Reasoning Dynamically about What One Says *

Nicholas Asher
IRIT, Université Paul Sabatier, Toulouse

Alex Lascarides
School of Informatics, University of Edinburgh

Abstract. In this paper we make SDRT's glue logic for computing logical form dynamic. This allows us to model a dialogue agent's understanding of what the update of the semantic representation of the dialogue would be after his next contribution, including the effects of the rhetorical moves that he is contemplating performing next. This is a pre-requisite for developing a model of how agents reason about what to say next. We make the glue logic dynamic by using a dynamic public announcement logic (PAL). We extend PAL with a particular variety of default reasoning suited to reasoning about discourse—this default reasoning being an essential component of inferring the pragmatic effects of one's dialogue moves. We add to the PAL language a new type of announcement, known as *ceteris paribus* announcement, and this is used to model how an agent anticipates the (default) pragmatic effects of his next dialogue move. Our extended PAL validates certain intuitive patterns of default inference that existing PALs for practical reasoning do not. We prove that the dynamic glue logic has a PSPACE validity problem, and as such is no more complex than PAL with multiple $\Box$ operators.

Keywords: Dialogue, glue logic, dynamic logic of public announcement

1. Introduction

Speakers in dialogue anticipate their interlocutors' interpretations and adjust their utterances accordingly. Anticipation crucially involves hypothetical reasoning—reasoning about what content the interlocutors will infer, if one says one thing and not another. This in turn involves the consideration of possibly several interpretations of a potential exchange, and choosing the optimal one. Researchers have adopted planning (e.g., Stone [1998]), decision theory (e.g., Williams and Young [2007]) or game theory (e.g., van Rooij [2001]) to model such decisions.

But these approaches tend to use models of semantics that don't capture constraints on interpretation stemming from syntax and logical structure. For the past thirty years, linguists have developed dynamic semantics to study the interaction of syntax and logical structure in *dynamic discourse update*—how a discourse context C can affect the interpretation of a novel bit of information I and how I can be integrated into C to produce a new context

* The authors would like to thank the editors of this *Synthèse* volume, Emiliano Lorini and Laure Vieu, and anonymous reviewers for their comments. They also acknowledge the support of the ERC Advanced Researcher Grant 269427 STAC.

$C*$ (e.g., Kamp [1981], Kamp and Reyle [1993]). Linguists and philosophers have subsequently enriched dynamic semantics with accounts of rhetorical and discourse structure [Asher, 1993, Asher and Lascarides, 2003], which researchers in artificial intelligence [Hobbs et al., 1993, Grosz and Sidner, 1986] showed was essential to the interpretive process.

Marrying these two strands of work has not been easy. While planning, decision theory or game theory can in principle provide an account of optimal dialogue moves and dynamic discourse update (see Benotti [2010] for the beginnings of an account using planning), the richness of the information sources used in dynamic discourse update on the one hand and the paucity of information about the cognitive states of other dialogue agents on the other make it difficult to develop planning, decision theoretic or game theoretic models adequate to the task of modelling optimal interpretation. In this paper we tackle the problem from the other end, by providing an additional step to dynamic discourse update so as to bring it closer to a model of how speakers reason about what they are going to say next.

Current models of discourse update that attempt to provide rhetorically structured updates either do not provide a clear logical basis for constructing logical form [Poesio and Traum, 1998], or fail to model *dynamic* discourse update. Hobbs et al. [1993] and Asher and Lascarides [2003] use a *static* logic (e.g., weighted abduction or nonmonotonic deduction) to calculate rhetorical relations between new information and the discourse context. But the *process* of updating the discourse context with the new information and its rhetorical structure cannot be modelled within the static logic.[1] Thus, Hobbs et al. [1993] and Asher and Lascarides [2003] can't model the process of hypothetical reasoning that a speaker goes through when he compares possible contributions to the dialogue: such static logics can't model a speaker's reasoning about what to say next.

This paper aims to remedy this defect in models of dynamic discourse update. We start with Segmented Discourse Representation Theory (SDRT, Asher and Lascarides [2003]), a model of discourse where interpretation depends on logical and rhetorical structure. We make SDRT's existing, static *glue logic* for constructing logical forms dynamic by integrating the glue logic and its underlying nonmonotonic reasoning component within an off the shelf *Public Announcement Logic*. We have not investigated whether such

[1] [Stone and Thomason, 2002] develop a dynamic weighted abduction to reflect the fact that weights must be calculated in a context-sensitive way to interpret discourse. But they also argue for logical principles underlying computing weights of the following form: what is normally the case in certain circumstances is overridden if other circumstances obtain. Stone and Thomason [2002] cannot express such principles even within dynamic weighted abduction. In a nonmonotonic framework, this requires *embedded defaults* of the form *If (in A, normally B) and C, then normally D*, something which is possible in the logic of [Asher and Lascarides, 2003].

an integration would be possible with other dynamic epistemic logics because that falls outside the scope of our aims—we are merely interested here in providing *one* system in which dynamic discourse update can be modelled. This in turn will allow us to model dynamic discourse update within the glue logic itself. We thus achieve a pre-requisite for making strategic decisions during conversation—the speaker can reason about an interlocutor's interpretation of a hypothetical next dialogue move—but we leave to future work the task of interfacing these expected outcomes with the speaker's preferences or goals (although for some very preliminary work see Asher and Lascarides [2008], Asher et al. [2010]).

Section 2 motivates and describes the logical form of dialogue in SDRT and Section 3 presents its existing *static glue logic* and the accompanying discourse update for constructing these logical forms from linguistic and contextual information. Section 4 replaces the static glue logic with a dynamic one that incorporates discourse update into the axiomatisation: we extend a dynamic public announcement logic with the capability to perform default reasoning, a necessary feature of dialogue interpretation. By providing reduction axioms, we prove in Section 5 that the dynamic glue logic can be reduced to the static glue logic (through the use of much longer formulas!). So it is decidable, just as the static glue logic is, ensuring that constructing logical form in the dynamic glue logic remains computable. In fact, we go further and in the Appendix we sketch a proof that the logic is at most PSPACE in complexity. Our work also has a pleasing side effect for public announcement logic: it provides it with a natural yet powerful notion of nonmonotonic consequences for public announcement.

2. Logical Forms for Dialogue

We start by recounting briefly established arguments for including *rhetorical relations*—e.g., *Explanation*, *Background*, *Contrast* etc.—in the semantic representation of discourse (see Hobbs [1979], Kehler [2002], Asher and Lascarides [2003] for more detailed evidence).

Rhetorical relations offer an inventory of things that a speaker might be doing by performing communicative actions in conversation. Speakers typically do not present information in isolation; rather, they expand on earlier contributions, by explaining, continuing a narrative, drawing a contrast, and so forth. Interpreters expect speakers to organise discourse to highlight these meaningful relationships among successive contributions, and examples like (1), as discussed by Hobbs [1979] and Kehler [2002], show how far interpreters go to draw inferential connections between juxtaposed material so as to establish discourse coherence.

(1) a. John took a train from Paris to Istanbul. He has family there.

b. John took a train from Paris to Istanbul. He likes spinach.

Discourse (1a) makes sense. Visiting family gives a natural reason for John to make the trip, and it's natural for a speaker to continue talk of John's trip by giving an explanation for it. On the other hand, (1b) is highly anomalous: even though both sentences offer straightforward descriptions of John, one feels that something is missing—perhaps some exceptional situation that would make John go to Istanbul for spinach. What's missing, following Hobbs [1979], is *coherence*:

> ...the very fact that one is driven to such explanations indicates that some desire for coherence is operating, which is deeper than the notion of a discourse just being "about" some set of entities. [Hobbs, 1979, p67]

The interpretive effects of coherence are evident in the resolution of reference. In (1a), both Paris and Istanbul are mentioned in the first sentence, but the anaphoric expression *there* refers to Istanbul. If we make these alternative readings explicit, by replacing the second sentence with *He has family in Istanbul* or *He has family in Paris*, we continue to find the first more natural. By the same token, we ordinarily expect an explanation of a trip to account for its destination. So the resolution of *there* to Istanbul *maximises coherence*.

Logical models of discourse coherence typically include rhetorical relations like *Explanation* in logical form. For instance, in Segmented Discourse Representation Theory (SDRT, Asher and Lascarides [2003]), the logical form of a coherent discourse is a Segmented Discourse Representation Structure (SDRS): this is a rooted hierarchical set of discourse segments or *labels*, with each label π associated with some content ϕ (written $\pi : \phi$). We may also refer to the content of a segment π as K_π. The contents ϕ are expressed in a language $\mathcal{L}$ of SDRS-formulae. The vocabulary of this language includes constructors that are needed for representing the content of individual clauses or *elementary discourse units*: individual and event terms, predicate symbols, quantifiers, Boolean connectives, tense and modal operators, and so on. In addition, $\mathcal{L}$'s vocabulary also includes an inventory of *rhetorical relations* like *Explanation* and *Continuation* that each take two labels as arguments. The hierarchical structure of an SDRS is thus borne from the fact that a label π_0 can be associated with a rhetorical connection—e.g., π_0 : *Explanation*(π_1, π_2). Making an SDRS a *rooted* hierarchical structure is tantamount to ensuring that a coherent discourse is a single segment consisting of rhetorically connected subsegments. For example, the logical form of (1a) is given in (2) (for the sake of space this is highly simplified, ignoring tense among other issues):

(2) $\pi_0 : \mathit{Explanation}(\pi_1, \pi_2)$
$\pi_1 : \exists e_1\ (\mathit{travel}(e_1, \mathit{john}, \mathit{paris}, \mathit{istanbul}))$
$\pi_2 : \exists e_2 \exists x\ (\mathit{family}(x) \wedge \mathit{have}(e_2, y, x) \wedge \mathit{location}(x, z) \wedge$
$y = \mathit{john} \wedge z = \mathit{istanbul})$

We have assumed in this logical form that the compositional semantics of the pronouns *he* and *there* respectively introduce free variables y and z; it is the context-sensitive aspects of the process for *constructing* logical form that then yields the additional context-specific content that their antecedents are respectively John and Istanbul, and likewise the rhetorical connection *Explanation* between the segments (see Section 3). The result, then, is a logical form that represents a completely specific and context-aware interpretation of the discourse that goes beyond the compositional semantics of its clauses, including the resolved interpretation of pronouns and rhetorical structure.

