THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Synchronizing Diachronic Uncertainty

OPEN ACCESS

Synchronizing Diachronic Uncertainty

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, this uncertainty can only obtain at a particular moment in time. We resolve this conceptual tension by providing a transformation from models with diachronic uncertainty relations into "equivalent" models with only synchronic uncertainty relations. The former are interpreted as capturing the *causal structure* of a situation, while the latter are interpreted as capturing its *epistemic structure*. The models are equivalent in the sense that agents pass through the same information sets in the same order, In this paper, we investigate how such a transformation may be used to define an appropriate notion of equivalence, which we call *epistemic equivalence*. Although our project is motivated by problems which have arisen in a variety of disciplines, especially philosophy and game theory, our formal development takes place within the general and flexible framework provided by epistemic temporal logic.

Keywords: Diachronic uncertainty, Epistemic temporal logic, The absent-minded driver, Absent-mindedness, Sleeping beauty

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 (Piccione and Rubinstein, 1997); in the philosophical literature, much discussion has centered around the Sleeping Beauty problem (Elga, 2000). The crucial observation behind our contribution to these debates is that 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) Is there an "equivalent" problem involving only synchronic uncertainties? (ii) What is the relevant notion of equivalence here? (iii) How do

the probabilities in this second model relate to those in the original model?

We will provide specific answers to (i) and (ii), and conclude with a brief sketch of a solution to (iii). Although a number of different formal frameworks may be used to investigate this issue, we will restrict ourselves to epistemic temporal logic, one of the most rich and flexible formalisms for modeling diachronic uncertainty. After introducing ETL in section 2 and discussing a fundamental ambiguity in the interpretation of ETL models, we will return to extensive form games in section 3 in order to motivate our concerns. Section 4 introduces the main concepts of the paper, *synchronic completion* and *epistemic equivalence*. In section 5, we provide a brief formal development, incluing a preservation theorem: the fragment of the ETL language containing only knowledge and future operators (but not their duals) is preserved under synchronic completion. We conclude in section 6 with a discussion of the issues that arise when extending synchronic completion to a probabilistic context.

2. Interpreting ETL models

We wish to investigate the relationship between synchronic and diachronic uncertainty. In order to make our investigation more precise, we must work within a unified formal 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, no one formalism dominates discussion; a linchpin assumption of the debate, however, is that propositions 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 logic* ((Parikh and Ramanujam, 1985), (Fagin, et al., 1995), (Hodkinson and Reynolds, 2006)) lies at a happy meeting ground between these two approaches. Syntactically, epistemic temporal languages are rich enough to express uncertainty about where the agent falls in the temporal order. Semantically, both epistemic temporal logic and extensive form games use partitioned tree structures as models. Only minor adjustments are needed to embed extensive form games into the space of ETL models. After a brief discussion of the basic structure of ETL models, we examine a fundamental ambiguity within the formalism: models within ETL can be interpreted as capturing either the *causal* or the *epistemic* structure of a situation.

2.1. Models for epistemic temporal logic

Fix a countable set Σ of events and a finite set of agents $\mathcal{A}$. We write Σ^* for the set of finite sequences of elements in Σ. A *history* is an element in Σ^*. Given $h \in \Sigma^*$ and $e \in \Sigma$, we write he for the history h followed by the event e. Also given $h, h' \in \Sigma^*$, we write $h \preceq h'$ if h is a finite prefix of h'. If $h \preceq h'$ and $h \neq h'$, we write $h \prec h'$. We denote the length of a given sequence σ by $\mathsf{len}(\sigma)$.

Definition 2.1 (ETL models) An *ETL model* is a tuple $(H, \sim, V)$ where:

- H is a subset of Σ^* that is closed under finite prefixes,

- $\sim$ is a function that assigns to an agent in $\mathcal{A}$ an equivalence relation on H.

- V is a function that assigns a subset of H to an atomic proposition in in the set of atomic propositions At. ◁

Since $\sim_i$ is an equivalence relation, it induces a partition on H. Given an ETL model $\mathcal{H} = (H, \sim, V)$ and a history $h \in H$, we denote by $[h]_i^{\mathcal{H}}$ the equivalence class under $\sim_i$ to which h belongs. When there is no confusion about which model is under discussion, we omit the superscript. We will refer to these equivalence classes as "information sets" or "uncertainty partitions."

Strictly speaking, ETL models only contain information about temporal ordering, they do not come with an explicit notion of absolute time. In many applications, however, it is helpful to assume that all histories are *synchronized, i.e.* that nodes at the same depth represent possibilities which might obtain at the same moment. The rest of the paper assumes synchronization in this sense. Since this assumption allows us to investigate the temporal relationship between nodes on different histories, it greatly simplifies the ensuing discussion. It also suggests a simple definition of diachronic uncertainty.

Definition 2.2 (Diachronic uncertainty) Let $\mathcal{H} = (H, \sim, V)$ be an ETL model and h a history in H. An equivalence class $[h]_i$ is *diachronic* if there is some $h' \sim_i h$ such that $\mathsf{len}(h') \neq \mathsf{len}(h)$. $\mathcal{H}$ is diachronic if there is some $h \in H$ such that $[h]_i$ is diachronic. Otherwise, $\mathcal{H}$ is *synchronic.* ◁

2.2. Causal models *vs.* epistemic models

Usually, ETL models are interpreted as capturing the causal structure of the world: if $he, hf \in H$ then the events e and f are *possible* after the history h. Divergent histories represent alternate possibilities allowed by the causal structure of the world. Each node is a choice point at which nature, or some agent, may act, causing some new event to occur. This causal tree is then decorated with relations $\sim_i$ for each agent i; these relations partition the set of worlds H into equivalence classes. The interpretation here is that agent i cannot distinguish between any two worlds falling within the same partition as induced by $\sim_i$. Call a model interpreted in such a fashion a *causal model*.

In a causal model, an information set may include nodes of different depths (or histories of different lengths). In such cases, the model exhibits *diachronic uncertainty*. An agent experiences diachronic uncertainty if she is uncertain about where she falls in time. From a conceptual standpoint, there is an important distinction to be made here. If an agent is uncertain between two histories h and h', but $h \not\prec h'$ and $h' \not\prec h$, then she visits the relevant information set only once. As such, even though the agent exhibits diachronic uncertainty, it is easy to conceive of this uncertainty as occurring at a single point in time. For example, I may hear an alarm clock go off, but be uncertain about whether my partner set it five minutes fast the night before or not. Then I am uncertain about where I fall in time, but I am uncertain at a single moment and not in any way forgetful. If it was me who who set the alarm, then I am forgetful, but my state upon hearing the alarm is still visited only once (this is sometimes called *imperfect recall*). If, however, $h \sim h'$, and $h \prec h'$ or $h' \prec h$, then I am *absent-minded*. Absent-mindedness is conceptually problematic in causal models because the agent passes through the same information set more than once (at h and h'), exhibiting the same state of uncertainty at multiple points in time (this issue is discussed more thoroughly in section 3.2).

By constructing a tree which captures causal possibilities and decorating it with epistemic relations, we have prioritized causal structure in our modeling choices. Suppose instead that we prioritize epistemic structure. What happens if we insist that models characterize the sequence of epistemically possible states rather than the sequence of causally possible states? Then $he, hf \in H$ would imply that the *epistemic states* e and f are possible after the history h. We might call a model interpreted in this way an *epistemic model*. An epistemic model of an absent-minded agent would explicitly distinguish the epistemic states left implicit in the causal model.

