

Information Systems as a Social Science

An Alternative to the FRISCO Formalism

Ronald K. Stamper

MEASURE Ltd, 29 Dartmouth Road Park, London NW5 1SU

Tel: +44 0207 813 7359 Fax: +44 0207 482 2968

Email: rstamper@compuserve.com

University of Twente, P.O. Box 217, NL-7500 AE Enschede

Email: r.k.stamper@sms.utwente.nl

Abstract: The FRISCO Report has initiated an important debate on the discipline of information. Chapters 1, 2 and 6 contribute ideas for all shades of opinion, a Broad View contrasting with Chapters 3, 4 and 5, “the core of the report”. This paper presents a Social View as an alternative to this Core View, which it criticises in two ways. Firstly, it limits the scope of information systems by treating it as an adjunct to software engineering, having no place for key properties of information in organisations and society. Secondly, it does not satisfactorily link its framework of concepts to the empirical world, but relies on the mentalistic notions of ‘perceptions’ and ‘conceptions’ in the mind of some, unspecified interpreter. The strength of the Core View lies in its formal precision. But a Social View can lead to an alternative formal framework that can place a much wider range of information systems concepts on a firm empirical basis. Thus provides a foundation for information systems as a social science, rather than a branch of applied mathematics. The Core View includes the belief in an objective reality to which we have direct access via the supposedly transparent languages of words, numbers and diagrams, whereas the Social View, without rejecting that position, requires us to explain how we come to construct such knowledge. It also compels us to examine all the other functions we perform with information, especially our construction of social reality. Forced to examine these questions, we have to address many elusive problems, such as the creation of meanings, the role of intentional communication, the construction of time, and the systems of norms we call ‘organisation’. Finally, the paper re-works the Japan Wines case study used in the FRISCO Report. This shows that the Social View leads to a more detailed analysis of the business problems than the Core View, with as much formal precision, which lends itself to computer interpretation, while also yielding models that are easier for users to interpret.

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35500-9_30](https://doi.org/10.1007/978-0-387-35500-9_30)

E. D. Falkenberg et al. (eds.), *Information System Concepts: An Integrated Discipline Emerging*

© IFIP International Federation for Information Processing 2000

If the FRISCO Report has a lasting contribution it will result from the debate it has provoked about the central concepts of Information Systems, and its effort to place the subject on a scientific foundation. The initial hope that we would produce a stable framework of information systems concepts turned out to be far too ambitious. The people who are working in this field now have a duty to make positions clear and to discuss them in a rigorous, critical, constructive and open way.

AN IMPORTANT DEBATE

FRISCO did not consider seriously, or present in its report all the views deserving attention. Nevertheless, it presents a broad, discursive treatment of information systems concepts in chapters 1, 2 and 6, which succeed in bringing together most of the key ideas that should inhabit any framework of IS concepts. I shall refer to this as the Broad View. On the other hand, one view is given priority. It is asserted on p. 14 that “the core of the report is formed by chapters 3 and 4”. So I shall refer to this as the Core View.

Many ideas carefully articulated in the Broad View find no place in the narrowly focused Core View, which I was tempted to might call the Formal View. It has the major virtues of precision, formality and consistency and it strives for clarity. However, I prefer not to suggest that formality is the preserve of the Core View. Indeed, every worthy alternative must aim for such formal clarity so that it can be subjected to constructive criticism in ways that, for instance, the Broad View does not invite. I shall focus my discussion on the Core View and attempt to keep my criticism as constructive as possible by drawing together more ideas from the Broad View, into an alternative also capable of precision, formality and consistency. This alternative I shall refer to as a Social View of information systems.

In Chapter 7 of the Report, my co-authors and I expressed personal positions. Some favoured the Broad View and indicated a little dissatisfaction with the Core View. Wolfgang Hesse emphasises on p. 175 the value of a framework “ranging from the physical level through the linguistic levels up to the social level [for dealing] with all the relevant aspects of information systems.” Paul Lindgreen on p. 183 puts his finger on the danger of the formal approach in the Core View by quoting its definition of a language as, “A non-empty set of permissible arrangements of symbols taken from an alphabet.”, remarking “For me a language is a social convention.” Björn Nilsson comments on the narrow definition of “the information in a message [as] the new knowledge gained by an interpreter thereof.” He points to three relevant missing part of the definition, which relate to intentions, meanings and effects of the message. Alex Verrijn-

Stuart, though content with the Core, points out, on p. 196, that the semiotic ladder reveals a different meaning of “meaning” on each level, something which the formal Chapter 4 fails to achieve.

On the contrary, semiotics provides a precise, formal and consistent basis for a discipline of information systems. The FRISCO group never examined this question. My own contribution to Chapter 7, “A Dissenting Position” attempted in the allotted six pages to sketch an answer. As the member responsible for introducing *six* semiotic levels (Stamper 1973 and 1989), I feel obliged to take this discussion further. Hence one of my objectives in this paper is to explain how to handle the three, more problematic, non-technological semiotic levels - semantic, pragmatic and social – in to the formal and precise standards that I share with my colleagues.

Two positions

My motivation for applying semiotics to information systems, besides removing vagueness about our basic concepts, has always been to bridge the gap between the technical and social orientation to information systems problems. This gap appears to me not only to be damaging to our scholarly community, but also to our problem-solving ability, and, most of all to our attempts to create a science of information systems. There still appear to be two, somewhat antagonistic communities representing these two approaches. During the 1970s my research programme was funded by two research councils, one on each side of the divide. We were trying to be precise about organisations using the concept of a norm and using legal norms as empirical materials, to define information requirements for computer applications,. The reviewers on the science and engineering side said, “This isn’t real science, it has something to do with the law.” Those on the economic and social side said, “We must not fund this because the computer should have nothing to do with the law, which is not just a system of rules.” In the debate, which we shall embark upon, I hope that this sterile set of attitudes will not be visible.

The Social View I shall present does not exclude the Core View of FRISCO but incorporates it as a special case. It is also overcomes the inadequacy of treating information systems as a branch of applied mathematics: a formal discipline without empirical foundations. Machines and formalisms for handling information are human artefacts and they should be incorporated naturally into a social model of information systems. This approach gives us a far better chance of building an empirically based science instead of a formalism that has only a rather nebulous connection with reality.

When I became involved in information systems, first in practice then academically, we taught methods and notations, supported by a wealth of practical experience, to enable people to define computer applications in the business world. There was no engineering because there was no supporting science. We were, and still are, the blacksmiths of the information age. The time has come to choose whether to continue that way or to re-discover information systems as a science that can stand comparison with any other in the quality of its theories and the stringency of the empirical tests they have to pass.

Therefore, this paper has three main parts. *First* it explains why the Core View, despite its excellence from a formal perspective, fails as a basis for an empirical science. It lacks the conceptual richness to tackle many of the issues introduced in the Broad View. *Second*, it presents an alternative solution that, while capable of achieving a level of formality comparable with the Core View, also handles the broader subject matter of information systems. *Third*, it applies methods derived from this alternative Social View to the analysis of the Japan Wines case study used in the FRISCO Report, producing a more thorough requirements analysis that is also easier to understand and implement.

THE CORE VIEW OF FRISCO – FORMAL BUT NOT SCIENTIFIC

At the end of the 20th century, the study of information systems is regrettably not yet a scientific. It does not even study information systems in any broad sense. Rather, it subsists as the handmaiden of software engineering, consisting of notations and methods to serve that end. The Broad View of FRISCO moves us towards an empirical science but the Core View, while correctly embracing formality, leaves the empirical questions aside. In effect, Chapter 4 treats information systems as a branch of applied mathematics, closely allied to software development.

Software engineering is important but far more important is the shaping of our information society, for which we need an empirical science that can link software to the social processes it should serve. We have the opportunity to create one of the most exciting and important disciplines for the next century.

What is essential to an empirical scientific discipline?

A massive literature (including Popper 1959 and 1963, Lakatos 1970, Feyerabend 1975) deals with the question of how to draw a line between science and non-science. All the authors relate their answers to the methods

of enquiry adopted by the discipline. These range from the strict demarcation line Popper imposes with his refutationist method to the tolerant Feyerabend slogan that “anything goes”. However, on one point the philosophers of science do agree. *They all expect scientists to check their theories, their universal laws, against particular empirical observations* (“The only pressing reason for changing a theory is disagreement with the facts.” Feyerabend p. 38). To achieve this they must state their theories clearly and unambiguously, using a terminology that allows one to deduce statements that can be tested operationally. Unless a theory or a model or a computer program can meet these conditions, there can be no way of checking it against experience of the real world.

There are two conditions to consider. The first is the standing of the discipline as a deductive system. It must be formulated in a language that allows unambiguous logical conclusions to be drawn from statements of theory and fact. Formality is necessary for exposing and eliminating *syntactic ambiguity*. However that does not deal with *semantic ambiguity* which is achieved by meeting the second condition, that of making operationally reliable connections between reality and the language of the theory. “Ambiguity” is itself ambiguous. In his commentary on FRISCO Alex Verrijn-Stuart (1998) falls into this trap when implicitly he claims that formality is sufficient: “formal, that is to say, unambiguous and precise . . .”. Unfortunately formality is only a protection against ambiguity and guarantee of precision in their *syntactic* senses.

Logicism

Logic, understandably, has a high standing in any field associated with computing. Pre-eminently, the computer is a machine for automating logical operations, so logic has, inevitably much to contribute to computing theory, and, one might expect, also to information systems. It has to play a central role in any alternative approach.

But logic inspired a tradition of thought over this century, originating in the works of Frege and of Russell and Whitehead who were convinced that the best way of placing mathematics on a sound foundation would be to show that it arises from logic. Logicism, as this line of thinking is called, is still influential and appears to have inspired the Core View. The goal of Chapter 4 appears to be a kind of “Principia Informatica” intended to give our subject a secure foundation. Its aim is to derive all the essential concepts relevant to information systems by building a logical edifice on some self-evident and totally reliable assumptions. Is this approach likely to succeed?

Logicism as a strategy is losing its hold. It was shaken by Gödel in the 1930s. Lakatos (1976) examined the creation of mathematics as a social

process, contributing to a growing body of opinion (for example, Kitcher 1984, Hersch 1997, Ernest 1998) which favours empirical activity and social constructivism as alternative and more reliable foundations for mathematics. This new line of thinking about the formal disciplines bears some affinity with the Social View of information systems.

Although, at first sight, it may seem reasonable to hope that one formal discipline, mathematics, might be given foundations in another (logic), it is doubtful that any messy, complex empirical discipline can be rooted in logic without losing touch with reality. Physics appears to be an exception. Newton's laws encode much of our knowledge of the field with astonishing economy. But what those laws encode is a vast amount of concrete knowledge of relevant domains, from everyday mechanics to astronomical observation, knowledge that was accumulated and assimilated long before the formalised summary in modern physical theory was conjectured. The Core View of FRISCO shows less evidence of being based on vast knowledge of organisations and human information systems, than upon deep familiarity with computers and software engineering. Whatever logic may be discovered in organised behaviour should arise from the study of organisations rather than imported from computing.