The syntax of SDRSs is given in Definition 1, taken from Asher and Lascarides [2003]:

DEFINITION 1. **Segmented Discourse Representation Structure (SDRS)**

The set $\mathcal{L}$ of SDRS*-formulae is constructed in the usual manner from a modal language* **vocabulary** *consisting of predicate symbols, terms, Boolean operators, quantifiers, and modal operators that serve to express the content of individual clauses. The vocabulary of $\mathcal{L}$ also includes labels $\pi_1, \pi_2, \ldots$ and a set of rhetorical relation symbols (e.g.,* Explanation, Narration*) that take labels as arguments (e.g.,* Explanation$(\pi_1, \pi_2) \in \mathcal{L}$*).*

An **SDRS** *is a triple $\langle \Pi, \mathcal{F}, last \rangle$, where:*

- Π *is a set of labels from the vocabulary of $\mathcal{L}$.*

- $last$ *is a label in Π (intuitively, this is the label of the content of the last clause in the discourse); and*

- $\mathcal{F}$ *is a function which assigns each member of Π a member of $\mathcal{L}$.*

- *The relation $\succ$ that is the transitive closure of the* immediately outscopes *relation on labels Π as defined by $\mathcal{F}$ (i.e., π immediately outscopes π' iff $\mathcal{F}(\pi)$ contains the literal $R(\pi'', \pi')$ or $R(\pi', \pi'')$ for some R and π'') is a well-founded partial order with a unique root (so there is a unique label π_0 such that $\forall \pi \in \Pi$, $\pi_0 \succeq \pi$).*

In Section 3, we will discuss how SDRSs are *constructed*. But first, we briefly describe how they are *interpreted*. SDRT assigns SDRSs a dynamic semantic model theory. Rhetorical connections are treated as types of *speech acts*: like actions generally, they *transform* an input context C into a distinct output one C'. For example, as shown in (3), the dynamic interpretation of a veridical relation $R(\pi_1, \pi_2)$ (e.g., *Explanation*) consists of interpreting the contents K_{π_1} and K_{π_2} of its component segments π_1 and π_2 in dynamic succession, and adding a set of conditions $\varphi_{R(\pi_1,\pi_2)}$ that encode the particular illocutionary effects of R (detailed definitions of what constitutes a context

C can be found in Asher and Lascarides [2003] but don't concern us here):

$$C[\![R(\pi_1,\pi_2)]\!]_m C' \text{ iff } C[\![K_{\pi_1} \wedge K_{\pi_2} \wedge \varphi_{R(\pi_1,\pi_2)}]\!]_m C' \quad (3)$$

The m in $[\![.]\!]_m$ stands for 'monologue'; we will turn to the logical forms of dialogue and their dynamic interpretation shortly.

Meaning postulates on $\varphi_{R(\pi_1,\pi_2)}$ stipulate illocutionary effects: for instance, they ensure that $\varphi_{\mathit{Explanation}(\pi_1,\pi_2)}$ entails that K_{π_2} is an answer to the question *Why* K_{π_1}*?*, with the semantics of why-questions encoded roughly according to Achinstein [1980] and Bromberger [1962]. On the other hand, a divergent relation like *Correction*(π_1,π_2) entails the *negation* of the first segment:

$$\begin{array}{l} C[\![\mathit{Correction}(\pi_1,\pi_2)]\!]_m C' \text{ iff} \\ \qquad C[\![(\neg K_{\pi_1}) \wedge K_{\pi_2} \wedge \varphi_{\mathit{Correction}(\pi_1,\pi_2)}]\!]_m C' \end{array} \quad (4)$$

The dynamic interpretation of an SDRS is unpacked recursively from the content of the root label, which by definition is unique.

Lascarides and Asher [2009] argue that rhetorical relations are a vital element in the logical form of *dialogue*, as well as monologues such as (1). A fundamental decision that a speaker must make about his next move is its effects on agreement. Lascarides and Asher [2009] argue that rhetorical relations are crucial for capturing *implicit agreement*: representing the illocutionary contribution of an agent's utterance via rhetorical relations reflects his commitments to another agent's commitments, even when this is linguistically implicit. For example, Karen's utterance (5c), taken from [Sacks et al., 1974, p717], commits her to (5b) thanks to the semantic consequences of the relational speech act *Explanation*$(5b, 5c)$ that she has performed:

(5) a. *Mark (to Karen and Sharon):* Karen 'n' I're having a fight,

b. *Mark (to Karen and Sharon):* after she went out with Keith and not me.

c. *Karen (to Mark and Sharon):* Wul Mark, you never asked me out.

Arguably, by committing to (5b) Karen also commits its *illocutionary effects*—(5b) explains (5a). These commitments aren't *monotonically* entailed by (5c)'s compositional semantics nor by Karen's asserting it. Rather, Karen's implicit acceptance of Mark's contribution is dependent on the relational speech acts they perform and their semantics. Therefore, with agreement defined to be shared public commitment, dialogues like (5) are evidence that the representation of an agent's commitments should include rhetorical relations.

More generally, Lascarides and Asher [2009] propose that the commitments of each agent at a given dialogue turn (where a turn boundary occurs whenever the speaker changes) is an SDRS, as shown in each cell of Table I,

Table I. A representation of dialogue (5).

Turn	Mark's SDRS	Karen's SDRS
1	π_{1M} : *Explanation*(π_1, π_2)	$\emptyset$
2	π_{1M} : *Explanation*(π_1, π_2)	π_{2K} : *Explanation*$(\pi_1, \pi_2) \wedge$ *Explanation*(π_2, π_3)

the proposed logical form for (5). For simplicity, we have omitted from Table I the contents of the clauses (5a) to (5c), corresponding to labels π_1 to π_3, and adopted a convention that the root label of agent a's SDRS for turn j is π_{ja}. Definition 2 is the formal definition of logical form for dialogue, as illustrated in Table I.

DEFINITION 2. **Dialogue SDRSs**

Let D be a set of dialogue agents. Then a Dialogue SDRS *(*DSDRS*) is a tuple $\langle n, T, \Pi, \mathcal{F}, last\rangle$, where:*

– n *is a natural number ($j \leq n$ is the j^{th} turn);*

– Π *is a set of labels;*

– $\mathcal{F}$ *maps* Π *to* SDRS*-formulae* $\mathcal{L}$*;*

– T *maps* $[1, n]$ *to a function from* D *to* SDRS*s, such that each* SDRS *is drawn from* Π *and* $\mathcal{F}$*. That is,* $T(j)(a) = \langle \Pi^a_j, \mathcal{F}^{d_i}_j, last^a_j\rangle$ *where* $\Pi^a_j \subseteq \Pi$, $\mathcal{F}^a_j = \mathcal{F} \upharpoonright \Pi^a_j$ *(i.e.,* $\mathcal{F}$ *restricted to* Π^a_j*), and* $last^a_j \in \Pi^a_j$*.*

– $last = last^d_n$*, where* d *is the (unique) speaker of the last turn* n*.*

We may refer to the SDRS $T(j)(a)$ *as* $T^a(j)$*.*

The logical form of dialogue (e.g., Table I) is called a Dialogue SDRS (DSDRS). The agents' SDRSs can share labels, and each label is always associated with the same content. Thus an agent can commit to the content expressed by prior speech acts, even if they were performed by another agent. For example, *Explanation*(π_1, π_2) is a part of Karen's SDRS, making her committed to her and Mark having a fight because she went out with Keith.

Lascarides and Asher [2009] define a dynamic semantics $\models_d$ for DSDRSs in terms of the semantics $\models_m$ for SDRSs that we briefly described earlier and that is given in detail in [Asher and Lascarides, 2003]. Informally, $\models_d$ captures shared commitments. A context of interpretation $\mathcal{C}$ for interpreting a DSDRS is a tuple of contexts C_a for interpreting SDRSs, one for each agent

a in the group of dialogue agents D. And a dialogue turn $T = \{S_a : a \in D\}$, consisting of an SDRS S_a for each agent $a \in D$, transforms an input dialogue context $\mathcal{C}$ into an output one $\mathcal{C}'$ only if for each $a \in D$, $C_a[\![S_a]\!]_m C'_a$. Thus, with agreement being shared public commitment, the logical form of (5) makes the following agreed upon: Mark and Karen were having a fight because she went out with Keith and not Mark.

3. SDRT's Glue Logic

Discourse interpretation involves *constructing* logical forms like (2) (for monologue) and Table I (for dialogue), as well as evaluating those logical forms. For constructing logical form, Asher and Lascarides [2003] propose a number of default axioms for inferring rhetorical connections among segments in a text that, in principle at least, support an inference that the contents of the clauses in (1a) are related by *Explanation* (and that *he* refers to John and *there* to Istanbul).

In addition, Lascarides and Asher [2009] propose a number of default axioms for constructing the logical forms of dialogue that are designed to predict the semantic scope of implicit and explicit endorsements and challenges, and provide the basis for adding *Explanation*(π_1, π_2) to Karen's SDRS in (5). In general, these principles are designed to maximise one's ongoing commitments from prior turns, subject to them being consistent with default inferences about the illocutionary contribution one makes in the current turn. This principle therefore predicts that Karen makes different commitments if (5a) is replaced with (5a$'$) (and the word *after* is removed from (5b)):

(5) a$'$. Karen is a bitch.

In this case, as before, Karen is committed to *Explanation*(b, c) and therefore is committed to the content of (5b). However, this inference that (5c) explains (5b) supports a further default inference about the speech act that Karen performed: intuitively, Karen uses this *Explanation* segment to *justify* her choices as conveyed in (5b), making those choices reasonable rather than vindictive and thereby undermining (5b) as an explanation of (5a$'$)—the content that Mark committed to. Thus, default inferences about the illocutionary contribution of (5c) makes Karen committed, via *Counterevidence*, to the negation of (5a$'$) (since *Counterevidence* is a divergent relation). This is inconsistent with a commitment to *Explanation*(a', b), and so the logic for constructing logical form does not add Mark's commitments from the first turn to Karen's commitments for the second turn, even though Karen is committed to a part of what Mark committed to: namely, (5b).

More generally, Asher and Lascarides [2003] emphasise two features that any adequate procedure for constructing a discourse interpretation must have.

First, it must support defeasible inferences about which rhetorical relations connect the segments of the discourse together. The inference is defeasible because as (1) and (5) illustrate, rhetorical connections are often implied on the basis of assumptions about non-linguistic contextual information (including assumptions about the speaker's cognitive state to which there's no direct access), rather than purely on the basis of explicit linguistic information. Secondly, they argue that *constructing* a discourse interpretation should be at least decidable, even if computing the consequences of that interpretation (e.g., deciding whether to believe it) needn't be. Consider, for example, discourse (6):

(6) a. There are some unsolvable problems in number theory.

b. Every even number greater than two is expressible as the sum of two primes is undecidable, for instance.

Sentence (6b) plainly elaborates (6a). If we were to infer this via a system of defeasible reasoning that has full access to the information content of the two clauses, then this would involve, according to all extant accounts of nonmonotonic reasoning, a consistency test: given that the semantics of *Elaboration* is such that it's true only if the *information content* of the two constituents it connects are also true (see (3)), testing the *Elaboration* connection for consistency will entail a test as to whether (6a) and (6b) are satisfiable. That is, we would need to test whether Goldbach's Conjecture is in fact undecidable or not, something which we have no idea how to do! But even the most mathematically inept interpreter can easily understand the discourse structure of (6) and construct its logical form; one has a clear picture of *what* is being said without being able to *evaluate* whether what is said is true.

The logical form one computes from the linguistic signal ought to be something that people can get right or wrong. It's at the level of logical form that phenomena like anaphoric accessibility, presupposition attachment and temporal structure are determined, phenomena about which competent speakers have reliable and robust intuitions. They should therefore be able to compute the 'theorems' of the logic of information packaging; there should be a guarantee that proofs of 'valid' inferences for building the logical form can be constructed. While the logic of information packaging should be computable in this sense, the logic of information content as described in Section 2 clearly is not: the language $\mathcal{L}$ has at least the expressive power of first order logic, making validity at least recursively enumerable. Thus, the logic of information packaging will be tractable only if it's distinct from this logic of information content. Of all the extant logical approaches to modelling discourse coherence (e.g., Hobbs et al. [1993], Stone and Thomason [2002]), SDRT is the only framework that makes this clear division between the logic for *constructing* logical form and the logic for *interpreting* logical form. We call the logic for constructing logical form the *glue logic* (GL).

