad f_x \sim f_x^R \tag{3}$$

(there is an analogous restriction for the case $G \succ_E R$). 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.

Notice that, unlike Eq. 2, the above assumption does not impose any restriction on the relative ranking of f_x^R and f_x^G; thus, a DM can satisfy (Reduction, Consequentialism and) Sophistication while at the same time exhibiting the modal preferences in the Ellsberg paradox. Yet, in this example, Sophistication is sufficient to "evaluate" the tree f_x in terms of plans (i.e. find one or more plans equivalent to it), and therefore (if Reduction is also assumed) in terms of simple Savage acts. The main result of this paper, Theorem 4, shows that, *for any collection of complete and transitive prior and conditional preferences, Sophistication and a related tie-breaking axiom discussed below hold if and only if decision trees are evaluated according to Backward Induction.*

The above discussion assumes that the DM's conditional preferences are given. However, it turns out that a weak form of Sophistication, together with basic structural assumptions that guarantee the existence of certainty equivalents, is sufficient to elicit or define conditional preferences from prior preferences: see Proposition 3 in Sec. 3.1.

Sophistication may be thought of as an assumption about the DM's prior beliefs concerning the behavior of her future selves: specifically, it reflects one implication of the assumption that such beliefs are *consistent with "rationality".* In this respect, although it is a fully behavioral

[5] As will be made clear, whether or not Reduction holds is *irrelevant* for the main results of this paper. Reduction is only assumed in the application to MEU preferences (cf. Sec. 4). As for Consequentialism, please refer to Sec. 5.

axiom, Sophistication is reminiscent of epistemic assumptions in the literature on the foundations of game-theoretic solution concepts.

1.1.4 Weak Commitment and Consistent Planning

Interpreting Sophistication as an assumption about beliefs also clarifies the limits to the kind of restrictions that it can generate. Consider the decision tree in Fig. 3, denoted f.

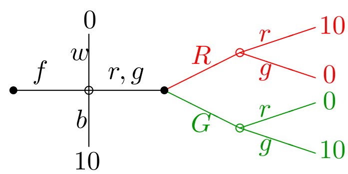

Figure 3: Sophistication and Weak Commitment

The state space comprises four states, denoted r, g, b, w; suppose that the DM has MEU preferences, with priors $C = \{q : q(r) = q(b),\, q(g) = q(w)\}$. Also let $E = \{r, g\}$ and suppose that her posterior preferences $\succcurlyeq_E$ are determined by full Bayesian updating; the set of posteriors is then $C_E = \{q : q(r) = 1 - q(g)\}$.

Notice that $R \sim_E G$; moreover, observe that the DM strictly prefers the Savage act that yields 10 if g or b obtain, and 0 otherwise, to the Savage act that yields 10 if r or b obtain, and 0 otherwise. In other words, a priori the DM would like to "commit" to choosing G at her second decision node.

Intuitively, this is consistent with the epistemic interpretation of Sophistication. The DM may equally "rationally" choose R or G upon reaching her second node; therefore, the (prior) belief that G will be chosen at the second node is fully consistent with rationality. Hence, Sophistication should allow for preferences consistent with the above desire for commitment—and indeed it does.

However, this example also illustrates that Sophistication alone may not be sufficient to pin down the DM's behavior in all decision trees: it must be complemented by some tie-breaking assumption. One possible approach, which will be fully developed in Section 3.2, is to assume that the DM is able to commit to specific future choices, if her future self does not have opposite strict preferences. This assumption, referrred to as *Weak Commitment*, formalizes the notion of *consistent planning* first proposed by Strotz [30]:

> [The DM's] problem is then to find the best plan among those that [s]he will actually follow. (Strotz [30], p. 173).

In the present setting, consider the trees f^R, f^G obtained from f by pruning the green and, respectively, the red subtrees. As in the dynamic Ellsberg model, these pruned acts reflect the possibility of commitment. Weak Commitment implies the following restriction:

$$G \sim_E R, \quad f^G \succ f^R \quad \Rightarrow \quad f \sim f^G$$

(and similarly if $f^R \succ f^G$). 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 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 Savage (or compound) acts is required; the "epistemic" character of this behavioral assumption should also be apparent.

Notice that different beliefs of the DM about her behavior at the second decision node in the tree f can have behavioral implications. In particular, if the DM believes that she will choose R instead of G, then a priori she will be indifferent between f and the pruned tree f^R (and strictly prefer f^G to either of these).

This suggests that, even if Weak Commitment is not imposed, it may nevertheless be possible to elicit the DM's beliefs about her future selves' choices from her prior preferences. This idea is developed in Section 3.3, and shown to characterize a notion of *Subgame Perfection*.

1.1.5 Backward Induction vs. "recursion"

The Backward Induction algorithm characterized by Sophistication and Weak Commitment involves the iterative removal of conditionally dominated actions. The algorithm does *not* involve replacing subtrees with the conditional certainty equivalent of the optimal continuation plans. In terms of the numerical representation of preferences, this implies that Sophistication and Weak Commitment do not characterize *recursion*, or "value substitution".

To see this, return to the tree f_x in Fig. 1. Suppose that the DM is endowed with MEU prior preferences, characterized by the set of priors $C = \{q : q(r) = \frac{1}{3}\}$; suppose further that her conditional preferences are derived from these by full Bayesian updating, so they are MEU with posteriors $C_E = \{q : q(r) = 1 - q(g) \geq \frac{1}{3}\}$. To ensure the existence of certainty equivalents, assume that the prize set is $X = [0, 10]$ and, for simplicity, that utility is linear.

It is then easy to verify that $R \succ_E G$, so Sophistication implies that $f_0 \sim f_0^R$. However, note that the conditional certainty equivalent of R is $\frac{10}{3}$; the prior MEU evaluation of the Savage act f' that yields $\frac{10}{3}$ in states r and g, and 0 otherwise, is $\frac{10}{9}$, whereas the prior MEU evaluation of f_0^R is $\frac{10}{3}$. Thus, the DM is *not* indifferent between f_0 and f'.

For value substitution to be legitimate, the full force of Dynamic Consistency is required. Very informally, Backward Induction (hence, Sophistication) enables the DM to evaluate her own future choices "from today's perspective" (i.e. in terms of her current preferences); value substitution, on the other hand, is only meaningful if today's and tomorrow's perspectives coincide.

1.2 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 [21] 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.

Epstein and Schneider [5] characterize recursive MEU preferences over Savage-style acts. In terms of the axioms in the "folk theorem" of Subsection 1.1.1, 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 [32].

Consistently with the preceding discussion, this results in a restriction on prior preferences; in particular, the "folk theorem" in Subsection 1.1.1 implies that, under the assumptions in [5], Savage's Postulate P2 and Eq. 1 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 [5] rule out Ellberg-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 [17] drops Reduction, and also introduces a form of state-dependence of preferences over prizes. More recently, Hanani and Klibanoff [13] have proposed a family of updating rules for MEU preferences; their analysis drops Consequentialism (but implicitly maintains Reduction).

This paper should be viewed as complementary to these contributions. Incidentally, the main results of this paper (i.e. Proposition 3 and Theorems 4 and 5) do *not* assume Reduction; the latter is only used in the application to conditional MEU preferences. Consequentialism is further discussed in Section 5. It should also be noted that the present paper does not restrict attention to MEU preferences, or to specific updating rules. Indeed, as was mentioned in the beginning of this Introduction, the results in this paper indicate that the analysis of *updating* is essentially orthogonal to the characterization of *sophisticated behavior* in decision tres.

Assuming that preferences are defined over decision trees is much in the same spirit as Kreps's seminal contribution on menu choice (Kreps [19]) and the temporal resolution of uncertainty (Kreps [20]). Gul and Pesendorfer [10] 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. A menu choice framework has been recently adopted by Epstein [3] to study "non-Bayesian" updating. Klibanoff and Ozdenoren [18] characterize a subjective version of Kreps's [20] model of recursive expected utility; they do not consider ambiguity-sensitive preferences, but their decision setting is related to the one adopted here.

Hammond [11, 12] and Cubitt [1] provide an analysis of Consequentialism and Dynamic Consistency in decision trees for EU preferences; also see the references in Footnote 3. Sarin and Wakker [25] consider non-EU behavior in decision trees, and in particular investigate the consequences of the assumption that prior and conditional preferences belong ot the same class of models (e.g. MEU, CEU, etc.).

2 Decision Setting

2.1 Histories and Trees

The notation employed in this paper adapts the notion of "perfect-information game tree" in Osborne-Rubinstein [22]. The basic building block in the description of a decision tree 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; chance moves represent *information* that the DM may receive. A decision tree can then be represented by a suitable collection of histories, together with an assignment of prizes to terminal (i.e. complete) histories.

This paper focuses on finite trees; see Sections 3.3 and 5 for remarks on extensions to infinite trees. Mainly for notational simplicity, the state space is also assumed to be finite (the extension to finite decision trees with an arbitrary underlying state space is trivial).

Formally, fix a finite set Ω of states, a countable collection A of action labels, and a collection X of prizes; assume the latter is a connected separable topological space. A **history of length** $T \geq 0$ **starting at** $E \subset \Omega$ is a sequence of ordered pairs, denoted

$$h = [(a_1, E_1), \ldots, (a_T, E_T)],$$

such that, for all $t = 1, \ldots, T$, $a_t \in A$, $E_t \subset E$. Denote the set of all histories starting at E by $\mathcal{H}_E$. Throughout this note, it is notationally convenient to also consider the "empty history" $h = \emptyset$. If h has length T and $1 \leq t \leq T$, then h_t denotes the history consisting of the first

t elements of h; furthermore, $h_0 = \emptyset$ for every history h. Similarly, $h_{t_1:t_2}$ denotes the history $[(a_{t_1}, E_{t_1}), (a_{t_1+1}, E_{t_1+1}), \ldots, (a_{t_2}, E_{t_2})]$.

The **length** of history h is denoted by $\lambda(h)$; also, $\lambda(\emptyset) = 0$. The action and event that appear in the last ordered pair of a history h (i.e. a_T and E_T above) are denoted by $a(h)$ and $E(h)$ respectively; it is also convenient to let $E(\emptyset) = E$ for every history h starting at E. Compositions of histories are denoted in an obvious way as follows: $[h, h']$, $[h, (a, E)]$; in particular, $[\emptyset, h] = [h, \emptyset] = h$. Finally, a history h is a **subhistory** of another history h', written $h \leq h'$, if $h = h'_t$ for some $t \in \{0, \ldots, \lambda(h')\}$; it is a **strict subhistory** of h' if $t < \lambda(h')$.

Definition 1 A **decision tree starting at** E is a tuple $f = (E, H, x)$, where $E \subset \Omega$, H is a finite collection of histories starting at E such that

1. if $h \in H$ and $\lambda(h) > 0$, then $h_{\lambda(h)-1} \in H$;
2. for every $h, h' \in H$ such that $\lambda(h), \lambda(h') \geq 1$, $a(h_1) = a(h'_1)$;
3. for every $h \in H$ and every $a \in A$ such that $[h, (a, F)] \in H$ for some $F \subset E$, the collection $\{F : [h, (a, F)] \in H\}$ is a non-trivial partition of $E(h)$;
4. for every $h \in H$, $a, a' \in A$ and $F \subset E$: if $\{h' \in \mathcal{H}_F : [h, (a, F), h'] \in H\} \neq \{h' \in \mathcal{H}_F : [h, (a', F), h'] \in H\}$, then $a \neq a'$;
5. for every $\omega \in \Omega$, there exists $h \in H$ such that $E(h) = \{\omega\}$;

and $x : \{h : |E(h)| = 1\} \to X$. Histories h such that $|E(h)| = 1$ are called **terminal**; all other histories are **non-terminal**. Denote the set of all decision trees starting at E by F_E. For every non-terminal history $h \in H$, the set of **actions available at** h is $A_f(h) = \{a \in A : \exists F \subset E, [h, (a, F)] \in H\}$; for every $a \in A_f(h)$, the **information partition** following h and a is $\mathcal{F}_f(h, a) = \{F \subset E : [h, (a, F)] \in H\}$.[6]

By Condition 1, if a history h can occur in a tree, then all its subhistories can also occur; in particular, the empty history can occur ($\emptyset \in H$). By Condition 2, in every decision tree, there is a unique action available at the initial node. This turns out to be convenient for the purposes of "tree surgery", but is otherwise inconsequential. In particular, in order to model the DM's behavior in a situation where she has more than one initial action available, each corresonding to a decision trees, it is clearly enough to consider the DM's preferences over the corresponding collection of trees.

[6]The elements of $\mathcal{F}_f(h, a)$ can informally be thought of as "chance moves" following h and a.

Condition 3 implies that events in a history form a decreasing subsequence that, by 5, ends with a singleton. Finally, Condition 4 ensures that no two distinct actions available at a history h have the same label.[7]

To clarify the notation, the tree f_x in Fig. 1 can be described as follows: $\Omega = \{r, g, b\}$, $X = \{0, 10\}$, $A = \{*, R, G\}$; throughout this paper, unless otherwise noted, the dummy action at the initial node will be denoted by the label "$*$". Then, $H = \big\{\emptyset$, $[(*, \{r, g\}]$, $[(*, \{b\}]$, $[(*, \{r, g\})$, $(R, \{r\})]$, $[(*, \{r, g\})$, $(R, \{g\})]$, $[(*, \{r, g\})$, $(G, \{r\})]$, $[(*, \{r, g\})$, $(G, \{g\})]\big\}$. The only non-terminal histories are $\emptyset$ and $h = [(*, \{r, g\}]$; finally, $A_{f_x}(\emptyset) = \{*\}$, $A_{f_x}(h) = \{R, G\}$, $\mathcal{F}_{f_x}(\emptyset) = \{\{r, g\}, \{b\}\}$, and $\mathcal{F}_{f_x}(h, R) = \mathcal{F}_{f_x}(h, G) = \{\{r\}, \{g\}\}$.