Note that the formalism of ETL does not distinguish causal from epistemic structure. Whether we take a model as a causal model or an epistemic model is purely a matter of interpretation. However, given our conceptual analysis, we may wish to restrict these interpretations to distinguished classes of ETL models. In particular, the above observation, that diachronic uncertainty is always uncertainty at a point in time, implies a crucial desideratum for characterizing the subset of ETL models which may be interpreted as capturing epistemic structure:

Desideratum A: Uncertainty relations $\sim_i$ must only occur between histories of the same length.

Before we can specify exactly which subset of ETL models are best interpreted as epistemic, however, we must consider a second desideratum, this one originating in the game theory literature.

3. Uncertainty in extensive form games

When playing a game, an agent has moves available to her at various decision points. If an agent is uncertain at which point she falls in a game, she must not be able to distinguish between the nodes in her uncertainty partition on the basis of which moves are available to her. In section 3.1 we introduce some useful definitions for understanding how this consideration has been modeled in the literature. In section 3.2, we consider a specific model of causal structure, the absent-minded driver, and discuss how it could be transformed into a model of the corresponding epistemic structure.

3.1. Game Theoretic Indistinguishability

We begin with several definitions. Although these definitions are motivated by models of uncertainty in extensive form games, we offer them in terms of ETL models for the sake of consistency with the rest of our discussion.

Definition 3.1 (strong synchronicity) An ETL model $(H, \sim, V)$ exhibits *strong synchronicity* if, for any agent i, and histories h, h', $h \sim_i h'$ implies $\mathsf{len}(h) = \mathsf{len}(h')$. ◁

Definition 3.2 (weak synchronicity) An ETL model $(H, \sim, V)$ exhibits *weak synchronicity* if for any agent i and histories h, h', $h \sim_i h'$ implies $h \nprec h'$. ◁

Strong synchronicity is the formal equivalent of desideratum A, and coincides exactly with the definition of synchronicity given in Definition 2.2. We list it again to contrast it with weak synchronicity, which obtains whenever there is no absent-mindedness. Although weak synchronicity has been the more popular constraint in the game theory literature (see discussion below), our insight that uncertainty is always uncertainty *at a point in time* demands that we interpret synchronicity in the strong sense.

The following two definitions may seem unduly strong to the logician, but one must bear in mind that in extensive form games, one can only be uncertain between nodes at which one is able to act.

Definition 3.3 (agent-dependent) An ETL model $(H, \sim, V)$ is *agent-dependent* if for any agent i, event e, and histories h, h', if $h \sim_i h'$ and $he \in H$, then $h'e \in H$ ◁

Definition 3.4 (cardinality-dependent) An ETL model $(H, \sim, V)$ is *cardinality-dependent* if for any agent i and histories h, h', if $h \sim_i h'$, then $|\{h'' \in H \mid \exists e(h'' = he)\}| = |\{h'' \in H \mid \exists e(h'' = h'e)\}|$, where $|x|$ denotes the cardinality of the set x. ◁

Information sets were initially introduced into extensive form games in order to model a player's ignorance about the moves of her opponents. A player may be uncertain about how many other players have played since her last move, what moves they made, or even her own past moves (*i.e.* she may exhibit *imperfect recall*). However, these players always know where they fall in the temporal structure of the game, in particular, they remember whether they have played or not, even if they can't remember the move they made. Models of games involving such agents are just those which satisfy weak synchronicity. However, from the game theoretic viewpoint, demanding weak synchronicity is not enough. We must ensure that there are no clues in the structure of the game which undermine the coherence of the uncertainty partitions. What additional constraint will address this worry?

(Kuhn, 1953) defines information sets in extensive form games such that two constraints are met. First, no two worlds in the information set may lie on the same branch (*i.e.* players do not forget if they have moved). Second, at each world in the partition, the cardinality of the set of potentially occurring events must be the same. (Piccione and Rubinstein, 1997) drops the first constraint in order to allow for absent-minded players; however, it 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, (Kuhn, 1953) defines

models which are cardinality-dependent and satisfy weak synchronicity, while (Piccione and Rubinstein, 1997) 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. Therefore, 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. These concerns motivate our second desideratum for characterizing the subset of ETL models which can be interpreted as capturing epistemic structure:

Desideratum B: The same (number of) events must follow any world in an uncertainty partition.

Formally, we realize this desideratum by constraining attention to agent-dependent models. However, as will become clear, this definition captures a slightly different constraint within our models than that which motivated Kuhn and Rubinstein. In particular, we must consider the discrepancy between an agent's beliefs and the actual state of the world: the agent may consider events possible which are not in fact possible. In order to make the distinction between possible events, and events the agent believes are possible, we need to distinguish worlds from the events which produced them. We illustrate this point with an example.

3.2. EXAMPLE: *the absent-minded driver*

In the example of the absent-minded driver (Piccione and Rubinstein, 1997), a man leaves a bar drunk and forgets while driving home whether he has already reached 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 satisfy agent-dependence.
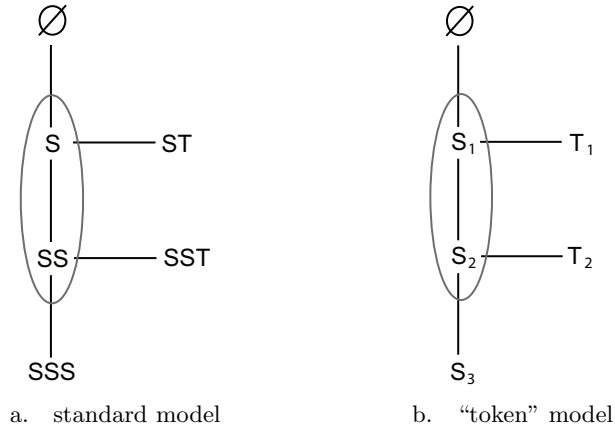
a.   standard model   b.   "token" model

Figure 1. The absent-minded driver

The driver desires to turn at the second intersection, as this is the road to his home. He wishes not to turn at the first intersection as it leads to the bad part of town. If he misses both turns and simply drives straight, he must spend the night at a motel. If we follow the conventions of epistemic temporal models as described in section 2.1, our tree model is generated by the appropriate sequences of events (see figure 1.a). So, for example, the node associated with reaching the bad part of town involves the events "go straight" and "turn", so it will be characterized by the history st, while the node associated with staying in a motel involves the event "go straight" occurring three times, so it will be characterized by the history sss.

However, if we attempt to construct a model of the driver's sequence of epistemic states, we discover that the state of uncertainty he exhibits concerning the intersection which he is passing through occurs *twice*: once at the intersection leading to the bad neighborhood, a second time at that leading to his home. If we attempt to construct a model of the epistemic structure of this situation, then, we must include the worlds corresponding to histories s and ss twice. How can we capture the idea that the same world, or epistemic possibility occurs twice? We cannot simply duplicate the relevant strings as our underlying set theory satisfies extensionality. Furthermore, length of string is what determines depth in the tree and, correspondingly, time of occurrence. Thus, there is no room here to say that the same world occurs twice at different times, as this would require a single string to be two distinct lengths. What we need is a trick to circumvent these facts about our formalism, a way to identify worlds other than considering the entire string associated with them.

The issue here is really a confusion between tokens and types. We have considered a set of event types, allowing the same type of event ("go straight") to occur over and over. Really, however, when an agent is uncertain, she cannot distinguish between event tokens—she confuses *this* instance of going straight with *that* instance of going straight. In order, then, to construct a clear model of the epistemic structure of the absent-minded driver, we need to first convert our standard model of the problem into one involving only token events. This amounts to the stipulation that each event occurs only once. We call ETL models satisfying this condition *token-event models* (Definition 5.1, below).