The logic GL is separated from the logic $\models_d$ of DSDRSs, but it must nevertheless be in some way informed by it. Lexical and compositional semantics play a crucial part in inferring rhetorical relations as we saw in (1) and (5). So GL needs access to information about the contents of the sentences. But to remain decidable, it must have only *limited* access to that information. This is the only way to preserve its computability, given the inferential flexibility we want this logic to have.

SDRT achieves this separation of GL from the logic $\models_d$ of information content by using *underspecified semantics* (e.g., Egg et al. [2001]). An underspecified logical form (ULF) is a *partial description* of the *form* of the logical form, here a DSDRS. SDRT's glue logic builds ULFs, and so its language $\mathcal{L}_{ulf}$ describes DSDRSs: each n-ary constructor in $\mathcal{L}$ corresponds in $\mathcal{L}_{ulf}$ to an $(n+1)$-ary predicate symbol that takes labels as arguments. Labels denote scopal positions—in other words, they determine the scopes of various operators—in the DSDRS being described. So the first n arguments to the predicate P in $\mathcal{L}_{ulf}$—where P corresponds to the n-place constructor P in $\mathcal{L}$—denotes respectively the n argument positions of P in the described DSDRS(s), and the $(n+1)^{th}$ label is the scopal position of P itself. This label is written to the left of a colon with its predicate and other arguments to the right. So $l{:}\mathsf{dog}(l') \wedge l'{:}\mathsf{d}$ in $\mathcal{L}_{ulf}$ describes $dog(d)$ in $\mathcal{L}$. Note that $\mathcal{L}_{ulf}$ also includes the Boolean connective $\wedge$, to be distinguished from $\mathcal{L}$'s $\wedge$, which in $\mathcal{L}_{ulf}$ corresponds to a *predicate symbol*.

$\mathcal{L}_{ulf}$ also includes variables $?_1, ?_2, \ldots$ to indicate that the value of a constructor in the described DSDRS(s) is unknown. For instance, the compositional semantics of a pronoun introduces into the ULF the formula (7)—an equality between a variable y from $\mathcal{L}$'s vocabulary and some unknown variable from $\mathcal{L}$'s vocabulary:

$$l_0{:}{=}(l_1, l_2) \wedge l_1{:}\mathsf{y} \wedge l_2 :?_2 \qquad (7)$$

GL's vocabulary also includes the weak conditional $>$ that's used to formalise defaults ($A > B$ means "If A then normally B") and equality (again to be distinguished from the *predicate symbol* = that corresponds to equality in $\mathcal{L}$). For instance $?_2 = \mathsf{john}$ is a formula of $\mathcal{L}_{ulf}$, and together with the $\mathcal{L}_{ulf}$ formula (7) they provide a partial description of any DSDRS that features the $\mathcal{L}$-formula $y = john$ as a part. Indeed, in constructing the logical form (2) of (1a), GL will validate a nonmonotonic inference from (7) and other contextual information to the consequent $?_2 = \mathsf{john}$.

The following formal definition of $\mathcal{L}_{ulf}$'s syntax is taken from Asher and Lascarides [2003]:

DEFINITION 3. **Syntax of the GL language $\mathcal{L}_{ulf}$**

Vocabulary: *The vocabulary consists of:*

1. *A set of labels* $l_1, l_2 \ldots, \pi_1, \pi_2 \ldots$. *We conventionally use the former in the* ULF *of a minimal unit of a* DSDRS *(typically, this is the* ULF *of a clause), and the latter for discourse segments (intuitively, they correspond to the labels* π *in the* DSDRS *being described).*

2. *A set of constants* $d_1, d_2 \ldots$ *(these denote dialogue agents) and* $m_1, m_2 \ldots$ *(these denote the turn of the dialogue).*

3. *A set of predicate symbols derived from the vocabulary of* $\mathcal{L}$ *(see Definition 1): for each* n*-place constructor* P *in* $\mathcal{L}$*'s vocabulary there is a unique* $(n+1)$*-place constructor* P *in* $\mathcal{L}_{ulf}$*. We also assume a three place predicate* T *where intuitively* $T(d, m, \pi)$ *means that* π *is a part of the* SDRS $T^d(m)$ *in the* DSDRS *being described.*

4. *A set of label variables and a set of predicate symbol variables. We will refer to labels and label variables as label terms, and predicate symbols and predicate variables as predicate terms.*

5. *The two-place predicates* $=$ *and* outscopes *(*outscopes(l_1, l_2) *means that the scopal position denoted by* l_1 *outscopes that of* l_2 *in the* DSDRS *being described).*

6. *Boolean connectives, the weak conditional* $>$ *and a modal operator* $\Box$.

Well-formed Formulae:

1. *If* P *and* P' *are* $(n+1)$*-place predicate terms and* $l_1, \ldots, l_{n+1}$ *are label terms, then* $l_{n+1} : P(l_1, \ldots, l_n)$, outscopes$(l_1, l_2)$, $l_1 = l_2$ *and* $P = P'$ *are all well-formed formulae.*

2. *If* ϕ, ψ *are well-formed formulae then so are* $\phi \wedge \psi$, $\phi > \psi$, $\neg\phi$ *and* $\Box\phi$.

For the sake of readability, we will often re-write ULFs so that each label that's associated with a unique constructor is replaced by that constructor. For instance $l : \neg(l_1) \wedge l_1 : \mathsf{dog}(l_2) \wedge l_2 : \mathsf{d}$ may be re-written $l : \neg(\mathsf{dog}(\mathsf{d}))$.

The glue logic derives a logical form (or, in fact, a ULF) via a finite background theory of default axiom schemata that attempt to capture constraints on interpretation that are imposed by discourse coherence. These axioms predict pragmatically preferred values for underspecified semantic elements. They in particular predict rhetorical connections, and have the form given in (8):

$$(\lambda{:}?(\alpha, \beta) \wedge \mathsf{Info}(\alpha, \beta, \lambda)) > \lambda{:}\mathsf{R}(\alpha, \beta, \lambda) \quad (8)$$

$$(\lambda{:}?(\alpha, \beta) \wedge \alpha{:}\mathsf{int}) > \lambda{:}\mathsf{IQAP}(\alpha, \beta) \quad (9)$$

In words, if segment β is rhetorically connected to α as part of a segment λ but the relation is unknown, and moreover $\mathsf{Info}(\alpha, \beta, \lambda)$ holds of the content labelled by λ, α and β, then normally the rhetorical relation is R. The conjunct $\mathsf{Info}(\alpha, \beta, \lambda)$ is proxy for particular ULF formulae, and the axioms are justified on the basis of linguistic knowledge, world knowledge, or knowledge of cognitive states. Rule (9) is an example from Asher and Lascarides [2003] where the predicate symbol int stands for interrogative mood; so (9) states that a response to a question is normally an indirect answer (*IQAP* stands for Indirect Question Answer Pair).

Definition 4 gives the model theory for the glue language $\mathcal{L}_{ulf}$: intuitively, $\mathcal{M}, s \models_g \phi$ means that ϕ is a (perhaps partial) description of the DSDRS s. This model theory is static; in Section 4 we make it dynamic.

DEFINITION 4. **Static Glue Model Theory**

A model $\mathcal{M} = \langle S, R_\Box, *, V\rangle$ for $\mathcal{L}_{ulf}$ is a set of states S where each $s \in S$ is a DSDRS, a transitive, reflexive and symmetric equivalence relation $R_\Box$ on S (for interpreting the S5 modality $\Box$), a function $*$ from a state and a set of states to a set of states (for interpreting $>$), and a function V for interpreting $\mathcal{L}_{ulf}$'s non-logical constants (we'll see shortly that V is constrained by the partial transfer of $\models_d$-entailments from $\mathcal{L}$ described above). Then:

– $\mathcal{M}, s \models_g \phi$ *iff* $s \in V(\phi)$ *for atomic* ϕ
$\mathcal{M}, s \models_g \phi \wedge \psi$ *iff* $\mathcal{M}, s \models_g \phi$ *and* $\mathcal{M}, s \models_g \psi$
$\mathcal{M}, s \models_g \neg\phi$ *iff* $\mathcal{M}, s \not\models_g \phi$
$\mathcal{M}, s \models_g \Box\phi$ *iff for all* $s' \in R_\Box(s)$, $\mathcal{M}, s' \models_g \phi$
$\mathcal{M}, s \models_g \phi > \psi$ *iff* $*^{\mathcal{M}}(s, [\![\phi]\!]^{\mathcal{M}}) \subseteq [\![\psi]\!]^{\mathcal{M}}$,
where $[\![\phi]\!]^{\mathcal{M}} = \{s' \ : \ \mathcal{M}, s' \models_g \phi\}$

As we've said, some (but not all) of the $\models_d$-consequences from the dynamic semantics of DSDRSs are transferred into the glue logic via constraints on the function V that restrict set of admissible models of $\models_g$. The $\models_d$-consequences on DSDRSs arising from substitution of equalities, and rules of propositional logic are validated by $\models_g$—in other words, if ϕ and ϕ' are DSDRSs such that $\phi \models_{fd} \phi'$, where $\models_{fd}$ is the (decidable) fragment of $\models_d$ that validates the inference rules we've just mentioned, and s and s' are respectively the states in the GL model $\mathcal{M}$ that are the DSDRSs ϕ and ϕ' then $\mathcal{M}, s' \models_g \psi$ only if $\mathcal{M}, s \models_g \psi$ as well. The formal details of this transfer function (or, equivalently, the definition of how the interpretation function V in admissible models is constrained), together with a proof that it is decidable, are described in [Asher and Lascarides, 2003, p195–197]. For reasons of space we do not repeat the details here because, while straightforward, they are long and tedious and divert us from the main point of this paper—to demonstrate that we can make the glue logic dynamic while maintaining its decidability. The only important issue for the purposes of this paper is that this information transfer is decidable and we don't change this transfer process in the dynamic version of the glue logic.

The glue logic also has monotonic and nonmonotonic relations $\vdash_g$ and $\mid\!\sim_g$ [Asher and Morreau, 1991, Asher, 1995]. $\vdash_g$ has a complete axiomatisation using a canonical model construction, while $\mid\!\sim_g$ yields default inferences, via axioms like (9), about discourse interpretation, including particular resolutions of anaphora. $\vdash_g$ abides by axioms of classical logic, the characteristic S5 axioms and rules for $\Box$ plus axioms on $>$-formulae such as those in (10)—(13):

$$A > A \tag{10}$$

$$\Box(A \leftrightarrow B) \rightarrow ((A > C) \ \leftrightarrow \ (B > C)) \tag{11}$$

$$(\Box(B \rightarrow C) \wedge (A > B)) \rightarrow (A > C) \tag{12}$$

$$(\Box(A \rightarrow B) \wedge (A > C) \wedge (B > \neg C)) \rightarrow (B > \neg A) \tag{13}$$

$$\Box A \rightarrow (B > A) \tag{14}$$

Corresponding constraints are imposed on the function $*$ in the model theory (where p, q etc. are subsets of S): these constraints make the logic sound and complete [Asher, 1995].

$$*(s,p) \subseteq p \tag{15}$$

$$(*(s,p) \subseteq q \wedge *(s,q) \subseteq p) \rightarrow *(s,p) = *(s,q) \tag{16}$$

$$(p \subseteq q \ \wedge \ *(s,p) \cap *(s,q) = \emptyset) \rightarrow *(s,q) \subseteq S \setminus p \tag{17}$$

$$(s,p) \subseteq R_{\Box}(s) \tag{18}$$

Nonmonotonic inferences $\mid\!\sim_g$ are computed via $\vdash_g$ by converting $>$-formulae into $\rightarrow$-ones, provided the result is $\vdash_g$-consistent. $\mid\!\sim_g$ validates many intuitively compelling inferences such as those below, and the logic is sound, complete and decidable [Asher, 1995].

