Diachronic Uncertainty and Equivalence Notions for ETL Models of Extensive Form Games

Alistair Isaac
Stanford University

Abstract

Diachronic uncertainty, uncertainty about where an agent falls in time, poses interesting conceptual difficulties. Although the agent is uncertain about where she falls in time, nevertheless, she can only be uncertain at a particular moment in time. This conceptual paradox can be resolved by providing an equivalence notion between models with diachronic uncertainty and models with synchronic uncertainty. The former are interpreted as capturing the causal structure of a situation, while the latter are interpreted as capturing its epistemic structure. We consider some of the properties of models for epistemic temporal logic which make them a suitable formalism for investigating such equivalence notions. We conclude with a simple example.

1. Introduction

Philosophers and Game Theorists have become increasingly interested in problems of diachronic uncertainty. In particular, if the agent knows at one state in a decision problem that at a later state she will forget or otherwise lose awareness of where she is in time, how should the agent compute appropriate actions and / or beliefs? In the game theory literature, the paradigmatic case of such forgetfulness is the Absent-Minded Driver ([Pi97]); in the philosophical literature, much discussion has centered around Sleeping Beauty ([El00]). Conceptually, however, agents can only be uncertain at a point in time; in other words, *all uncertainty is synchronic*. Any realistic model of a decision making agent should describe the succession of epistemic states through which the agent passes. Each one of these states will be synchronic, in the sense that it occurs at a distinct point in time, although these synchronic uncertainties may be uncertainties about where the agent falls in time. Given a specification of a decision problem involving diachronic uncertainty, we may ask: (i) How can we convert this into a problem involving only synchronic uncertainties? (ii) How should probabilities be assigned within the new information partitions? This document will focus on the properties of extensive form games when interpreted as ETL models as considerations towards an answer to (i).

2. Interpreting ETL Models

In order to make these questions more precise, we must work within a unified framework. In the game theory literature, all modeling of such problems uses the formalism of extensive form games. In the philosophical literature, although a vanilla Bayesianism lurks in the background of the debate, no one formalism dominates discussion; a crucial ingredient to the points of contention, however, is the use of propositions which can change truth value through time (in particular, "self-locating" propositions, which refer indexically to the agent's position in the temporal structure of the world). Epistemic temporal logics lie at a happy meeting ground between these two approaches. Syntactically, epistemic temporal languages are powerful enough to express uncertainty about where the agent falls in the temporal order. Semantically, the models of epistemic temporal logics are rich enough to include extensive form games as a special case. Furthermore, epistemic temporal logics have natural probabilistic extensions. In this section, we characterize pertinent subsets of the space of epistemic temporal models.

Epistemic temporal models are forests partitioned into equivalence classes for each agent. The interpretation of these partitions is that the agent is unable to distinguish between worlds in a partition. We refer to these as uncertainty partitions or information sets. Given a set of events Σ, Σ^* is the set of strings over Σ. Elements of Σ^* are called histories, states, or worlds. A set $\Pi \subseteq \Sigma^*$ is a protocol if it is closed under finite prefixes. So, a protocol Π is just a forest, and if Π contains the empty set, it is a tree. Call the set of agents A. With each agent $i \in A$, we identify an equivalence relation $\sim_i$. These equivalence relations partition the nodes of Π into sets of worlds which are indistinguishable for agent i.

DEFINITION 1: An **ETL frame** is a tuple $\langle \Sigma, \Pi, \{\sim_i\}_{i \in A} \rangle$ where Σ is a set of events, Π is a protocol, and for each $i \in A$, $\sim_i$ is an equivalence relation on Π.

DEFINITION 2: An **ETL model** is a tuple $\langle \Sigma, \Pi, \{\sim_i\}_{i \in A}, V \rangle$ where Σ, Π, and $\{\sim_i\}_{i \in A}$ are an ETL frame and *V* is a valuation function from the set of atomic formulae *At* into the power set of Π, $V : At \longrightarrow 2^{\Pi}$.