In the case of the absent-minded driver, we have two event types, but five event tokens. If we reconstruct the model using only tokens, instead of one event $s =$ "go straight", we now have three events $s_1 =$ "go straight from the bar"; $s_2 =$ "go straight at the first intersection"; and $s_3 =$ "go straight at the second intersection". Likewise for the two turns. Nodes in the tree are still defined by sequences of events. Now, however, the history associated with arriving at the bad part of town is not captured by the string st, but by s_1t_1, and that associated with arriving at the motel corresponds not to sss, but instead to $s_1s_2s_3$. Since we have stipulated that events occur only once in the tree, we have effectively produced a distinct label for each history. If we wish, we may identify each history with its distinct final (*i.e.* rightmost) event. Figure 1.b reproduces the token-event model for the absent minded driver labeled in this fashion.

What is the epistemic structure of this situation? Well, our analysis should satisfy desiderata A and B above. According to desideratum A, uncertainty should only be synchronic. This means that the agent experiences two distinct states of uncertainty. The first occurs when he passes through the first intersection, the second when he passes through the second intersection. Both states have the same content; however, the agent cannot tell whether the previous event was "go straight from the bar" or "go straight at the first intersection." *So, these event-tokens are* types *when considered from the epistemic standpoint.* By beginning with a token-event model isomorphic to the original model, we can identify multiple nodes in the generated model with a single, token state of the world by simply repeating the corresponding token-event.

According to desideratum B, the same set of events should be possible from all histories in an information set. Possibility, now, is interpreted as epistemic possibility—*i.e.* what epistemic states are possible from a given epistemic state? The answer to this question will still be constrained by the causal structure of the world. At the first intersection, for example, the agent is in a state of uncertainty concerning which intersection he is passing through. Two epistemic states are possible
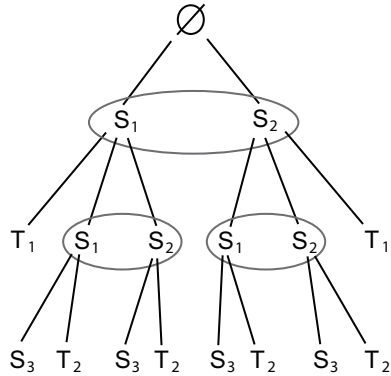
Figure 2. The absent-minded driver (*epistemic*)

in the agent's future: one is identical to that he just left, a state of uncertainty about which intersection he is passing through; the second is that induced by the agent's arrival in the bad part of town.

With these two desiderata in mind, consider figure 2. This tree exhibits only synchronic uncertainties. For any worlds in an information set, the exact same possibilities follow. Furthermore, if we connect the labeling of nodes with that in figure 1.b, we can see that the agent passes through the same epistemic states in the same order as in the absent-minded driver. We claim that figure 2 characterizes the epistemic structure of the absent-minded driver.

Notice that our labeling trick in figure 1.b allows us to distinguish possible events from those the agent merely believes are possible. The subsequent worlds from any node in figure 2 represent *epistemic possibilities*, *i.e.* states of knowledge at which the agent may find himself at the next point in time. The labels on worlds (*i.e.* the rightmost event on the corresponding history) keep track of the events the agent (perhaps wrongly) believes are possible. This brings out the conceptual difference between the role of agent dependence in (Piccione and Rubinstein, 1997) and in our own discussion. For Piccione and Rubinstein, agent dependence captures the constraint that the agent must believe he has the same actions available at each world in an information set or else he could distinguish between them. For us, agent dependence captures the constraint that from a single epistemic state, there is only a single uniform set of possible epistemic states which may occur. An information set represents a single epistemic state, so from each world in the information set, the same set of possibilities must follow. Thus, the same formal constraint does distinct conceptual work in our approach.

4. Epistemic equivalence

We now have the apparatus to answer our motivating question: when are two models epistemically equivalent? Usually, models are defined as equivalent with respect to some specific criterion they both satisfy. For example, (Thompson, 1952) investigates the idea that two extensive form games are equivalent if they share the same strategic form. We would like to find a concept that will play the same role as strategic form did for Thompson, an epistemic criterion that can be used to partition the space of ETL models into equivalence classes.

Let's begin by singling out the class of models which satisfy desiderata A and B above. These models are in *synchronic normal form*:

Definition 4.1 (Synchronic normal form) An ETL model is in *synchronic normal form* if

- — it exhibits *strong synchronicity*

- — it is *agent-dependent* ◁

For every countable diachronic model, there exists an isomorphic token-event model. Furthermore, for every token-event model, there is a unique model in synchronic normal form such that agents pass through the same information sets in the same order. We call such a normal form model the *synchronic completion* of the corresponding token-event model. For example, figure 2 depicts the synchronic completion of the model in figure 1.b. The intuitive characterization of this concept (which we formalize in section 5) can be expressed as follows:

Synchronic completion For any ETL token-event model $\mathcal{H}$, the *synchronic completion* of that model, $SC(\mathcal{H})$, is just the unique model in synchronic normal form such that agents pass through the same information sets in the same order as in the original model.

It is crucial to note here that the concept of synchronic completion applies only to token-event models. This is explained by the conceptual analysis in section 3.2. To summarize: when considering diachronic uncertainty from the epistemic standpoint, we realize that agents are not confused between event types, but rather event tokens. Thus, we can only make sense of the epistemic structure of a situation if we first relabel all event types in its model with corresponding tokens. These new events are tokens from the causal perspective, but will be treated as types in the corresponding epistemic model. Formally speaking, there is no distinction; the difference is purely one of interpretation.

We now have all the apparatus to define our key concept: *epistemic equivalence*. The idea here is that two models are epistemically equivalent if isomorphic token-event models share the same synchronic completion (up to isomorphism); intuitively, agents pass through the same epistemic states in the same order in both models.

Epistemic Equivalence Two token-event ETL models $\mathcal{H}$ and $\mathcal{G}$ are *epistemically equivalent* if the synchronic completions of $\mathcal{H}$ and $\mathcal{G}$ are isomorphic.

The concept of epistemic equivalence partitions the space of all ETL models into equivalence classes of models which share synchronic completions (either directly or via an isomorphic token-event model). There is a unique (up to isomorphism) model in synchronic normal form which represents the epistemic structure of each equivalence class. The preservation theorem proved below is surprising precisely because it illustrates a breakdown between this intuitive notion of epistemic equivalence and the basic epistemic language.

5. Formalization

Now we will make these ideas more precise. We first formalize the basic notions described in the previous sections and then introduce the model transformation that captures the idea of synchronic completion. Once we have formalized this concept, we will be able to show that a certain class of ETL formulas are preserved under synchronic completion.

5.1. BASIC NOTIONS

Fix a countable set of events Σ. For simplicity, we only deal with the single agent case and omit indices for the equivalence relation $\sim$ in ETL models. Given an ETL model $\mathcal{H} = (H, \sim, V)$ and a history h, we denote by $[h]^{\mathcal{H}}$ the equivalence class under $\sim$ to which h belongs. When there is no confusion about which model is under discussion, we omit the superscript. Given $h \in \Sigma^*$, we denote by $\Sigma(h)$ the set of events constituting h and by $r(h)$ the right-most element of h. Given a set x, we denote the cardinality of x by $|x|$.

