

A Formal Characterisation of Hamblin's Action-State Semantics

Chris Reed (chris@computing.dundee.ac.uk)

School of Computing, University of Dundee, Dundee DD1 4HN, Scotland UK

Timothy J. Norman (tnorman@csd.abdn.ac.uk)

Department of Computing Science, University of Aberdeen, Aberdeen AB24 3UE, Scotland UK

Abstract. Hamblin's Action State Semantics provides a sound philosophical foundation for understanding the character of the imperative. Taking this as our inspiration, in this paper we present a logic of action, which we call **ST**, that captures the clear ontological distinction between being responsible for the achievement of a state of affairs and being responsible for the performance of an action. We argue that a relativised modal logic of type **RT** founded upon a ternary relation over possible worlds integrated with a basic tense logic captures intuitions of the Hamblinian model of imperatives. The logic implements a direct mapping of each of Hamblin's key concepts: strategies, partial strategies and wholehearted satisfaction.

1. Introduction

Our aim is to develop a logical theory of action that supports a number of interesting features: (1) an intuitive characterisation of the concepts of responsibility and delegation through imperatival utterance; (2) a semantic underpinning based on Hamblin's (1987) model (rather than upon more conventional branching time logic approaches or their variants); (3) a distinction between bringing about states of affairs and executing actions that echoes von Wright's (1968) distinction between *Seinsollen* and *Tunsollen*; and (4) an integration with a simple tense logic.

Our motivation follows a similar pattern: delegation and responsibility for action do not, to our knowledge, have elsewhere an adequate formal foundation that can account for the notion of fulfilling an imperative. It is our contention that such a foundation needs to distinguish between achieving and doing, and Hamblin's Action State Semantics offers an elegant underpinning for such a distinction. Finally, any such implementation will also typically need to account for a temporal component, and again Hamblin's model provides a means for developing such an account. We aim to strike a balance in this research between demonstrating formal rigour on the one hand and simplicity in explication on the other, in an attempt to facilitate broader appeal across



© 2006 Kluwer Academic Publishers. Printed in the Netherlands.

theoretical and applied research areas. This presentation, therefore, assumes nothing more than an understanding of elementary modal logic.

We take as our point of departure (Chellas, 1992). Though numerous more recent works have developed aspects of the logical theories of action, of knowledge, and of time (many of which we shall introduce as appropriate), it is Chellas's paper that lays out the greatest part of our agenda here. Specifically, the remainder of this paper proposes a model of action coupled with a tense logic that employs a simple semantic model and supports a clear characterisation of responsibility which in turn allows for a rich and flexible model of delegation between agents. The first component of this model is a Kripkean-style characterisation of Hamblin's rich semantic foundation: action-state semantics.

2. Action-State Semantics

Numerous proposals have been laid out in both philosophical and computational literature for the classification of utterance types, or, more specifically, of illocutionary acts. Austin (1962, p. 150) and Searle (1969, pp. 66–67) are perhaps the two most prominent.

Although widely available, these schemes have at least one thing in common: not all utterances are indicative. This is not in itself remarkable, until it is considered that the logics employed to handle and manipulate utterances are almost always exclusively based upon the predominant formal tradition of treating only the indicative. The interrogative and imperative utterances (which figure amongst Austin's Exercitives and Expositives, and include Searle's Request, Question, Advise and Warn) rarely benefit from the luxury of a logic designed to handle them.

Interrogative logics for handling questions have been proposed by Åqvist (1975) and Hintikka *et al.* (1996) among others, and these form an interesting avenue for future exploration. The focus of the current work, however, is on imperative logic. Hamblin's (1987) book *Imperatives* represents the first thorough treatment of the subject, providing a systematic analysis not only of linguistic examples, but also of grammatical structure, semantics and the role imperatives play in dialogue.

Hamblin states (1987, p. 137) that to handle imperatives there are several features, "usually regarded as specialised", which are indispensable for a formal model: (1) a time-scale; (2) a distinction between actions and states; (3) physical and mental causation; (4) agency and action-reduction; and (5) intensionality. Following the second feature

listed above, both events and states of affairs are explicitly represented in the Action State Semantics: a world is a series of states connected by events. The states can be seen as collections of propositions. Events are of two types: deeds, which are performed by specific agents, and happenings, which are world effects.

Hamblin notes that the model is unusually 'lavish', commenting that, "It would be possible, for some purposes, to conceive the world less lavishly, either as a *world of states* or, say, a *world of events*, where events are deeds plus happenings. Neither of these more restricted conceptions would be adequate for the representation of imperatives, [...] but the point should be made, perhaps, that most logical models are built on one or other of them." (1987, p. 144).

This rich underlying model is important in several respects. Firstly, it has the syntactic structures that support the expression of imperatives concerned with agents bringing about states of affairs as well as imperatives concerned with agents performing actions. Secondly, it avoids both ontological and practical problems of having to interrelate states and events. Ontologically, the relationship between timeless states, instantaneous and non-instantaneous events, and the time-line along which they proceed, is an ancient philosophical problem going back at least as far as the Stoics, and manifest more recently in the tense logic revolution of the 1960s. To choose between conceiving of the world as sets of states, and conceiving of the world as sets of events, is to demand that a model of causality (and possibly deterministic causality) be built in to both the conception and the logic, from the start. Further practical problems often become manifest in both theory and subsequent application of theory (in computational systems, for example) as having to keep track of 'done events' in every state (Dignum and Meyer, 1990). If all that is required is to answer queries relating to what single, primitive, acts have just been done by what agents, there is no practical problem. This problem, however, becomes manifest when the system requires that answers to general queries of whether it is true that some arbitrary (partial) program has been executed in order to attain this state of affairs. In this case, all past acts by all agents for all time must be captured within the model of each state. Managing this information and associated truth maintenance machinery is cumbersome from both theoretical and implementation perspectives. Finally, this construction of a world as a chain of states connected by deeds and happenings makes it possible to distinguish those worlds in which a given imperative i is satisfied (in some set of states). Thus the imperative "Shut the door!" is satisfied in those worlds in which the door is shut (given appropriate deixis).

This model is then used by Hamblin in developing an account that contrasts the ‘extensional’ satisfaction with a stronger notion, of ‘wholehearted’ satisfaction, which characterises an agent’s involvement and responsibility in fulfilling an imperative. The remainder of this section provides a précis of Hamblin’s account, summarising his description, but preserving both the spirit and notational detail of (1987, Ch. 4).

Wholehearted satisfaction is based upon the notion of a *strategy* for an imperative i . A strategy for a particular agent is the assignment of a deed to each time point. A partial i -strategy is then a set of incompletely specified strategies, all of which involve worlds in which i is extensionally satisfied. The wholehearted satisfaction of an imperative i by an agent x is then defined as being x ’s adoption of a partial i -strategy and the execution of a deed from that strategy at every time point after the imperative is issued.

A Hamblinian world $w \in W$ is defined such that for every time point in T there is:

1. a state from the set of states S ,
2. a member of the set H of ‘big happenings’ (each of which simply collects together all happenings from one state to the next), and
3. a deed (in D) for every agent (in X), i.e. an element from D^X .

The set W of worlds is, therefore, defined as $(S \times H \times D^X)^T$. The states, happenings and deed-agent assignments of a given world w are given by $S(w)$, $H(w)$ and $D(w)$.

The next step is to let j_t be a history of a world up to time t , including all states, deeds and happenings of the world up to t . Thus j_t is equivalent to a set of worlds which have a common history up to (at least) time t . J_t is then the set of all possible histories up to t ; i.e. all the ways in which the universe could have evolved up to time t . A strategy q_t is then an allocation of a deed to each $j_{t'} \in J_{t'}$ for every $t' \geq t$.¹ Q_t then denotes the set of all possible strategies at time t .

Let the possible worlds in which the deeds of agent x are those specified by strategy q_t be $W_{\text{strat}}(x, q_t)$, and the worlds in which an imperative, i , is extensionally satisfied be W_i . A strategy for the satisfaction of an imperative i (i.e. an i -strategy) can, therefore, be defined as follows: A strategy $q_t \in Q_t$ is an i -strategy for agent x if and only if the worlds in which x does the deeds specified by q_t are also worlds in which i is extensionally satisfied: $W_{\text{strat}}(x, q_t) \subseteq W_i$.

¹ This notion of a strategy has an intensional component, since it prescribes over a set of possible w , rather than picking out, at this stage, the actual world.

In practice, however, it is not feasible for an agent to select a particular strategy in Q_t at time t that specifies every deed for every time t' after t . For this reason, an agent will adopt a *partial i -strategy*. A partial i -strategy is a disjunction of i -strategies, $Q'_t \subseteq Q_t$, and the set of worlds in which x adopts this partial i -strategy is $W_{\text{strat}}(x, Q'_t)$.

With this grounding, the wholehearted satisfaction of an imperative, i , can now be defined. An agent x may be said to wholeheartedly satisfy an imperative i issued at t if and only if for every $t' \geq t$:

1. x has a partial i -strategy, $Q'_{t'}$; and
2. x does a deed from the set of deeds specified by that $Q'_{t'}$.

For further details, the reader is referred, of course, to Hamblin's original monograph (Hamblin, 1987), and also to (Walton and Krabbe, 1995), that provides more detail on the role of such a model in the wider context of dialogue and, in its appendix, a more complete set-theoretic précis of Hamblin's model.

The summary outlined in this section sets our goalposts. Our objective is to develop a theory of agentive activity that is sufficient for modelling imperatives in the way that Hamblin has laid out. First, an axiomatisation of two new modalities S and T is presented. Second, we develop a semantics for these operators using a Kripkean-style characterisation of Action State Semantics. We then return to Hamblin's model of the imperative and show how it is captured within the Kripkean framework. Finally, we compare and contrast our semantics with existing models of action.

3. The axioms of a model of action and time

The syntactic and axiomatic presentation is divided into three parts. First, we treat the action components, S and T, then we go on to combine them with the temporal components, and, thirdly, we address the issues of delegation and responsibility by discussing relevant theorems and axioms of the language ST. Before exploring the axioms of our action modalities S and T, we specify the notation used within this paper and detail the well-formed formulae of our language.