Logicism assumes that the particular form of logic being used is necessarily a sound instrument, quite free from any pre-suppositions, and capable of guiding our thinking through purely rational deductive arguments. Logic, as any formal system, functions mechanically and can do a wonderful job in the right circumstances. What are the circumstances where classical logics are appropriate? They evolved to deal with problems in the natural sciences and related forms of engineering. We should at least be cautious about applying to the study of information systems forms of logic that evolved to serve disciplines that regard the information in everyday discourse and scientific discussion as a transparent medium through which to 'see' the natural reality that interests them. We can safely adopt the same attitudes of natural science and mathematics when we study the physics, empirics and syntactic aspects of information. But we also have a professional duty to examine those aspects of information that any natural scientist can take for granted. We have no obligation to do more natural science but one of our tasks, for example, is to understand how natural scientists and mathematicians use information to explain their success. We cannot take for granted the informal meta-languages that others use unconsciously, because they are our what our subject is about. Therefore we should be cautious about being trapped within the limits of a formalism that evolved for totally different purposes from our own. I shall return to this issue later, because an alternative approach arises from a critical questioning of the assumptions behind classical logic.

A framework of information systems concepts can be held together *internally* using logical methods but it will only stand up to external, empirical testing if the foundations are well chosen and the framework firmly fastened down. Thus it will only serve its purpose if it is based on all the key concepts relevant to our discipline and if it also links its treatment of them to reality (begging the question about the nature of reality, which natural science can also usually ignore).

The link with reality

Formal systems do not necessarily have any connection with reality. Anyone using information of any kind should be careful to check out its connection to reality. This applies just as much to theories and models as to any other information. Such connections may be irrelevant because all we want is a self-contained formal system that will stimulate the imagination – computer games are like that. Sometimes the link to reality is implicit, as I hope it is with systems to control nuclear power stations or run social security. We just hope we can rely on the intuition of the designers and programmers in the absence of any explicit treatment of that link: the semantics of the data and programs. Surely, any information systems professional should always be able to advise on the quality of the relationship between information and reality.

To set up its links to reality, the Core View uses a sequence of definitions, which, after getting lost in some circularity, eventually lead us back to some primitive notions. The choice of primitive notions is crucial. There is no point in defining everything in words because that process has no end. At some point we have to say, “Well, beyond these definitions we rely on a few primitive terms that anyone can understand by learning to operate with them in practice.” Physics uses notions of a physical body and the simplest phenomena involving them. Newton’s laws talk of bodies, their motion, change of motion and forces applied to them. A child who has spent time kicking a football on the street will have an *operational* grasp of those concepts *and a way of demonstrating them convincingly to other people*. Physics is an ideal paradigm for a science built upon well-founded, primitive notions. How does the Core View of FRISCO fare in this respect?

Unfortunately, when we reach the primitive notions, they are ‘perceptions’ and ‘conceptions’ linked to things in the real domain by processes within the minds of unspecified individuals. The inaccessibility of these mental processes is admitted and dismissed in the same breath, on p. 35: “While, the conceptions, as such, are not accessible, indirectly they are available for discussion and hence operationalisable.” If information systems were no more than the discipline of manipulating character strings or other

sign-tokens, we could have formulated the Core View without invoking ‘perceptions’ and ‘conceptions’. The physical tokens and their syntactic types can all be observed accurately and handled by operations that can be demonstrated reliably. In that case, information systems would be firmly within the domain of computer science, and would continue to lack the capacity to say anything about people and organisations.

By introducing ‘perceptions’ and ‘conceptions’ and declaring that they can be treated operationally, simply because we can discuss them, we have admitted that information systems does raise problems that a computer scientist can safely ignore. But all we are offered to see and work with empirically by examining “any available discussion” will be representations, the signs that our information system will manipulate. We are back in a world that can be handled solely in terms of manipulating tokens, signs, representations, data and so on – the world of computing. We have not begun to get to grips with a real world in any sense. When one designs an information system, one has to be careful how the sign-tokens and sign-types are linked to reality. The Core View would allow us to operationalise astrology and physics in just the same way. Astrologers can discuss their ‘perceptions’ and ‘conceptions’ equally convincingly as physicists but they easily bamboozle anyone who is satisfied by the operational criterion expressed in the Core View. Modern science is a huge practical success because the physicists do more than discuss their perceptions and conceptions. They impose the strict criteria of operational meanings on their theories, which must imply particular statements that can be tested empirically.

Meaning

“Meaning” is a word with many meanings, as noted at length in FRISCO’s Broad View in Chapter 6. A relationship between a representation and what it represents provides one meaning of that ambiguous word. A general account of how such relationships are formed constitutes a *semantics*. Computer scientists use a completely formal semantics which defines the meaning of one formal structure (eg: a program) as a representation of another (another program or a set-theoretic structure). They inhabit a world of pure symbolic expressions in which their meanings can be defined by rules translating them into other expressions. Their world is tightly controlled and, within mathematically clear limits, uncertainty can be totally eliminated. The world studied by information systems, as envisaged in the Broad View, is quite different; it is full of people and uncertainty, unlike the limited world studied by computer science.

Meaning, in a sense adequate for a science of information systems, must account for three things: for *reality*, the things we represent using information, for the *signs* that do the representing and for the *link* between them. Each of these three components of the semantic model:

sign ----- MEANING ----- thing
 (representation) (link) (reality)

It is no good lamenting that the subject matter of information systems is “fuzzy” and complaining that semantics is one of the “nebulous” levels of the semiotic ladder and, therefore ignoring it. Semantics is one of the central issues for a science of information systems and any relevant theory must come to grips with it. Unfortunately, FRISCO’s Core View clearly fails to do this. Firstly, it contains no clear commitment concerning the nature of reality. Secondly, instead of the excellent, ostensibly definable primitive, ‘sign’, on which semiotics is based, it uses ‘representation’ (defined as an ‘actand’ (defined as a ‘thing’ (defined as a ‘conception’))). So, thirdly, the only link between the sign and reality we are given is some process in a mind. We are not even told how to identify the relevant mind in any situation. Should we do so, the link between sign and signified is accounted for by private, subjective processes of perceiving and conceiving, processes that are inaccessible except by introspection.

Incompleteness of the Core View

Although semantics is an issue of central importance to anyone using information, the Core View has nothing relevant to say on this subject. It is also unable fully to handle issues on two other semiotic levels: pragmatic and social.

Newton’s laws apply equally well to the motion of Mars round the Sun as to the motion of a boot hurled by an angry spouse at a partner. The intention behind the throwing is irrelevant to the motion of the boot, so physics rightly ignores it. However, if nothing more than a verbal insult were hurled, the intention would be highly relevant to understanding this message, as the Broad View makes clear on p. 22. Any framework of concepts that cannot deal with *intention* is not going to be relevant to information systems involving people. The Core View deals with *intension* (the set-theoretic counterpart to an extension) but does not mention “intention”. The notion of a “goal” is a partial substitute. On p. 72 “communication” is defined as “an exchange of messages” with an explanation that it may involve several human agents in many transactions “and the goal of the communication is made present in all communication partners.” This final proposition is not clear, but it talks of “*the*” goal, implying that one goal alone can govern such a complex process. Often each participant has several goals in mind and

changes them in the course of the conversation. The claim that this approach to communication “corresponds to the reasoning applied in speech act theory” is false. The use of goals and omission of intentionality make the Core View unable to treat the separate motivations behind each transaction in a conversation. It cannot therefore handle the analysis of the detailed pragmatics of an information system. This is confirmed by the explanation of the goals in the case study, on p. 132. A science of information systems must be able to accommodate the rich details of speech act theory. Again the Core View of FRISCO is irrelevant to an essential part of the subject.

Now we ask whether the Core View is relevant to the social aspects of information. Assumption [h] on p. 84 acknowledges that “inter-subjective reality” depends on agreement across a society. This idea is basic to a science of information systems. The information, representations, messages, signs we deal with – call them what you will – do not grow on trees. They are normally social products that fall beyond the scope of natural science. We only have information when we have something (a representation) that can stand for something else (what it represents). Such linkages have natural origins in part-whole relationships and in similarities but are now almost entirely products of human cultures. Society is necessarily a core concept of information systems but not only because it is the origin of information. Society is also the product of information. The Broad View, on pp. 17 and 18, makes clear that a science of information systems should account for the value that communication adds through its consequences for society. If we ask what information gives us, it is clear we cannot eat it, wear it nor travel about on it. Indirectly it helps us obtain food, protection and transport, but only by virtue of disposing people to serve us and act in a co-ordinated fashion. Information has no other immediate value than the social effects it causes. Its secondary consequences may be physical and economic. The notions of transition and state, on p. 86, are relevant but they do not relate to changes of a social kind. The Core View has only a treatment of state-transitions lifted directly from computer science, a special version of cause-and-effect relationships in the physical world. An extension is possible. The effect of a communication is to change the propositional attitudes of the people who have some responsible role to play. In the realm of information and society, transitions always entail responsibility on the parts of the players. Responsibility is *mentioned* in the Core View but not *used*. No theory of information systems can be complete if it does not include an adequate means of treating responsibility.

Shortcomings summarised

The Core View of FRISCO belongs in the tradition of the formal disciplines, including computer ‘science’, which studies the mechanical

manipulation of formal structures. I am using “science” in the sense of an empirical discipline that is grounded in the experiences we share in the public domain that our information strives to represent. Empirical sciences can benefit by using formal methods but only if the formal language is firmly linked to the reality being studied. For this link, the Core View relies on empirically inaccessible mental processes.

Logical formality is desirable because it helps us to eliminate *syntactic ambiguity*. In this, the Core View succeeds. But the kind of logic employed still leaves the problem of *semantic ambiguity* untouched. Good definitions rooted in simple, operationalised, primitive notions can solve this problem. We are given definitions of such key terms as “information”, “communication” and “knowledge” based on the primitive concepts of “perception” and “conception”. Moreover, these fail to distinguish the multiple meanings of these terms that are acknowledged elsewhere in the FRISCO Report..

Users of information are generally indifferent whether the information they receive has been produced by computers or by paper and pencil. But they do want to know that it is *meaningful* and be able to anticipate any semantic problems that may arise. They need to know the true *intentions* and any hidden ones lurking behind the messages they receive. They want to know its *effects on people*, who can only respond by adjusting their values, beliefs, obligations, expectations and so on, before using the information to guide their own actions. And, if pressed to account for those actions, they will try to explain their degree of *responsibility* as well as the responsibilities of any others implicated with them. The Core View does not provide adequate tools for handling these semantic, pragmatic and social aspects of information.

AN ALTERNATIVE FRAMEWORK

To criticise is one thing but to propose an alternative is altogether more difficult. For any alternative worth considering one must ask a number of questions. The Report, p. 159, raises some relevant criteria related to “expressiveness”, “arbitrariness” and “suitability”.

- 1) Does it deal adequately with the relevant IS concepts identified in the Broad View, including meaning, intention, responsibility and society?
- 2) Does it handle the multiple meanings of key IS concepts such as information, communication, meaning and knowledge?
- 3) Does its structure of definitions lead back to operationally sound primitive notions so that its propositions can be tested empirically?
- 4) Does it lend itself to a formal method for modelling or describing information systems that ensures internal consistency?