Defeasible Modus Ponens: $\phi, \phi > \psi \ \mid\!\sim \ _g\psi$

Penguin Principle: $\Box(\phi \rightarrow \psi), \phi, \phi > \chi, \psi > \neg\chi \mid\!\sim_g \chi$

Nixon Diamond: $\neg\Box(\phi \rightarrow \psi), \phi, \psi, \phi > \chi, \psi > \neg\chi \not\mid\!\sim_g \chi$ (and $\not\mid\!\sim_g \neg\chi$)

Weak Deduction: if (a) $\Gamma, \phi \mid\!\sim_g \psi$, (b) $\Gamma \not\mid\!\sim_g \psi$ and (c) $\Gamma \not\mid\!\sim_g \neg(\phi > \psi)$, then (d) $\Gamma \mid\!\sim_g (\phi > \psi)$

The Penguin Principle is valid because (13) is a rule of $\vdash_g$. As Weak Deduction makes evident, our logic distinguishes between $>$, which is used to express defaults about interpretation, from the nonmonotonic consequence relation $\mid\!\sim_g$. This is important, because it allows us to express embedded defaults, which are needed to encode constraints on interpretation [Asher and

Lascarides, 2001, 2003]. We assume henceforth that all $\mathcal{L}_{ulf}$ models verify the background axioms of GL; e.g., (9). We furthermore assume that there are a finite number of these axioms.

Definition 5 makes the updated ULF include all the $\mid\!\sim_g$-consequences of the old and the new information (see **Simple Update**). So update always *adds* constraints to what the dialogue means. If there is a choice of labels that new content can attach to, then update is conservative and generalises over all the possibilities (see **Discourse Update**).

DEFINITION 5. **Discourse Update for DSDRSs**

Simple Update *of a context with new content β, given a particular attachment site α.*

Recall that $T(d, m, \lambda) \in \mathcal{L}_{ulf}$ means that the label λ is a part of the SDRS $T^d(m)$ *in the* DSDRS *being described. So the* ULF*-formula* $\lambda{:}?(\alpha, \beta) \wedge T(d, m, \lambda)$ *specifies that the new information β attaches to the* DSDRS *as a part of the* SDRS $T^d(m)$. *Let σ be a set of (fully-specified)* DSDRS*s, and let $Th(\sigma)$ be the set of all* ULF*s that partially describe the* DSDRS*s in σ—That is, a formula ϕ is in $Th(\sigma)$ iff for each $s \in \sigma$, $s \models_g \phi$. Let ψ be either (a) a* ULF $\mathcal{K}_\beta$, *or (b) a formula* $\lambda{:}?(\alpha, \beta) \wedge T(d, m, \lambda)$, *where* $Th(\sigma) \vdash_g \mathcal{K}_\beta$. *Then:*

$$\begin{aligned} \sigma + \psi = {} & \{\tau : \textit{ if } Th(\sigma), \psi \mid\!\sim_g \phi \textit{ then } \tau \vdash_g \phi\}, \\ & \textit{provided this is not } \emptyset; \\ \sigma + \psi = {} & \sigma \textit{ otherwise} \end{aligned}$$

Discourse Update. *Suppose that A is the set of available attachment points in the old information σ. $update_{\text{SDRT}}(\sigma, \mathcal{K}_\beta)$ is the* union *of* DSDRS*s that results from a sequence of $+$-operations for each member of the powerset $\mathcal{P}(A)$ together with a stipulation that the last element of the updated* DSDRS *is β.*

The powerset $\mathcal{P}(A)$ represents all possible choices for what labels in σ the new label β is attached to, so $update_{\text{SDRT}}$ is neutral about which member of $\mathcal{P}(A)$ is the 'right' choice.

Discourse update typically doesn't yield a specific enough ULF to identify a *unique* logical form or DSDRS. But intuitively, some DSDRSs that satisfy the $\mid\!\sim_g$-consequences are 'preferred' because they are *more coherent* (see the earlier discussion of (1a), with *there* interpreted as Istanbul vs. Paris). SDRT makes degree of coherence influence interpretation by *ranking* the DSDRSs in the update into a partial order. This partial order adheres to some very conservative assumptions about what contributes to coherence; Definition 6 is from Asher and Lascarides [2003].

DEFINITION 6. **Maximise Discourse Coherence (MDC)**

Discourse is interpreted so as to maximise discourse coherence, where the (partial) ranking $\prec$ among interpretations adheres to the following:

1. *All else being equal, if* DSDRS ϕ *has more rhetorical connections between two labels than* DSDRS ψ, *then* $\psi \prec \phi$.

2. *All else being equal,* $\psi \prec \phi$ *if* ϕ *features more atomic predications that are an essential part of the premises in a* $\mid\!\sim_g$*-inference of a particular rhetorical relation. In other words,* Max fell because John pushed him. *is more coherent than* Max fell because Mary's hair is black, *because the former can use the atomic predication featuring* because *to* $\mid\!\sim_g$*-infer* Explanation *and it can also use the combination of the atomic predications featuring* fall *and* push *(cf* Max fell. John pushed him.*). In contrast, the latter can use only* because.

3. *Some rhetorical relations are inherently scalar. For example, the quality of a Continuation is dependent on the specificity of its common topic (compare* Max is tall. He has blonde hair *with* Max is tall. John has blonde hair*). All else being equal,* $\psi \prec \phi$ *if* ϕ *features higher quality rhetorical relations.*

4. *All else being equal,* $\psi \prec \phi$ *if* ϕ *has fewer labels but no semantic anomalies: e.g.,* $\pi_0{:}Contrast(\pi_1, \pi_2) \wedge Condition(\pi_2, \pi_3)$ *is anomalous because the first speech act 'asserts'* K_{π_2} *and the second doesn't. Conversely,* $\pi_0{:}Contrast(\pi_1, \pi)$, $\pi{:}Condition(\pi_2, \pi_3)$ *isn't anomalous.*

4. Dynamic Commitments in SDRT

Definition 5 uses the entailment relation $\mid\!\sim_g$ but it is external to the glue logic; so is MDC. But they must become a part of the axiomatisation of GL if we are to support strategic decisions about what to say (see discussion in Section 1). Luckily, the tools of dynamic epistemic logic, and in particular Public Announcement Logic (PAL) [Baltag et al., 1999], make this relatively straightforward to achieve. Particularly interesting from our point of view is that PAL introduces dynamicity *within the form of operators*, while keeping the underlying logic and consequence notion very close to that of the original glue logic. This is an extremely useful feature, given that we think it an important conceptual point to keep our glue logic of low computational complexity. Given that we can prove reduction axioms for our use of PAL relative to the original glue logic, which is decidable, we have an easy proof of the decidability of the dynamic glue logic. This would not be nearly so straightforward were we to dynamicise our glue logic in the spirit of dynamic semantics: we would have to devise a new dynamic consequence relation incorporating nonmonotonicity. This could be undecidable, and in any case it would be a complex affair. Using PAL avoids these problems.

A PAL features the action of announcing a formula ϕ, which changes the model $\mathcal{M}$ by restricting the states in the output model $\mathcal{M}^{\phi}$ to those which satisfy the announced formula (see Definition 7 below, which follows the standard PAL definition, for the formal definition of $\mathcal{M}^{\phi}$). SDRT's glue logic can thus be recast in terms of the effects of announcing a formula: the states of the model are still DSDRSs (see Definition 4), with announcements eliminating DSDRSs from the input model that fail to satisfy the announcement.

As Definition 5 suggests, we need to specify the effects of **three sorts of announcements**:

1. $\mathcal{K}_{\beta}$—the ULF of an utterance or segment.
2. $\lambda{:}?(\alpha, \beta) \wedge T(d, j, \lambda)$—a choice of where to attach a new segment.
3. $last = \beta$—the segment β is the last entered element.

If all consequences of one's announcements were monotonic, then simple PAL would do. But Section 3 makes plain that *nonmonotonic* consequences of announcements determine the DSDRSs, since discourse interpretation is generally a product of commonsense reasoning.

Extensions to PAL that support nonmonotonic reasoning already exist. For instance, van Benthem [2007] and Baltag and Smets [2006] propose dynamic PALs for modelling belief revision; they incorporate into a standard PAL conditional doxastic models, with logics equivalent to AGM belief revision theory [Alchourrón et al., 1985]. Like their PALs, ours is extended by introducing a weak conditional connective. But there are some important differences. Our extension to PAL does not incorporate doxastic models; we define the nonmonotonic reasoning directly via a transformation of the conditional connective, not via an operation of belief revision. Underlying this difference is a difference in motivation. We are not engaged in the same enterprise as van Benthem [2007] and Baltag and Smets [2006]: we are *not* reasoning about what agents explicitly believe and what they nonmonotonically infer from those beliefs, but rather how a dialogue agent reasons monotonically and nonmonotonically about what discourse content is expressed by the linguistic signal given its discourse context. All the reasoning in the glue logic takes place, as it were, within the beliefs of the agent. And while we believe that a dialogue agent does reason about the cognitive states of other agents in choosing an optimal interpretation, the notion of discourse update can and should be defined independently of how one represents such cognitive information.

Our logic differs from that of van Benthem [2007] and Baltag and Smets [2006] in other ways. We distinguish the weak conditional from the consequence relation it gives rise to. This allows us to express embedded defaults, which are essential in encoding constraints on interpretation. Asher [1995] shows that the presence of embedded defaults makes the consequence relation $\mid\!\sim$ itself not subject to the AGM postulates. Furthermore, the distinction between the connective $>$ and the nonmonotonic consequence relation from Commonsense Entailment [Asher and Morreau, 1991, Asher, 1995] enables us to construct models validating the monotonic axiom (13) and hence verify the Penguin Principle. The Penguin Principle incorporates an important and intuitively compelling principle of nonmonotonic inference that we have shown extensively elsewhere is vital for accurately predicting the logical form of coherent discourse [Lascarides and Asher, 1993, Asher and Lascarides, 2003]. So it is essential that our dynamic PAL version of the glue logic continue to support this type of inference.

Our strategy, then, is to introduce another sort of announcement—not simple announcement but *announcement ceteris paribus* or ACP—and we will define ACPs in terms of the conditional $>$ from Commonsense Entailment [Asher and Morreau, 1991, Asher, 1995], so that ACPs support similar inference patterns to those supported in the static version of the Glue Logic.