We use the modalities S and T to represent 'being responsible for the achievement of a state of affairs', and 'being responsible for the performance of an action' respectively. These modalities are relativised to specific agents (x, y, \dots) and we adopt the convention of using upper case Roman letters (A, B, C, \dots) to stand for arbitrary states of affairs and lower case Greek letters to stand for arbitrary actions

($\alpha, \beta, \gamma, \dots$). Where it is important to do so, agent x is associated with action α in the following manner: α^x . Where actions are not thus specified, it is assumed that the agent is unbound (that is, the action is carried out by some agent, but it is not important which one).² In this way, $S_x A$ refers to agent x being responsible for the achievement of the state of affairs A , and $T_x \alpha$ refers to x being responsible for the performance of action α . Note that in both cases (and particularly the second, which is strongly suggestive), the statements do not specify any particular action for agent x itself. So for example, x may ensure that α is executed by ordering some other agent to carry it out. We may, however, capture the imperative that agent x is responsible for x doing α thus: $T_x \alpha^x$.

In specifying well-formed formulae of our language, we divide these basic atoms into two classes: (i) those that consist entirely of propositional expressions of action (bound or unbound), which we term *event formulae*, and (ii) all others, which we term *state formulae*. All such basic atoms are *wffs*. Thence by conventional *PL*, for any two *wffs*, ϕ and ψ , that are state formulae, $\phi \vee \psi$, $\phi \wedge \psi$, $\phi \rightarrow \psi$ and $\neg \phi$ are also *wffs* that are state formulae. Similarly for any two *wffs* that are event formulae, any *PL* combination of them is also a *wff* that is an event formula. For the action modalities, any *wffs* that are event formulae can be used to form a further *wff* with the T modality: $T_x \alpha$, $T_x \alpha^x$, $T_x \alpha^y$, $T_x(\alpha \vee \beta^y)$, etc., which are themselves state formulae. Any *wff* that is a state formula can be used to form a further *wff* with the S modality $S_x A$, $S_x(A \vee B)$, $S_x T_y \alpha^z$, etc.

To emphasise, note therefore that expressions such as $\alpha \wedge A$ are not well formed: descriptions of actions and states cannot be combined in a single statement of propositional logic. Responsibility for such things, however, are expressions of state and can therefore be combined in an intuitive manner: $T_x \alpha \wedge S_x A$, and so on.

Finally, well formed temporal expressions are synthesised easily, with any *wff* yielding a further *wff* when prepended with any of the temporal modalities discussed in section 3.2; i.e. F , G , P , H . In thus way, both $F\alpha$ and $FS_x A$ are well formed, for example. Any such temporal expression is a state formula.

² We are not aiming to build a predicate logic here, but the attribution of action to agent is clearly a vital component of the architecture. For this reason we introduce a small extension to the logic that allows specified agency. From a syntactic point of view, we require axioms for generalisation and specialisation, *viz.* if α then $\exists x \in X$ s.t. α^x and if α^x then α . The semantics of specified agency are defined in section 4, where the very limited extent of this agentive predication is much clearer.

RES	$\frac{A \leftrightarrow B}{S_x A \leftrightarrow S_x B}$
TS	$S_x A \rightarrow A$
CS	$(S_x A \wedge S_x B) \rightarrow S_x(A \wedge B)$
MS	$S_x(A \wedge B) \rightarrow (S_x A \wedge S_x B)$
RS	$S_x(A \wedge B) \leftrightarrow S_x A \wedge S_x B$
RMS	$\frac{A \rightarrow B}{S_x A \rightarrow S_x B}$
KS	$S_x(A \rightarrow B) \rightarrow (S_x A \rightarrow S_x B)$
DS	$S_x A \rightarrow \neg S_x \neg A$

Figure 1. Rules of inference and axiom schemas of modality S.

R¬NS	$\frac{A}{\neg S_x A}$
RNS	$\frac{A}{S_x A}$
5S	$\neg S_x A \rightarrow S_x \neg S_x A$
4S	$S_x A \rightarrow S_x S_x A$

Figure 2. Some rules of inference and axiom schemas that are *not* valid for S.

3.1. SEINSOLLEN AND TUNSOLLEN

We turn now to an axiomatisation of the logic of S and T. Figures 1 and 2 summarise key rules of inference and axiom schema that will be referred to in this and subsequent sections of the paper.

One of the most fundamental disagreements between theories of agency concerns the rule of necessitation (RN as it would appear for modality S is given in Figure 2). This arises from a deep intuitive dilemma. The argument for adopting the reverse, R¬N, is summarised succinctly by Jones and Sergot: “Whatever else we may have in mind . . . on no account could we accept that an agent brings about what is logically true” (Jones and Sergot, 1996, p. 435). Thus it could be argued that Jones and Sergot, like Belnap and Perloff (whose *negative condition* entails R¬N) do capture an element of the notion of responsibility, in the sense that no agent can be said to be ‘responsible’ for a tautology. Chellas’s intuitions, by contrast, run rather differently. He

is happy to accept RN, a much more conventional rule of a normal modal logic, and his argument too is tabled very briefly: “Can it ever be the case that someone sees to it that something logically true is so? I believe the answer is yes. When one sees to something, one sees to anything that logically follows, including the easiest such things, such as those represented by \top . One should think of seeing to it that, for example, $0 = 0$ as a sort of trivial pursuit, attendant upon seeing to anything at all.” (Chellas, 1992, p. 508). Chellas’s decision, in particular, is motivated by the logical consequences of the rule, and on the availability of schemata C and M.

The *outward* distributivity of an action modality is adopted in the axiom schema C. If an agent brings about two *wffs*, then that agent has thereby brought about the conjunction comprising those two *wffs*. Schema C is adopted by Chellas, Jones and Sergot, Belnap and Perloff, and, similarly, in the work presented here (see Figure 1 for CS); it is difficult to argue on an intuitive basis how C might fail.

The *inward* distributivity axiom schema, M, however, is more troublesome. MS, like CS, seems intuitively appealing, but, for Jones and Sergot (and other systems adopting R \neg N), it is pathological, since, with RES, it yields the rule RMS³ (see Figure 1). Taking the tautology $A \rightarrow \top$, RM gives $S_x A \rightarrow S_x \top$. Since R \neg N gives $\neg S_x \top$ directly, any $S_x A$ is thus a contradiction. Jones and Sergot, therefore, reject M because they are committed to the notion of responsibility captured by R \neg N; Chellas on the other hand, accepts RN and, thereby, the loss of the Jones and Sergot conception of agentive responsibility, but is able, as a result, to maintain M.

The solution proposed for the modalities S and T represents a half-way house, eschewing both the restrictive nature of a (smallest) classical modal logic, and the counterintuitive results of a normal modal logic, in favour of a (smallest) regular modal logic. Both modalities thus include the rule RE and the axiom schema R (and, consequently, M, C and K), but they require neither the rule of necessitation (RN), nor the rule of anti-necessitation (R \neg N). In this way, the logic of S and T responds to

³ The proof that RMS may be simply derived from the rule of inference RES and axiom schema MS is given below where PL refers to propositional logic, and other axioms and rules are as in Figures 1 and 2. The following presentation is in the style of Chellas (1980, p. 236):

- | | | |
|--|------------|---|
| 1. $A \rightarrow B$ | hypothesis | |
| 2. $A \leftrightarrow (A \wedge B)$ | 1, PL | |
| 3. $S_x A \leftrightarrow S_x (A \wedge B)$ | 2, RES | |
| 4. $S_x (A \wedge B) \rightarrow (S_x A \wedge S_x B)$ | MS | |
| 5. $S_x A \rightarrow S_x B$ | 3, 4, PL | □ |

Chellas’s call for construction and investigation of a regular logic with neither RN nor R¬N (Chellas, 1992, p. 515).

We have touched upon the intuitive appeal of M and C. It is also worth digressing to offer an intuitive gloss on the schema K to demonstrate its role. An imperative with the form of an implication is, linguistically, quite straightforward: “Make sure that if you go out then you lock the door”. If an agent brings it about that the implication holds then K states that if the agent in addition brings about the antecedent then it is, as a result, logically responsible for bringing about the consequent. This does not impinge upon the autonomy of an agent to decide not to fulfil some imperative; rather, it states only that if the agent brought about the antecedent, then it can only also be said to have brought about the implication if it is responsible for the consequent. For a whole host of good axiomatic and semantic reasons, K is adopted in all modal systems of seeing-to-it-that, so it is encouraging that it accords with intuitions in this way.

The axioms 4 and 5 are commonly employed in mentalistic modalities, and, less frequently, in agentic modalities. First, consider schema 5S (Figure 2). This is explicitly rejected for several reasons, not least of which is that with TS, it would yield RNS, which we wish to avoid.⁴ Schema 5 is similarly rejected across the board by Jones and Sergot, Belnap and Perloff, and Chellas.

Axiom schema 4, however, is accepted by Belnap and Perloff. Consider schema 4S (Figure 2). With TS, this yields the following equivalence, which we reject: $S_x A \leftrightarrow S_x S_x A$. Chellas too rejects this equivalence, and the axiom 4 from which it is derived (p492), though for him it is because his normal model operator, Δ , is based on an accessibility relation that is not transitive. Here, we wish to keep as simple as possible the relationships between nested modalities so as to make characterisation of delegation — which is founded upon such nesting —

⁴ The proof that RNS (or, equivalently, $A \rightarrow S_x \top$) may be derived from axioms TS and 5S is simple, but included here for completeness:

1.	A	hypothesis	
2.	$S_x \neg A \rightarrow \neg A$	TS	
3.	$A \rightarrow \neg S_x \neg A$	2, PL	
4.	$\neg S_x \neg A$	1, 3, PL	
5.	$\neg S_x \neg A \leftrightarrow \top$	4, PL	
6.	$S_x \neg S_x \neg A \leftrightarrow S_x \top$	5, RES	
7.	$\neg S_x \neg A \rightarrow S_x \neg S_x \neg A$	5S	
8.	$S_x \neg S_x \neg A$	4, 7, PL	
9.	$S_x \top$	6, 8, PL	□

as straightforward as possible. So, we agree with Chellas that “English renderings [of 4] strike the ear at least ambivalently” — but note that as Migita and Hosoi (1997) have shown, the logic that would result if 4 was accepted would still be regular, and would preserve the other properties we have discussed.

Finally, the adoption of T in the models of Jones and Sergot, of Belnap and Perloff, and of Chellas entails the inclusion of axiom schema D, and similarly for the current work (see Figure 1 for DS).

