Situation Theory

Jerry Seligman
Lawrence S. Moss

Contents

1 Introduction

With the book *Situations and Attitudes*, Barwise and Perry (1983) initiated a program of research that has occupied a diverse group of researchers for many of the subsequent years. The original motivation was to provide a richer, more honest semantic theory of natural languages, one that called a spade a spade, a situation a situation, and (eventually) a proposition a proposition. The emphasis was on defending a naïve ontology containing all the entities to which linguistic expressions appear to refer.

A natural division of labour developed. Linguists were encouraged to provide semantic analyses that used whatever entities they needed, without worrying too much about the technical matter of how such entities should be modelled. Logicians were driven to investigate foundational questions about the emerging ontology, with the hope of providing a unified mathematical framework in which the linguists' work could be interpreted. The linguistic project became known as *Situation Semantics* and the logical project was called *Situation Theory*. Situation Theory was intended to stand to Situation Semantics as Type Theory stands to Montague Grammar.

This chapter is written as a self-contained introduction to the main themes and technical contributions of Situation Theory. We have not attempted to give a survey Situation Semantics, although key references are included in the bibliography. This is a serious omission because many of the ideas in Situation Theory originated arose from linguistic considerations. One excuse is that neither of us is a linguist; a second is that the chapter is too long as it stands.

A more defensible excuse is that *Situations and Attitudes* was not only concerned with semantics. The program of research it initiated has spilled into many neighboring disciplines, such as philosophy, computer science, psychology, and even sociology. The influence of these other fields has been significant—so much so that it would be difficult to account for the purely linguistic factors.

In writing this chapter, we aimed both to present the existing literature on the subject and to provide a coherent statement of Situation Theory. Unfortunately, these goals conflict. The literature contains many technical contributions, some quite sophisticated; but they have not yet been put together. We judged that a mere survey of the existing results would be of limited interest, and decided to re-present the material in as coherent a way as possible.

In any foundational study one has to decide whether to build models or theories. The strategies are distinct if not divergent, and the ideal of a canonical theory is rarely achieved. Most of the research on Situation Theory has been directed towards constructing models. There are many models; they differ both on subject matter (which parts of the theory they model) and on substance (which axioms they satisfy), and no clear consensus has emerged. The work on developing a formal theory has been even less unified: approaches using classical logic, partial logic and even 'illative logic' have been tried. The subject matter of the theories and the models do not coincide completely, and when they do overlap there are many differences on matters of detail. In an effort to make this chapter compatible with all this, we have adopted a fairly abstract, loose approach. Axioms and models are discussed side by side.

In Section 2 we introduce the class of 'simple information structures'. Structures in this class are intended to provide a naïve model of the structure of information, as captured by the relational structures of first-order model theory. The class is axiomatized, and each of the axioms is discussed with a view to generalization. Many of the novelties of Situation Theory can be seen as generalizations of this sort.

Section 3 introduces techniques for constructing models of the generalizations considered in Section 2. All the techniques may be found in the literature but we have re-organized them into a framework called 'the Theory of Structural Relations'. This is a new approach, with which those working on Situation Theory will be unfamiliar.

In Sections 2 and 3 the reader will not meet a single situation. Section 4 remedies the disappointment by introducing a host of ontological categories, including facts, restricted abstracts, situations, propositions, types, and constraints. The theory of structural relations is used throughout, but the focus in this section is on the motivation for diverging opinions in Situation Theory, critical evaluation, and pointers to open problems.

Finally, in Section 5 we give a brief guide to the literature on Situation Theory and related areas of research, followed by a selective bibliography.

2 The Structure of Information

Situation Theory has its fair share of neologisms, the most inspired of which is surely the term *infon*, coined by Keith Devlin. An infon is an item of information. The term is intended to be as neutral as possible about the form in which the information is represented. The common item of information separately conveyed by the statements made in the following little conversation is an example of an infon.

Raymond: (to Paul, displaying an omelette) I have cooked this for you.

Paul: (proudly, to Delia) Raymond cooked this omelette for me.

Delia: (to Albert) The chef cooked the omelette for Paul.

Of course, much more information can be gleaned about the situation from each of these statements, but that it is one in which the two people, Raymond and Paul, and an omelette stand in a certain relationship, that of Raymond having cooked the omelette for Paul, is information expressed by all three. The very same information may also be represented by other means, in a film, or a cartoon-strip, for example, or in the minds of Raymond, Paul, Delia and even Albert.[1]

The first goal of this chapter is to show how information of this simple form is modelled. To specify the structure of the infon in the example, we must say which individuals are involved and how they are related. We shall abbreviate 'Raymond' and 'Paul' to 'R' and 'P', and introduce the name 'O' to refer to the omelette. We use 'cooked' to name the relation that holds between a cook, a diner and a dish just in case the the cook has cooked the dish for the diner. Let I be information that R,P, and O stand in the relation cooked. The individuals R, P, and O are called *arguments* of the infon I, and the relation cooked, is called the *relation* of I.

By specifying the relation and arguments of I we have done a lot toward characterizing everything of informational significance, but not quite all. Care must be taken to distinguish the information that Raymond cooked the omelette for Paul from the information that Paul cooked the omelette for Raymond. For this purpose, we require the concept of a *role*. In any infon involving the relation cooked there are three roles to be filled: that of the cook, the diner and the dish. In the present case, the cook is R, the diner is P, and the dish is O; if instead Paul cooked the omelette for Raymond, then the roles of cook and diner would be reversed. We say that an individual is an *argument* of an infon if it fills one of its roles. The omelette O is an argument of the infon I by virtue of filling the dish role of I's relation cooked.

We have established that there are two basic ingredients determining the structure of I: that of cooked being the relation of I, and that of the three individuals R, P, and O filling the three roles of I, which we name cook, diner, and dish, respectively. We introduce the predicates Rel and Arg to express these structural relations, writing $\mathsf{Rel}(\mathsf{cooked}, \mathsf{I})$ to mean that cooked is the relation of I, and $\mathsf{Arg}(\mathsf{O}, \mathsf{dish}, \mathsf{I})$ to mean that O, the omelette, fills the dish role of our infon I.

[1] The philosophically oriented reader may think to identify this item of information with the Russellian proposition expressed by 'R cooked O for P', with the understanding that 'R', 'O', and 'P' are proper names for Raymond, the omelette, and Paul, respectively. One branch of the theory follows this path (see Section 4) but the identification would be premature.

More generally, we shall take simple infons like I to be 'structured' by the two relations, Rel and Arg. The first of these, Rel, determines the relation of the infon and the other, Arg, determines which arguments fill which roles. We define a 'simple infon' to be anything that has a relation or an argument. Similarly, a 'relation' is anything that is a relation of some infon, and a 'role' is anything that is filled by an argument in some infon. In other words,

σ is a *(simple) infon*	if $\exists r\ \mathsf{Rel}(r,\sigma)$ or $\exists a, i\ \mathsf{Arg}(a,i,\sigma)$,
r is a *relation*	if $\exists \sigma\ \mathsf{Rel}(r,\sigma)$, and
r is a *role*	if $\exists a, \sigma\ \mathsf{Arg}(a,i,\sigma)$.

These definitions form the basis of the abstract account of the structure of information that will be the focus of Section 2. The approach is 'abstract' because we consider any structure in which the predicates 'Rel' and 'Arg' are defined to be a candidate model for the theory of information, with the derived predicates 'relation', 'role' and 'infon' interpreted according to the above definitions. Our strategy is to isolate various axioms describing those properties of infons that have been discussed in the literature on Situation Theory, generalizing from the properties of certain 'standard' models, to be introduced shortly.

First, a short digression.

2.1 Relational Structures

It is a commonplace that in foundational theories one must be very careful to distinguish between meta-language and object-language. Situation Theory is no exception, and special difficulties arises with the word 'relation'. Before going further it will prove useful to make some preliminary definitions concerning relational structures. The following is slightly non-standard but enables us to avoid certain ambiguities with our use of 'relation' and similar words.

A *(relational) signature* is a function $\nu: S \to \mathbb{N}$ from a finite set S to the set $\mathbb{N}$ of positive integers. The elements of S are called *primitive relations* and the number $\nu(R)$ associated with a primitive relation R is a called its *arity*. A *relational structure* $\mathfrak{A}$ of signature ν consists of a class $|\mathfrak{A}|$, called the *universe* of $\mathfrak{A}$, and for each primitive relation R, a class $[\![R]\!]^{\mathfrak{A}}$ of finite sequences of elements of $|\mathfrak{A}|$, each of length $\nu(R)$, called the *extension* of R in $\mathfrak{A}$.[2] For example, we may model the situation of Raymond having cooked an omelette for Paul as a relational structure $\mathfrak{N}_1$ of signature $\nu: \{\mathsf{cooked}\} \to \mathbb{N}$ defined by $\nu(\mathsf{cooked}) = 3$, with universe $|\mathfrak{N}_1| = \{\mathsf{R}, \mathsf{P}, \mathsf{O}\}$ and $[\![\mathsf{cooked}]\!]^{\mathfrak{N}_1}$ containing, at least, the sequence ROP.

[2] We allow relational structures whose universe is a proper class because models of Situation Theory are often very large. The exact meaning of the word 'set' will be discussed below, but for the moment the reader may assume that we are working with any theory of sets that accounts for their mundane properties, supports a distinction between 'set' and 'class', and allows collections of ordinary objects to be called 'sets'.

In discussing relational structures, we say that $\mathfrak{A}$ *is of type* $[A, R_1^{\nu_1}, \ldots, R_n^{\nu_n}]$ if $\mathfrak{A}$ has universe A and primitive relations $R_1, \ldots, R_n$, with arities $\nu_1, \ldots, \nu_n$ respectively. For example, $\mathfrak{N}_1$ is of type $[\{\mathsf{R}, \mathsf{P}, \mathsf{O}\}, \mathsf{cooked}^3]$. The notation is stretched to provide a way of saying that that one relational structure extends another. If $\mathfrak{A}$ is of type $[A, R_1^{\nu_1}, \ldots, R_n^{\nu_n}]$ and $n < m$ then we say that $\mathfrak{B}$ *is of type* $[\mathfrak{A}, R_{n+1}^{\nu_{n+1}}, \ldots, R_m^{\nu_m}]$ if it is of type $[A, R_1^{\nu_1}, \ldots, R_n^{\nu_n}, \ldots, R_m^{\nu_m}]$ and $[\![R_i]\!]^{\mathfrak{B}} = [\![R_i]\!]^{\mathfrak{A}}$ for $1 \leq i \leq n$.

We use the names of the primitive relations of a structure to form sentences about the structure as if they were predicates in a formal language. Given elements $a_1, \ldots, a_{\nu_i}$ of $|\mathfrak{A}|$, the sentence '$R_i(a_1, \ldots, a_{\nu_i})$' when used to describe the structure $\mathfrak{A}$, is true if and only if the finite sequence $a_1, \ldots, a_{\nu_i}$ lies in $[\![R_i]\!]^{\mathfrak{A}}$. For example, the sentence '$\mathsf{cooked}(\mathsf{R},\mathsf{P},\mathsf{O})$' provides yet another way of saying that Raymond cooked the omelette for Paul.

We also need to talk about specific elements of the universe of a relational structure. Given a set C, a relational structure *with constants* C is a relational structure $\mathfrak{A}$ together with an element $[\![c]\!]^{\mathfrak{A}}$ of $|\mathfrak{A}|$ for each $c \in C$.

Relational structures structures $\mathfrak{A}$ and $\mathfrak{B}$, both of signature $\nu\colon S \to \mathbb{N}$ and with constants C, are *isomorphic* if there is a bijective function $f\colon |\mathfrak{A}| \to |\mathfrak{B}|$ such that for each $R \in S$ and $a_1, \ldots, a_{\nu(R)} \in |\mathfrak{A}|$,

$$R(a_1, \ldots, a_{\nu(R)}) \text{ in } \mathfrak{A} \text{ iff } R(f(a_1), \ldots, f(a_{\nu(R)})) \text{ in } \mathfrak{B}$$

and for each $c \in C$, $[\![c]\!]^{\mathfrak{B}} = f([\![c]\!]^{\mathfrak{A}})$.

2.2 Simple Information Structures

We are interested in relational structures in two quite different ways. First, we may use them to build models of a given situation. The structure $\mathfrak{N}_1$, defined in the previous section, models the situation discussed by Raymond and his friends in our initial example. It is the sort of model that underlies standard 'model-theoretic' semantics. In the manner of Tarski, we may provide a definition of truth in $\mathfrak{N}_1$ for a language L having the predicate 'cooked' and perhaps the names 'R', 'O', and 'P'. By virtue of this semantic function, structures like $\mathfrak{N}_1$ may also be thought to provide a model of the *information* conveyed in the conversation, but in this capacity they are rather limited. In constructing $\mathfrak{N}_1$ we are forced to determine the complete extension of cooked and so include more information that was conveyed. Consider the sequence ROR. If this lies in $[\![\mathsf{cooked}]\!]^{\mathfrak{N}_1}$ then the situation modelled is one in which Raymond also cooked the omelette for himself, and O is a two-person omelette. If, on the other hand, ROP is in not in $[\![\mathsf{cooked}]\!]^{\mathfrak{N}_1}$ then Paul has the feast to himself. The necessary overdetermination of information is one of the deficiencies that Situation Theory is intended to overcome.

Our second, more important use of relational structures is as models of fragments of Situation Theory itself. Our strategy is to define successive classes of relational structures, that may be regarded as approximations to a suitable model for the theory as a whole. An example of the second use is given by our first substantial definition.

Definition 2.1 A relational structure $\mathfrak{A}$ of type $[A, \mathsf{Rel}^2, \mathsf{Arg}^3]$ is a *(simple) information structure* if it satisfies the following conditions:

A 1 *If* $\mathsf{Rel}(r, \sigma)$ *and* $\mathsf{Rel}(r', \sigma)$ *then* $r = r'$*.*

A 2 *If* $\mathsf{Arg}(a, i, \sigma)$ *and* $\mathsf{Arg}(b, i, \sigma)$ *then* $a = b$*.*

A 3 *No relation is a role.*

A 4 *No role is an infon.*

A 5 *No infon is a relation.*

Axioms A1 to A5 are not particularly stringent; they merely impose some discipline on the domain. Together, they say nothing more than that the relations, roles and infons form disjoint classes, and that infons may have at most one relation and at most one argument filling each role.

The previously mentioned ambiguity in the word 'relation' can be seen clearly with the help of this definition.[3] Firstly, we use the word to talk about the relations of any relational structure. For example, the theoretical predicates 'Rel' and 'Arg' refer to the primitive *relations* of simple information structures. Relations of this kind will be dubbed '*external*'. Any other relation between elements of a relational structure whose extension is determined by our theoretical language is also called 'external'. Examples include relations defined from primitive relations, such as the property of being a simple infon.[4]

The second use of 'relation' is governed by the extension of Rel in a specific information structure. An element r of the universe of a simple information structure $\mathfrak{A}$ is an *internal* relation of $\mathfrak{A}$ if $\mathsf{Rel}(r, \sigma)$ for some infon σ. The ambiguity, once understood, should not cause the reader any difficulty. Where necessary the sense will be made clear using the adjectives 'internal' and 'external'.

Aside from illustrating ambiguities, simple information structures provide our first, crude approximation to the models of a fragment of Situation Theory. Their principal use is to show how the two ways in which we use relational structures are related. Suppose, for example, we use the structure $\mathfrak{N}_1$ to model the situation discussed by Raymond and his friends. There is a sense in which

[3] Other words we use, such as 'role' and 'argument', have a similar ambiguity.

[4] Typically, an n-ary external relations is identified with its *graph*, the set, or class, of sequences $x_1 \ldots x_n$ of objects that stand in the relation. We make no such assumption here. A meta-theory in which intensional distinctions are made between relations—as they are in our object-theory—is quite compatible with our approach.

$\mathfrak{N}_1$ gives us all we need to know about the common informational content of the three culinary pronouncements. The relation of the infon, cooked, is included as a primitive relation of $\mathfrak{N}_1$; the three roles of this relation, cook, diner, and dish, are modelled by the first, second, and third positions in the sequences in the extension of cooked; and the arguments—Raymond, Paul and the omelette—are also directly included in the universe of $\mathfrak{N}_1$. The thought that we have thereby obtained a simple information structure is made more precise as follows.

Construction 2.1 Given a relational structure $\mathfrak{M}$ of type $[M, R_1^{\nu_1}, \ldots, R_n^{\nu_n}]$, possibly with constants, we construct an information structure $\mathrm{SInf}(\mathfrak{M})$. First, let

$$\begin{array}{lcl} A_1 & = & \{1\} \times \{R_1, \ldots, R_n\} \\ A_2 & = & \{2\} \times \bigcup_{1 \le i \le n} \{1, \ldots, \nu_i\} \\ A_3 & = & \{\langle 3, R_i, \alpha \rangle \mid 1 \le i \le n \text{ and } \alpha \colon \{1, \ldots, \nu_i\} \to M\} \end{array}$$

The basic idea is to model simple infons as pairs of the form $\langle R_i, \alpha \rangle$ in which α is a function from $\{1, \ldots, \nu_i\}$ to the universe of $\mathfrak{M}$. But we must be careful to keep relations, roles and infons distinct, and so we also include a numerical prefix indicating the sort of object modelled: 1 for relations, 2 for roles, and 3 for simple infons. The information structure $\mathrm{SInf}^-(\mathfrak{M})$ has universe $M \cup A_1 \cup A_2 \cup A_3$ and extensions of Rel and Arg given by

$$\begin{array}{ll} \mathsf{Rel}(r, \sigma) & \text{iff } \exists R, \alpha \text{ with } r = \langle 1, R \rangle \text{ and } \sigma = \langle 3, R, \alpha \rangle, \text{ and} \\ \mathsf{Arg}(a, i, \sigma) & \text{iff } \exists j, R, \alpha \text{ with } i = \langle 2, j \rangle,\ \sigma = \langle 3, R, \alpha \rangle, \text{ and } \alpha(j) = a. \end{array}$$

The constants of $\mathrm{SInf}^-(\mathfrak{M})$ include all the constants of $\mathfrak{M}$. In addition, we take each primitive relation R of $\mathfrak{M}$ to be a constant denoting $\langle 1, R \rangle$. It is a simple matter to check that the resulting structure $\mathfrak{A}$ satisfies axioms A1 to A5.

A minor complication arises because of the need to ensure that isomorphic relational structures give rise to isomorphic information structures—we don't want accidental features of the representation to encode more information than there is in the original structure. We would get more if, for example, M happened to contain an element $m = \langle 1, R \rangle$ in which R is one of the primitive relations of $\mathfrak{M}$. In the structure $\mathrm{SInf}^-(\mathfrak{M})$, this pair would serve both as the object m of $\mathfrak{M}$ and as the internal relation denoted by the constant 'R'.

We overcome this difficulty by first making a copy of $\mathfrak{M}$. Let $\mathfrak{M}_0$ be the isomorphic copy of $\mathfrak{M}$ in which each element m of M is replaced by the pair $\langle 0, m \rangle$, and let $\mathrm{SInf}(\mathfrak{M}) = \mathrm{SInf}^-(\mathfrak{M}_0)$. Unlike $\mathrm{SInf}^-(\mathfrak{M})$, the structure $\mathrm{SInf}(\mathfrak{M})$ keeps the old elements distinct from the newly constructed elements, and so our construction has the desired property of preserving isomorphism: if $\mathfrak{M}_1$ and $\mathfrak{M}_2$ are isomorphic, so are $\mathrm{SInf}(\mathfrak{M}_1)$ and $\mathrm{SInf}(\mathfrak{M}_2)$.[5]

[5] Note also that the construction depends only on the type of the relational structure. If $\mathfrak{M}_1$ and $\mathfrak{M}_2$ are isomorphic relational structures of the same type then indeed $\mathrm{SInf}(\mathfrak{M}_1) = \mathrm{SInf}(\mathfrak{M}_2)$.

For example, consider once again the model $\mathfrak{N}_1$. The information structure $\mathrm{SInf}(\mathfrak{N}_1)$ has individuals $\langle 0, \mathsf{R} \rangle$, $\langle 0, \mathsf{P} \rangle$, and $\langle 0, \mathsf{O} \rangle$, copied from $\mathfrak{N}_1$, a single relation $\langle 1, \mathsf{cooked} \rangle$, three roles $\langle 2, 1 \rangle$, $\langle 2, 2 \rangle$, and $\langle 2, 3 \rangle$, and infons of the form $\langle 3, \langle \mathsf{cooked}, \alpha \rangle \rangle$, where α is a function from $\{1, 2, 3\}$ into $\{\langle 0, \mathsf{R} \rangle, \langle 0, \mathsf{P} \rangle, \langle 0, \mathsf{O} \rangle\}$. The structure also has a constant cooked that denotes the internal relation $\langle 1, \mathsf{cooked} \rangle$. The information that Raymond cooked the omelette O for Paul is modelled by the element $\langle 3, \langle \mathsf{cooked}, \alpha \rangle \rangle$ in which α is the function with domain $\{1, 2, 3\}$, and such that $\alpha(1) = \langle 0, \mathsf{R} \rangle$, $\alpha(2) = \langle 0, \mathsf{P} \rangle$, and $\alpha(3) = \langle 0, \mathsf{O} \rangle$.

Information structures constructed in this way serve as the starting point in our investigation into the structure of information. But we must proceed with care. Various aspects of the construction are artifacts of the set-theoretic tools used. For example, the choice of the numbers 0 to 3 as indicators of the sort of object modelled is quite arbitrary. In making theoretical claims about the properties of these structures we must consider only on the essential aspects of the construction. This is easily achieved by widening our focus to include all structures isomorphic to one obtained by Construction 2.1.

Definition 2.2 A simple information structure is *standard* if it is isomorphic to $\mathrm{SInf}(\mathfrak{M})$ for some relational structure $\mathfrak{M}$.

In the remainder of Section 2 we shall study standard information structures in some detail. We shall see that some properties of the standard structures capture our basic intuitions about the structure of information, while others reflect only the limitations of the construction. By the end of the section we shall have a list of axioms characterizing the class of standard structures, and a number of alternative axioms designed to overcome their limitations.[6]

2.3 Roles

Suppose $\mathfrak{N}_2$ is a relational structure of type $[N_2, \mathsf{stir\text{-}fried}^2, \mathsf{braised}^2]$ and that R and P are elements of M. Furthermore, let α_1 and α_2 be functions with domain $\{1, 2\}$ and such that $\alpha_1(1) = \mathsf{R}$ and $\alpha_2(1) = \mathsf{P}$. The information structure $\mathrm{SInf}(\mathfrak{N}_2)$ contains the simple infons $\sigma_1 = \langle 3, \mathsf{stir\text{-}fried}^2, \alpha_1 \rangle$ and $\sigma_2 = \langle 3, \mathsf{braised}^2, \alpha_2 \rangle$. Thus, in $\mathrm{SInf}(\mathfrak{N}_2)$, we have that

$$\mathsf{Arg}(\langle 0, \mathsf{R} \rangle, \langle 2, 1 \rangle, \sigma_1) \text{ and } \mathsf{Arg}(\langle 0, \mathsf{P} \rangle, \langle 2, 1 \rangle, \sigma_2).$$

The elements $\langle 0, \mathsf{R} \rangle$ and $\langle 0, \mathsf{P} \rangle$ have the common property of filling the role $\langle 2, 1 \rangle$. What, if anything, does this signify?

In information structures obtained using Construction 2.1, roles are pairs of the form $\langle 2, n \rangle$, where n is a positive integer. These roles are simply indices

[6] The axioms for standard information structures are labelled **A1**, **A2**, and so on. The alternatives to Axiom **A**n are labelled **A**n**.1**, **A**n**.2**, and so on. Some of the alternatives are weaker than the corresponding standard axioms, generalizing some aspect of standard structures. Others are stronger, asserting the existence of objects not present in standard structures.

recording the relative position of arguments in original relational structure, and it is difficult to see how the identity of roles—mere indices—in different infons can be given much import. Yet we must not be mislead by the peculiarities of the coding used in the construction. The statement that Raymond stir-fried the frogs' legs and Paul braised the monkfish tells us, among many other things, that both Raymond and Paul are cooks; at least, it tells us that they played the role of cook in the situation described. Perhaps we can regard $\langle 2, 1\rangle$ as modelling this role in $\mathrm{SInf}(\mathfrak{N}_2)$.

Whatever the merits of this suggestion, the behaviour of roles in standard information structures is severely constrained. The roles of $\mathsf{stir\text{-}fried}$ and $\mathsf{braised}$ in $\mathfrak{N}_2$ are the same two, $\langle 2, 1\rangle$ and $\langle 2, 2\rangle$, not for culinary reasons but simply because both are binary relations and so use the same numbers, 1 and 2, to index their arguments. This rather artificial limitation of standard structures is captured by the following rather artificial axiom.

A 6 *If σ has a role that is not a role of τ, then every role of τ is a role of σ.*

In effect, Axiom A6 says that the sets of roles of infons are linearly ordered by inclusion. In moving away from standard structures, this axiom is sure to be dropped. Theoretical considerations concerning roles have been studied in connection with the linguistic concept of 'thematic role' by Engdahl (1991). Another restriction obeyed by standard structures is that

A 7 *each infon has only finitely many roles.*

Relations with infinite arities are conceivable, and perhaps of some theoretical use, but it is important to be able to focus on those items of information that are finitely expressible. Indeed, it has been proposed (Devlin, 1991) that the finiteness of infons is an essential characteristic. A much less reasonable restriction is that

A 8 *There are only finitely many relations.*

This is satisfied by all standard information structures but it will be dropped when we generalize.

2.4 Identity of Basic Infons

In all standard information structures, the following criterion of identity is satisfied.

A 9 *Suppose σ and τ are infons such that for all r, a, and i,*

1. $\mathsf{Rel}(r, \sigma)$ *iff* $\mathsf{Rel}(r, \tau)$, *and*
2. $\mathsf{Arg}(a, i, \sigma)$ *iff* $\mathsf{Arg}(a, i, \tau)$.

Then $\sigma = \tau$.

In many presentations of Situation Theory there is a further condition, relating to the 'polarity' of an infon. The information expressed by 'Amelia is loquacious' and 'Amelia is not loquacious' are considered to be on an equal footing, instead of taking the negative form to be a constructed from the positive by an operation of negation, as is done in propositional logic. The infons have the same relation (being loquacious) and the same argument (Amelia) filling the same role. They are distinguished by their *polarity*, which can be either *positive* or *negative*.

We can incorporate this proposal into the present account in various ways. For example, we could introduce a new unary relation, Pos, that holds of the positive infons only, and axioms to give the modified identity conditions and to ensure that for every positive infon there is a corresponding negative infon. For the moment, we shall keep our basic definitions uncluttered by matters of polarity; but we return to the topic when discussing complex infons in Section 3.6.

Considerations of polarity aside, Axiom A9 has found widespread support. Moreover, it allows the use of a convenient *functional notation*. For any finite infon σ there is a relation r and a finite sequence $\langle i_1, a_1\rangle, \ldots, \langle i_n, a_n\rangle$ of role-argument pairs, such that $\mathsf{Rel}(r, \sigma)$ and $\mathsf{Arg}(a_j, i_j, \sigma)$ for $1 \leq j \leq n$ (and no other pair $\langle i, a\rangle$ is such that $\mathsf{Arg}(a, i, \sigma)$). Given Axiom A9, the identity of σ is completely determined by this information, and so we may write σ unambiguously as

$$\langle\!\langle\ r; i_1 \colon a_1, \ldots, i_n \colon a_n\ \rangle\!\rangle$$

The set of role-argument pairs in a basic infon is called an *assignment*. Although confined by the linearity of text, the elements of the assignment are intended to be unordered, so that, for example, $\langle\!\langle\ r; i : a, j : b\ \rangle\!\rangle$ and $\langle\!\langle\ r; j : b, i : a\ \rangle\!\rangle$ denote the same infon—this follows from Axiom A9, of course.[7]

The functional notation and its variants are widely used. Indeed, it is tempting to base the subsequent development of Situation Theory on an infon-building *function* that maps each pair consisting of a relation r and an assignment $\langle\langle i_1, a_1\rangle, \ldots, \langle i_n, a_n\rangle\rangle$ to the infon $\langle\!\langle\ r; i_1 \colon a_1, \ldots, i_n \colon a_n\ \rangle\!\rangle$. A problem is that the infon-building function is partial. Even in standard information structures, not every relation-assignment pair determines an infon; an infon will result only if the length of assignment is the same as the arity of the relation.

Faced with partiality there are three natural responses. The first (Plotkin, 1991), is to have an infon-building function that is total but does not always deliver an infon. The advantage is an uncomplicated model; the disadvantage is the existence of many 'pseudo-objects' in the universe of the model that fail to model anything 'real'. The second response (Plotkin (1991), Barwise and Cooper (1991)) is to accept that the infon-building function is partial, and make

[7] The use of angle brackets in the functional notation is suggestive of the set-theoretic model of infons used in Construction 2.1; it is important not to confuse the two.

do with a partial language (and logic) for describing situation-theoretic objects. The third response (Muskens (1995), Robin Cooper (1991) , and this chapter) is to base the theory on relations instead of functions. The advantage is that we may retain the services of classical logic and our models remain uncomplicated. The disadvantage is that we must abandon the functional notation, at least for theoretical purposes.