We will convert the static model theory from Definition 4 into a dynamic one for interpreting ACPs. This involves (a) extending the language $\mathcal{L}_{ulf}$ (see Definition 3) to express announcements; and (b) defining how models are transformed by such announcements in interpretation. As is standard in PAL, we add a modality $[!\phi]$ to $\mathcal{L}_{ulf}$ to express the announcement that ϕ. The formula $[!\phi]\psi$ means that ψ follows from announcing ϕ. The above three sorts of announcements, as suggested by Definition 5, are all $>$-free. So we make announcements $>$-free.[2] We extend standard PAL by introducing a new modality $[!\phi]^{cp}$ for ACPs, where $[!\phi]^{cp}\psi$ means that ψ normally follows from announcing ϕ.

SDRT's Glue Logic (GL) contains a number of axioms—for instance axiom (9)—that characterise SDRT's constraints on attachment of new constituents to a given SDRS, and in particular characterise what rhetorical relations can be inferred between two given points of attachment. These axioms form the *background theory* of GL. The background theory of GL is, crucially, finite; and so the premises for any GL inference, which includes the background theory together with information about the discourse context, forms a finite set as well. For such finite sets of sentences, there exists a **prime implicate** or strongest finite formula that is a nonmonotonic consequence of the theory and from which all other nonmonotonic consequences follow [Asher, 1995]. The prime implicate of the background theory together with ϕ, which we shall write as $\Im_\phi$, characterises the nonmonotonic closure of ϕ. Because $\mid\!\sim$ is supra-classical (this follows from the rules for $\Box$ and axioms (10) and (12)), we have as an axiom $\Im_\phi \rightarrow \phi$. Indeed, the following is a **fact**:

$$\vdash \Im_\phi \rightarrow \psi,\ \text{for all } \psi \text{ such that } \phi \mid\!\sim \psi$$

[2] This restriction is not needed for the logic; it's just that we don't see the utility of including announcements of $>$-formulae.

Because we can make use of the prime implicate, we can reduce $[!\phi]^{cp}$ to a standard PAL operator; the nonmonotonic logic just tells us which formula is *ceteris paribus* announced. We now assign this extended language $\mathcal{L}_{ulf}$ a dynamic model theory, with announcements transforming models. Note the use of the prime implicate in the definition of ACPs.

DEFINITION 7. **Dynamic Glue Model Theory**
Let $\mathcal{M} = \langle S, R_{\Box}, *, V\rangle$ be a model as in Definition 4. We define $\mathcal{M}^{\phi}$ in the standard way for 'classical' PAL, and we define $\mathcal{M}^{cp(\phi)}$ using the non-monotonic closure of ϕ given the background truths of GL:

$$\begin{aligned}\mathcal{M}^{\phi} &= \langle S^{\phi},\ *^{\mathcal{M}}|S^{\phi}, R^{\mathcal{M}}|S^{\phi}, V\rangle,\ \text{where } S^{\phi} = S^{\mathcal{M}} \cap [\![\phi]\!]^{\mathcal{M}}\\ \mathcal{M}^{cp(\phi)} &= \langle S^{cp(\phi)}, *^{\mathcal{M}}|S^{cp(\phi)}, V\rangle,\ \text{where } S^{cp(\phi)} = S^{\mathcal{M}} \cap [\![\Im_{\phi}]\!]^{\mathcal{M}}\end{aligned}$$

The interpretations of $[!\phi]\psi$ and $[!\phi]^{cp}\psi$ are:

$$\begin{aligned}&\mathcal{M}, s \models [!\phi]\psi \quad \text{iff}\ \ \mathcal{M}, s \not\models \phi \text{ or } \mathcal{M}^{\phi}, s \models \psi\\ &\mathcal{M}, s \models [!\phi]^{cp}\psi \ \ \text{iff}\ \ \mathcal{M}, s \not\models \Im_{\phi} \text{ or } \mathcal{M}^{cp(\phi)}, s \models \psi\end{aligned}$$

In words, model $\mathcal{M}^{\phi}$ is formed from $\mathcal{M}$ by eliminating all states that don't satisfy the monotonic consequences of announcing ϕ; and $\mathcal{M}^{cp(\phi)}$ is formed by eliminating all states that don't satisfy the nonmonotonic consequences of announcing ϕ. Note that because $>$ is supra-classical,[3] ACPs, like 'simple' announcements, make the announcement true; i.e. $S^{cp(\phi)} \subseteq [\![\phi]\!]$. This simply means that the ULF of an utterance is always a ULF for the entire dialogue; or to put this another way, a partial description of the logical form of an utterance is also a partial description of the logical form of the whole dialogue. It does *not* mean that the utterance is true or even that the speaker is committed to it.

Glue logic axioms like (9) make the consequences of ACPs express information about rhetorical connections or specific values for other underspecified elements introduced by linguistic syntax, like antecedents to pronouns. The axioms from Lascarides and Asher [2009] also ensure that ACPs predict which commitments from prior turns are current commitments. For instance, for dialogue (5), the GL axioms detailed in Asher and Lascarides [2003] and Lascarides and Asher [2009] ensure that, in words, the consequence of announcing *ceteris paribus* the content of (5c) followed by an assumption that it is attached, within Karen's SDRS for the second turn, to the segment (5b), is that Karen is committed within that turn to the content that (5b) explains (5a). More formally, the satisfaction relation in (19) holds, where $\mathcal{M}$ is the model constructed by updating with utterances π_1 and π_2 in that order and $s \in \mathcal{M}$:

$$\begin{aligned}\mathcal{M}, s \models [!\mathcal{K}_{\pi_3}]^{cp}([!(\pi_{2K}{:}?(\pi_2, \pi_3) \wedge T(K, 2, \pi_{2K}))]^{cp}[last = \pi_3]^{cp}\\ \pi_{2K}{:}\text{Explanation}(\pi_1, \pi_2))\end{aligned} \tag{19}$$

[3] For note that from axioms (10) and (12) it follows that $\Box(A \rightarrow B) \rightarrow (A > B)$ is an axiom too.

We note that the background truths of GL are modally free apart from a $\Box$ with widest scope (i.e., they are of the form $\Box(A > B)$ where A and B are $\Box$-free), and so they are automatically preserved as we move from $\mathcal{M}$ to $\mathcal{M}^{\phi}$ or to $\mathcal{M}^{cp(\phi)}$, thanks to the particularly simple nature of the PAL update operation.[4]

In equation (19), the announcement of segment π_3's content ($\mathcal{K}_{\pi_3}$) is before the announcement of π_3's attachment to π_2. This ordering in the announcements mirrors the static definition of discourse update (see Definition 5). Update must proceed in this order to ensure that we explore inferences about rhetorical connections in a context that already satisfies the *content* of the segment to be connected. Otherwise, the GL background theory is in danger of either undergenerating or overgenerating inferences about which rhetorical connection holds. This two-step approach to announcing an update with new information is reflected in the formal Definition 9 of dynamic discourse update. Indeed, it is now straightforward to define discourse update *within the logic*, as Definition 8 shows. We make the set of DSDRSs σ the set of states of a model $\mathcal{M}_\sigma$:

DEFINITION 8. **Dynamic Simple Update:**
$\sigma + \phi \models \psi$ *iff* $M_\sigma \models [!\phi]^{cp}\psi$

To define full DSDRS update, we take Boolean combinations of ACP updates so as to match the update process from Definition 5 (second paragraph).

DEFINITION 9. **Dynamic Discourse Update:**
Let $\mathcal{M}_\sigma$ be 'old' information and the ULF $\mathcal{K}_\beta$ *be new information. Let $\Sigma_1, \ldots \Sigma_n$ be all the jointly compossible attachment sites of β, chosen from the set A of all possible attachment sites for each* DSDRS *in σ. Let k_i be an enumeration of the compossible attachment sites in Σ_i, $1 \leq i \leq n$. And let k_i be the sequence of assumptions about attachment provided by the enumeration k_i of sites in Σ_i:*

$$\mathsf{k}_i = \lambda^i_1{:}?(\alpha^i_1, \beta) \wedge T(d, j, \lambda^i_1) \wedge \ldots \wedge \lambda^i_{k_i}{:}?(\alpha^i_{k_i}, \beta) \wedge T(d, j, \lambda^i_{k_i})$$

Then

$$Update(\mathcal{M}_\sigma, \mathcal{K}_\beta) \models \psi \text{ iff } \forall s \in S^{\mathcal{M}_\sigma},$$
$$\mathcal{M}_\sigma, s \models [!(\mathcal{K}_\beta \wedge last = \beta)]^{cp}(\textstyle\bigwedge_{i=1}^n [!\mathsf{k}_\mathsf{i}]^{cp}\psi)$$

We now also axiomatise MDC from Definition 6 within the glue logic. We will take MDC as imposing a coherence order on ULFs ϕ and ψ that entail

[4] Given that $\Box$ verifies $\Box A \rightarrow A$, we should also check that if $A > B$ is true throughout $\mathcal{M}$ it is also true throughout $\mathcal{M}^{cp(\phi)}$. Recall that $A > B$ is true at s in $\mathcal{M}$ iff $*^{\mathcal{M}}(s, [\![A]\!]^{\mathcal{M}}) \subseteq [\![B]\!]^{\mathcal{M}}$. Abstracting away from s, we consider the relation between $*^{\mathcal{M}^\phi}[\![A]\!]^{\mathcal{M}^\phi}$ and $[\![B]\!]^{\mathcal{M}^\phi}$. $*^{\mathcal{M}^\phi}[\![A]\!]^{\mathcal{M}^\phi} \subseteq *^{\mathcal{M}}[\![A]\!]^{\mathcal{M}}$ and $[\![B]\!]^{\mathcal{M}^\phi} \subseteq [\![B]\!]^{\mathcal{M}}$ (this follows directly from Definition 7). For any $s \in *^{\mathcal{M}^\phi}[\![A]\!]^{\mathcal{M}^\phi}$, $s \in *^{\mathcal{M}}[\![A]\!]^{\mathcal{M}} \cap [\![\phi]\!]^{\mathcal{M}} \subseteq [\![B]\!]^{\mathcal{M}} \cap [\![\phi]\!]^{\mathcal{M}} = [\![B]\!]^{\mathcal{M}^\phi}$.

the old information $Th(\sigma)$ and the new information $\mathcal{K}_\beta$. So ϕ and ψ resolve some aspect of underspecified content in the discourse: e.g., an antecedent to a pronoun is fixed to a specific value, or a particular assumption about a rhetorical connection is made. Thus, without loss of generality, we can assume that ϕ and ψ are conjunctive formulae where at least one of the conjuncts is an assumption λ:?(α, β) about attachment of the new information. This is because for ϕ and ψ to describe a coherent logical form at all β must be attached to some available label α in $Th(\sigma)$ (and if ϕ and ψ differ in their assumptions about what β is attached to, then ϕ will contain a conjunct λ :?(α, β), say, and ψ the conjunct λ' :?(α', β), where $\alpha \neq \alpha'$). We can therefore re-cast the states ϕ and ψ to be ordered according to degree of coherence as ACPs containing at least one assumption about what β attaches to. We also assume a partial ordering $\leq$ on rhetorical relations: $R \leq R'$ means that R is a less coherent relation than R' (see clause 3 from Definition 6). For example, $\mathsf{Background} \leq \mathsf{Explanation}$ would make MDC prefer interpreting new information as an *Explanation* rather than as a *Background*, all else being equal.