The axiomatisation of the T modality follows exactly the same pattern, and enjoys the same rules of inference and axiom schema of Figure 1. To summarise then, the logics of S_x and T_x are relativised classical regular modal logics of type RT (Chellas, 1980, p. 237). To set this briefly in the broader context of other axiomatic theories, in the Lewis system, RT logics correspond to E2 (Bowen, 1979, pp. 2–3), a Lemmon system that is “quite a powerful system”, and was intended by Lewis as a “possible epistemic counterpart” to S2 (Lemmon, 1957, pp. 181–182). Hughes and Cresswell (1966, p. 208) describe, very briefly, the axiomatisation and ‘semantical point of view’ for E2, but otherwise it occurs rarely in the literature. By comparison, the effect of removing necessitation from the stronger S1-S5 systems has been studied by Feys (1965).

3.2. A LOGIC OF TENSE

To talk about action, and particularly, about the responsibility for future actions, it is vital to integrate a model of time, as Hamblin’s argument makes abundantly clear. As we shall see in the next section, adopting Hamblin’s semantic model makes it straightforward to follow a conventional approach to the logical machinery required for handling time. Though one might adopt a more or less complex approach to time, depending on context and goals, our driving aim here is to be able to express Hamblin’s notions using a familiar logic of time that is, in the first instance, as simple as possible.

For this, a relatively simple Priorian FGPH tense logic will suffice. As laid out in the next section, the underlying structure of our model, however, is not a single, linear time-line and so we may not employ the standard minimal tense logic \mathbf{K}_t (Thomason, 1984). Instead, we need an interpretation of the behaviour of $F\varphi$ that can be squared with the rich, branching time of Hamblin’s semantics. Prior’s (1967, p. 133) account of a Peircean model fits the bill:

$$\begin{array}{ll} Gp \rightarrow q \rightarrow (Gp \rightarrow Gq) & Gp \rightarrow \neg G\neg p \\ Hp \rightarrow q \rightarrow (Hp \rightarrow Hq) & Hp \rightarrow \neg H\neg p \end{array}$$

$$\begin{array}{ll}
p \rightarrow (Hp \rightarrow (Gp \rightarrow GHp)) & Gp \rightarrow GGp \\
& Hp \rightarrow HHp
\end{array}$$

The interpretations for G and H are simply, always true in all futures, and always true in all pasts. Though, as Prior discusses, there are limitations of the Peircean approach, it is sufficient to be able to express future impossibility of a proposition on all possible future histories (i.e. $G\neg p$), which is the key to interpreting Hamblin's model in section 5. We are not advocating such a tense logic as a panacea, nor even as a perfect fit for Hamblin's model. Action-state semantics provides a framework that is both rich enough to support either Ockhamist or Peircean approaches, and also compatible with the transformations described by Goranko and Zanardo (2004). The goal in this section is to demonstrate how one of Goranko and Zanardo's simple Kripkean-based Peircean logics fits onto the underlying action-state semantics. From there, their translations could be applied to yield alternative structures for the temporal modality.

Tensed sentences are states of affairs. To say that something will be true, or that some action has been done, etc. is, as a sentence, a state of affairs (that is, the sentence is certainly not in itself an event). It is, however, entirely reasonable to permit tense operators to range over both states and events: states of affairs may have held in the past, and events may happen in the future, etc. Together, these two requirements establish that sentences such as $P\alpha$ (at some point in the past, action α was executed), $FS_x A$ (it will be the case in some future that x brings about the state of affairs A), $HT_x \alpha$, (it has always been the case that x has seen to it that action α has been executed) $S_x GA$ (x sees to it that it will always be the case that A holds) should all be well formed formulae expressing states of affairs.

It is thus through the tense logic component that events and states are linked at the syntactic level: states precede events which precede subsequent states, and so on. As the S modality ranges only over states of affairs, it handles propositional states (A , B , etc.), S-modal states ($S_x A$, $S_y B$, etc.) and tense-modal states (FA , $GS_x A$, $HT_x \alpha$, etc.).

3.3. DELEGATION AND RESPONSIBILITY

In common with the model of agentive action proposed by Chellas (1992), but contrary to Von Wright's (1968) characterisation and, to some extent, that of Belnap *et al.* (2001), the theory presented here offers scope for nesting the two modalities in building a rich notion of responsibility. Thus, in this section we discuss briefly the theorems and axioms of delegation summarised in Figure 3.

TSS	$S_x S_y A \rightarrow S_y A$	QS	$S_x S_y A \rightarrow S_x A$
TST	$S_x T_y \alpha \rightarrow T_y \alpha$	QT	$S_x T_y \alpha \rightarrow T_x \alpha$
TFS	$FS_x A \rightarrow FA$	QFS	$S_x FS_y A \rightarrow S_x FA$
TPS	$PS_x A \rightarrow PA$	QPS	$S_x PS_y A \rightarrow S_x PA$
TFT	$FT_x \alpha \rightarrow F\alpha$	QFT	$S_x FT_y \alpha \rightarrow S_x F\alpha$
TPT	$PT_x \alpha \rightarrow P\alpha$	QPT	$S_x PT_y \alpha \rightarrow S_x P\alpha$

Figure 3. Theorems and Axioms of Delegation.

Schema QT is worthy of particular note: if agent x sees to it that agent y sees to it that action α is done, then x sees to it that α is done. The adoption of this schema is intuitively appealing: agent x , through seeing to it that y is responsible for α is itself, through delegation of the act, responsible for the performance of α . Thus, following Chellas *inter alia*, we accept the axiom schemas QS and QT (and so capture the legal principle *Qui facit per alium facit per se*). It is interesting to note, however, that these axioms cannot be accepted by Belnap *et al.*; Chellas (1992, p. 506) shows that the *something happens* condition (for all possible combinations of choices of all agents there is at least one history) leads to (using our notation) $S_x S_y A$ being false whenever $x \neq y$.

We further accept the specialisations of the TS schema, TSS and TST (Figure 3). The theorem TSS expresses the idea that for two different agents, x and y , if x *successfully* sees to it that y sees to it that A holds by issuing a command in an appropriate social context, for example, then agent y , to whom the goal (i.e. the achievement of A) is delegated, itself sees to it that A holds.

These axiom schemata and theorems — the variants of Q and T respectively — lay the foundation for characterising acts of delegation. Up to this point, however, we have considered only atemporal theorems and axioms of delegation. Through our combination of a logic of action and a simple tense logic, we are able to capture a more refined notion of delegation and agentive responsibility. Given the rules for legal formulations of *wffs*, there are in total eight basic combinations of the base tense and action modalities:

$$\begin{array}{cccc}
 S_x FA & S_x F\alpha & S_x PA & S_x P\alpha \\
 \\
 FS_x A & FT_x \alpha & PS_x A & PT_x \alpha
 \end{array}$$

These eight fall in to two categories, those that are action modalities ranging over tensed propositions (the first line presented above), and those that are tensed statements referring to action modalities (the

second line). The first category is unremarkable except to note that S_xFA entails FA by axiom T. It is also worth noting that from S_xFA neither S_xA nor FS_xA are derivable, nor — according to our intuitions — should they be.

The second group is a little more interesting. In each of the four cases, we need an analog of axiom T to capture successful action (as outlined in an atemporal fashion above). The successful action now, however, needs carrying over our tense modalities. Thus from FS_xA , we want to conclude FA , and so on; these axioms are listed in Figure 3 as TFS, TPS, TFT and TPT. There are no further intuitive relationships between these theorems beyond those already captured by the existing axiomatic machinery.

In the same vein, there are in total 12 well-formed combinations that involve three alternating modalities:

$$\begin{array}{cccc}
 S_xFS_yA & S_xFT_y\alpha & S_xPS_yA & S_xPT_y\alpha \\
 FS_xFA & FS_xF\alpha & FS_xPA & FS_xP\alpha \\
 PS_xFA & PS_xF\alpha & PS_xPA & PS_xP\alpha
 \end{array}$$

Again, for many of these, all the deductions one would wish to be able to draw are already catered for using the axioms from section 3.1, or using those of successful action over time, presented above in section 3.2. What is missing, however, is characterisation of delegation of responsibility over time. Thus S_xFS_yA yields not only FS_yA via axiom T, and thence FA via axiom TFS, but in addition, we want to capture the fact that S_xFS_yA also has a more intimate connection (i.e. x 's responsibility for) the future occurrence of A . Specifically, by analogy to the atemporal Q axioms, we would want to be able to derive S_xFA . In this way, we construct four new analogs of the Q axiom (QFS, QPS, QFT and QPT) that carry over tense modalities, in just the same way as we have done for analogs of the T axiom above. To illustrate the role of these axioms, let us consider just one of them and summarise its meaning in the same terms as we did with axiom QT above. Axiom QFT, for instance, captures the idea that if agent x sees to it that, at some point in a possible future, agent y sees to it that action α is done, then x sees to it that, at some point in a possible future, α is done.

The meanings of axioms QPS and QPT are, possibly, less intuitively clear. Putting the same gloss on QPT that we did for QFT above, for example, would be to say that if agent x sees to it that, at some point in a possible past, agent y sees to it that action α is done, then x sees to it that, at some point in a possible past, α is done. It may be a little difficult to accept this axiom in terms of our intuitions about time (and causality), but it should be noted that what we are capturing is the

notion of agentive responsibility in a temporal context. Thus, QPT may be better understood using the following characterisation: if agent x is responsible for a state of affairs in which, at some point in a possible past, agent y sees to it that α is done, then x is responsible for, at some point in a possible past, the doing of α .

With this axiomatisation in place, we turn now to the construction of the semantics that ground our model of delegation and responsibility.

4. Semantics of the model for **S** and **T**

The characterisation of the formal model is non-trivial because Hamblin's rich model requires careful interpretation in order to be viewed as a Kripke structure upon which modal operator semantics might be defined. The difficulty lies in the profligacy of action state semantics, in that both states and events are represented explicitly. It is, therefore, not sufficient to see either state descriptions alone or event descriptions alone as the contents of possible worlds. Rather, a possible worlds structure must involve both components. At the same time, the formal model should be sufficiently expressive to reflect and preserve this richness at the syntactic level.