Certain distinguished classes of decision trees deserve a closer look. First, if $H = \{[(*, \{\omega\})] : \omega \in E\}$, where "$*$"$\in A$, then $f = (E, H, x)$ corresponds to the **Savage act** conditional on E that yields the outcome $x(\omega)$ in state ω; in particular, if $x(\cdot)$ is constant and equal to some $\bar{x} \in X$, f is a constant act, and will be denoted by $\bar{x}$.

As noted in the Introduction, **plans** are decision trees wherein, for any non-terminal history h of f, there is a unique $a \in A$ such that $[h, (a, E)]$ for some $E \subset \Omega$. Notice that a Savage act, and in particular a constant act, is necessarily a plan, but the converse is not true. On the other hand, every plan "induces" a Savage act: see Section 4 for details.

For a more interesting example, fix a sequence $\mathcal{F}_1, \ldots, \mathcal{F}_T$ of progressively finer partitions of Ω, such that $\mathcal{F}_T = \{\{\omega\} : \omega \in \Omega\}$. A **partitional tree** is a tree $f = (\Omega, H, x)$ such that all terminal histories have length T, and furthermore $\mathcal{F}_f(h, a) = \{F \in \mathcal{F}_{\lambda(h)+1} : F \subset E(h)\}$ for all non-terminal history h and action $a \in A_f(h)$. Thus, in a partitional tree, Nature's moves are independent of the DM's choices. The notation adopted here 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.

2.2 Tree Surgery

First, consider a tree $f = (E, H, x) \in F_E$ and a non-terminal history $h \in H$. Then any action $a \in A(h)$ identifies a **continuation tree**, denoted $f(h, a)$ and defined by the tuple $(E', H', x') \in F_{E'}$, where:

- $E' = E(h)$;
- $H' = \{h' \in \mathcal{H}_E : [h, h'] \in H \text{ and } a(h'_1) = a\}$;
- $x'(z') = x'([h, z'])$ for all z' such that $[h, z']$ is terminal.

[7] Regardless of their label, two actions are distinct if they yield different sets of possible continuation histories for some realization of the uncertainty. In this case, Condition 4 ensures that distinct actions are labelled differently.

That is: every continuation h' of h that begins with the choice of a becomes a history in the continuation tree. Notice that, according to this definition, for $h = \emptyset$, $f(h,a) = f$.

Second, consider $f = (E, H, x) \in F_E$, $\{f_1, \ldots, f_n\} \subset F_{E'}$ such that $f_i = (E', H_i, x_i)$ for all $i = 1, \ldots, n$, and $h \in H$ such that $E(h) = E'$. It is often useful to consider a tree that differs from f in that the continuation trees available at h are replaced with the trees $f_1, \ldots, f_n$. Formally, the **replacement tree** $\{f_1, \ldots, f_n\}_h f$ is the element of F_E defined by the tuple $(E, \bar{H}, \bar{x})$, where:

- $a_1^*, \ldots, a_n^* \in A$ are n distinct action labels;
- $\bar{H} = \{h' \in H \,:\, h \not< h'\} \cup \bigcup_i \{[h, (a_i^*, E(h_1')), h'_{2:\lambda(h')}] \,:\, h' \in H_i\}$;
- $\bar{x}(z) = x(z)$ if z is a terminal history for f and $h \not< z$; $\bar{x}([h, (a_i^*, E(z_1')), z'_{2:\lambda(z')}]) = x_i(z')$ if z' is a terminal history for f_i.

The key point is that histories in the replacement tree consist of histories from f that do not strictly follow h, plus histories that begin with h and continue with some history from one of the trees f_i. Since two or more trees in the collection $f_1, \ldots, f_n$ may have the same initial action label, continuation histories are modified so as to ensure that a distinct label a_i^* is used for each f_i. Notice also that, for $h = \emptyset$, $\{f_1, \ldots, f_n\}_h f = f$.

3 Axioms and Results

Throughout this section, the main object of interest is a collection $\{\succcurlyeq_E\}_{\neq E \subset \Omega}$ of binary relations; in particular, for all non-empty $E \subset \Omega$, $\succcurlyeq_E$ is a preference relation on F_E. For notational simplicity, $\succcurlyeq_\Omega$ will be denoted by $\succcurlyeq$.

Notice that each conditional preference $\succcurlyeq_E$ is defined over the set of decision trees starting at E. Thus, implicitly, ***Consequentialism*** *is assumed throughout*. While this can be formalized explicitly in simple decision trees (e.g. those arising from filtrations), doing so in the full generality of the present framework is notationally demanding.[8] Section 5 considers possible extensions of the results to non-Consequentialist preferences.

Subsection 3.1 shows that, under suitable axioms, conditional preferences can be either elicited or defined from unconditional ones; this result clarifies the behavioral significance of the analysis carried out in the remainder of this paper.

[8] In non-partitional trees, different histories h, h' may correspond to the same event—say, $E(h) = E(h') = E$. Suppose the trees f, g both contain h and h'; then the notation $f \succcurlyeq_E g$ is ambiguous.

Subsection 3.2 formalizes the key notions of sophistication and weak commitment, and establishes the main result of this section, the promised characterization of Backward Induction for decision trees.

The material in Subsections 3.2 and 3.1 is essentially orthogonal; readers who are more interested in the analysis of Backward Induction can safely read §3.2 first and §3.1 later.

Subsection 3.3 provides an alternative to the weak commitment assumption, and characterizes a single-person version of Subgame perfection. This is related to the "multi-selves" (or game-theoretic) approach to dynamic decision making under certainty in the absence of full commitment.

All proofs are in the Appendix.

3.1 Eliciting Conditional Preferences

Recall that, as discussed in the Introduction, the "folk theorem" for dynamic decision making shows that, if Consequentialism and Dynamic Consistency hold, then conditional preferences can be *elicited* from prior preferences via Eq. 1, the Bayesian updating relation; conversely, if one wishes to *define* conditional preferences from prior preferences, Bayesian updating (together with Savage's postulate P2) is necessary to ensure that Consequentialism and Dynamic Consistency hold.

This subsection establishes a similar result in the setting of choice among trees. The key ingredients of the approach described in this section are *conditional certainty equivalents* and a weak form of *sophistication*. It should be emphasized that, while notions of sophistication play a key role throughout this paper, certainty equivalents merely provide one convenient way to elicit preferences—but not necessarily the only one. Alternative approaches may be viable, e.g. in a setting where the state space is "rich".

3.1.1 Axioms on Conditional Preferences

To aid intuition, I first assume that the DM is characterized by a preference system $\{\succcurlyeq_E\}_{\emptyset \neq E \subset \Omega}$, and consider axioms regarding these relations. First, assume that preferences over prizes (constant trees) are unaffected by conditioning.[9]

Axiom 3.1 (Stable Tastes) *For all $x, x' \in X$, and all non-empty $E \subset \Omega$: $x \succcurlyeq_E x'$ if and only if $x \succcurlyeq x'$.*

Next, one must ensure that conditional certainty equivalents exist. Since X is assumed to be a connected and separable topological space, standard dominance (or monotonicity) and continuity requirements suffice.

[9] For 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 3.1 to constant streams.

Axiom 3.2 (Conditional Dominance) *For all non-empty* $E \subset \Omega$, $f = (E, H, x) \in F_E$, *and all* $x', x'' \in X$*: if* $x' \succcurlyeq x(z) \succcurlyeq x''$ *for all terminal histories* z, *then* $x' \succcurlyeq_E f \succcurlyeq_E x''$.

Axiom 3.3 (Conditional Prize-Act Continuity) *For all non-empty* $E \subset \Omega$ *and all* $f \in F_E$, *the sets* $\{x \in X : x \succcurlyeq_E f\}$ *and* $\{x \in X : x \preccurlyeq_E f\}$ *are closed in* X.

Remark 3.1 Consider a non-empty event $E \subset \Omega$ and a complete and transitive relation $\succcurlyeq_E$ on F_E that satisfies Axioms 3.1, 3.2 and 3.3. Then, for all $f \in F_E$, there exists $x \in X$ such that $x \sim_E f$.

Finally, conditional and unconditional preferences over trees must be axiomatically related: this is the objective of the following weak sophistication requirement. Suppose that, upon reaching the partial history $h = [(*, E)]$, the DM faces a choice between the tree $f \in F_E$ and the prize x. If x is strictly preferred to f conditional on E, then, a priori, the DM should recognize that the alternative f is actually irrelevant: she will not choose it if the history h is reached. Similar considerations hold in case f is strictly preferred to x conditional on E. This motivates the following axiom.

Axiom 3.4 (Weak Sophistication) *For all* $E \subset \Omega$, *all* $g = (\Omega, H, x) \in F_\Omega$ *such that* $h = [(*, E)] \in H$, *all* $f \in F_E$, *and all* $\bar{x} \in X$*:*

(i) if $\bar{x} \succ_E f$, *then* $\{f, \bar{x}\}_h g \sim \{\bar{x}\}_h g$*; and*

(ii) if $\bar{x} \prec_E f$, *then* $\{f, \bar{x}\}_h g \sim \{f\}_h g$.

Notice that the stronger sophistication axiom considered in Section 3.2 implies Axiom 3.4. On the other hand, Axiom 3.4 is not sufficient to yield backward-induction decisions, except in very simple trees: it is "just enough" to ensure that conditional preferences can be retrieved from unconditional ones. Thus, the present approach makes it possible to address the distinct issues of elicitation and sophistication in a relatively independent way.

3.1.2 Axioms on Unconditional Preferences

Consistently with the introductory discussion, this subsection assumes that only the DM's unconditional preferences are given, and proposes axiomatic restrictions that enable one to elicit or define conditional preferences. Begin with a simple continuity requirement.

Axiom 3.5 (Prize Continuity) *For all* $\bar{x} \in X$, *the sets* $\{x \in X : x \succcurlyeq \bar{x}\}$ *and* $\{x \in X : x \preccurlyeq \bar{x}\}$ *are closed in* X.

Next, say that an event $E \subset \Omega$ is **immediately relevant** for the act $g \in F_\Omega$ if $[(*, E)]$ is a history of g; say that E is **$\succcurlyeq$-essential** if, whenever $h = [(*, E)]$ is a history of $g \in F_\Omega$, and $x, x' \in X$ satisfy $x \succ x'$, $\{x\}_h g \succ \{x'\}_h g$. Denote the set of acts $g \in F_\Omega$ for which the non-empty event $E \subset \Omega$ is immediately relevant by $F_\Omega(E)$; denote the set of $\succcurlyeq$-essential events by $\mathrm{E}(\succcurlyeq)$. Observe that, if $g \in F_\Omega(E)$, then clearly $M_h g \in F_\Omega(E)$ for any finite set $M \subset F_E$.

It is now possible to state the key assumption on unconditional preferences. Consider the following situation: conditional upon the ($\succcurlyeq$-essential) event E, the certainty equivalent of the continuation tree $f \in F_E$ is $\bar{x}_{f|E}$; notice that, if conditional preferences are not directly observable, it is not possible to directly determine $\bar{x}_{f|E}$, so a restriction on *un*conditional preferences involving it would not be verifiable (i.e. "fully behavioral"). Thus, it is necessary to work around this observability restriction.

Consider an *arbitrary* prize $\bar{x}$. Suppose first that $\bar{x} \succcurlyeq \bar{x}_{f|E}$, and consider a prize $x \succ \bar{x}$ and a tree wherein, if the history $h = [(*, E)]$ is reached, the DM faces a choice between f and x. *If* her ranking of prizes is not altered by conditioning, then it is also the case that $x \succ_E \bar{x} \succcurlyeq_E \bar{x}_{f|E} \sim_E f$; thus, $x \succ_E f$. Furthermore, *if* the DM correctly anticipates her future preferences (or at least her ranking of prizes vs. trees), she should be able to conclude that, upon reaching h, she will choose x rather than f. Hence, a priori, she should be indifferent between the tree under consideration, and a modified tree wherein the choice f is removed at h. Notice that this conclusion must hold for *any* prize $x \succ \bar{x}$, and *any* tree wherein the choices available at h are f and x.

Now suppose instead that $\bar{x} \preccurlyeq x_{f|E}$, and consider an arbitrary prize $x \prec \bar{x}$. By analogous considerations, one can conclude that the DM should be indifferent between a tree wherein the choices available at h are f and x, and a modified tree wherein the choice x is removed. Again, this must be the case regardless of what the tree specifies at histories that do not follow h (i.e. "outside the event E"), and for any prize $x \prec \bar{x}$.