In less formal discussions we shall still use terms of the form '$\langle\!\langle\, r; i : a, j : b \,\rangle\!\rangle$', with the presupposition that there is an infon to which the term refers. When the roles are clear from the context, or are simply numerical indices in the canonical order ($i = 1$, $j = 2$), we use the abbreviated form '$\langle\!\langle\, r; a, b \,\rangle\!\rangle$'.

The identity criteria discussed above are only intended to apply to infons of the simplest kind. Quite different criteria will be needed for complex infons, such as conjunctions and disjunctions. These will be discussed in Section 3.4.

2.5 Arguments

Standard information structures have the property that

A 10 *no argument is an infon, relation or role.*

In other words, the information modelled is of a purely 'non-reflexive' kind: no information *about* infons, relations, or roles is included. This is a severe limitation whose removal has motivated a number of important developments in Situation Theory.

A reason for rejecting Axiom A10 in favour of a more liberal framework is the need to express 'higher-order' information—information about infons. Ascriptions of propositional attitudes and conditionals are the obvious examples. For example, if σ is the information that Albert is replete, then the information that Delia knows that Albert is replete may be thought to involve some relation between Delia and σ (and possibly other things). A first candidate for modelling this would be an infon of the form $\langle\!\langle\, \mathsf{knows}; \mathsf{Delia}, \sigma \,\rangle\!\rangle$. Likewise, if τ is the information that Albert is snoring, then we might regard the conditional 'if Albert is snoring then he is replete' as expressing the information $\langle\!\langle\, \text{if}, \sigma, \tau \,\rangle\!\rangle$.[8]. We might also wish to model information about Situation Theory itself. For example, the information that r is the relation of σ may be modelled by an infon of the form $\langle\!\langle\, \mathsf{Rel}; r, \sigma \,\rangle\!\rangle$.

Recalling Construction 2.1, it is easy to see why Axiom A10 is satisfied. In $\mathrm{SInf}(\mathfrak{M})$ all arguments of infons are of the form $\langle 0, m\rangle$ but infons, relations and roles are of the form $\langle i, a\rangle$ for $i \neq 0$. As we have seen, these sortal restrictions are needed to ensure that the construction preserves isomorphism. They are enforced in the final move from $\mathrm{SInf}^{-}(\mathfrak{M})$ to $\mathrm{SInf}(\mathfrak{M})$ in which a copy is made of $|\mathfrak{M}|$. We can obtain models constraining higher-order infons by dropping this step. For example, suppose $\mathfrak{M}$ is of type $[M, R_1^1, R_2^2]$ and M happens to contain

[8] Conditionals are reconsidered in Section 4.9

an element $\sigma = \langle 3, R_2, \alpha\rangle$, in which α is a function from $\{1,2\}$ to M. Then σ is classified as an infon because of its set-theoretic form. But it is also an element of M and so may occur as a argument to another infon. If β is the function with domain $\{1\}$ and $\beta(1) = \sigma$ then $\tau = \langle 3, R_1, \beta\rangle$ is such a 'higher-order' infon. In the functional notation, $\tau = \langle\!\langle\ R_1; \sigma\ \rangle\!\rangle$. The nearest counterpart in $\mathrm{SInf}(\mathfrak{M})$ is $\langle\!\langle\ R_1; \langle 0, \sigma\rangle\ \rangle\!\rangle$, which is not the same thing at all.

The method may be made more systematic by iteratating the construction of $\mathrm{SInf}^-(\mathfrak{M})$.

Construction 2.2 Given a relational structure $\mathfrak{M}$, we define an infinite sequence $\mathfrak{A}_0, \mathfrak{A}_1, \mathfrak{A}_2, \ldots$ of information structures as follows. Let $\mathfrak{A}_0$ be $\mathfrak{M}$, and for each integer n, let $\mathfrak{A}_{n+1}$ be $\mathrm{SInf}^-(\mathfrak{A}_n)$. Let $\mathfrak{A}$ be the 'iterated' structure defined by taking unions:

$$|\mathfrak{A}| = \bigcup_{n\in\mathbb{N}} |\mathfrak{A}_n|, \quad [\![\mathsf{Rel}]\!]^{\mathfrak{A}} = \bigcup_{n\in\mathbb{N}} [\![\mathsf{Rel}]\!]^{\mathfrak{A}_n}, \quad \text{and} \quad [\![\mathsf{Arg}]\!]^{\mathfrak{A}} = \bigcup_{n\in\mathbb{N}} [\![\mathsf{Arg}]\!]^{\mathfrak{A}_n}.$$

It is simple to check that $\mathfrak{A}$ is an information structure, and that it is a fixed point: $\mathrm{SInf}^-(\mathfrak{A}) = \mathfrak{A}$. Moreover, like Construction 2.1, this construction preserves isomorphism.

Axiom A10 fails in the iterated structures. In addition to the infons of $\mathrm{SInf}(\mathfrak{M})$, the iterated structures contains infons of the forms $\langle\!\langle\ \mathsf{Rel}; r, \sigma\ \rangle\!\rangle$ and $\langle\!\langle\ \mathsf{Arg}; a, i, \sigma\ \rangle\!\rangle$. Other 'higher-order' infons can be incorporated by extending the definition of $\mathrm{SInf}^-(\mathfrak{M})$ appropriately.[9]

2.6 Circularity

Despite escaping Axiom A10, the arguments of infons in iterated information structures are not entirely without restriction. The following axiom is satisfied by all iterated information structures, as well as by all standard ones.

A 10.1 *There is no infinite sequence $a_0, a_1, a_2, \ldots$ such that, for each integer n, a_{n+1} is an argument of a_n.*

Axiom A10.1 is a weaker than Axiom A10, which disallows even finite sequences of nested arguments, but it is strong enough to prohibit 'circular' information. For example, let σ be the information expressed by the following sentence: 'The information expressed by this sentence is expressible in English'. A simple-minded analysis of the structure of σ is that it has a relation $\mathsf{expressible\text{-}in\text{-}English}$ with a single role i, and is such that $\mathsf{Arg}(\sigma, i, \sigma)$. In other words, σ satisfies the equation:

[9] For example, in addition to A_1, A_2, and A_3, let the universe include

$$\begin{aligned} A_4 &= \{\langle 4, \mathsf{knows}, \alpha\rangle \mid \alpha\colon \{1,2\} \to M \text{ and } \exists \mathrm{x}, \mathrm{y}\ \alpha(1) = \langle 0, \mathrm{x}\rangle,\ \alpha(2) = \langle 3, \mathrm{y}\rangle\} \text{ and} \\ A_5 &= \{\langle 5, \mathsf{if}, \alpha\rangle \mid \alpha\colon \{1,2\} \to M \text{ and } \exists \mathrm{x}, \mathrm{y}\ \alpha(1) = \langle 3, \mathrm{x}\rangle,\ \alpha(2) = \langle 3, \mathrm{y}\rangle\}. \end{aligned}$$

$$\sigma \quad = \quad \langle\!\langle \,\mathsf{expressible\text{-}in\text{-}English}; \sigma\, \rangle\!\rangle$$

Then then infinite path $\sigma, \sigma, \sigma, \ldots$ is of kind forbidden by Axiom A10.1. Infons with the self-referential properties that conflict with Axiom A10.1 have been called 'hyperinfons'. There are other examples.

Suppose that Paul and Albert both know that the soufflé Paul cooked has failed to rise. Let τ be this information. On one analysis of common knowledge (Barwise, 1989), we would also expect that Paul knows that Albert knows τ and vice versa. Each of these items of information are also known by Paul and by Albert, and so on *ad infinitum.* If we allow ourselves to use conjunctive infons (Section 3.4), the shared information σ may be thought to satisfy the following equation:

$$\sigma \quad = \quad \tau \wedge \langle\!\langle \,\mathsf{knows}, \mathsf{Albert}, \sigma\, \rangle\!\rangle \wedge \langle\!\langle \,\mathsf{knows}, \mathsf{Paul}, \sigma\, \rangle\!\rangle$$

Next, suppose that being incorrect is a property of infons, and let '$\mathsf{incorrect}$' name this property. Any infon σ satisfying the equation

$$\sigma \quad = \quad \langle\!\langle \,\mathsf{incorrect}; \sigma\, \rangle\!\rangle$$

leads to some tricky problems for a theory of the relationship between information and truth (to be discussed in Section 4). Perhaps there are no such infons, but then one wonders what information is expressed by the statement that the information expressed by this statement is incorrect. Many similar puzzles are discussed in Barwise and Etchemendy (1987). Koons (1990, 1992) uses a similar approach to tackle doxastic puzzles involving self-reference.

The elucidation of the structure of hyperinfons and other 'circular phenomena' is one of Situation Theory's main achievements. It is also one of the main reasons for the present, relational approach. Although Axiom A10.1 is satisfied by all standard information structures, there are non-standard structures in which hyperinfons exist. This is easy to see: the structure $\mathfrak{A}_1$ with just three elements $\{a, r, i\}$ and with extensions $[\![\mathsf{Arg}]\!] = \{\langle a, i, a\rangle\}$ and $[\![\mathsf{Rel}]\!] = \{\langle r, a\rangle\}$ is a simple information structure in which the element a satisfies the equation $a = \langle\!\langle\, r; a\, \rangle\!\rangle$.

When abandon Axiom A10.1 in order to model hyperinfons, Axiom A9 is no longer sufficient to determine identity. Consider the structure $\mathfrak{A}_2$ with universe $\{a, b, r, i\}$ and extensions $[\![\mathsf{Arg}]\!] = \{\langle a, i, a\rangle, \langle b, i, b\rangle\}$ and $[\![\mathsf{Rel}]\!] = \{\langle r, a\rangle, \langle r, b\rangle\}$. In this structure there are two infons a and b that satisfy the equations $a = \langle\!\langle\, r; a\, \rangle\!\rangle$ and $b = \langle\!\langle\, r; b\, \rangle\!\rangle$. Intuitively, there is no difference in the structure of a and b, and yet they are distinct, despite the fact that Axiom A9 is satisfied. What went wrong?

The solution to this problem was a turning point in the development of Situation Theory. The crucial concept of 'bisimulation' first appeared in modal logic and in work in theoretical computer science, where it is used to define equivalence of processes. This latter development lead directly to Aczel (1988), where it is used to formulate the Anti-Foundation Axiom (AFA), an alternative to the Axiom of Foundation in Set Theory. The resulting theory of sets offered a rich new class of structures with which to construct models of circular objects, such as hyperinfons.

Applied to simple information structures, 'bisimulation' is defined as follows.

Definition 2.3 Given an information structure $\mathfrak{A}$, a binary relation R on A, is a *bisimulation* iff for all σ, τ in A, if $R(\sigma, \tau)$ then

1. if either σ or τ are not infons then $\sigma = \tau$,
2. if σ and τ are infons then they have the same relation and roles,
3. $\forall i, a, b$ if $\mathsf{Arg}(a, i, \sigma)$ and $\mathsf{Arg}(b, i, \tau)$ then $R(a, b)$, and
4. $R(\tau, \sigma)$.

For any elements a, b in A, we say that a is *bisimilar* to b iff there is a bisimulation R such that $R(a, b)$. It is easy to show that the relation of being bisimilar is itself a bisimulation, the largest bisimulation in $\mathfrak{A}$.

A more general definition of bisimulation is given in Section 3.1, but the one above will do for now. It allows use to state an improved principle of identity for infons:

A 10.2 *If a is bisimilar to b then $a = b$.*

Returning to our example, we see that the structure $\mathfrak{A}_2$ fails to satisfy Axiom A10.2 because a and b are bisimilar but not identical. Axiom A10.2 is strictly stronger than Axiom A9, but equivalent to it when restricted to the class of structures satisfying Axiom A10.1, which includes the standard structures. At this stage it is important to realize that we have only shown the feasibility of constructing models of hyperinfons. The question of whether all 'circular phenomena' in Situation Theory can be modelled in this way is as yet undecided. A positive answer will be given in Section 3.2.

2.7 Appropriateness

Let's say that an object is *ordinary* if it is neither a role, nor an infon, nor a relation. In standard structures, for each relation r of arity n and each sequence $a_1 \ldots a_n$ of ordinary objects, the infon $\langle\!\langle\, r; 1{:}\, a_1, \ldots, n{:}\, a_n \,\rangle\!\rangle$ exists. We call this property of standard structures 'generality'. To state it precisely, we say that infons σ and τ each with role i are *i-variants* if for each role $j \neq i$ and each a, $\mathsf{Arg}(a, j, \sigma)$ iff $\mathsf{Arg}(a, j, \tau)$. With this terminology, the principle of generality, which is respected by all standard information structures, may be stated as follows.

A 11 *For each infon σ with role i and each ordinary object a, there is an i-variant τ of σ such that* $\mathsf{Arg}(a, i, \sigma)$.

There are reasons for thinking that Axiom A11 is both too weak and too strong. It is too weak because it says nothing about the appropriateness of arguments that are not ordinary. Removing the restriction to ordinary objects, we get the following:

A 11.1 *For each infon σ and with role i, and each object a, there is an i-variant τ of σ such that* $\mathsf{Arg}(a, i, \sigma)$.

Axiom A11.1 is not satisfied by standard information structures. The iterated structures of Construction 2.2 obey a version of the axiom when restricted to infons with relation Rel or Arg.

The reason for thinking that Axioms A11 and A11.1 are both too strong is that they do not permit sortal restrictions. For example, the information expressed by 'Albert tasted the *crème brûlée*' presupposes that Albert is an agent, perhaps even that he is subject of conscious experience. Good evidence for this is that the presupposition accompanies the claim that Albert did not taste the *crème brûlée*, and we would be surprised to the point of incomprehension if we were told later that Albert is a Moulinex food processor. Sortal presuppositions like these may be explained by supposing that the argument roles of infons carry certain restrictions on what can fill them. The omelette O of our original example simply cannot fill the role of cook played by Raymond, and vice versa. This may be captured by ensuring that there is no infon τ such that $\mathsf{Arg}(0, \mathsf{cook}, \tau)$. Since there is an infon σ such that $\mathsf{Arg}(\mathsf{R}, \mathsf{cook}, \tau)$ and O is an ordinary object, this is inconsistent with Axiom A11.

We say that an object a is an *appropriate filler* of role i just in case there is an infon σ such that $\mathsf{Arg}(a, i, \sigma)$. To allow for role-linked sortal restrictions, we may modify Axiom A11 to the principle of 'sortal generality':

A 11.2 *for each infon σ with role i and each appropriate filler a of i, there is an i-variant τ of σ such that* $\mathsf{Arg}(a, i, \sigma)$.

This is a great improvement. In effect, we have associated a domain of objects with each role—their appropriate fillers—and allowed unrestricted replacement of one argument by any other argument in the domain. In standard information structures, the domain of appropriate fillers for each role is just the class of ordinary objects, but other structures may have different domains for different roles. The iterated structures of Construction 2.2 all satisfy Axiom A11.2. They have two domains: roles of the original structure may be filled by any ordinary object, and roles of infons whose relation is Arg of Rel may be filled by any object at all.

A theory of appropriateness, stating in general terms which arguments may be fill which roles, has not yet been given. Axiom A11.2 accords with the way roles are usually thought to restrict their fillers, but there may be restrictions on the formation of infons that are not attributable to a single role. For example, consider a sketch of two people standing side-by-side, and call the depicted relationship between the people '$\mathsf{next\text{-}to}$'. It is arguable that $\mathsf{next\text{-}to}$ is not the same as the relation expressed by English preposition 'next to' because any picture of two people standing side-by-side presupposes that there are *two* people rather than one. We may analyze the difference by saying that the two roles of $\mathsf{next\text{-}to}$ place a joint restriction on their fillers, namely, that they cannot be filled by the same argument. This kind of restriction is not permitted by Axiom A11 or by Axiom A11.2.

2.8 Partiality

In standard information structures every infon has a relation and a set of roles determined by that relation. This is captured by the following two axioms.

A 12 *For each infon σ there is a relation r such that* $\mathsf{Rel}(r, \sigma)$.

A 13 *If* $\mathsf{Rel}(r, \sigma)$ *and* $\mathsf{Rel}(r, \tau)$ *and* $\mathsf{Arg}(a, i, \sigma)$ *then for some b,* $\mathsf{Arg}(b, i, \tau)$.

One reason for dropping these axioms is to allow infons to be *unsaturated*. For example, suppose you overhear someone saying 'Mary saw Elvis in Tokyo' but a distracting noise prevents you from hearing the word 'Elvis'. The information conveyed is unsaturated because the filler of one of the roles is missing. We may represent the unsaturated infon as $\langle\!\langle\ \mathsf{saw};\ \mathsf{seer}\colon\ \mathsf{Mary},\ \mathsf{location}\colon\ \mathsf{Tokyo}\ \rangle\!\rangle$, to be contrasted with the infon

$$\langle\!\langle\ \mathsf{saw};\ \mathsf{seer}\colon\ \mathsf{Mary},\ \mathsf{seen}\colon\ \mathsf{Elvis},\ \mathsf{location}\colon\ \mathsf{Tokyo}\ \rangle\!\rangle$$

which would have been conveyed had you heard the whole utterance. The coexistence of these infons is forbidden by Axiom A13, which must therefore be dropped if they are to be modelled successfully. Another way in which an infon can be unsaturated is by lacking a relation—consider, for example, the information conveyed if the word 'saw' had been obscured. This counts against Axiom A12.

In the absence of the constraints of Axioms A12 and A13, it is useful to define an ordering of infons that captures the degree to which they are saturated.

Definition 2.4 Infon τ is *part of* infon σ, written $\sigma \sqsubseteq \tau$, if

1. for all r, if $\mathsf{Rel}(r, \sigma)$ then $\mathsf{Rel}(r, \tau)$, and
2. for each role i and object a, if $\mathsf{Arg}(a, i, \sigma)$ then $\mathsf{Arg}(a, i, \tau)$.

An infon σ is *unsaturated* if there is another infon τ such that $\sigma \sqsubseteq \tau$. If there is no such infon, σ is said to be *saturated*.

Despite our need to have unsaturated infons, there is something to be said for the intuition that an infons *should* have a relation that determines its roles. We can recover the force of this idea by restricting Axioms A12 and A13 to apply only to saturated infons.

A 13.1 *Every saturated infon has a relation.*

A 13.2 *For saturated infons, σ and τ, having the same relation, and for each object a and role i, if $\mathsf{Arg}(a, i, \sigma)$ then for some b, $\mathsf{Arg}(b, i, \tau)$.*

In this way, we see that the standard structures are a limiting special case. Axioms A12 and A13 follow from A13.1 and A13.2 given the additional assumption that every infon is saturated.

The ordering $\sqsubseteq$ is clearly a pre-order of infons (reflexive and transitive). In information structures satisfying Axiom A9 it is a partial order (reflexive, transitive, and antisymmetric), and in standard structures it is trivial ($\sigma \sqsubseteq \tau$ iff $\sigma = \tau$) because every infon is saturated. The properties of the information ordering in non-standard information structures may be constrained by the assumptions we make concerning the existence of upper bounds.

Definition 2.5 Infons σ and τ are *compatible* if they have the same relation and for each role i and objects a and b, if $\mathsf{Arg}(a, i, \sigma)$ and $\mathsf{Arg}(b, i, \tau)$ then $a = b$. They are *unifiable* if they possess a least upper bound in the $\sqsubseteq$ ordering.[10]

For example, the information that Mary saw Elvis is compatible with the information that Elvis was seen in Tokyo. Moreover, it seems reasonable to suppose that these two infons may be unified to produce the information that Mary saw Elvis in Tokyo.[11] This suggests the following unification axiom.

A 13.3 *Every strictly compatible pair of infons is unifiable.*

Questions concerning the order of infons become a lot more difficult to answer if the arguments of infons are also partial objects. This is unavoidable if we wish to model 'higher-order' infons, as discussed in Section 2.5. There are a variety of ways of extending the order between infons by taking the order of arguments into account, but none of them is entirely convincing. Consider, for example, the idea that if $a \sqsubseteq a'$ then the infon $\langle\!\langle\, r; i\!: a, j\!: b \,\rangle\!\rangle$ is part of the infon $\langle\!\langle\, r; i\!: a', j\!: b \,\rangle\!\rangle$. This seems plausible at first, until one recalls the idea that the information expressed by a conditional is an infon of the form $\langle\!\langle\, \mathsf{if}; \sigma, \tau \,\rangle\!\rangle$. If $\sigma \neq \sigma'$ then we would *not* wish to say that $\langle\!\langle\, \mathsf{if}; \sigma, \tau \,\rangle\!\rangle$ is part of $\langle\!\langle\, \mathsf{if}; \sigma', \tau \,\rangle\!\rangle$ even if $\sigma \sqsubseteq \sigma'$. Consequently, we shall stick to the above definition, even for infons with other infons as arguments.

[10] I.e., if $\exists x$ such that $\sigma \sqsubseteq x$, $\tau \sqsubseteq x$, and $\forall y$, if $\sigma \sqsubseteq y$ and $\tau \sqsubseteq y$, then $x \sqsubseteq y$.

[11] There is no requirement that the unification of two infons is 'entailed' by the infons being unified; just that the unification of the two infons is the least infon containing them.

With discussions of unification, comparison with the literature on feature structures (cf. Rounds' chapter in this volume) is unavoidable. There is a strong similarity between feature structures and information structures. Indeed, it is easy to see how one can construct information structures using features structures to model the infons. Conversely, one can use information structures to model feature structures, by identifying features with roles and values with arguments. The order $\sqsubseteq$ defined for infons is relatively uninteresting when applied to feature structures, precisely because it does not take into account the order of arguments. Rounds (1991) discusses the relationship between the two subjects.

2.9 Representation Theorems

In addition to the axioms discussed above, we need one more axiom to account for the behaviour of constants in standard information structures.

A 14 *Every relation but no role and no infon is denoted by a constant.*

This axiom is entirely artificial, and of little consequence to the general theory.

Theorem 2.1 *A simple information structure is standard if and only if it satisfies Axioms A1 to A14.*

PROOF: We have already seen that each of these axioms is satisfied by all standard information structures. For the converse, let $\mathfrak{A}$ be an information structure satisfying the axioms, with constants C. We shall construct a relational structure $\mathfrak{M}$ and an isomorphism $f\colon |\mathfrak{A}| \rightarrow |\,\mathrm{SInf}(\mathfrak{M})|$.

Let R be the class of relations in $\mathfrak{A}$, I the class of roles, Σ the class of infons, and M the class of ordinary objects. By Axioms A3 to A5, these four classes are pairwise disjoint and exhaust the universe of $\mathfrak{A}$.

The class R of relations is finite by Axiom A8. For each r in R, let

$$I_r \quad = \quad \{i \mid \exists a, \sigma \ \mathsf{Rel}(r,\sigma) \text{ and } \mathsf{Arg}(a,i,\sigma)\}.$$

We shall show that I_r is finite. For each infon σ of $\mathfrak{A}$, let $I_\sigma = \{i | \exists a \ \mathsf{Arg}(a,i,\sigma)\}$. By Axiom A13, for all infons σ and τ such that $\mathsf{Rel}(r,\sigma)$ and $\mathsf{Rel}(r,\tau)$, $I_\sigma = I_\tau$. Thus for each σ having relation r, $I_r = I_\sigma$. Moreover, by Axiom A7, I_σ is finite for each infon σ, and so I_r is finite also.

Let $S \subseteq C$ be the class of those constants that denote relations in $\mathfrak{A}$. By Axiom A14, S is the same size as R. This enables us to define a relational signature ν with primitive relations S and such that for each R' in S, $\nu(R')$ is equal to the size of $I_{[\![R']\!]^{\mathfrak{A}}}$. Let $\mathfrak{M}$ be any relational structure with signature ν, universe M, and constants $C - S$, each with the same denotation it has in $\mathfrak{A}$; this is possible because these constants denote ordinary objects, by Axiom A14. We shall construct an isomorphism f from $\mathfrak{A}$ to $\mathrm{SInf}(\mathfrak{M})$.

First, we must enumerate the set I of roles of $\mathfrak{A}$. By Axioms A1 and A12, for each infon σ of $\mathfrak{A}$ there is a unique relation r_σ such that $\mathsf{Rel}(r_\sigma, \sigma)$. By the above, $I_\sigma = I_{r_\sigma}$. But there are only a finite number of relations (Axiom A8), and so there are only a finite number of sets of the form I_σ, even if the number of infons is infinite. Moreover, by Axiom A6, these sets are linearly ordered by inclusion. Consequently, there is an enumeration $r_1, \ldots, r_n$ of R such that $I_{r_k} \subseteq I_{r_{k+1}}$, for for $1 \leq k \leq n-1$. Now

$$I = I_{r_1} \cup \ldots \cup I_{r_n} = I_{r_1} \cup \ldots \cup (I_{r_{k+1}} - I_{r_k}) \cup \ldots \cup (I_{r_n} - I_{n-1}).$$

and so we may enumerate I by enumerating each of the sets $(I_{r_{k+1}} - I_{r_k})$ in order, and without repetition. Let $i_1, \ldots, i_N$ be such an enumeration. It has the property that for $1 \leq k \leq n$, $I_{r_k} = \{i_1, \ldots, i_{\nu(r_k)}\}$.

Next, we must consider the arguments of infons. For each infon σ, we have seen that

$$I_\sigma = I_{r_\sigma} = \{i_1, \ldots, i_{\nu(r_\sigma)}\}$$

and by Axiom A2, for $1 \leq k \leq n$, there is a unique object a_k such that $\mathsf{Arg}(a_k, i_k, \sigma)$. By Axiom A10, a_k is ordinary and so is in M. Define the function $\alpha_\sigma\colon \{1, \ldots, \nu(r_\sigma)\} \to \{0\} \times M$ by $\alpha_\sigma(k) = \langle 0, a_k \rangle$.

Now we are ready to define the function f from $|\mathfrak{A}|$ to $|\,\mathrm{SInf}(\mathfrak{M})|$.

$$f(x) = \begin{cases} \langle 0, x \rangle & \text{if } x \text{ is in } M \\ \langle 1, x \rangle & \text{if } x \text{ is in } R \\ \langle 2, k \rangle & \text{if } x \text{ is in } I \text{ and } i_k = x \\ \langle 3, \langle r_x, \alpha_x \rangle \rangle & \text{if } x \text{ is in } \Sigma \end{cases}.$$

This is a good definition, because every element of $|\mathfrak{A}|$ is in exactly one of the classes M, R, I or Σ, and for each x in I there is a unique k such that $i_k = x$. To show that f is one-one, suppose that $f(x) = f(y)$. We can show that $x = y$ by cases according to whether x and y are in M, R, I or Σ. The only non-trivial case is that in which x and y are infons. Then $f(x) = \langle 3, \langle r_x, \alpha_x \rangle \rangle$ and $f(y) = \langle 3, \langle r_y, \alpha_y \rangle \rangle$. Thus, $r_x = r_y$ and $\alpha_x = \alpha_y$, and so by Axiom A9, $x = y$, as required.

To show that f is onto the universe of $\mathrm{SInf}(\mathfrak{M})$, consider any element $\langle j, x \rangle$ of that universe. We need to find a y such that $f(y) = \langle j, x \rangle$. The only non-trivial case is that in which $j = 3$. Then $x = \langle r, \alpha \rangle$ for some relation r and function $\alpha\colon \{1, \ldots, \nu(r)\} \to \{0\} \times M$. For $1 \leq k \leq \nu(r)$, let m_k be the element of M for which $\alpha(k) = \langle 0, m_k \rangle$. Now, r is a relation, so there is some infon σ of $\mathfrak{A}$ such that $\mathsf{Rel}(r, \sigma)$. The roles of σ are those in the set $I_\sigma = \{i_1, \ldots, i_{\nu(r)}\}$ and so, applying Axiom A11 repeatedly ($\nu(r)$ times), we obtain an infon σ' such that $\mathsf{Arg}(m_k, i_k, \sigma')$ for $1 \leq k \leq \nu(r)$. Spelling out the definition of $\alpha_{\sigma'}$ we see that this function is just α, so $\langle r_{\sigma'}, \alpha_{\sigma'} \rangle = \langle r, \alpha \rangle = x$, and so $f(\sigma') = \langle 3, \langle r_{\sigma'}, \alpha_{\sigma'} \rangle \rangle = \langle 3, x \rangle$, as required.