Conceptually, we can think of an ETL model as a specification of the causal structure of the world (the ordering of possible events), which is then decorated with epistemic relations. We are interested, however, in what will happen if we prioritize epistemic structure rather than causal structure. What happens if we insist that models characterize the sequence of the agent's epistemic states, even when this sequence conflicts with the sequence of events? This is the case with diachronic uncertainty. If an agent *i* is uncertain at t_2 whether the time is t_1 or t_2, then some events which the agent considers possible will not in fact be possible (namely, those events which can only immediately follow t_1). Our goal is to consider this conceptual transformation within the framework of epistemic temporal models by considering equivalence classes of ETL models with respect to intuitively motivated notions of situation equivalence.

The space of ETL models is quite rich, and characteristics of its fine structure have been charted in [vB08] and [vB06]. [vB08] characterizes the subset of ETL models which are equivalent to models for dynamic epistemic logic (DEL). From [vB06] we know that this fragment of ETL preserves some nice computational properties (in particular, so long as we limit ourselves to a future modality which can only see ahead one step in time (this is the essence of DEL), we preserve decidability). [vB08] distinguishes two types of DEL-generated protocols: uniform protocols and state-dependent protocols. These notions are of interest for our purposes in the constraints they place on permissable uncertainty partitions $\sim_i$. In order to emphasize this aspect of the situation, we define four classes of ETL models.

DEFINITION 3: An ETL model $\langle \Sigma, \Pi, \{\sim_i\}_{i \in A}, V \rangle$ is

(i) **state-dependent** *iff* there is no general restriction on the events that can occur after any history

(ii) **agent-dependent** *iff* for any agent i, event e, and histories h, h', if $h \sim_i h'$ and $he \in \Pi$, then $h'e \in \Pi$

(iii) **cardinality-dependent** *iff* for any agent i and histories h, h', if $h \sim_i h'$, then $|\{h'' \in \Pi | \exists e\, (h'' = he)\}| = |\{h'' \in \Pi | \exists e\, (h'' = h'e)\}|$

(iv) **uniform** *iff* if $p \in At$ is a *precondition* of event *e* and $h \in V(p)$, then $he \in \Pi$

In standard DEL models, the events possible at a world are characterized by a function $E : At \longrightarrow \Sigma$. If $E(p) = e$, then the proposition *p* represents a *precondition* of the event e, and e is possible at any world $h \in V(p)$. We write $pre(e)$ for $E^{-1}(e)$, *i.e.* the set of preconditions of e. The notion of a state-dependent DEL protocol generalizes this idea by replacing the function from atomic formulae into the space of events with a function from histories *h* into the space of events (*i.e.* from Π into Σ). This is the appropriate interpretation of a state-dependent ETL model: it is a model in which the events following a given history are not constrained in any systematic way by other features of the model. Agent-dependent, cardinality-dependent, and uniform models are all special cases of state-dependent models where the function from histories into events is somehow constrained. These notions will help us distinguish various definitions of extensive form games in the following section. Before embarking on that discussion, let us expand our repertoire with some further notions from [vB08].

DEFINITION 4: An ETL model $\langle \Sigma, \Pi, \{\sim_i\}_{i \in A}, V \rangle$ satisfies

(i) **strong synchronicity** *iff* for all histories h, h', if for some agent i, $h \sim_i h'$, then $len(h) = len(h')$, where $len(h)$ is just the number of events in h

(ii) **weak synchronicity** *iff* for all histories h, h', if for some agent i, $h \sim_i h'$, then h is not a proper prefix of h'[1]

(iii) **perfect recall** *iff* for all histories h, h' and events e, e', if $he \sim_i h'e'$, then $h \sim_i h'$

(iv) **local no miracles** *iff* for all histories h, h', h'', h''', agents i, and events e, e', if $he \sim_i h'e'$, $h' \sim^* h''$, and $h'' \sim_i h'''$, then $h''e \sim_i h'''e'$, where $\sim^*$ is the reflexive transitive closure of the $\sim_i$ relations

[vB08] shows that the class of ETL models generated from uniform DEL protocols is just that which satisfies strong synchronicity, perfect recall, local no miracles, and local epistemic bisimulation invariance, and the class of ETL models generated from state-dependent DEL protocols is just that which satisfies propositional stability, strong syn-

[1] "Weak synchronicity" does not appear in [vB08], but it will be helpful in our discussion of games below. Strong synchronicity implies weak synchronicity, but not vice versa. In some situations, we can transform a model satisfying weak synchronicity into one satisfying strong synchronicity by simply introducing dummy nodes which bring asynchronous uncertainty partitions into sync (*c.f.* the introduction of "'dummy' chance moves with one alternative" in [Ku53], 51). For an example of a game which cannot be brought into synchrony using this method, see [Pi97], example 6.

chronicity, perfect recall, and local no miracles.[2] Tomohiro Hoshi (this conference) has investigated these distinctions in more detail for the subset of DEL known as public announcement logic (PAL).