Definition 5.1 (Token-event ETL models) An ETL model $\mathcal{H} = (H, \sim, V)$ is a *token-event ETL model* if

  1. for every $h \in H$, $\mathsf{len}(h) = |\Sigma(h)|$ and

2. for every $h, h' \in H$, $\Sigma(h) \cap \Sigma(h') - \Sigma(g) = \emptyset$

where g is the longest (possibly empty) initial segment of h and h', *i.e.* $g \preceq h$ and $g \preceq h'$ plus there is no g' such that $g \prec g'$, $g' \preceq h$, and $g' \preceq h'$. ◁

We next demonstrate that consideration can be restricted to token-event models without undermining the generality of our results. ETL models, $\mathcal{H} = (H, \sim, V)$ and $\mathcal{H}' = (H', \sim', V')$, are *isomorphic*, if there is a one-to-one surjective map $f : H \to H'$ such that, for all $h, h' \in H$,

  — $h \sim h'$ iff $f(h) \sim' f(h')$ and

  — $h \in V(p)$ iff $f(h) \in V'(p)$.

Given the definition, it is clear that, for any ETL model $\mathcal{H} = (H, \sim, V)$, if $|H| \leq |\Sigma|$, then there is a token-event ETL model isomorphic to $\mathcal{H}$. This ensures that we can always find an appropriate token-event model for every (countable) ETL model.

5.2. Synchronic completion

We now give the formal definition of synchronic completion of token-event ETL models. We need some definitions. We denote by $h_{(n)}$ ($1 \leq n \leq \mathsf{len}(h)$) the initial segment of h of length n and by h_n the n-th element of h. Given a set $H \subseteq \Sigma^*$, we define $r(H) := \{r(h) \mid h \in H\}$. Given a number $n \geq 0$, we denote by H_n the set of elements h in H such that $\mathsf{len}(h) = n$. The next definition provides an operation that "synchronizes" an ETL models up to a finite level n.

Definition 5.2 (Synchronization up to n) Let $\mathcal{H} = (H, \sim, V)$ be a token-event ETL model. Let $n \geq 1$. We define $SC^n(H)$, $SC^n(\sim)$, and $SC^n(V)$ by induction up to n as follows:

1. $SC^1(H) := H_1 \cup \{r(h) \mid \exists h' \in H_1 : h \in H \text{ and } h \in [h']\}$:

2. $h \in SC^{k+1}(H)$ iff   there are $h' \in SC^k(H)$ and $e \in \Sigma$ such that:

    — $h = h'e$ and

    — $\exists g \in H_k : r(\{h'' \mid (h, h'') \in SC^k(\sim)\}) = r([g])$ and $ge \in H_{k+1}$

3. $(h, h') \in SC^1(\sim)$ iff $h, h' \in SC^1(H)$ and there are $g, g' \in H$ such that:

    — $h = r(g)$,

   – $h' = r(g')$ and

   – $g \sim g'$.

4. $(he, h'e') \in SC^{k+1}(\sim)$ iff   (i) $h, h' \in SC^k(H)$, (ii) $e, e' \in \Sigma$, (iii) $h = h'$, and (iv) there are $g, g' \in H$ such that:

   – $g \sim g'$,

   – $e = r(g)$ and

   – $e' = r(g')$

5. $h \in SC^1(V)(p)$ iff $h \in SC^1(H)$ and $\exists g \in H : h = r(g)$ and $g \in V(p)$

6. $h \in SC^{k+1}(V)(p)$ iff $\exists g \in H : r(g) = r(h)$ and $g \in V(p)$ $\triangleleft$

This definition makes use of the following fact: each event e that constitutes a token-event ETL model corresponds to the unique history in that model whose right-most element is e. Conditions 1 and 2 ensure that the k-th level of the synchronized ETL model $SC^k(H)$ contains both the nodes that are in the k-th level *and* the nodes that are indistinguishable from them in the original model $\mathcal{H}$. Intuitively, if an agent is uncertain between a node at time k and a node at time $j \neq k$, we erase the diachronic uncertainty relation and add a copy of the node from time j at time k. This new node is "the same" as the original one in the sense that it satisfies the same propositional letters as the original; this is ensured by conditions 5 and 6. Finally, conditions 3 and 4 guarantee that the new indistinguishability relation $SC^n(\sim)$ respects the original $\sim$ in the sense that nodes in the new model are indistinguishable only if the corresponding nodes in the original model are indistinguishable.

    We can now state the definition of synchronic completion.

Definition 5.3 (Synchronic Completion) Let $\mathcal{H} = (H, \sim, V)$ be a token-event ETL model. The synchronic completion of $\mathcal{H}$, $SC(\mathcal{H}) = (SC(H), SC(\sim), SC(V))$, is defined as follows:

   – $SC(H) := \bigcup_{k=1}^{\infty} SC^k(H)$

   – $SC(\sim) := \bigcup_{k=1}^{\infty} SC^k(\sim)$

   – $SC(V)(p) := \bigcup_{k=1}^{\infty} SC^k(V)(p)$ $\triangleleft$

    A brief inspection of the above definitions reveals that the synchronic completion of an ETL model satisfies the two desiderata, synchonicity and agent-dependency, discussed in the previous section. The following is a simple consequence of the definitions; the proof of the proposition is given in appendix 7.1.

Proposition 5.4 *The synchronic completion of any ETL model is in synchronic normal form.*

Next, we present the result that underscores the intuition that a token-event model $\mathcal{H}$ and its synchronic completion are *epistemically equivalent* in the sense that an agent's experience is the same in both models. To do so, we need the following definition.

Definition 5.5 (Epistemic history along h in $\mathcal{H}$) Let $\mathcal{H} = (H, \sim, V)$. The epistemic history $\mathcal{I}(\mathcal{H}, h) = (\mathcal{I}_1(\mathcal{H}, h), \ldots, \mathcal{I}_{\mathsf{len}(h)}(\mathcal{H}, h))$ is a sequence such that, for all $1 \le k \le \mathsf{len}(h)$ $\mathcal{I}_k(\mathcal{H}, h) = r([h_{(k)}])$. ◁

This definition formalizes the concept of an agent's "experience" by equating it with the sequence of information sets an agent passes through. (This analysis is closely related to that offered in (Piccione and Rubinstein, 1997).) The following theorem states that, given a token-event ETL model $\mathcal{H}$ and a history h in $\mathcal{H}$, an agent passes through the same sequence of information sets along h in the synchronic completion of $\mathcal{H}$ as in $\mathcal{H}$. The proof of the theorem is given in appendix 7.2.

Theorem 5.6 *Let $\mathcal{H} = (H, \sim, V)$ be a token-event ETL model. For all $h \in H$,*
$$\mathcal{I}(\mathcal{H}, h) = \mathcal{I}(SC(\mathcal{H}), h).$$

Now that we have Definition 5.5 and Theorem 5.6, we may be confident that the previously defined concept of epistemic equivalence captures an intuitive notion of epistemic sameness. For completeness, we restate the definition here.

Definition 5.7 (Epistemic Equivalence) Two token-event ETL models $\mathcal{H}$ and $\mathcal{G}$ are *epistemically equivalent* if $SC(\mathcal{H})$ and $SC(\mathcal{G})$ are isomorphic.

5.3. TRUTH-PRESERVATION UNDER SYNCHRONIC COMPLETION

By Definition 5.7, an ETL model and its synchronic completion are epistemically equivalent. Together with Proposition 5.4, we see that the operation of synchronic completion transforms a (token-event) ETL model into an epistemically equivalent ETL model that is synchronic and agent-dependent.