The approach we take is to stratify the semantic model into two layers. By doing this, it becomes possible to simplify one of the challenges facing characterisation of a logic of action and time, namely to bind together the temporal component with the action component. For the former, a given application domain will guide selection of an appropriate logic. For many (and for the presentation here), a simple Priorian tense logic will suffice; for other domains, more sophisticated temporal logics may be more suitable. For the action component, as we have already argued from an axiomatic point of view, the best choice is a regular logic. The challenge, then, is of uniting semantics for these two fundamentally different mechanisms.

The first step is to preserve the distinction between states and events simply by dividing them into separate sets of possible worlds. That is, some worlds contain state descriptions, and other worlds contain event descriptions. Then we define an accessibility relation that holds between these worlds. Rather than adopting a conventional binary relation, Hamblin's semantics demands a ternary relation that links a world of state descriptions, with a world of event descriptions and another world of state descriptions.⁵ Though ternary accessibility relations are unusual they are discussed in the context of general modal similarity

⁵ It could equally well be linking two event description worlds with one state description worlds; ultimately nothing of importance hangs upon this decision.

types by Blackburn et al. (2001), and in the context of a Smith-style logic by Macleod and Schotch (2000).

In our formal characterisation of the imperative, we build upon a Hamblinian frame of reference, $\mathcal{F}_{\mathcal{H}}$.

DEFINITION 1 ($\mathcal{F}_{\mathcal{H}}$). $\mathcal{F}_{\mathcal{H}} = \langle \mathcal{W}, \mathcal{R}_{\mathcal{H}} \rangle$ such that

- (i) \mathcal{W} is a non-empty set that collects together our ‘state worlds’ and ‘event worlds’; and
- (ii) $\mathcal{R}_{\mathcal{H}}$ is a ternary relation, where $\langle x, y, z \rangle \in \mathcal{R}_{\mathcal{H}}$ should be read as ‘state z is accessible from state x by way of y ’.

To give an intuitive grasp of the ternary relation $\mathcal{R}_{\mathcal{H}}$ consider the act of requesting the bill at a restaurant. The state of not having the bill and wishing to pay for the meal (x) may, through the act of requesting the bill, be transformed into a state in which the waiter addressed is to provide the bill (z). This request may be issued in a number of ways including saying to the waiter “May I have the bill please?” or by indicating that this is your wish by motioning a signature. These are just two possible ways in which state x may be transformed into state z , and hence would form two elements of the ternary relation $\mathcal{R}_{\mathcal{H}}$ for this restaurant domain.

Note that it is not necessary to enforce types upon the worlds explicitly in the semantics — this can be handled implicitly by $\mathcal{R}_{\mathcal{H}}$. The set of all the worlds that take the first and third places in tuples that are instances of the relation are completely disjoint for the set of worlds that take place two. That is, we can define the sets $\mathcal{W}_{\text{states}}$ and $\mathcal{W}_{\text{events}}$ of ‘state worlds’ and ‘event worlds’ respectively as follows:

DEFINITION 2 ($\mathcal{W}_{\text{states}}$).

$$\mathcal{W}_{\text{states}} = \{w \mid \langle w, x, y \rangle \in \mathcal{R}_{\mathcal{H}}\} \cup \{w \mid \langle x, y, w \rangle \in \mathcal{R}_{\mathcal{H}}\}$$

DEFINITION 3 ($\mathcal{W}_{\text{events}}$).

$$\mathcal{W}_{\text{events}} = \{w \mid \langle x, w, y \rangle \in \mathcal{R}_{\mathcal{H}}\}$$

There are two conditions on $\mathcal{R}_{\mathcal{H}}$ required to accurately capture the Hamblinian picture and the intuitions associated with it. The first is that each world in \mathcal{W} is either a world containing state descriptions (in this case it is a member of the set $\mathcal{W}_{\text{states}}$ as defined above) or a world containing event descriptions (in which it is a member of the set $\mathcal{W}_{\text{events}}$ as defined above), but not both; i.e. $\mathcal{W}_{\text{states}} \cap \mathcal{W}_{\text{events}} = \emptyset$. The second constraint on our relation $\mathcal{R}_{\mathcal{H}}$ is that the structure of interconnected

‘state worlds’ (via ‘event worlds’) is a directed acyclic graph.⁶ Beyond these two conditions, the relation $\mathcal{R}_{\mathcal{H}}$ does not portray any of the more typical characteristics of many binary accessibility relations used by classical modal logics: it is neither symmetric nor reflexive, transitive nor Euclidean.

The semantic structure defined by $\mathcal{R}_{\mathcal{H}}$ forms the lower of the two layers in our stratified model. On the basis of this layer, it is then possible to define accessibility relations (and later, necessitation functions) which are more familiar. First, for the temporal component we define an accessibility relation R_F , expressing the earlier-later relation.

DEFINITION 4 (R_F).

$$\begin{aligned} \langle u, v \rangle \in R_F \text{ iff } & \exists w \in \mathcal{W} \text{ s.t. } \langle u, v, w \rangle \in \mathcal{R}_{\mathcal{H}} \text{ or} & \text{(i)} \\ & \exists w \in \mathcal{W} \text{ s.t. } \langle u, w, v \rangle \in \mathcal{R}_{\mathcal{H}} \text{ or} & \text{(ii)} \\ & \exists w, x \in \mathcal{W} \text{ s.t. } (\langle u, w, x \rangle \in \mathcal{R}_{\mathcal{H}} \text{ and } \langle x, v \rangle \in R_F) & \text{(iii)} \end{aligned}$$

There are several interesting points to note concerning the relationship between $\mathcal{R}_{\mathcal{H}}$ and R_F . Most importantly, where $\mathcal{R}_{\mathcal{H}}$ distinguishes between state worlds and event worlds, R_F does not. So, in clause (i) of Definition 4, state world u is explicitly related temporally to an event world v that immediately follows it. Clauses (ii) and (iii) of Definition 4 are more straightforward in that (ii) is relating consecutive state worlds, and (iii) is then building transitivity into the relation through a recursive definition (a world is in the future from u if it is in the future for the immediately next state world, x).

Normally, we expect to see accessibility relations capturing both earlier-later, and its dual, later-earlier. Equally conventional is the definition of one simply in terms of a syntactic rearrangement of the other, i.e. $\text{before}(x, y) \equiv \text{after}(y, x)$. The Hamblinian foundation, however, requires a slightly more complex characterisation, since we must avoid the derivation of temporally modal statements in event worlds. Sentences expressing temporal facts are not, in themselves, events. Intuitively, we want to capture the picture illustrated in Figure 4, in which, from the point of view of a state world, there are both state worlds and event worlds that are in the future, and, similarly, that there are both state worlds and event worlds in the past. From event worlds, however, there are no other accessible possible worlds at all (according to $\mathcal{R}_{\mathcal{F}}$ or $\mathcal{R}_{\mathcal{P}}$).

⁶ We may express this constraint formally as follows: $\forall w \in \mathcal{W}_{\text{states}}, \text{parents}(w) \cap \text{descendants}(w) = \emptyset$ such that $\text{parents}(w) = \{u \in \mathcal{W}_{\text{states}} \text{ such that } \langle u, v, w \rangle \in \mathcal{R}_{\mathcal{H}}\}$ and $\text{descendants}(w) = \{u \in \mathcal{W}_{\text{states}} \text{ such that there is a path from } w \text{ to } u\}$, and w_1, \dots, w_n is a path if w_1, w_2, \dots, w_n are distinct members of the set $\mathcal{W}_{\text{states}}$ and v_1, v_2, \dots, v_{n-1} are distinct members of the set $\mathcal{W}_{\text{events}}$ and $\langle w_i, v_i, w_{i+1} \rangle \in \mathcal{R}_{\mathcal{H}}$ for $1 \leq i \leq n$. These paths correspond closely to Hamblinian histories (but we postpone a precise definition until section 5).




Figure 4. Accessibility relations R_F and R_P illustrated over $\langle u, v, w \rangle \in \mathcal{R}_{\mathcal{H}}$.

Formally, then, we define R_P separately in Definition 5, which functions as the direct analog of R_F . R_F and R_P thus together form a conventional temporal frame for a traditional, transitive Priorian tense logic (Prior, 1967).

DEFINITION 5 (R_P).

$$\begin{aligned} \langle u, v \rangle \in R_P \text{ iff } & \exists w \in \mathcal{W} \text{ s.t. } \langle w, v, u \rangle \in \mathcal{R}_{\mathcal{H}} \text{ or} & \text{(i)} \\ & \exists w \in \mathcal{W} \text{ s.t. } \langle v, w, u \rangle \in \mathcal{R}_{\mathcal{H}} \text{ or} & \text{(ii)} \\ & \exists w, x \in \mathcal{W} \text{ s.t. } (\langle x, w, u \rangle \in \mathcal{R}_{\mathcal{H}} \text{ and } \langle x, v \rangle \in R_P) & \text{(iii)} \end{aligned}$$

The action component is then tackled in a slightly different way. Given that the logic of S and T is non-normal, it demands a minimal model, defined upon necessitation functions. Those necessitation functions must act upon a different substrate: for the S modality, the substrate is state descriptions, for T, event descriptions. The necessitation functions are relativised to individual agents in the usual way (that is, the way in which one agent’s behaviour is described is independent of how other agents’ behaviour is described). Thus \mathcal{S}^x is the necessitation function for the modality S, relativised to some agent, x . The functions map from worlds to sets of worlds. So, $\mathcal{S}^x : \mathcal{W}_{\text{states}} \rightarrow \wp(\wp(\mathcal{W}_{\text{states}}))$, as usual (thereby picking out worlds by which necessity is defined).

The T modality is a little less straightforward, in that $\mathcal{T}^x : \mathcal{W}_{\text{states}} \rightarrow \wp(\wp(\mathcal{W}_{\text{events}}))$. \mathcal{T}^x is, therefore, picking out particular events that are, loosely, “actionable” by x from a ‘state world’ ω . Furthermore, the $\mathcal{W}_{\text{events}}$ worlds are not simply propositional. To accurately model Hamblin’s conception of “deed-agent assignments”, these worlds are filled (exclusively) with statements of the form *agent x performs action α* , that we represent with the typographic shorthand α^x , and *wffs* constructed from such statements using *PL* (see section 3).