3. Uncertainty in Extensive Form Games

Much of the game theory literature on uncertainty has focused on uncertainty about other players' moves. Since one usually assumes that players alternate turns in extensive form games, this uncertainty can be modeled via weakly synchronic equivalence relations for each agent. Furthermore, modeling choices have been constrained by conceptual analysis of the notion of indistinguishability itself: *what constraint appropriately captures the idea that the agent cannot distinguish between two states?*

[Th52]and [Ku53] define information partitions in extensive form games such that two constraints are met. First, no two worlds in the information partition may lie on the same branch. Second, at each world in the partition, the cardinality of the set of potentially occurring events must be the same. [Pi97] drops the first constraint, in order to allow for forgetfulness, yet strengthens the second constraint by stipulating not just that the cardinality of the set of possible events be the same for each world in an information set, but that the set of possible events be *identical* for each world. Thus, [Th52]and [Ku53] define models which are cardinality-dependent and satisfy weak synchronicity, while [Pi97] defines models which are agent-dependent.

These constraints are motivated by the idea that an information set models a situation in which an agent must *act*, although she does not know the current state of the world. Actions are just distinguished events, events caused by some particular agent. If different (or different numbers of) actions are available to an agent at two nodes in the game tree, then the agent can use her knowledge of which actions are available to her to distinguish these histories from each other. Thus, if two states of the world are indistinguishable to an agent, then the agent must have the same actions available to her at each one. Agent-dependency and cardinality-dependency are attempts to capture this intuition.

As noted above, if an agent is uncertain between two worlds at different points in time t_1 and t_2, then it must be the case that different events are possible at the two worlds. However, it may nevertheless be the case that the agent has the same set of available actions. In the Absent Minded Driver example, a man has left a bar drunk and forgets while driving home whether he has already made his turn or not. The problem is usually modeled with an information set including two indistinguishable intersections. The driver must pass through these intersections in sequence, so he will encounter them at different times. Thus, there must be some events possible at one which are not possible at the other. However, in terms of actions, the driver only has two options: turn or go straight. So, if our model only includes the actions available to the agent, excluding other events, it will satsify agent-dependence.

The constraints on modeling in the game theory literature are motivated by extrinsic considerations. In thinking about agents performing actions in a game, and what it might mean for an agent to be uncertain between possible states of play, concept analysis dictates that either agent-dependency or cardinality-dependency constrain permissable models. In state-dependent ETL models, however, we have as modeling tools both a valuation function *V* and an event function *E*. If we retain the notion of preconditions at the conceptual level, then we can characterize a situation in which an agent believes e to be possible, when in fact it is not.

DEFINITION 5: The **possible events** $E_i^P(h)$ for an agent i at a history h in an uncertainty partition $I = \{w_1, ..., h, ..., w_n\}$ are just those events e such that $\bigcup I \subseteq V(pre(e))$

If we are in a state-dependent ETL model $\left\langle \Sigma, \Pi, \{\sim_i\}_{i \in A}, V \right\rangle$, then there is no constraint that for any $h \in \Pi$, and agent i, $E_i^P(h) = \{e \in \Sigma | he \in \Pi\}$. In other words, the set of possible events *from the agent's perspective* need not equal the *actual* possibilities allowed by the model. Of course, this move deflates the role of preconditions; they no longer play a structural role in constraining the model, but merely act as a bookkeeping device for tracking agent expectations.

4. An Example: the Absent Minded Driver

Perhaps the simplest example of a game with diachronic uncertainty is the Absent Minded Driver (*fig. 1*).

The driver begins at Ø and drives straight. He passes through two intersections, w_1 and w_2. At each intersection, he can either continue to drive straight, or turn. If he turns at the first intersection, he arrives in the bad part of town, *B*. If he turns at the second intersection, he arrives home, *H*, as desired. If the driver continues straight through both intersections, he must stay at a motel, *M*. It is stipulated, however, that the driver cannot distinguish the first and second intersections; in other words, he cannot remember whether he has turned or not.