Although synchronic completion preserves epistemic structure, interpreted as the sequence of information sets that the agent passes through, it changes other aspects of the original ETL model's structure. Natural questions we can ask here are: What properties of an ETL

structure are preserved by synchonic completion? What properties are not? Examining these questions will shed some light on the nature of knowledge under diachronic uncertainty.

First we discuss what structures are preserved by the operation. For this, we investigate the kinds of formulas whose truths are preserved by synchronic completion. We emphasize here that this is another advantage of working within the ETL framework.

Definition 5.8 (Preserved formulas) We define the set $Pres$ of formulas inductively as follows:

$$\varphi \ ::= P \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid K\varphi \mid F\varphi$$

where P is a propositional formula. ◁

The intended readings of $K\varphi$ and $F\varphi$ are respectively "φ is known" and "There is some future sequence of events after which φ." The duals, $\hat{K}$ and $\hat{F}$, of K and F are defined in the standard way. The intended readings of $\hat{K}\varphi$ and $\hat{F}\varphi$ are "φ is considered possible" and "After any future sequence of events, φ."

Let $\mathcal{H} = (H, \sim, V)$ be a token-event ETL model. The truth of formulas in $Pres$ is defined in a standard way in ETL. Here we list the cases for the two modalities.

$$\mathcal{H}, h \models K\varphi \quad \text{iff} \quad \forall h' \in H : h \sim_i h' \Rightarrow \mathcal{H}, h' \models \varphi$$
$$\mathcal{H}, h \models F\varphi \quad \text{iff} \quad \exists \sigma \in \Sigma^* : h\sigma \in H \text{ and } \mathcal{H}, h\sigma \models \varphi$$

We can show the following result. The proof is given in appendix 7.3.

Theorem 5.9 (Preservation under synchronic completion) *Let φ be a formula in $Pres$. Then, for all $h \in H$,*

$$\mathcal{H}, h \models \varphi \quad \Rightarrow \quad SC(\mathcal{H}), h \models \varphi.$$

Theorem 5.9 demonstrates that formulas in which the knowledge operator K and the future operator F appear *but not their duals* are preserved under synchronic completion. More colloquially, knowledge of future possibility and the future possibility of knowledge are preserved. Many types of formulas are not preserved, however; why these particular ones? The crucial point to remember here is that we have reinterpreted what a history signifies. In a model interpreted as causal, nodes signify states of the world; in a model interpreted as epistemic, nodes signify states of information. With the change in our interpretation of the models, comes a corresponding change in our interpretation of modal operators.

F, for example, no longer corresponds to physical possibility, but to epistemic possibility. As an example, consider again the absent-minded driver, as illustrated in figure 1.b. Call this $\mathcal{H}$; $SC(\mathcal{H})$ is illustrated in figure 2. The absent-minded driver can reason about the histories available to him if we stipulate that each world has a distinguished proposition letter true at that world and no other. For the sake of discussion, let's consider the second intersection (s_2), the driver's home (t_2), and the hotel (s_3). Stipulate that $\mathcal{H}, s_2 \models a$, $\mathcal{H}, t_2 \models b$, and $\mathcal{H}, s_3 \models c$, but no other worlds satisfy a, b, or c.

Now, consider a formula such as KFb; intuitively, the driver knows that in some future he arrives home. This formula is in $Pres$ and is true at the root node of both $\mathcal{H}$ and $SC(\mathcal{H})$. However, this formula could mean two distinct things depending upon whether we interpret the model causally or epistemically. If interpreted causally, the formula says that there is a future way *the world* could be such that the driver arrives home, and he knows it. Interpreted epistemically, the formula says there is a future way *the driver's epistemic state* could be such that it contains the information he is home, and he knows this. What theorem 5.9 tells us is that even in cases of diachronic uncertainty, the epistemic and causal interpretations of this formula have the same truth value. This is important because only when causal and epistemic possibility match up does the driver have knowledge *that can help him decide how to act.*

In contrast, consider the formula $\varphi = K\hat{F}(a \rightarrow \hat{F}(b \vee c))$. Since this formula contains the dual of F, it is not in $Pres$. Investigation reveals that $\mathcal{H}, \emptyset \models \varphi$ but $SC(\mathcal{H}), \emptyset \not\models \varphi$. How can we make sense of this? Interpreted causally, the formula says that the driver knows that in every future way the world could be, if he is at the second intersection, then in every future way the world could be he is either at home or at the hotel. Interpreted epistemically, the formula says that the driver knows that in every future epistemic state he could be in, if he is at the second intersection, then in every future epistemic state he is either at home or at the hotel. In this case, the epistemic and causal interpretations come apart. The reason, of course, is that the driver considers possible scenarios which are not physically possible; in particular, he considers it possible that he is at the second intersection when he is only at the first, and thus there are future epistemic possibilities in which he is not at home or the hotel, but instead in the bad part of town (t_1).

This is a case where knowledge comes apart from *knowledge which can guide action*, but only because it is knowledge which quantifies over all futures. By quantifying over all futures we can uniquely identify temporal states, but that is exactly what an agent who is diachronically

uncertain cannot do. So, in general, such "knowledge" is not *useful knowledge* for the absent-minded agent.

6. Probability and sleeping beauty

The "paradox" of the absent-minded driver is a puzzle about how knowledge of future absent-mindedness affects the driver's strategy. In general, however, diachronic uncertainty also poses a puzzle for the assignment of probabilities: if I am in a diachronic information set, what probability should I assign to each world I consider possible? This question does not arise in the absent-minded driver (for reasons to be discussed below), but it does arise in the problem of sleeping beauty (Elga, 2000). We conclude with a brief discussion of this problem and potential strategies for analyzing it with synchronic completion.

Sleeping Beauty voluntarily participates in a bizarre psychological experiment. On Sunday night, the experimenters flip a fair coin, then put Beauty to sleep without telling her the outcome. If the coin came up heads, they wake Beauty on Monday, but do not expose her to any information about what day of the week it is. At the end of the day, they erase her memory and put her back to sleep. On Wednesday, they wake Beauty and the experiment is over. If the coin came up tails, the experimenters wake Beauty on both Monday and Tuesday. Each day, Beauty is prevented from accessing information about the day of the week and at the end of the day, her memory is erased. On Wednesday, the experimenters wake Beauty and the experiment is over. Since Beauty knows the exact structure of the experiment, we can ask how she should assign probabilities (for example, to the outcome of the initial coin toss) once the experiment is underway.

The key feature of the sleeping beauty scenario is that Beauty experiences an event (waking) once or twice depending upon the outcome of a random process, but her diachronic uncertainty prevents her from distinguishing the single occurrence from the double occurrence case. A structurally analogous problem was introduced in (Piccione and Rubinstein, 1997), but without the elaborate backstory. Figure 3.a is analogous to the model in (Piccione and Rubinstein, 1997), but it omits moves which would allow Beauty to opt out of the experiment in accordance with the Elga scenario. The strategy taken by Elga here is analogous to that taken by (Aumann, *et al.*, 1997) with respect to the absent-minded driver. In order to analyze the structure of belief change in the absent-minded driver, (Aumann, *et al.*, 1997) removes all possibility for the driver to act, constructing a scenario called "the