These arguments identify two alternatives, at least one of which *must* hold for any $\succcurlyeq$-essential event E, act f, and prize $\bar{x}$; neither alternative involves the (possibly unobservable) conditional certainty equivalent $\bar{x}_{f|E}$. Formally:

Axiom 3.6 (Conjectural Separability) *Consider $E \in \mathrm{E}(\succcurlyeq)$, $f \in F_E$, and $\bar{x} \in X$; let $h = [(*, E)]$. Then one (or both) of the following statements hold:*

(i) for all $x \in X$ such that $x \succ \bar{x}$, and all $g \in F_\Omega(E)$, $\{f, x\}_h g \sim \{x\}_h g$;

(ii) for all $x \in X$ such that $x \prec \bar{x}$, and all $g \in F_\Omega(E)$, $\{f, x\}_h g \sim \{f\}_h g$.

The reference to *separability* is justified by noting that, in accordance with the preceding intuitive discussion, both statements (i) and (ii) in Axiom 3.6 impose a restriction on unconditional preferences that must hold regardless of choices available at histories that do not follow h (in the axiom, such choices are specified by the tree g). In this respect, Axiom 3.6 can be viewed as a counterpart to Savage's Postulate P2 in the present setting.

The third and final axiom on unconditional preferences serves as a counterpart to Axiom 3.2. Suppose that the prize $\bar{x}$ is strictly preferred to any prize that may be obtained if the continuation tree $f \in F_E$ is chosen; then, when the DM contemplates a tree that provides a choice between f and $\bar{x}$ conditional upon the (essential and immediately relevant) event E, she should deem f an irrelevant alternative. Similar cosiderations hold if $\bar{x}$ is strictly inferior to any prize that may be obtained in the continuation tree f.

Axiom 3.7 (Conjectural Dominance) *Consider $E \in \mathrm{E}(\succcurlyeq)$, $g \in F_\Omega(E)$, $f \in F_E$, and $\bar{x} \in X$; let $h = [(*, E)]$. Then*

(i) if $x(z) \succ \bar{x}$ for all terminal histories z of f, then $\{f, \bar{x}\}_h g \sim \{f\}_h g$;

(ii) if $x(z) \prec \bar{x}$ for all terminal histories z of f, then $\{f, \bar{x}\}_h g \sim \{\bar{x}\}_h g$.

3.1.3 E-Certainty Equivalents, Equivalence Result and Updating

Consider the following definition:

Definition 2 Consider $E \in \mathrm{E}(\succcurlyeq)$, $f \in F_E$, and $\bar{x} \in X$; let $h = [(*, E)]$. Then $\bar{x}$ is a **E-certainty equivalent of** f iff Statements (i) and (ii) in Axiom 3.6 both hold.

Observe that E-certainty equivalents are defined solely in terms of the unconditional preference $\succcurlyeq$; as such, they are "conjectural" constructs—they reflect the conjectured evaluation of a continuation tree conditional upon E, from the perspective of "time 0". Proposition 3 below implies that this "conjectural" evaluation coincides with the actual "conditional" evaluation of the same act.

First of all, however, the existence of E-certainty equivalents must be established; the following Proposition (which, unlike Remark 3.1, does not follow from entirely standard arguments) provides the first main result of this subsection.

Proposition 2 *Suppose that $\succcurlyeq$ is a complete and transitive relation that satisfies Axioms 3.5, 3.6 and 3.7. Then, for all $E \in \mathrm{E}(\succcurlyeq)$ and $f \in F_E$, the set of E-certainty equivalents of f is non-empty, and forms an indifference class of prizes for $\succcurlyeq$.*

The generic E-certainty equivalent of the tree $f \in F_E$ will be denoted by $\bar{x}_{f|E}$.

It is now possible to establish the equivalence of the two approaches to conditional preferences discussed in this subsection.

Proposition 3 *Consider the conditional preference system* $\{\succcurlyeq_E\}_{\emptyset \neq E \subset \Omega}$*. Assume that* $\succcurlyeq$ *is a complete and transitive relation on* F_Ω*. Then the following statements are equivalent.*

1. $\succcurlyeq$ *satisfies Axioms 3.5, 3.6 and 3.7; furthermore, for all* $E \in \mathrm{E}(\succcurlyeq)$ *and all* $f, g \in F_E$, $f \succcurlyeq_E g$ *if and only if* $\bar{x}_{f|E} \succcurlyeq \bar{x}_{g|E}$*.*

2. $\{\succcurlyeq_E\}_{E \in \mathrm{E}(\succcurlyeq)}$ *is a collection of complete and transitive relations that satisfies Axioms 3.1, 3.2, 3.3 and 3.4.*

Furthermore, if either condition holds, then for every $E \in \mathrm{E}(\succcurlyeq)$*, any prize* $x \in X$*, and any tree* $f \in F_E$, $x \sim_E f$ *if and only if* x *is an* E*-certainty equivalent of* f*.*

Thus, if conditional preferences are defined as in 1 above, they satisfy the regularity and consistency axioms of subsection 3.1.1; conversely, if one assumes the existence of conditional preferences that satisfy the axioms in subsection 3.1.1, then *un*conditional preferences must satisfy the axioms in subsection 3.1.2 (in particular, Conjectural Separability), and it must be the case that the conditional ordering of any two acts coincides with the (unconditional) ordering of their respective E-certainty equivalents—which can be determined by observing *un*conditional preferences. The final claim in Proposition 3 confirms that, although E-certainty equivalents are a "conjectural" construct defined in terms of unconditional preferences alone, they do coincide with the intuitively more direct notion of conditional certainty equivalents.

3.2 A decision-theoretic analysis of Backward Induction

Subsection 3.2.1 formalizes the Sophistication and Weak Commitment axioms discussed in the Introduction; the characterization result is provided in §3.2.2.

3.2.1 Sophistication and Weak Commitment

Recall that *Sophistication* prescribes that the DM be ex-ante indifferent between a tree that makes available a collection of actions at some history h, and another tree that differs from the first only in that one or more ex-post strictly dominated actions are removed.

The formulation of this axiom is straightforward. At any non-terminal history h, any action $a \in A_f(h)$ available in the tree f at h, corresponds to a continuation tree $f(h, a) \in F_{E(h)}$; thus, the DM's conditional preference $\succcurlyeq_{E(h)}$ readily induces an ordering over actions available at h, which makes it possible to formalize the assumption that a subset B of such actions are conditionally strictly dominated.

Axiom 3.8 (Sophistication) *For all* $f = (E, H, x) \in F_E$*, all non-terminal* $h \in H$*, and all* $B \subset A_f(h)$*: if, for all* $b \in B$ *and* $w \in A_f(h) \setminus B$, $f(h, b) \succ_{E(h)} f(h, w)$*, then* $f \sim_E \{f(h, b) : b \in B\}_h f$*.*

As noted in the Introduction, a *tie breaking* rule may be requred in addition to Sophistication. Here, it will be assumed that, if the DM would like to commit to some action a at a subsequent history h, she *can* do so, provided a is conditionally optimal.

To formalize this assumption it is necessary to introduce additional notation. Consider an event $E \subset \Omega$, a tree $f \in F_E$, $h \in H$, and an action $a \in A_f(h)$. Recall that, after choosing a at h, the DM learns that one of the events in the set $\mathcal{F}_f(h,a)$—say, E'—contains the true state. If E' is a singleton, then the DM receives a prize and the decision problem terminates; otherwise, she will be able to choose any one of the actions available at $[h,(a,E')]$. I first define a continuation tree in $F_{E(h)}$ that differs from $f(h,a)$ only in that, following the choice of a, the DM commits to choosing one particular action at every non-terminal history of the form $[h,(a,E')]$, with $E' \in \mathcal{F}_f(h,a)$. Formally, assume that $E_1,\dots,E_n$ are the only non-singleton events in $\mathcal{F}_f(h,a)$, and consider $b_m \in A_f([h,(a,E_m)])$, for $m = 1,\dots,n$; finally, let $f^0 = f(h,a)$ and $f^m = \{f([h,(a,E_m)],b_m)\}_{[h,(a,E_m)]} f^{m-1}$ for $m = 1,\dots,n$. Then f^n differs from f only in that the choice b_m is made at $[h,(a,E_m)]$, as required; to make this explicit, let $f(h,a;b_1,E_1;\dots;b_n,E_m) = f^n$, and denote by $F_f(h,a)$ the collection of all continuation trees thus obtained: that is, $F_f(h,a) = \{f(h,a;b_1,E_1;\dots;b_n,E_n) \;:\; \forall m,\, b_m \in A_f([h,(a,E_m)]\}$. Finally, let $M_f(h,a)$ be the set of $\succcurlyeq_{E(h)}$–maximal elements of $F_f(h,a)$: that is,

$$M_f(h,a) = \{g \in F_f(h,a) \;:\; \forall g' \in F_f(h,a),\; g \succcurlyeq_{E(h)} g'\}; \tag{4}$$

Intuitively, the continuation trees in $M_f(h,a)$ represent courses of actions the DM *would* want to follow at history h, *if* she could commit to appropriate history-$[h,(a,E_m)]$ choice for every m.

The appropriate notion of commitment can now be formalized. Consider the following situation: at every history of the form $[h,(a,E')]$, with $E' \in \mathcal{F}_f(h,a)$ as above:

(a) all actions available to the DM correspond to *plans*, so any commitment problem at any history following the DM's choice of a at h must necessarily pertain to choices at histories *immediately* following a; and

(b) the DM is indifferent among all actions available at every history $[h,(a,E')]$ as above, so in fact there is *no* commitment problem.

Then, it makes sense to assume that:

(i) the DM is able to commit at h to the choice(s) she likes best at histories that immediately follow a; thus, she must be indifferent at h between continuing with $f(h,a)$ and choosing any plan $g \in M_f(h,a)$. Furthermore:

(ii) ex-ante, the DM should understand that her "history-h self" *will* have this capability to commit, and will take advantage of it.

Axiom 3.9 (Weak Commitment) *For all* $f = (E, H, x) \in F_E$ *and all histories* $h \in H$*: if, for some* $a \in A_f(h)$*, and for every non-singleton* $E' \in \mathcal{F}_f(h, a)$*,*

(a) $f([h,(a,E')], a')$ *is a plan for every* $a' \in A_f([h,(a,E')])$*, and*

(b) $f([h,(a,E')], a') \sim_{E'} f([h,(a,E')], b')$ *for all* $a', b' \in A_f([h,(a,E')])$*,*

then

(i) $f(h, a) \sim_{E(h)} g$ *for all* $g \in M_f(h, a)$*, and*

(ii) $f \sim_E \{\{f(h,a') : a' \in A_f(h) \setminus \{a\}\} \cup M_f(h,a)\}_h f$*.*

3.2.2 Formulation and Characterization of Backward Induction

Backward induction in decision trees can now be characterized as the iterative application of Axioms 3.8 and 3.9; upon each iteration, a suitably smaller and simpler equivalent tree is obtained. The procedure terminates when the reduced tree is a plan, i.e. a decision tree wherein a single action is available at every non-terminal history. Furthermore, in order to carry out the procedure, it is enough to specify the DM's conditional preferences over plans. One way to do so is to assume that the DM's preferences satisfy Reduction; in this case, any system of conditional preferences over Savage acts can be uniquely extended to a system of conditional preferences over trees. This approach will be exemplified in the next section. It should be noted, however, that the results in this section are independent of the actual way the DM evaluates plans: they do *not* require that plans be evaluated by reducing them to acts. In particular, they can accommodate a preference for early or late resolution of uncertainty.

The main result of this section also shows that, if Backward Induction is used to extend preferences from plans to arbitrary trees, then the resulting conditional preference system satisfies Axioms 3.8 and 3.9. Hence, these axioms fully characterize a complete behavioral theory of dynamic choice in the presence of ambiguity.

Consider a tree $f = (E, H, x)$; say that an action a at a history $h \in H$ is **reducible** if $f(h, a)$ is *not* a plan and, for every non-singleton element E' of $\mathcal{F}_f(h, a)$, and every $b \in A_f([h,(a,E')])$, $f([h,(a,E')], b)$ is a plan. That is, the action a is reducible if the only choices the DM needs to make are at histories immediately following h and a. Notice that, since $f(h, a)$ is not a plan, it is a fortiori not an act, which implies that one or more of the elements of $\mathcal{F}_f(h, a)$ *must* be non-singleton.

Algorithm 1 (Backward Induction) Let $E \subset \Omega$ be nonempty and $f \in F_E$.

1. Find a history h and a reducible action $a \in A_f(h)$. IF there is none, then STOP and RETURN the set $\{f\}$.

2. Denote the non-singleton elements of $\mathcal{F}_f(h, a)$ by $E_1, \ldots, E_n$. Inductively construct an act $f' \in F_E$ as follows. Let $f^0 = f$ and, for $m = 1, \ldots, n$,

(a) $h^m = [h, (a, E_m)]$,

(b) $M_m = \{b' \in A_f(h^m) \,:\, \forall a' \in A_f(h^m),\ f(h^m, b') \succcurlyeq_{E_m} f(h^m, a')\}$,

(c) $f^m = \{f(h^m, b') \,:\, b' \in M_m\}_{h^m} f^{m-1}$.

The required act f' is then f^n.

3. IF $h = \emptyset$, then STOP and RETURN the set $M_{f'}(h, a)$.

4. Let $f'' = \Big\{\{f(h, a') \,:\, a' \in A_f(h) \setminus \{a\}\} \cup M_{f'}(h, a)\Big\}_h f$.

5. REPLACE f WITH f'' and GO TO step 1.

At every iteration of the Backward Induction algorithm (BI henceforth), the continuation tree $f(h, a)$ corresponding to a reducible action a is replaced with a collection of plans. This replacement is carried out in two phases. First, consistently with Sophistication, all dominated actions at histories h^m that immediately follow h and a are eliminated; thus, at any such history h^m, the DM is indifferent among all remaining actions. Second, BI constructs plans[10] that select actions at each such history h^m in a way that is optimal from the perspective of history h: this step is justified by Weak Commitment. If there are no further reducible actions, BI outputs the resulting set of plans; otherwise, $f(h, a)$ is replaced with the set of plans thus constructed, and a new iteration begins.