In this way, the model as a whole is defined as $\langle \mathcal{W}, \mathcal{X}, \mathcal{I}, \mathcal{R}_{\mathcal{H}}, \mathcal{S}^x, \mathcal{T}^x \rangle$ for a set of possible worlds \mathcal{W} , a set of agents \mathcal{X} , an interpretation function \mathcal{I} , the ternary Hamblinian accessibility relation, $\mathcal{R}_{\mathcal{H}}$ and the relativised necessitation functions for the modalities S and T, \mathcal{S}^x and \mathcal{T}^x (for each $x \in \mathcal{X}$), respectively.

The necessitation functions, in combination with the accessibility relations then offer a straightforward way of characterising the semantics of the logic as a whole:

$$\begin{aligned} \models_{\omega}^{\mathcal{M}} A &\text{ iff } \mathcal{I}(A, \omega) = \top \\ \models_{\omega}^{\mathcal{M}} \alpha^x &\text{ iff } \mathcal{I}(\alpha^x, \omega) = \top \\ \models_{\omega}^{\mathcal{M}} \alpha &\text{ iff } \exists x \in \mathcal{X} \text{ s.t. } \mathcal{I}(\alpha^x, \omega) = \top \\ \models_{\omega}^{\mathcal{M}} \mathcal{S}_x A &\text{ iff } \|A\|^{\mathcal{M}} \in \mathcal{S}^x(\omega) \\ \models_{\omega}^{\mathcal{M}} \top_x \alpha &\text{ iff } \|\alpha\|^{\mathcal{M}} \in \mathcal{T}^x(\omega) \end{aligned}$$

The truth set is constructed normally:

$$\|\varphi\| = \{\omega \mid \models_{\omega}^{\mathcal{M}} \varphi\}$$

(Notice that the truth set is thus constructed in the same manner for both states and events; this symmetry is a result of the typing of possible worlds, so that increased complexity in the model structure yields increased simplicity in the connection between that structure and the syntactic surface. In an earlier draft of this work (Norman and Reed, 2002), this balance was reversed, but the desire to provide a usable logical system drove the development of this simpler definition of truth sets.)

Then, for the tense component, the characterisation is straightforward:

$$\begin{aligned} \models_{\omega}^{\mathcal{M}} \mathcal{F}\phi &\text{ iff } \omega \in \mathcal{W}_{\text{states}} \wedge \exists \omega' \in \mathcal{W} \text{ s.t. } (R_F(\omega, \omega') \text{ and } \models_{\omega'}^{\mathcal{M}} \phi) \\ \models_{\omega}^{\mathcal{M}} \mathcal{P}\phi &\text{ iff } \omega \in \mathcal{W}_{\text{states}} \wedge \exists \omega' \in \mathcal{W} \text{ s.t. } (R_P(\omega, \omega') \text{ and } \models_{\omega'}^{\mathcal{M}} \phi) \end{aligned}$$

The only complexity is that tense statements are restricted to state worlds — it is counterintuitive to consider tensed state descriptions as events. The duals of these two are slightly more complex, as we must bear in mind that for something to always be true (or to always have been true), it does not need to hold in all worlds, only in all state worlds or all event worlds. So, for example, to say that $G\alpha$ — that event α is always going to happen, corresponds only to α holding in all future event worlds. Thus,

$$\begin{aligned} \models_{\omega}^{\mathcal{M}} \mathcal{G}\phi &\text{ iff } \omega \in \mathcal{W}_{\text{states}} \wedge \left((\forall \omega' \in \mathcal{W}_{\text{states}} \text{ s.t. } R_F(\omega, \omega'), \models_{\omega'}^{\mathcal{M}} \phi) \vee \right. \\ &\quad \left. (\forall \omega' \in \mathcal{W}_{\text{events}} \text{ s.t. } R_F(\omega, \omega'), \models_{\omega'}^{\mathcal{M}} \phi) \right) \end{aligned}$$

$$\models_{\omega}^{\mathcal{M}} \mathbf{H}\phi \text{ iff } \omega \in \mathcal{W}_{\text{states}} \wedge \left(\left(\forall \omega' \in \mathcal{W}_{\text{states}} \text{ s.t. } R_P(\omega, \omega'), \models_{\omega'}^{\mathcal{M}} \phi \right) \vee \left(\forall \omega' \in \mathcal{W}_{\text{events}} \text{ s.t. } R_P(\omega, \omega'), \models_{\omega'}^{\mathcal{M}} \phi \right) \right)$$

It is then a relatively straightforward matter to combine these tense modalities with action modalities. So, *wffs* constructed from action modalities can be tensed, for example:

$$\models_{\omega}^{\mathcal{M}} \mathbf{PT}_x \alpha \text{ iff } \omega \in \mathcal{W}_{\text{states}} \wedge \exists \omega' \in \mathcal{W} \text{ s.t. } \left(R_P(\omega, \omega') \wedge \left(\|\alpha\|^{\mathcal{M}} \in \mathcal{T}^x(\omega') \right) \right)$$

and similarly, tense statements can form the contents of action statements, for example:

$$\models_{\omega}^{\mathcal{M}} \mathbf{S}_x \mathbf{F}\alpha \text{ iff } \left\{ \omega' \mid \omega' \in \mathcal{W}_{\text{states}} \wedge \exists \omega'' \in \mathcal{W} \text{ s.t. } R_F(\omega', \omega'') \wedge \models_{\omega''}^{\mathcal{M}} \alpha \right\} \in \mathcal{S}^x(\omega)$$

5. Characterising Hamblin's Model

With the semantics of our model in place, it then becomes possible to construct an interpretation of the various key notions in Hamblin's presentation. In doing so, we represent the summary of Hamblin's model given in section 2, grounding the notions of a world, a history, strategies and, most importantly, partial *i*-strategies for the wholehearted satisfaction of imperative *i*, in the formal model developed in section 4.

The first of these notions are Hamblin worlds and Hamblin histories (henceforth H-worlds and H-histories, respectively). Some H-world, w , is a string of states tied together by particular events, all in a specific order identified by a contiguous sequence of natural numbers⁷ and W denotes the set of all possible H-worlds (see Definition 6).

DEFINITION 6 (W). *The set of valid H-worlds, W , is defined such that for each $w \in W$, every step in that H-world must be licensed by the accessibility relation $\mathcal{R}_{\mathcal{H}}$, and, furthermore, if there is a next step, it must start from the state reached at the end of this step.*

$$W = \{w \mid w \in \Omega \text{ s.t. } (\forall \langle n, s, d \rangle \in w, (\exists s' \in \mathcal{W}_{\text{states}} \text{ s.t. } \mathcal{R}_{\mathcal{H}}(s, d, s') \text{ and } (\text{if } \langle n+1, s'', d' \rangle \in w \text{ then } s'' = s'))))\}$$

where $\Omega = \wp(\mathbb{N} \times \mathcal{W}_{\text{states}} \times \mathcal{W}_{\text{events}})$.

⁷ Formally: $\forall \langle n, s, d \rangle, \langle o, s', d' \rangle \in w, (\forall m \in \mathbb{N} \text{ s.t. } n < m < o, \exists \langle m, s'', d'' \rangle \in w) \wedge (\text{if } n = o \text{ then } s = s' \text{ and } d = d')$

At this point, it is worth making explicit the relationship between the states, happenings and deed-agents assignments of a given H-world, w , denoted $S(w)$, $H(w)$ and $D(w)$ respectively in section 2, and the sets $\mathcal{W}_{\text{states}}$ (the ‘state worlds’ in \mathcal{W} of our semantic frame $\mathcal{F}_{\mathcal{H}}$) and $\mathcal{W}_{\text{events}}$ (‘event worlds’) as specified in section 4. $\mathcal{W}_{\text{states}}$ is the set of all states $S(w)$ for each world w in the set of all Hamblinian worlds W .

$$\forall w \in W, S(w) \subseteq \mathcal{W}_{\text{states}}$$

$\mathcal{W}_{\text{events}}$ is the set of all deed-agent assignments, $D(w)$, and happenings, $H(w)$, for each world w in the set of all Hamblinian worlds W .

$$\forall w \in W, D(w) \cup H(w) \subseteq \mathcal{W}_{\text{events}}$$

We therefore conflate deed-agent assignments and happenings into the ‘event worlds’ — each happening can be viewed as a deed done by a special ‘world’ agent.

Above, we defined an H-world as an ordered series of state-event steps using the components of our semantic frame $\mathcal{F}_{\mathcal{H}}$. Moving on now to histories, we may define an H-history, j_t , of a world w simply as a restriction of w up to a given step.

DEFINITION 7 (j_t).

$$\forall \langle n, s, d \rangle \in w, \text{ if } n \leq t \text{ then } \langle n, s, d \rangle \in j_t$$

The set of all H-histories up until some time point, J_t , is then simply defined by placing the same restriction on all $w \in W$.

Of course, if we read t as time, this implicitly associates the natural number identities of the steps with time points. This may or may not be desirable (there is no restriction so far in place to ensure a global system of such identifiers). Given that the underlying temporal semantics is Priorian, introducing fixed time points to which reference can be made requires additional machinery to limit the Priorian model. One way would be to enforce numbered time steps in the semantics; another would be to represent a special value in each world to indicate some clock. There may also be alternative approaches that are more suited to particular applications. Here we do not suppose to anticipate these design decisions, and simply leave it as an implementation choice. We thus interpret Hamblin’s “history up to a time” as “history up to a given step in the world”.

Now, finally, we may define the H-worlds in which an imperative i is extensionally satisfied: W_i . We consider two forms of the imperative: seeing to it that a state of affairs A holds, $S_x A$, and seeing to it that an action α is executed, $T_x \alpha$. In our definition of W_i , both these forms are captured; see Definition 8.

DEFINITION 8 (W_i). *The worlds in which an imperative $i = W_{S_x A}$ is extensionally satisfied, $W_{S_x A}$, will be the set containing all those H-worlds $w \in W$ such that A holds at some point in that H-world.*

$$\forall w \in W, \text{ if } \exists \langle n, s, d \rangle \in w \text{ s.t. } \models_s^M A \text{ then } w \in W_{S_x A}$$

Similarly, the H-worlds in which the imperative $i = T_x \alpha$ is extensionally satisfied, $W_{T_x \alpha}$, will be the set containing all those H-worlds $w \in W$ such that α is done at some point in that H-world.

$$\forall w \in W, \text{ if } \exists \langle n, s, d \rangle \in w \text{ s.t. } \models_d^M \alpha, \text{ then } w \in W_{T_x \alpha}$$

We are now in a position to capture the notion of a strategy. In section 2 we noted that a strategy at some time t , denoted q_t , is an allocation of a deed to every history up to every time t' such that $t' \geq t$. In other words, a strategy specifies what should be done now and in the future in all possible H-worlds (see Definition 9).