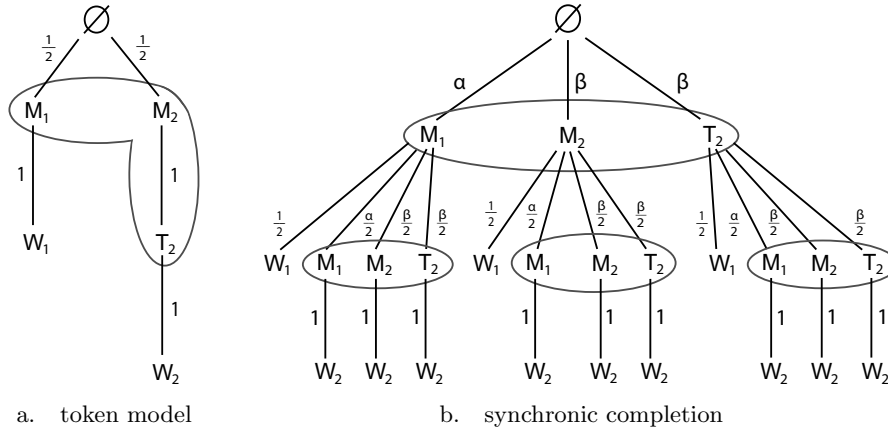a.   token model   b.   synchronic completion

Figure 3. Sleeping beauty. For *even* synchronization, $\alpha = \beta = 1/3$; for *historic* synchronization, $\alpha = 1/2$ and $\beta = 1/4$.

forgetful passenger." Likewise, Elga has removed the actions available to Beauty in order to analyze her beliefs in isolation.

Once she is awoken during the experiment, Beauty knows that there are three possible scenarios: 1. it is Monday and she won't be awoken again; 2. it is Monday and she will be awoken again; 3. it is Tuesday. Furthermore, she can easily calculate from the model the probability with which she arrives at each of these worlds. If probabilities are assigned to edges, we can find the probability that one is at a particular world by simply multiplying the values assigned to all edges leading up to that world. So, from outside the model, it is easy to see that $p(M_1) = 1/2$, $p(M_2) = 1/2$, and $p(T_2) = 1/2$. But Beauty cannot distinguish between any of these three worlds, so we must somehow normalize the probabilities within her information set.

(Piccione and Rubinstein, 1997) discusses two distinct strategies for normalizing probabilities in the sleeping beauty scenario. In the philosophical literature, the respective virtues of these two strategies have been hotly debated; much of the discussion has centered around details specific to the backstory and the evidential rules Beauty should apply in light of these details. In the context of the present framework, however, the two options for assigning probabilities in sleeping beauty are simply special cases of two general rules for extending synchronic completion to probabilistic models.

One strategy is to simply normalize over the worlds, treating them all equally. Since the probability we arrive at each of the worlds is $1/2$, this analysis assigns $p(M_1) = p(M_2) = p(T_2) = 1/3$. A second strategy normalizes separately over histories. Since there is only one world in the "heads" history, $p(M_1) = 1(1/2) = 1/2$. Since there are two

worlds with equal probabilities in the "tails" history, $p(M_2) = p(T_2) = (1/2)(1/2) = 1/4$. (Piccione and Rubinstein, 1997) calls the former strategy "consistent" and the latter strategy "Z-consistent." This terminology is somewhat loaded, however, so let's refer to evenly weighted normalization as "*even* synchronization" and the normalizaton over histories separately as "*historic* synchronization." Formal definitions for *even* and *historic* synchronization are given in appendix 7.4. The synchronic completion of sleeping beauty is given in figure 3.b.

(It should be obvious now why this debate does not arise in the context of the absent-minded driver. Both worlds in the absent-minded driver's diachronic information set fall on the same history. Since only one history is under consideration, normalizing over all worlds evenly and normalizing separately within histories produce the same result.)

Although a number of arguments have been given for either strategy, we will point out only two considerations here. If one follows the second strategy and assigns a larger probability to M_1 than M_2 and T_2, one is subject to a "money pump," *i.e.* a second agent can take advantage of the fact that Beauty will place a bet twice in one scenario and only once in the other to make bets with her such that she will always lose. Assigning probabilities in accordance with the first strategy avoids the money pump, but at the cost of an apparent inconsistency with previous beliefs. If Beauty normalizes over worlds and not histories, she seems doomed to conclude that the probability of the outcome heads changes from 1/2 to 1/3 just because she has been woken during the experiment.

The advantage of examining this debate in the context of synchronic completion is simple: the method for assigning probabilities within a diachronic information set must be explicitly stipulated in the transformation itself. This allows for the general investigation of the consequences of each strategy, without an undo dependence on the details of this specific example.

The authors suspend judgment on the "correct" transformation until after a more thorough investigation. At present, however, we lean slightly towards historic synchronization for two reasons. Although even synchronization sounds like the simpler procedure when described in ordinary language, it turns out that its formal statement as a transformation is slightly more complex than that of historic synchronization (compare Definition 7.16 and Definition 7.17). This is because during even synchronization an additional normalization procedure needs to occur to ensure the information set as a whole does not take up any more than the permitted probability mass at a given depth in the tree. In historic synchronization, this normalization is performed by the probability of the history at that depth.

Second, consider again figure 3.a and its synchronic completion, figure 3.b. The probability of reaching W_1 is clearly $1/2$. Both even and historic synchronization preserve the probability of W_1 (and indeed, the probabilities of all histories outside of diachronic information sets). But historic synchronization assigns $1/2$ to α while even synchronization assigns $1/3$ to α. As mentioned above, Beauty exhibits an apparent inconsistency with her knowledge of the past coin toss if she assigns M_1 probability $1/3$. What is not often emphasized in the literature is her apparent inconsistency with future events as well. Once Beauty emerges from the experiment, the probability that it is Wednesday and the coin came up heads is still $1/2$. Deviating from the known probabilities of both one's past *and one's future* within an information set seems unwise in principle.

Finally, we conclude with a suggestion for future research. One of the most compelling arguments for even synchronization, going all the way back to (Piccione and Rubinstein, 1997) is the money pump argument. But accepting a bet is an action, and ETL can model actions just as well as chance events. A natural move would be to build this action into the model and see how it is transformed by even and historic synchronizations. Of course, a serious attempt at exploring this possibility should move past the simple probabilistic ETL models discussed here to models which distinguish intentional actions from chance. The authors hope that the present work paves the way for future investigations along these lines.

References

Aumann, Robert, Sergiu Hart, and Motty Perry: 1997, 'The Forgetful Passenger,' *Games and Economic Behavior* **20**, pp. 117–120

Cao, Zining: 2006, 'Model Checking for Epistemic and Temporal Properties of Uncertain Agents', Z. Shi and R. Sadananda (Eds.): *PRIMA 2006*, LNAI 4088, pp. 46–58

Elga, Adam: 2000, 'Self-locating Belief and the Sleeping Beauty Problem, *Analysis* **60**:2, pp. 143–147

Fagin, R., J. Halpern, Y. Moses, and M. Vardi: 1995, *Reasoning about Knowledge*, Boston, MA; MIT Press.

Hodkinson, I. and M. Reynolds: 2006. 'Temporal Logic', in *Handbook of Modal Logic*, P. Blackburn, J. van Benthem, and F. Wolter (eds.), Amsterdam; Elsevier. pp. 655–720.

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.); pp 46–68.

Parikh, R. and R. Ramanujam: 1985, 'Distributed Processes and the Logic of Knowledge', in *Logic of Programs* (*Lecture Notes in Computer Science* **193**), London; Springer. pp. 256–268.

Piccione, Michele and Ariel Rubinstein: 1997, 'On the Interpretation of Decision Problems with Imperfect Recall, *Games and Economic Behavior* **20**, pp. 3–24

Thompson, F. B.: 1952, 'Equivalence of Games in Extensive Form, RAND Memo RM-759, in *Classics in Game Theory* (1997), Harold W. Kuhn (ed.); pp. 36–45.