Finally, we must show that f preserves the structure of $\mathfrak{A}$. This follows from the following two chains of equivalences. Firstly for Rel:

$$\begin{gathered}\mathsf{Rel}(f(r), f(\sigma)) \text{ in } \mathrm{SInf}(\mathfrak{M}) \\ \mathsf{Rel}(\langle 1, r\rangle, \langle 3, \langle r_\sigma, \alpha_\sigma\rangle\rangle) \text{ in } \mathrm{SInf}(\mathfrak{M}) \\ r = r_\sigma \\ \mathsf{Rel}(r, \sigma) \text{ in } \mathfrak{A}.\end{gathered}$$

And then for Arg:

$$\begin{gathered}\mathsf{Arg}(f(m), f(i), f(\sigma)) \text{ in } \mathrm{SInf}(\mathfrak{M}) \\ \mathsf{Arg}(\langle 0, m\rangle, \langle 2, k\rangle, \langle 3, \langle r_\sigma, \alpha_\sigma\rangle\rangle) \text{ in } \mathrm{SInf}(\mathfrak{M}) \text{ and } i_k = i \\ \alpha_\sigma(k) = \langle 0, m\rangle \text{ and } i_k = i \\ \mathsf{Arg}(m, i, \sigma) \text{ in } \mathfrak{A}.\end{gathered}$$

That f preserves the denotation of the constants in C follows from Axiom A14 and the definition of $\mathrm{SInf}(\mathfrak{M})$. QED

So much for standard information structures. We have argued that these structures are too rigid to provide a comprehensive framework for modelling the structure of information, even information of the simple kind we have been considering. Nonetheless, several of the axioms of standard structures lead to attractive generalizations. Specifically, Axioms A10.2 and A11.2 provide a useful direction. Both will be explored in greater depth in the next section.

3 A Theory of Structural Relations

Information processed, inferred, conveyed, expressed, or otherwise represented need not be of the simple kind considered above; it may come in larger, more complex chunks.

Logical combinations present an obvious example: we need to be able to model conjunctions, disjunctions, and perhaps negations of infons. Yet even with the first item on the list there is a potential problem. The identity condition for (well-founded) basic infons was very clear: σ and τ are identical if and only if they have the same relation, and the same arguments filling the same roles. It is much less clear what should be said about conjunctions. A strong condition, analogous to the one for basic infons, would be that $\sigma_1 \wedge \tau_1$ is identical to $\sigma_2 \wedge \tau_2$ if and only if σ_1 is identical to σ_2 and τ_1 is identical to τ_2. The problem arises if we combine this condition with natural logically-motivated requirements on conjunction, such as commutativity, associativity and idempotence. Then the condition becomes too strong. For example, consider an arbitrary conjunction $\sigma \wedge \tau$. By idempotence, $\sigma \wedge \tau = (\sigma \wedge \tau) \wedge (\sigma \wedge \tau)$, and so, by the above condition, $\sigma = \sigma \wedge \tau = \tau$.

The heart of the problem is an inherent tension in the concept of information. On the one hand, information is *representation-independent*: the same information may be represented in many different ways. On the other hand, information is *fine-grained*: two pieces of information may be logically equivalent without being identical. Consequently, the identity conditions for information represented by complex signs must lie somewhere between those for the syntactic form of the sign and those for its semantic content. Striking the right balance is, in general, very difficult.

Another problem is that there has been little agreement as to which combinations of infons are needed. Finite conjunction and disjunction are commonly adopted; infinite conjunctions and disjunction, quantified infons (with variable binding), various negations and conditionals, have also been proposed; and for applications in computer science other forms may be found useful— for example, is not clear that the information stored as a list is simply a conjunction of the items of information stored in each cell, and even if it is some kind of conjunction, it is not clear that it is the same conjunction as expressed by an unordered set of the same items.

In view of the above, it would seem sensible to pursue a fairly open-minded policy about complex infons. The approach adopted here is intended to cover various proposals made in the literature, as well as offering a framework in which other approaches may be tried.

3.1 Extensional Structures

Our idea is to generalize the notion of bisimulation to apply in a wider context. In the sequel, we will have (relational) structures of type

$$[A, S_1, \ldots, S_m; R_1, \ldots, R_n]$$

The relations in the first group, S_1, ..., S_m, are called *structural relations* because they capture the structure of elements of the domain. A structural relation S_i of arity $n+1$ is to be thought of as relating a list of n objects to a single structured object. If $S_i(x_1, \ldots, x_n, y)$ then y is a structured object with components $x_1, \ldots, x_n$, which may or may not be structured themselves.

More generally, we say that b is a *component* of a in $\mathfrak{A}$ if there is a structural relation S_i of arity $n+1$ and elements $x_1, \ldots, x_n$ of A such that $S_i(x_1, \ldots, x_n, a)$ and $x_j = b$ for some $1 \leq j \leq n$. For technical reasons, we require that the number of components of any is not a proper class—that is to say, the class of all components of a given object can be placed in one-to-one correspondence with some set. An object a is an *atom* of $\mathfrak{A}$ if it has no components.

In an information structure, the only structured objects are the infons. Relations, roles and ordinary objects are all atomic, but infons have a component structure captured by the relations Rel and Arg. These are the structural relations of information structures. Information structures have no other primitive relations, but the defined relation $\sqsubseteq$ is an example of a non-structural relation—albeit one whose extension is determined by the structure of infons. The important distinction between Rel and Arg, on the one hand, and $\sqsubseteq$, on the other, is that the identity of infons is determined by the former, by virtue of adherence to Axiom A10.2.

But now consider an arbitrary relational structure. What conditions must a relation satisfy to qualify as a structural relation? Our answer is based on the following definition.

Definition 3.1 Given a relational structure $\mathfrak{A}$ of type

$$[A, S_1, \ldots, S_m; R_1, \ldots, R_n],$$

a binary relation E on A is said to be a *bisimulation* on $\mathfrak{A}$ if for all $a, b \in A$, if $E(a, b)$ then the following three conditions hold:

1. if a is atomic then $a = b$,
2. for $1 \leq i \leq m$, if $\nu_i = k$, then for all $y_1, \ldots, y_k$ such that $S_i(y_1, \ldots, y_k, a)$, there are $z_1, \ldots, z_k$ such that $S_i(z_1, \ldots, z_k, b)$, and $E(y_j, z_j)$ for $1 \leq j \leq k$, and
3. $E(b, a)$

a is *bisimilar* to b in $\mathfrak{A}$ iff there is a bisimulation E of $\mathfrak{A}$ such that $E(a, b)$. The structure $\mathfrak{A}$ is *extensional* if it has no distinct, bisimilar objects, i.e. if a is bisimilar to b in $\mathfrak{A}$ then $a = b$.

If two objects are bisimilar then they are structurally equivalent, so there must be something else that distinguishes them. But if they are structural objects, then there can be nothing else to distinguish them. In this way, the extensionality condition enforces structural conditions of identity on the non-atomic objects of the domain, with the structure of an object determined solely by the relations structural relations. The non-structural relations, $R_1, \ldots, R_n$, on the other hand, are not constrained in any way. They are used to model additional facts about the objects of the domain, which need not be dependent on their structure alone.

Extensional structures will be used throughout the rest of this chapter to model a variety of situation-theoretic objects. Our strategy will be to define different kinds of extensional structures to model different parts of the situation-theoretic universe, but to do so in a modular fashion. The modularity is facilitated by the notion of a 'sort'.

Definition 3.2 If $\mathfrak{A}$ is an extensional structure, then for each structural relation S of $\mathfrak{A}$, let S^* be the class of those objects a such that $S(\vec{x}, a)$ in $\mathfrak{A}$ for some sequence $\vec{x}$ (possibly empty) of elements of $|\mathfrak{A}|$. In other words, S^* is the projection of S along its last co-ordinate. We call these classes the *(structural) sorts* of $\mathfrak{A}$.

Every generalized information structure is an extensional structure with structural relations Rel and Arg. This is given by Axiom A10.2. The standard structures and even the iterated structures are also well-founded, by Axiom A10.1. In a standard information structure $\mathsf{Rel}^* = \mathsf{Arg}^*$ is the class of infons, the only structured objects. In generalized information structures, Rel^* and Arg^* may be slightly different if there are infons with a relation but no arguments or with arguments but no relation. In any case, the class of infons is $\mathsf{Rel}^* \cup \mathsf{Arg}^*$. The relations, roles and ordinary objects are all atoms.

Another, familiar example of an extensional structure is the structure $\mathfrak{V} = \langle V, \in^2; \rangle$ where V is the class of all sets, and $\in$ is the binary relation of membership. The class $\in^*$ consists of all non-empty sets. Extensionality is assured by the set-theoretic axiom of the same name, together with the Axiom of Foundation or Anti-foundation, depending on whether V is assumed to satisfy the axioms of well-founded or anti-founded set theory.

Definition 3.3 An extensional structure $\mathfrak{A}$ of of type $[A, \mathsf{Set}^1, \in^2;]$ is called a *set structure* if $\in^* \subseteq \mathsf{Set}^*$ and $\mathfrak{A}$ satisfy the axioms of ZFC$^-$ with quantifiers restricted to Set^*.

The ambiguity in our use of the word 'set' and the predicate '$\in$' is of the now familiar kind. Internal sets are those elements of the structure in Set^*, which may or may not be (external) sets. But for any internal set a, we may define a corresponding external set $a^* = \{b \mid \in (b, a)\}$.

Our third example of an extensional structure presents functions as structured objects whose components are their arguments and values.

Definition 3.4 An extensional structure $\mathfrak{A}$ of type $[A, \mathsf{App}^3, \mathsf{Fun}^1;]$ is a *function structure* if $\mathsf{App}^* \subseteq \mathsf{Fun}^*$ and for all α, x, y and i in A, if $\mathsf{App}(i, x, \alpha)$ and $\mathsf{App}(i, y, \alpha)$, then $x = y$.

If $\alpha \in \mathsf{App}^*$ then α is an (internal) function of $\mathfrak{A}$. We associate it with an external function α^* whose domain consists of those elements i of A for which there is an x such that $\mathsf{App}(i, x, \alpha)$. For each such i, $\alpha^*(i) = x$, which is uniquely specified, by the defining condition of function structures.

What is the import of the extensionality requirement on function structures? It amounts to the condition that different internal functions represent different external functions. For if $a^* = b^*$, then $I_A \cup \{\langle a, b\rangle, \langle b, a\rangle\}$ would be a bisimulation relating a and b. If the structure is *flat* (i.e., if no function is also an argument), then extensionality is equivalent to this condition. In other structures, extensionality is a stronger condition.

3.2 Structural Descriptions and Anti-Foundedness

The extensionality condition ensures that our structures are not over-populated. In this section we guard against the danger of under-population, by showing how extensional structures can be populated to capacity.

Every object in an extensional structure can be identified by listing its structural relationships with other objects. Given an extensional structure $\mathfrak{A}$ and a 'parameter' x, which may be any object not necessarily an element of $|\mathfrak{A}|$, a *unary structural description* is a pair $\langle x, T_x \rangle$ consisting of a set T_x of tuples of the form $\langle S, b_1, \ldots . b_k \rangle$ where S is a structural relation of arity $k+1$ and each of $b_1, \ldots, b_k$ is either the parameter x or an object of $\mathfrak{A}$. For each a in $|\mathfrak{A}|$, we form the set $a.T$ of tuples of the form $\langle S, b_1[x/a], \ldots . b_k[x/a] \rangle$ for each $\langle S, b_1, \ldots . b_k \rangle$ in T_x, where $b_i[x/a] = a$ if $b_i = x$ and is b_i otherwise. The object a *satisfies* the description T if for all S and $b_1, \ldots b_k$,

$$S(b_1, \ldots b_k, a) \text{ in } \mathfrak{A} \quad \text{iff} \quad \langle S, b_1, \ldots b_k \rangle \text{ is in } a.T$$

For example, in an information structure the infon $\sigma = \langle\!\langle\ r; i : \sigma, j : a\ \rangle\!\rangle$ satisfies the description $T = \langle x, \{\langle \mathsf{Rel}, r \rangle, \langle \mathsf{Arg}, i, x \rangle, \langle \mathsf{Arg}, j, a \rangle\}\rangle$ because $\sigma.T = \{\langle \mathsf{Rel}, r \rangle, \langle \mathsf{Arg}, i, \sigma \rangle, \langle \mathsf{Arg}, j, a \rangle\}$ and $\mathsf{Rel}(r, \sigma)$, $\mathsf{Arg}(i, \sigma, \sigma)$, and $\mathsf{Arg}(j, a, \sigma)$; and these are all the structural relationships in which σ participates.

Every object satisfies some description. An object a of $\mathfrak{A}$ satisfies the description $\langle a, T_a \rangle$ where T_a is the set of tuples of the form $\langle S, b_1, \ldots b_k \rangle$ such that $S(b_1, \ldots b_k, a)$ in $\mathfrak{A}$.[12] By extensionality, no two objects satisfy the same description. So the description of an object serves as its unique identifier. In most cases, however, there many more descriptions than there are objects described. The problem of under-population is that of there being two few objects to satisfy the descriptions that Situation Theory requires. Before we show how it can be solved, the definition of 'structural description' must be extended to cover the polyadic case.

Definition 3.5 Suppose $\mathfrak{A}$ is an extensional structure. A *structural description* T in $\mathfrak{A}$ is an indexed family $\{T_x\}_{x \in X}$ of sets T_x of tuples of the form $\langle S, b_1, \ldots . b_k \rangle$ such that S is a structural relation of $\mathfrak{A}$ of arity $k+1$ and $b_1, \ldots, b_k$ are elements of $|\mathfrak{A}| \cup X$. The elements of the index class X are called *parameters*. Given a mapping s from X to $|\mathfrak{A}|$ let $s.T$ be the family of sets $s.T_x$ of tuples of the form $\langle S, b_1[s], \ldots . b_k[s] \rangle$ for each $\langle S, b_1, \ldots . b_k \rangle$ in T_x where $b_i[s] = s(b_i)$ if b_i is in X and is b_i otherwise. The mapping s *satisfies* T if for each x in X, $s(x)$ satisfies the unary description $\langle s(x), s.T_x \rangle$.

To see how this definition works, consider its application to set structures. A set a satisfies the unary description $T = \langle x, T_x \rangle$ just in case

$$T_a = \{\langle \in, b \rangle \mid b \text{ in } a^*\} \cup \{\langle \mathsf{Set} \rangle\}.$$

Thus if we define $T^*_a = \{b \mid \langle \in, b \rangle \text{ in } T_a\}$, then a satisfies T iff $a^* = T^*_a$. Likewise, if T is a polyadic structural description, then a mapping s satisfies T iff $s(x)^* = (s.T_x)^*$ for each parameter x of T. This relationship enables us to represent a structural description in a set structure as a system of simultaneous equations. For example, consider the system of equations shown on the left below

[12] T_a cannot be a proper class because we have assumed that every object has only a set-sized class of components.

$$\begin{array}{rcl|rcl} x & = & \{y, z\} & T_x & = & \{\langle \in, y\rangle, \langle \in, z\rangle, \langle \mathsf{Set}\rangle\} \\ y & = & \emptyset & T_y & = & \{\langle \mathsf{Set}\rangle\} \\ z & = & \{x, a\} & T_z & = & \{\langle \in, x\rangle, \langle \in, a\rangle, \langle \mathsf{Set}\rangle\} \end{array}$$

corresponds to the structural description shown on the right. Moreover, a mapping s that satisfies the description on the right determines (and is determined by) a solution to the equations. It gives us sets $s(x)$, $s(y)$, and $s(z)$, such that

$$\begin{array}{rcl} s(x) & = & \{s(y), s(z)\} \\ s(y) & = & \emptyset \\ s(z) & = & \{s(x), a\} \end{array}$$

Thus, in set structures the population density is determined by the number of solutions to systems of equations.[13] The usual models of Set Theory are rather sparse by this standard. The above system, for example, cannot be solved because any solution would conflict with the Foundation Axiom. If s is a solution, then there is an infinite descending sequence of elements:

$$s(x) \ni s(z) \ni \cdots \ni s(x) \ni s(z) \ni \cdots$$

Fortunately, Aczel (1988) showed that if the Foundation Axiom is dropped from Zermelo-Frankel Set Theory (ZFC), the resulting theory (ZFC$^-$) is consistent with the claim that *every* system of equations has a unique solution. That claim—called the Solution Lemma in Barwise and Etchemendy (1987)—is equivalent to the Anti-Foundation Axiom (AFA) used by Aczel.[14] Aczel proposes ZFC$^-$ plus AFA as a new theory of sets, perhaps more suitable than

[13] Structural descriptions correspond to a wider range of systems of equations than one might expect. Consider $X = \{x, y\}$ and

$$\begin{array}{rcl} e(x) & = & \{\langle x, y\rangle\} \\ e(y) & = & \{x, y, \emptyset\} \end{array}$$

This system is unlike the ones we've seen so far because $e(x)$ contains the *ordered pair* $\langle x, y\rangle$. In set theory, this pair is standardly taken to be $\{\{x\}, \{x, y\}\}$. So when to solve a system involving pairs, we must add more variables to X and more equations. Here we would take $X' = \{x, y, z_0, z_1, z_2\}$ and

$$\begin{array}{rcl} e(x) & = & \{z_0\} \\ e(y) & = & \{x, y, \emptyset\} \\ e(z_0) & = & \{z_1, z_2\} \\ e(z_1) & = & \{x\} \\ e(z_2) & = & \{x, y\} \end{array}$$

A solution to the expanded system gives us a solution to the original system. This would be a map s defined on X with the property that $s(x) = \{\langle s(x), s(y)\rangle\}$ and $s(y) = \{s(x), s(y), \emptyset\}$.

[14] The Anti-Foundation Axiom is defined using 'decorations' of a pointed graph, rather than solutions of equations or the satisfiers of structural descriptions. Another approach is to look at the fixed points of monotone operators on an upper semi-lattice. We advise the interested reader to consult Aczel (1988) and Barwise and Moss (1996) for details.

ZFC when it comes to modelling the various circular phenomena encountered in applied logic.

Returning to the general case, we can formulate the two extremes of population density by generalizing Foundation and Anti-Foundation to apply to arbitrary extensional structures.

Definition 3.6 An extensional structure $\mathfrak{A}$ is *well-founded* if there is no infinite sequence $a_1, a_2, \ldots$ of elements of A such that a_{n+1} is a component of a_n for each positive integer n. It is *anti-founded* if every structural description in $\mathfrak{A}$ is satisfied.

In particular, a set structure is well-founded if it satisfies the Foundation Axiom of ZFC and is anti-founded if it satisfies the Anti-Foundation Axiom. The well-founded extensional structures are those with the lowest population density: only those objects explicitly requested by other parts of the theory are required to be present. Anti-founded structures, by contrast, are full to the brim: any object you can describe is sure to be there. The next theorem shows that every extensional structure lies between these two such extremes.

Theorem 3.1 *For every extensional structure $\mathfrak{A}$ there is a well-founded extensional structure $\mathfrak{A}_{\mathrm{wf}}$ and an anti-founded extensional structure $\mathfrak{A}_{\mathrm{af}}$ such that $\mathfrak{A}_{\mathrm{wf}}$ is isomorphically embedded in a $\mathfrak{A}$, and $\mathfrak{A}$ is isomorphically embedded $\mathfrak{A}_{\mathrm{af}}$.*

PROOF: The existence of a well-founded part of $\mathfrak{A}$ is easy to show. We can just restrict the structure to those objects that are not at the start of an infinitely descending chain of components, and check that the extensionality condition still holds. For the anti-founded extension of $\mathfrak{A}$, we shall assume that we are working in a universe of sets that satisfies ZFC$^-$+AFA, and that there is a proper class of ur-elements (non-sets) available. (Given the relative consistency of AFA with ZFC$^-$, this is a legitimate strategy.)

With each structural relation S of $\mathfrak{A}$, we associate an ur-element z_s. We also associate an ur-element z_a with each atom of $\mathfrak{A}$. Let U be the collection of all these ur-elements. We assume that the association is one-to-one. For any set a, the 'support' of a the set of objects b_0 for which there is a sequence $b_0 \in \ldots \in b_n$ with $b_n = a$. We assume that no ur-elements in U occur in the support of $|\mathfrak{A}|$.

From Aczel (1988) we know that every system of equations in a model of ZFC$^-$+AFA has a unique solution. By a 'system of equations' we mean a function e whose domain is a set X of ur-elements and which maps each x in X to a set whose members are either elements of X or do not contain elements of X in their supports. A 'solution' to e is a function s mapping each x in X to a set $s(x)$ such that $s(x)$ is the set $\{y[s] \mid y \in e(x)\}$ where $y[s] = s(y)$ if y is in X and is y otherwise.

Now, with each structural description T in $\mathfrak{A}$ with parameter set X, we associate the function e_T mapping each parameter x to the set

$$\{\langle z_s, b_1, \ldots, b_k\rangle \mid \langle S, b_1, \ldots, b_k\rangle \in T_x\}.$$

In other words, each structural relation S replaced by the corresponding z_s. As it stands, e is not a system of equations in the strict sense defined above because the elements of $e_T(x)$ are tuples of parameters and sets, rather than the parameters and sets themselves. But this is easily converted into a system of equations e'_T with parameter set X' (containing X) by the method illustrated earlier, such that a solution for e' is a solution for e and vice versa.

By AFA, the system e'_T has a unique solution, s_T. Let

$$B \quad = \quad \{s_T(x) \mid T \text{ is a structural description in } \mathfrak{A} \text{ and } x \text{ is in } X'\}.$$

We'll see how to interpret the relations of $\mathfrak{A}$ on this set to get the desired structure $\mathfrak{A}_{\mathrm{af}}$. But first, we need a subsidiary fact about structural descriptions.

Recall that each $a \in A$ satisfies a description T_a in one variable x. Let $i : A \to B$ be given by $i(a) = s_{T_a}(x)$ for a a non-atom of $\mathfrak{A}$, and $i(a) = z_a$ if a is an atom. Then is one-to-one: if $s_{T_a}(x) = s_{T_b}(x)$, then we could get a bisimulation relation on $\mathfrak{A}$ relating a and b. Since $\mathfrak{A}$ is extensional, we would have $a = b$. And if a is a non-atom, then $i(a)$ is a set, so it cannot equal $i(b) = z_b$ for b a non-atom

Now we can see how to make B into a structure $\mathfrak{A}_{\mathrm{af}}$ of the same type as $\mathfrak{A}$. Let S be a structural relation. We interpret it by

$$S(c_1, \ldots, c_n, c) \text{ in } \mathfrak{A}_{\mathrm{af}} \text{ iff } \langle z_S, c_1, \ldots, c_n\rangle \in c.$$

The definition of the principal descriptions T_a implies that

$$S(b_1, \ldots, b_k, a) \text{ in } \mathfrak{A} \text{ iff } S(i(b_1), \ldots, i(b_k), i(a)) \text{ in } \mathfrak{A}_{\mathrm{af}}.$$

The non-structural relations extend to the image of $\mathfrak{A}$ via the isomorphism; we'll say that on tuples outside the image of i, each non-structural relation is false.

At this point we have $\mathfrak{A}_{\mathrm{af}}$. This structure is extensional by AFA. A bisimulation of structures relating c and c' would give us a bisimulation of *sets* doing the same thing. And according to AFA, bisimilar sets are equal.

Finally, we sketch a proof that $\mathfrak{A}_{\mathrm{af}}$ is anti-founded. Before we do this, we need to recall that for each $c \in B$ there is a structural description T_c over the original $\mathfrak{A}$ and a parameter x of T such that $s_T(x) = c$. Let T' be a structural description in $\mathfrak{A}_{\mathrm{af}}$. Now each T_y may contain elements $c \in |\mathfrak{A}_{\mathrm{af}}|$, so we replace each such c by a new ur-element (say z_c) and then adjoin all of the structural descriptions T_c to T. Of course, we must use new parameters to do this. In this way, T' is equivalent to a structural description in the original $\mathfrak{A}$. So by definition of $\mathfrak{A}_{\mathrm{af}}$, there is some $a \in |\mathfrak{A}_{\mathrm{af}}|$ instantiating T'. QED

As a corollary, we shall apply this result to simple information structures. Suppose $\mathfrak{A}$ is a simple information structure with a set R of relations, a set I of roles, and a set O of ordinary objects. A structural description of an infon in $\mathfrak{A}$ with parameter set X is determined by a system of equations of the form $e(x) = \langle\!\langle\, r; i_1 : b_1, \ldots, i_n : b_n \,\rangle\!\rangle$ where $r \in R$, $i_j \in I$ and $b_i \in O \cup X$ for $1 \leq i \leq n$.

Theorem 3.1 tells us that there is an anti-founded information structure $\mathfrak{A}_{\mathrm{af}}$ extending $\mathfrak{A}$. In this structure all structural descriptions are satisfied and so there are solutions to the equations. In particular, all the hyperinfons we care to define using relations R, roles I, and ordinary objects O can be found in this structure.

3.3 Partiality and order

We have seen how infons may be ordered according to their arguments and relations, enabling us to model the partiality of information. In fact, the definition of $\sqsubseteq$ is quite general.

Definition 3.7 Given an extensional structure $\mathfrak{A}$ of type

$$[A, S_1, \ldots, S_m; R_1, \ldots, R_n],$$

we define a relation $\sqsubseteq$ on the non-atomic elements of A as follows: $a \sqsubseteq b$ if for $1 \leq i \leq m$, and each sequence $\vec{x}$ of appropriate length, if $S_i(\vec{x}, a)$ then $S_i(\vec{x}, b)$.

For information structures, the definition of $\sqsubseteq$ agrees with our previous definition. For set structures, $a \sqsubseteq b$ iff $a^* \subseteq b^*$, and for function structures, it is the usual order on partial functions: $f \sqsubseteq g$ iff each i in the domain of f^* belongs to the domain of g^* and $f^*(i) = g^*(i)$. The relation $\sqsubseteq$ is a partial order on structured (non-atomic) objects.[15]

Other ordering of structured objects are possible and have been studied extensively. The most important of these is the 'hereditary' order, which takes into account the order of components. Roughly, an object a is a hereditary part of b if for every component a' of a there is a component b' of b such that either $a' = b'$ or a' is a hereditary component of b'. Stated thus, the definition is circular. In some cases we can overcome the problem using an inductive layering of the domain. Say that an object is of order 0 if it is atomic, and of order $n+1$ if it has a component of order n but no components of order greater that n. This enables us to give a recursive definition of the hereditary order on objects of order n in terms of the hereditary order of objects of order less than n.

This strategy works for well-founded structures, in which every object has an order. But in structures with a 'circular' component relation there are objects that do not have an order. This is a ubiquitous problem in the study of circular phenomena. The solution is as follows. First, we regard the 'definition' of the order as an operation on binary relations. Given one binary relation R, we let R^+ be defined by

> $R^+(a, b)$ if for every component a' of a there is a component b' of b such that either $a' = b'$ or $R(a', b')$.

[15] It is clearly reflexive and transitive, and anti-symmetry follows from extensionality: if $a \sqsubseteq b$ and $b \sqsubseteq a$ then the union of $\{\langle a, b\rangle, \langle b, a\rangle\}$ and the identity relation is a bisimulation of $\mathfrak{A}$, and so $a = b$.

Now, the operation taking R to R^+ is *monotone*: if $R \subseteq S$ then $R^+ \subseteq S^+$. A basic result of the theory of order is that every monotone operation has a *greatest fixed point*. In other words, there is an R such that $R^+ = R$ and for any S such that $S^+ = S$, $S \subseteq R$. This observation enables to define the hereditary order in the general case as follows.

Definition 3.8 The *hereditary order*, $\sqsubseteq^h$, on $\mathfrak{A} : [A, S_1, \ldots, S_m; R_1, \ldots, R_n]$ is the largest binary relation between non-atomic objects of $\mathfrak{A}$ such that if $a \sqsubseteq^h b$ then for $1 \leq i \leq m$, if $S_i(x_1, \ldots, x_k, a)$ then there are $y_1, \ldots, y_k$ such that $S_i(y_1, \ldots, y_k, b)$ and $x_j \sqsubseteq^h y_j$ for each $1 \leq j \leq k$.

The hereditary order is also a partial order on non-atomic objects for the same reason that $\sqsubseteq$ is. Other orders may be obtained by defining different monotone operations.

It is important to realize that both $\sqsubseteq$ and $\sqsubseteq^h$ are purely structural concepts. In the case of simple infons, if $\sigma \sqsubseteq \tau$ then it is possible to go further and interpret the order in terms of epistemological advantage: knowing σ is not as valuable as knowing τ. This gloss is not available in the general case. For example, in Section 3.4 we shall show how to obtain a disjunction $\sigma_1 \vee \sigma_2$ whose components are σ_1 and σ_2, and such that $\sigma_1 \vee \sigma_2 \sqsubseteq \sigma_1 \vee \sigma_2 \vee \sigma_3$. The question of how to relate the partiality of structure to logical operations is a difficult one; it is shared by those studying databases and feature logic (Rounds, 1991, and Chapter 8, this volume).

The discussion of unification in information structures may also be generalized, although no uniform definition of compatibility is available. Instead, we may introduce the following abstract version.

Definition 3.9 A relational structure

$$\mathfrak{A} : [A, S_1, \ldots, S_m; R_1, \ldots, R_n, C^2, P^2]$$

is a *unification structure* with *order* P and *compatibility relation* C if

1. $\mathfrak{A}$ is extensional,
2. C is symmetric and reflexive on non-atoms,
3. P is a partial order of non-atoms, and
4. If $C(a, b)$ then a and b have a least upper bound in P.