In light of the considerations raised above, we might ask whether there is a distinct game tree, equivalent to the Absent Minded Driver *in the relevant respects*, but prioritizing epistemic states. Such a game tree would satisfy synchronicity, in line with the analysis of uncertainty as a fundamentally synchronic notion, yet would preserve *in some*

[2] We have omitted definitions of local epistemic bisimulation invariance and propositional stability as they are not discussed in the sequel.

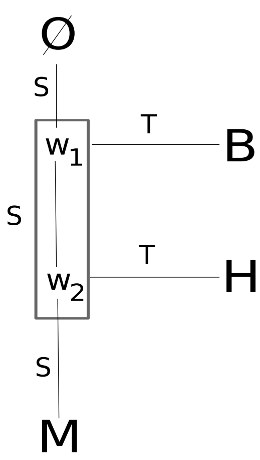

Figure 1. The Absent Minded Driver

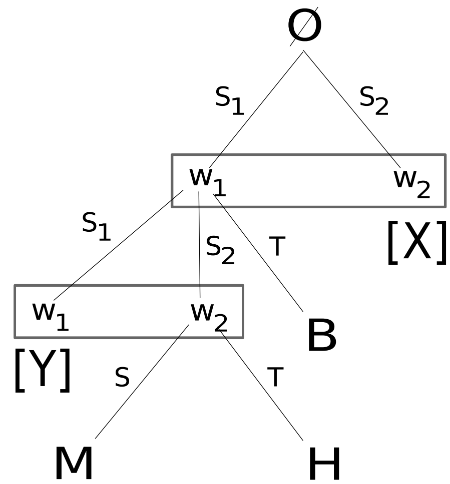

Figure 2. The *Epistemic* Absent Minded Driver

sense the causal structure of the original model. We might call such a pair of models *epistemically equivalent*.

DEFINITION 6: Two ETL models M_1 and M_2 are **epistemically equivalent** *iff*
(i) all agents $i \in A_1 \bigcap A_2$ pass through uncertainty states in the same order (with possible duplications) in M_1 and M_2
(ii) all events $e \in \Sigma_1 \bigcap \Sigma_2$ occur in the same order (with possible duplications) in M_1 and M_2

Consider, for example, the ETL model depicted in *figure 2*. In this model, the agent passes through the same sequence of epistemic states as in the Absent Minded Driver; the uncertainty partition from *figure 1* has merely been duplicated to capture the fact that the agent will experience it at two distinct points in time. There are two questionable modeling choices here, however. First, what is the significance of moves s_1 and s_2? Second, how should one interpret [X] and [Y]? s_1 and $s_2$represent the epistemic disjunct between the choice of a single action (go straight), and the two resulting possibilities, w_1 and w_2. Rather than consider these as two distinct moves, or events, we may instead wish to include separate moves by the driver and a chance player. The driver chooses *s*, but *chance* (or, perhaps it would be better to call him *confusion*) plays to increase the possibilities the agent countenances. This "move" must be interpreted as an *epistemic*, rather than *physical*, event: *the event of forgetting*.

The worlds above [X] and [Y] are those which the agent erroneously believes possible. We have several modeling options available to us here, though we consider only three. First, we might simply leave these as terminal nodes. Second, we may replace [X] and [Y] with unitary chance moves to some distinguished world; this world might be interpreted as an impossible state. Both these strategies would capture the agent's error, or the physical impossibility of any events occurring at [X] and [Y]. Both these options fail to connect with game theoretic models, however. The first, because game theoretic models never allow uncertainties over terminal states; the second, because transitions to an "impossible" world would involve adding a new terminal state, but one without any coherent notion of payoff attached to it.

A third option would connect quite nicely with the game theoretic literature, in particular [Th52]. Thompson defines equivalence classes of extensive form games with respect to the corresponding strategic game. He suggests four transformations on game trees which preserve strategic form. Any two extensive form games which share a strategic form can be transformed into each other via some sequence of these four transformations. Since Thompson only considers models which satisfy weak synchronicity, the Absent Minded Driver does not fall within his paradigm. One strategy for dealing with [X] and [Y] suggested by Thompson's transformations is simply to copy the game tree from under the other node in the uncertainty partition to the position under the "impossible" node. Conceptually, we might interpret this as capturing the fact that the *actual* possibilities are the same from both states in the driver's uncertainty partition as he is only *actually* at one of them. If we apply this strategy plus that described above for adding moves by a chance player, we derive *figure 3*. In *figure 3*, the sequence of epistemic states in *figure 1* is preserved, as is the sequence of actual events. We have had to add a chance move, interpreted as the epistemic process of forgetting, but doing so has allowed us to produce a model which is susceptible to the transformations described in [Th52].