The principles that govern degree of coherence are then stipulated in Definition 10. Clause 1 says that an ACP ψ that yields a less coherent rhetorical connection compared with the ACP ϕ is normally less coherent. Clause 2 says that an ACP ψ that resolves fewer underspecified elements than ACP ϕ is normally less coherent. Finally, clause 3 says that an ACP ψ that results in a logical form with more segments than ACP ϕ is normally less coherent. To express clauses 3. requires us to extend the object language by introducing formulas $Labels_n$ that are true at a state just in case the SDRS that is that state has at least n labels for discourse units.[5] To express clause 2. we must deal with 'degrees' of underspecification. Underspecifications yield disjunctions of fully-specific descriptions in a DSDRS; the more underspecification, the more disjunctions in the disjunctive normal form (DNF) of the DSDRS-description. We thus need to extend the language with terms $Indeterminates_n$ and allow states to contain underspecified variables as well as normal SDRS conditions. Then $Indeterminates_n$ is true at a state s iff S contains at most n indeterminates.[6]

DEFINITION 10. **Coding up MDC**

1. $(Th(\sigma) \wedge \mathcal{K}_\beta) > \psi \prec \phi$ *if there's a permutation* f *on* $\Pi_{Th(\sigma)} \cup \Pi_{\mathcal{K}_\beta}$ *st* $R' \leq R \wedge \forall \pi_1 \pi_2 \in \Pi_{Th(\sigma)} \cup \Pi_{\mathcal{K}_\beta}$, $[!\psi]^{cp} R'(\pi_1, \pi_2) \rightarrow [!\phi]^{cp} R(f(\pi_1), f(\pi_2)))$

2. $(Th(\sigma) \wedge \mathcal{K}_\beta) > \psi \prec \phi$ *if* $[!\phi]^{cp}\mathrm{Indeterminates}_n \rightarrow [!\psi]^{cp}\mathrm{Indeterminates}_n$

[5] We will need axioms of the form $(\pi_1 : \mathcal{K}_1 \wedge \ldots \wedge \pi_n : \mathcal{K}_n) \leftrightarrow Labels_n$.

[6] The permutations are equally bounded over the finite number of labels in the DSDRS, and so all of the formulas in Definition 10 reduce to quantifier-free, though unwieldy, formulas.

3. $(Th(\sigma) \wedge \mathcal{K}_\beta) > \psi \prec \phi$ *if* $[!\psi]\text{Labels}_n \rightarrow [!\phi]\text{Labels}_n$

This is still only an approximation of MDC from Definition 6. We have not, for instance, encoded the principle that interpretations with more rhetorical connections are more coherent than those without (see clause 1 from Definition 6). But given that the number of rhetorical relation symbols in $\mathcal{L}$ is finite, it would be very straightforward to express this factor via the ordering relation $\leq$ on predicate symbols in $\mathcal{L}_{ulf}$.

The axioms in Definition 10 can be used to influence interpretation. We do this by adding to our existing update function *Update*$(\mathcal{M}_\sigma, \mathcal{K}_\beta)$, which abstracts away entirely from how interpretation is influenced by degree of coherence, a new update function *Best-update*$(\mathcal{M}_\sigma, \mathcal{K}_\beta)$. Its definition is exactly like that of *Update*, save that the consequences of the announcement $[!\mathcal{K}_\beta \wedge \textit{last} = \beta]^{cp}$ are restricted to those conjuncts about attachment that are maximal on the partial ordering $\prec$ given by Definition 10. Thus a speaker can anticipate what the most coherent interpretation of his announcement will be, as well as the range of possible coherent interpretations, as given by *Update*.

Now, suppose that the consequences of the ACP ϕ is a logical form with more segments than that which follows from the ACP ψ, but ψ's consequences feature lower quality rhetorical connections. Then the default axioms in Definition 10 whose antecedents are satisfied will conflict, with one having the consequence $\phi \prec \psi$ and the other having the conflicting consequence $\psi \prec \phi$. This results in a Nixon Diamond within the logic of the conditional $>$, and hence no inferences about the relative coherence of the announcements ϕ and ψ. So just like Definition 6 of MDC, the axioms in Definition 10 validate a default inference that $\psi \prec \phi$ only if ϕ is at least as coherent as ψ in all three respects, and more coherent in at least one of them.

5. The Complexity of the Dynamic Glue Logic

Our version of PAL has **reduction axioms and rules** that reduce our PAL extension of GL back down to GL's base logic, described in Section 3. The reduction axioms for the $[!\phi]$ operator are quite standard (though we offer a proof for the reduction axiom for $>$).

I $[!\phi]p \leftrightarrow (\phi \rightarrow p)$

II $[!\phi](\psi \wedge \chi) \leftrightarrow ([!\phi]\psi \wedge [!\phi]\chi)$

III $[!\phi]\neg\psi \leftrightarrow (\phi \rightarrow \neg[!\phi]\psi)$

IV $[!\phi](\psi > \chi) \leftrightarrow (\phi \rightarrow ((\phi \wedge [!\phi]\psi) > [!\phi]\chi))$

V $[!\phi]\Box\psi \leftrightarrow (\phi \rightarrow \Box[!\phi]\psi)$

Proof for Axiom IV. Note that $[\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}} = [\![\psi]\!]^{\mathcal{M}^\phi}$, since $s \in [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}}$ iff $s \in [\![\phi]\!]^{\mathcal{M}} \cap [\![[!\phi]\psi]\!]^{\mathcal{M}}$ iff $s \in [\![\psi]\!]^{\mathcal{M}^\phi}$. To prove Axiom IV, either $\mathcal{M}, s \models \phi$ or not. If not, then both the left and right hand sides of Axiom IV are true. So assume now that $\mathcal{M}, s \models \phi$. Observe that by definition,

$$*^{\mathcal{M}^\phi}(s, [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}}) \subseteq *^{\mathcal{M}}(s, [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}})$$

And so,

$$*^{\mathcal{M}^\phi}(s, [\![\psi]\!]^{\mathcal{M}^\phi}) \subseteq *^{\mathcal{M}}(s, [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}})$$

This suffices to prove Axiom IV, since given that $\mathcal{M}, s \models \phi$, the following holds:

$$\begin{aligned}
\mathcal{M}^\phi, s \models \psi > \chi \text{ iff } & *^{\mathcal{M}^\phi}(s, [\![\psi]\!]^{\mathcal{M}^\phi}) \subseteq [\![\chi]\!]^{\mathcal{M}^\phi} \\
\text{iff } & *^{\mathcal{M}}(s, [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}}) \subseteq [\![\chi]\!]^{\mathcal{M}^\phi} \\
\text{iff } & *^{\mathcal{M}}(s, [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}}) \subseteq [\![\phi \wedge [!\phi]\chi]\!]^{\mathcal{M}} \\
\text{iff } & *^{\mathcal{M}}(s, [\![\phi \wedge [!\phi]\psi]\!]^{\mathcal{M}}) \subseteq [\![[!\phi]\chi]\!]^{\mathcal{M}} \\
\text{iff } & \mathcal{M}, s \models (\phi \wedge [!\phi]\psi) > [!\phi]\chi
\end{aligned}$$

(The penultimate step follows again because of the reflexivity of $*$ in GL). $\Box$.

We note that Axiom IV is equivalent to one given by van Benthem [2007] and similar to that given by Baltag and Smets [2006]. Nevertheless, our proof of it is somewhat different from theirs. They both use an ordering exogenous to the particular theory whose nonmonotonic consequences one is computing; they need to assume this ordering is both reflexive and transitive to prove Axiom IV. In our work on generics and discourse interpretation, we have found it difficult if not impossible to specify an ordering independent of the facts given by the discourse context; and for this very reason, we made the discourse context and the background theory together provide the ordering. The way we have defined $>$ and $\mid\!\sim$ as distinct but interrelated concepts allows us to do this. As a consequence our proof of Axiom IV requires only that the selection function underlying the conditional $>$ be reflexive; in fact we don't want it to be transitive.[7] We think it an advantage that roughly the same reduction axiom can be arrived at by quite different means.

The more interesting question concerns the reduction axioms for ACPs. Using the prime implicate $\mathfrak{I}_\phi$, which are just Boolean conjunctions of literals of the background theory and whose size is no bigger than the number of literals in the background theory [Asher, 1995], we get reduction axioms for $[!\phi]^{cp}\psi$ that are exactly analogous to those for $[!\phi]\psi$ and which confirm the reduction of ACPs to normal announcements in PAL.[8] This is because

[7] For a discussion see Pelletier and Asher [1997].
[8] This is also not the case for van Benthem [2007] and Baltag and Smets [2006].

the prime implicate encapsulates the nonmonotonic reasoning inherent in $\mid\!\sim$. Equation (20) stipulates the connection between $\mathcal{M}$ and $\mathcal{M}^{cp(\phi)}$:

$$[\![\psi]\!]^{\mathcal{M}^{cp(\phi)}} = [\![\mathfrak{I}_\phi \wedge [!\phi]^{cp}\psi]\!]^{\mathcal{M}} \tag{20}$$

The relevant reduction axioms are Axioms VI to X:

VI $[!\phi]^{cp}p \leftrightarrow (\mathfrak{I}_\phi \rightarrow p)$

VII $[!\phi]^{cp}\neg\psi \leftrightarrow (\mathfrak{I}_\phi \rightarrow \neg[!\phi]^{cp}\psi)$

VIII $[!\phi]^{cp}(\psi \wedge \chi) \leftrightarrow ([!\phi]^{cp}\psi \wedge [!\phi]^{cp}\chi)$

IX $[!\phi]^{cp}(\psi > \chi) \leftrightarrow (\mathfrak{I}_\phi \rightarrow ((\mathfrak{I}_\phi \wedge [!\phi]^{cp}\psi) > [!\phi]^{cp}\chi))$

X $[!\phi]^{cp}\Box\psi \leftrightarrow (\mathfrak{I}_\phi \rightarrow \Box[!\phi]^{cp}\psi)$

Given the restrictions that hold of GL's background theory, prime implicates exist and can be decidably computed [Asher, 1995]. Furthermore, the base logic of GL is also decidable [Lascarides and Asher, 1993]. These reduction axioms thus ensure that our extension of PAL is decidable as well.

In principle we are happy with a simple decidability result. But one might still wonder what the complexity of dynamic GL is. While the reduction axioms guarantee the decidability of the question whether any formula ϕ follows from a set of premises in dynamic GL given that static GL is decidable, this does not put any further bounds on the complexity of the problem, because the reduction axioms may produce formulas that are exponentially longer than those in our PAL language [Lutz, 2006]. We can get a PSPACE complexity result by using the tableau system of [Balbiani et al., 2010], replacing their rules for K_i with the single S5 modality $\Box$.

THEOREM 1. *Given $\mathfrak{I}_\phi$, $\phi\mid\!\sim\psi$ for dynamic* GL *is at most* PSPACE *in complexity.*

THEOREM 2. *The amount of memory needed to calculate $\mathfrak{I}_\phi$ for dynamic* GL *is a* PSPACE *function of the (finite) background* GL *theory* Γ.