5) Does it lead us to draw distinctions and investigate problems central to information systems that are now overlooked by the Core View?

6) Does it support the continual accretion of testable knowledge of the information systems domain, as we would expect of any other modern science?

7) Does it enable us to do the routine work on information systems conveniently and effectively?

I shall present an alternative and demonstrate that it meets these criteria. Detailed, critical assessment of this alternative will be welcome.

To make the transparent visible

The people working in the natural sciences aspire to discover the eternal laws of nature that have governed the world since its creation and will continue to do so whether or not there are any people around. Relativity and quantum theories recognise the role of the observer but, in general, natural science would become unnecessarily over-complicated if it introduced human observers and experimenters into their explanations at every point. They want to look at the natural world *through* their theories and observational records without having to enquire about all this information as well. For natural scientists, the information systems they use should be transparent.

However, the opposite should be true for us as students of information systems. Interesting though they are, the natural sciences fall outside the scope of our discipline. Of central importance to us are questions about how those natural scientists know anything, how they communicate their knowledge to others, how they justify accepting what anyone says, and indeed, what is the nature of their knowledge and information. Moreover, we must ask similar questions about business people, lawyers, politicians and citizens, concerning how they use information, how they create the kinds of knowledge they need and how they justify the ways they use information and knowledge. We must make the transparent visible.

Transparency is achieved by a mental trick. Asked to explain how we know anything, the person studying the natural world would tell us to open our eyes and look at the world where we shall find the individual things that compose it. We can mentally order them and place them into sets to define properties, relationships and other predicates that we can assert about the individuals. This is the familiar objectivist position.

The Broad View in FRISCO adopted a radical solution on p. 10. “The constructivist approach rejects the objectivist assumption that reality simply exists as ready-made individual things and replaces it with a reality that is

built through the operations of the members of a community through the social consensus about subjective conceptions that serve their culture.” However, the Core View retreats from this subjectivist position on p. 28: “**Constructivist:** somebody who also believes that ‘reality’ exists independently of any observer, . . .” and clings on to the old way of thinking. By introducing the perceiving and conceiving mind, it does not reject objectivism. Rather, it provides an excuse for maintaining the orthodox view of reality.

Reluctance to abandon the security of established habits of thought is understandable, especially as they have served the natural sciences so well. For information systems to be studied scientifically, however, we must bring the old tacit assumptions to the surface because they beg exactly those questions that lie at the heart of our subject: they are assumptions about meanings and how we obtain and represent knowledge of the world. To analyse the transparent layer:

signs ----- MEANING ----- real things

we have to have to think very carefully about what we understand by signs, meanings and, perhaps most difficult of all, real things. This last problem can only be solved by adopting an *ontology*, a solution to the kind of metaphysical question we would prefer never to ask.

Ontological assumptions accompanying predicate logic and set theory

Our duty as information systems professionals includes that of making clear and explicit the relationship between information and reality. This applies to the hidden assumptions on which a theory may be based. Examining the Core View reveals several assumptions:

- 1) An objective reality exists independently of anyone, so that, despite differences of opinion, it is always the same reality that we experience.
- 2) That reality is composed of individual things.
- 3) The expressions of a language are always references to individuals or sets (or, recursively, sets of sets and individuals).
- 4) These true meanings are independent of anyone, despite the possibility of individuals having muddled perceptions and conceptions.

These are accompanied by constructivist caveats cautioning us that different people might have rather different views of this one reality and have to negotiate a shared view, but without realising that the nature of these negotiations is a proper subject for a science of information systems.

Nevertheless, subject to that negotiation, they take for granted one *real* reality about which we can express facts using propositions in the form:

predicate (individual, individual, . . .)

This is the basis of the formalism used in Chapter 4. Although it has the virtues of formality, it allows us to ignore two of the most obvious assumptions that govern any enquiry into how we know anything or represent what we know.

Those two assumptions are rather certain.

- 1) We know nothing without a knowing agent being involved.
- 2) That knowing agent has to do things to acquire that knowledge.

They are the basis of what I call an “actualist” ontology because it assumes only the actual world to which we have access via our actions. As far as information systems studies are concerned, predicate logic and set theory are not neutral. They embody assumptions. These do not *prevent* our referring to knowing agents and their actions, but they fail to *ensure* the agents are included in the picture. We could introduce propositional attitudes, predicates applied to propositions, that specify that someone knows [believes, sees, thinks, . . .] that some proposition is the case, but this strategy creates all kinds of difficulties that ruin the mathematical elegance of first order logic.

If we look just a little more carefully at how we know things, we find two routes. The knowing agent can get involved directly to obtain substantive knowledge of the world, alternatively the agent can know something semiologically, by vicariously using the information that others supply about their direct experiences. These correspond to working with the two ends of the line:

signs ----- MEANING ----- real things

We can get to know real things directly or vicariously through the use of signs (information). This distinction is basic to any science of information systems which must explain how we use information (signs) to know about the world and order our lives, and where the knowledge comes from in the first place.

Affordances – reality as invariant repertoires of behaviour

Every animal can acquire direct knowledge by encountering the world through its own actions. Human beings learn even more about the world through the use of signs. Signs are also real things with the additional property that they stand for other things. So we should begin by thinking about direct knowledge, first.

Such knowledge we can represent using an alternative to the well-formed-formulas of predicate logic, I proposed in 1984 (Stamper 1985) a wff with two terms:

Agent behaviour

This involves none of the tacit assumptions of an objectivist ontology. No one has any operational knowledge of reality more basic than of himself or herself, as the real thing that does the knowing (excluding some possible doubts about dreaming). Making operational sense of the 'behaviour' term is more challenging. (For a more detailed explanation see, Stamper 1996.) We can start with James Gibson's Theory of Affordances, which he developed, in his new account of perception. He pointed out that the objectivist, opening her eyes to see the world, looks at something and forms an inverted image on her retina; to perceive this image, another, inner eye must do something similar, beginning an infinite regress that succeeds in explaining nothing about perception. Gibson rejected the idea that animals perceive a ready-made world. Instead, they have to make it up by discovering, within the limitless flux of signals impinging on them, whatever can help them to survive, preferably in comfort and without experiencing harm. What every organism has to discover are the useful *invariant repertoires of behaviour* that its own physiology combined with the structure of the environment afford them. Appropriately, Gibson termed these invariants that we perceive "affordances". Affordances are the only real things we can perceive directly. A rabbit, as Gibson points out, must learn where to obtain the repertoire of behaviour we call 'hiding' if it is to escape from predators, some kinds of terrain, with trees, bushes, deep grass or soft earth afford this repertoire of behaviour and others do not.

To understand adequately the notion of an affordance may take some while initially. If you want to see an animal acquiring them, just watch your babies creating their own reality. A baby, given a cup, will do everything imaginable with it, from pouring a drink all over the floor, throwing the cup at Mummy, banging itself on the head with it, and, at some point, if you are lucky, drinking with it. The cup is not a ready-made objective thing, sitting there waiting for the baby to see it as a cup, but it is the full repertoire of behaviour she discovers. Social forces will lead that all-too-rich affordance, the baby's notion of a cup, to be limited within the bounds that the rest of the family will countenance – but that is to run ahead of our story. Everything we recognise constitutes an invariant repertoire of behaviour; I have found no exception, so far. You will find plenty of examples of affordances in the work of Gibson himself (1979) and a useful work by Michaels and Carello (1981), or you can make up your own.

[In this formal system nothing simply *is the case*, as propositional logic would allow us to say, but everything is made or observed to be that way by some responsible agent or agents doing something, most basically, discovering the relevant affordances. (I emphasised this point in my dissenting reflection in the FRISCO Report by writing it without using the verb “to be”. Anything with an operational meaning that is worth saying can be said in this constrained language by writing about what agents do, with no assumption that things “are” as they are, of their own accord.)]

How does society create reality? An information systems problem

The recognition of affordances can be handed down genetically, as instincts, but all animals will also learn some by themselves (especially our human babies). Humans are unusual because we only construct a very small part of our reality independently. Most of the reality we know has been created by our enveloping cultures over the course of centuries or millennia. The affordance we call a ‘cup’ is a human artefact and the way it can be used is not only a matter of what behaviour it makes possible physically but also of what it affords us socially. The baby will learn that, in general, you do not throw cups at people or tip them over when full except to pour the contents to another vessel. In later life she may find that cup throwing is permitted in some fairground side-shows or has its use in a dispute with a drunken husband. The world we create is as much social in character as physical. Even the physical world known to natural science is predominantly a social construct, and science is the social activity constructing it.

Moreover, the social construction of reality takes place, not in the way baby first learned the affordance we call ‘cup’, but in the context of our information systems. We may learn the *name* of a well-established affordance, long before we obtain sufficient direct knowledge of it. I had certainly heard of a reef knot before I learned to tie one and language and diagrams played a significant role in my acquiring the ability to realise that affordance. Many social affordances are constructed with the aid of that most important information system, the legal system: marriage, copyright, and insolvency, for example. Rules, but not legal rules give us football and etiquette. But first class football, which calls for more than sticking to the rules, and courtesy, which goes far beyond etiquette, involve tacit norms of a subtlety that no explicit rules can capture, and we learn them too. These are all products of our social use of information. For these reasons, we should take Society as the ultimate agent at the root of all our common sense knowledge of reality.

Whenever I have advanced this radically subjectivist position, I have provoked two principal kinds of objections. Some arguments have an ideological basis and others are of a rational nature.

Information systems are central to the effective functioning of society

The mention of Society has often elicited the most intemperate reactions, apparently for political reasons, ever since Thatcher and Regan gave Society a bad name, Thatcher even went so far as to assert that it does not exist. Perhaps I should not bother to counter this irrationality, as their outmoded ideology is losing its power to distort perceptions. My argument requires no more than a common sense acceptance that society, represented by the culture in which we are raised, contributes hugely to our perception and understanding of the world. Indeed, today the progress of the natural sciences depends increasingly on social mechanisms.

Even the perception of many important phenomena is beyond any individual. Openness and sharing information gives us enhanced, collective powers of perception. Consider some of the things no single person could possibly perceive without massive contributions of observations and theories from other people – the hole in the ozone layer, a large computer operating system and tectonic plates, for example. Only groups of people supported by sophisticated information systems can experience these huge repertoires of behaviour. If you have difficulty accepting that without society we should be puny things, try to find an individual able to perform these perceptual tasks alone with no prior relevant knowledge. We may credit an individual with discovering the hole in the ozone layer or tectonic plates but they certainly did not perceive them, their achievements were the synthesis and analysis of masses of information generated by large numbers of their colleagues. Without their scientific communities they would have been 'blind'. This does not diminish their achievement in fitting the last piece of the jigsaw-puzzle.