7. Appendix

7.1. PROOF OF PROPOSITION 5.4

The following facts are straightforward consequences of Definition 5.2.

Observation 7.1 *For all $n \geq 1$ and $h, h' \in SC^n(H)$, $\mathsf{len}(h) = \mathsf{len}(h') = n$.*

Observation 7.2 *For every $n \geq 1$, $H_n \subseteq SC^n(H)$.*

Proposition 7.3 *Let $n \geq 1$. For every $h, g \in SC^n(H)$ and $\sigma \in \Sigma^*$, if $h\sigma \in SC^n(H)$ and $(h, g) \in SC^n(\sim)$, then $g\sigma \in SC^n(H)$.*

Proof. By straightforward induction on the length of σ, based on the inductive clause of $SC^n(H)$ in Definition 5.2.   QED

Now we show some simple but important properties of synchronic completions, which will be used also for the proofs below. Let $\mathcal{H} = (H, \sim, V)$ be an ETL model and $\mathcal{G} = (G, \approx, U)$, the synchronic completion of $\mathcal{H}$.

Proposition 7.4 $H \subseteq G$.

Proof. The claim is an immediate consequence of Definition 5.3 and Observation 7.2.   QED

Proposition 7.5 $\mathcal{G}$ *is synchronic.*

Proof. For each $h, g \in G$, if $h \approx g$, then $(h, g) \in \approx$ for some n. However, by the inductive clause for $\approx$ in Definition 5.2, we have $\mathsf{len}(h) = \mathsf{len}(g)$. QED

Proposition 7.6 *For every $h, g \in G$ and $e \in \Sigma$, if $he \in G$ and $h \approx g$, then $ge \in H'$.*

Proof. The claim follows from Proposition 7.3.   QED

Proposition 5.4 The synchronic completion of any ETL model is in a synchronic normal form.

Proof. The proposition is an immediate consequence of Propositions 7.5 and 7.6. <div style="text-align: right">QED</div>

7.2. Proof of Theorem 5.6

Theorem 5.6 Let $\mathcal{H} = (H, \sim, V)$ and $\mathcal{G} = (G, \approx, U)$ be respectively a token-event ETL model and its synchronic completion. For all $h \in H$,

$$\mathcal{I}(\mathcal{H}, h) = \mathcal{I}(\mathcal{G}, h).$$

Proof. We show by induction on n ($1 \leq n \leq \mathsf{len}(h)$) that $\mathcal{I}_n(\mathcal{H}, h) = \mathcal{I}_n(\mathcal{G}, h)$. The base case follows immediately from Proposition 7.5 and the definition of $SC^1(H)$ and $\sim^1$ in Definition 5.2.

Let $[h] = \{g \mid h \sim g\}$ and $[h]' = \{g \mid h \approx g\}$. Assume as IH that $\mathcal{I}_n(\mathcal{H}, h) = \mathcal{I}_n(\mathcal{G}, h)$. Suppose that $a \in \mathcal{I}_{n+1}(\mathcal{H}, h)$. Then there is some g such that $h_{(n+1)} \sim ga$. Now, by Proposition 7.4, we have $h_{(n)} \in G$ and in particular $h_{(n)} \in SC^n(H)$ by the definition of $SC^n(H)$ (in Definition 5.2). By IH, we have $r([h_{(n)}]) = r([h_{(n)}]')$. Furthermore, by inspecting the definition of $SC^n(\sim)$, we see $r([h_{(n)}]_n) \subseteq r([h_{(n)}]')$. So, by the definition of $SC^{n+1}(H)$, we have $h_{(n)}a \in SC^{n+1}(H)$. Finally we have $a \in \mathcal{I}_{n+1}(\mathcal{G}, h)$ by the definition of $SC^{n+1}(\sim)$. Hence $\mathcal{I}_{n+1}(\mathcal{H}, h) \subseteq \mathcal{I}_{n+1}(\mathcal{G}, h)$.

Next, suppose $a \in \mathcal{I}_{n+1}(\mathcal{G}, h)$. Then there is some g such that $h_{(n+1)} \approx ga$. By the definition of $SC^{n+1}(\sim)$, there are some l, k such that $lh_{n+1} \sim ka$. Since $\mathcal{H}$ is a token-event ETL model, l is unique. It must be the case, then, that $l = h_{(n)}$ and therefore $h_{(n+1)} \sim ka$. Thus we have $a \in \mathcal{I}_{n+1}(\mathcal{H}, h)$. Hence $\mathcal{I}_{n+1}(\mathcal{H}, h) \subseteq \mathcal{I}(\mathcal{G}, h)$. <div style="text-align: right">QED</div>

7.3. Proof of Theorem 5.9

To obtain the desired preservation result, we need to prove some facts about synchronic completions of token-event ETL models. Below let $\mathcal{H} = (H, \sim, V)$ be a token-event ETL model and $\mathcal{G} = (G, \approx, U)$, the synchronic completion $SC(\mathcal{H})$ of $\mathcal{H}$.

Observation 7.7 *Let P be a propositional formulas. Then, for every $h \in H$ and $g \in G$, if $r(h) = r(g)$, $\mathcal{H}, h \models P$ iff $\mathcal{G}, g \models P$.*

Proof. The claim immediately follows from the definition of U in Definition 5.2. <div style="text-align: right">QED</div>

Observation 7.8 *Let $K\varphi \in Pres$. If $\mathcal{G}, h \models K\varphi$, then $\mathcal{G}, g \models K\varphi$ for every g such that $h \sim g$.*

Proof. The claim is an immediate consequence of the fact that $\sim$ is an equivalence relation on G. QED

Definition 7.9 (ETL model after h) For every $h \in G$, *the model* $\mathcal{G}_h = (G_h, \approx_h, U_h)$ *after* h is defined by:

  − $g \in G_h$ iff $\exists l : lg \in G$ and $h \approx lg_1$.

  − $\approx_h = \approx \cap (G_h \times G_h)$.

  − $U_h(p) = U \cap G_h$. ◁

Proposition 7.10 *For every* $h, g, g' \in \mathcal{G}$ *such that* $\mathsf{len}(g), \mathsf{len}(g') \geq \mathsf{len}(h)$,

$$g \approx g' \;\; iff \;\; g \approx_h g'.$$

Proof. The left-to-right direction is clear. The other direction follows from the fact that, for every $g, g' \in G$, if $g \approx g'$, then $g = le$ and $g' = le'$ for some $l \in \Sigma^*$ and $e, e' \in \Sigma$ (as defined in the inductive clause of $\approx$ in Definition 5.2.) QED

Proposition 7.11 $\mathcal{G}, h \models \varphi$ *iff* $\mathcal{G}_h, h \models \varphi$.

Proof. The proof is by induction on φ. The base case is immediate from Definition 7.9. The cases for $\vee, \wedge, F$ are straightforward. For the knowledge modality case, assume $\mathcal{G}, h \models K\psi$. This means that $\mathcal{G}, g \models \psi$ for all g with $h \approx g$. By Proposition 7.10, this is equivalent to $\mathcal{G}_h, g \models \psi$ for all g with $h \approx_h g$. QED

Proposition 7.12 *Let* $F\varphi \in Pres$. *If* $\mathcal{G}, h \models F\varphi$, *then* $\mathcal{G}, g \models F\varphi$ *for every* g *such that* $h \approx g$.