The proofs for Theorems 1 and 2 are given in the Appendix.

From Theorem 1 and Theorem 2, since a PSPACE applied to a PSPACE function of Γ is still a PSPACE function of Γ, Theorem 3 follows.

THEOREM 3. *The amount of memory needed to calculate $\phi\mid\!\sim\psi$ for dynamic* GL *is a* PSPACE *function of the background* GL *theory* Γ.

We close this section with some remarks on the consequences of these results. Theorem 1 together with the completeness proof translates into a more proper deduction theorem for our PAL extension of GL:[9]

THEOREM 4. $\vdash [!\phi]^{cp}\psi$ *iff* $\phi {\mid\!\sim} \psi$

Typically, such Deduction Theorems for nonmonotonic consequence relations are difficult to get and fragile.[10] One cannot 'drop' some of the information in the premises of the nonmonotonic consequence relation and preserve the equivalence. That is, something like (21) is not valid, for finite sets of premises Γ:

(21) $\Gamma {\mid\!\sim} [!\phi \wedge \bigwedge \Gamma]^{cp}\psi$ only if $\Gamma, \phi {\mid\!\sim} \psi$

(21) isn't valid because like many nonmonotonic logics ours 'suffers' from the Drowning Problem [Benferhat et al., 1993]—defaults from ϕ 'drown out' those from Γ when they are mixed together.[11] The problem comes from nested conditionals: in (22), $\Gamma {\mid\!\sim} [!\phi \wedge \bigwedge \Gamma]^{cp}\psi$, but $\Gamma, \phi \not{\mid\!\sim} \psi$—the defaults from ϕ 'drown out' those from Γ when they are mixed together.

(22) a. $\Gamma\colon \{A, A > D, A > ((B \wedge E) > C)\}$

b. $\phi\colon B \wedge E \wedge ((A \wedge E) > \neg D)$

c. $\psi\colon C$

To see that (22) is a counterexample to (21), observe first that $\Gamma {\mid\!\sim} (B \wedge E) > C$. We can also show that in $\mathcal{M}_0^{cp(\Gamma)}$ all states in fact verify $(B \wedge E) > C$.

[9] **Proof of Theorem 4.** To prove the deduction theorem, we need only to establish $\models [!\phi]^{cp}\psi$ iff $\phi {\mid\!\sim} \psi$, which we do now. Right to left direction: If $\phi {\mid\!\sim} \psi$, then the canonical model for the base logic GL $\mathcal{M}_0$ is such that for every state $s \in \mathcal{M}_0^{cp(\phi)}$, it is the case that $\mathcal{M}_0^{cp(\phi)}, s \models \psi$. Note that by Theorem 1, $\phi {\mid\!\sim} \psi$ entails that $\Im_\phi \to \psi$. So now consider any state s of a model $\mathcal{M}$. $\mathcal{M}, s \models [!\phi]^{cp}\psi$ iff if $\mathcal{M}, s \models \Im_\phi$, then $\mathcal{M}^{cp(\phi)}, s \models \psi$. If $\mathcal{M}, s \not\models \Im_\phi$, we are done. So assume $\mathcal{M}, s \models \Im_\phi$. Then $\mathcal{M}, s \models \psi$; so $\mathcal{M}^{cp(\phi)}, s \models \psi$.

The left to right direction: Let us look at the canonical model $\mathcal{M}_0$ and any state s such that $\mathcal{M}_0, s \models [!\phi]^{cp}\psi$. Assume ϕ is consistent. Then we can assume there are states in $\mathcal{M}_0$ that verify $\Im_\phi$ since $\Im_\phi$ is consistent if ϕ is, thanks to the properties of Commonsense Entailment (see Asher [1995] for details). Pick any such state s_0. Then by the definition of the semantics of ACPs, $\mathcal{M}_0^{cp(\phi)}, s_0 \models \psi$. But this is true of *every* state in $\mathcal{M}_0^{cp(\phi)}$, and so we have for every state s_1 in $\mathcal{M}_0^{cp(\phi)}$, $\mathcal{M}_0^{cp(\phi)}, s_1 \models \psi$. So $\phi {\mid\!\sim} \psi$.

□

[10] The strengthening in Corollary 1 of the left hand side of the biconditional in Theorem 4 has no effect, since ${\mid\!\sim}$ has a fixed point definition.

COROLLARY 1. $\phi {\mid\!\sim} [!\phi]^{cp}\psi$ *iff* $\vdash [!\phi]^{cp}\psi$

[11] Actually, we don't think this is a problem for the logic but rather a virtue. But justifying this is a matter for another time.

Now consider the formula $[!\phi \wedge \bigwedge \Gamma]^{cp}\psi$ in any such state s of $\mathcal{M}_0^{cp(\Gamma)}$. Such states will not satisfy $\Im_{\phi\wedge\bigwedge\Gamma}$, since they must satisfy D and $\Im_{\phi\wedge\bigwedge\Gamma} \rightarrow \neg D$. And so they trivially satisfy $[!\phi \wedge \bigwedge \Gamma]^{cp}\psi$. Thus, $\mathcal{M}_0^{cp(\Gamma)} \vdash [!\phi \wedge \bigwedge \Gamma]^{cp}\psi$. However, $\Gamma, \phi \not\mid\!\sim \psi$ since $\Gamma, \phi \not\mid\!\sim (B \wedge E) > C$.

The following apparent strengthening of Theorem 4 does not hold:

(23) $\Gamma \vdash [!\phi]^{cp}\psi$ only if $\Gamma, \phi \mid\!\sim \psi$

To see why, choose a Γ such that any state verifying Γ perforce fails to satisfy $\Im_\phi$. Then, $\Gamma \vdash [!\phi]^{cp}\psi$ for arbitrary ψ, but clearly if ϕ is consistent, then there is a ψ such that $\Gamma, \phi \not\mid\!\sim \psi$.

6. Conclusion

In this paper we have made SDRT's glue logic for computing the logical form of conversation dynamic. We have done this by extending a dynamic public announcement logic (PAL, Baltag et al. [1999]) with the default conditional $>$ from commonsense entailment [Asher and Morreau, 1991, Asher, 1995], thereby allowing us to import into a dynamic logic SDRT's nonmonotonic axioms for inferring rhetorical connections from linguistic knowledge and contextual information [Asher and Lascarides, 2003]. Making the glue logic dynamic allows a dialogue agent to reason about what the update of the content of the dialogue will be after his contribution, including the effects of his candidate rhetorical moves. This is a pre-requisite for *planning* one's next move, but so is reasoning about attitudes like preferences—a major research issue that we don't address here.

In making a basic PAL suitable for reasoning about dialogue content, we have endowed basic PAL with a pleasing notion of the defeasible consequences of an announcement, indicating that logicians may find interesting extensions of their theories from those that are developed to achieve tasks within formal semantics and pragmatics. Finally, we have also shown that by using the resources of commonsense entailment, our PAL extension for supporting default reasoning is decidable. In fact its complexity is no greater than that of PAL with multiple knowledge operators. And unlike previous default PALs it also supports intuitively compelling patterns of inference like the Penguin Principle without having to assume inference orderings that are extraneous to the logic.

The next step is to examine SDRT's other shallow logic, the logic of cognitive modelling, so as to optimise the trade offs between expected interpretations, as modelled in this paper, and speaker preferences. This interaction will complicate the move from $\mathcal{M}$ to $\mathcal{M}^{cp(\phi)}$. Asher and Lascarides [2008], Asher et al. [2010] provide the first steps towards this major task, but further progress requires a deeper understanding of human decision making in a wide variety of cooperative and non-cooperative dialogue domains. We plan to address this in future work.

Appendix

To prove Theorem 1, we first note a lemma that links $\phi {\mid\!\sim} \psi$ to monotonic validity (see Asher [1995] for its proof):

LEMMA 1. $\models \mathfrak{I}_\phi \rightarrow \psi$ *iff* $\phi {\mid\!\sim} \psi$

Thus, what we must do to capture ${\mid\!\sim}$ in dynamic GL is to provide a complete axiomatisation for the monotonic validity notion in dynamic logic. This we do by adapting the tableau procedure from [Balbiani et al., 2010]. Using their notation, we explain now the tableau rule for $\phi > \psi$. Given $\langle(\psi_1, \ldots \psi_n), n, (\phi > \chi)\rangle \in L$ and $(\phi, n, n') \in S$—i.e. given a labelled formula which consists in a sequence of PAL updates $\psi_1, \ldots \psi_n$ that get us to a node n at which we have the formula $\phi > \chi$, and the node n' which is a ϕ-normal node given n, we have $B = \{\langle L \cup \{\langle \epsilon, n', [\psi_1] \ldots [\psi_n](\phi \rightarrow \chi)\rangle\}, S\rangle\}$—i.e. the set of branches B created by the rule consists in just one in which $[\psi_1] \ldots [\psi_n](\phi \rightarrow \chi)$ is true at node n'. The form of the rule is then:

$$\frac{\langle(\psi_1, \ldots \psi_n), n, (\phi > \chi)\rangle \in L \text{ and } \langle \phi, n, n'\rangle \in S}{B = \{\langle L \cup \{\langle \epsilon, n', \langle[\psi_1] \ldots [\psi_n](\phi \rightarrow \chi)\rangle\}, S\rangle\}}$$

We have a special rule for $\neg(\phi > \chi)$:[12]

$$\frac{\langle(\psi_1, \ldots \psi_n), n, \neg(\phi > \chi)\rangle \in L}{B = \{\langle L \cup \{\langle \epsilon, n', \langle\psi_1\rangle \ldots \langle\psi_n\rangle(\phi \wedge \neg\chi)\rangle\}, S \cup \{\langle \phi, n, n'\rangle\}\rangle\} \text{ for some } n' \notin S}$$

These tableaux rules already verify axioms (10) and (12). We leave to the reader the construction of tableaux rules guaranteeing rule (11). The tableau rule to get the Specificity axiom (13) is:

$$\frac{\langle \mu, n, (\phi > \chi)\rangle \in L \text{ and } \langle \mu, m, (\Box(\psi \rightarrow \phi)\rangle \in L, \text{ and } \langle \mu, k, (\psi > \neg\chi)\rangle \in L}{B = \{\langle L - \{\langle \mu, n, (\phi > \chi)\rangle\} \cup \{\langle \mu, n, (\phi > (\neg\psi \wedge \chi)\rangle\}, S\rangle\}}$$

The 'Mix' axiom (14) is captured by:

$$\frac{\langle \mu, n, (\phi > \psi)\rangle \in L \text{ and } \langle \mu, m, \Box\chi\rangle \in L}{B = \{\langle L - \{\langle \mu, n, (\phi > \psi)\rangle\} \cup \{\langle \mu, n, (\phi > (\chi \wedge \psi))\rangle\}, S\rangle\}}$$

These tableau rules allow us to derive all the axioms of static modalised GL and are clearly sound with respect to the semantics. To prove completeness,

[12] corresponding to the rule for $\hat{K}$ in Balbiani et al. [2010].

we add saturation rules for $>$ as Balbiani et al. [2010] do for their operators K_a. The discussion in Balbiani et al. [2010] shows how to construct a model for any satisfiable formula ϕ. Given that ACP introduces no new complexities into the logic of PAL assuming that we have the computation of $\Im_\phi$, this tableau system suffices for the completeness (of the monotonic logic) of dynamic GL.