The rational arguments are prompted by a desire to maintain our faith in an absolutely objective reality. I am arguing that we shall not begin to investigate the central scientific problems of information systems until we drop the pretence that all we need do is open our eyes to see the objective world that surrounds us. To admit that we are making a huge assumption brings us face to face with such problems as the following. *How should we use information so that, collectively, we can form reliable perceptions and understanding of the world on a scale and to a degree of detail that surpasses the abilities of even the most able individual?* If we are concerned with scientific information this is often a central issue, and one that calls for our specialist information systems expertise. When we can produce

information systems that are reliable in this sense, we can safely resume our simplifying assumption that we can look out upon an objective reality. Objectivity is socially constructed by a consensus that meets certain criteria of clarity and honesty.

There are similar problems where information is used in other fields. Unfortunately we are currently, and for obvious historical reasons obsessed with only one of the many central information systems problems: *How can we be more efficient in office and factory?* Important though this problem may be, it should no longer be our principal or sole professional purpose. The way we are employed or funded, the common reluctance to confront issues with an ethical flavour, as well as the history of our discipline, all conspire to keep our noses on the purely economic grindstone. But we have responsibilities to the wider society, not only to the accountants. We should enquire: What information systems can make us feel safe when living on the lee side of a nuclear power station? In health care we should ask: What information systems enable us to make (economically, ethically, . . .) sound choices when given limited finance but faced with unlimited demands for expensive treatments and drugs? In the law, our use of information is critical for implementing policies democratically arrived at, rather than any arbitrary and prejudiced decisions of those administering them. The list could be extended, and contemplating it should prompt a reappraisal of our subject.

Although we remain responsible individuals, the starting point for our perceptions, cognitions, values and customary behaviour is found in society where they evolve, are maintained and put to use through the communication of information. Our discipline should elucidate both the healthy and pathological functioning of the necessary information systems.

Note that I am not suggesting that an information science should attempt to make the choices governing our knowledge and culture. Everyone should be involved in making these decisions aided by the insights from the fields of epistemology, jurisprudence, ethics, and so on. Our task is limited. It is a task that scarcely existed before the 20th century brought its rapid development of electronic information technologies. Our limited task is to examine how best to give effect to the choices society makes, given the available information technology.

Justifiable and unjustifiable social relativism

Strong rational objections to my position come from opponents of relativism. No better example may be cited than the widely read and powerful attack by Sokal and Bricmont (1998) on the epistemic relativism propounded by a number of leading, French post-modernists (Lacan,

Kristeva, Irigaray, Latour, Baudrillard, Deleuze, Guattari and Virilio). At first sight, it will appear that their criticisms apply to the position I am adopting, especially as those under attack are, to varying degrees, writing on semiology.

Sokal and Bricmont attack the post-modernist claim (p. 50) “that the truth or falsity of any statement is relative to an individual or to a social group.” They also complain (p. 2), of relativism, that “after mocking the old-fashioned ‘dogma’ that ‘there exists an external world, whose properties are independent of any individual human being and indeed of humanity as a whole’, it proclaims categorically that ‘physical “reality”, no less than social “reality”, is at the bottom a social and linguistic construct.” Obviously, I have some explaining to do in this paper!

In fact, personally, I support Sokal and Bricmont in their criticisms of lax relativism. I am just as ready as they (and my FRISCO colleagues) to trust that a single reality does exist independently of us all. But they also accept, along with everyone since at least the 18th century, that *absolute knowledge* of the natural world is not available because fallible observers are always involved, their senses being less than perfect and what they see being influenced by all kinds of presuppositions. The difference between their position and mine lies in their absolute belief that, despite all the difficulties, an objective reality does exist, whereas I want always to emphasise that *belief* in an objective reality is just that, a belief, no matter how sound it appears. If I were working on a natural science, I would probably have little interest in this ‘quibble’. But I am not. As an information systems scientist, I want to focus our attention on questions about those beliefs. Why and how do we acquire them? How do we *justify* such beliefs? How can we create a body of knowledge we are prepared to treat as “independent of any individual human being and indeed of humanity as a whole.”? Questions of this kind are further core issues for a science of information systems.

Take only our contribution to the natural sciences. In this era of electronic information technology we should be concerned to build and maintain the quality of information systems that will give both scientists and the ordinary citizen confidence that claims to have scientific knowledge are justified. We have enough examples of relevant problems today: mad cow disease, genetically modified organisms, global warming, to mention only three. We are not dealing with 19th century scientific work conducted by individuals or small teams in single laboratories, placing their findings before learned societies many of whose members could subject them to informed criticism. No. We are dealing with teams of various experts scattered over the globe, some working in the open, but many in highly confidential projects in commercial organisations. The openness and

informed criticism typical of earlier science is loaded with political and economic pressures which often find their expression in the designs of information systems to gather, select, interpret, propagate, keep secret, publicise etc, not only facts and theories but also funding proposals. The lax relativists would seem not object to each chemical company, each ecological pressure group, each religious group, and each ethnic community, having a diversity of incompatible opinions on the scientific knowledge pertaining to problems of the kinds cited above. Our job, while working on the information systems employed in these domains, is to ensure that the information is being used in ways that can be *justified*. And, of course, understanding what *justification* and its close relative, *responsibility*, should mean are further problems for us to investigate.

Of course, mentioning responsibility brings us back to the ideological distaste some people have for the notion of Society. They usually claim that talk of Society allows people to evade personal responsibility. Indeed that may happen with sloppy political argument, but we are here concerned with scientific argument about information systems, and from that point of view, society and responsibility are closely linked. If, as I argue, it is wise to accept that knowledge is socially constructed, its justification depends on our being able to make people answerable for their contributions towards creating it, maintaining it and using it. A society with coherent perceptions, beliefs, patterns of behaviour and values depends on the exercise of responsibility. The objectivist is tempted to take too much for granted. The notion that knowledge of a world can be “independent of any individual human being and indeed of humanity as a whole”, underlies the impersonal style required of us when at school we reported our laboratory experiments. Good science comes from constant testing by those observers and experimentalists responsible for maintaining it, and from the rest of us checking up that they are exercising their responsibility.

Perceptual norms operationalised

The view I present is a rejection of the Core View and its definition of a “**thing** [as] any part of a conception of a domain (being itself a “part” or “aspect” of the “world”).” ‘Conception’ is a vague, mentalistic term and incapable of operationalisation. Instead, by adopting an actualist ontology, the radically subjectivist position I recommend, we can be far more precise. In this case, things, entities, properties, relationships, attributes, everything indeed, is an instance of an affordance. Real things (in this sense of being invariant repertoires of behaviour) can be understood in operational terms – you can take someone by the hand and introduce them to the behaviour, demonstrate it to them, talk about it and test their competence in applying it. We dispense with the inaccessible processes of perception and conception. It

is all in the public domain and open to empirical study (though not always without difficulty).

The term 'norm' is also mentioned but not used in the Core View: "Norms are socially agreed rules affecting and to a large extent directing the actions within an organisational system." This suggests that norms are built using more primitive components. Sometimes this is the case, but many of our perceptual norms, especially the physical affordances we distinguish, are basic components of our constructed reality. Some legal norms, which are expressed by explicit rules, create parts of our social reality, copyright and marriage, for example. When certain things are done, instances of these come into existence. If a couple makes specified vows in an appointed place before a recognised person and some witnesses, a marriage will start. If a literary work is published, subject to a few conditions, a copyright will start together with the ownership of it, usually by the author; the copyright will finish its existence 70 years after the death of the author, by virtue of another rule. Thus we can extend Gibson's notion of affordances by adding norms which define the invariant repertoires of behaviour that constitute our social world.

There are other kinds of norms beside the perceptual ones. Even more fundamental than perceptual norms are the evaluative norms, which determine whether or not some invariant repertoire of behaviour is worth recognition. The temporal boundaries to the existence of a cup will be fixed by our judgement whether or not it serves its valuable role as a drinking vessel. (Note that a curator of ceramics will still regard an antique Wedgewood cup as a cup, even when broken, because her criteria concern no the use of the cup for drinking now but its having been used that way so that now it tells us about those historical times. The curator's meaning for 'cup' is thus different from the restaurant manager's meaning.) If the link between sign and reality belongs in someone's mind then it may be crucial to know which mind is involved. We shall look at some other kinds of norms later because they play a central role in this alternative framework of IS concepts.

Whose values? Authority and responsibility

The new formalism I am offering, called "Norma" as it represents norms and affordances, is built on the structure

Agent behaviour

This makes it clear that an agent is responsible for the world constructed by recognising useful repertoires of behaviour. The ultimate agent is Society. If we want to operationalise a common sense meaning then we can appeal to

Society to instruct us. The law has to be interpreted ultimately in terms of common sense meanings and who better to provide them than ‘the man on the Clapham omnibus’, to use the words of in a famous dictum by an English judge. If we can employ a dozen, representative, ordinary members of society to demonstrate what some expression means, there is no need to introduce perceptions and conceptions into our theory.

Of course we often need meanings that are far more precise than any random passenger on the Clapham omnibus can furnish. When a person is sick, an insurance company may need the judgement of a qualified medical practitioner to determine the start of this period, so we use the judgement of a person trained for ten years in this skill to provide it. In this case we cannot find explicit rules; we have to rely on the norms that have been internalised by someone. Different people, with different values, may give us different meanings. Indeed only when we know well enough the agent responsible for making the judgement behind the choice of a term or governing its interpretation, can we grasp the exact meaning being employed.

Sometimes the judgements of several people will be needed. For example, the start of a contract will require the consent of both parties. In other cases, we use explicit rules to determine when things start and finish: copyright, for example. These norms have conditions that refer to many other affordances, which, in turn, depend on the judgements of other agents or other norms, which, in turn, . . . This recursive pattern spreads responsibility for meanings over many individuals. For example, copyright law depends on the notions of a literary work and of publication and other norms help to define them, but ultimately we have to turn to the passenger on the Clapham omnibus, perhaps in the role of a jury member. The meaning in such cases also depends on those who framed the laws and those who administer them. What is so important in a this formalism for modelling information systems is that the structure:

Agent behaviour
and its associated start and finish events,

compels us to be explicit about the temporal boundaries of everything, and about the agents who are responsible. The formalism used in the Core View does not even raise these issues, let alone provide a solution.

The start or finish of an affordance may be determined by a person, a group or a norm. This we call the start or finish authority. If we want an explanation for any state of affairs we should turn to these authorities and question them or those who made them because they are responsible or answerable. In more complex cases, an affordance established by a norm can involve many other affordances. The agents determining *their* starts and

finishes also carry some responsibility for the collective meaning. The making of reality and its representations is far more complex than the Core View suggests, nonetheless these complexities can be analysed with formal precision.

Time

Time has been introduced into the discussion. The importance of knowing when an affordance starts and finishes, has been noted. If one wants to be clear about the meaning of a term, then think of it as naming an invariant repertoire of behaviour, and next ask what conditions determine its start and finish. In the case of a cup, I suggested two different meanings, one for a museum curator, and another for a restaurant manager. The approach allows us to be as precise about meanings as we require for a given purpose. If you examine a piece of legislation (copyright law will do) then you will find that a large part of it will be concerned with providing clear, operational semantics, and consist of norms that specify when things start and finish. Time, as we know it, is bound up with our perception of affordances.