For example, taking $C(\sigma, \tau)$ iff σ and τ are compatible simple infons, and $P(\sigma, \tau)$ iff $\sigma \sqsubseteq \tau$, we see that generalized information structures are unification structures; for set structures, take P to be $\sqsubseteq$ and C to be the universal relation on Set^*—the unification of any two sets is just their union; and for function structures with P taken to be the $\sqsubseteq$ order, let $C(f, g)$ iff there is no i in the domains of both f^* and g^* and such that $f^*(i) \neq g^*(i)$. The hereditary orders may also be used to define unification structures with weaker, hereditary compatibility relations.

3.4 Complex Infons

We have seen that infons, sets and functions may all be modelled using extensional structures. In each of these cases, the identity conditions are relatively straightforward. Now we shall show how structured objects with more complicated identity conditions can also be modelled.

Suppose we wish to have a binary operation of conjunction with the following properties:

> (commutativity) $\sigma \wedge \tau = \tau \wedge \sigma$,
>
> (associativity) $\sigma \wedge (\tau_1 \wedge \tau_2) = (\sigma \wedge \tau_1) \wedge \tau_2)$, and
>
> (idempotence) $\sigma \wedge \sigma = \sigma$.

We shall suppose that we already have a stock of infons Σ in some information structure $\mathfrak{A}$ from which to form conjunctions. Our task is to extend $\mathfrak{A}$ with new elements and relations designed to capture the structure of the conjunctions.

The solution is to define a new binary structural relation $\mathsf{ConjunctOf}$ that holds between conjuncts and conjunction. Using this relation we can define the operation of conjunction as follows:

> given infons σ and τ, the conjunction $\sigma \wedge \tau$ is the least x in the $\sqsubseteq$ order such that $\mathsf{ConjunctOf}(\sigma, x)$ and $\mathsf{ConjunctOf}(\tau, x)$.

In other words, $\sigma \wedge \tau$ is the structurally smallest conjunction containing σ and τ as conjuncts. Clearly any such conjunction is commutative. Additional properties of $\wedge$ may be imposed by careful construction of $\mathsf{ConjunctOf}$. By way of example, we show that it is possible to extend an information structure with a conjunction that is associative and idempotent, and another that has neither of these properties.

Construction 3.1 Let $\mathfrak{A}$ be an extensional structure with a class Σ of non-atomic objects, which we shall call 'infons'. We shall define an extensional structure $\mathfrak{B}$ extending $\mathfrak{A}$ with new elements and a new binary structural relation $\mathsf{ConjunctOf}$ such that the conjunction $\wedge$ defined, as above, on objects in Σ is both associative and idempotent. Let

$$\begin{array}{rcl} B_0 & = & 0 \times (|\mathfrak{A}| - \Sigma) \\ B_1 & = & 1 \times \mathrm{pow}(\Sigma) \end{array}$$

and for each $a \in |\mathfrak{A}|$ let

$$a^* \quad = \quad \begin{cases} \langle 0, a \rangle & \text{if } a \notin \Sigma \\ \langle 1, \{a\} \rangle & \text{if } a \in \Sigma \end{cases} .$$

Now let $\mathfrak{B}$ be the structure with universe $B_0 \cup B_1$, and with

$$R(a_1^*, \ldots, a_n^*) \text{ in } \mathfrak{B} \text{ iff } R(a_1, \ldots, a_n) \text{ in } \mathfrak{A}$$

for each primitive relation R of $\mathfrak{A}$ and

$$\mathsf{ConjunctOf}(\langle 1, x\rangle, \langle 1, y\rangle) \text{ iff } x \subseteq y.$$

The new relation ConjunctOf is treated as a structural relation of $\mathfrak{B}$ so we must check that $\mathfrak{B}$ is extensional.[16] The binary conjunction $\wedge$ is defined, as above, for all 'infons', which in this case refers to those objects in B_1. Now for any $x, y \subseteq \Sigma$, it is easily checked that $\langle 1, x\rangle \wedge \langle 1, y\rangle = \langle 1, x \cup y\rangle$. Thus conjunction is both idempotent and associative in $\mathfrak{B}$.

Construction 3.2 For the second construction, let $\mathfrak{A}$ and Σ be as before. Let $\Sigma_1 = \{1\} \times \mathrm{pow}(\{0\} \times \Sigma)$, $\Sigma_{(n+1)} = \Sigma_n \cup \{1\} \times \mathrm{pow}\, \Sigma_n$, and $\Sigma^* = \bigcup_{n=1}^{\omega} \Sigma_n$. Now let $\mathfrak{B}$ be the structure with universe $\Sigma^* \cup \{0\} \times |\mathfrak{A}|$, and with

$$R(\langle 0, a_1\rangle, \ldots, \langle 0, a_n\rangle) \text{ in } \mathfrak{B} \text{ iff } R(a_1, \ldots, a_n), \text{ in } \mathfrak{A}$$

for each primitive relation R of $\mathfrak{A}$, and

$$\mathsf{ConjunctOf}(x, \langle 1, y\rangle) \text{ iff } x \in y.$$

Again we must check that $\mathfrak{B}$ is extensional.[17] The 'infons' of $\mathfrak{B}$ are those objects of non-zero rank together with those objects of the form $\langle 0, x\rangle$ for which $x \in \Sigma$. It is easy to check that $a \wedge b = \langle 1, \{a, b\}\rangle$ and so conjunction is neither idempotent nor associative.

[16] If a is bisimilar to b then $a = \langle i, a_0\rangle$, $b = \langle j, b_0\rangle$ for some a_0 and b_0. First observe that $i = j$. For if $i \neq j$, say $i = 0$ and $j = 1$, then b_0 is a subset of Σ, and so $\mathsf{ConjunctOf}(\langle 1, \emptyset\rangle, b)$ whereas there is no x such that $\mathsf{ConjunctOf}(x, a)$, and this contradicts the assumption that a is bisimilar to b. In the case that $i = j = 0$, a_0 is bisimilar to b_0 in $\mathfrak{A}$, so $a_0 = b_0$ by the extensionality of $\mathfrak{A}$, and so $a = b$. And if $i = j = 1$ then for every $x \in a_0$, there is a $b_1 \subseteq b_0$ such that $\langle 1, \{x\}\rangle$ is bisimilar to $\langle 1, b_1\rangle$. Then for each $y \in b_1$, $\langle 1, \{y\}\rangle$ is bisimilar to $\langle 1, \{x\}\rangle$—it cannot be bisimilar to $\langle 1, \emptyset\rangle$, the only other conjunct of $\langle 1, \{x\}\rangle$. But then either x and y are atoms of $\mathfrak{A}$, which is ruled out because $x, y \in \Sigma$, or x is bisimilar to y in $\mathfrak{A}$ and so $x = y$ by the extensionality of $\mathfrak{A}$. Thus $b_1 = \{x\}$ and so $a_0 \subseteq b_0$. By a symmetrical argument, it can be shown that $b_0 \subseteq a_0$, so $a_0 = b_0$ and so $a = b$.

[17] The proof requires the notion of the 'rank' of an element of $|\mathfrak{B}|$. An element $b \in |\mathfrak{B}|$ is of rank 0 if $b = \langle 0, a\rangle$ for some $a \in |\mathfrak{A}|$ and is of rank $n > 0$ if $b \in \Sigma_n$ but $b \notin \Sigma_m$ for each $m < n$. We show that if a is bisimilar to b then $a = b$, by induction on the maximum of the ranks of a and b. Suppose that a is bisimilar to b. There are three cases: (i) both x and y have rank 0; (ii) one of x or y has rank 0 and the other has rank greater than 0; and (iii) both x and y have ranks greater than 0. Case (i) is the base of the induction. In this case, $a = \langle 0, x\rangle$ and $b = \langle 0, y\rangle$ for some x and y, so x is bisimilar to y in $\mathfrak{A}$, so $x = y$ by the extensionality of $\mathfrak{A}$, and so $a = b$. Case (ii) is impossible. Suppose, for example, that $a = \langle 0, x\rangle$ and $b = \langle 1, y\rangle$. Then either y is empty, making b an atom of $\mathfrak{B}$ and so only bisimilar to itself, or y has an element y' and so $\mathsf{ConjunctOf}(y', b)$. There is no x' such that $\mathsf{ConjunctOf}(x', a)$, and so (ii) contradicts the assumption that a is bisimilar to b. Case (iii) is the inductive case. We may suppose that $a = \langle 1, x\rangle$ and $b = \langle 1, y\rangle$ for some x and y. Then for each $x' \in x$ there is a $y' \in y$ such that x' is bisimilar to y'. But the ranks of x' and y' are strictly less than those of x and y respectively, and so by the inductive hypothesis, $x' = y'$. Thus $x \subseteq y$, and by a similar argument $y \subseteq x$, so $x = y$, and so $a = b$.

Both constructions give a commutative conjunction. In fact, commutativity follows from the definition of $\wedge$ from $\mathsf{ConjunctOf}$. To get a non-commutative conjunction, we would need use different structural relations. For a non-associative, non-commutative conjunction we could use two binary structural relations, say, $\mathsf{ConjunctOf}_1$ and $\mathsf{ConjunctOf}_2$, with $\wedge$ defined such that $\mathsf{ConjunctOf}_1(\sigma, \sigma \wedge \tau)$ and $\mathsf{ConjunctOf}_2(\tau, \sigma \wedge \tau)$. For an associative, non-commutative conjunction, we could use a ternary structural relation $\mathsf{ConjunctOf}_*$, with integer roles, and $\wedge$ defined so that

$$\mathsf{ConjunctOf}_*(i, \sigma_i, (\sigma_1 \wedge \ldots \wedge \sigma_n)).$$

The goal in defining a structural notion of conjunction is to capture those properties of conjunctive information that are immediately available to anyone receiving the information, without further reflection on the significance of the information itself. For this reason, it may be that conjunctive information represented in different forms may be best modelled by different conjunctions. For example, information about the relative heights of Indiana's Hoosiers may be conveyed in a list of sentences of the form 'x is taller than y', or by a team photograph. In the former case, the information conveyed may be best represented by a non-commutative, associative conjunction; in the latter case, a commutative, associative conjunction would be preferable.

Similar issues affect the choice of structural relations for modelling other complex infons, such as disjunctions and conditionals. Negations may be modelled as complex infons, or as basic infons, as indicated in Section 2.4. For disjunctions, constructions almost identical to Constructions 3.1 and 3.2 may be used. We simply re-name '$\mathsf{ConjunctOf}$' as '$\mathsf{DisjunctOf}$' and '$\wedge$' as '$\vee$' throughout. This further underlines the point that the issues pertaining to the modelling of complex infons are structural, not logical. Definitions of several kinds of structures having complex infons are given in Section 3.6, below.

The question of how to represent quantifiers is even more vexed (Robin Cooper (1991) and Richard Cooper (1991)). One approach is to use infinitary conjunctions and disjunctions. In both structures constructed above we may define an infinitary conjunction as follows:

> given a set X of infons, the conjunction $\bigwedge X$ is the least x in the $\sqsubseteq$ order such that $\mathsf{ConjunctOf}(\sigma, x)$ for each $\sigma \in X$.

This operation has similar properties to the corresponding binary conjunction. In the structure defined in Construction 3.1 it is both associative an idempotent, but in the structure defined in Construction 3.2 it is neither. Infinitary conjunctions are always commutative.

Another approach is to model a quantified infon as a pair $\langle Q, \lambda x.\sigma \rangle$, in which Q is a quantifier and $\lambda x.\sigma$ is a 'infon-abstract'. We shall discuss abstracts in Section 3.7.

3.5 Substitution

In any extensional structure there is a well-defined operation of substitution. The following definition is influenced by the treatment of substitution in Aczel (1990) and Aczel and Lunnon (1991).

Definition 3.10 Let $\mathfrak{A} : [A, S_1, \ldots, S_m; R_1, \ldots, R_n]$. A function f is a *substitution* in $\mathfrak{A}$ if its arguments and values are all elements of A. Given a substitution f, a binary relation E on A is an *f-simulation* if for all $a, b \in A$, if $E(a, b)$ then $b = f(a)$ if a is in the domain of f, and if a is not in the domain of f then

1. if a is atomic then $a = b$
2. for $1 \leq i \leq m$, if S_i is of arity k then for all $x_1, \ldots, x_k$ such that $S_i(x_1, \ldots, x_k, a)$,
   there are $y_1, \ldots, y_k$ such that $S_i(y_1, \ldots, y_k, b)$, and $E(x_j, y_j)$ for $1 \leq j \leq k$, and
3. for $1 \leq i \leq m$, if S_i is of arity k then for all $y_1, \ldots, y_k$ such that $S_i(y_1, \ldots, y_k, b)$,
   there are $x_1, \ldots, x_k$ such that $S_i(x_1, \ldots, x_k, a)$, and $E(x_j, y_j)$ for $1 \leq j \leq k$.

a is *f-similar* to b in $\mathfrak{A}$ if there is an f-simulation E of $\mathfrak{A}$ such that $E(a, b)$. For a given f and a, there need not be an element of $\mathfrak{A}$ that is f-similar to a, but if there is one then there is only one, by the extensionality of $\mathfrak{A}$. We write $f.a$ for the unique element f-similar to a, should it exist. An extensional structure is a *substitution structure* if $f.a$ exists for every f and a.

For example, suppose that there are solutions to the equations

$$\begin{array}{rcl} v & = & \langle\!\langle\, r; a, b \,\rangle\!\rangle \\ w & = & \langle\!\langle\, r; b, b \,\rangle\!\rangle \\ x & = & \langle\!\langle\, r; y, a \,\rangle\!\rangle \\ y & = & \langle\!\langle\, r; x, b \,\rangle\!\rangle \\ z & = & \langle\!\langle\, r; z, b \,\rangle\!\rangle \end{array}$$

in an extensional information structure $\mathfrak{A}$. Let $[a \mapsto b]$ be the function with domain $\{a\}$ such that $[a \mapsto b](a) = b$, and let E be the relation with graph

$$\{\langle v, w\rangle, \langle x, z\rangle, \langle y, z\rangle, \langle a, b\rangle\} \cup I_{\{b,r,z\}},$$

where $I_{\{b,r,z\}}$ is the identity relation on $\{b, r, z\}$. It is easy to check that E is an $[a \mapsto b]$-simulation on $\mathfrak{A}$, and so $[a \mapsto b].v = w$ and $[a \mapsto b].x = [a \mapsto b].y = z$.

Set structures are substitution structures, by the set-theoretic Axiom of Replacement, together with Foundation or Anti-Foundation, as the case may be; but many of the structures in which we are interested are not. For example, any simple information structure satisfying Axioms A7 (finite roles) and A10.2 (extensionality) is a substitution structure if and only if it satisfies the generality principle given in Axiom A11.1. No room is made for 'legitimate' grounds for the non-existence of $f.a$, such as the sortal restrictions considered in Section 2.7. We would like to generalize the principle of sortal generality, given by Axiom A11.2, but it refers to specific features of the structure of simple infons (their roles), and so cannot be generalized to arbitrary extensional structures. Instead, we rest with the fact that all extensional structures are at least *partial* substitution structures, and interpret the non-existence of $f.a$ as the inappropriateness of the substitution f for a.[18]

Next, we prove a result that shows us how to extend an arbitrary extensional structure to a substitution structure.

Theorem 3.2 *Every anti-founded extensional structure is a substitution structure.*

PROOF:Suppose that $\mathfrak{A} : [A, S_1, \ldots, S_m; R_1, \ldots, R_n]$ be an anti-founded extensional xtructure, and that f is a substitution in $\mathfrak{A}$. Let X be the set $\{x_a : a \text{ a non-atom of } \mathfrak{A}\}$. We consider a structural description $T = \{T(x)\}_{x \in X}$, using x as the set of parameters.

For a non-atom a, $T(x_a)$ is obtained from the canonical description T_a as follows: For each $\langle S, b_1, \ldots, b_k \rangle$ that holds in $\mathfrak{A}$, let $y_i = f(b_i)$ if b_i is in the domain of f $y_i = x_{b_i}$ if b_i is not in the domain of f. We put the corresponding tuple $\langle S, y_1, \ldots, y_k \rangle$ into $T(x_a)$. This defines the structural description T.

Since $\mathfrak{A}$ is anti-founded, let s satisfy T in $\mathfrak{A}$. For each $a \in \mathfrak{A}$, the construction has arranged that each $s(x_a)$ is f-similar to a. This proves that $f.a$ exists. QED

As an immediate corollary of Theorem 3.2 and Theorem 3.1, we see that every extensional may be embedded in an extensional substitution structure.

3.6 Infon Structures

We are now in a position to improve our model of infons in a way that allows for the smooth incorporation of complex infons and a uniform treatment of abstraction (to be given in Section 3.7). An advantage of the models used in Section 2—the 'simple information structures'—is their simplicity, with only two structural relations. But there are disadvantages. First, the predicate 'infon', as applied to these structures, is defined so as to refer only to simple infons. This makes the move to structures containing complex infons a bit awkward. Second, the functions associating objects to roles in the infons of a simple information structure are not *internal* functions. This complicates the treatment of substitution and abstraction.

[18] An interesting, weaker, substitutional property of extensional structures is as follows: if $f.a$ and $g.a$ both exist and f and g are compatible functions then $(f \sqcup g).a$ also exists. This is implied by sortal generality if the domains of f and g do not contain roles.

Now that we know how to represent the structure of functions (in function structures) in the framework of the theory of structural relations, we can solve this problem. The basic idea is to represent the infon $\sigma = \langle\!\langle\, r; i:a, j:b, k:c \,\rangle\!\rangle$ by as a pair consisting of the relation r and an internal function α with domain $\{i, j, k\}$ and such that $\alpha^*(i) = a$, $\alpha^*(j) = b$, and $\alpha^*(k) = c$. The function α is called the *assignment* of σ. We adapt the functional notation for assignments and write

$$\alpha = [i:a, j:b, k:c] \text{ and } \sigma = \langle\!\langle\, r; \alpha \,\rangle\!\rangle.$$

The structure of infons is now represented using three structural relations, Rel^2, Ass^2 and Inf^1, and one non-structural relation, Approp^1. $\mathsf{Inf}(\sigma)$ iff σ is an infon; $\mathsf{Rel}(r, \sigma)$ and $\mathsf{Ass}(\alpha, \sigma)$ iff $\sigma = \langle\!\langle\, r; \alpha \,\rangle\!\rangle$. An assignment $\alpha = [i:a, j:b, k:c]$ is represented as an internal function, so that $\mathsf{Fun}(\alpha)$ and $\mathsf{App}(i, a, \alpha)$, and α is appropriate iff $\mathsf{Approp}(\alpha)$.

Definition 3.11 A function structure of type $[\mathfrak{A}, \mathsf{Ass}^2, \mathsf{Rel}^2, \mathsf{Inf}^1; \mathsf{Approp}^1]$ is an *infon structure* if it satisfies the following conditions:

1. (sorts) $\mathsf{Ass}^* \subseteq \mathsf{Inf}^*$, $\mathsf{Rel}^* \subseteq \mathsf{Inf}^*$, and $\mathsf{Fun}^* \cap \mathsf{Inf}^* = \emptyset$,
2. (basic infons) if $\mathsf{Rel}(r, \sigma)$ then $\exists \alpha\ \mathsf{Ass}(\alpha, \sigma)$,
3. (appropriateness) $\mathsf{Approp}(\alpha)$ iff $\exists \sigma\ \mathsf{Ass}(\alpha, \sigma)$, and
4. (substitution) if $\mathsf{Ass}(\alpha, \sigma)$ and if f is substitution such that $f.\alpha$ exists and $\mathsf{Approp}(f.\alpha)$, then $f.\sigma$ also exists.

An infon σ is *basic* if there is an α such that $\mathsf{Ass}(\alpha, \sigma)$. An infon structure is a *basic-infon structure* if every infon is basic.

Clause 3 (appropriateness) links the appropriateness of an assignment to the existence of an infon with that assignment. If, for example, a is not an appropriate filler for the role i then there is no α in Fun^* such that $\mathsf{Approp}(\alpha)$ and $\alpha^*(i) = a$. This is a generalization of the treatment of appropriateness in information structures. Clause 4 (substitution) ensures that the only issue governing substitution of infons is the appropriateness of the resulting assignment. This is an abstract form of the various principles of generality discussed in Section 2.7. To go further, we must look at the unification structure of appropriate assignments.

Definition 3.12 A unification structure is an *infon unification structure* if it is an infon structure and

1. $P(x, y)$ iff $x \sqsubseteq y$,
2. (function compatibility) if $\mathsf{Fun}(\alpha_1)$ and $\mathsf{Fun}(\alpha_2)$, then $C(\alpha_1, \alpha_2)$ iff for every i, a and b, if $\mathsf{App}(i, a, \alpha_1)$ and $\mathsf{App}(i, b, \alpha_2)$ then $a = b$,

3. (infon compatibility) Suppose that $\mathsf{Ass}(\alpha_1, \sigma_1)$, $\mathsf{Ass}(\alpha_2, \sigma_2)$, $\mathsf{Rel}(r_1, \sigma_1)$, and $\mathsf{Rel}(r_1, \sigma_1)$. Then $C(\sigma_1, \sigma_2)$ iff $r_1 = r_2$ and $C(\alpha_1, \alpha_2)$, and

4. (unification of appropriate assignments) if $C(\alpha_1, \alpha_2)$, $\mathsf{Approp}(\alpha_1)$ and $\mathsf{Approp}(\alpha_2)$, then $\mathsf{Approp}(\alpha_1 \sqcup \alpha_2)$.[19]

Infon unification structures give us our best theory of appropriateness. Part 4 ensures that the appropriateness of assignments depends only on 'local' issues: if α_1 and α_2 are compatible appropriate assignments, then so is $\alpha_1 \sqcup \alpha_2$; no additional grounds for inappropriateness can be introduced. Moreover, by clause 4 (substitution) of the definition of infon structures, substitutions that result in appropriate assignments may always be used to obtain new infons. For example, if

$\sigma = \langle\!\langle\, r; \alpha \,\rangle\!\rangle$ and $\alpha = [i : a, j : b, k : c]$,

f is the function with domain $\{a, b\}$ such that $f(a) = x$ and $f(b) = y$, and

g is the function with domain $\{b, c\}$ such that $g(b) = y$ and $f(c) = z$,

and both $f.\alpha$ and $g.\alpha$ exist and are appropriate—that is,

$$\mathsf{Approp}([i : x, j : y, k : c]) \quad \text{and} \quad \mathsf{Approp}([i : a, j : y, k : z]),$$

then not only is $[i : x, j : y, k : z]$ an assignment (by function unification), and appropriate (by unification of appropriate assignments) but the infon $f.\sigma = \langle\!\langle\, r; i : x, j : y, k : z \,\rangle\!\rangle$ exists (by substitution). This is still weaker than the principle of sortal generality, according to which appropriateness is dependent only on restrictions applying to individual roles. In particular, in infon unification structures there may be restrictions on related roles, as countenanced in the example of $\mathsf{next\text{-}to}$ at the send of Section 2.7. If $\mathsf{next\text{-}to}$ has roles nt_1 and nt_2, then we may deem assignments of the form $[\mathsf{nt}_1 : x, \mathsf{nt}_2 : x]$ inappropriate, without jeopardizing the unification principle.[20]

The second feature of infon structures not present in information structures, is the unary structural relation Inf. We say that σ is a *basic infon* if $\mathsf{Ass}(\alpha, \sigma)$ for some assignment α —by clause 2 (basic infons) all infons having a relation are required to have an assignment, possibly the empty assignment $[\,]$. In simple information structures, by definition of the term 'infon', all infons are basic, but in infon structures we may have non-basic infons. Of course, by extensionality and clause 1 (sorts), if there are no more structural relations in $\mathfrak{A}$ then there is at most one non-basic infon. If there are more structural relations, such

[19] In effect, both $\mathfrak{A}$ and the structure in which Fun^* is restricted to appropriate assignments are required to be function unification structures.

[20] This would imply that either $[\mathsf{nt}_1 : x]$ or $[\mathsf{nt}_2 : x]$ is also inappropriate, which may appear counterintuitive until one realizes that this does not prevent $[\mathsf{nt}_1 : x, \mathsf{nt}_2 : y]$ from being appropriate when $x \neq y$.

as $\mathsf{ConjunctOf}$, then we can add the axiom that $\mathsf{ConjunctOf}^* \subseteq \mathsf{Inf}^*$ without contradicting the axioms for infon structures.

This is a good example of how the 'holistic' nature of the extensionality property, which quantifies over *all* structural relations, allows a modular approach to the theory of structured objects. Another example is the treatment of polarity, discussed earlier in Section 2.4. We need only add new structural relations to deal with polarity and the extensionality condition takes care of the rest.

Definition 3.13 An infon structure of type $[\mathfrak{A}, \mathsf{Pos}^1, \mathsf{Neg}^1;]$ is *bi-polar* if it satisfies the following conditions:

1. σ is a basic infon iff either $\mathsf{Pos}(\sigma)$ or $\mathsf{Neg}(\sigma)$, and
2. if $\mathsf{Pos}(\sigma)$ then not $\mathsf{Neg}(\sigma)$.

Infons in Pos^* are *positive* and those in Neg^* are *negative*.[21]

Finite conjunctions and disjunctions are introduced effortlessly, as follows:

Definition 3.14 An infon structure of type $[\mathfrak{A}, \mathsf{ConjunctOf}^2, \mathsf{DisjunctOf}^2;]$ is $\wedge\vee$-*closed* if it satisfies the following conditions:

1. $\mathsf{ConjunctOf}^* \subseteq \mathsf{Inf}^*$ and $\mathsf{DisjunctOf}^* \subseteq \mathsf{Inf}^*$,
2. if $\mathsf{Inf}(\sigma)$ and $\mathsf{Inf}(\tau)$ then there is a $\sqsubseteq$-minimal object $\sigma \wedge \tau$ such that $\mathsf{ConjunctOf}(\sigma, \sigma \wedge \tau)$ and $\mathsf{ConjunctOf}(\tau, \sigma \wedge \tau)$
3. if $\mathsf{Inf}(\sigma)$ and $\mathsf{Inf}(\tau)$ then there is a $\sqsubseteq$-minimal object $\sigma \vee \tau$ such that $\mathsf{DisjunctOf}(\sigma, \sigma \vee \tau)$ and $\mathsf{DisjunctOf}(\tau, \sigma \vee \tau)$.

The principles of existence for infinite conjunctions and disjunctions are similar. We would like to say that for each set X of infons there is a conjunction $\bigwedge X$ and a disjunction $\bigvee X$. In stating these we must be careful because the properties of (external) sets are determined by the meta-theory, which—as we have seen—is not fixed. It is a good idea to bring recognize this indeterminacy explicitly by using internal sets instead, leaving open the question of how many internal sets there are. Thus we define:

Definition 3.15 An infon structure of type

$$[\mathfrak{A}, \mathsf{Set}^1, \in^2, \mathsf{ConjunctOf}^2, \mathsf{DisjunctOf}^2;]$$

is $\bigwedge\bigvee$-*closed* if it is also a set structure, and

[21] For *bi-polar infon unification structures* we must modify clause 3 of the definition of infon unification structures to take polarity into account:

3′. if $\mathsf{Ass}(\alpha_1, \sigma_1)$, $\mathsf{Ass}(\alpha_2, \sigma_2)$, $\mathsf{Rel}(r_1, \sigma_1)$, and $\mathsf{Rel}(r_1, \sigma_1)$, then $C(\sigma_1, \sigma_2)$ iff $r_1 = r_2$, $C(\alpha_1, \alpha_2)$, and $\mathsf{Pos}(\sigma_1) \leftrightarrow \mathsf{Pos}(\sigma_2)$.

1. $\mathsf{ConjunctOf}^* \subseteq \mathsf{Inf}^*$ and $\mathsf{DisjunctOf}^* \subseteq \mathsf{Inf}^*$

2. if $\mathsf{Set}(X)$ and $\mathsf{Inf}(\sigma)$ for each σ such that $\in(\sigma, X)$, then there is a $\sqsubseteq$-minimal object $\bigwedge X$ such that $\mathsf{ConjunctOf}(\sigma, \bigwedge X)$ for each σ such that $\in(\sigma, X)$,

3. if $\mathsf{Set}(X)$ and $\mathsf{Inf}(\sigma)$ for each σ such that $\in(\sigma, X)$, then there is a $\sqsubseteq$-minimal object $\bigvee X$ such that $\mathsf{DisjunctOf}(\sigma, \bigvee X)$ for each σ such that $\in(\sigma, X)$.[22]

Negations may be modelled using a structural relation NegOf^2 with the following existence condition: if $\mathsf{Inf}(\sigma)$ then there is a unique object $\neg\sigma$ such that $\mathsf{NegOf}(\sigma, \neg\sigma)$. Another approach is to base our model on a bi-polar infon structure and treat NegOf^2 as a *non*-structural relation defined as follows:

Definition 3.16 A bi-polar infon structure of type $[\mathfrak{A}; \mathsf{NegOf}]$ is a *de Morgan infon structure* if it is $\bigwedge\bigvee$-closed and satisfies the following conditions:

1. if $\mathsf{NegOf}(\sigma, \tau)$ then $\mathsf{Inf}(\sigma)$ and $\mathsf{Inf}(\tau)$,

2. if $\mathsf{NegOf}(\sigma, \tau)$ and $\mathsf{NegOf}(\sigma, \tau')$ then $\tau = \tau'$,

3. if $\mathsf{NegOf}(\sigma, \tau)$ then $\mathsf{NegOf}(\tau, \sigma)$,

4. if $\mathsf{Rel}(r, \sigma)$ and $\mathsf{Arg}(\alpha, \sigma)$ then there is an infon $\neg\sigma \neq \sigma$ such that $\mathsf{NegOf}(\sigma, \neg\sigma)$ $\mathsf{Rel}(r, \neg\sigma)$ and $\mathsf{Arg}(\alpha, \neg\sigma)$, and

5. if $\mathsf{Set}(X)$, $\mathsf{Set}(Y)$ and

   (a) for each $\sigma \in X^*$ there is a $\tau \in Y^*$ such that $\mathsf{NegOf}(\sigma, \tau)$

   (b) for each $\tau \in Y^*$ there is a $\sigma \in X^*$ such that $\mathsf{NegOf}(\sigma, \tau)$

   then $\mathsf{NegOf}(\bigwedge X, \bigvee Y)$.[23]

A remaining question is how to distinguish between saturated and unsaturated infons in infon structures. In basic-infon structures, the answer is clear from the discussion in Section 2.8: saturated basic infons are those that are maximal in the $\sqsubseteq$-ordering. For complex infons, it is not so easy. In a $\wedge\vee$-closed infon structure, if σ and τ are distinct basic infon then $\sigma \sqsubseteq \sigma \wedge \tau$ and so σ is not maximal, no matter how many roles are filled. We decide the matter as follows.