To prove the PSPACE result, we adapt slightly the algorithmic procedure of [Balbiani et al., 2010] to take account of the tableau rules for $>$. If the tableau is not already closed or saturated, the algorithm must try to 'saturate' the right hand sides of $>$-statements using the tableau rule for Specificity and then mark the premises used. The rest of algorithm proceeds as described in [Balbiani et al., 2010], and we get a termination result for the algorithm. Using the space saving strategies for the tableau system described in [Balbiani et al., 2010], the PSPACE result for $\mathrel{|\!\!\sim}$ follows.

The procedure described above depends on the availability of the prime implicate $\Im_\phi$ for announcements ϕ. To compute $\Im_\phi$, we need to compute the formula from which we can deduce in the monotonic logic all the nonmonotonic consequences of ϕ together with the background truths of GL. Note that ϕ in these cases will be modally free, and the background truths do not mention announcements, as they are axioms for computing discourse relations and resolving underspecifications. But this means that we must sketch an algorithm for computing $\mathrel{|\!\!\sim}$ with finite sets of premises.

To do the latter, we give more detail about the procedure for establishing when a finite set of premises Γ nonmonotonically implies ψ, or $\Gamma \mathrel{|\!\!\sim} \psi$. $\mathrel{|\!\!\sim}$ is defined using the notion of an extension sequence relative to an ordering ζ over the antecedents of $>$-formulae in Γ. We here build finite extension sequences, which we can do, because our starting point Γ is finite. We use what Asher [1995] called a *prime implicate* (perhaps he should have called it *prime implicant*), which is a finite formula.

DEFINITION 11. *Φ is the prime implicate of a set S iff for all $\psi \in S$, $\vdash \Phi \rightarrow \psi$.*

We can produce a formula entailing all the formulas that are normally true when some formula ϕ is true in the following way. Given a finite set of formulas Γ, and a formula ψ that is an antecedent of a conditional in Γ, we first saturate Γ under applications of Specificity to get all the relevant $>$-statements. That is, if we have for subformulas of Γ ϕ, ψ, and χ, $\Gamma \vdash (\phi > \psi) \wedge (\chi > \neg\psi) \wedge \Box(\phi \rightarrow \chi)$, then we add $\chi > \neg\phi$. Call the result Γ^+. Using the property of Closure on the Right for $>$ (i.e., axiom (12), we now conjoin all the subformulas of Γ that are $>$-consequences of Γ. This the prime implicate of $\{\psi : \Gamma \vdash \phi > \psi\}$. The size needed for the computation of the prime implicate is at most a PSPACE function of the size of Γ given Theorem 1.

DEFINITION 12. *Finite Extension Sequences*

$$\begin{array}{lcl} Ext^0(\Gamma,\zeta) & = & \Gamma \\ Ext^{\beta+1}(\Gamma,\zeta) & = & Ext^{\beta}(\Gamma,\zeta) \cup \{\psi \rightarrow \Phi : \Phi \textit{ is the prime implicate of} \\ & & \qquad \{\varphi : Ext^{\beta}(\Gamma,\zeta) \vdash \psi > \varphi\}\}, \\ & & \textit{where } \zeta(\psi) = n \textit{ and } \beta + 1 = \lambda + mn, \\ & & \textit{provided } Ext^{\beta}(\Gamma,\zeta) \textit{ is consistent with } \Phi. \\ & = & Ext^{\beta}(\Gamma,\zeta), \textit{otherwise}. \\ Ext^{\lambda}(\Gamma,\zeta) & = & \cup_{\beta<\lambda} Ext^{\beta}(\Gamma,\zeta), \textit{ for limit ordinals } \lambda. \end{array}$$

Note that the formula $\psi \rightarrow \Phi$ computed for each $\mathrm{Ext}^{\beta}(\Gamma,\zeta)$ is the prime implicate of that stage.

Every such finite extension sequence has a fixpoint.

DEFINITION 13. *The fixpoint of a finite extension sequence of* Γ *relative to ordering* ζ *is a* $\Gamma\zeta$ *extension.*

DEFINITION 14. $\Gamma \mid\!\sim \phi$ *iff for all orderings* ζ *the fixpoint* $\Gamma*$ *of each* $\Gamma\zeta$ *extension sequence is such that* $\Gamma* \vdash \phi$

For a finite theory, Asher [1995] shows that the fixpoint of any extension sequence occurs in at most $nm + 1$ stages where n is the number of $>$-antecedents in Γ and m is the greatest depth of $>$-embedding in Γ. If we simply enumerate all possible orderings for $>$ antecedents where we may consider the same antecedent more than once, we get $(nm + 1)!$ extension sequences, which would clearly exceed PSPACE bounds if we had to keep all extensions in memory. However, we need not do this. We can build two extensions E_1 and E_2 relative to two of the possible orderings, and then take the common entailments $E*$ of E_1 and E_2. We now construct a third extension and take the common entailments of the new extension with $E*$. We repeat this proceedure until we have examined all the extensions and all the relevant orderings. In this way, we keep only at most 3 sets of formulas in memory at a time, and all of these sets are linear in the size of Γ. We also need space to calculate each extension and this calls for a satisfiability test in the base logic and the computation of a prime implicate. But we can erase each calculation of satisfiability as well as the calculation of the prime implicate needed to construct an extension stage after they have been performed. Thus, this calculation requires only an at most PSPACE amount of memory.

So we can fix the upper bound of the space needed to calculate $\Im_\Gamma$ to be that given in equation (24), where λ is the linear parameter giving the size of the prime implicate of the extension of Γ relative to Γ, $f(|\Gamma|)$ is an at most PSPACE function needed to construct the prime implicate of any stage of any extension sequence (in the base modal logic with a single S5 modality $\Box$ and $>$), and $spsat(\lambda|\Gamma|)$ is the maximal amount of space needed to check the satisfiability in the monotonic logic of GL of any stage of an extension sequence of Γ.

$$\text{Max}\{spsat(|\Gamma|), f(|\Gamma)|\} + 3\lambda(|\Gamma|) \tag{24}$$

(24) is a PSPACE function of the size of Γ. This establishes Theorem 2.

This proof of the result has a syntactic restriction on ϕ to modally free formulas (those that figure in announcements). However, we can relax this assumption to allow ϕ to contain formulas of the form $A > B$ as subformulas without changing the proof. Can announcements contain other ACPs? It seems so, and it also seems that we can apply the computation of $\Im_\phi$ recursively using the background truths of GL. However, this too doesn't harm the PSPACE result as we can erase the space needed to compute $\Im_\phi$ at level n to reuse it to calculate $\Im_\psi$ at level $n+1$. Thus Theorem 3 generalises to an unrestricted ϕ in dynamic GL.

References

P. Achinstein. *The Nature of Explanation*. Oxford University Press, 1980.

CE Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meeting contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

N. Asher. *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers, 1993.

N. Asher. Commonsense entailment: A conditional logic for some generics. In G. Crocco, L. Farinas, and A. Herzig, editors, *Conditionals: From Philosophy to Computer Science*, pages 103–145. Oxford University Press, Oxford, 1995.

N. Asher and A. Lascarides. Indirect speech acts. *Synthese*, 128(1–2):183–228, 2001.

N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.

N. Asher and A. Lascarides. Commitments, beliefs and intentions in dialogue. In *Proceedings of the 12th Workshop on the Semantics and Pragmatics of Dialogue (Londial)*, pages 35–42, London, 2008.

N. Asher and M. Morreau. Commonsense entailment. In John Mylopoulos and Raymond Reiter, editors, *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 387–392, Los Altos, California, 1991. Morgan Kaufmann.

N. Asher, E. Bonzon, and A. Lascarides. Extracting and modelling preferences from dialogue. In *Proceedings of the International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems (IPMU)*, pages 542–553, Heidelberg, 2010.

P. Balbiani, HP. van Ditmarsch, A. Herzig, and T. de Lima. Tableaux for public announcement logic. *Journal of Logic and Comuptation*, 20(1): 55–76, 2010.

A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. In G. Mints and R. de Queiroz, editors, *Electronic Notes in Theoretical Computer Science, 165*, pages 5–21. 2006.

A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. Technical Report SEN-R9922, Centrum voor Wiskunde en Informatica, 1999.

S. Benferhat, D. Dubois, and H. Prade. Argumentative inference in uncertain and inconsistent knowledge bases. In D. Heckermann and A. Mamdani, editors, *Proceedings of the 9th International Conference on Uncertainty in Artificial Intelligence*, pages 411–419. Morgan Kaufmann, 1993.

L. Benotti. *Implicature as an Interactive Process*. PhD thesis, Université Henri Poincaré, INRIA Nancy Grand Est, France, 2010.

S. Bromberger. An approach to explanation. In R. Butler, editor, *Analytical Philsophy*, pages 72–105. Oxford University Press, 1962.

M. Egg, A. Koller, and J. Niehren. The constraint language for lambda structures. *Journal of Logic, Language, and Information*, 10:457–485, 2001.

B. Grosz and C. Sidner. Attention, intentions and the structure of discourse. *Computational Linguistics*, 12:175–204, 1986.

J. R. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142, 1993.

J.R. Hobbs. Coherence and coreference. *Cognitive Science*, 3(1):67–90, 1979.

H. Kamp. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*, pages 277–322. Mathematisch Centrum, Amsterdam, 1981.

H. Kamp and U. Reyle. *From Discourse to the Lexicon: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Publishers, 1993.

A. Kehler. *Coherence, Reference and the Theory of Grammar*. CSLI Publications, Cambridge University Press, 2002.

A. Lascarides and N. Asher. Temporal interpretation, discourse relations and commonsense entailment. *Linguistics and Philosophy*, 16(5):437–493, 1993.

A. Lascarides and N. Asher. Agreement, disputes and commitment in dialogue. *Journal of Semantics*, 26(2):109–158, 2009.

K. Lutz. Complexity and succintness of public announcement logic. In P. Stone and G. Weiss, editors, *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 137–144, 2006.

J. Pelletier and N. Asher. Generics and defaults. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 1125–1177. MIT Press, 1997.

M. Poesio and D. Traum. Towards an axiomatisation of dialogue acts. In J. Hulstijn and A. Nijholt, editors, *Proceedings of the Twente Workshop on the Formal Semantics and Pragmatics of Dialogue*. 1998.

H. Sacks, E. A. Schegloff, and G. Jefferson. A simplest systematics for the organization of turn-taking in conversation. *Language*, 50(4):696–735, 1974.

M. Stone. *Modality in Dialogue: Planning, Pragmatics and Computation*. PhD thesis, University of Pennsylvania, 1998.

M. Stone and R. Thomason. Context in abductive interpretation. In *Proceedings of the 6th International Workshop on the Semantics and Pragmatics of Dialogue (Edilog)*, Edinburgh, 2002.

J. van Benthem. Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.

R. van Rooij. Relevance of communicative acts. In *Proceedings of Theoretical Aspects of Rationality and Knowledge*, pages 83–97, 2001.

JD Williams and SJ Young. Partially Observable Markov Decision Processes for Spoken Dialog Systems. *Computer Speech and Language*, 21(2):231–422, 2007.