This observation is important because the analysis of affordances does not produce a static model but it also includes the dynamics. Everything, but everything, has a finite period of existence. Some things may last a very long time but they have a start and a finish in every case. This is evidently true of numbers and other, supposedly eternal mathematical concepts, once we remember they are socially constructed. We may produce *theories* about eternal things but those theories belong to a socially constructed reality which, therefore, cannot exist outside the existence of the society which has constructed them. We are so accustomed to thinking that the natural world is eternal and subject to laws that never change that this notion will be a strange one. I am quite as ready to adopt this belief about the natural world as anyone. However, as someone even more interested in information systems, I am also prepared to admit that all the science I know is a symbolic structure with a history of how it was made, maintained and changed, including the discovery and recording of facts, and the conjecture and refutation of theories. These are information system issues that fall within the scope of our scientific field. Having to be clear about time makes us consider important issues that otherwise escape our attention.

Ontological dependence

Norma, the formalism I am introducing, provides us with another powerful instrument for understanding time. It is the logical relationship of ontological dependence, which, as far as I have been able to ascertain, has no place in any other formalism. It arises because certain affordances are only

available when others have been realised. Thus being upright makes possible a range of affordances such as walking, shuffling about, running, jumping, none of which is available if one is seated. But being seated allows one to take one's feet off the ground for extended periods, an affordance denied to an upright person. So we can elaborate our simple wff:

(Agent behaviour) behaviour

where the original agent is modified by the first affordance to become a different kind of agent for whom the second behavioral repertoire is available. Thus we can write

Stamper upright walk

We also have some affordances that depend for their existence on the coexistence of two others. Thus I may separately realise the affordances of paper and of pen but by having one while I have the other available, I have the affordance of writing

Stamper (paper *while* pen) writing

Using a comma as an abbreviation for *while*, we can specify the well-formed structuref as

Agent (behaviour, behaviour) behaviour

The behaviours to the right are ontologically dependent on the antecedent affordances to their left. (It may seem strange, until one is accustomed to the idea, that ordinary objects such as paper and pen can be regarded as behaviours. Remember that every single one of the affordances we discover or learn is an invariant repertoire of behaviour. The categories of entities, attributes and relationships, beloved of database people, vanish, subsumed under the one category of invariant repertoires of behaviour.)

When an affordance goes out of existence, so do all its dependants. These dependencies express important time constraints that can be expressed in predicate logic but only with notational clumsiness. The relationship of ontological dependency applies recursively, of course, so quite complex structures can arise, and in the case of social reality they usually do. But before taking this concept further, it will be convenient to introduce another.

Universals and particulars – the model includes the meta-model

On a particular occasion one may stand. On another one may sit. When one is seated one still has the *ability* to stand and when standing there is the prospect of sitting once again. The *ability* to realise some repertoire of behaviour is itself a repertoire of behaviour, having its own start and finish, each with its own responsible authorities.

An ability relates to a realisation as a universal relates to a particular, or, in data-processing terminology, as a type to an instance. If we want to construct a universal model for an information system, then we do so using types. Similarly, we can introduce a simple notation for ontological dependencies between abilities or universals that are inherited by their particular realisations.

Put these two ideas together. First, an ability is an invariant repertoire of behaviour just like any other affordance. Second, a universal is the ability to realise its particular instances. The result is a unification of all the levels of a meta-model hierarchy as deep as one requires.

This is best explained by an illustration. Take the social institution of marriage.

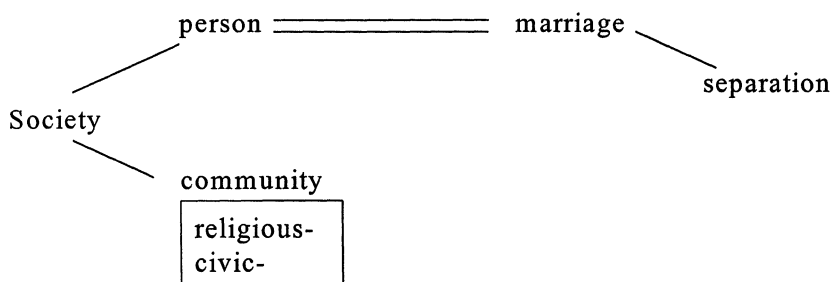


Figure 1: The elements of the marriage institution

The graph is directed with the implied arrows representing the ontological dependencies being left-to-right. Note that marriage requires two persons as antecedents (two antecedents is the maximum permitted) and that they can be identified, in that relationship, by role names. If one partner dies, the marriage finishes its existence. There is also the notion of a separation, which, as well as its common sense meaning, has acquired a number of others, through a series of common law judgements and Matrimonial Acts, which created and evolved the legal concept of being judicially separated. This universal affordance came into being in the 19th century in the UK. There are many other universals that have recognisably finite lives, created not only by the law but by science and technology, as well as by inventors, artists and the flux of fashion. Many of these can be precisely and formally distinguished by specifying the relevant authority.

The time constraints of ontological dependency apply at both the universal and particular levels. Take any affordance and trace through all its antecedents back to Society at the root, and one will have a lattice, called its 'stem'. For the chosen instance of the affordance to exist, appropriate

instances of every antecedent and the corresponding universals must also exist. This complex set of temporal constraints is a powerful check on the meaningfulness of data supplied according to this model.

Notice that, when reading this diagram, one should keep in mind that associated with every universal and every particular are start and finish times and start and finish authorities. These can be provided separately and play a role in a social system analogous to a program for a computer. Indeed, we have taken such a model of a social system and derived directly from it a computer system to handle relevant data and support the functioning of the social system. This model though not devised for software engineering, turns out to be suitable for that purpose. As a by-product of modelling social behaviour, one also carries out the information engineering.

There is also a generic-specific hierarchy showing the various communities that have their own meanings for ‘marriage’ supplied by exercising their authority over the start and finish of that relationship. It is quite feasible for a couple to have multiple marriages. For instance I have friends who have three different marriages, two under the authorities of the different religious communities they come from as well as a UK civil marriage. Thus, Christian marriages cannot predate the 1st century, Islamic marriages, the 7th century and the different kinds of UK civil marriages, the various Acts of Union with Scotland, Ireland and Northern Ireland.

Thus semantics is not a “nebulous” idea. When approached in this manner, any desired degree of semantic detail and precision can be achieved and expressed formally. Although it would be easy, I have not bothered with the superfluous, space-consuming task of translating these specifications of meaning into an impressive mathematical symbolism. I approve of formality but mere formalisation is only decorative. Chapter 4 adds nothing to Chapter 3, beyond a final check on its formal potential.

Here and now

Look again at the model of the marriage institution. You will notice that the elements on the right can only exist during a period of time that falls within the existence of elements to their left. If we use this model to give ourselves a picture of the world that actually exists, it can tell us only about the here-and-now. Although we can use the model to organise data about the past and the future, the model, with its strict temporal constraints given by ontological dependencies, allows us to say nothing about how the past and future are related!

This dreadful limitation appears, at first, to be a defect in the modelling method. In fact the opposite is true, as it forces us to attend to a central problem in the science of information systems: the creation of time.

The Core View's notation has no problem dealing with times past and future. Times, they would say, are either instants or intervals, which we can include with other objectively existing individuals in predicate expressions. For example, the two-place predicate

born (Matthew Stamper, 10 January 1987)

relates my son to an interval of a day. Days exist in possible worlds if not here and now. Objectivist semantics gets into trouble otherwise. For example, what is meant by 'person' is given *extensively* by the set of all persons, so if we do not also count those persons who exist in possible worlds but only those in the present world, the exact meaning of 'person' continually changes as individuals are born and die. As one who always demands an empirical foundation for a science, I feel uncomfortable trying to explain meanings in terms of a purely imaginary domain of possibilities.

The actualist alternative I am offering has no place in its reality for anything except the here-and-now. The past and the future, and indeed things over the horizon, do not exist here-and-now. *We have to create them using signs to represent them in the here-and-now.* This is the most miraculous, magical thing about information! It is our job to explain the processes involved. Rather than to copy the natural scientists who can remain contentedly working with the old illusion of a real, physical time-dimension, which they "see" through the transparent lens of information systems they use, we must investigate the "optical" properties of the information lens.

We construct time by imposing an order on the names of events. For example: I know my son, Matthew, and I know he had a start, so I can construct an expression, "the start of Matthew". This is a name for an event. Many such expressions representing start and finish events can be combined into an ordered pattern. Obviously a collection of these expressions would be a hideously obscure way of talking about time in most circumstances, although you will often find yourself among family and friends talking about time this way ("When X cut his first tooth, just before we moved to York." etc). For general communications about time, we use a calendar and chronometer to generate a stream of conveniently named events or intervals such as "10 January 1987". Today, my knowledge of the past depends upon my being able, here and now, to construct such marks on paper or in memory (neurological notepaper) as, "the start of Matthew" or "10 January 1987". To obtain the time dimension we use in physics, we arrange those events or intervals in an order or partial order, with the appropriate

mathematical properties. That is how we construct time. We can have no knowledge of time unless we use signs to represent the starts and finishes of what we know here and now. We cannot experience an event as an affordance or repertoire of behaviour. Once it is here, it is gone!

Quite different are the processes of beginning and ending. These can be experienced, as I witnessed, and his mother experienced, Matthew's birth, the beginning process that successfully started him. This raises the issue of causality.

Causality

Before tackling this subject, let me remove a common misunderstanding about the schemas illustrated in Figures 1, 2, 4, 5, 6 and 7. These do not represent causal relationships but existence dependencies. The affordances on the left do not cause those on the right. Take any affordance, and you will see that it makes *possible* those to its right, while those to its left are *necessary*. People working with computers, who are raised on a diet of flowcharts, dataflow diagrams, logical sequences and so on are in greatest danger of falling into this error.

“So where does causality fit into this schema?” you may ask. Causality, quite different from, but related to time, is another semiological construct. Our familiarity with physics tempts us to believe that causality is quite a simple, physical notion. In physics causal connections take place in the here-and-now, governed by simple physical laws.

Causality is a complex notion in the social world. Social cause-and-effect depends upon information (signs) to link the past to the future. For example, when we look for the cause of an accident, we do so in order to avoid similar harm in future by changing the behaviour of people and their environment, and often to recover payment for the damage done. Consider the judicial enquiry that found a Swiss chemical company so negligent that it caused the escape of dioxin and a tragic loss of life Seveso. Far more was required than tracing chains of physical cause and effect; such a decision called for an investigation of who knew what when, what assumptions should have been made by reasonable persons, who was responsible for training those who controlled the industrial plant and others who handled emergencies. Causality in the social domain goes beyond physical variables into the problems of human responsibility.

Returning to physics, we should note that it deals with causality by formulating laws of a universal kind. A particular cause and its particular effect (one billiard ball setting another in motion) can never be experienced again. We can only create a similar situation and predict the consequences

by applying the law. Physical laws and theories are signs. So, even physical causality is a semiological construct.

In our everyday struggle to make sense of causality, we use our experiences of processes that can result in things starting or finishing. We learn to relate the *beginning* and *ending* processes to the *start* and *finish* events. We must be careful not to confuse the processes of change with their results. But we must learn to relate them in order to develop a notion of causality. The birth was the beginning process that successfully concluded with the start of Matthew, his mother experienced it and others witnessed it. That birth process also started and finished, and had its own beginning and ending. Beginning and ending are common sense ideas that enable us to talk about everyday causality. Clearly the beginning of Matthew was a cause of his existence. But note that if we count his existence starting with the finish of his birth, he never coexisted with his birth process. So common sense causality cannot be captured in an ontological dependency schema such as that in Figure 1. The theory needs to be extended.