The sequence of trees constructed in successive iterations of BI will be referred to as a **run** of the algorithm. In addition, in the last iteration, BI produces one or more plans as output.

Notice that, in general, an action a at a history h that is not reducible in the initial tree f may become reducible if actions at subsequent histories are reduced first. Thus, Algorithm 1 does indeed proceed "backwards", beginning with near-terminal histories and working towards the root of the decision tree. The following properties of BI are also noteworthy:[11]

1. Since decision trees are assumed to be finite, BI clearly terminates in a finite number of steps: that is, a BI solution always exists;

2. In each iteration, BI only requires comparisons among *plans*, not general trees;

3. If more than one reducible action is available in a given iteration, BI does not specify any particular order of reduction. However, the output of the procedure (the set of plans it produces in the last step) is independent of the order of elimination.

[10] Since a is reducible, the DM has a unique action available at every history that follows h^m. Hence, if a single action is selected at each h^m, the resulting tree is a plan.

[11] Properties 1 and 2 are obvious by inspection of Algorithm 1; Properties 3 and 4 are established in the course of the proof of Theorem 4.

4. The output of BI consists of plans that belong to the same indifference class for the DM.

The above conclusions do *not* depen upon any of the axioms mentioned in the preceding subsection: they follow directly from the definition of the algorithm. As will be clear momentarily, the role of these axioms is to ensure that the output of the BI procedure consists of plans that the DM deems indifferent to the input tree f.

The main result of this section follows.

Theorem 4 *Consider a system of preferences $\{\succcurlyeq_E\}_{\emptyset \neq E \subset \Omega}$ such that, for every non-empty $E \subset \Omega$, $\succcurlyeq_E$ is a complete and transitive binary relation on F_E. Then the following statements are equivalent.*

1. *Axioms 3.8 and 3.9 hold;*

2. *For every non-empty $E \subset \Omega$ and every $f \in F_E$, if BI outputs the collection $\{g_1, \ldots, g_n\}$ of plans, then $f \sim_E g_1 \sim_E \ldots \sim_E g_n$.*

3.3 Point Beliefs and Subgame Perfection

This subsection develops an alternative to the Weak Commitment axiom. The motivation for this is twofold. First, Weak Commitment serves as a tie-breaking rule in BI, but it is clearly not the only possible way to resolve indifferences. Recall that, according to Algorithm 1, if the DM is indifferent between two or more actions at a history h of length T, the tie is resolved by taking into account the preferences of her history-h_{T-1} self; if indifferences persist are the preferences of the history-h_{T-2} self are considered, and so on. Thus, in particular, BI favors the *last* non-indifferent self on the path to history h. While this assumption can be reasonably defended,[12] it is possible to define variants of BI that adopt different tie-breaking rules (such as favoring the *first* non-indifferent self on the path to h). This subsection considers a different alternative approach: essentially, a tie-breaking rule is *derived* from preferences, not imposed axiomatically.

Second, an explicitly game-theoretic approach is often adopted in the analysis of time-inconsistent choice, following the key contribution of Peleg and Yaari [23]. In this approach, the intertemporal decision problem is formally treated as an extensive game played by multiple selves. Under suitable regularity conditions, game-theoretic solutions exist for choice problems with infinitely many actions or decision epochs; by way of contrast, reasonably "regular" decision problems may fail to

[12] Loosely speaking, favoring the last non-indifferent self entails the least assumption of commitment abilities, which seems consistent with the basic motivation of this paper—as well as of Strotz [30] and related contributions. Additionally, in the context of choice under uncertainty, the last non-indifferent self has objectively superior (finer) information, which suggests that Weak Commitment may also have a "rationality" connotation.

admit consistent-planning solutions à la Strotz. The analysis in this subsection proposes a version of this game-theoretic approach to choice under uncertainty, and provides a decision-theoretic foundation that again involves Axiom 3.8, Sophistication. Results are stated here for finite decision problems; however, an extension to infinite trees is possible: see below for details. Thus, adopting the approach described in this subsection, sophisticated choice behavior is also possible in infinite decision problems.[13]

As was just noted, the key axiom in this Subsection is again Sophistication (Axiom 3.8). However, Weak Commitment (Axiom 3.9) is replaced with a weaker assumption: whenever the DM has two or more actions available at some history h of length T, the DM's selves at histories h' that precede h agree on (i.e. share the same beliefs about) the history-h self's choices; in particular, their beliefs are concentrated on a single action—they hold "point (or degenerate) beliefs".[14]

To formulate the axiom, and for the development to follow, it is convenient to denote the collection of non-empty, non-singleton, strict subsets of Ω by NT; such sets correspond to "non-trivial" conditional decision points. Also, denote the set of plans in F_E by F^p_E.

Axiom 3.10 (Point Beliefs) *Consider an event $E \in \mathrm{NT}$ and a finite set of plans $A \subset F^p_{E'}$. Then there exists $f \in A$ such that, for all non-empty events $E' \subsetneq E$, all $g = (E', H, x)$, and all $h \in H$ such that $E(h) = E$, $A_h g \sim_{E'} \{f\}_h g$.*

Strictly speaking, the Axiom requires that the beliefs concerning the DM's choice among a collection of plans $A \in F^p_E$ be the same in *any* decision tree where E can be reached and the set of plans A is available. This, however, is less restrictive than it may appear; since Reduction is *not* assumed, two plans that correspond to the same Savage act, but assign different labels to at least one action, are formally different. Thus, consider two decision trees that both allow E to be reached, and make available two sets of plans at E that reduce to the same set of Savage acts. Axiom 3.10 does *not* require consistency of beliefs across the two decision trees.

If Weak Commitment is not invoked as a tie-breaking rule, the Backward Induction algorithm must be modified; in particular, it must now accept as input a specification of the DM's beliefs about the choices of her future selves. To formalize this notion, for every $E \in \mathrm{NT}$, and given a

[13] If no restriction is placed on the functional representation of preferences and/or on the updating rule, it is possible to construct trees featuring a continuum of choices at one or more histories, for which Algorithm 1 fails to yield an equivalent plan. However, I have been unable to construct examples assuming e.g. MEU preferences and full Bayesian updating; if such examples exist, they are likely to be relatively involved. By way of contrast, examples of non-existence of consistent-planning solutions in the context of time-inconsistent choice under certainty can be relatively simple: see e.g. Gul and Pesendorfer [10].

[14] Gul and Pesendorfer [10] consider a similar (but conceptually more demanding) axiom, which they call "No Compromise".

preference relation $\succsim_E^*$ over F_E or F_E^p, let $I_E(\succsim_E^*)$ denote the collection of non-empty, finite subsets B of F_E^p such that $f, f' \in B$ implies $f \sim_E^* f'$. [The additional flexibility in the specification of $\succsim_E^*$ is required to streamline the description of the modified BI procedure below, and the corresponding characterization theorem.]

Now fix $E \in \mathrm{NT}$ and a preference relation $\succsim_E^*$ over F_E or F_E^p. A **belief over F_E^p for the preference** $\succsim_E^*$ is a map $\mathbf{B}_E : I_E(\succsim_E) \to F_E^p$ such that $\mathbf{B}_E(A) \in A$ for all $A \in I_E(\succsim_E^*)$. A **belief system** is a collection $\{\mathbf{B}_E\}_{E\in\mathrm{NT}}$, where each $\mathbf{B}_E$ is a belief over $I_p(\succsim_E^*)$, for some preference $\succsim_E^*$ over F_p or F_p^p.

The following algorithm extends BI in a manner that is reminiscent of *Subgame Perfection*.

Algorithm 2 (Subgame Perfection) Fix a belief system $\{\mathbf{B}_E\}_{E\in\mathrm{NT}}$. Let $E \in \mathrm{NT}$ and $f \in F_E$.

1. Find a history h and a reducible action $a \in A_f(h)$. IF there is none, then STOP and RETURN the plan f.

2. Denote the non-singleton elements of $\mathcal{F}_f(h, a)$ by $E_1, \dots, E_n$. Inductively construct an act $f' \in F_E$ as follows. Let $f^0 = f$ and, for $m = 1, \dots, n$,

   (a) $h^m = [h, (a, E_m)]$,

   (b) $M_m = \{b' \in A_f(h^m) : \forall a' \in A_f(h^m),\ f(h^m, b') \succsim_{E_m} f(h^m, a')\}$,

   (c) $f^m = \Big\{\mathbf{B}_{E(h^m)}\Big(\big\{f(h^m, b') : b' \in M_m\big\}\Big)\Big\}_{h^m} f^{m-1}$.

   The required act f' is then f^n.

3. IF $h = \emptyset$, then STOP and RETURN f'.

4. REPLACE f WITH f' and GO TO step 1.

Unlike BI, Algorithm 2 (SP henceforth) yields a single plan upon termination. Like BI, SP does not specify the order in which reducible actions are to be replaced; however, as in the case of BI, order turns out not to matter.

The characterization result can now be stated.

Theorem 5 *Consider a system of preferences $\{\succsim_E\}_{\emptyset\neq E\subset\Omega}$ such that, for every non-empty $E \subset \Omega$, $\succsim_E$ is a complete and transitive binary relation on F_E. Then the following statements are equivalent.*

1. *Axioms 3.8 and 3.10 hold;*

2. *There exists a belief system $\{\mathbf{B}_E\}_{E\in\mathrm{NT}}$ such that, for every non-empty $E \subset \Omega$ and every $f \in F_E$, if SP outputs the plan g, then $f \sim_E g$.*

Moreover, if another belief system $\{\mathbf{B}'_E\}_{E \in \mathrm{NT}}$ *satisfies (2), then, for all* $E \in \mathrm{NT}$, *all* $A \in I_E(\succsim_E)$, *all non-empty* $E' \subsetneq E$, *all* $g = (E', H, x)$ *and all* $h \in H$ *with* $E(h) = E$, $\{\mathbf{B}_E(A)\}_h g \sim_{E'} \{\mathbf{B}'_E(A)\}_h g$.

Theorems 4 and 5 are clearly similar in structure. It is worth emphasizing that the belief system $\{\mathcal{B}_E\}_{E \in \mathrm{NT}}$ can be *elicited from preferences* if the Point Beliefs axiom holds; indeed, beliefs satisfy a "uniqueness" property. Thus, belief systems are a fully behavioral construct.

Axioms 3.8 and 3.10 entail well-posed and non-degenerate restrictions on preferences in both finite and infinite decision problems.[15] The techniques in Goldman [9] and Harris [14] can be adapted to show that, for suitably regular infinite decision problems, SP is well-defined (in particular, the sets M_m of maximizers constructed in Step 2.b are non-empty at every iteration); this, in turn, can be shown to imply that Theorem 5 holds for such infinite decision problems.

4 An application: Full Bayesian Updating for MEU preferences

This section considers preferences consistent with the MEU decision model (Gilboa and Schmeidler [7]) and provides a characterization of Backward Induction under prior-by-prior, or "full" Bayesian updating. By Theorem 4, it is sufficient to provide a characterization of this updating rule for preferences over plans. Furthermore, the standard Reduction assumption will be adopted in this section: that is, the DM will be assumed to be indifferent between a plan and the corresponding Savage act. As a consequence of this assumption, the analysis in this section can be carried out almost exclusively in the familiar setting of preferences over Savage acts. However, once again, neither Proposition 3 nor Theorems 4 and 5 depend upon Reduction.

The results in this section are also meant to exemplify the general approach to dynamic decision making under ambiguity suggested in this paper. It is straightforward to adapt the analysis to different representations of preferences (e.g. Choquet-expected utility) and different updating rules (e.g. the Dempster-Shafer rule).

4.1 Setup

4.1.1 Acts and Plans

Recall that, in the present setting, a tree $f = (E, H, x) \in F_E$ is an act if H consists of the empty history, and of histories of the form $[(*, \{\omega\})]$, for all $\omega \in \Omega$. To simplify the notation, the prize assigned by f to the terminal history $[(*, \{\omega\})]$, corresponding to the event that ω is the prevailing

[15] This is *not* the case for Axiom 3.9: the set $M_f(h, a)$ may be empty for some history h and action a in a suitably pathological decision tree f; the latter Axiom yields no restrictions corresponding to h and a.

state, will be denoted simply by $f(\omega)$, insted of $x([(*, \{\omega\})])$. This is in accordance with standard notation for Savage acts.

Also recall that a plan is an act wherein a single action is available at every non-terminal history. As noted in §2, every act is a plan, but the converse is not true; however, there is a "natural" map from plans to acts. It is easy to verify that, in any plan, a unique prize can be assigned to every state $\omega \in \Omega$ by following the unique choices made by the DM at each history, and Nature's moves corresponding to ω: formally, if $f = (E, H, x) \in F_E$ is a plan, then for every $\omega \in \Omega$ there is a unique terminal history $z_\omega \in H$ such that $E(z_\omega) = \{\omega\}$.[16] Hence, one can define an act, denoted $\hat{f} = \{E, \hat{H}, \hat{x}), characterized by $\hat{H} = \{\emptyset\} \cup \{[(*, \{\omega\})] : \omega \in E\}$, and $\hat{x}([(*, \{\omega\})]) = x(z_\omega)$. This new act $\hat{f}$ will be referred to as the **act derived from the plan** f.