Final remark: the strategy described for transforming the model in *figure 1* to the model in *figure 3* will not work in all cases. Consider, for example, *crossed* ETL models.

> DEFINITION 7: An ETL model $\left\langle \Sigma, \Pi, \{\sim_i\}_{i \in A}, V \right\rangle$ is **crossed** *iff* there exists an agent $i \in A$ and histories h, h', h'', and h''' with $h \neq h''$ such that $hh', h''h''' \in \Pi$, $h \nsim_i h'$, $h'' \nsim_i h'''$, $hh' \sim_i h''$, and $h \sim_i h''h'''$

Example 6 of [Pi97] is a crossed model. If one attempts to implement the transformation strategy described above on a crossed ETL model, one will produce an infinite tree which cycles through the two partitions $\{h, h''h'''\}$ and $\{hh', h''\}$. It remains to be seen what precise class of constraints on ETL models characterizes just those susceptible to the above described transformation. At the very least, such models must be *uncrossed*.

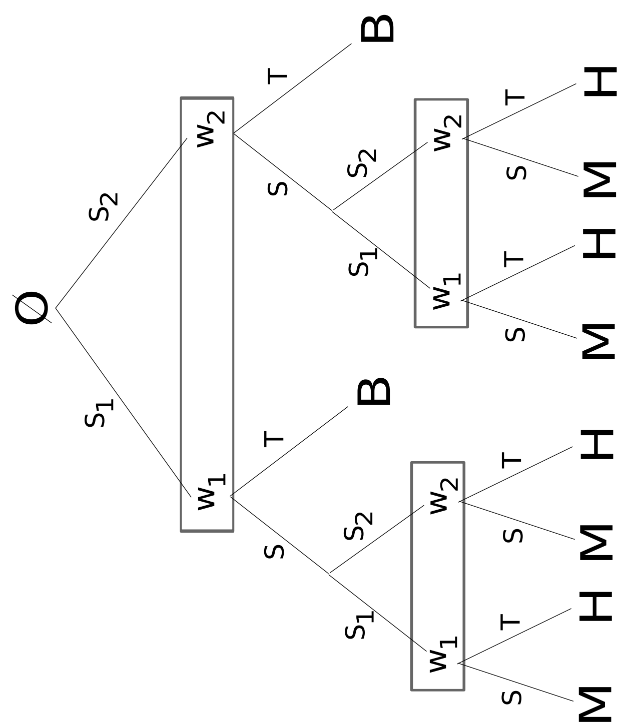

Figure 3. The *Epistemic* **Absent Minded Driver (***final***)**

References

[El00] Elga, Adam (2000) "Self-locating Belief and the Sleeping Beauty Problem," in *Analysis* 60.2; 143-7.

[Ku53] Kuhn, Harold W. (1953) "Extensive Games and the Problem of Information," reprinted from *Contributions to the Theory of Games II* in *Classics in Game Theory* (1997), Harold W. Kuhn (ed.); 46-68.

[Pi97] Piccione, Michele and Ariel Rubinstein (1997) "On the Interpretation of Decision Problems with Imperfect Recall," in *Games and Economic Behavior* 20; 3-24.

[Th52] Thompson, F. B. (1952) "Equivalence of Games in Extensive Form," RAND Memo RM-759, in *Classics in Game Theory* (1997), Harold W. Kuhn (ed.); 36-45.

[vB06] van Benthem, Johan and Eric Pacuit (2006) "The Tree of Knowledge in Action: Towards a Common Perspective," in *Proceedings of Advances in Modal Logic* (AiML 2006), King's College Press, Guido Governatori, Ian Hodkinson and Yde Venema (eds.).

[vB08] van Benthem, Johan, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit (2008) "Merging Frameworks for Interaction," unpublished manuscript available online at staff.science.uva.nl/˜johan/deletl-march2008.pdf