Let me reiterate that such a schema represents only what exists at present. Unless a cause and its effect coexist, they cannot be treated as the antecedents of a causal relationship. This is because *the schema defines only existence dependencies*. Sometimes there is a moment of coexistence but in general there is none. “So how do we deal with causality?” you may ask, “How do we link the past to the present and the present to the future?”

Making sense of the real world certainly requires a notion of causality. The ontological theory outlined here succeeds in forcing us to recognise the central role played by signs (information) in creating a notion of causation. Whenever we talk about a cause, even a common sense cause, we construct a little theory. The natural sciences consist, to a large extent, of theories of causality. In physical laws, cause and effect coexist. But in a social context, the recognised cause and effect seldom do coexist. For example, causality is often a legal issue, as in the notorious Seveso case. Just as we can take an affordance, *x*, realised, here and now, and name its start, “start-of-*x*”, so we can name a process, the “beginning-of-*x*”. This enable us always to satisfy our common sense belief that every effect has a cause. You will find that this analysis in terms of signs, whether legal findings or physical theories, removes much of the mystery of causality by forcing us to look carefully at how we use information.

For anyone constructing an information system, this analytical device is a great help. The making and changing of the social world is accomplished by information. Indeed, the only thing that information produces is a change in one or more person’s disposition to act, and normally this amounts to a social change. More often than not, this is accomplished by a conversation.

Pragmatics – time and intention

Pragmatics comes higher on the semiotic ladder than semantics, even so it is far from nebulous, being capable of precise analysis. We have seen how we can understand time and causality in terms of our use of signs. Now let me show you how the pragmatic properties of signs are essential for creating the past and the future and binding them together.




Figure 2: A future created by intentional acts making use of meaningful signs

Take the marriage example. Marriage is not spontaneous and random, unlike the mating of rabbits. Someone has to propose that a marriage takes place and one or more appropriate persons have to accept the proposal. So we have to extend the diagram to include these “communication acts”.

This shows marriage as an antecedent to a proposal but linked by a broken line. This is because the marriage itself is not the antecedent (there is no point in proposing an existing marriage!) but the antecedent is a sign representing a possible future marriage. The future marriage exists here and now only as information. The various intentions are expressed by existing persons uttering actual words, such as “I propose that John and Mary should get married.”. The antecedents are the speaker and a *sign* representing the marriage. We must be careful to distinguish the speech act, which is a process, from the intention that it begins. The proposal remains on the table until, continuing its existence until, perhaps, withdrawn by another speech act. The proposal forms a small but significant element of the social world. It represents an intention on the part of the person who created it. The nature of this intention is highly sensitive to the social role of the maker – John, Mary, one of their parents, another relative, a marriage broker, a friend. The existence of the proposal may well bring into existence all kinds of other attitudes of expectation, obligation and evaluation, which have not been included in the analysis. All these details could be formalised as norms if that is the point of your analysis.

An acceptance, or a number of acceptances, will be necessary before a wedding is arranged. A proposal and acceptance by youngsters under 18 would require acceptances also by their parents or guardians. This is typical

of how, in the here-and-now, we plan, to create, some future world. It is not accomplished by natural processes nor by passing meaningless, intentionless messages from computer to computer. It necessarily involves responsible people using signs that have meanings that only groups of people can supply. It is a social process. Information is the instrument we use for making, preserving and using our social reality. To study information systems scientifically is to engage in a social science. Interestingly, we can make it quite formal and precise, just as precise as anything in FRISCO Chapter 4.

The social level of the semiotic ladder

Now we come to the social level where Falkenberg thinks we are most at risk of being nebulous. Let us test this opinion.

I propose that we adopt a formal and precise definition of society as people and groups of people behaving according to shared norms. Most norms are shared. Even when norms govern the behaviour of individuals, other people have to know about those norms. For example, a head of state may be the only person subject to certain norms, but that role only functions because it is widely understood and taken into account by everyone else.

Norms only exist when they govern the behaviour of people. They can be represented by rules but, remember: the map is not the territory. So, although some norms may be explicit, as in rules and laws, others are implicit, unspoken and even incapable of being expressed but, nevertheless, observable in use. We assume that they all have the common structure:

IF condition THEN norm subject
ADOPTS attitude TOWARDS some thing or proposition

The condition specifies the knowledge of a state of affairs that triggers the norm. At such a point the person or group agent subject to the norm will adopt a direct or indirect disposition to act towards a thing or proposition about the world. Based on the formalism I have introduced, we can express any norm quite precisely and formally, as though we were writing a logical version of a legal rule. But less formal precision may be substituted by greater clarity if a person, rather than a machine is the subject of the norm.

There are four important categories of norms of which we have so far encountered two. *Perceptual norms* generate the attitudes of deeming or accepting the existence of something (contract, copyright, marriage etc). *Evaluative norms* require attitudes of liking, disliking, hoping, fearing and so on. *Cognitive norms*, like those of science or our common sense ideas of how the world functions, generate attitudes of belief or expectation. *Behavioral norms* dispose the subject more directly to act by the attitudes of obligation, permission, prohibition and liberty.

There is also the subcategory of norms having no condition and a fixed attitude. These are called ‘attitudes’, quite simply. “Everyone abhors murder.” for example.

Society is very complex and subtle. So is the physical world, as any engineer will tell you. The elegant simplicity of physical theories can help us to solve engineering problems but more goes on in everyday reality than a physicist can model. Similarly, I would not pretend that whole, real societies could be modelled using norms, as I outline here. However, certain problems can be elucidated using this kind of model, especially in cases where it is normal to use explicit rules, as when we design legal or contractual systems or when we want to apply a computer. Typically such limited social systems we call ‘organisations’ or, in their more general form, ‘institutions’. As information scientists, in addition to formulating broad and general theories about information, it is our duty to identify interesting and instructive institutional phenomena, by analogy with the work of physicists. We can then apply these results to engineering particular organisations in which computers and other technology may be able to play a useful role.

Information fields and knowledge

We are now in a position to replace the vague, unoperationalised, definition of knowledge in the Core View: “**Knowledge** is a relatively stable and sufficiently consistent set of conceptions possessed by single human actors.” by one that is operationally clear and precise.

Instead: **Knowledge** comprises norms and attitudes, whether explicit or tacit. This includes knowledge of what exists, what is right or wrong, good or bad, of what to believe about the world, and of how to behave.

This definition is far broader than the philosophers’ “justified, true belief”. Otherwise it will not serve the purpose of an information systems science..

The Core View has the definition: “**Information** is the knowledge increment brought about by a receiving action in a message transfer, i.e. it is the difference between the conceptions interpreted from the received message and the knowledge before the receiving action.” This indicates that we probably need one meaning for “information” that relates it directly to norms (alias knowledge). There are many other precise meanings of “information” (Stamper 1973) but they cannot be found using notions as vague as “conceptions”. Instead, following the Broad View, we prefer the notion of a sign as a primitive because it can be defined ostensively. Signs interact with norms, as this figure shows:

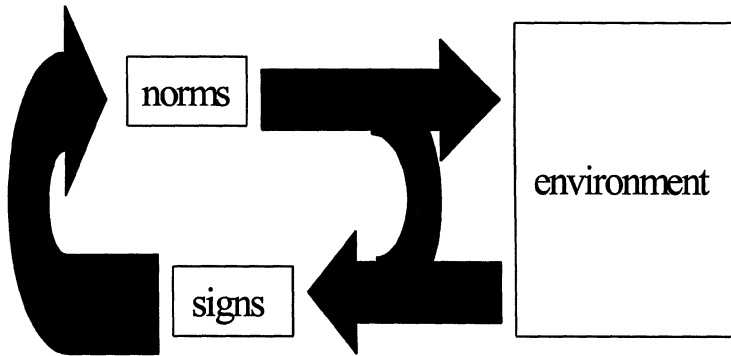


Figure 3: Information field - the interplay of norms, signs and environment

To recognise when a norm should be triggered, the subject needs information (signs) relating to the condition. The resulting attitude may not produce an immediate outcome but sooner or later will be revealed in words or comportment, or sometimes translated eventually into action. In either case the result will be more signs. In addition, observation of the environment (the application of perceptual norms) will also generate a stream of signs. As an analogy, we might consider knowledge or norms as a form of potential semiological energy and signs as the kinetic form, constantly being transformed into one another.

This observation is of major importance to the engineering of information systems. It allows us to introduce the notion of *information fields*. We normally use models of messages flowing between operating nodes, or the dual models of data-structures, which we use to scan storage media. An information field is a social group governed by a system of norms. The norms determine what should happen in that domain and they, in turn, determine what signs are required to fire them and what signs they produce. Thus the information flow is just a logical consequence of the norm structure.

The fields are more fundamental to the information system than the information flows. For any given information field, there may be thousands of different but adequate solutions to the message flow and record search problems. The typical information systems model produced to define the requirements for a computer-based system will tend to reproduce or closely resemble an existing dataflow, often one that has evolved to suit a different technology, and, most likely containing all kinds of irrelevant features for historical reasons. The information field approach cuts through the accumulations of time to capture the essence of what is required to do the job, making it far easier to find innovative solutions to suit whatever new

technology is available. Better still, we can generate supporting IT systems directly from the norms specified in the field model.

Summary and comments on formalisms and nebulous ideas

This alternative to framework of information systems concepts in Chapter 4, forming FRISCO's Core View, could have been expressed with as much formality. However, striving at symbolic niceties does not add to the content. The Social View leads to an alternative because it tackles head on the difficult ontological and epistemic issues which the Core View evades. The ontology adopted is that the only world to which we have access consists of the invariant repertoires of behaviour (affordances) we recognise. The epistemological principle is that our knowledge is founded on direct perception of some basic affordances, supplemented by an understanding of many others that Society has discovered over the centuries. The adequacy of such knowledge depends on its practical effectiveness. To know anything we must act. The only world directly accessible to us is the here-and-now. Information allows us to construct the rest.

The following syntax governs most of an ontological structure, or a realisation, based on these assumptions:

- 1) realisation = Agent | Agent affordance
- 2) affordance = (affordance *while* affordance) | realisation *sign* | affordance *ability*

The concatenations defined in 1) are the ontological dependencies. The ontological antecedent can be derived from others, as expressed in 2). It may be restricted to the joint existence of two antecedents, it may be anything that functions as a sign for a realisation, or it may be the universal or ability to realise a type of affordance. When we draw an ontology chart, we work with the universals or abilities.

This syntax does not cover the generic-specific hierarchy, which consists of cognitive norms, nor the determiners, indicated by the prefix #, which must have an ability as an antecedent. To extend this syntax to cover the construction of norms requires other functions in addition to *while*. These include *orwhile*, *whilenot*, *before*, *after*, *whenever*, for example. That topic lies outside the scope of this paper.