[22] Every $\bigwedge\bigvee$-closed infon structure is also $\wedge\vee$-closed because for infons σ and τ there is an internal set X with $X^* = \{\sigma, \tau\}$ and so $\sigma \wedge \tau = \bigwedge X$; and similarly for disjunction.

[23] No mention is made of polarities here, but from the definition of bi-polar infon structures we know that if σ is a basic infon then it is either positive or negative (but not both), and from clause 4, $\sigma \neq \neg\sigma$; so if σ is positive then $\neg\sigma$ is negative, and vice versa.

Definition 3.17 A basic infon σ is *unsaturated* if there is a distinct basic infon τ such that $\sigma \sqsubseteq \tau$. A non-basic infon is *unsaturated* if it has a component that is unsaturated. An infon is *saturated* if it is not unsaturated.[24]

By now, the reader will have an impression of some of the possibilities for modelling basic and complex infons using the theory of structural relations. Observers of the development of Situation Theory have been somewhat frustrated because of the lack of consensus on which structures—or even on which kinds of structures, obeying which axioms—are to be chosen for the theory (see Barwise 1989). We will see other reasons for diversity in Section 4, but the basic reason is clear from the above: logical and structural properties of infons are separate but interdependent, and striking the right balance is a difficult matter.

3.7 Abstraction

Abstraction is a natural operation in semantics. By abstracting Paul from the information σ that Raymond cooked the omelette for Paul, we obtain the property of being someone for whom Raymond cooked the omelette. By abstracting both Paul and the omelette from σ, we obtain the relation that holds between x and y just in case Raymond cooked x for y. This relation and the previous property differ from the relations and properties we have considered so far in that they are clearly not atomic. In this section we see how infon structures can be used to model these complex relations. [25]

There are two parts to abstraction over a structural object: the object after abstraction, which contains 'holes' where the abstracted object used to be, and something indicating where the holes are so that subsequently they may be filled in. We call these the 'abstract' and the 'pointer', respectively.[26] The relationship between abstract and pointer may be captured by a binary structural relation Abs with $\mathsf{Abs}(a, x)$ meaning that x is a pointer to the abstract a.

Definition 3.18 An *abstraction structure* $\mathfrak{A}$ is an extensional structure with structural relation Abs^2 such that if $\mathsf{Abs}(a, x)$ and $\mathsf{Abs}(a', x)$ then $a = a'$. The objects of sort Abs^* are called *pointers*. If $\mathsf{Abs}(a, x)$ then x is said *to point* to the *abstract* a.

The objects of interest in abstraction structures are the abstracts, but they do not form a structural sort; rather, it is the pointers that are structurally

[24] This is not a complete definition. In particular, it does not legislate on the saturation of a non-basic infon that is a component of itself but all of whose other components are saturated. In all such cases we say that the infon is saturated.

[25] In Situation Theory, abstraction has been treated as a structural opertation, and here we shall attempt to model it using structural relations. This goes against the dominant trend in semantics, which is to use functional types in a type theory, described by means of the lambda-calculus (see Chapter 9, this volume). Interestingly, Lunnon (1996), building on the work of Aczel and Lunnon, integrates these two approaches.

[26] Pointers are often called 'indeterminates' and 'parameters'; but that description of them is bad for the philosopher's digestion, and leaves a sour taste on the metaphysical palate.

determined. The intention is that each pointer x of an abstract a is a hereditary component of a. We obtain the results of 'applying' an abstract a to some argument b by substituting b for a pointer x in a. To see how this achieves the desired effect we must consider how abstracts arise by abstraction from other objects in the domain.

Suppose, for example, that we want to abstract b from the infon $\sigma = \langle\!\langle\, r; i : b, j : c \,\rangle\!\rangle$. We would expect to obtain an abstract a with exactly one pointer x such that the result of substituting b for x in a is just σ. Furthermore, the result of substituting b' for x in a should be the infon $\langle\!\langle\, r; i : b', j : c \,\rangle\!\rangle$, if this infon exists. This is captured in the following definition.

Definition 3.19 Suppose $\mathfrak{A}$ is an abstraction structure. Given elements a and b in $|\mathfrak{A}|$, an abstract $\lambda b.a$ is *the abstraction of b from a* if there is an x such that

1. $[b \mapsto x].a = \lambda b.a$,
2. $\mathsf{Abs}(\lambda b.a, x)$, and
3. the only sort of x is Abs^*.

The definite article is justified in the above definition, because if $\mathsf{Abs}([b \mapsto x].a, x)$ and $\mathsf{Abs}([b \mapsto y].a, y)$ then, by extensionality, $x = y$. For example, if $\sigma = \langle\!\langle\, r; i : b, j : c \,\rangle\!\rangle$ and $\lambda b.\sigma$ exists, then there is an x such that $\mathsf{Abs}(\langle\!\langle\, r; i : x, j : c \,\rangle\!\rangle, x)$. But suppose there was a y such that $\mathsf{Abs}(\langle\!\langle\, r; i : y, j : c \,\rangle\!\rangle, y)$. Since this is these are the only structural relations determining x and y, they are obviously bisimilar, and so equal. A similar argument shows that

$$\lambda b.\langle\!\langle\, r; i : b, j : c \,\rangle\!\rangle \quad = \quad \lambda b'.\langle\!\langle\, r; i : b', j : c \,\rangle\!\rangle,$$

assuming that b and b' are not hereditary components of r, i, j, or c. This is the expected principle of α-equivalence.

Notice also that under similar assumptions the order of abstraction is insignificant. If, for example, b and c are both atomic then

$$\lambda c.\lambda b.\langle\!\langle\, r; i : b, j : c \,\rangle\!\rangle \quad = \quad \lambda b.\lambda c.\langle\!\langle\, r; i : b, j : c \,\rangle\!\rangle.$$

In both cases, the abstract obtained has two pointers, pointing to the positions previously occupied by b and c respectively, but there nothing to distinguish the pointers apart from that, and so no trace of the order in which the abstraction was performed. This motivates a slightly more general notion of abstraction.

Definition 3.20 Suppose $\mathfrak{A}$ is an abstraction structure. Given an element a of $|\mathfrak{A}|$ and a set B of elements of $|\mathfrak{A}|$, an abstract $\Lambda B.a$ is *the simultaneous abstraction of B from a* if there is a function $\pi\colon B \rightarrow |\mathfrak{A}|$ such that

1. $\pi.a = \Lambda B.a$,

2. $\mathsf{Abs}(\pi.a, \pi(b))$ for each $b \in B$, and

3. for each $b \in B$, the only sort of $\pi(b)$ is Abs^*.

Despite the previous observation about commuting lambdas, one should be careful to distinguish simultaneous abstraction from successive abstractions. When the elements of B are structurally dependent, the results obtained may differ (see Aczel and Lunnon (1991), Ruhrberg (1994)). Nonetheless, it is clear that simultaneous abstraction is a generalization of abstraction because $\lambda b.a = \Lambda\{b\}.a$.

We have already observed that abstraction obeys an identity principle analogous to α-equivalence in the λ-calculus, at least in special cases. To state the general result observe that if $\sigma = \langle\!\langle\, r; b, c \,\rangle\!\rangle$ and $\sigma' = \langle\!\langle\, r; b', c \,\rangle\!\rangle$ and b and b' are structurally independent then $[b \mapsto b'].\sigma = \sigma$ and $[b' \mapsto b].\sigma' = \sigma$. These are the conditions needed to show that the two abstracts are bisimilar and so justify the claim that $\lambda b.\sigma = \lambda b'.\sigma'$. We state, without proof, the generalization of this result.

Theorem 3.3 *(α-identity) If $\Lambda B.a$ and $\Lambda B'.a'$ exist in an abstraction structure and there is a one-to-one correspondence f between B and B' such that $f.a = a'$ and $f^{-1}.a' = a$ then $\Lambda B.a = \Lambda B'.a'$.*

We now show that extensional structures with abstraction really exist.

Definition 3.21 An abstraction structure $\mathfrak{A}$ is a *Lambda structure* if for every element a of $|\mathfrak{A}|$ and a set B of elements of $|\mathfrak{A}|$, the simultaneous abstraction $\Lambda B.a$ exists.

Theorem 3.4 *Every extensional structure can be extended to a Lambda structure.*

Proof: Given an extensional structure $\mathfrak{A}$, extend it trivially to an abstraction structure by adding the structural relation Abs^2 with an empty extension. By Theorem 3.1 this can be extended to an anti-founded structure $\mathfrak{B}_0$. Throw away all pointers in this structure that point to more than one abstract to obtain an abstraction structure $\mathfrak{B}$. We show that $\mathfrak{B}$ contains a Lambda structure extending $\mathfrak{A}$. Suppose a is an element of $|\mathfrak{B}|$ and B is a set of elements of $|\mathfrak{B}|$. We exhibit a structural description whose satisfier gives the needed element $\lambda B.a$. For each $b \in B$, let $T_b = \{\langle \mathsf{Abs}, a\rangle\}$. Let H be the set of hereditary components of a that are not in B, and for each $c \in H$, let T_c be the canonical description of c in $\mathfrak{B}$. Now the structural description we want is $T = \{T_x\}_{x \in B \cup H}$. Since $\mathfrak{B}$ is contained in an anti-founded structure, there is a function s mapping into $\mathfrak{B}_0$ that satisfies T. But, clearly, the range of s does not contain any of the elements we threw away, so s maps into $\mathfrak{B}$. Thus for each b in B, $s(b)$ satisfies the description $s.T_b = \{\langle \mathsf{Abs}, s.a\rangle\}$ and so $\mathsf{Abs}(s.a, s(b))$. Moreover, $s(b)$ is only of sort Abs^*, and so $s.a$ is the required abstract $\Lambda B.a$. QED

A short digression on the ontology of abstracts is in order. If σ is an infon with argument b then $\lambda b.\sigma$ is also of sort Inf^*. But, intuitively, it is a property not an infon. Moreover, it has a pointer as an argument, filling the role left by b, and so interfering with whatever implicit appropriateness conditions there may have been on this role. It is therefore useful to distinguish between genuine infons and these new entities, which we may call *infon abstracts*. Infon abstracts are just properties (or relations if they have more than one pointer) and so ontologically acceptable. That is not quite the end of the matter. Consider the infon $\tau = \langle\!\langle\, r, \sigma \,\rangle\!\rangle$ and the abstract $\lambda b.\tau$ with pointer x. The abstract $\lambda b.\tau$ is an infon abstract, not a genuine infon. But it has an argument $x.\sigma$ that is also of sort Inf^* but neither an infon abstract nor, intuitively, a genuine infon. The only reasonable position to adopt is that $x.\sigma$ is a *part* of an abstract. Fortunately, it is easily determined which abstract it is a part of, because the pointer x points to it. In this way, we may satisfy ourselves that no unexpected ontological categories have been created.[27]

3.8 Application

In the previous section we saw that an extensional structure may be extended with abstracts of the form $\Lambda B.a$, obtained by abstracting the elements of B from the object a. The purpose of these abstracts is to model generalizations across a class of structured objects. In particular, infon abstracts may be used to model properties and relations. To see how this works, we need an account of 'application'.

Definition 3.22 Suppose $\mathfrak{A}$ is an abstraction structure. If a is an abstract then a function $f\colon X \rightarrow |\mathfrak{A}|$ is an *(external) assignment for* a if every x in X is a pointer of a. It is *appropriate for* a if $f.a$ exists. If f is appropriate for a then *the application of* a *to* f, written $a^*(f)$, is $f.a$. In this way, every abstract a is associated with a (second-order) function a^* mapping appropriate assignments to elements of $|\mathfrak{A}|$.

Theorem 3.5 *(β-identity) if $\Lambda B.a$ exists and f is an appropriate assignment for $\Lambda B.a$ then $(\Lambda B.a)^*(f) = f.a$.*

PROOF: This follows directly from the definitions of simultaneous abstraction and application. QED

How, then, should we incorporate complex properties and relations into infon structures? The key point is that the appropriateness conditions of infons having complex relations should be determined by those of basic infons.

[27] The ontological status of 'parametric objects', which are closely related to our abstracts and their parts, is the subject of much debate. See, for example, Westerståhl (1990).

Definition 3.23 Let $\mathfrak{A}$ be an infon structure that is also an abstraction structure. $\mathfrak{A}$ has *coherent appropriateness conditions* if for each infon $\sigma = \langle\langle\, r; f \,\rangle\rangle$ in $|\mathfrak{A}|$, if r is an infon abstract then f^* is an appropriate assignment for r. $\mathfrak{A}$ is *relation-closed* if for each infon σ and each set B of elements of $\mathfrak{A}$, the infon abstract $\Lambda B.\sigma$ exists and for each appropriate assignment f for $\Lambda B.\sigma$, the infon $\langle\langle\, \Lambda B.\sigma; f \,\rangle\rangle$ also exists.

We state without proof the following result:

Theorem 3.6 *Every infon structure is extendible to a relation-closed infon structure.*

By way of closing this section, we note that application has only been modelled in an 'external' way. We may ask what the properties of an 'internal' application operation should be. For more on this, see Aczel and Lunnon (1990), and Lunnon (to appear).

4 Truth and Circumstance

Notoriously, the word 'information' is ambiguous. Traditionally, the possession of information that Smith is an anarchist like the knowledge of the same implies that Smith *is* an anarchist. With the advent of mechanisms for the gathering, storing and retrieval of large amounts of information, the word has taken on a more neutral meaning. According to more modern usage, to be adopted here, the information that Smith is an anarchist may be stored on a computer file in ASCII code, transmitted across the globe, converted to speech and heard over a loud speaker, without Smith's ever having had a subversive political thought in his life.

A problem with the modern usage is that a new word is required to separate genuine items of information, in the traditional, truth-implying sense, from mere infons. Provisionally, we will use the adjective 'factual' for this purpose. Thus, the infon $\langle\langle\, \mathsf{anarchist}; \mathsf{Smith} \,\rangle\rangle$ is *factual* just in case Smith is indeed an anarchist; otherwise it is non-factual, but information nonetheless.

The present use of the word 'factual' is provisional because the distinction on which it rests is an important contention in situation theory. No one would deny of the statement that Smith is an anarchist, that it is true if and only if Smith is an anarchist. It is less clear that this biconditional may be used to determine whether the corresponding infon is factual. To see the difficulty, let us suppose that Smith had a rebellious youth, but has now settled down into conformist middle age. An utterance of 'Smith is an anarchist' during Smith's early period would have made a true statement, but an utterance of the same sentence today would not. On the assumption that both utterances express the infon $\langle\langle\, \mathsf{anarchist}, \mathsf{Smith} \,\rangle\rangle$, we arrive at an impasse: which statement do we use to decide whether the infon is factual?

Of course, the problem can be resolved by denying the assumption that both utterances express the same information. We may say that the infons expressed by the two utterances are distinguished by a temporal role that is filled by the time of each utterance. Instead of one infon, we have two:

$\langle\!\langle\, \mathsf{anarchist}; \mathsf{subject}\colon\ \mathsf{Smith}, \mathsf{time}\colon\ \mathsf{Friday\ 15th\ May,\ 1970}\,\rangle\!\rangle$ and
$\langle\!\langle\, \mathsf{anarchist}; \mathsf{subject}\colon\ \mathsf{Smithtime}\colon\ \mathsf{Monday\ 18th\ July,\ 1994}\,\rangle\!\rangle$.

The first infon is presumed to be factual, because of the truth of the statement made in 1970, and the second non-factual, because of the falsity of the statement made today. The $\langle\!\langle\, \mathsf{anarchist}; \mathsf{Smith}\,\rangle\!\rangle$ is either excused from considerations of factuality because it is unsaturated, or supposed not to exist at all. We appropriate the terminology of Barwise and Etchemendy (1987) by calling this account *Russellian*.

A quite different solution is to take the truth or falsity of a statement to depend on more than just the information expressed. Following Austin (1961), two components of a statement are distinguished: the situation, event or state of affairs that the statement is about, and what is said about it. On this, the *Austinian* account, the difference between the two statements arises not because of a difference in the information they express, but because different situations are described: the utterance in 1970 correctly describes Smith's attitudes at that time, whereas today's utterance falsely describes his present state of bourgeois conformity.[28] The infon $\langle\!\langle\, \mathsf{anarchist}; \mathsf{Smith}\,\rangle\!\rangle$ is factual if taken to be about Smith's beliefs and actions in 1970, but not if it is taken to be about his current views.

Proponents of both accounts agree that the common information expressed by the two utterances is neither factual nor straightforwardly non-factual. They differ on the reason for this failure. According to the Russellian, the infon $\langle\!\langle\, \mathsf{anarchist}; \mathsf{Smith}\,\rangle\!\rangle$, if it exists at all, is merely the unsaturated common part of the distinct, saturated infons expressed by the two utterances. On the Austinian account, all claims to factuality if intelligible at all are relative to the situation the information is taken to be about. The infon $\langle\!\langle\, \mathsf{anarchist}; \mathsf{Smith}\,\rangle\!\rangle$ is factual if it is taken to be information about the situation described by the first utterance, but not if it is taken to be about the situation described by the second.

On matters of detail, there is room for interplay between the two positions. For example, the Austinian may agree that the two statements express different information by virtue of the indexical nature of tense, but still maintain that the factual status of the information expressed is determined by the described situation. Moreover, the debate may be conducted in different areas. Instead of focusing on tense, the Austinian may claim an implicit relativity to a perspective

[28] For Barwise and Etchemendy the terms 'Russellian' and 'Austinian' refer to different views about *propositions*, but the distinction is closely related to the present one. Propositions will be discussed in Section 4.8. Of course, neither Russell nor Austin were concerned with information *per se*, but basic infons play a very similar role in Situation Theory to the basic facts of Logical Atomism.

from which the assessment of Smith's political views is to be made—it may be Smith's own, his fellow citizens', 'the authorities', or even that of the person making the statement. The Russellian account may respond by insisting either that the term 'anarchist' is ambiguous, or that the infon expressed has an extra role for the perspective from which the property of being an anarchist is judged.

In the next few sections we shall consider the Russellian account in more detail, before returning to the Austinian account in Section 4.5.

4.1 Fact Structures

From a purely technical standpoint, the Russellian account is the more straightforward of the two. The information expressed by an unambiguous statement is modelled as a saturated infon, which may or may not be factual. If it is factual then the statement is true; if not, it is false. Those roles of the infon that are not filled by arguments given explicitly by linguistic features of a statement must be determined by other features of the context in which the statement is made. These additional roles are called 'hidden parameters'. For example, the temporal 'parameter' of the statement that Smith is an anarchist is filled by the time at which the statement is made. If there are further hidden parameters, such as the perspective from which this assessment is made, then the statement is either ambiguous or the additional role must be filled by some other aspect of the context.

This account can be modelled using a 'fact structures', defined as follows.

Definition 4.1 An infon structure $\mathfrak{F}$ of type $[\mathfrak{A}; \mathsf{Fact}]$ is a *fact structure* if for each $\sigma \in A$, if $\mathsf{Fact}(\sigma)$ then σ is a saturated infon. An infon σ is a *fact* if $\mathsf{Fact}(\sigma)$. $\mathfrak{F}$ is *trivial* if every saturated infon is a fact.

Standard fact structure may be constructed from arbitrary relational structures in a straightforward manner. If $\mathfrak{M} = \langle M, R_1, \ldots, R_n \rangle$ is a relational structure and $\mathrm{SInf}(\mathfrak{M})$ information structure constructed from it using Construction 2.1, then let $F(\mathfrak{M})$ be the fact structure extending $\mathrm{SInf}(\mathfrak{M})$ with

$$\mathsf{Fact}(\langle 3, R_i, \alpha \rangle) \text{ iff in } \mathfrak{M}_0,\ R_i(\alpha(1), \ldots, \alpha(\nu_i))$$

for each $i \leq n$ and each $\alpha\colon \{1, \ldots, \nu_i\} \to M$. Every infon in a standard structure is saturated, so there is nothing more to check. We call $F(\mathfrak{M})$ the *standard* fact structure generated by $\mathfrak{M}$.[29]

The main problem for the Russellian is that there are many statements that would not usually be called ambiguous, but for which hidden parameters remain

[29] Unlike the construction of standard basic-infon structures, the construction of standard basic-fact structure is one-one: every relational structure is uniquely represented as a basic-fact structure by the above construction. Conversely, by Theorem 2.1, every fact structure satisfying Axioms A1 to A14 is isomorphic to a standard fact structure; but in this direction the representation is not unique, because of the arbitrary ordering of relations in a relational structure.

unfilled. These are statements that only express unsaturated information, even when all contextual factors are considered. They are more common that one might suspect. Suppose that in the course of telling the story of Smith's involvement in a riot, we say that he panicked and ran. Superficially, the content of the statement is the information that Smith panicked and ran, which might be modelled by the infon

$$\langle\!\langle\, \mathsf{panic}; \mathsf{Smith} \,\rangle\!\rangle \;\wedge\; \langle\!\langle\, \mathsf{run}; \mathsf{Smith} \,\rangle\!\rangle.$$

This infon is unsaturated because panicking and running are events, and so occur at a particular time that has not yet been specified. Unlike in the previous example, the tense of the statement (simple past) is not sufficient to determine a saturated content. Unless there are aspects of the circumstances in which the statement is made that provide a way of saturating the infon, it is difficult to see how a truth-value is to be determined.

There are various ways in which the missing information may be provided. The date of Smith's alleged actions may be given explicitly at some earlier point in the narrative ("On the 15th May, 1970, Smith was involved in a riot...') or provided by some other means: the date of the riot may be common knowledge, or the narrator may have special access to the facts, perhaps by being witness to the original event.

The hunt for contextual factors that determine the content of statements is ultimately restricted by the lack of precision in ordinary language. A threat to the Russellian view is that suggestion that ordinary language is *indefinitely imprecise*: there are always ways of making a given statement more precise, with successive approximations to the 'missing' information, and no definite level of precision beyond which such improvements cannot be made.

For example, it there is no way of finding the exact date of the riot from the context—if for example, the context only specifies that it occurred in May, 1970—then we would be forced to conclude that the statement is ambiguous. But even if the date is supplied by context, we may still be in trouble. Supposing the date to be determined to be 15 May, 1970, the content of the statement would be the infon

$$\langle\!\langle\, \mathsf{panic}; \mathsf{subject} : \mathsf{Smith}, \mathsf{date} : 5/15/70 \,\rangle\!\rangle \;\wedge$$
$$\langle\!\langle\, \mathsf{run}, \mathsf{subject} : \mathsf{Smith}, \mathsf{date} : 5/15/70 \,\rangle\!\rangle.$$

Unfortunately, this infon is still unsaturated, as we can see by supposing that Smith was involved in two separate clashes with the police on the same day. On the first occasion, he bravely stood his ground; on the second he panicked and ran. The mere possibility of two clashes is sufficient to show that the infon is unsaturated and so the Russellian is left with ambiguity, even if there was only one clash.

In an effort to avoid the conclusion that there is widespread ambiguity in our use of ordinary language, we may search for more elaborate ways in which the context fills hidden parameters. If the exact time is not supplied in the linguistic context or by common knowledge, it may be determined by a causally related chain of events starting with Smith's actions in the riot.

The danger in this move is that information expressed by a statement may be unknown, even to the person who makes the statement. The narrator may well remember the day of the riot, and even whether it occurred in the morning or afternoon, but not the exact time. If this time is to be a component of the infon expressed by the statement then this is information that the narrator does not possess.

If, on the other hand, we are to resist the move to considering factors that go beyond the knowledge of the narrator, there are other problems. If we take the narrator's beliefs to determine the temporal parameter, then falsehoods are liable to appear in surprising places. Suppose, for example, that the narrator falsely believed the incident took place at around noon on the 15th May, 1970; in fact, it occurred at about 11 a.m. It is awkward to maintain that this false belief, which is otherwise irrelevant to the narrative, makes his statement that Smith panicked and ran false.

Finally, there is a problem in making the demand for ever-greater precision compatible with our somewhat roughshod ontology of events and processes. Although we think of the event of Smith panicking and running as having duration, it is not the sort of interval that can be measured in seconds.

We may hope that for each kind of description there is a fixed tolerance with which information about the values of hidden parameters must be specified. In this way, we could avoid the problems of requiring perfect precision. In many cases, the truth of a statement in which a time is not mentioned explicitly is fairly robust: small variations do not alter the truth-value. But there are circumstances in which small changes do matter—Smith's two clashes with the police may have happened within half-an-hour of each other—and so, the precision required for the saturation of a given piece of information is indefinite, depending on particular facts about the circumstances of the original event, and not just about the later circumstances in which the statement is made.

A way out of this web of difficulty is to embrace the conclusion that many, if not all, of the statements we make are ambiguous; or, better, to cut the link between lack of a truth-value and ambiguity. We may say that a statement is unambiguous if it expresses a determinate item of information, even if this information is unsaturated and the statement neither true nor false.

A quite different approach is to claim that an unsaturated infon is factual if there is some saturated fact of which it is a part. By quantifying over the possible ways of saturating the infon, we arrive at the rather charitable position that a statement is true just in case, in the circumstances, there is some way of construing it as expressing a saturated fact. This is reminiscent of Davidson's treatment (1967) of the same problem. There are two main difficulties with this solution. Firstly, an unrestricted quantification over possible arguments is clearly too generous. If in the course of discussing the riot in 1970 we state that

Smith panicked and ran, then this statement cannot be made true by Smith's panicking and running last week. Thus, we must appeal to the context again, this time to provide suitable restrictions to the range of quantification. Secondly, the restrictions on quantifiers are often interdependent. For example, if in interpreting the text 'Smith panicked. He ran.' we may use of the quantificational strategy, then the restrictions on the possibly ways of saturating the two unsaturated infons $\langle\!\langle\ \mathsf{panic};\ \mathsf{Smith}\ \rangle\!\rangle$ and $\langle\!\langle\ \mathsf{run};\ \mathsf{Smith}\ \rangle\!\rangle$ must be related: if $\langle\!\langle\ \mathsf{panic};\ \mathsf{Smith}, \mathsf{time}\colon t_1, \ldots\ \rangle\!\rangle$ and $\langle\!\langle\ \mathsf{run};\ \mathsf{Smith}, \mathsf{time}\colon t_2, \ldots\ \rangle\!\rangle$ are among the permissible saturations then t_1 must be before t_2. Such technical obstacles are not insuperable (see the chapter on dynamics in this volume) but introduce complications that we shall not go into here.[30]

4.2 Logic in Fact Structures

In Section 3.4 we showed how to model compound infons using the theory of structural relations. There we were concerned only with the structural properties of those objects; now we shall examine their logical properties. For example, any reasonable model of conjunctive facts should satisfy the condition that

$\mathsf{Fact}(\sigma \wedge \tau)$ iff $\mathsf{Fact}(\sigma)$ and $\mathsf{Fact}(\tau)$.

Certain special cases of this condition may be determined by structural properties of conjunction. For example, in the infon structure of Construction 3.1 conjunction is idempotent, and so in any fact structure extending this, $\mathsf{Fact}(\sigma \wedge \sigma)$ iff $\mathsf{Fact}(\sigma)$. But structural properties alone will not usually suffice to determine every instance of the above condition, and so we shall need to impose additional axioms. Typically, Situation Theory uses a de Morgan infon structure (from Section 3.6) $\mathfrak{F}$ that satisfies the following conditions:

1. $\mathsf{Fact}(\bigwedge \Sigma)$ iff $\mathsf{Fact}(\sigma)$ for each $\sigma \in \Sigma$, and
2. $\mathsf{Fact}(\bigvee \Sigma)$ iff $\mathsf{Fact}(\sigma)$ for some $\sigma \in \Sigma$.
3. $\mathsf{Fact}(\neg\sigma)$ iff $\mathsf{Inf}(\sigma)$ and not $\mathsf{Fact}(\sigma)$.