DEFINITION 9 (q_t). *For all histories terminating at this and all future time points, there is some deed, α , in the set of all possible deeds, D , that is associated with that history within the strategy.*

$$\forall t' \geq t, \forall j_{t'} \in J_{t'}, (\exists \alpha \in D \text{ s.t. } \langle j_{t'}, \alpha \rangle \in q_t)$$

This provides us with a definition of one possible strategy, but we also want to define the set of all possible strategies at some time point t , denoted Q_t . This set enumerates all those possible action assignments to histories. To do this, first we must define the set of all possible histories up to all $t' \geq t$, denoted J_{t+} .

DEFINITION 10 (J_{t+}). *Given that j_t is the history of an H-world up to t and J_t is the set of all histories (of all H-worlds) up to t , the set of all histories (of all H-worlds) up to all t' such that $t' \geq t$ is defined as follows:*

$$J_{t+} = \bigcup_{t' \geq t} J_{t'}$$

We wish to assign deeds to histories to construct strategies. The cross product of J_{t+} (all histories up to all time points at or after t) and D (all possible deeds) gives us all the possible combinations of these histories and deeds. Now, we define Σ to be the set of all subsets of these history-deed assignments.

DEFINITION 11 (Σ).

$$\Sigma = \wp(J_{t+} \times D)$$

Q_t , the set of all possible strategies at time t will be the *partition* of Σ such that every history in J_{t+} appears with some deed assignment as an element in each member of the partition. The criteria for an element of Σ to be an element of the partition of Σ that is Q_t are: (i) the number of elements in Q_t must be maximal with respect to J_{t+} , i.e. must be the same as the number of elements in J_{t+} , and (ii) all histories in J_{t+} must have an assignment.

DEFINITION 12 (Q_t). *The set of all possible strategies at time t is defined as follows:*

$$\forall \sigma \in \Sigma \quad \begin{array}{l} \sigma \in Q_t \quad \text{if } |\sigma| = |J_{t+}| \text{ and } \{j \mid \langle j, \alpha \rangle \in \sigma\} = J_{t+} \\ \sigma \notin Q_t \quad \text{otherwise} \end{array}$$

This set Q_t contains all possible *generalised* strategies; i.e. strategies that are independent of a specific agent. We wish to get to a position in which we can talk about *specific* agents adopting *partially specified* strategies for the *wholehearted satisfaction* of an imperative. In getting to this goal, we must first be able to isolate those H-worlds in which a specific agent complies with a specific strategy. Then we can go on to discuss partial strategies and conditions for wholehearted satisfaction of an imperative.

So, let $W_{\text{strat}}(x, q_t)$ be exactly that set of H-worlds in which agent x complies with strategy q_t ; i.e. those H-worlds that are prescribed by the strategy q_t when adopted by agent x . To define this set, we need to instantiate a general strategy q_t to a specific agent and use this to obtain a subset of the set of H-worlds, W .

DEFINITION 13 ($W_{\text{strat}}(x, q_t)$).

$$W_{\text{strat}}(x, q_t) = \left\{ w \mid w \in W, \forall \langle j_{t'}, \alpha \rangle \in q_t \text{ s.t. } j_{t'} \subseteq w \text{ and } \models_d^M \alpha^x \right. \\ \left. \text{where } \langle t', s, d \rangle \in j_{t'} \right\}$$

As argued by Hamblin (1987, p. 157), “in practice, no one ever chooses or is allocated a strategy in the minute detail that specifies every deed; and certainly not to the end of time”. We wish to leave an agent’s strategy *underspecified*, and so we need to capture Hamblin’s notion of a *partial* strategy for an agent. A partial strategy is, simply, a disjunction of strategies; i.e. a subset of the set of all strategies Q_t . So, if $Q'_t \subseteq Q_t$ then $W_{\text{strat}}(x, Q'_t)$ denotes the set of H-worlds such that each H-world in $W_{\text{strat}}(x, Q'_t)$ is prescribed by at least one of the strategies in Q'_t .

DEFINITION 14 ($W_{\text{strat}}(x, Q'_t)$).

$$W_{\text{strat}}(x, Q'_t) = \bigcup_{q_t \in Q'_t} W_{\text{strat}}(x, q_t)$$

For an imperative i , a partial i -strategy for an agent x is a partial strategy, Q'_t , such that all those H-worlds prescribed by Q'_t are also worlds in which i is extensionally satisfied.

DEFINITION 15 (**partial i -strategy**).

Q'_t is a partial i -strategy iff $W_{\text{strat}}(x, Q'_t) \subseteq W_i$

If an agent has maintained a partial strategy for the satisfaction of an imperative i and that imperative is extensionally satisfied, then the agent has *whole-heartedly* satisfied the imperative. This, therefore, provides us with our sought-for interpretation of wholehearted satisfaction founded upon a formal model of action. Take, for example, the following imperative issued to an agent x : $S_x A!$ For A to be extensionally satisfied, of course, all that is required is that the state of the actual world at some point includes A . But for x to wholeheartedly fulfil the imperative, x must both maintain a partial strategy for the satisfaction of $S_x A$, and at every time point, select an action for itself according to that strategy. The need for both intensional and extensional components is thus respected.

The model enables an understanding of responsibility for states and events enacted through imperatives in two steps. When one agent communicates an imperative to another, it requires instantaneous acquiescence (consider, for example, that if an imperative is to be rejected it should most usually be rejected forthwith). Yet if some kind of acquiescence or contract is struck immediately, then it is clearly not (usually) the completion of the demand. A trivial "Sit down!" might be fulfilled (almost) instantaneously, but "Make sure the report's on my desk tomorrow morning!" clearly is not. So to what, then, is the recipient acquiescing at the point of the communicative exchange? This conundrum is resolved easily when the imperative is analysed as a demand for an agent to "see to it that". It is to such a modal expression that the recipient is committing. Analysing imperatives in the form $S_x A$ or $\top_x \alpha$ is thus the first step.

The second is in the recipient's subsequent action. So an agent x can be said to have seen to it that some state of affairs holds, (or that some action is performed), just in the case that x wholeheartedly satisfies the appropriate modal expression.

These two steps together offer a definition of what it means for an agent to be responsible for the fulfilment of an imperative. So in response to an imperative to "Make sure that the report's on my desk" (make sure that $r!$), x is only responsible for its fulfilment if x wholeheartedly satisfies $S_x r$. Pure extensional satisfaction alone (such as someone else completing the report unbeknown to x) is, as we would hope, not enough.

6. Related Work

There are a number of landmark studies in the development of logical theories of action, and in this section we discuss a few of the principal theories and compare them to the model presented in this paper. The notable studies that we discuss are those reported by Belnap *et al.* (Belnap and Perloff, 1988; Belnap and Perloff, 1992; Horty and Belnap, 1995; Horty, 2001; Belnap *et al.*, 2001), Chellas (1992), Jones *et al.* (Jones and Sergot, 1996; Santos *et al.*, 1997), Elgesem (1997), Segerberg (1989), Aumann (1976), d'Altan *et al.* (1996) and Singh (1991; 1993).

Belnap and Perloff's (1988) model of individual agency marks the starting point of a substantial body of research into theories of agentive action. Belnap and Perloff (1992) present a theory of individual agency built upon a branching time structure with moments (states of affairs) related by an "earlier-than" relation and histories as paths through moments. This model, notable for the use of the "negative condition" (it is not the case that an agent sees to something that is necessarily true), is refined by Horty and Belnap (1995), and given a complete axiomatisation by Xu (1995). Both Horty (2001) and Belnap *et al.* (2001) then present further refinements and analyses.

There are a number of distinctions that can be made between this seminal body of research and the model presented in this paper. The first, and most important, is that in the original Belnap and Perloff model, and in all subsequent refinements, there is no means to express that an agent sees to it that some action is done; there is no direct counterpart to the T modality. Horty and Belnap (1995) (see also Horty (2001)) focus on two forms of *seeing to it that*: the achievement *stit* and the deliberative *stit*. The distinction between these two notions lies in their interpretation within a *stit* frame. The achievement *stit* is interpreted with reference to two moments: the moment on some history at which some state of affairs holds and the immediately prior moment at which the agent made a decision on what to do. An agent, α can be said to have seen to it that A at moment m on history h if and only if A holds at that moment on that history and α made an action choice in the moment immediately prior to m such that A was *guaranteed* at m on history h . Rather than saying anything about seeing to it that an action is performed, the truth condition of the achievement *stit* is an attempt to capture the notion of *wholehearted satisfaction* of an imperative to see to it that some state of affairs holds. Thus, this specification of the achievement *stit* serves to tie together the action choices of an agent with the actual outcomes that are produced in the world as a result of those choices in the same way that an agent's partial i -strategy is intended to capture this notion within the model

presented in this paper. The difference, however, is that whereas the satisfaction of the achievement *stit* depends only upon the immediately prior action choice, wholehearted satisfaction of an imperative within the model proposed by Hamblin and formalised here depends on the selection of actions from a partial strategy over an arbitrary number of time steps. The deliberative *stit* has no such means of associating the action choices of an agent with outcomes; it is interpreted only on the basis of the outcomes, and so there is no way to distinguish between the wholehearted and the extensional satisfaction of an imperative.

Chellas (1992) discusses in some detail the “metaphysical backdrop” of Belnap *et al.*’s model. As he explains, that backdrop is different from his own, and, though there are some similarities, Chellas’s own is different again from Hamblin’s as it is interpreted here. So for Chellas, “a history’s past is unique, whereas its future may be manifold” (p. 489) whereas by definition for Hamblin, a history has exactly one past and exactly one future. For Chellas, each history maps from a single time line to a state of affairs, for Hamblin, multiple histories can pass through a single state of affairs at a given time. To an extent, these differences may be terminological, but perhaps a closer translation might be achieved in comparing Hamblin’s histories with more complex structures offered by both Belnap *et al.* and Chellas in response to criticisms of their respective theories. Specifically, they use *chains* to tie together moments in their histories, and then use these chains to provide complex non-atomic witnesses. These chains are very close to the definition of Hamblin histories. For Chellas and Belnap *et al.*, chains are composed from more basic structures and introduce significant additional complexity (with associated challenges); for Hamblin’s theory, these structures are basic, and support a simplicity in the formal account constructed from it. There are a number of other differences between the model presented here and that in Chellas (1992), such as the definition of instigative alternatives (which for Chellas can be defined using an accessibility relation because his model is standard, whereas the equivalents in the model here are the necessitation functions \mathcal{S} and \mathcal{T} , because of our need for a minimal model), and the use of temporal modalities (Chellas needs to identify a number of new temporal modalities, whereas our attempt has been to decouple them from the underlying framework as appropriate). But of course the key difference lies in the adoption, as Chellas suggests (p. 515), of a regular logic with neither RN nor R¬N, upon which to build an account of agency.