It is convenient to denote the set of Savage acts and plans in F_E by F^a_E and F^p_E respectively. Also, it is convenient to introduce simplified notation for **eventwise combinations** of acts. Specifically, given $E \subset \Omega$, $f \in F^a_E$ and $g \in F^a_{\Omega \setminus E}$, denote by fEg the act in F_Ω such that

$$\forall \omega \in E, \qquad fEg(\omega) = \begin{cases} f(\omega) & \omega \in E; \\ g(\omega) & \omega \in \Omega \setminus E. \end{cases}$$

Finally, it is also useful to introduce notation for **truncations** of acts. Consider $f \in F^a_\Omega$ and $E \subsetneq \Omega$ non-empty; the act $f|_E \in F^a_E$ is defined by

$$\forall \omega \in E, \quad f|_E(\omega) = f(\omega).$$

4.1.2 MEU representation; Reduction

It will now be assumed that unconditional preferences over acts have a MEU representation.

Assumption 4.1 (MEU) There exists a closed,[17] convex set of probability measures $C \subset \Delta(\Omega)$ and a continuous function $u : X \to \mathbb{R}$ such that

$$\forall f, g \in F^a_\Omega, \qquad f \succcurlyeq g \quad \Leftrightarrow \quad \min_{q \in C} \int_S u(f(s)) q(d\omega) \geq \min_{q \in C} \int_S u(g(s)) q(d\omega).$$

Moreover, there exist $f, g \in F^a_\Omega$ such that $f \succ g$.

[16] Suppose there are two, say z, z'; let h be the longest history such that $h \leq z$, $h \leq z'$ (perhaps $h = \emptyset$). Suppose that in fact $h < z, z'$. Since f is a plan, $A_f(h) = \{a_h\}$ is a singleton, so there must be $E, E' \subset E(h)$ such that $E \cap E' = \emptyset$and $[h, (a_h, E)] \leq z$, $[h, (a_h, E')] \leq z'$. But this implies that $E(z) \subset E$ and $E(z') \subset E'$, so $E \cap E' \neq \emptyset$: hence, it cannot be that $h < z, z'$. The argument for $h < z$, $h = z'$ is analogous.

[17] If the analysis is extended to infinite state spaces Ω, then C must be assumed to be weak*-closed.

The last part of Assumption 4.1 is a non-triviality requirement. Notice that no assumption is made concerning the representation of conditional preferences.

For simplicity, it will also be assumed that all non-empty events are $\succcurlyeq$-essential. The following requirement is slightly weaker than the notion of essentiality introduced in the preceding section, because it only pertains to acts; however, if Axiom 4.2 holds, together with Axioms 3.8 and 3.9, the two notions turn out to coincide.

Assumption 4.2 (Essential Events) For all non-empty $E \subset \Omega$, and for all $x, x' \in X$ such that $x \succ x'$, $x \succ xEx' \succ x'$.

Assumption 4.2 holds if and only if $\min_{q \in C} q(E) > 0$ for all non-empty events $E \subset \Omega$.

Finally, the assumption that the DM evaluates plans by reducing them to acts is explicitly formulated. The following statement also provides a first example of the notation $\hat{f}$ introduced above to indicate the act derived from the plan f.

Definition 3 (Reduction) For all non-empty $E \subset \Omega$, the preference $\succcurlyeq_E$ is said to *satisfy Reduction* if and only if, for all $f \in F_E^p$, $f \sim_E \hat{f}$.

4.2 Dynamic Consistency and Constant-Act Dynamic Consistency

It is useful to consider the standard Dynamic Consistency axiom as a starting point. The following is essentially the axiom stated in Subsection 1.1.1, modified so as to account for the fact that, consistently with the decision setting in Sections 2 and 3, each conditional preference $\succcurlyeq_E$ is defined over partial decision trees, and in particular partial acts, i.e. maps from E to X.

Axiom 4.1 (Dynamic Consistency) *For all non-empty $E \subsetneq \Omega$, all $f \in F_E^a$, and all $g \in F_\Omega^a$:*

$$f \succcurlyeq_E g|_E \iff fEg \succcurlyeq g.$$

Dynamic Consistency may be inconsistent with full Bayesian updating and non-neutral attitudes towards ambiguity. Consider the three-color-urn example described in the Introduction, and refer to the tree f_x in Fig. 1. Recall that $\Omega = \{r, g, b\}$, representing the draw of a red, green, or blue ball respectively, and that X contains the prizes $x = 0$ and $x = 10$. Assume that $\succcurlyeq$ is a MEU preference characterized by a utility function u with $u(10) > u(0)$ and the set of priors $C = \{q \in \Delta(\Omega) : q(r) = \frac{1}{3}, q(g) \geq \epsilon > 0, q(b) \geq \epsilon\}$, with $\epsilon \in (0, \frac{1}{3})$. Recall that f_x^R and f_x^G are plans obtained from the tree f_x in Fig. 1 by removing the green and red subtrees respectively. It is easy to verify that $\hat{f}_{10}^R \prec \hat{f}_{10}^G$: note that $\hat{f}_{10}^R$ (the act derived from the plan f_{10}^R) represents a bet on "red or blue", as featured in the usual, static version of the Ellsberg paradox, whereas $\hat{f}_{10}^G$ represents a bet on "green or blue". Now let $E = \{r, g\}$ and assume that $\succcurlyeq_E$ is a MEU preference, characterized by the set of posteriors C_E derived from C by full Bayesian updating:

$$C_E = \left\{ q \in \Delta(S) \ : \ \frac{\frac{1}{3}}{1-\epsilon} \leq q(\rho) = 1 - q(\upsilon) \leq \frac{\frac{1}{3}}{\frac{1}{3}+\epsilon} \right\}.$$

Then $R \succ_E G$;[18] since $\hat{f}^R_{10} = R\,E\,\hat{f}^G_{10} \prec \hat{f}^G_{10}$, Dynamic Consistency is violated.[19]

As in the static version of the Ellsberg paradox, the preference for $\hat{f}^G_{10}$ over $\hat{f}^R_{10}$ (a bet on "green or blue" vs. a bet on "red or blue") is rationalized by the consideration that the urn is guaranteed to contain 60 non-red balls, but may contain very few blue balls. On the other hand, once this ambiguity-averse DM learns that the ball drawn is not blue, the possibility that the urn contains very few green balls looms larger; hence the conditional preference for R over G.

In loose but suggestive terms, the elementary events $\{g\}$ and $\{b\}$ are "complementary": acts that deliver good outcomes on their union $\{g, b\}$ are especially valuable to this ambiguity-averse DM. Learning that $\{b\}$ has not occurred essentially prevents the DM from taking advantage of this complementarity, so it is not surprising to observe a "preference reversal".

The following axiom identifies situations in which no such reversals should occur, even if the DM is ambiguity averse.

Axiom 4.2 (Constant-Act Dynamic Consistency) *For every non-empty $E \subset \Omega$, all acts $f \in F^a_E$, and all prizes $x \in X$:*

$$f \succcurlyeq_E x \quad \Longleftrightarrow \quad fEx \succcurlyeq x;^{20}$$

furthermore, if one preference is strict, so is the other.

Axiom 4.2 differs from Axiom 4.1 only in that the act $g \in F^a_\Omega$ in the latter is required to be a constant act $x \in X$ in the former. With this modification, the preceding discussion suggests that $f \succcurlyeq_E x$ should imply $fEx \succcurlyeq x$; loosely speaking, the sets $\{\omega \subset E : f(\omega) \succcurlyeq x\}$ and $\Omega \setminus E$ may be complementary, in which case the DM's prior preference for fEx over x should be at least as strong as her conditional preference for f over x. Similarly, if $f \succ_E x$, then $fEx \succ x$.

[18] Note that the subtrees R and G are acts.

[19] This violation may of course also be seen as an immediate consequence of the analysis in the Introduction, which does not rely upon specific functional-form assumptions. However, this example is useful to motivate a weakening of Axiom 4.1 that is compatible with arbitrary MEU preferences and full Bayesian updating.

[20] This statement exploits the usual abuse of notation for constant acts. In $fEx \succcurlyeq x$, "x" is the constant act in F^a_Ω that yields the prize $x \in X$ in every state. In $f \succcurlyeq_E x$, the "x" on the right-hand side should be thought of as the restriction of the constant act $x \in F^a_\Omega$ to E; but since this restriction is itself a constant act in F^a_E, the simpler notation is adopted in lieu of "$x|_E$".

The opposite implication is more interesting, and suggests a connection with full Bayesian updating for MEU preferences. To avoid trivialities, assume that $f(\omega) \succ x \succ f(\omega')$ for some $\omega, \omega' \in E$. Very loosely speaking, an ambiguity-averse DM will evaluate the act fEx by assigning as much "weight" as possible to states where, in particular, this composite act yields prizes worse than x. These states necessarily all belong to the set E. Therefore, if $fEx \succcurlyeq x$, then, despite the possible presence of ambiguity, the DM still considers the relative likelihood of states in E where f yields a prize equal to or better than x, vs. states where f yields prizes worse than x, to be suitably "large". Thus, it seems plausible to require that $f \succcurlyeq_E x$. More precisely: if $\succcurlyeq_E$ is derived from $\succcurlyeq$ by means of an updating rule that *preserves relative likelihood assessments concerning disjoint subsets of* E, then $fEx \succcurlyeq x$ should imply $f \succcurlyeq_E x$. Full Bayesian updating is of course one such rule; indeed, the main result of this section shows that Axiom 4.2 completely characterizes it.[21]

Proposition 6 *Under Assumptions 4.1 and 4.2, the following statements are equivalent:*

1. *Every $\succcurlyeq_E$ (E non-empty) is complete and transitive on F^a_E, and Axiom 4.2 holds;*

2. *For every non-empty $E \subset \Omega$, and for every $f, g \in F^a_E$,*

$$f \succcurlyeq_E g \quad \text{if and only if} \quad \min_{q \in C} \int_E u \circ f \; q(d\omega|E) \geq \min_{q \in C} \int_E u \circ g \; q(d\omega|E) \tag{5}$$

   where u and C are as in Assumption 4.1.

It should be observed that this characterization result generalizes to arbitrary state spaces.

Notice that Proposition 6 does *not* provide a recursive MEU representation of preferences over acts (unlike the main result in Epstein and Schneider [5]). However, if preferences satisfy Reduction, it guarantees that *every decision problem can be solved using the BI algorithm, invoking the conditional MEU representation in Eq. 5.*

The straightforward details are as follows. First, BI must be modified to employ the conditional MEU representation; notice that this is possible because, in each step and at each iteration, BI relies solely upon comparisons of *plans*, which can be reduced to acts if Reduction holds.

Algorithm 3 (Backward Induction for MEU preferences) Let $E \subset \Omega$ be nonempty and $f = (E, H, x) \in F_E$. Let u and C be as in Assumption 4.1. Furthermore, assume that $\succcurlyeq_{E(h)}$ satisfies Reduction for every non-terminal history $h \in H$.

---

[21] Full Bayesian Updating has been independently characterized by Jaffray [15] and Pires [24] (as well as in Siniscalchi [29]; also, although it is not explicitly decision-theoretic, the earliest characterization is probably due to Walley [31]). In lieu of Axiom 4.2, these authors employ "fixpoint condition" that also appears as a step in the proof of Proposition 6. Axiom 4.2 has the advantage of emphasizing similarities and differences with Dynamic Consistency. But, again, the main purpose of this Section is to exemplify the approach adopted in this paper; in this respect, Theorem 7 is the more "significant" contribution here.

1. Find a history h and a reducible action $a \in A_f(h)$. IF there is none, then STOP and RETURN the set $\{f\}$.

2. Denote the non-singleton elements of $\mathcal{F}_f(h,a)$ by $E_1, \ldots, E_n$. Inductively construct an act $f' \in F_E$ as follows. Let $f^0 = f$ and, for $m = 1, \ldots, n$,

   (a) $h^m = [h, (a, E_m)]$,

   (b) M_m is the collection of $b' \in A_f(h^m)$ such that

   $$\forall a' \in A_f(h^m), \min_{q \in C} \int_{E_m} u \circ \hat{f}(h^m, b') dq(\omega|E_m) \geq \min_{q \in C} \int_{E_m} u \circ \hat{f}(h^m, a') dq(\omega|E_m),$$

   (c) $f^m = \{f(h^m, b') : b' \in M_m\}_{h^m} f^{m-1}$.

   The required act f' is then f^n.

3. Define the set

   $$M_{f'}(h,a) = \Big\{ g \in F_{f'}(h,a) \,:\, \forall g' \in F_{f'}(h,a), \min_{q \in C} \int_E u \circ \hat{g} dq(\omega|E) \geq \min_{q \in C} \int_E u \circ \hat{g}' dq(\omega|E) \Big\}.$$

4. IF $h = \emptyset$, then STOP and RETURN the set $M_{f'}(h,a) = M_{f'}(\emptyset, *)$.

5. Let $f'' = \Big\{ \{f(h,a') : a' \in A_f(h) \setminus \{a\}\} \cup M_{f'}(h,a) \Big\}_h f$.

6. REPLACE f WITH f'' and GO TO step 1.

The counterpart to Theorem 4 can then be stated.