However, we must note that every realisation has a number of mappings to realisations. These are

sort that maps to the corresponding ability or universal
antecedent 1 that maps to one ontological antecedent

antecedent 2 that maps to a second ontological antecedent, if there is one
start that maps to a realisation simultaneous with the affordance's start event
finish that does the same for the finish event

start authority that maps directly to the responsible agent(s) or indirectly via norm

finish authority that does the same to identify the responsibility for the finish

One point of great importance that has been mentioned deserves to be stressed again. There is a maximum of two ontological antecedents. This constraint is very demanding and the principal reason for the ontological structure being a kind of "semantic normal form" or canonical form for a given problem. It is canonical, not in the sense of a mathematical or logical canonical form, but in a way that can only be justified empirically in the final analysis.

I trust that the above account of the framework arising from the Social View will shoot out of the sky the *canard* that the semantic, pragmatic and social levels of semiotics are nebulous. Admittedly they are frequently discussed in a rather woolly terminology, but that is not the case when we employ this formalism that captures the ontological dependencies among norms and affordances. The case study helps to illustrate this point.

A REWORKING OF THE JAPAN WINES CASE

The Core View of FRISCO concludes, in Chapter 5 9 (q.v.), with a case study. It is now be reworked on the basis of the Social View to demonstrate the practical advantages of the new method. You will notice several improvements. First, for the information system theorist, the new approach models the meanings and the intentions of the messages as well as the relevant norms used in running the business. Second, the solution, while including models of data-structures and dataflows that the Core View produces, also captures, in a very compact form, the minimal, substantive organisational requirements, which helps users to understand their problem in information terms. And third, in the realm of software engineering, the specification is precise and formal enough for it to generate the necessary computer system directly.

The problem statement you will find in Chapter 5, section 1 of the Report. Its formulation betrays a mind-set that treats a business as message-passing system, paper-based in this instance. To avoid the presuppositions this implies and use the information field approach, we need to distinguish between substantive and semiological affordances. The substantive things

are what the messages are about, while the semiological things tell us about them – they are representations. For example, the wine in *stock* (which people drink) is substantive and we signify it using an *inventory* (which people read). Another substantive would be a retailer’s need for a quantity of wine (based on his assessment of the customers’ desire to drink) which he will represent in on an order-form. The contract is a substantive component of the social world, a mutual commitment to exchange wine for money, which will probably also be represented in a document confusingly also called a “contract”. Briefly, the map is not the territory, and we want to model separately the business of dealing in wine and its associated paper-chase.

We have a bad habit of confusing the substantive and the semiological. Anyone trained in “inventory management” when they really mean “stock management” is likely to overlook this distinction. The case description includes this misunderstanding in S2: “The business of this centre is to manage the inventory of products and to distribute products to retail shops corresponding to their orders.” Thinking is easily confused by sloppy semantics.

The affordances

We begin by listing, in no significant order, some of the obvious, relevant, substantive affordances appearing or implied by the problem description:

#Japan Wines Inc. (a particular, while the rest are universals)

company	nation	distribute	wine	winery	shop
product	stock	item	quantity	warehouse	deliver
place	truck				

Next we make a list of affordances concerned with passing messages and operating on them:

phone	inventory	order	process	check	item
file	orderline	delivery ticket		back order	

Order may need justification for its inclusion in this second list, but statements 4 to 11 make it clear that they concern a document having a syntactic structure (composed of lines or items which refer to products) and a physical representation. There is a quite different substantive notion of order as an intention expressed by a communication act.

The information flow or message logistics

If we turn to the solution in Chapter 5, beginning on p. 126, we find an analysis entirely in terms of flows. Flow charts 5.2-1 and 5.2-2 give vague outlines of the logistics of goods and messages but chart 5.2-3 eliminates the goods. The partial nature of these models is explained as a way of circumventing “modelling uncertainty”, in other words, the arbitrariness inherent in the modelling method. As you will note, the new approach leaves hardly any room for arbitrariness. The vagueness appears in descriptions of actions such as “Stock administration” or “Orders-to-wineries handling”. We find mention of actors such as “Stock handlers”, “Truck drivers”, “Order handlers”. However, when we come to the detail, supplied by entity-relationship chart 5.2-5 plus the definitions of relationships, these actors have vanished. The messages themselves are not defined but are implicit in the details of the relationships. Nowhere can you see who is responsible for anything. The only place time appears in the model is the date of an order.

The result is a rather vague impression of the logistics, mainly concerning the various documents. Even my business were some kind of document handling (banking, for example), I would at least want to know who sends them, who receives them and where they are at any time. The response I would expect to this complaint is “Go and look in the order file if you want to know where the order is!” Of course, this is a valid implementation detail but a model of the information system at this level of logistics should include an implementation-free specification of document location. This is missing from the Core model.

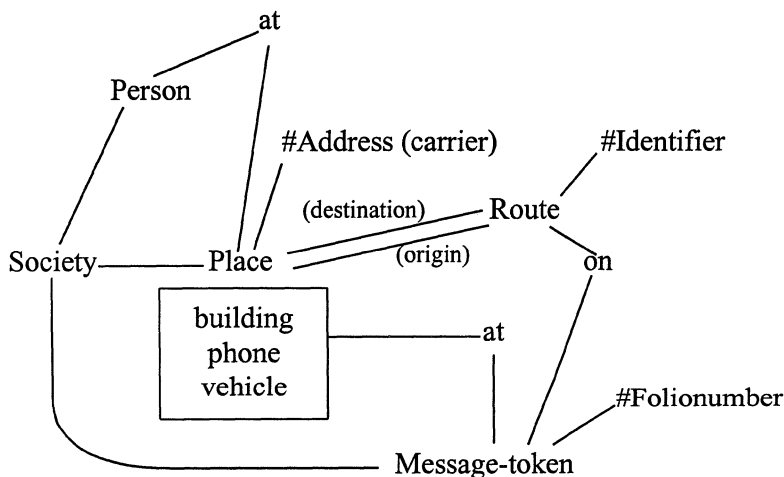


Figure 4: Message logistics

To build the alternative model based on Norma, we construct an ontological schema using the above list of affordances concerned with messages and their flows. There are three levels of analysis we need to include. The most obvious concerns the movements of the physical message-tokens, as shown in figure 4, which omits some obvious details.

Some of the affordances are places, such as a phone, to which I have added other specifics such as building and vehicle, but also person. Places do not have to be fixed. We need to know when a delivery note is on a vehicle just as we need to know when a person has pocketed a message. The arrow from person to place indicates a specific-generic relationship. The name of a place is an address or an equivalent identifier; note the antecedent “carrier” in parentheses, the antecedent agent responsible for assigning the address. We also need to know routes, which are fixed by two places as antecedents with many identifiable alternatives. Sufficient identifiers may select a carrier such as post / courier / email / fax etc. Every message (order, delivery ticket etc.) can be registered as at a place or on a route. It also has a unique identifier, the folio number. “Sign-token” would have been technically a better term to use than “message” because it treats files as large messages, including inventory, back-order file, order file.

When reading the schema in figure 4, one must remember that it represents the universals only, the equivalents of the headings in FRISCO figure 5.2-4 of Chapter 5. You must imagine the sample data for some particulars. Every universal and particular also has its start and finish times and start and finish authorities. So, this Social model pins down all times and responsibilities (such as when a delivery note was raised and by whom). Unlike the perceptions and conceptions of the Core View, this tells us exactly who is making the relevant interpretations.

Also when reading the schema, do not expect that the words should fall into certain syntactic categories. It is not a verbal schema but a pattern of repertoires of behaviour, which are linguistically neutral. The labels attached to the nodes can be in any language or in several. The structure of the graph contains enough information to disambiguate the chosen labels. It is, we believe, an interesting point of departure for research into a semantic grammar.

Syntactic structures of the messages and records

After modelling the message logistics, we have to specify the structures of the messages, stating the syntactic-type (order, delivery ticket) and the components they include (item, order line, company details and so on). See figure 5. The dotted line leading from syntactic-type to message-token is a universal to particular relationship, which their own particulars inherit, so

that a particular syntactic-type, such as an order, is the universal to a subset of message-tokens. Finally, the terminal elements in these structures also have to be specified. Notice that the company is the antecedent for its own syntactic structures, although increasingly inter-organisational standard message formats are being agreed. Formats are not completely specified because this level of detail omits the spatial sequences of the types and the syntactic 'pictures' of the elements. These should be treated separately as implementation decisions.

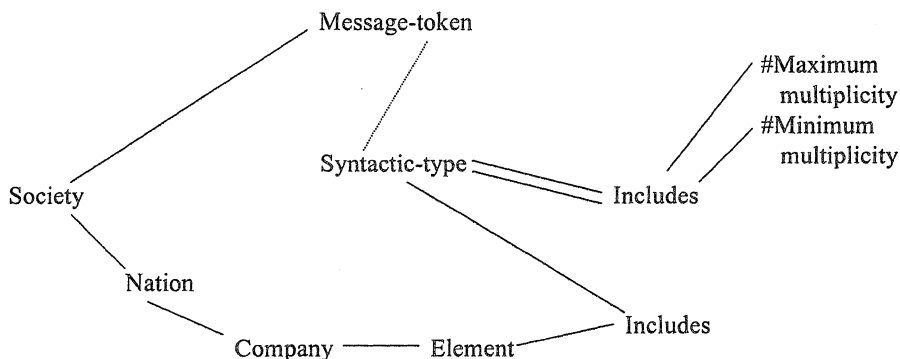


Figure 5: Composition of messages

We could also specify all the processes (check, duplicate, and so on) that have to be performed on the messages, making use of their syntactic structures. All this lies at the implementation level, so it needs to be separated from the core of the specification. Moreover, using this new method, the processes are specified in a different and more fundamental way when we model the substantive system.

The information flow paradigm as used here draws attention to the physical and syntactic levels. We could add a specification at the empiric level by modelling the frequencies of messages of different kinds, the probabilities of errors and so on. Norma and the information field model accommodate all levels of the semiotic properties of a system. However, running a business is more concerned with the semantic, pragmatic and social levels, which are the main focus of the information field paradigm.

The substantive information fields

The Core View models the message logistics but says little or nothing about the core of the business, which is the purchasing, sale and distribution of various products and the management of stocks of them (*not* inventories of them, of course). Any message transactions should be purely a logical consequence of having to perform these substantive tasks. The Core View has no way of modelling them. It relates to selling a computer application to

automate the existing message flow, rather than to the more serious problems of making and discharging of unambiguous contracts that concern the a person running this business. Let us now develop the schema of ontological dependencies for these fundamental, substantive business affordances.