Jones and Sergot’s (1996) stance is similar to both Belnap *et al.* and Chellas in that they discuss at some length possible axioms of a logic of agentive action; captured by their E modality. Following this

analysis, they settle on a small classical modal logic of type ECT, which includes R-N, and so, in this respect at least, they agree with Belnap *et al.* Their analysis goes deeper than this, however, because their goal is to capture a notion of institutionalized power. In doing so they introduce a relation \Rightarrow_s , which is read “counts as in society s ”. This idea is important because it is an attempt to capture the societal effects of agents bringing about certain states of affairs, and in so doing Jones and Sergot are addressing the important issue (not addressed in the current work) of the social context within which an imperative is issued. Santos *et al.* (1997) present a logic in the style of Jones and Sergot (1996) that introduces novel distinctions between direct and indirect agency (their G operator represents indirect agency), and between successful and attempted action (their H operator represents “not necessarily successful” action); their E operator captures direct successful action as in Jones and Sergot (1996). Beyond TS and TT, we do not here need to make any distinction between successful and attempted action, because Hamblin’s underlying framework cleanly decouples an action from its effects. Similarly, direct and indirect agency is handled at a deeper level by the decoupling of responsibility for action (captured by T) and the binding of action to agency (such as α^x). Elgesem (1997) builds goal-directedness into the foundation of his logic, and tackles the same issues, which he refers to as *Success*, *Non-accidence* and *Avoidability*. Success corresponds to the same notion of successful action discussed in Jones and Sergot (1996), Santos *et al.* (1997), and here. Non-accidence corresponds very closely to Hamblin’s exposition of the need for a counter-factual, extensional component built in to his definition of wholehearted satisfaction. Finally, Avoidability is a characterisation of the intuition motivating the adoption of R-N, which forms a cornerstone of Elgesem’s approach.

Seegerberg’s (1989) concise formal account of “bringing it about” rests upon propositional dynamic logic, PDL. It aims to explicate seeing to it that A , or δA , through transitions expressed in PDL. Specifically, using this interpretation of a “program”, Seegerberg builds a logic to express action predicated upon achieving states of affairs in which the action operator can be axiomatised without programs being represented explicitly and separately in the object language.

Seegerberg’s logic also differs from the logic of S and T presented here in several ways. The first suggestion of divergence is apparent very quickly, when he describes that “in the case of John eating (all of) an apple, John selects and runs a routine such [that] at the end of that routine John has eaten an apple” (p. 327). This suggests that Seegerberg is aiming towards an operational, postconditional, STRIPS-like view of the world (Fikes and Nilsson, 1971). Despite this, he presents a

clear syntactic division between states and events, in which formulae of propositions in PDL represent states, and terms composed from action descriptions represent events. This division is preserved equally clearly in the semantics, where Segerberg wraps up with: “... the intension of a formula is a proposition while the intension of a term is an action.” (p. 337). But in the syntax, terms (i.e. action descriptions) are constructed solely from the δ operator (similar to Horty and Belnap’s (1995) *a-stit*) and state formulae; semantically, action is construed as a binary relation over states. Our suspicions are confirmed formally when Segerberg presents a modal interpretation explicitly;

$$\models_x [\alpha]C \text{ iff } \forall y(xR(\alpha)y \rightarrow \models_y C)$$

(Segerberg, 1989, p. 337). This is precisely the sort of conflation Hamblin is trying to avoid. For Segerberg, as for most other authors in the area, the interest is upon building a logic of successful action (what he calls “reliable doings”), but, like Chellas this leads him to necessitation, presented on page 333 as Lemma 4.1; if $\vdash A$ then $\vdash [\alpha]A$. Segerberg’s approach is in this way more like Chellas than, say Jones and Sergot. Though there are further and interesting similarities between the logic of S and T and Segerberg’s presentation (such as the “transitivity” of his $[\delta A]B \rightarrow ([\delta B]C \rightarrow [\delta A]C)$ reminiscent of the axioms of delegation presented here), necessitation and the state/event conflation mark fundamental differences in the approach.

Aumann structures (Aumann, 1976) are interesting in that they provide a semantics oriented towards events rather than propositions, and thus offer an alternative means of building a semantic model in place of the Kripkean approach adopted by Segerberg, Jones *et al.* and Chellas. Where Kripkean structures use an accessibility relation to characterise modalities ranging over propositions, Aumann structures partition the set of worlds to yield modal subsets ranging over events. Fagin *et al.* (1995) show how the approach provides a useful basis (and very real alternative to the Kripke style) for characterising multi-agent knowledge. In their description they also portray clearly why the approach is insufficient for our needs; in order to represent knowledge within a system, they refer to “the event of [agent] *i* knowing *e*” (Fagin *et al.*, 1995, p. 38). At best, this is a generous interpretation of something with state-like character construed as an event. The unique property of the Aumann approach is that events can be handled natively, but not that both events and states can be so handled (indeed Fagin *et al.* go on to show how Aumann structures can be reinterpreted in a traditional state-oriented manner).

d’Altan *et al.* (1996) return to von Wright’s (1968) original distinction, exploring the relationship between the deontic notions of ought-

to-be and ought-to-do. They offer powerful arguments for representing both notions separately in an integrated framework (having demonstrated serious problems with attempts to reduce one notion to the other). We adduce those arguments here as further evidence for the utility of Hamblin’s “unusual lavishness”. At the same time, the objectives of d’Altan et al. are quite different in that their aim is to extend Standard Deontic Logic so that it captures the Seinsollen/Tunsollen distinction (though they do not frame it in those terms). Whilst their resulting $PDeL^{AM}$ logic successfully meets that aim, it (i) requires dynamic logic which ties actions to their consequences; (ii) divorces actions from the actors that carry them out; and (iii) excludes an explicit model of time. Furthermore, though S5 and deontic S5 (i.e. KD45) (Chellas, 1980) may arguably be appropriate for deontic characterisations of out-to-do and ought-to-be, respectively, they cannot — as we have shown — be appropriate for more general accounts of agentive responsibility. Any extension of the logic of S and T to deontic expressions, however, should take careful account of the results in (d’Altan et al., 1996) which represents one of the most significant steps forward in the area since von Wright.

Singh (1991; 1993) aims to provide a formal characterisation of wholehearted satisfaction in order to provide a semantics for the successful fulfillment of different classes of speech act including the imperative. His canvas is extremely broad, including not only a specific and novel communication framework replete with intentions and other mental states, but also a modal (but non-Kripkean) interpretation of wholehearted satisfaction, a rich syntax and an analysis not only of speech acts, but also of linguistic mechanisms for expressing the subsequent fulfillment of each type of speech act. As a result of this breadth, the account is necessarily less detailed in places. In particular, the definition of the WSAT modality (the relation that captures his notion of wholehearted satisfaction) is different for different illocutionary forces such as the assertive, the directive and the commissive. For this reason, these definitions offer the basis for the specification of a number of different modalities, none of which, however, is characterised in terms of the axioms underpinning their logics. It is not clear, therefore, how these different notions of wholehearted satisfaction relate to each other. A further significant distinction between Singh’s work and the present account is in the modelling of Hamblin’s intermediary concepts. For although Singh aims at an intuitive counterpart to wholehearted satisfaction, it is not at all clear that he succeeds in capturing Hamblin’s original notion, since it omits not only Hamblin’s “metaphysical backdrop”, but also any notion of partial strategy or the extensional component strategy maintenance. Although there are a number of limitations of Singh’s research in this

area, it should be noted that this body of work represents a significant contribution to the understanding of communication, particularly in a computational setting — a contribution that is not always sufficiently appreciated. We see the work presented in this paper as complimentary to Singh's characterisation of different illocutionary forces in that it offers a more in-depth analysis of just one type of illocutionary force: the imperative.

In closing our review of related work, we turn finally to work in Artificial Intelligence aimed at engineering systems that reason and act in either simulated or real environments. All but the most trivial systems need to be able to represent a model of the world that supports exploration of hypothetical situations. Unlike Hamblin's Action State Semantics with its explicit handling of events and states, most formal models designed for such systems have explicit representation of only one or the other, defining either states in terms of the sequences of events (true of most action and temporal logics) (Davidson, 2001; Kowalski and Sergot, 1986), or else events in terms of a succession of states such as in classical AI planning (Fikes and Nilsson, 1971) and in languages such as Concurrent METATEM (Barringer et al., 1995). Even more recent planning methods, such as (Blum and Furst, 1997) make a fundamental assumption that the world can be represented using a view of actions as state-modifiers. (One very recent thread of planning work (Fox et al., 2005) seems to be moving away from the ubiquity of such a view in tackling action selection against a backdrop of continuous change — although even here, non-environmental actions are classical, discrete operators defined in terms of pre- and post-condition descriptions of state).

The situation calculus (McCarthy and Hayes, 1969) allows both states and events to be represented, but the commonly adopted “axioms of arboreality” (Shanahan, 1997) restrict the flexibility so that a given sequence of events is associated with a single, unique situation. These axioms characterise both theoretical accounts (Pirri and Reiter, 1999) and more practical work (such as the GOLOG programming language (Levesque et al., 1997)). Even if all the fluents in two situations have identical values, under the axioms of arboreality, those two situations are only the same if the events leading to them have also been the same. In Hamblin's work, however, there can be several different histories up to a given state and the histories are not themselves a part of those states. A later advance, the event calculus (Kowalski and Sergot, 1986) has many representational advantages, but is even more restrictive with respect to the problem at hand, in that it is a linear logic, rather than a branching time logic (Shafer et al., 2000).