Proof. Assume $\mathcal{G}, h \models F\varphi$. Then we have $\mathcal{G}, h\sigma \models \varphi$ for some $\sigma \in \Sigma^*$. By Proposition 7.11, we have $\mathcal{G}_{h\sigma}, h\sigma \models \varphi$. By Proposition 7.6, we have $g\sigma \in G$. By Proposition 7.6 and Definition 5.2, $\mathcal{G}_{h\sigma}$ and $\mathcal{G}_{g\sigma}$ are isomorphic. Thus $\mathcal{G}_{g\sigma}, g\sigma \models \varphi$. By applying Proposition 7.11, we obtain $\mathcal{G}, g\sigma \models \varphi$. QED

Theorem 5.9 Let φ be a formula in $Pres$. Then, for all $h \in H$,

$$\mathcal{H}, h \models \varphi \quad \Rightarrow \quad \mathcal{G}, h \models \varphi.$$

Proof. We show the claim by induction on φ. The base case is clear from Observation 7.7. The boolean cases, $\wedge$ and $\vee$, and the future modality case are straightforward. For the knowledge modality, assume $\mathcal{H}, h \models K\varphi$. Then we have $\mathcal{H}, h \models \varphi$ for every g with $h \sim g$. In

particular, since $\sim$ is an equivalence relation, we have $\mathcal{H}, h \models \varphi$. By IH, $\mathcal{G}, h \models \varphi$. Now we would like to show that $\mathcal{G}, g \models \varphi$ for all g with $h \approx g$.

Now we go by cases. First, assume that φ is propositional. By Definition 5.2, $r([h]) = r(\{h' \mid h \approx h'\})$. Thus, by Observation 7.7, if $\mathcal{H}, h' \models \varphi$ for all h' with $h \sim h'$, then $\mathcal{G}, g \models \varphi$ for all g with $h \approx g$.

Second, assume that φ is of the form $K\psi$ or $F\psi$. Then we are done by Observation 7.8 or Proposition 7.12.

Finally, assume that φ is not of the form in the previous two cases. Then, transform φ into a conjunctive normal form $\bigwedge \Phi_i$. Each Φ_i is a disjunction of propositional formulas and formulas of the form $K\alpha$ or $F\alpha$. Let g be such that $h \sim g$. If some non-propositional disjunct is true at h, then Φ_i is true in all g by Observation 7.8 or Proposition 7.12. Thus, assume that non-propositional disjuncts are all false. In this case, we apply the argument from the case that φ is propositional. Hence, $\mathcal{H}, g \models \bigwedge \Phi_i$ for all g with $h \sim g$. QED

7.4. SYNCHRONIC COMPLETION OF PROBABILISTIC ETL MODELS

There are a number of natural ways to extend epistemic temporal logic with probabilities. (Cao, 2006), for example, develops a sophisticated approach in which probabilities are indexed by worlds. The precise probabilistic version of ETL chosen will determine what types of transformation rule are possible for assigning new probabilities during synchronic completion. In this section, we examine a very simple probabilistic extension of ETL, one suitable for analyzing the beliefs of "passive" players in a game where all moves are made by chance.

Definition 7.13 (probabilistic ETL models) A *probabilistic ETL model* is a tuple $(H, \sim, V, \mu)$ where:

1. $(H, \sim, V)$ is an ETL model

2. $\mu : H \longrightarrow [0, 1]$ is a function assigning real numbers to histories such that,

   - $\mu(\emptyset) = 1$
   - for any $h \in H$, $\sum_{\{h' : \exists e \in \Sigma(h' = he)\}} \mu(h') = 1$ ◁

μ assigns transition probabilities to histories. The correct interpretation here is not that $\mu(h)$ tells us the probability of being at h simpliciter, but rather the probability of transitioning to h from the previous node in the tree.

Definition 7.14 (probability at time t) Given a probabilistic ETL model $(H, \sim, V, \mu)$ and a time t, the probability $P(h)$ of a history h such that $\mathsf{len}(h) = t$ is $P(h) = \Pi_{\{h' : h' \preceq h\}} \mu(h')$. $\qquad \triangleleft$

It is easy to prove that P is a probability distribution over histories of length t.

Within a synchronic information set, it is uncontroversial how to assign probabilities to worlds.

Definition 7.15 (synchronic probability) Given a synchronic probabilistic ETL model $(H, \sim, V, \mu)$ and an information set $[h]$, the probability $\bar{P}(h)$ that the agent is at history $h \in [h]$ is

$$\bar{P}(h) = \frac{P(h)}{\sum_{\{h' : h' \in [h]\}} P(h')}$$

$\qquad \triangleleft$

Next we define two strategies for extending the definition of synchronic completion to probabilistic ETL models. *Even* synchronization normalizes over all worlds in an information set. *Historic* synchronization normalizes separately over worlds on each history within an information set.

Definition 7.16 (*Even* synchronization up to n) Let $\mathcal{H} = (H, \sim, V, \mu)$ be a token-event probabilistic ETL model. Let $n \geq 1$. We define $SC^n(H)$, $SC^n(\sim)$, $SC^n(V)$, and $SC^n(\mu)$ by induction up to n as follows:

1. $SC^n(H)$, $SC^n(\sim)$, and $SC^n(V)$ are just as in Definition 5.2

2. for $h \in SC^k(H)$,

$$SC^k(\mu)(h) = \frac{\left(\sum_{\{h' : h' \in [h] \wedge h' \in H_k\}} P(h') \right) \left[\frac{P(h)}{\sum_{\{h' : h' \in [h]\}} P(h')} \right]}{\sum_{\{h' : h' \in SC^k(H) \wedge \exists e, f \in \Sigma \exists h''(h = h''e \wedge h' = h''f)\}} P(h')}$$

$\qquad \triangleleft$

The most important constraint on μ is that the μ of all histories leading out of single node must sum to 1. In synchronic completion, however, we attach new edges to histories at each level. So, we must renormalize over all histories which share a parent at each depth k, this is the role of the denominator. The left half of the numerator ensures that a diachronic information set is normalized with respect to already synchronic events.

In other words, we need to ensure that the probability mass assigned to a particular information set at depth k doesn't change when we add in the (formerly) diachronic nodes. Finally, the right side of the numerator performs the evenly weighted normalization over all nodes in an information set.

Definition 7.17 (*Historic* synchronization up to n) Let $\mathcal{H} = (H, \sim, V, \mu)$ be a token-event probabilistic ETL model. Let $n \geq 1$. We define $SC^n(H)$, $SC^n(\sim)$, $SC^n(V)$, and $SC^n(\mu)$ by induction up to n as follows:

1. $SC^n(H)$, $SC^n(\sim)$, and $SC^n(V)$ are just as in Definition 5.2

2. for $h \in SC^k(H)$,

$$
SC^k(\mu)(h) = \frac{P(h) \left[\dfrac{P(h)}{\sum_{\{h' : h' \in [h] \wedge (h \preceq h' \vee h' \preceq h)\}} P(h')} \right]}{\sum_{\{h' : h' \in SC^k(H) \wedge \exists e, f \in \Sigma \exists h''(h = h''e \wedge h' = h''f)\}} P(h')}
$$

$\triangleleft$

When normalizing over histories rather than complete information sets, we do not need to adjust the probabilities of worlds h in an information set which does not contain worlds h' such that $h \preceq h'$ or $h' \preceq h$. This is because the total probability weight on all the worlds in the information set which do lie in the same history is always equal to the weight on the single world which actually falls at the relevant time. This is ensured by the left side of the numerator. The right side of the numerator normalizes over all nodes in an information set falling on the same history. Just as in even synchronization, the denominator normalizes over all worlds descending from a single parent.