Theorem 7 *Consider a system of preferences $\{\succcurlyeq_E\}_{\emptyset \neq E \subset \Omega}$; suppose that Assumptions 4.1 and 4.2 hold, and that $\succcurlyeq$ satisfies Reduction. Then the following statements are equivalent.*

1. *Every $\succcurlyeq_E$ is complete and transitive, and satisfies Reduction; moreover, Axioms 3.8, 3.9 and 4.2 hold;*

2. *For every non-empty $E \subset \Omega$ and every $f \in F_E$, if Algorithm 3 outputs the collection $\{g_1, \ldots, g_n\}$ of plans, then $f \sim_E g_1 \sim_E \ldots \sim_E g_n \sim x$, where $x \in X$ and*

   $$u(x) = \min_{q \in C} \int_E u \circ \hat{g}_1 q(d\omega|E).$$

Notice that the assumptions of Theorem 7 pertain solely to the prior preference $\succcurlyeq$. Statement 2 asserts that, conditional upon any event E, (i) acts are evaluated according to full Bayesian updating, (ii) plans are reduced to acts, and (iii) trees are reduced to plans by BI. In particular, (i) and (ii) follow from the observation that, if f is a plan or an act, BI outputs f itself.

5 Concluding Remarks

As was pointed out several times, Consequentialism is implicitly assumed in the formal setup of this paper: conditional preferences are defined over continuation trees, not full decision trees.

If instead prior and conditional preferences are defined over full decision trees[22] and Consequentialism is *not* assumed, a version of Theorem 4 still holds, with one important modification. Recall that, in the (consequentialist) setting of this paper, if two different runs of the BI algorithm yield two plans g_1 and g_2 given the same input $f \in F_\Omega$, then $g_1 \sim g_2$, *regardless of whether or not Axioms 3.8 and 3.9 hold* (the latter are of course needed to ensure that $f \sim g_1 \sim g_2$). In a non-consequentialist setup, without further assumptions on conditional preferences, different runs of BI with the same input might yield plans that are *not* equivalent for the DM. On the other hand, natural extensions of Axioms 3.8 and 3.9 *do* rule out such indeterminacies. In other words, the indifference $g_1 \sim g_2$ would be part of the characterization result, rather than simply a consequence of the definition of the BI algorithm.

The implications of this fact deserve further investigation, and are left to future research.

The analysis in this paper focuses on finite trees. This makes it possible to investigate the role of Sophistication (and Weak Commitment) in dynamic decision making without restricting attention to any parametric class of decision models: in particular, recall that Theorems 4 and 5 only require that conditional preferences be complete and transitive.

However, as was mentioned in Section 3.3, the SP algorithm and its characterization in Theorem 5 can be extended to suitably regular infinite decision trees.

For specific infinite choice problems, the existence of solutions consistent with both Sophistication and Weak Commitment must be established directly. Future research may provide general regularity conditions that guarantee existence for interesting parametric classes of decision models and updating rules (e.g. MEU with full Bayesian updating).

A Proofs

A.1 Eliciting Conditional Preferences

A.1.1 Proof of Remark 3.1 (Conditional Certainty Equivalents)

Proof: Denote by U and L the closed sets in Axiom 3.3. Since Ω is finite, there exist x', x'' such that $x' \succcurlyeq x(z) \succcurlyeq x''$ for all terminal histories z of f. By Axioms 3.1, 3.2 and transitivity, $x' \in U$ and $x'' \in L$. By completeness, $U \cup L = X$. Since X is separable, the non-empty, closed sets U and

[22] In particular, in view of the considerations in Footnote 8, restrict attention to *partitional* trees as defined in §2.

L must have non-empty intersection; any $x \in U \cap L$ clearly satisfies $x \sim_E f$. ■

A.1.2 Proof of Proposition 2 (E-certainty equivalents)

Proof: Consider the sets

$$U(f|E) = \Big\{\bar{x} \in X \;:\; \forall x \in X \text{ s.t. } x \succ \bar{x},\, \forall g \in F_\Omega(E),\; \{f,x\}_h g \sim \{x\}_h g\Big\}, \tag{6}$$

$$L(f|E) = \Big\{\bar{x} \in X \;:\; \forall x \in X \text{ s.t. } x \prec \bar{x},\, \forall g \in F_\Omega(E),\; \{f,x\}_h g \sim \{f\}_h g\Big\}. \tag{7}$$

Clearly, the set of E-certainty equivalents of f is $U(f|E) \cap L(f|E)$.

Note that, if $x \in U(f|E)$ [resp. $x \in L(f|E)$] and $x' \sim x$, then clearly $x' \in U(f|E)$ [resp. $x' \in L(f|E)$]. Also, suppose $x \in L(f|E)$, $x' \in U(f|E)$, and $x \succ x'$. By Axiom 3.5, there exist $x_1, x_2 \in X$ such that $x \succ x_1 \succ x_2 \succ x'$.[23] For $i = 1, 2$, $x \in L(f|E)$ and $x_i \prec x$ imply $\{f, x_i\}_h g \sim \{f\}_h g$, whereas $x' \in U(f|E)$ and $x_i \succ x'$ imply $\{f, x_i\}_h g \sim \{x_i\}_h g$. Thus,

$$\{x_1\}_h g \sim \{f, x_1\}_h g \sim \{f\}_h g \sim \{f, x_2\}_h g \sim \{x_2\}_h g,$$

which implies that E cannot be essential. Hence, if E is essential, then $x \in L(f|E)$ and $x' \in U(f|E)$ imply $x' \succcurlyeq x$; in particular, either $U(f|E) \cap L(f, E)$ is empty, or it is an indifference class of prizes for $\succcurlyeq$.

Moreover, it is clear that $x \in U(f|E)$ and $x' \succ x$ imply $x' \in U(f|E)$; similarly, $x \in L(f|E)$ and $x' \prec x$ imply $x' \in L(f|E)$. Now consider $U(f|E)$; since Ω is finite,[24] there is $\bar{x} \in X$ such that $\bar{x} \succcurlyeq x(z)$ for all terminal histories $z \geq h$. If there is no $x \in X$ with $x \succ \bar{x}$, then $\bar{x} \in U(f|E)$ holds vacuously; otherwise, Axiom 3.7 ensures that, for all $x \succ \bar{x}$, $\{f, x\}_h g \sim \{x\}_h g$, so again $\bar{x} \in U(f|E)$. Similarly, there is $\bar{x}' \in X$ such that $\bar{x}' \preccurlyeq x(z)$ for all $z \geq h$, and again $\bar{x}' \in L(f|E)$. Hence, in particular, both $U(f|E)$ and $L(f|E$ are non-empty.

It follows that either $U(f|E) = \{x : x \succ x_U\}$ or $U(f|E) = \{x : x \succcurlyeq x_U\}$ for some $x_U \in X$. Also, in either case, for any $x \succ x_U$, $\{f, x\}_h g \sim \{x\}_h g$: to see that, recall that, as argued above, there exists x' such that $x \succ x' \succ x_U$, so in either case $x' \in U(f|E)$, and the required conclusion holds. But then by definition $x_U \in U(f|E)$, i.e. it must be the case that $U(f|E) = \{x : x \succcurlyeq x_U\}$. Similarly, $L(f|E) = \{x : x \preccurlyeq x_L\}$ for some $x_L \in X$.

Finally, by Axiom 3.6, $U(f, E) \cup L(f, E) = X$. By Axiom 3.5, $U(f|E)$ and $L(F|E)$ are closed, so they cannot be disjoint because X is connected. Clearly, any $x \in U(f|E) \cap L(f|E)$ is an E-certainty

[23] Suppose that, for all x'', either $x'' \succcurlyeq x$ or $x'' \preccurlyeq x'$: by Axiom 3.5, $\{x'' : x'' \succcurlyeq x\}$ and $\{x'' : x'' \preccurlyeq x'\}$ are closed, so X is the union of two disjoint closed sets, which contradicts the assumption that X is connected. Thus, there is $x_1 \in X$ with $x \succ x_1 \succ x'$. Now repeat the argument with x_1 in lieu of x to get a suitable x_2.

[24] If it is not, we assume trees are "bounded" in the usual sense, so the argument still goes through.

equivalent of f. ■

A.1.3 Proof of Proposition 3 (Equivalence result)

Proof: Let $U(f|E)$ and $L(f|E)$ be as in Eqs. (6) and (7). Observe that $x \succcurlyeq \bar{x}_{f,E} \succcurlyeq x'$ for all $x \in U(f|E)$ and $x' \in L(f|E)$.

$\Rightarrow$: it is clear that, if $f \succcurlyeq_E g$ iff $\bar{x}_{f|E} \succcurlyeq \bar{x}_{g|E}$, then $\succcurlyeq_E$ is a weak order. Furthermore, consider a constant act $\bar{x}$. By Axiom 3.7, $\bar{x} \in U(\bar{x}|E) \cap L(\bar{x}|E)$; hence, $\bar{x}_{\bar{x},E} \sim \bar{x}$ for any constant $\bar{x}$; this implies that $\succcurlyeq_E$ satisfies Axiom 3.1.

To verify that Axiom 3.2 holds, consider $x \in X$ such that $x' \succcurlyeq x(z)$ for all terminal z. If there is no $x \succ x'$, then $x' \in U(f|E)$ holds vacuously; otherwise, Axiom 3.7 ensures that $\{f,x\}_h g \sim \{x\}_h g$ for all $x \succ x'$ and $g \in F_\Omega(E)$, so again $x' \in U(f|E)$. This implies that $x' \succcurlyeq \bar{x}_{f,E}$, and hence $x' \succcurlyeq_E f$, as required. The argument for $x'' \preccurlyeq x(z)$ is analogous.

As for Axiom 3.3, the first set in the statement can be written as $\{x \in X : x \succcurlyeq_E \bar{x}_{f,E}\}$, i.e. $\{x \in X : x \succcurlyeq \bar{x}_{f,E}\}$, which is closed by Axiom 3.5; similarly for the other set.

Finally, to see that Axiom 3.4 holds, note that $x \succ_E f$ implies $x \succ_E \bar{x}_{f,E}$, i.e. $x \succ \bar{x}_{f,E}$; since $\bar{x}_{f,E} \in U(f|E)$, this implies $\{f,x\}_h g \sim \{x\}_h g$ for all $g \in F_\Omega(E)$; similarly for (ii).

$\Leftarrow$: consider $E \in \mathrm{E}(\succcurlyeq)$. Since $\succcurlyeq_E$ satisfies Axiom 3.1 and 3.3, $\succcurlyeq$ must satisfy Axiom 3.5: just take $f = \bar{x}$ in Axiom 3.3. Next, for Axiom 3.6, let $x_{f|E}$ be such that $x_{f|E} \sim f$ (one such prize must exist by Remark 3.1). Suppose $\bar{x} \succcurlyeq x_{f,E}$: then $x \succ \bar{x}$ implies $x \succ_E f$, so Axiom 3.4 implies $\{f,x\}_h g \sim \{x\}_h g$ for any $g \in F_\Omega(E)$; if instead $\bar{x} \preccurlyeq_E x_{f,E}$, $x \prec \bar{x}$ implies $x \prec_E f$, and again Axiom 3.4 yields the required conclusion.

Now consider Axiom 3.7, case (i). Let x_L be such that $x_L \preccurlyeq x(z)$ for all terminal histories z of f, and $x_L = x(z_L)$ for some such z_L. Axiom 3.2 implies that $f \succcurlyeq_E x_L$, and by assumption $x_L \succ \bar{x}$. By Axiom 3.1 and transitivity, $f \succ_E \bar{x}$, so Axiom 3.4 implies $\{f,x\}_h g \sim \{f\}_h g$ for all $g \in F_\Omega(E)$, as required. Case (ii) is analogous.

Finally, it will be shown that $x \sim_E f$ iff $x \sim \bar{x}_{f|E}$, which is the last claim of Proposition 3; this also implies that $f \succcurlyeq_E g$ iff $\bar{x}_{f|E} \succcurlyeq \bar{x}_{g|E}$, which completes the proof that 2 implies 1. Suppose $x \sim_E f$: then $x' \succ x$ implies $x' \succ_E f$ by Axiom 3.1 and transitivity, and hence $\{f,x'\}_h g \sim \{x'\}_h g$ whenever $h = [(*,E)]$ and $g \in F_\Omega(E)$, by Axiom 3.4. Similarly, if $x' \prec x$, then $x' \prec_E f$, and hence $\{f,x'\}_h g \sim \{f\}_h g$. Thus, $x \sim_E f$ implies that x is an E-certainty equivalent of f. Conversely, since it was just shown that $\succcurlyeq$ satisfies Axioms 3.5, 3.6 and 3.7, Proposition 2 can be invoked to conclude that the set of E-certainty equivalents is an indifference class for $\succcurlyeq$; by the argument just given it contains all $x \in X$ such that $x \sim_E f$: thus, $\bar{x}_{f|E} \sim x$, which implies $\bar{x}_{f|E} \sim_E f$ by Axiom 3.1 and transitivity, as required. ■

A.2 Backward Induction and Subgame Perfection

A.2.1 Preliminary Results (common to BI and SP)

It is useful to establish two facts concerning the BI and SP algorithms.

Notice first that, in any given run $f^0, \ldots, f^K$ of BI, every history h of the original tree f is also a history of $f^0, \ldots, f^{k^*}$, where k^* is the iteration at which the last reducible action at h is replaced.[25]

Define the **height** $\eta(h)$ of a history h as follows. Fix a tree $f = (E, H, x) \in F_E$); if h is terminal, then $\eta(h) = 0$; otherwise, $\eta(h) = \max\{\eta(h') \; : \; h < h'\} + 1$.