Let us say that a fact structure $\mathfrak{F}$ is *classical* if it does satisfy the above.

The logical behaviour of abstract relations must also be considered. A fact structure that is also an abstraction structure should satisfy the

> *Principle of β-equivalence for Facts*: if $\langle\!\langle\ r, f\ \rangle\!\rangle$ is an infon whose relation r is an infon-abstract then $\mathsf{Fact}(\langle\!\langle\ r, f\ \rangle\!\rangle)$ iff $\mathsf{Fact}(r^*(f))$.

[30] The role of indexicals constituent of the content of statements is discussed in John Perry's seminal article Perry (1979), and something close to the Russellian approach is pursued by Perry and David Israel in a series of papers about information beginning with Israel and Perry (1990).

We could go on to define a host of different kinds of fact structures, catalogued by their logical properties. Such a list would be tediously long and necessarily incomplete, and so there is no merit in producing it here. In any case, we wish the approach developed here to apply not just to the multitude of existing logics whose compounds have the familiar syntactic structure of (formal) sentences, but to future logics acting on more intricate structures—circular structures, infinite structures, and anything else that can be modelled using the theory of structural relations. Our solution is to give a general method for imposing more or less arbitrary logical structure on fact structures.

Our work is based on the notion of a consequence relation, as it is studied in philosophical logic. For a survey of this large field, see Makinson (1991).

Definition 4.2 A *consequence relation* $\vdash$ on a class C is a binary relation between subclasses of C such that $\Gamma \vdash \Delta$ iff $\Gamma' \vdash \Delta'$ for every partition $\langle \Gamma', \Delta' \rangle$ of C such that $\Gamma \subseteq \Gamma'$ and $\Delta \subseteq \Delta'$.[31] Given a fact structure $\mathfrak{F}$, a consequence relation $\vdash$ on the class of infons in $\mathfrak{F}$ is *sound* if whenever $\Gamma \vdash \Delta$ and every infon in Γ is a fact then some infon in Δ is a fact.

Consequence relations have the following familiar properties:

(identity) $\sigma \vdash \sigma$,

(weakening) if $\Gamma \vdash \Delta$ then $\Gamma', \Gamma \vdash \Delta, \Delta'$, and

(cut) if $\Gamma \vdash \Delta, \sigma$ and $\sigma, \Gamma \vdash \Delta$ then $\Gamma \vdash \Delta$.[32]

Moreover, if $\vdash$ is compact then these three conditions are jointly sufficient to ensure that $\vdash$ is a consequence relation (Barwise and Seligman, to appear).[33] Also note that a consequence relation is entirely determined by the class of partitions that lie in it. It is easy to see that a consequence relation is sound in a fact structure just in case the partition of Inf^* into facts and non-facts does not lie in the consequence relation.

In any de Morgan infon structure $\mathfrak{F}$, if $\vdash$ is a sound consequence relation on Inf^* and $\mathfrak{F}$ is classical then $\vdash$ must satisfy the following conditions:

1. $\bigwedge \Sigma, \Gamma \vdash \Delta$ if $\Sigma, \Gamma \vdash \Delta$,
2. $\Gamma \vdash \Delta, \bigwedge \Sigma$ if $\Gamma \vdash \Delta, \sigma$ for each σ in Σ,
3. $\Gamma \vdash \Delta, \bigvee \Sigma$ if $\Gamma \vdash \Delta, \Sigma$,
4. $\bigvee \Sigma, \Gamma \vdash \Delta$ if $\sigma, \Gamma \vdash \Delta$ for each σ in Σ,
5. $\Gamma, \neg\sigma \vdash \Delta$ if $\Gamma \vdash \sigma\Delta$,

[31] $\langle X, Y \rangle$ is a *partition* of C if $X \cup Y = C$ and $X \cap Y = \emptyset$.
[32] We use the usual abbreviations, e.g., Δ, σ stands for $\Gamma \cup \{\sigma\}$ when an argument of $\vdash$.
[33] $\vdash$ is *compact* if whenever $\Gamma \vdash \Delta$ there are finite sets $\Gamma_0 \subseteq \Gamma$ and $\Delta_0 \subseteq \Delta$ such that $\Gamma_0 \vdash \Delta_0$.

6. $\Gamma \vdash \neg\sigma\Delta \quad$ if $\quad \Gamma, \sigma \vdash \Delta$.

The conditions are analogues of the standard inference rules of the sequent calculus for (infinitary) classical propositional logic; let us say that $\vdash$ is *classical* if it satisfies all these conditions. Then it is easy to check that if $\mathfrak{F}$ is non-trivial and has a classical consequence relation then $\mathfrak{F}$ must be classical. These observations amount to the following characterization of de Morgan fact structures: a non-trivial de Morgan fact structure is classical iff it has a sound classical consequence relation.[34]

In this way, the logical properties of complex infons may be investigated by studying their consequence relations, without explicit reference to fact structures. A multitude of consequence relations have already been catalogued by logicians, and so we can import whichever logic we wish, defining the corresponding class of fact structures on which the consequence relation is sound.

4.3 Restriction

In Section 3.8 we saw that properties defined by means of abstraction inherit the appropriateness conditions of the infons over which they are defined. For example, the property $\lambda x.\langle\!\langle\, \mathsf{fed}; \mathsf{Martha}, x \,\rangle\!\rangle$ of being eaten by Martha is appropriate for just those objects a for which $\langle\!\langle\, \mathsf{fed}; \mathsf{Martha}, a \,\rangle\!\rangle$ is an infon.

For semantics, it has proved useful to introduce a more flexible notion of abstraction by which a much wider range of restrictions may be placed on appropriate assignments. On some semantic analyses, the information expressed by the sentence 'Martha fed him' is that the referent of 'him' has the property of being fed by Martha, modelled by the abstract $\lambda x.\langle\!\langle\, \mathsf{fed}; \mathsf{Martha}, x \,\rangle\!\rangle$. But the sentence also conveys the information that the person fed by Martha is male. One way of incorporating this information is to maintain that the abstract is restricted so that it may only be applied to males. In other words, we want a property p such that $\langle\!\langle\, p; a \,\rangle\!\rangle$ is

1. an infon iff a is male and $\langle\!\langle\, \mathsf{fed}; \mathsf{Martha}, a \,\rangle\!\rangle$ is an infon
2. factual iff a is male and $\langle\!\langle\, \mathsf{fed}; \mathsf{Martha}, a \,\rangle\!\rangle$ is factual

In some formulations, the restriction that a be male attaches to the abstracted variable (or 'parameter') x, and a theory of *restricted parameters* is developed (Gawron and Peters 1990, Fernando (1991); in others, the situation-theoretic universe is extended to include *restricted objects* quite generally: such objects are identical to their unrestricted counterparts when the (often contingent) restricting condition is met, and are otherwise 'undefined' or of some other ontologically inferior status. Plotkin (1990) adapts Curry's 'illative' approach

[34] Furthermore, there is a smallest classical consequence relation $\vdash_c$ on any de Morgan infon structure $\mathfrak{A}$, and so a non-trivial fact structure $\mathfrak{F}$ extending $\mathfrak{A}$ is classical iff $\vdash_c$ is sound on $\mathfrak{F}$.

to logic, treating restriction by means of a connective $\restriction$ whose formation rules require that the restriction be met: if σ is an infon and τ is a fact then $\sigma \restriction \tau$ is an infon; otherwise the expression '$\sigma \restriction \tau$' is not even well-formed.

Barwise and Cooper (1993) introduce an elegant graphical notation for restrictions and abstracts called Extended Kamp Notation, in honour of Kamp's Discourse Representation Structures (see the chapter on DRT in this volume). They write situation theoretic objects as boxes. A box directly to the right of another box acts as a restriction, and a box directly above lists abstracted objects. For example, the restricted property p mentioned above would be written as

<table>
<tr><td colspan="2">a</td></tr>
<tr><td>$\langle\!\langle\, \mathsf{fed}; \mathsf{Martha}, a \,\rangle\!\rangle$</td><td>$\langle\!\langle\, \mathsf{male}; a \,\rangle\!\rangle$</td></tr>
</table>

The approach adopted here is to extend the theory of abstraction and application developed in Sections 3.7 and 3.8.

Definition 4.3 A fact structure $\mathfrak{F}$ has *restricted abstraction* if it is also an abstraction structure and has an additional structural relation Res^2, such that if $\mathsf{Res}(\sigma, a)$ then a is an abstract and for each assignment f for a, $f.\sigma$ is an infon. We say that a is a *restricted abstract* and that σ is a *restriction* of a. For each object a, set B of objects in $|\mathfrak{F}|$, and set Σ of infons in $|\mathfrak{F}|$, an object $\Lambda B[\Sigma].a$ is the *abstraction of* B *from* a *restricted by* Σ if there is a function $\pi\colon B \to |\mathfrak{A}|$ such that

1. $\mathsf{Abs}(\Lambda B[\Sigma].a, \pi(b))$ for each b in B,
2. for each b in B, $\pi(b)$ is only of sort Abs^*, and
3. $\Lambda B[\Sigma].a$ satisfies the description $\langle \Lambda B[\Sigma].a, T_{\pi.a} \cup \{\langle \mathsf{Res}, \pi.\sigma\rangle \mid \sigma \in \Sigma\}\rangle$.[35]

$\mathfrak{F}$ is a *restricted-Lambda structure* if $\Lambda B[\Sigma].a$ exists for each a, B, and Σ.

The above nails down the structure of restricted abstracts but we still need to say how they may be applied. For this the key definition is the following:

> an assignment f for a is *appropriate for* a if $f.a$ exists and $f.\sigma$ is a fact.

As before, if f is appropriate for a then *the application of* a *to* f, written $a^*(f)$, is just $f.a$. $\mathfrak{F}$ has *coherent appropriateness conditions* if for each infon $\sigma = \langle\!\langle\, r; f \,\rangle\!\rangle$, if r is a (restricted) infon abstract then f^* is an appropriate assignment for r. $\mathfrak{F}$ is *restricted-relation-closed* if for each infon σ, each set B of elements of $\mathfrak{A}$, and each set Σ of infons, the infon abstract $\Lambda B[\Sigma].\sigma$ exists and for each appropriate assignment f for $\Lambda B[\Sigma].\sigma$, the infon $\langle\!\langle\, \Lambda B[\Sigma].\sigma; f \,\rangle\!\rangle$ also exists.

[35] This clause is slightly more complicated than the corresponding clause for unrestricted abstraction. It says that the restricted abstract $\Lambda B[\Sigma].a$ has the component structure of $\pi.a$, together with the additional structural relationships $\mathsf{Res}(\pi.\sigma, \Lambda B[\Sigma].a)$ for each σ in Σ.

Technically there are no difficulties here. Restricted Lambda structures and restricted-relation-closed structures may be constructed using methods similar to those used in Theorems 3.4 and 3.6.[36]

4.4 Internal Definability

We have already seen how an internal function f represents an external function f^* in a function structure, and how in an abstraction structure an abstract a represents a second-order function a^* mapping assignments to objects of the domain. In fact structures, internal relations represent external relations. Suppose $\mathfrak{F}$ is a fact structure and R is an n-ary relation on $|\mathfrak{F}|$. The basic idea is that an internal relation r of $\mathfrak{F}$ represents R if for each $a_1, \ldots, a_n$ in $|\mathfrak{F}|$, there is a saturated infon $\langle\!\langle\, r; a_1, \ldots, a_n \,\rangle\!\rangle$ in $\mathfrak{F}$ and $R(a_1, \ldots, a_n)$ iff $\mathsf{Fact}(\langle\!\langle\, r; a_1, \ldots, a_n \,\rangle\!\rangle)$. There are, however, several complications. Firstly, our use of the functional notation for infons hides an assumed correlation between the roles of r and the integers $1, \ldots, n$; this must be made explicit. Secondly, the condition requires that the infon $\langle\!\langle\, r; a_1, \ldots, a_n \,\rangle\!\rangle$ exists for every sequence $a_1, \ldots, a_n$ of elements of $|\mathfrak{F}|$. This will rarely be satisfied because most internal relations have non-trivial appropriateness conditions—it fails even in standard fact structures because the relations of the underlying relational structure only take ordinary objects, not infons and roles, as appropriate arguments. The definition must be revised so that the domain of each relation is restricted in some way.

Definition 4.4 Given a fact structure $\mathfrak{F}$, an n-ary external relation R on $|\mathfrak{F}|$ is *represented by* elements $r, i_1, \ldots, i_n$ of $|\mathfrak{F}|$ on the domain $A \subseteq |\mathfrak{F}|$ if

1. f is appropriate for r iff $f^*\colon \{i_1, \ldots, i_n\} \to A$, and
2. if f is appropriate for r then $\mathsf{Fact}(\langle\!\langle\, r; f \,\rangle\!\rangle)$ iff $R(f^*(i_1), \ldots, f^*(i_n))$.

In standard fact structures, every relation r has a fixed finite arity, n say, and (by the generality principle) the infon $\langle\!\langle\, r; a_1, \ldots, a_n \,\rangle\!\rangle$ exists for each sequence $a_1, \ldots, a_n$ of ordinary objects. Thus every internal relation of a standard fact structure represents an external relation on the domain of ordinary objects.

In the non-standard fact structures, internal relations may fail to represent for a variety of reasons: they may have infinite or variable arity; they may generate unsaturated infons; or they may have sortal—or even more complex—appropriateness conditions. A simple generalization of the above definition handles the last of these cases.

Definition 4.5 Given a fact structure $\mathfrak{F}$, an n-ary external relation R on $|\mathfrak{F}|$ is *represented by* elements $r, i_1, \ldots, i_n$ of $|\mathfrak{F}|$ *relative to* another n-ary relation D on $|\mathfrak{F}|$ if

[36] In the literature on Situation Theory, most authors have been concerned with objects restricted by true propositions rather than facts. We have explained the mechanism for the latter case only because we have yet to introduce propositions—see Section 4.8.

1. f is appropriate for r iff $D(f^*(i_1), \ldots, f^*(i_n))$, and

2. if f is appropriate for r then $\mathsf{Fact}(\langle\!\langle\, r; f \,\rangle\!\rangle)$ iff $R(f^*(i_1), \ldots, f^*(i_n))$.

In effect r is taken to represent R only in the context of D; outside of D it does not matter which sequences are in R, and this is captured by making the corresponding assignments for r inappropriate. This gives three different ways of classifying a sequence $\vec{a}$: either $\vec{a}$ in D and R, or $\vec{a}$ in D but not in R, or $\vec{a}$ is not in D. Correspondingly, there are three ways of classifying an assignment f: either $\langle\!\langle\, r; f \,\rangle\!\rangle$ is a fact, or $\langle\!\langle\, r; f \,\rangle\!\rangle$ is not a fact, or f is not appropriate for r and so there no infon of the form $\langle\!\langle\, r; f \,\rangle\!\rangle$.

Another approach is to say that, in general, internal relations represent *partial* relations on $|\mathfrak{F}|$. A partial relation has an extension and an anti-extension that are disjoint but that need not exhaust the domain. Facts with relation r represent sequences lying in the extension of R, and infons that are not facts represent sequences lying in the anti-extension of R.

The issue of infinite and variable arities can be handled with similar generalizations of the meta-theoretic concept of 'relation'.[37] Whichever characterization of representation is adopted, it is natural to ask which external relations are represented in a given fact structure. For example, we may wonder if any of the structural relations Arg, Rel, Inf and Fact are represented.

Plotkin (1990) has various negative results. Suppose that Fact is represented by F in a Lambda structure with classical negation. Let p be the abstract $\lambda x.\neg\langle\!\langle\, \mathsf{F}; x \,\rangle\!\rangle$, which is sure to exist in such a structure. A simple argument shows that p cannot have a fixed point. Suppose that $\sigma = \neg\langle\!\langle\, \mathsf{F}; \sigma \,\rangle\!\rangle$. If σ is a fact then so is $\neg\langle\!\langle\, \mathsf{F}; \sigma \,\rangle\!\rangle$ and so $\langle\!\langle\, \mathsf{F}; \sigma \,\rangle\!\rangle$ is not a fact. This contradicts the assumption that F represents Fact and so σ must not be a fact. It follows that $\langle\!\langle\, \mathsf{F}; \sigma \,\rangle\!\rangle$ is not a fact and so that $\neg\langle\!\langle\, \mathsf{F}; \sigma \,\rangle\!\rangle$ is a fact. But $\sigma = \neg\langle\!\langle\, \mathsf{F}; \sigma \,\rangle\!\rangle$ and so we have shown that σ both is and is not a fact—a contradiction.

In the structures used by Plotkin, every abstract has a fixed point, and Fact cannot be represented unless at the expense of classical logic. Plotkin's uses Frege structures (Aczel 1980), which are constructed from models of the untyped λ-calculus in which fixed-point combinators such as Y exist. Similar problems arise in structures with internalized application, because they also contain fixed-point combinators. Even without such combinators fixed points will exist in any anti-founded fact structure. For example, the description $\langle x, \neg\langle\!\langle\, r; x \,\rangle\!\rangle\rangle$ is satisfied in any such structure, and this solution is a fixed point of p.

These considerations illustrate an important trade-off in the construction of models for Situation Theory. There are two ways in which we may measure

[37] In the most general case, an internal relation may be said to represent a partial predicate of indexed sets. First specify a class I of indices. Say that an *I-set* is an indexed set $\{x_j\}_{j \in J}$ for some $J \subseteq I$; equivalently, it is a function with domain J. Then say that r represents the partial predicate P of I-sets if there is a one-one function g from I to A such that the extension (anti-extension) of P consists of those I-sets $\{f^*(g(j))\}_{j \in J}$ such that $\langle\!\langle\, r; f \,\rangle\!\rangle$ is a fact (infon but not a fact) and $J = g^{-1}\,\mathrm{dom}\, f^*$.

the power of our structures as modelling tools. On the one hand, we may see which operations are possible, desiring closure under such operations as abstraction, infon-formation, and restricted abstraction. On the other hand, we may see which external entities are internalized, which functions and relations are represented, whether application can be internalized, and so on. We have seen in Section 3.2 that structures in which many descriptions are satisfied—in particular, anti-founded structures—are very useful for ensuring closure properties under structural operations. Unfortunately, the above results show that if too many descriptions are satisfied, then some predicates may be not be representable. The goals of structural closure and internalization are in conflict. Where the boundary lies is still far from clear.

4.5 Situation Structures

Sections 4.1 to 4.4 were primarily concerned with what we have called the Russellian approach to the relationship between information and truth. Now we turn to the alternative, Austinian approach.

Consider again the example of Smith's involvement in a riot. The straightforward analysis of the statement that Smith panicked and ran is that it expresses the information $\langle\!\langle\, \mathsf{panic}; \mathsf{Smith} \,\rangle\!\rangle \wedge \langle\!\langle\, \mathsf{run}; \mathsf{Smith} \,\rangle\!\rangle$. The puzzle is to identify the property of this infon that makes the statement true. To be a fact an infon must be saturated, but this infon is not. So, if the truth of the statement depends on its expressing a fact, the straightforward analysis cannot be correct. The Russellian solution is to claim that the statement, if it is unambiguous, expresses a saturated infon whose additional roles are filled by various contextually determined parameters. The task of identifying and cataloging the various contextual parameters is a matter for further semantic analysis.[38]

On the Austinian account, the gap between information and truth is attributed to the neglect of a fundamental component of descriptive statements: the described situation. The basic idea is that a statement using sentence S involves reference to a particular situation s, and that the truth-value of the statement depends on whether or not the S is a correct description of s. In our example, the described situation is a particular course of events twenty-six years ago in which Smith panicked and ran. The sentence 'Smith panicked and ran' is therefore a correct description of the described situation, and so the statement is true.

The same sentence may be used to describe other situations. For example, later in the narration we may be told of a second riot on the same day in which Smith participated but at which he stood his ground. This is obviously *not*

[38] An alternative is to explain semantic relationships directly in terms of information expressed rather than truth. This approach is summed up nicely by Barwise's slogan 'Information conditions not truth conditions' and has become a core idea in Dynamic Semantics (see the Chapter 10, this volume).

correctly described by 'Smith panicked and ran', and so a statement using this sentence to describe the later situation would be false.

Another possibility is that the described situation is one that encompasses many different events, some correctly described and others not. In this case the statement may be genuinely ambiguous, or even unintelligible. Imagine an account of all the riotous events of 1970, including both the riot in which Smith panicked and the one in which he did not, and itemized by the people involved. If Smith's entry is the single sentence 'Smith panicked and ran', then it is difficult to assign any definite truth-value.

The different possibilities are accounted for by introducing a new predicate, 'supports', to refer to the relation holding between a situation and the information expressed by a true statement about it. In other words, a statement expressing the information σ made about a situation s is true if and only if s supports σ. The straightforward account of information is now sufficient. In our example, the statement that Smith panicked and ran expresses the information

$$\sigma \quad = \quad \langle\!\langle\, \mathsf{panic}; \mathsf{Smith} \,\rangle\!\rangle \wedge \langle\!\langle\, \mathsf{run}; \mathsf{Smith} \,\rangle\!\rangle.$$

and is true because the described situation—the course of events during the first riot—supports σ. A statement made about the second riot using the same sentence expresses the same information σ but is false because the situation to which it refers does not support σ.

The example shows why a simple distinction between fact and non-fact is not sufficient to accommodate the referential account. Two statements, one true and the other false, express the same information. Clearly no property of the information expressed will suffice to explain the difference in truth-value.

Criticism of the Austinian account has centered on two related problems. Firstly, it is difficult explain how reference to situations is achieved. Austin (1961) proposed that a statement is related to the situation it describes by 'demonstrative conventions', which he contrasted with the 'descriptive conventions' relating the statement to what it says about the situation—the information it expresses, in our terminology. In simple present tense statements about one's immediate surroundings, the relation may be akin to that relating demonstratives ('this', 'that', etc.) to their referents; but in general, the appeal to demonstration is clearly inadequate. The problem is especially acute if the described situation is temporally and spatially remote from the the person making the statement.

The determination of the reference of proper names is beset with similar difficulties, and so it is not surprising that solutions to the present problem parallel the familiar moves in the philosophy of language concerning singular reference. The intentions of the person making the statement may be involved; there may be some kind of causal relationship between the statement and described situation; and conventions within a linguistic community may be appealed to; but, in the end, no generally satisfactory account has been given, and this remains a serious lacuna of the theory. Special problems arise in the present case because the context of a statement has considerable importance in determining the situation to which it refers. A past tense sentence, such as the one we have been discussing, may well be referentially ambiguous when uttered out of context, attaining a definite reference only in the context of a discourse or narrative text (see Seligman and ter Meulen (1993)).

The first problem is compounded by the second: that it is unclear to what a statement refers when reference is achieved. The word 'situation' is intended to be neutral between 'event' and 'state', so that it makes sense to speak both of the situation in Bosnia, meaning the war (an event or course of events), and the situation on Capitol Hill, meaning the state of deadlock over the budget. The theoretical uniformity of this terminology is bought at a price. Questions about the identity of situations appear much more pressing than the traditional (and equally baffling) metaphysical questions about the identity of material objects, persons and members of a host of other ontological categories.

A short, bold answer to the question of the identity of situations is that distinct situations support different infons. Contraposing we get the

> *Principle of Extensionality:* if for each infon σ, s supports σ iff s' supports σ, then $s = s'$.

If we accept this principle then we may think of situations, and model them, as structured objects whose structure is determined by the infons they support. This is the basis of a model of situations using the theory of structural relations. We introduce the structural relation Sit^1 of being a situation, and $\mathsf{HoldsIn}^2$, which holds between an infon and a situation just in case the situation supports the infon.

Definition 4.6 An extensional structure $\mathfrak{S}$ of type $[\mathfrak{A}, \mathsf{HoldsIn}^2, \mathsf{Sit}^1;]$ is a *situation structure* if $\mathfrak{A}$ is an infon structure and the following conditions are satisfied:

1. if $\mathsf{HoldsIn}(\sigma, s)$ then $\mathsf{Inf}(\sigma)$ and $\mathsf{Sit}(s)$, and
2. if $\mathsf{Sit}(s)$ then s is an ordinary object of $\mathfrak{A}$.

The elements of Sit^* are called *situations.* We say that *s supports σ*, written $s \models \sigma$, if $\mathsf{HoldsIn}(\sigma, s)$. An *ordinary* object of $\mathfrak{S}$ is an ordinary object of $\mathfrak{A}$ that is not a situation.

The move from fact structures to situation structures is a generalization in two directions. First, the non-relational property of being factual is replaced by the relational property of holding in, or being supported by, a situation. Second, this property may be had by infons that are unsaturated. In other words, situations are *partial* in two respects: they need not support all of the facts, and the infons they do support may be only unsaturated parts of facts.

Any fact structure $\mathfrak{F}$ may be used to define a situation structure $\mathfrak{F}^{\#}$ by adding a new element s_Σ for each set Σ of facts, and defining: $\mathsf{HoldsIn}(\sigma, s_\Sigma)$ iff

σ is in Σ. We say that a situation structure is *standard* if it is isomorphic to $\mathfrak{F}^{\#}$, for some standard fact structure $\mathfrak{F}$.

Standard situation structures depart from fact structures only in the first respect. Non-standard situation structures may be constructed by modelling situation using sets of infons taken from an infon structure with unsaturated infons. This is a very general method if we work with a theory of sets satisfying AFA, so that no problems are encountered with situations supporting infons about themselves. There are many examples, including Barwise (1987), Fernando (1991), Westerståhl (1990), Aczel (1990), Barwise and Cooper (1991), and Lunnon (1991). Here, as previously, we use the theory of structural relations to discuss these models at general level, without going in to the details of the constructions proposed in the literature.

4.6 Parts of the World

How smaller situations relate to larger situations is perhaps the most controversial issue in Situation Theory. To see why, consider the two riots, riot_1 and riot_2, in which Smith participated, and let riots be the larger situation, encompassing all of Smith's riotous activities that day. We have seen that the infon $\sigma = \langle\!\langle\, \mathsf{panic}; \mathsf{Smith}\, \rangle\!\rangle$ is supported by riot_1 but not by riot_2. Indeed, riot_2 supports the infon $\neg\sigma$. Clearly, the larger situation riots cannot support both σ and $\neg\sigma$, and so it does not support all of the infons supported by its parts.

If these intuitions are to be accepted, it follows that the part-whole relation between situations is not correctly modelled by the $\sqsubseteq$-order, obtained as an instance of the general theory of partiality in Section 3.3. According to this account, $s_1 \sqsubseteq s_2$ iff all infons holding in s_1 also hold in s_2. So, if $\mathsf{riot}_1 \sqsubseteq \mathsf{riots}$ and $\mathsf{riot}_2 \sqsubseteq \mathsf{riots}$ then riots would have to support the contradictory infons σ and $\neg\sigma$.

This difficulty may be resolved by supposing that riots supports the infons $\langle\!\langle\, \mathsf{panic}; \mathsf{Smith}, \mathsf{10{:}15am}\, \rangle\!\rangle$ and $\neg\langle\!\langle\, \mathsf{panic}; \mathsf{Smith}, \mathsf{2{:}30pm}\, \rangle\!\rangle$, which contain the infons σ and $\neg\sigma$ as parts. This suggests the following definition:

Definition 4.7 s_1 is *a part of* s_2, written $s_1 \trianglelefteq s_2$, if

$$\forall\sigma \text{ if } s_1 \models \sigma \text{ then } \exists\tau\ (s_2 \models \tau \text{ and } \sigma \sqsubseteq \tau).$$

That $\mathsf{riot}_1 \trianglelefteq \mathsf{riots}$ and $\mathsf{riot}_1 \trianglelefteq \mathsf{riots}$ is consistent with the assumption that $\mathsf{riot}_1 \models \sigma$ and $\mathsf{riot}_1 \models \neg\sigma$, and so we see that riot_1 and riot_2 are compatible after all. Note that the $\trianglelefteq$-order satisfies the

Principle of Inclusion: if $s_1 \sqsubseteq s_2$ then $s_1 \trianglelefteq s_2$,

but not necessarily its converse, which is equivalent to the

Principle of Persistence: if $s_1 \trianglelefteq s_2$ and $s_1 \models \sigma$ then $s_2 \models \sigma$

In every standard situation structure, Persistence is satisfied, making $\trianglelefteq$ equivalent to $\sqsubseteq$, but this is only because all infons are saturated in standard structures. If we consider situations modelled by sets of possibly unsaturated infons, then Persistence may be violated, as in the example given above. A consequence of Extensionality and Persistence is the

Principle of Anti-symmetry: if $s_1 \trianglelefteq s_2$ and $s_2 \trianglelefteq s_1$ then $s_1 = s_2$,

which makes $\trianglelefteq$ a partial order. This is a very desirable consequence because it ensures that the (partial) operation of joining two situations together is uniquely defined, as intuition demands it should be.

Anti-symmetry may be violated in non-standard situation structures that do not satisfy Persistence. For example, suppose that σ_1 is the information that Kara eats and σ_2 is the information that Kara eats trout, so that $\sigma_1 \sqsubseteq \sigma_2$ but $\sigma_1 \neq \sigma_2$. If s_2 is the (necessarily unique) situation that supports only σ_2, and s_1 is the situation that supports only σ_1 and σ_2, then $s_1 \trianglelefteq s_2$ and $s_2 \trianglelefteq s_1$ but $s_1 \neq s_2$.