More recently, an interesting system has been proposed for representing causality directly in a nonmonotonic logic (Giunchiglia et al., 2004). Specifically, causal rules are expressed as $F \Leftarrow G$: F is caused if G is true (notice that this is not quite the same as saying that G causes F directly, merely that if G is true, then F has a cause). The propositional underpinning to the logic is extended in several ways including, most appositely for our current purpose, the “timing” of propositions in situation calculus style (with, syntactically, $t : p$ standing for the truth of proposition p at time point t). Although the theory offers an elegant way of capturing “inertia”, thereby tackling the frame problem (amongst other things), it demands that all actions be treated in the same propositional framework as state descriptions, yielding Hamblin’s “pseudo states” of having just performed (or in the case of (Giunchiglia et al., 2004), being just about to perform) an action. Their language $C+$ provides a high-level, syntactically attractive way of describing domains, but one which still suffers from this basic conflation that Hamblin was trying to avoid.

7. Conclusions

In this paper, our aspiration was to develop a logical theory of action that captures the action-state semantics proposed by Hamblin in 1987. In working towards this goal, our attention has had two principal foci: (i) maintaining in both the semantics and syntax of the language a clear distinction between the notion of seeing to it that a state of affairs is achieved (our S modality) and that of seeing to it that an action is performed (our T modality); and (ii) providing a formal characterisation of wholehearted satisfaction.

In relation to the first focus, we have demonstrated how Hamblin’s approach to the state/event distinction can be preserved syntactically, yielding a simple axiological system, whilst at the same time maintaining clarity in the underlying possible worlds semantics, albeit one based upon a minimal model. As a regular logic, it is not equipollent to previous systems, but offers an intuitive and consistent way of expressing the notion of responsibility. Hamblin’s requirement for the inclusion of time offers a challenge that has here been met in a straightforward manner by defining one accessibility relation in terms of another, an approach designed to admit of alternative modifications to suit different temporal models.

In relation to the second focus, this paper offers, for the first time, a full and accurate characterisation of Hamblin’s conception of wholehearted satisfaction. The characterisation maps Hamblin’s ideas di-

rectly into a Kripkean semantics and RT axiom system, from where those underlying powerful ideas can be harnessed as a modal logic of responsibility.

With this characterisation in place, the system provides a foundation for several new directions, and in particular, for further investigation of imperatival utterance and for exploration of the mechanics of delegation. In both of these areas, the approach supports not only theory building, but also the opportunity for the construction of computational systems that exploit the results of such theory. One of the next steps along this route is to explore the computational character of the logic, including its complexity and operational properties. In this way, we hope to be able to put to work in artificial intelligence at least some of the components of Hamblin's elegant, far-reaching and visionary model.

References

- Åqvist, L.: 1975, *A new approach to the logical theory of interrogatives*. Tübingen, TBL Verlag Gunter Barr.
- Aumann, R. J.: 1976, 'Agreeing to Disagree'. *Annals of Statistics* **4**(6), 1236–1239.
- Austin, J.: 1962, *How to do things with words*. Oxford University Press.
- Barringer, H., M. Fisher, D. Gabbay, G. Gough, and R. Owens: 1995, 'METATEM: An introduction'. *Formal Aspects of Computing* **7**(5), 533–549.
- Belnap, N. and M. Perloff: 1988, 'Seeing to it that: A canonical form for agentives'. *Theoria* **9**, 175–199.
- Belnap, N. and M. Perloff: 1992, 'The way of the agent'. *Studia Logica* **51**(3/4), 463–484.
- Belnap, N., M. Perloff, and M. Xu: 2001, *Facing the Future*. Oxford University Press.
- Blackburn, P., M. de Rijke, and Y. Venema: 2001, *Modal Logic*, Vol. 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.
- Blum, A. L. and M. L. Furst: 1997, 'Fast planning through planning graph analysis'. *Artificial Intelligence* **90**(1–2), 279–298.
- Bowen, K. A.: 1979, *Model Theory for Modal Logic*. D. Reidel Publishing Company.
- Chellas, B. F.: 1980, *Modal logic: An introduction*. Cambridge University Press.
- Chellas, B. F.: 1992, 'Time and modality in the logic of agency'. *Studia Logica* **51**(3/4), 485–517.
- d'Altan, P., J.-J. C. Meyer, and R. J. Wieringa: 1996, 'An integrated framework for out-to-be and ought-to-do constraints'. *Artificial Intelligence and Law* **4**(2), 77–111.
- Davidson, D.: 2001, *Essays on Actions and Events*. Oxford: Clarendon Press, 2nd edition.
- Dignum, F. and J.-J. C. Meyer: 1990, 'Negations of transactions and their use in the specification of dynamic and deontic integrity constraints'. In: M. Kwiatkowska, M. W. Shields, and R. M. Thomas (eds.): *Semantics for Concurrency*. Berlin, pp. 61–80, Springer-Verlag.
- Elgesem, D.: 1997, 'The Modal Logic of Agency'. *Nordic Journal of Philosophical Logic* **2**(2), 1–46.

- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi: 1995, *Reasoning about Knowledge*. MIT Press.
- Feys, R.: 1965, *Modal Logics*. Louvain, E. Nauwelaerts.
- Fikes, R. E. and N. J. Nilsson: 1971, 'STRIPS: A new approach to the application of theorem proving to problem solving'. *Artificial Intelligence* **2**, 189–208.
- Fox, M., R. Howey, and D. P. Long: 2005, 'Validating Plans in the Context of Processes and Exogenous Events'. In: *Proceedings of the Twentieth National Conference on Artificial Intelligence*. pp. 1151–1156.
- Giunchiglia, F., J. Lee, V. Lifschitz, N. McCain, and H. Turner: 2004, 'Nonmonotonic causal theories'. *Artificial Intelligence* **153**(1–2), 49–104.
- Goranko, V. and A. Zanardo: 2004, 'Combining Linear Orders with Modalities for Possible Histories'. In: W. A. Carnielli, F. M. Dionísio, and P. Mateus (eds.): *The Workshop on Combination of Logics: Theory and Applications (CombLog 04)*, Vol. 4(5). http://www.cle.unicamp.br/e-prints/vol_4,n_5,2004.html, CLE e-Prints.
- Hamblin, C. L.: 1987, *Imperatives*. Basil Blackwell.
- Hintikka, J., I. Halonen, and A. Mutanen: 1996, 'Interrogative Logic as a General Theory of Reasoning'. unpublished manuscript.
- Horty, J. F.: 2001, *Agency and deontic logic*. Oxford University Press.
- Horty, J. F. and N. Belnap: 1995, 'The deliberative stit: A study of action, omission, ability, and obligation'. *Journal of Philosophical Logic* **24**, 583–644.
- Hughes, G. E. and M. J. Cresswell: 1996, *A New Introduction to Modal Logic*. Routledge.
- Jones, A. I. J. and M. J. Sergot: 1996, 'A formal characterisation of institutionalised power'. *Journal of the IGPL* **4**(3), 429–445.
- Kowalski, R. A. and M. J. Sergot: 1986, 'A Logic-Based Calculus of Events'. *New Generation Computing* **4**, 67–95.
- Lemmon, E. J.: 1957, 'New foundations for Lewis modal systems'. *Journal of Symbolic Logic* **22**, 176–186.
- Levesque, H. J., R. Reiter, Y. Lesperance, F. Lin, and R. B. Scherl: 1997, 'GOLOG: A logic programming language for dynamic domains'. *Journal of Logic Programming* **31**, 59–84.
- Macleod, M. C. and P. K. Schotch: 2000, 'Remarks on the Modal Logic of Henry Bradford Smith'. *Journal of Philosophical Logic* **29**, 603–615.
- McCarthy, J. and P. Hayes: 1969, 'Some philosophical problems from the standpoint of artificial intelligence'. In: D. Michie and B. Meltzer (eds.): *Machine Intelligence*, Vol. 4. Edinburgh University Press, pp. 463–502.
- Migita, T. and T. Hosoi: 1997, 'Gentzen-Types Formulation of some Modal Logics'. *SUT Journal of Mathematics* **33**(2), 207–237.
- Norman, T. J. and C. A. Reed: 2002, 'Group delegation and responsibility'. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*. pp. 491–498.
- Pirri, F. and R. Reiter: 1999, 'Some contributions to the Metatheory of the Situation Calculus'. *Journal of the ACM* **46**(3), 325–361.
- Prior, A.: 1967, *Past, Present and Future*. Oxford: Clarendon Press.
- Santos, F., A. J. I. Jones, and J. Carmo: 1997, 'Responsibility for Action in Organisations: A Formal Model'. In: G. Holmström-Hintikka and R. Tuomela (eds.): *Contemporary Action Theory*, Vol. 2. Dordrecht: Kluwer Academic Publishers, pp. 333–350.
- Searle, J. R.: 1969, *Speech acts: An essay in the philosophy of language*. Cambridge University Press.

- Segerberg, K.: 1989, 'Bringing it about'. *Journal of Philosophical Logic* **18**, 327–347.
- Shafer, G., P. R. Gillet, and R. Sckerl: 2000, 'The logic of events'. *Annals of Mathematics and AI* **28**(1–4), 315–389.
- Shanahan, M.: 1997, *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. Cambridge, MA: MIT Press.
- Singh, M. P.: 1991, 'Towards a formal theory of communication for multi-agent systems'. In: *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*. pp. 69–74.
- Singh, M. P.: 1993, 'A semantics for speech acts'. *Annals of Mathematics and Artificial Intelligence* **8**, 47–71.
- Thomason, R. H.: 1984, 'Combinations of Tense and Modality'. In: D. Gabbay and F. Guentner (eds.): *Handbook of Philosophical Logic*, Vol. II. Dordrecht: D. Reidel Publishing Company, pp. 135–166.
- von Wright, G. H.: 1968, *An essay in deontic logic and the general theory of action*, Vol. 21 of *Acta philosophica Fennica*. Amsterdam: North-Holland.
- Walton, D. N. and E. C. W. Krabbe: 1995, *Commitment in dialogue: Basic concepts of interpersonal reasoning*. New York: SUNY.
- Xu, M.: 1995, 'On the basic logic of STIT with a single agent'. *Journal of Symbolic Logic* **60**(2), 459–483.

Address for Offprints: Kluwer Prepress Department
P.O. Box 990
3300 AZ Dordrecht
The Netherlands