Remark A.1 Consider two runs $f^0, \ldots, f^K$ and $\bar{f}^0, \ldots, \bar{f}^{\bar{K}}$ of BI or SPI (possibly with $f^0 \neq \bar{f}^0$). If there exist $k, \bar{k}$ such that $f^k = \bar{f}^{\bar{k}}$, then both runs yield the same output.

Proof: It is clear that $f^k, \ldots, f^K$ and $\bar{f}^{\bar{k}}, \ldots, \bar{f}^{\bar{K}}$ are also runs of BI/SP (with inputs f^0 and $\bar{f}^0$ respectively), so it is wlog to take $k = \bar{k} = 0$; thus, $f^0 = \bar{f}^0 = f$. Notice that, if the two runs have different outputs, there must be some history of f and some action available at that history that is replaced with different (sets of) plans in the two runs; consider the *longest* history h where such an action a can be found. Formally, let ℓ and $\bar{\ell}$ be such that (1) h is a history of f and a is reducible in f^ℓ and $\bar{f}^{\bar{\ell}}$ (hence, as noted above, h is also a history of f^ℓ and $\bar{f}^{\bar{\ell}}$); (2) a is replaced in the ℓ-th and $\bar{\ell}$-th iterations of the two runs respectively, with *different* sets of plans; and (3) at all histories of f longer than h, the two runs replace reducible actions with the same plan or plans.

Since a is reducible, for all $E' \in \mathcal{F}_{f^\ell}(h, a) = \mathcal{F}_{\bar{f}^{\bar{\ell}}}(h, a) = \mathcal{F}_f(h, a)$, $f^\ell([h, (a, E')], a')$ is a plan for all $a' \in A_{f^\ell}([h, (a, E')])$, and $\bar{f}^{\bar{\ell}}([h, (a, E')], a')$ is a plan for all $a' \in A_{\bar{f}^{\bar{\ell}}}([h, (a, E')])$.

Furthermore, there must be E^* such that $\{f^\ell([h, (a, E^*)]), a') \; : \; a' \in A_{f^\ell}([h, (a, E^*)])\} \neq \{\bar{f}^{\bar{\ell}}([h, (a, E^*)]), a') : a' \in A_{\bar{f}^{\bar{\ell}}}([h, (a, E^*)])\}$, or BI/SP would prescribe the same replacement for a in both runs. But this implies that the two runs of the algorithm perform different replacements for some action $a^* \in A_f([h, (a, E^*)])$, which contradicts the assumption that equal replacements are performed at all histories longer than h. ■

Corollary 8 *Suppose that BI or SP output the collection* $\{g_1, \ldots, g_N\}$ *for the continuation act* $f(h, a)$;[26] *then BI or, respectively, SP produce the same output for both* f *and* $\bar{f} = \{\{f(h, b) : b \in$

[25] It may also be a history of subsequent elements of the run, in case $a(h)$ is not immediately replaced in the $(k+1)$-th iteration. This, however, is unimportant for the present purposes.

[26] In the case of SP, $N = 1$.

$A_f(h) \setminus \{a\}\} \cup \{g_1, \dots, g_N\}\}_h f$.

Proof: The claim is obvious if $f(h,a)$ is itself a plan; thus, assume it is not. Consider a run $f^0, \dots, f^K$ with input $f^0 = f$, and a run $\bar{f}^0, \dots, \bar{f}^{\bar{K}}$ with input $\bar{f}^0 = \bar{f}$. Suppose that the runs have the following features: in $f^0, \dots, f^K$, only histories that follow h and a are considered (i.e. histories h' such that $h_t = h$ and $a(h_{t+1}) = a$ for some $t < \lambda(h')$), until, for some $\ell > 0$, either $f^\ell(h,a)$ is a plan, or a is reducible at h in f^ℓ. Now observe that $f^0(h,a), \dots, f^\ell(h,a) \in F_{E(h)}$ constitute a run of BI/SP with input $f(h,a)$. If $f^\ell(h,a)$ is a plan, then BI/SP must output $\{f^\ell(h,a)\}$ for the input $f(h,a)$, so in fact $f^\ell = \bar{f}^0$, and Remark A.1 yields the result. In the other case, an additional iteration of BI/SP for the input $f(h,a)$ yields $\{g_1, \dots, g_N\}$; wlog, assume that the run $f^0, \dots, f^K$ with input f also replaces the reducible action a at the $(\ell+1)$-st iteration; the set of plans substituted for $f^\ell(h,a)$ is clearly $\{g_1, \dots, g_N\}$. But then $f^{\ell+1} = \bar{f}^0$, and again the claim follows from the Remark. ■

A.2.2 Proof of Theorem 4 (BI)

Proof: (1) $\Rightarrow$ (2): It must first be shown that, in any run $f^0, \dots, f^K$ of BI, $f^k \sim_E f^{k-1}$ for all $k > 0$; then, it will be argued that, starting with f^K, an additional iteration of BI yields a collection of plans $\{g_1, \dots, g_N\}$ that satisfies the required property.

Consider an arbitrary non-terminal iteration $k \in \{1, \dots, K\}$ (if BI terminates immediately, there is no such iteration, of course). By convention, the input to this iteration is f^{k-1}, and the output is f^k. To simplify notation, consistently with the description of Algorithm 1, denote the input f^{k-1} to the k-th iteration by f, and the resulting output f^k by f''; thus, it must be shown that $f \sim_E f''$ if the algorithm does not terminate in Steps 1 or 3.

Since the algorithm does not terminate, there must be a reducible action a at a history h in Step 1. Now consider Step 2. Clearly, for $m = 1, \dots, n$, Axiom 3.8 implies that $f^m \sim_E f^{m-1}$ (take $B = M^m$), so $f \sim f'$. Since BI does not terminate at the k-th iteration, consider Step 4 next. Then (ii) in Axiom 3.9 ensures that $f'' \sim_E f' \sim_E f$ (note that f' and f coincide at histories that do not weakly follow h). This proves the first claim.

We thus have $f^0 \sim_E f^K$. If now, at the $(K+1)$-th iteration, BI terminates in Step 1, the output is $\{f^K\}$ and the result follows immediately. Otherwise, BI must terminate in Step 3; this means that the only reducible action is $*$ at $\emptyset$. In this case, again denote the input f^K to the last iteration of BI by f; as above, the tree f' constructed in Step 2 satisfies $f \sim_E f'$. Furthermore, (i) in Axiom 3.9 implies that, for any $g \in M_{f'}(\emptyset, *)$, $f' = f'(\emptyset, *) \sim_E g$; therefore, $f \sim_E g$ as well, and the proof of this direction is complete.

(2) ⇒ (1): Consider Axiom 3.8 first, and let f and h be as in the statement of the latter. For every $a \in A_f(h)$, denote by $G(h,a)$ the collection of plans in $F_{E(h)}$ that BI produces for the input $f(h,a)$. Iterated applications of Corollary 8 show that BI yields the same output for f and for the tree where each $f(h,a)$ is replaced with $G(h,a)$. The same is true of the act $\{f(h,b) : b \in B\}_h f$. Moreover, by assumption, $f(h,b) \succ_{E(h)} f(h,a)$ iff $g_{h,b} \succ_{E(h)} g_{h,a}$ for any (hence all) $g_{h,b} \in G(h,b)$ and $g_{h,a} \in G(h,a)$. Hence, it is sufficient (and notationally simpler) to prove the result under the assumption that every $f(h,a)$ is a plan.

Denote by $f^0, \dots, f^K$ and $\bar{f}^0, \dots, \bar{f}^{\bar{K}}$ the runs of BI with inputs $f^0 = f$ and $f^0 = \{f(h,b) : b \in B\}_h f$ respectively. By Remark A.1, it is wlog to assume that, in both runs, BI begins by replacing reducible actions at histories that follow $h_{\lambda(h)-1}$; clearly, no reduction will occur at h and at histories that follow h, because every action therein corresponds to a plan by assumption. Notice that f and $\{f(h,b) : b \in B\}_h f$ are identical at histories that follow $h_{\lambda(h)-1}$, but do not weakly follow h:[27] hence, after ℓ iterations, the two runs will produce trees f^ℓ and $\bar{f}^\ell = \{f(h,b) : b \in B\}_h f^\ell$; furthermore, either $a(h)$ is reducible in such trees, or the continuation tree corresponding to $a(h)$ in f^ℓ and $\bar{f}^\ell$ are actually plans, and indeed $f^\ell(h_{\lambda(h)-1}, a) = \bar{f}^\ell(h_{\lambda(h)-1}, a)$.

In the latter, case, $f^\ell = \bar{f}^\ell$, so Remark A.1 implies that $f^K \sim_E \bar{f}^{\bar{K}}$, and the claim follows. Otherwise, it is wlog to assume that BI replaces $a(h)$ at the $(\ell+1)$-th iteration in both runs; but notice that the *same* modified acts will be constructed in Step 2 (denoted f' there) for both f^ℓ and $\bar{f}^\ell$: in particular, the continuation plans selected at $h = [h_{\lambda(h)-1}, (a(h), E(h))]$ will correspond to one or more actions in the set B. Hence, either BI terminates with the same output in Steps 1 or 3 (so $\ell = K = \bar{K}$), or $f^{\ell+1} = \bar{f}^{\ell+1}$, in which case, again, Remark A.1 ensures that the output for the two runs will be identical. Hence, $f \sim_E \{f(h,b) : b \in B\}_h f$.

Next, consider Axiom 3.9, and let f, h, a be as in the statement of the latter. Note that, by assumption, for every non-singleton $E' \in \mathcal{F}_f(h,a)$ and $a' \in A_f([h,(a,E')])$, $f([h,(a,E')],a')$ is a plan, so either $f(h,a)$ is a plan, or a is reducible.

If $f(h,a)$ is a plan, then (i) and (ii) hold because $M_f(h,a) = \{f(h,a)\}$, so the two acts under consideration in (i) and (ii) actually coincide. Otherwise, for (i), apply BI to $f(h,a)$. Since the DM is indifferent among all continuation plans at histories following h, the act f' constructed in Step 2 coincides with $f(h,a)$; moreover, BI will terminate in Step 3 and output $M_{f(h,a)}(\emptyset, *) = M_f(h,a)$. The required indifference then follows from the assumption that the DM is indifferent between a tree and any element of the BI output for that tree (the "extension assumption"). For (ii), apply BI to f; since a is reducible, it is wlog to assume that the algorithm begins by replacing a at h. Again, the act f' constructed in Step 2 coincides with f; in Step 4, the act $\{\{f(h,a') : a' \in A_f(h) \setminus \{a\}\} \cup M_f(h,a)\}_h f$ is obtained from f; the next iteration then

[27] That is, histories of the form $[h_{\lambda(h)-1}, a(h), E']$, with $E' \neq E(h)$.

begins. But this is precisely the act appearing in the r.h.s of the indifference in (ii), so Remark A.1 implies that this run of BI for f and any run of BI for the modified act will yield the same output. Again, by the extension assumption, this implies that the required indifference will hold. ■

A.2.3 Proof of Theorem 5 (SP)

Proof: (1) ⇒ (2): Define $\{\mathbf{B}_E\}_{E \in \mathrm{NT}}$ by letting $\mathbf{B}_E(A) = f \in A$, where $A \in I_E(\succcurlyeq_E)$ and f is the plan whose existence is asserted in Axiom 3.10. Note that then $A_h g \sim_E \{\mathbf{B}_E(A)\}_h g$ for all suitable g and h.

It must first be shown that, in any run $f^0, \ldots, f^K$ of SP, $f^k \sim_E f^{k-1}$ for all $k > 0$; then, it will be argued that, starting with f^K, an additional iteration of SP yields a plan g that satisfies the required property.

Consider an arbitrary non-terminal iteration $k \in \{1, \ldots, K\}$ (if SP terminates immediately, there is no such iteration, of course). As in the proof of Theorem 4, at every iteration k, the input f^{k-1} is denoted f and the output f^k is denoted f'; thus, it must be shown that $f \sim_E f'$ if the algorithm does not terminate in Steps 1 or 3.

Since the algorithm does not terminate in the current iteration, there must be a reducible action a at a history h in Step 1. Now consider Step 2. For $m = 1, \ldots, n$, Axiom 3.8 implies that $f^m \sim_E \{f(h^m, b') : b' \in M_m\}_{h^m} f^{m-1}$. Furthermore, the definition of $\mathbf{B}_{E(h)}$ ensures that $\{f(h^m, b') : b' \in M_m\}_{h^m} f^{m-1} \sim_E \{\mathbf{B}_{E(h^m)}(\{f(h^m, b') : b' \in M_m\})\}_h f^{m-1} = f^m$. Thus, $f^{m-1} \sim_E f^m$, and therefore $f \sim_E f'$ as required. Observe that $f'(h, a)$ is now a plan.

Conclude that $f = f^0 \sim_E f^K$. Now consider the terminal iteration of SP. If SP terminates in Step 1, its output equals its input f^K, which must be a plan, as there are no reducible actions. Otherwise, SP terminates in Step 3; but if $h = \emptyset$, then $f' = f'(\emptyset, *)$ is a plan, and, arguing as above, $f^K \sim_E f'$. In either case, the proof is complete.