Even if one does not accept the Principle of Persistence, one might reject such counterexamples to Anti-symmetry as artificial. One way of ensuring that $\trianglelefteq$ is a partial order without a commitment to Persistence is to argue that all situations are *relatively saturated*: if s supports σ then there is no unsaturated infon $\tau \sqsubseteq \sigma$ also supported by s, although such an infon may be supported by some other situation.

The discussion here must remain inconclusive. Definition 4.7 is by no means universally accepted as the correct definition of the part-whole relation between situations, and a number of theorists have either stuck with the $\sqsubseteq$-order or taken the part-whole relation to be a new primitive relation, not determined by the structural properties of situations, but perhaps constrained by one or more of the principles discussed above. In what follows we shall assume that $\trianglelefteq$ is used, but the questions we pose must be faced by all of the rival accounts.

Two situations are *compatible* if they have an upper bound in the $\trianglelefteq$-ordering. We shall consider several interpretations of compatibility below. But first some more definitions. We say that a situation s is *maximal* if for each situation s', if $s \trianglelefteq s'$ then $s = s'$. A set of situations may be *joined* if they have a least upper bound in the $\trianglelefteq$-ordering, called the *join* of S. The $\trianglelefteq$-ordering in standard situation structure is relatively simple. Every standard situation structure satisfies the following conditions:

S1 any two situations are compatible,

S2 every set of pairwise-compatible situations may be joined, and so

S3 every set of situations may be joined.

In standard situation structures there is a simple relation between situations and facts: $\mathsf{Fact}(\sigma)$ iff $\exists s.s \models \sigma$. Moreover, if the class of facts is a set, there is a unique maximal situation w of which every other situation is a part and so $\mathsf{Fact}(\sigma)$ iff $w \models \sigma$. Non-standard situation structures are not so constrained, but to interpret them we must be able to make sense of incompatible situations. We shall consider two interpretations of compatibility that give rise to very different ways of understanding Situation Theory as a whole.

Compossibility interpretation Maximal situations are completely saturated, in the sense that information about the filling of roles in infons is specified, and there is no compatible way of adding further infons. This suggests the interpretation of maximal situations as possible worlds. Under such an interpretation we would insist that

S4 every situation is part of some maximal situation.

Consequently, compatibility becomes compossibility: two situations are compossible iff they are part of the same possible world. Under this interpretation, one maximal situation w_a must be distinguished as the *actual world*, and the other incompatible maximal situations are ways the world might have been. There is a sharp metaphysical distinction between those situations that are part of w_a and those that are not; the former are *actual situations*, the latter are merely *possible situations*.

The compossibility interpretation allows one to introduce many of the tools of the theory of possible worlds.[39] For example, standard analyses of modality and conditionals may be given. Such analyses tend to run into difficulties because of the partiality of situations; we shall not go into the details here.

Relativist interpretation Incompatible situations are regarded as embodiments of different perspectives on the same, actual world. For example, suppose we are facing each other across the dinner table, so that for you the salt is to the left of the pepper. The situation s_{you} concerning the arrangement of objects on the table from your perspective supports the infon $\langle\!\langle\, \mathsf{LeftOf}; \mathsf{salt}, \mathsf{pepper} \,\rangle\!\rangle$, whereas the situation s_{me} from my perspective supports $\langle\!\langle\, \mathsf{LeftOf}; \mathsf{pepper}, \mathsf{salt} \,\rangle\!\rangle$. On a relativist conception, these situations may be incompatible because they capture the structure of the world from essentially different perspectives. No one can see the objects on the table from a perspective from which the salt is to the left of the pepper and the pepper is to the left of the salt.

Whether or not the situations are incompatible depends on the substantial issue as to whether the relation expressed by the phrase ‘to the left of’ is a perspectival one or not; and, more generally, whether there are such things as perspectival facts. Without prejudging the outcome of metaphysical dispute

[39] Our construction of standard situation structures may be modified in a fairly obvious way to build standard ‘modal’ situation structures from first-order Kripke structures, by generating situations from sets of compossible facts.

on these matters, we can see that the present framework is able to make the pertinent distinctions.

For example, the perspectives on either side of the dinner table may be made compatible by appealing to a hidden parameter. If there are infons $\langle\!\langle\ \mathsf{LeftOf}; \mathsf{salt}, \mathsf{pepper}, \mathsf{me}\ \rangle\!\rangle$ and $\langle\!\langle\ \mathsf{LeftOf}; \mathsf{salt}, \mathsf{pepper}, \mathsf{you}\ \rangle\!\rangle$, in which a role for an egocentric frame of reference is filled by me and you respectively, then there could be a situation supporting both of these infons without conflict, and of which s_{you} and s_{me} would be parts. In this way, we can account for the fact that there are different perspectives in this example (s_{me} and s_{you} have no upper bound in the $\sqsubseteq$-ordering), while showing how the two perspectives can be reconciled using hidden parameters.[40]

4.7 Logic in Situation Structures

The use of consequence relations to capture logical structure (Section 4.2) may also be applied to situation structures. Given a consequence relation $\vdash$ on a class C, a class $X \subseteq C$ *respects* $\vdash$ if the partition $\langle X, C - X\rangle$ is not in $\vdash$. This is a slight generalization of soundness: if $\mathfrak{F}$ is a non-trivial fact structure then $\vdash$ is sound on $\mathfrak{F}$ iff Fact^* respects $\vdash$.

The way in which a logic with consequence relation is applied to a situation structure depends on selecting a class of infons to respect it. There are two obvious candidates: the infons supported by a single situation and the infons supported by some situation or other.

Definition 4.8 Let $\mathfrak{S}$ be a situation structure and $\vdash$ a consequence relation on the class of infons of $\mathfrak{S}$. For each situation s, let F_s be the set of infons supported by s. $\vdash$ is *locally sound* on $\mathfrak{S}$ if F_s respects $\vdash$, for each situation s. Let $\mathsf{F}_{\mathfrak{S}}$ be the class of infons supported by some situation in $\mathfrak{S}$. $\vdash$ is *globally sound* on $\mathfrak{S}$ if $\mathsf{F}_{\mathfrak{S}}$ respects $\vdash$.

$\mathfrak{S}$ is *locally/globally classical* if there is a classical consequence relation that is locally/globally sound on $\mathfrak{S}$. From Section 4.2, we know that $\mathfrak{S}$ is locally classical if and only if it satisfies the following conditions:

1. $s \models \bigwedge \Sigma$ iff $s \models \sigma$ for each $\sigma \in \Sigma$,
2. $s \models \bigvee \Sigma$ iff $s \models \sigma$ for some $\sigma \in \Sigma$, and
3. $s \models \neg\sigma$ iff $s \not\models \sigma$.

For most applications these conditions are too strong. In particular, they imply that for each infon σ, every situation supports either σ or $\neg\sigma$, and this is quite contrary to the spirit of Situation Theory. Typically, the consequence relation chosen for local soundness is that of a partial logic such as Kleene's strong

[40] For more discussion of perspectives, see Barwise (1989c) and Seligman (1991).

three-valued logic. Another possibility, little studied in the literature but very appropriate to epistemological interpretations of Situation Theory (see Schultz 1996), is that of taking a locally intuitionistic situation structure: structures for which the consequence relation for intuitionistic logic is locally sound. Both of these possibilities are consistent with the situation structure being globally classical.

There are many other ways of evaluating a consequence relation in a situation structure that are somewhere between the local and the global. For example, in order to make sense of classical reasoning on the compossibility interpretation, we may restrict the condition of local soundness to maximal situations. One could also involve the $\trianglelefteq$-order in a more direct way. For example, for each situation s, let E_s be the class of infons σ such that for each $s' \trianglerighteq s$ there is a $s'' \trianglerighteq s'$ such that $s'' \models \sigma$. A situation structure $\mathfrak{S}$ *eventually sound* if E_s respects $\vdash$, for each situation s in $\mathfrak{S}$. An attractive combination is for a situation structure to be both locally intuitionistic and eventually classical.

Space prohibits a more comprehensive treatment of logical matters, so we shall close this section by touching on a question that brings together logical and order-theoretic considerations: how many situations are there? Logical considerations tend to restrict the number of situations, eliminating arbitrary combinations of infons that do not respect the local consequence relation. Consideration of the $\trianglelefteq$-order has the opposite tendency: we are driven to claim the existence of the joins of existing situations or of parts that support some subset of the infons supported by an existing situation.

For example, if $\mathfrak{S}$ is a situation structure on which $\vdash$ is locally sound, then an example of a fairly generous existence principle compatible with logical considerations is the following: if Σ is a set of infons that respects $\vdash$ and there is a situation s such that for each σ in Σ there is a τ supported by s for which $\sigma \sqsubseteq \tau$, then there is a situation that supports all and only the infons in Σ. This principle can be used to establish the existence of joins of compatible situations. If s_1 and s_2 are compatible situations then there is a situation s of which both are parts. Let Σ be the smallest set of infons satisfying the conditions of the existence principle and such that for each infon τ supported by either s_1 or s_2 there is a infon σ in Σ such that $\tau \sqsubseteq \sigma$. Then the join of s_1 and s_2 is the situation supporting all and only the infons in Σ.

4.8 Propositions

On the Austinian account, the truth-value of a statement is determined by whether the described situation supports the information expressed. Truth is a property of statements but not of the information they express, which may be supported by some situations and not by others. To formulate a theory of truth it is therefore necessary to introduce a new kind of entity to which the property of being true may be predicated. Statements are the obvious choice, but statements have many contingent properties that are irrelevant to their truth value, including many of the details of their production, and—most seriously—the fact that they have to be made to exist. Statements not made do not exist and so cannot be true.

This problem is solved by introducing a new kind of abstract object: propositions. Despite its long philosophical pedigree, including an implied role as the object of propositional attitudes, the word 'proposition' is used in Situation Theory purely as a term denoting that by virtue of which a statement has the truth-value it does.

The proposition associated with a given statement is its *propositional content.* On the Russellian account, this is the infon expressed by the statement—fully saturated, possibly by hidden, contextual parameters—and so a *Russellian proposition* is just an infon. On the Austinian account, the truth of a statement is determined by the situation s it describes and the (possibly unsaturated) information σ it expresses. Its propositional content is the *Austinian proposition* $(s \vDash \sigma)$, an abstract entity that combines the two components in a structurally extensional manner: if $(s_1 \vDash \sigma_1) = (s_2 \vDash \sigma_2)$ then $s_1 = s_2$ and $\sigma_1 = \sigma_2$.

Truth is taken primarily to be a property of propositions: a statement is true if and only if its propositional content is true. Thus a Russellian proposition σ is true if and only if σ is a fact. The theory of Russellian propositions is therefore just the theory of facts, and so we shall consider them no further.

An Austinian proposition $(s \vDash \sigma)$ is true if and only if $s \vDash \sigma$. Austinian propositions were studied extensively by Barwise and Etchemendy (1987) and King (1994), and used to give an analysis of various semantic paradoxes. The Austinian proposition $(s \vDash \sigma)$ is usually modelled as the ordered pair $\langle s, \sigma \rangle$. It is easy to see how to characterize these models using structural relations, and so we shall skip over the details.

Austinian propositions are true or false absolutely but (Austinian) infons may only be evaluated relative to a situation. The contrast suggests an immediate generalization of Austinian propositions to include similarly 'absolute' propositions. For example, whether or not a given object σ is an infon is structurally determined and so does not depend on the situation in which the claim is evaluated. We introduce a new proposition $(\sigma\colon \mathsf{Inf})$ that is true if and only if $\mathsf{Inf}(\sigma)$. Although no typographical distinction is made, it is important to realize that the two occurrence of 'Inf' in the previous sentence are grammatically distinct. The latter refers to a structural (meta-theoretic) relation, whereas the former refers to a new abstract object: the *type* of infons.

Extending this idea to other structural relations, we write and $(\alpha\colon R)$ for the proposition that the objects in the sequence α stand in the structural relation R. The simple Austinian proposition $(s \vDash \sigma)$ is identified with the proposition $(s, \sigma\colon \mathsf{HoldsIn})$. This is modelled in a situation structure extended with structural relations Seq, Type, Prop, and a non-structural relation True, having the following interpretations:

$\mathsf{Seq}(\alpha, p)$	α is a sequence of objects in the basic proposition p
$\mathsf{Type}(T, p)$	T is the type in the basic proposition p
$\mathsf{Prop}(p)$	p is a proposition
$\mathsf{True}(p)$	p is true.

The key to modelling propositions and types is to observe that the axioms governing the structure of propositions are exactly the same as those governing facts, with the following relabelling of primitive relations:

$\mathsf{Seq} \rightarrow \mathsf{Ass}$
$\mathsf{Type} \rightarrow \mathsf{Rel}$
$\mathsf{Prop} \rightarrow \mathsf{Inf}$
$\mathsf{True} \rightarrow \mathsf{Fact}$

With this in mind, we make the following definition.

Definition 4.9 An extensional structure $\mathfrak{P}$ of type

$$[\mathfrak{A}, \mathsf{Seq}, \mathsf{Type}, \mathsf{Prop}; \mathsf{Approp}, \mathsf{True}]$$

is a *proposition structure* if the result of relabelling the primitive relations according to the above scheme is a fact structure in which each structural relation in $\mathfrak{A}$ is represented.

The requirement that each structural relation of $\mathfrak{A}$ is represented is exactly what is needed. For example, if $\mathfrak{A}$ is a situation structure then it will have a binary structural relation $\mathsf{HoldsIn}$, and so there must be a type, also denoted by '$\mathsf{HoldsIn}$', roles $\mathsf{HoldsIn}_1$ and $\mathsf{HoldsIn}_2$, and for each situation s and infon σ, an appropriate assignment f such that $\mathsf{Fun}(\mathsf{HoldsIn}_1, \sigma, f)$ and $\mathsf{Fun}(\mathsf{HoldsIn}_2, s, f)$, and a proposition p such that $\mathsf{Seq}(f, p)$ and $\mathsf{Type}(\mathsf{HoldsIn}, p)$, and $\mathsf{True}(p)$ iff $\mathsf{HoldsIn}(\sigma, s)$. This shows that $\mathfrak{P}$ contains all the Austinian propositions.

A proposition is *basic* if it has a type. Compound propositions, such as conjunctions, disjunctions and negations may be modelled in the same way as compound infons. In a Lambda structure there may also be *proposition-abstracts* of the form $\lambda x.p$ for each proposition p. [41] The structure is then *type-closed* if for each proposition-abstract T and appropriate assignment f, there is a proposition $(f\colon T)$. Moreover, when these higher-order propositions exist, they should satisfy the

> *Principle of β-equivalence for Propositions*: if $(f\colon T)$ is a proposition whose type T is a proposition-abstract then $\mathsf{True}((f\colon T))$ iff $\mathsf{True}(T^*(f))$.

A common use of abstraction is the formation of 'situation types'. A *situation type T* is a type with a single argument role that forms a proposition when that argument role is assigned a situation. A situation s is *of type T* if

[41] The literature has a rich variety of notations for abstraction over propositions. In addition to the notations for abstraction mentioned earlier, '$[x|p]$' is quite common.

$(s{:}T)$ is a true proposition. In a Lambda proposition structure, we can form the situation type $\lambda x.(x \models \sigma)$ for any infon σ. β-equivalence for Propositions ensures that s is of type $\lambda x.(x \models \sigma)$ just in case $s \models \sigma$. A theoretic advantage of working with situation types and complex propositions is that it is possible to capture complex conditions on situations in a way that is independent of the structure of information. For example, suppose we are working with a model in which there is no conjunction of infons, and we need to internalize a condition of supporting both of the infons σ and τ. This can be done with the complex situation type $\lambda x.(x \models \sigma) \wedge (x \models \tau)$. A situation is of this type if and only if it supports both σ and τ.

Typically, situation theorists have proposed that the logic of propositions is classical.[42] This is compatible with any logic on infons, but a favourite combination is a classical proposition structure whose underlying situation structure locally respects a partial logic, such as Kleene's strong three-valued logic. To see that this is a coherent combination, note that the proposition $(s \models \sigma)$ is either true or not, and so in a classical proposition structure one of the propositions $(s \models \sigma)$ or $\neg(s \models \sigma)$ is true. If a situation s fails to support σ then $\neg(s \models \sigma)$ is true, but this does not imply that s supports $\neg\sigma$ unless the underlying situation structure is locally classical.

Another line of thought suggests that logical consequence is primarily a relation between propositions—they are the bearers of truth, after all. Some models of Situation Theory take infons to be basic infons only (possibly with complex relations) and model all logical combinations as yielding complex propositions. Whatever course is taken, the machinery of Section sec.structural-relations can be adapted to provide the necessary structures and the approach to logic in fact structures (Section sec.logic-on-facts) can be applied to yield an appropriate definition of consequence.

The account of propositions summarized here suggests a two-storey residence for logical structure.

[42] In other words, the class of true propositions respects a classical consequence relation on the class of propositions.

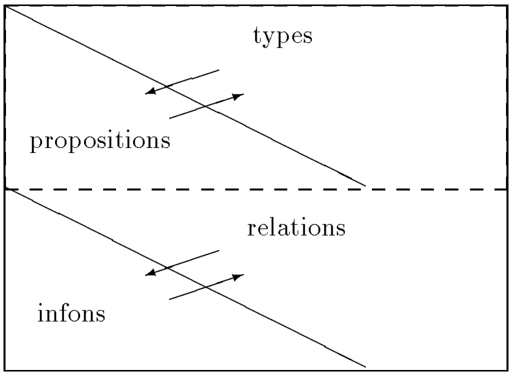

On the ground floor there are infons. Basic infons are made up of assignments and relations; compound infons are structured objects constructed from basic infons; and more complex relations are formed from infons by abstraction.

On the upper floor there are propositions. Basic propositions are made up of assignments and types; compound propositions are structured objects constructed from basic propositions; and more complex types are formed from propositions by abstraction.

The two floors may differ in matters of truth and factuality. Typically, the ground floor is *situated.* An infon needs a situation to determine if it is factual or not; likewise, the extension of a relation may vary from situation to situation. The upper floor is *absolute*, the lower *relative.* Propositions are either true or not, and types have a fixed extensions. The floors are linked by the relationship between a structural relation R on the ground floor and the type R^* that represents it on the upper floor.

Like all metaphors involving a duality, the two-story picture raises a question of redundancy: can we make do with just one of the two, and if so, which one? It should be clear that the notion of a proposition structure could serve as the basis for the whole theory. We could re-express the theory of infons and situations using types, with axioms involving constants referring to the internal representations of the structural relations Inf, Sit, and so on.

In the other direction we have an intriguing possibility. We have already observed that the same axioms govern infons and propositions, but so far we have regarded them as distinct kinds. The possibility of models in which the proposition $(\alpha\colon t)$ is *identical* to the infon $\langle\!\langle\, t, \alpha \,\rangle\!\rangle$ is open. In such a model, we would have to decide how the truth of propositions is related to their support by situations. Much would depend on whether we adopt a compossibility or a relativist interpretation of compatibility. On the relativist interpretation, there is no privileged perspective from which to judge the truth of perspective-independent truths, but we may require that they be recognized as true from all perspectives. This is in conflict with the partiality of situations, but the conflict can be overcome by resorting to the notion of eventual truth, introduced in Section 4.7, and claiming that a proposition is true if it is eventually supported by every situation.

On the compossiblity interpretation, we could choose whether to interpret truth of a proposition as necessary truth—truth in all possible worlds—or as truth in the actual world. As we have seen, propositions are used to capture structural relationships, such as Infand Rel, but also HoldsIn. If σ is an infon and r a relation, then surely they are necessarily so. It is less easy to evaluate the status of true Austinian propositions. On the one hand, they are true solely in virtue of structural relations holding between situations and the infons they support. [43] On the other hand, they are responsible for the truth of contingent statements, like 'Raymond cooked an omelette'. This points toward another fundamental contention. Allegiance to the the Principle of Extensionality suggests that a situation's support of infons is essential: if s_1 supports σ and τ and no other infons, and s_2 supports only σ, then it could not be that s_1 failed to support τ, for then s_1 would be s_2, by Extensionality.

This and many other important metaphysical issues raised by Situation Theory have yet to be debated seriously. In part this is because the technical tools have been slow in development. In a theory such as this, inconsistency lurks around every corner (as Plotkin (1991) has shown), and it is difficult to tell if a given list of axioms is consistent. For this reason, among others, research has focussed on the construction of models.

Although many models of Situation Theory have been presented, it is in the nature of model-building that all decisions be made: the result is a concrete structure that represents the world in a particular way. Yet in Situation Theory, the decisions turn on significant philosophical issues that have yet to be considered in enough detail for us to be sure of the battle lines, let alone the outcome.

Those building models in this philosophical vacuum either have taken a particular line, contested by other proponents of Situation Theory, or have steered clear of the more contentious issues. Barwise's solution (1989a) was to give a list of 'branch points' at which critical decisions must be made. This is all very well, but it leaves the critical observer of Situation Theory wondering what the theory *is*.

It is our hope that the model-construction tools are now sufficiently well understood to allow the looser, more abstract approach adopted in this chapter to be used to frame the important questions and sharpen the lines of future philosophical debate.

[43] They are 'structurally determined'—see Cooper (1991).

4.9 Constraints and the Flow of Information

Situations are user-friendly. Unlike worlds, situations may contain a limited amount of information, an amount that a person may acquire with limited resources. But information is promiscuous. The limited information acquired in perception breeds in the fertile soil of a person's mind, and in the cognitive organization of any other well-adapted animal and machine. Information about one's immediate environment carries further information about the more remote and hidden parts of the world. The information that the sky is darkening may be acquired by visual perception, but the further information that a storm is coming, is not. The latter is not information about the immediate environment, and so not part of the information supported by that situation. We say that a situation s *carries* an infon σ, written ($s \Vdash \sigma$), if σ follows somehow from the information supported by s but is not necessarily supported by s; indeed, it may be information about a different situation altogether.

The responsibility for this 'flow of information' (Dretske, 1981) is placed on the existence of law-like relationships between situations called 'constraints'. Barwise and Perry (1983) placed considerable importance on the role of constraints. The appeal to constraints is central to the claim that Situation Theory is a viable foundation for semantics, significantly different from Montague's. It also lies behind the supposed affinity between Situation Theory and Gibsonian psychology and the approach to cognition urged by Barwise (1986) in his debate with Fodor (1985). More recently, constraints form an essential part of the theory of incremental information developed by Perry and Israel(1990, 1993). And, constraints are essential to applications of Situation Theory to semantics and AI, of the kind considered by Glasbey (1994,1996) and Cavedon (1995,1996).

Despite its importance, and considerable research effort, the theory of constraints has remained elusive. Recent developments (Seligman 1991, Barwise 1991, Barwise and Seligman 1993, Barwise and Seligman, to appear) suggest that a quite different outlook is required, one that has not yet been integrated with other parts of Situation Theory. For this reason, we content ourselves here with a exposition of some of the basic problem.

What are Constraints? Constraints have be categorized roughly as follows:

1. *Necessary Constraints*, including taxonomic relations between properties like 'moles are mammals', appropriateness conditions like 'sipping involves sipping something', incompatibility restrictions between properties like 'red things aren't green', or even between individuals like 'Kriesel is not Keisler', and mathematical constraints like '$5 + 7 = 12$';

2. *Conventional Constraints*, including linguistic rules, syntax and semantics, the rules of backgammon or the convention of driving on the right side on the road;

3. *Nomic Constraints*, including all laws of natures, both the common place—such as the law that unsupported coffee mugs fall to the floor, and that the radio works if you hit it—and the more esoteric laws about chemical valency or electromagnetism. Most notoriously, 'smoke means fire' expresses a nomic constraint.

4. *Meta-theoretic Constraints*, including all those laws that form a part of Situation Theory itself, such as the law that if $s \models \sigma$ then s is a situation and σ is an infon.

This list is haphazard, and there is little convergence on an answer to the question: what is a constraint? Sentences, relations, rules, and laws are all given as examples. In semantics, constraints have been proposed as the referents of conditional and generic statements (see Barwise 1985, Cavedon, 1995,1996). The common thread is to be found in the epistemological role of constraints, and not their ontological status. A constraint is something that allows someone with information about one situation to gain information about a possibly different situation. It is because moles are mammals that a photograph of Ivor the mole carries the information that Ivor is warm-blooded, at least for anyone who knows than moles are warm-blooded. Likewise, it is the linguistic conventions of English-speaking sailors that ensure that the sentence 'Land ahoy!' carries the information that there land has been sighted.

To model constraints we must decide where to put them in the universe of situation-theoretic entities. The simplest approach is to take constraints to be true propositions. [44] Many constraints concern the dependency between types of situations, and can be modelled by basic propositions of the form $(T_1, T_2 \colon \Rightarrow)$, abbreviated to $(T_1 \Rightarrow T_2)$, whose arguments T_1 and T_2 are situation types. If $(T_1 \Rightarrow T_2)$ is a constraint (true proposition) then we say that T_1 *involves* T_2. Such constraints constrain by satisfying the

> *Principle of Involvement:* if s is a situation of type T_1 and T_1 involves T_2 then there is a compatible situation of type T_2.

The reason for the restriction to compatible situations is clear under the compossibility interpretation: constraints only constrain situations in the same possible world. On the relativist interpretation, it ensures that constraints only apply within and not between irreconcilable perspectives. There may be constraints that relate different perspectives, but they are not of the simple form considered above.

The Principle of Involvement gives a necessary but not a sufficient condition for the existence of a constraint. It is not for Situation Theory to decide what

[44] The choice is grammatically awkward because we talk of the *existence* of a constraint and not its *truth*. Other approaches include: constraints as facts; constraints as relations between situations, infons, or propositions; and constraints as sets of situations, worlds, infons or propositions. Some authors have used formal expressions purported to refer to constraints without attempting to answer the ontological question.

constraints there are in the world. Nonetheless, some sufficient conditions have been proposed, the foremost of which is the

> *Xerox Principle:* if T_1 involves T_2 and T_2 involves T_3 then T_1 involves T_3.

The name is taken from Dretske's (1981) principle about information flow. According to Dretske, the transfer of information is an all-or-nothing affair, with no degradation of quality, such as one would obtain from an ideal copying machine.[45] The link between constraints and information flow is given by the

> *Principle of Information Flow:* s carries the information that σ if s is of some type T that involves $\lambda x.(x \vDash \sigma)$.

These two Principles account for the 'long range' of information flow: if s carries the information that σ and $\lambda x.(x \vDash \sigma)$ involves $\lambda x.(x \vDash \tau)$ then s carries the information that τ. For this reason, the information carried by a situation may far exceed the information it supports.[46]

Other Constraints Let us call the approach to information flow discussed above the *Theory of Involvement.* Before discussing the deficiencies of this theory, we shall consider briefly some other kinds of constraints considered in the literature.

Reflexive Constraints: The Principle of Involvement is too weak to capture the constraints governing the informational dependencies *within* a situation. If a constraint $(T_1 \Rightarrow T_2)$ is *reflexive* then every situation of type T_1 is also of type T_2. For example, the analytic constraint that if something is a square then it is a rectangle, may be considered to be reflexive. In that case, any situation supporting $\langle\!\langle\, \mathsf{square}; a \,\rangle\!\rangle$ also supports $\langle\!\langle\, \mathsf{rectangle}; a \,\rangle\!\rangle$. Reflexive constraints may also be used to internalize the logic of situations (Section 4.7). The local soundness of a consequence relation $\vdash$ can be captured with the condition that if $\Gamma \vdash \Delta$ then $(\wedge_\Gamma \Rightarrow \vee_\Delta)$ is a reflexive constraint, where $\wedge_\Gamma$ is the type of situations supporting every infon in Γ and $\vee_\Delta$ is the type of situations supporting at least one infon in Δ.

General constraints: Let $\mathsf{kiss}(a, b)$ and $\mathsf{touch}(a, b)$ be the types of situation supporting the information that a kisses b and that a touches b, respectively. Then the constraint $(\mathsf{kiss}(a, b) \Rightarrow \mathsf{touch}(a, b))$ captures the dependency between a's kissing b and a's touching b, but not the general relationship between kissing

[45] The Xerox Principle further constrains the distribution of information in the world: if s is of type T_1 then from the Principle of Involvement alone it follows that if T_1 involves T_2 and T_2 involves T_3 then there is a situation s' of type T_3, but not that s' is compatible with s—this requires the Xerox Principle.

[46] The information carried by a situation is supported by some situation, albeit in a more saturated form: if $s \Vdash \sigma$ there there is a situation s' and an infon σ' such that $s \trianglelefteq s'$ and $\sigma \sqsubseteq \sigma'$ (by the Principles of Involvement and Information Flow).

and touching. An advantage of using propositions to model constraints is that the general constraint that kissing involves touching can be modelled as the complex proposition $\forall xy.(\mathsf{kiss}(x,y) \Rightarrow \mathsf{touch}(x,y))$. The logic of propositions together with the Principle of Involvement yield the desired relationship between kissing and touching.