(2) ⇒ (1): Consider Axiom 3.8 first, and let f and h be as in the statement of the latter. For every $a \in A_f(h)$, denote by $g_{h,a}$ the plan in $F_{E(h)}$ that SP produces for the input $f(h, a)$. Corollary 8 implies that SP yields the same output for f and for the tree where each $f(h, a)$ is replaced with $g_{h,a}$. The same is true of the act $\{f(h, b) : b \in B\}_h f$. Moreover, by assumption, $f(h, b) \succ_{E(h)} f(h, a)$ iff $g_{h,b} \succ_{E(h)} g_{h,a}$. Hence, as in the proof of Theorem 4, it is sufficient to prove the result under the assumption that every $f(h, a)$ is a plan.

Denote by $f^0, \ldots, f^K$ and $\bar{f}^0, \ldots, \bar{f}^{\bar{K}}$ the runs of SP with inputs $f^0 = f$ and $f^0 = \bar{f} = \{f(h, b) : b \in B\}_h f$ respectively. As in the proof of Theorem 4, assume that SP replaces actions at histories that follow $h_{\lambda(h)-1}$, and let ℓ be the lowest index such that either the continuation trees corresponding to $a(h)$ are plans and coincide, or $a(h)$ is reducible in both f^ℓ and $\bar{f}^\ell =$

$\{f(h,b) : b \in B\}_h f^\ell$. As in the above proof, the result follows immediately from Remark A.1 in the former case; otherwise, assume that $a(h)$ is considered next. Notice that the *same* modified acts (denoted f' there) will be constructed in Step 2 for both f^ℓ and $\bar{f}^\ell$: in particular, if $E(h) = E^m \in \mathcal{F}_f(h_{\lambda(h)-1}, a(h)) = F_{\bar{f}}(h_{\lambda(h)-1}, a(h))$, the same set M_m will be constructed for f and $\bar{f}$, because $M_m \subset B$, and consequently the same plan $\mathbf{B}_{E(h)}(\{f(h,b) : b \in M_m\})$ will be selected. Hence, either SP terminates (with the same output) at the $(\ell+1)$-th iteration, or $f^{\ell+1} = \bar{f}^{\ell+1}$, in which case, again, Remark A.1 ensures that the output for the two runs will be identical. Hence, $f \sim_E \{f(h,b) : b \in B\}_h f$.

Next, consider Axiom 3.10; let $E \in \text{NT}$ and $A \subset F^p_E$ (finite and non-empty). Also consider $E' \subsetneq E$, $g = (E', H, x)$, and $h \in H$ with $E(h) = E$.

Consider the trees $g' = A_h g$ and

$$g'' = \Big\{\mathbf{B}_{E(h)}\Big(\{f \in A : \forall f' \in A, f \not\succ_{E(h)} f'\}\Big)\Big\}_h g.$$

As above, it is wlog to assume that, for both g' and g'', SP begins by replacing reducible actions at histories that follow $h_{\lambda(h)-1}$, but do not weakly follow h; since g' and g'' coincide at such histories, the trees obtained in the course of the two SP procedures, say $\bar{g}'$ and $\bar{g}'$", must also coincide.

Notice that $a(h)$ is now reducible in $\bar{g}'$; it may or may not be reducible in $\bar{g}''$ (because $\bar{g}''(h_{\lambda(h)-1}, a)$ may now be a plan).

Suppose it is. Then, it is wlog to assume that SP next replaces the reducible action $a(h)$ at $h_{\lambda(h)-1}$ for both $\bar{g}'$ and $\bar{g}''$. Since $\bar{g}'$ and $\bar{g}''$ coincide at histories other than h, Step 2 is identical for these trees at all histories $[h_{\lambda(h)-1}, (a(h), E']$, for $E' \in \mathcal{F}_g(h_{\lambda(h)-1}, a(h))$ *other than* $E(h)$. Furthermore, at h, Step 2 does *nothing* for $\bar{g}''$, and replaces the choices available at h in $\bar{g}'$ precisely with $\mathbf{B}_{E(h)}(\{f \in A : \forall f' \in A, f \not\succ_{E(h)} f'\})$. This implies that, at the end of the iteration with input $\bar{g}'$ and $\bar{g}''$, SP will output the same tree. Remark A.1 then implies that $g' \sim_E g''$, as required.

If $a(h)$ is not reducible in $\bar{g}''$, then $g''(h_{\lambda(h)-1}, a)$ is a plan. This means that a non-degenerate action set is available only at h in $\bar{g}'$; by the argument given above, one iteration of SP for the latter tree must yield precisely $\bar{g}''$, and the result follows.

For the uniqueness claim, consider $E \in \text{NT}$ and $A \in I_E(\succcurlyeq_E)$. Choose $E' \subsetneq E$ non-empty, $g = (E', H, x)$, $h \in H$ with $E(h) = E$. It was just shown that (2) implies that Axiom 3.10 must hold, and in particular that, for a belief $\mathbf{B}_E$, f can be taken to be $\mathbf{B}_E(A)$ (there are no dominated trees in A by assumption): that is, $A_h g \sim_{E'} \{f\}_h g$. If there is another belief system $\{\mathbf{B}'_E\}_{E \in \text{NT}}$ that satisfies (2), let $f' = \mathbf{B}'_E(A)$; then the argument just given implies that also $A_h g \sim_{E'} \{f'\}_h g$. This proves the claim. ■

A.3 Backward Induction for MEU preferences

A.3.1 Proof of Proposition 6 (Full Bayesian updating

Proof: Observe first that a non-empty event $E \subset \Omega$ is essential iff $\min_{q \in C} q(E) \in (0,1)$. To see this, consider $x, x' \in X$ with $x \succ x'$. Assume wlog that $u(x) = 1$ and $u(x') = 0$. Then $x \succ xEx' \succ x'$ iff $1 > \min_{q \in C} q(E) > 0$, as needed.

(1) $\Rightarrow$ (2): Axiom 4.2 implies that $f \sim_E x$ iff $fEx \sim x$. [Suppose $f \sim_E x$; then $f \succcurlyeq_E x$, so $fEx \succcurlyeq x$, and furthermore it cannot be the case that $fEx \succ x$, for otherwise $f \succ_E x$ by the last part of Axiom 4.2: thus, $fEx \sim x$. Conversely, suppose $fEx \sim x$. Then $fEx \succcurlyeq x$, so $f \succcurlyeq_E x$, and again it cannot be the case that $f \succ_E x$: thus, $f \sim_E x$.] The result then follows if it can be shown that $fEx \sim x$ iff $u(x) = \min_{q \in C} \int_E u \circ f \, q(d\omega|E)$.

To see this, first introduce the following simplifying notation: for all acts $g \in F^a_\Omega$ or $g \in F^a_E$, and all $q \in C$, let $q(g) = \int_\Omega u \circ g \, q(d\omega)$ and $q(g|E) = \int_E u \circ g \, q(d\omega|E)$; note that $q(E) > 0$ because E is essential. Also let $q_f \in \arg\min_{q \in C} q(f|E)$ and $q_{fEx} \in \arg\min_{q \in C} q(fEx)$. Then, $fEx \sim x$ iff

$$u(x) = q_{fEx}(fEx) = q_{fEx}(E) q_{fEx}(f|E) + [1 - q_{fEx}(E)] u(x),$$

i.e. iff $u(x) = q_{fEx}(f|E)$, because $q_{fEx}(E) > 0$. Now, by the choice of q_{fEx}, $q_f(fEx) \geq q_{fEx}(fEx)$; if $fEx \sim x$, this implies $q_f(f|E) \geq u(x) = q_{fEx}(f|E)$, and hence $u(x) = q_f(f|E)$, as required. Conversely, note that, by the choice of q_f, $q_{fEx}(f|E) \geq q_f(f|E)$; if $u(x) = q_f(f|E)$, then clearly also $q_f(fEx) = u(x)$, so by the choice of q_{fEx}, $q_{fEx}(fEx) \leq u(x)$. But this implies $q_{fEx}(f|E) = u(x)$, because $q_{fEx}(E) > 0$.

(2) $\Rightarrow$ (1): clearly, $\succcurlyeq_E$ is complete and transitive, so it remains to be shown that Axiom 4.2 holds. The last part of the proof that (1) $\Rightarrow$ (2) shows that $u(x) = q_f(f|E)$ holds iff $fEx \sim x$; by Eq. (5), this implies that $f \sim_E x$ holds iff $fEx \sim x$. Now suppose that $f \succ_E x$; by Eq. (5), this implies $u(x) < q_f(f|E)$; hence, a fortiori $q_{fEx}(f|E) > u(x)$, so $q_{fEx}(fEx) > u(x)$ because $q_{fEx} > 0$: that is, $fEx \succ x$. Finally, suppose that $fEx \succ x$, so $q_{fEx}(fEx) > u(x)$ and a fortiori $q_f(fEx) > u(x)$; since $q_f(E) > 0$, this implies $q_f(f|E) > u(x)$, which, by Eq. (5) implies $f \succ_E x$. ■

A.3.2 Proof of Theorem 7 (BI with full Bayesian updating)

Proof: (1) $\Rightarrow$ (2): since every $\succcurlyeq_E$ is complete and transitive and satisfies Reduction, Assumptions 4.1 and 4.2 hold, and Axiom 4.2 hold, Algorithm 3 coincides with Algorithm 1, i.e. BI; in particular, the inequality in Step 2 is $\hat{f}(h^m, b') \succcurlyeq_{E_m} \hat{f}(h^m, a')$, which by Reduction is equivalent to $f(h^m, b') \succcurlyeq_{E_m} f(h^m, a')$; similarly, the definition of $M_{f'}(h, a)$ in Step 3 coincides with the preference-based one provided in the preceding section. Theorem 4 then implies that $f \sim_E g$

whenever g is an ouptut of Algorithm 3; Reduction implies that $g \sim_E \hat{g}$, and Proposition 6 yields the expression for $u(x)$.

(2) $\Rightarrow$ (1): note first that, if f is an act, then Algorithm 3 terminates in Step 1 with output f itself; the assumption in (2) then implies that $f \sim x$, where $u(x)$ is the conditional MEU evaluation of f. Thus, conditional preferences restricted to F_E^a are satisfy (2) in Proposition 6, so they are complete and transitive; furthermore, Axiom 4.2 holds. Similarly, if f is a plan, f is again the output of Algorithm 3, and the specification of $u(x)$ implies that $f \sim_E \hat{f}$: that is, Reduction holds. Therefore, Algorithm 3 coincides with BI, so Theorem 4 ensures that Axioms 3.8 and 3.9 also hold. ■

References

[1] Robin Cubitt. Rational dynamic choice and expected utility theory. *Oxford Economic Papers*, 48(1):1–19, January 1996. New Series.

[2] Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

[3] Larry Epstein. An axiomatic model of non-bayesian updating. University of Rochester, March 2004.

[4] Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, 61:1–22, 1993.

[5] Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113:1–31, 2001.

[6] Paolo Ghirardato. Revisiting Savage in a conditional world. *Economic Theory*, 20:83–92, March 2002.

[7] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

[8] Itzhak Gilboa and David Schmeidler. Updating ambiguous beliefs. *Journal of Economic Theory*, 59:33–49, 1993.

[9] S. Goldman. Consistent plans. *Review of Economic Studies*, 47:533–537, 1980.

[10] Faruk Gul and Wolfgang Pesendorfer. The revealed preference theory of changing tastes. *Review of Economic Studies*, forthcoming.

[11] Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.

[12] Peter J. Hammond. Consistent plans, consequentialism, and expected utility. *Econometrica*, 57:1445–1449, 1989.

[13] Eran Hanani and Peter Klibanoff. Updating multiple-prior preferences. Kellogg Graduate School of Management, July 2004.

[14] Christopher Harris. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica*, 53(3):613–628, May 1985.

[15] Jean-Yves Jaffray. Dynamic decision making with belief functions. In Yaeger et al., editor, *Advances in the Dempster-Shafer Theory of Evidence*. Wiley, 1994.

[16] Edi Karni and David Schmeidler. Atemporal dynamic consistency and expectded utility theory. *Journal of Economic Theory*, 54:401–408, 1991.

[17] Peter Klibanoff. Dynamic choice with uncertainty aversion. mimeo, Kellogg School of Management, March 1995.

[18] Peter Klibanoff and Emre Ozdenoren. Subjective recursive expected utility. mimeo, Kellogg School of Management, 2004.

[19] David M. Kreps. A representation theorem for "preference for flexibility". *Econometrica*, 47(3):565–577, May 1979.

[20] David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46:185–200, 1978.

[21] 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.

[22] Martin J. Osborne and A. Rubinstein. *A Course on Game Theory.* MIT Press, Cambridge, MA, 1994.

[23] 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.

[24] C. P. Pires. A rule for updating ambiguous beliefs. mimeo, 2001.

[25] Rakesh Sarin and Peter Wakker. Dynamic choice and non-expected utility. *Journal of Risk and Uncertainty*, 17:87–119, 1998.

[26] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

[27] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.

[28] Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, New Jersey, 1976.

[29] Marciano Siniscalchi. Bayesian updating for general maxmin expected utility preferences. Discussion Paper n. 1366, CMS-EMS, Northwestern University: `http://kellogg.northwestern.edu/research/math`, 2001.

[30] R.H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1955-1956.

[31] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London and New York, 1991.

[32] Tan Wang. Conditional preferences and updating. *Journal of Economic Theory*, 108:286–321, 2003.