Preclusion: Informational dependencies are not always positive. The information acquired in a given situation may rule out possibilities without indicating which of the remaining possibilities obtains. For example, the information that Helen's birthday is on a weekday precludes its being on a Saturday without determining the day of the week on which it falls. We may model a negative constraint as a proposition of the form $(T_1, T_2 \colon \perp)$, abbreviated to $(T_1 \perp T_2)$, and reported by saying that T_1 *precludes* T_2. The *Principle of Preclusion* is that if s is a situation of type T_1 and T_1 precludes T_2 then there is no compatible situation of type T_2.

The Problem of Specificity The Theory of Involvement tells us that if a situation carries the information σ then there is a situation that supports σ, but it does not tell us which one.

On the Austinian approach to Situation Theory, the situation a statement is about is a very important ingredient in determining the truth of the statement. Comprehension may be severely restricted if one knows that a statement expresses the information σ without knowing which situation the statement is about. We have seen very different statement can be made using the sentence 'Smith panicked and ran', all of which express the same information. Likewise, knowing that a situation carries an infon σ may be of little significance unless one also knows which situation supports σ. On observing the darkening sky, we know more than that *some* situation is stormy; if we did not then this could not serve as a reason to take cover.

One approach to solving this problem is to model constraints as complex propositions. We can capture the dependency between a situation s_1 being of type T_1 and a situation s_2 being of type T_2 by means of a proposition of the form $(s_1 \vDash T_1) \rightarrow (s_2 \vDash T_2)$, in which '$\rightarrow$' is a binary propositional connective, modelled using structural relations in the obvious way. The Principle of Involvement is replaced by the principle that if $p_1 \rightarrow p_2$ and p_1 is true then so is p_2; and the new Principle of Information Flow states that a situation s carries the *proposition* p if there is type T such that $(s \colon T) \Rightarrow p$. Call the result of these modifications the *Propositional Theory of Information Flow.*

The Propositional Theory errs in a direction opposite to that of the Theory of Involvement: it is *too* specific. One can understand a statement very well without knowing exactly which situation it is about, only that it lies within a certain range, or that it is the same situation that was described by an earlier statement. Likewise, the darkening sky does not indicate exactly which situation is stormy.

The middle ground has been explored by Seligman (1990 and 1991), Barwise (1991), and Perry and Israel (1991). Each approach involves the addition of a parameter to capture the dependency between specific situations. Very roughly, we may suppose that some situations are *connected*, while others are not. Then we may modify the Principle of Involvement as follows: if a situation is of type T_1 and T_1 involves T_2 then there is a connected situation of type T_2. Clearly, which situations are 'connected' will vary greatly depending on the nature of the constraint. The connection between the observed situation in which the sky is darkening and the later, stormy situation is quite different from the connection between the statement that Smith panicked and the event of his panicking.

Another consideration that bears on this issue, and that the above accounts fail to address, is the limitations on a person's access to constraints. If every law-like regularity in the world is a constraint, then the world is a very constrained place, and just as 'connected'. Yet a person's information about the world at large is limited, not just by her observations, but also by the constraints on which her knowledge is based. In determining the information carried by a situation, some constraints must be excluded.

Barwise and Perry (1983) adopt the Gibsonian metaphor of *attunement* for this purpose. The ability of a person (animal or machine) to extract information from the environment is determined by the constraints to which that person is attuned. The mere existence of a constraint is not sufficient.

Conditional Constraints The above discussion has proceeded on the assumption that constraints constrain without exception. This is a reasonable assumption for the necessary and meta-theoretic constraints of our original list, and perhaps for some of the nomic constraints concerned with fundamental laws of nature; but for the vast majority of constraints that underpin a person's beliefs and actions, the assumption is less secure. Smoke means fire only if there are no artificial smoke generators nearby; unsupported coffee mugs fall near the surface of the Earth but not in interstellar space; and an utterance of 'Land ahoy!' by a drunk ship's cook may not carry the information that land has been sighted.

Barwise and Perry (1983) introduced a separate category of constraint to account for information that flows only under special circumstances. A *conditional constraint* is a constraint that is only guaranteed to hold only when certain *background conditions* are met. For example, the background conditions of the constraint that smoke means fire must exclude those situations in which there are other sources of smoke nearby; those of the constraint that unsupported coffee mugs fall specify that the mug should be in the gravitational field of the Earth.

Continuing with the strategy of modelling constraints as true propositions, we take conditional constraints to be obtained from unconditional constraint by adding a *background situation type*. Conditional involvements are formed using a ternary type $\Rightarrow_c$ whose arguments are all situation types. The constraint $(T_1, T_2, B{:}\Rightarrow_c)$ is abbreviated to $(T_1 \Rightarrow T_2|B)$, and we say that *under conditions B, T_1 involves T_2*.[47] Such constraints conform to the

> *Principle of Conditional Involvement:* if under conditions B, T_1 involves T_2, then for each situation s of type B, if s is of type T_1 then there is a compatible situation of type T_2.

For example, suppose that B is the type of situation in which smoke does indeed mean fire, and that smoky and firey are the types of smoky and firey situations, respectively. Then there is a conditional constraint of the form $(\mathsf{smoky} \Rightarrow \mathsf{firey}|B)$. By the Principle of Conditional Involvement, if s is a situation of type B and s is of type smoky then there is a situation s' of type firey and s' is compatible with s.

The Xerox Principle must also be modified for conditional constraints. The basic idea is that it should hold under fixed background conditions, but not when background conditions vary.

> *Conditional Xerox Principle:* if $T_1 \Rightarrow T_2|B$ and $T_2 \Rightarrow T_3|B$ then $T_1 \Rightarrow T_3|B$.

This Principle was used by Barwise to explain the failure of inferences of the following form: from 'if φ then ψ' and 'if ψ then θ' infer 'if φ then θ'. The inference goes through only if the background conditions of the constraints denoted by each of the conditional statements are the same. Further connections between conditional logic and Situation Theory are discussed in Barwise (1985) and Cavedon (1995,1996).

The Elusive Background Condition The theory of conditional constraints faces a number of difficult problems. First, it is unclear how the background condition is specified. In psychological applications of constraints, the basic idea is that the background conditions of a constraints are those that hold in the 'environment' of the person (animal or machine) attuned to it, but that may fail to hold outside that environment. The background conditions are preconditions for the success of the organism's behaviour, and failure of a given response may therefore be attributed to the failure of the background conditions, rather than to a failure of attunement *per se*. Unfortunately, this idea has not been successfully incorporated into Situation Theory. In particular, the role of an 'environment' as the fixer of background conditions has proved very difficult to nail down.

In semantics, the determination of background conditions is just as mysterious. One idea is that they are is fixed by the described situation, either as the

[47] Conditional versions of reflexive constraints, general constraints and preclusions can be constructed similarly.

'environment' of that situation or as the situation itself.[48] But the relationship between a statement and the situation it describes (discussed in Section 4.5) is far from clear, and so no substantial progress is made.

A second problem is with the scope of background conditions. Consider the following example. A bathroom is wired so that a switch in the corridor outside may be used to control the light inside. If the switch is down then the light is on, and if the switch is up then light is off. These conditionals may be used by a person outside to know whether the light is on even when the door is closed. Let switch-down, switch-up, light-on, and light-off be the types of situations in which the switch and light are as indicated. The truth of the two conditional statements is attributed to the existence of constraints of the form (switch-down $\Rightarrow$ light-on $\mid B$) and (switch-up $\Rightarrow$ light-off $\mid B$), in which B is intended to capture the conditions under which the electrical circuit functions properly—no shorts, no loose wires, no blown fuses, etc.

Now suppose that Janet is outside the bathroom. The door is closed but she observes a situation s of type switch-down and infers that the light is on inside. Her conclusion that is warranted so long as s is of type B, even if Janet does not know that it is. It is difficult to see how the situation *outside* the bathroom can be of type B, a type that ensures the correct functioning of the wiring *inside*. A situation of type B must involve the whole system, within and without the bathroom walls, and perhaps incorporating components even further afield, such as the fuse box and the power supply. Such a situation may be comprehensive enough to ensure the success of Janet's inference, but it is not the situation she observes in the corridor outside the bathroom. Moreover, if the larger situation is large enough to guarantee that conditions are right, then it is difficult to see how it can fail to support the information that the light is on and the switch is down. In that case, even if Janet were able to observe it, and so obtain the information that the light is on, the conditional constraint would play no part in the process.

The problem of specificity, the need to consider connections between situations, the problems of identifying background conditions and determining their scope, and the need to incorporate a limitation on access suggest that there is a long way to go before Situation Theory is extended with a theory of constraints that satisfies the expectations raised by Barwise and Perry (1983).

Situation Theory, in its contemporary form, provides a rich, fine-grained ontology and the means of framing subtle questions about the the nature of truth and circumstance that could not otherwise be put. But it is a theory of stationary objects, of frozen structure, and has proved ill-suited for accounting for the flow of information. An alternative strategy is to aim for a theory of the flow of information that is not primarily concerned with its structure. The 'dynamic' approach to semantics (Chapter 10, this volume) is an example,

[48] Barwise (1995) takes the background condition of a constraint to be a situation instead of a type of situation.

and perhaps the results of dynamic semantics may lead to a better theory of situations.[49]

5 Guide to the Literature

The bulk of research papers in Situation Theory and Situation Semantics are contained in the collections entitled *Situation Theory and its Applications*, Vols. I to III, published by CSLI. These contain selections of papers presented at a biannual conference. The name of the conference was changed in 1994 to 'Information-Theoretic Approaches to Logic, Language, and Computation'. We have listed in the references a number of papers from these volumes. For an introduction to Situation Semantics one could read Cooper (1991). Other source are the books Fenstad et al. (1987) and Gawron and Peters (1990). The latter uses technically intricate concepts from Situation Theory, like restricted abstraction, to propose a theory of anaphora and quantification. In addition, ter Meulen (1995) studies the representation of temporal information in a related framework, as does Seligman and ter Meulen (1995). Conditionals are addressed in Barwise (1985) and Cavedon (1995, 1996).

Situation semantics has given rise to computational paradigms as well. For example, see Tin and Akman (1993) and Tin, Akman, and Ersan (1995). The general theory has also lead to work in the social sciences; see Devlin and Rosenberg (1993) for the first paper in a series of explorations intended to give a foundation for ethnomethodology.

Key philosophical works connected with Situation Theory are Barwise and Perry (1983), Barwise (1989), Barwise and Etchemendy (1987), and Devlin (1991). Perhaps the closest in spirit to the approach taken here is that of Barwise and Etchemendy (1987); it was the first to consider the circular structures and apply them in a diagnosis of the Liar Paradox. Another work that concerns circularity and the perspective of Situation Theory is Barwise (1987); this paper shows what how to use circular structures to model common knowledge.

The seminal work on non-well-founded sets is Aczel (1988), supplemented more recently by Barwise and Moss (1996). The latter book has less of an emphasis on set theory, and covers many applications of bisimulations to logic and computer science. The role of bisimulation in modal logic is thoroughly explored in de Rijke (1995).

Recent work on anti-foundation has focussed on sorted universes, abstraction and application. See, for example, Aczel and Lunnon (1991), Lunnon (1996) and Aczel (1996). The guiding idea behind these developments is that one

[49] The theory of information flow developed by Barwise and Seligman (1993, and to appear) is directly motivated by the difficulties in giving an account of constraints. Although sharing the philosophical orientation of earlier work in Situation Theory, the mathematical tools are quite different, and there is no concern with the structure of information, or with the ontological issues addressed in this chapter. For these reasons, it lies beyond the scope of this chapter.

should be able to combine methods from set theory and the lambda calculus. The techniques developed by Aczel and Lunnon have been applied directly to Situation Theory by Barwise and Cooper (1991).

The concept of information employed in Situation Theory has developed from that used in Dretske (1981) and Barwise and Perry (1983). Israel and Perry (1990, 1991) have more to say on the subject. The importance of studying information *flow* has been voiced by many writers (see Chapter 10, this volume) and connections to relevant logic and the Lambek calculus have been brought out by Barwise (1993), Restall (1996) and Barwise et al. (1996).

An recently culminating development is the theory of *channels* (Barwise and Seligman, to appear). The goal of this work is a framework for understanding how information flow results from law-like regularities in distributed systems. Key concerns are the modularity of information systems, the limitations imposed by an agent's perspective and the provision for an account of error. Early papers in this direction are Seligman (1991) and Barwise (1991).

References

Aczel, P. (1980). Frege Structures and the Notions of Proposition, Truth and Set. In J. B. et al. (Ed.), *The Kleene Symposion*. North Holland.

Aczel, P. (1988). *Non-Well-Founded Sets*. Number 14 in CSLI Lecture Notes. Stanford: CSLI Publications.

Aczel, P. (1990). Replacement Systems and the Axiomatization of Situation Theory. See Cooper, Mukai, and Perry (1990).

Aczel, P. (1996). Generalized Set Theory. See Seligman and Westerståhl (1996).

Aczel, P., D. Israel, Y. Katagiri, and S. Peters (Eds.) (1993). *Situation Theory and Its Applications, III*, Number 22 in CSLI Lecture Notes, Stanford. CSLI Publications.

Aczel, P. and R. Lunnon (1991). Universes and Parameters. See Barwise, Gawron, Plotkin, and Tutiya (1991).

Austin, J. L. (1961). Truth. In J. O. Urmson and G. J. Warnock (Eds.), *Philosophical Papers*. Oxford: Oxford.

Barwise, J. (1981). Scenes and other situations. *Journal of Philosophy*, 369–97.

Barwise, J. (1985). Conditionals and conditional information. Technical report, CSLI. Reprinted in J. Barwise, 1989, pp. 97–135.

Barwise, J. (1986). Information and Circumstance: A Reply to Fodor. *Notre Dame Journal of Formal Logic 27*(3), 324–338.

Barwise, J. (1987). Three Views of Common Knowledge. In M. Vardi (Ed.), *Theoretical Aspects of Reasoning about Knowledge, II.* Los Altos, CA: Morgan-Kaufmann. Reprinted with changes as "On the Model Theory of Common Knowledge" in J. Barwise, 1989, pp. 201–20.

Barwise, J. (1989a). Notes on branch points in situation theory. See Barwise (1989b), pp. 255–276.

Barwise, J. (1989b). *The Situation in Logic.* Number 17 in CSLI Lecture Notes. Stanford, CA: CSLI Publications.

Barwise, J. (1989c). Situations, Facts, and True Propositions. See Barwise (1989b), pp. 221–254.

Barwise, J. (1991). Information links in domain theory. Technical report, Indiana University Logic Group. Preprint #IULG-91-7.

Barwise, J. (1993). Constraints, Channels, and the Flow of Information. See Aczel, Israel, Katagiri, and Peters (1993), pp. 3–27.

Barwise, J. and R. Cooper (1991). Simple Situation Theory and its graphical representation. See Seligman (1991a). DYANA Deliverable R2.1.C.

Barwise, J. and R. Cooper (1993). Extended Kamp Notation. See Aczel, Israel, Katagiri, and Peters (1993).

Barwise, J. and J. Etchemendy (1987). *The Liar: An Essay on Truth and Circularity.* New York: Oxford University Press.

Barwise, J. and J. Etchemendy (1990). Information, Infons, and Inference. See Cooper, Mukai, and Perry (1990), pp. 33–78.

Barwise, J., D. Gabbay, and C. Hartonas (1994). On the Logic of Information Flow. Technical report, Indiana University Logic Group.

Barwise, J., D. Gabbay, and C. Hartonas (1996). Information Flow and the Lambek Calculus. See Seligman and Westerståhl (1996).

Barwise, J., J. M. Gawron, G. Plotkin, and S. Tutiya (Eds.) (1991). *Situation Theory and Its Applications, II,* Number 26 in CSLI Lecture Notes, Stanford. CSLI Publications.

Barwise, J. and L. S. Moss (1996). *Vicious Circles: On the Mathematics of Non-Wellfounded Phenomena.* CSLI Lecuture Notes. Stanford: CSLI Publications.

Barwise, J. and J. Perry (1983). *Situations and Attitudes.* Cambridge, MA: MIT Press.

Barwise, J. and J. Perry (1985). Shifting situations and shaken attitudes. *Linguistics and Philosophy 8*(1), 105–61.

Barwise, J. and J. Seligman. *Information Flow in Distributed Systems.* Cambridge University Press.

Barwise, J. and J. Seligman (1993). Imperfect Information Flow. In M. Vardi (Ed.), *Proceedings of the 8th Annual IEEE Symposium on Logic in Computer Science*, pp. 252–261. Los Alamito, CA.: IEEE Computer Society Press.

Barwise, J. and J. Seligman (1994). The Rights and Wrongs of Natural Regularity. In J. Tomberlin (Ed.), *Perspectives in Philosophy, 8*, pp. 331–364. Atascadero, California: Ridgeview.

Black, A. W. (1992). *A Situation Theoretic Approach to Computational Semantics.* Ph. D. thesis, Department of Artificial Intelligence, University of Edinburgh.

Black, A. W. (1993). *An Approach to Computational Situation Semantics.* Ph. D. thesis, Department of Artificial Intelligence, University of Edinburgh, Edinburgh, U.K.

Braisby, N. (1990). Situation word meaning. See Cooper, Mukai, and Perry (1990), pp. 315–341.

Braisby, N. and R. P. Cooper (1996). Naturalising Constraints. See Seligman and Westerståhl (1996).

Burke, T. (1990). Dewey on Defeasibility. See Cooper, Mukai, and Perry (1990), pp. 233–268.

Cavedon, L. (1995). *A Channel-Theoretic Approach to Conditional Reasoning.* Ph. D. thesis, Centre for Cognitive Science, University of Edinburgh.

Cavedon, L. (1996). A Channel-Theoretic Model for Conditional Logics. See Seligman and Westerståhl (1996).

Cooper, R. (1986). Tense and discourse location in situation semantics. *Linguistics and Philosophy 9.*

Cooper, R. (1991a). Three lectures on situation theoretic grammar. In M. Filgeiras, L. Damas, N. Moreira, and A. P. Tomás (Eds.), *Natural Language Processing: EAIA 90 Proceedings*, Lecture notes in Artificial Intelligence, no. 476, pp. 101–140. New York, NY: Springer-Verlag.

Cooper, R. (1992). Situation theoretic discourse representation theory. Manuscript, HCRC, Edinburgh.

Cooper, R. (1996). The Attitudes in Discourse Representation Theory and Situation Semantics. See Seligman and Westerståhl (1996).

Cooper, R. and J. Ginzburg (1996). A Compositional Situation Semantics for Attitude Reports. See Seligman and Westerståhl (1996).

Cooper, R. and H. Kamp (1991). Negation in Situation Semantics and Discourse Representation Theory. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 311–333.

Cooper, R., K. Mukai, and J. Perry (Eds.) (1990). *Situation Theory and Its Applications, I*, Number 22 in CSLI Lecture Notes, Stanford. CSLI Publications.

Cooper, R. P. (1991b). Persistence and structural determination. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 295–309.

Davidson, D. (1967). Truth and Meaning. *Synthese 17*, 304–323.

Devlin, K. (1990). Infons and Types in an Information-Based Logic. See Cooper, Mukai, and Perry (1990), pp. 79–95.

Devlin, K. (1991a). *Logic and Information*. Cambridge, U.K.: Cambridge University Press.

Devlin, K. (1991b). Situations as Mathematical Abstractions. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 25–39.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Massachusetts: MIT Press.

Engdahl, E. (1990). Argument Roles and Anaphora. See Cooper, Mukai, and Perry (1990), pp. 379–393.

Fenstad, J. E., P.-K. Halvorsen, T. Langholm, and J. van Benthem (1987). *Situations, Language, and Logic*. Dordrecht, Holland: Reidel.

Fernando, T. (1990). On the Logic of Situation Theory. See Cooper, Mukai, and Perry (1990), pp. 97–116.

Fodor, J. D. (1985). Situations and Representation. *Linguistics and Philosophy 8*(1), 13–22.

Gawron, J. M. (1986). Situations and prepositions. *Linguistics and Philosophy 9*, 327–382.

Gawron, J. M. and S. Peters (1990a). *Anaphora and Quantification in Situation Semantics*, Volume 19 of *CSLI Lecture Notes*. Stanford: CSLI Publications.

Gawron, J. M. and S. Peters (1990b). Some Puzzles About Pronouns. See Cooper, Mukai, and Perry (1990), pp. 395–431.

Georgeff, M., D. Morley, and A. Rao (1993). Situation Theory and Its Applications, III. See Aczel, Israel, Katagiri, and Peters (1993), pp. 119–140.

Ginzburg, J. (1991). Questions without Answers, Wh-Phrases without Scope: A Semantics for Direct Wh-Questions and their Responses. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 363–404.

Glasbey, S. (1994). *Event Structure in Natural Language Discourse*. Ph. D. thesis, Centre for Cognitive Science, University of Edinburgh.

Glasbey, S. (1996). Towards a Channel-Theoretic Account of the Progressive. See Seligman and Westerståhl (1996).

Healey, P. and C. Vogel (1993). A situation theoretic model of dialogue. In K. Jokinen (Ed.), *Pragmatics in Dialogue Management*, Gothenburg Monographs in Linguistics.

Hintikka, J. (1983). Situations, possible worlds, and attitudes. *Synthese 54*, 153–162.

Israel, D. and J. Perry (1990). What is Information? In P. P. Hanson (Ed.), *Information, Language, and Cognition*, pp. 1–28. Vancouver, Canada: The University of British Columbia Press.

Israel, D. and J. Perry (1991). Information and architecture. See Barwise, Gawron, Plotkin, and Tutiya (1991).

Jean Mark Gawron, John Nerbonne, a. S. P. (1991). The Absorption Principle and E-Type Anaphora. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 335–362.

Katagiri, Y. (1991). Perspectivity and the Japanese Reflexive 'zibun'. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 425–447.

Katagiri, Y. (1996). A Distributed System Model for Actions of Situated Agents. See Seligman and Westerståhl (1996).

King, P. J. (1994). Reconciling Austinian and Russellian Accounts of the Liar Paradox. *Journal of Philosophical Logic 23*, 451–494.

Koons, R. (1990). Three Indexical Solutions to the Liar Paradox. See Cooper, Mukai, and Perry (1990), pp. 269–296.

Koons, R. (1992). *Paradoxes of Belief and Strategic Rationality.* Cambridge, U.K.: Cambridge University Press.

Lewis, D. K. (1979). Score Keeping in a Language Game. In R. Bauerle (Ed.), *Semantics from different points of view*, pp. 172–187. Springer.

Lewis, D. K. (1983). General Semantics. In *Philosophical Papers, Vol. 1*, pp. 233–249. Oxford: Oxford.

Lewis, M. (1991). Visualization and Situations. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 553–580.

Lunnon, R. A Theory of Sets and Functions. *Journal of Symbolic Logic*.

Lunnon, R. (1991a). *Generalised Universes.* Ph. D. thesis, Manchester University.

Lunnon, R. (1991b). Many Sorted Universes, SRD's, and Injective Sums. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 51–79.

Makinson, D. (1994). General Non-Monotonic Logic. In D. Gabbay, C. Hogger, and J. Robinson (Eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming, vol. III.* Oxford: Oxford University Press.

Mori, T. and H. Nakagawa (1991). A Formalization of Metaphor Understanding in Situation Semantics. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 449–467.

Muskens, R. (1995). *Meaning and Partiality.* Studies in Logic, Language, and Information. Stanford: CSLI Publications.

Nakashima, H., H. Suzuki, P.-K. Halvorsen, and S. Peters (1988). Towards a Computational Interpretation of Situation Theory. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, Institute for New Generation Computer Technology, Tokyo, Japan, pp. 489–498.

Nakashima, H. and S. Tutiya (1991). Inferring *in* a Situation *about* Situations. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 215–227.

Perry, J. (1979). The Problem of the Essential Indexical. *Nous 13*, 3–21.

Perry, J. (1984). Contradictory Situations. In F. Landman and F. Veltman (Eds.), *Varieties of Formal Semantics: Proceedings of the 4th Amsterdam Colloquium*, Groningen-Amsterdam Series in Semantics. Dordrecht: Foris.

Perry, J. (1986). From Worlds to Situations. *Journal of Philosophical Logic 15*, 83–107.

Perry, J. (1993). *The Essential Indexical and Other Essays.* New York: Oxford University Press.

Perry, J. and E. Macken (1996). Interfacing Situations. See Seligman and Westerståhl (1996).

Plotkin, G. (1990). An Illative Theory of Relations. See Cooper, Mukai, and Perry (1990), pp. 133–146.

Poesio, M. (1991). Relational Semantics and Scope Ambiguity. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 469–497.

Poesio, M. (1993). A situation-theoretic formalization of definite description interpretation in plan elaboration dialogues. See Aczel, Israel, Katagiri, and Peters (1993), pp. 339–374.

Restall, G. A. (1996). Information flow and relevant logics. See Seligman and Westerståhl (1996).

Rounds, W. C. (1990). The Complexity of Paradox. See Cooper, Mukai, and Perry (1990), pp. 297–311.

Rounds, W. C. (1991). Situation-Theoretic Aspects of Databases. See Barwise, Gawron, Plotkin, and Tutiya (1991).

Rounds, W. C. and G.-Q. Zhang (1996). Attunement to Constraints in Nonmonotonic Reasoning. See Seligman and Westerståhl (1996).

Ruhrberg, P. (1996). A simultaneous abstraction calculus and theories of semantics. See Seligman and Westerståhl (1996).

Schulz, S. (1996). Minimal Truth Predicates and Situation Theory. See Seligman and Westerståhl (1996).

Schütze, H. (1991). The prosit language, version 0.4. CSLI Memo, Center for the Study of Language and Information, Stanford University, Stanford, California.

Seligman, J. (1990a). *Perspectives: a Relativistic Approach to the Theory of Information*. Ph. D. thesis, University of Edinburgh.

Seligman, J. (1990b). Perspectives in situation theory. See Cooper, Mukai, and Perry (1990), pp. 147–191.

Seligman, J. (Ed.) (1991a). *Partial and Dynamic Semantics III*. Centre for Cognitive Science, University of Edinburgh. DYANA Deliverable R2.1.C.

Seligman, J. (1991b). Physical Situations and Information Flow. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 257–292.

Seligman, J. (1996). The logic of correct description. In M. de Rijke (Ed.), *Advances in Modal Logic*, pp. 107–136. Kluwer.

Seligman, J. and A. ter Meulen (1995). Dynamic aspect trees. In L. Pólos and M. Masuch (Eds.), *Applied Logic: How, What and Why*. Kluwer.

Seligman, J. and D. Westerståhl (Eds.) (1996). *Logic, Language and Computation*, CSLI Lecture Notes, Stanford. CSLI Publications.

Sem, H. F., K. J. Saebo, , G. B. Verne, and E. J. Vestre (1991). Parameters: Dependence and Absorption. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 499–516.

Shin, S.-J. (1991). A Situation-Theoretic Account of Valid Reasoning with Venn Diagrams. See Barwise, Gawron, Plotkin, and Tutiya (1991).

Stalnaker, R. (1986). Possible Worlds and Situations. *Journal of Philosophical Logic 15*, 109–123.

Suzuki, H. and S. Tutiya (1991). A Strictly Incremental Approach to Japanese Grammar. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 517–532.

ter Meulen, A. (1995). *Representing Time in Natural Language. The dynamic interpretation of tense and aspect.* Bradford Books. Cambridge, MA: MIT Press.

Tın, E. and V. Akman (1993). BABY-SIT: A computational medium based on situations. In P. Dekker and M. Stokhof (Eds.), *Proceedings of the 9th Amsterdam Colloquium*. Amsterdam: ILLC.

Tın, E. and V. Akman (1994a). Computational Situation Theory. *ACM Sigart Bulletin 5(4)*, 4–17.

Tın, E. and V. Akman (1994b). Situated Processing of Pronominal Anaphora. In H. Trost (Ed.), *Tagungsband KONVENS '94 Verarbeitung Natürlicher Sprache*, Vienna, Austria, pp. 369–378. Informatik Xpress.

Tın, E. and V. Akman (1996). Information-Oriented Computation with BABY-SIT. See Seligman and Westerståhl (1996).

Tin, E., V. Akman, and M. .Ersan (1995). Towards situation-oriented programming languages. *ACM Sigplan Notices 30*(1), 27–36.

Tojo, S. and S. Wong (1996). A Legal Reasoning System Based on Situation Theory. See Seligman and Westerståhl (1996).

Vogel, C. (1992). A situation theoretic model of inheritance reasoning. In *Research Conference on Logic, Language, and Information: Toward an Integrated Theory of Linguistic Processing*. European Science Foundation. December 13-17, 1992.

Westerståhl, D. (1990). Parametric Types and Propositions in First-Order Situation Theory. See Cooper, Mukai, and Perry (1990).

Zadrozny, W. (1996). From Utterances to Situations: Parsing with Constructions in Small Domains. See Seligman and Westerståhl (1996).

Zaefferer, D. (1991). Probing the Iroquoian Perspective: Toawrds a Situated Inquiry of Linguistic Relativity. See Barwise, Gawron, Plotkin, and Tutiya (1991), pp. 533–549.

Zalta, E. (1991). A Theory of Situations. See Barwise, Gawron, Plotkin, and Tutiya (1991).