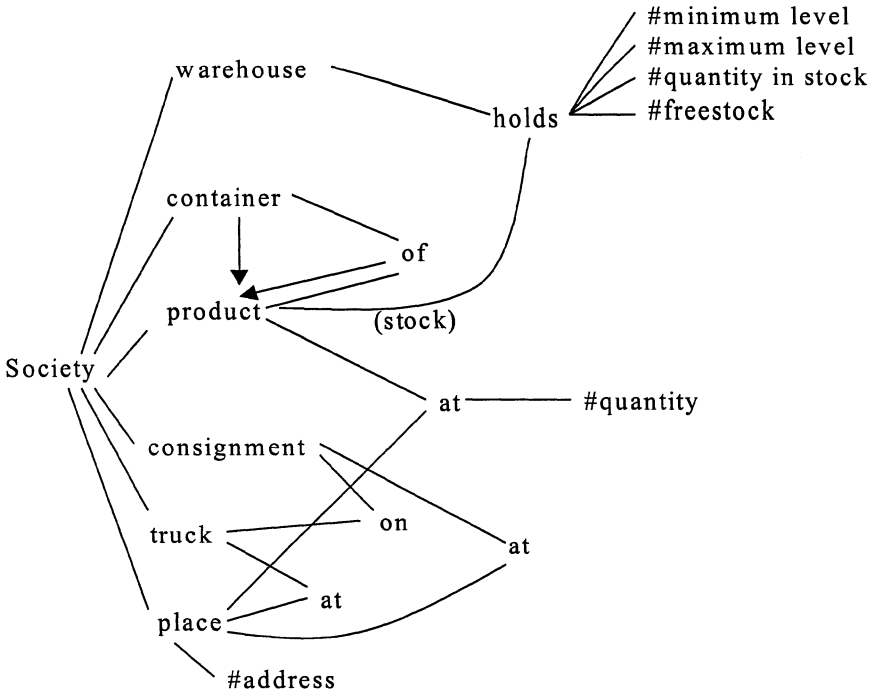


Figure 6: Some substantive logistic affordances

A draft for such an ontology chart is shown in Figure 6. While attempting to construct it, we are forced to recognise the existence of some additional affordances. For example, *product* may mean a type of object (a liquid in the case of wine) which no one can handle unless gathered into an identifiable batch by some form of container. The Core solution on p. 130, says “A specific product, e.g. a brand/vintage of wine, is an elementary thing.” which is an example of dangerous semantic vagueness. Once you are forced to explain meanings in terms of affordances or repertoires of behaviour, you are driven to realise that in this business, your products have to be discrete objects that your warehouse staff can handle. They need information at that level of precision. One cannot simply buy and transport Bordeaux. One has to deal in actual, identifiable barrels or bottles of the wine, otherwise the store worker or driver will not know what to do, and no court would ever be able to resolve a dispute over a contract about anything so vague as “Bordeaux wine”. Hence the schema shows containers of the product. You

can have a bottle of 1995 Chateau Fourtanay, or a one-dozen case of those bottles or a twelve-case pallet of them. You might purchase by the 1000litre cask and bottle it for the retailers. Incidentally, the arrows from container and container of product indicate that these are also specific forms of product.

Goods have to be gathered together as batches or stocks or consignments (three quite different invariants) in order to handle them. Note that a stock or a consignment will remain invariant, although its contents may change. In the despatch bay, a consignment may be built up during the course of a morning as goods are retrieved from their bins and then packed for lifting onto one or more trucks. The schema shows quantities of various products in a consignment at a place, so that the precise composition of the consignment can be handled. The consignment, as delivered, may not be what was sent or loaded. The Core model seems to use "delivery" as meaning consignment in this sense, but it cannot account for partial and incomplete deliveries. Note that a consignment is not a natural thing but a social construct.

On closer inspection, a number of other socially defined invariants become evident, such as *assign* when a quantity of a stock is assigned to a contract, prior to the warehouse adding it to a consignment. The transport company can *present* a consignment at a certain place with the intention to *deliver* it but it may have to try again. The consignment will be consigned (a communication act) by a company to be at a certain place, even before the consignment is made up. The consignment will go through a cycle from being made up to delivered with redeliver as one other status in the cycle. Of course, the 'redeliver' status should appear on one of the documents in the Core model, but it is far more important that the warehouse and transport staff know this status.

Now let us turn to the non-physical aspects of the job. Contracts have to be made and discharged. This always involves making an offer/order and having it accepted. The order will be for the supply of a quantity of a product at some place (and time, implicitly) but always under a contract containing many, often complex rules and conditions in addition to the payment details. So the order will be a sign of all those semantic essentials used by the company placing the order, the party standing in the role of purchaser in relation to that contract. All those semantic structures are shown in the schema.

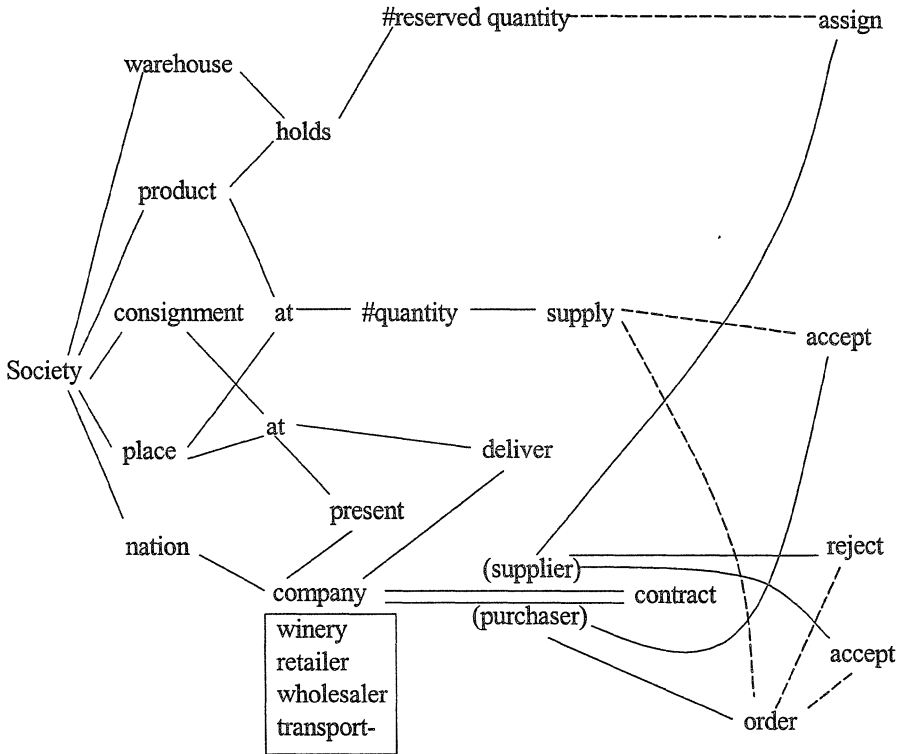


Figure 7: Some legally significant pragmatic affordances

Another sign, referring to that order has to be used by the supplier to accept the order. Then, under provisions of the relevant national law of contract, a contract will come into existence. Notice that this model supplies all the start and finish times and their related authorities that are of importance legally. The FRISCO Core model contains no means of handling these social constructs explicitly. One must have the relevant legal knowledge and apply it implicitly to the messages. In all but the most trivial cases, this is a dangerous strategy.

This analysis of *order* and *accept* transcends the restriction of reality to the here-and-now. For each of these, a broken line leads to one of the antecedents, indicating that it should be interpreted as a sign with that meaning. An order is made using a sign representing some future supply of a batch of a product to a specified place under some contract, an act that does not but belongs in a possible future. On this, substantive level of analysis, we need pay no attention to the physical sign-tokens that represent these communication acts, however, they can be related to the various messages which appear in the Core model. The Social model enables us to design

messages that express the essential business intentions unequivocally, and be quite explicit about their roles as authorities for starting and finishing social constructs, such as contracts or consignments. The Core model has nothing to say about these business issues and, instead, concerns itself with the syntax and the flow of documents or their electronic equivalents.

Note some further details: the box headed *company* defines a generic-specific hierarchy. Measurements, prescribed quantities and identifiers, which we call ‘determiners’ have an initial #. In the context of a certain dependency, an element can have a role name such as *purchaser* or *supplier*, which serves to select the relationship instance (the contract in this case) with the role occupant singled out for attention. Note that *stock* is the role of a *product* that a *warehouse holds*.

A canonical semantic form

The re-working of the case has already illustrated that the new method, based on Norma, is far more expressive than conventional models such as the Core View yields. Take one example of the kind of semantic problem that the new method can handle. There are two related notions: *deliver* and *supply*. In the first case, a company or other agent can *deliver* a batch of a product at some place without there being any legal significance to their action, but in the case of *supply* a contract is an antecedent because the goods are being delivered under its legal provisions. One possibility would be to have *deliver* as an antecedent to *supply*. The behaviour involved in making a delivery is essentially one of physically relocating the goods. *Supply*, it might be argued, may be achieved by the purely legal, communication act of asserting: “The goods at that place are now yours.” a step in transferring ownership rather than relocating material things. Incidentally, the point at which tax is collected will probably depend critically on this distinction. Other subtle distinctions come to the surface as one analyses the ontological dependencies, whereas they would be undreamed of, by anyone using the modelling methods of the Core View. Some important innovations in electronic commerce depend on distinguishing logistic delivery from legal supply. If a meta-model prevents our perceiving such things it will block creative thinking about organisations rather than encouraging it.

Far from being ‘nebulous’ the schema treats semantics with a degree of precision that other methods cannot enforce. The Core View focuses, as do other orthodox methods, on the syntactic level where hundreds of legitimate data-flow solutions for any business problem may be offered. Where the systems designer is trying to devise ways of using computing resources economically, she should be free to use any strategy. However, on the

semantic level we are not dealing with computers but with the repertoires of behaviour relevant to the business and how we name them. These are not things an analyst or programmer can dream up; usually they will have evolved over decades or centuries or even millennia, in some cases. Semantics cannot be toyed with by every imaginative systems analyst. Meanings do not belong to the computer industry but to society and its many communities who make them. The importance of the concept of ontological dependency is that it makes it possible to get to grips with the core of a semantic structure quite rigorously. By adding the norms that define start and finish times, the semantics can be expressed in refined form to any desired level of detail. If the analysis is faulty, careful testing should reveal this so that it can be improved. Orthodox modelling methods usually accept indifferently any number of arbitrary solutions, so they cannot ensure a process of cumulating improvements. The precision is achieved by asking whether something can exist, in the sense of the actualist ontology, finding its ontological time constraints, and then checking that it corresponds to the repertoires of behaviour employed by the people acting in the problem domain. In this way the Social View can generate a cumulative and continually improving corpus of semantics, as objective as our corpus of botanical knowledge, for example.

The enforced rigour is probably (judging by practical experience) enough to result in a canonical form which independent analysts should all reach if they are modelling the same system. This contrasts with the high degree of arbitrariness or “modelling uncertainty” permitted by the Core View. Identifying the socially constructed, canonical, stable, semantic core to the information system brings huge reductions in system costs in several ways. Establishing the semantic core makes it easier to maintain a system in line with continually changing, organisational requirements (again confirmed empirically). The stable semantics must remain invariant whatever syntactic, computer-orientated solution is selected. Hence it serves as a ‘hinge’ around which we can reconfigure an organisation or a collection of legacy systems when we are re-engineering them. The ever-shifting demands placed on organisations make this work continually necessary and costly. The most important re-usable component of any information system consists of the painstakingly negotiated meanings that make the electronic signs and signals useable. Given a few days of tuition, any person of average intelligence can put together plausible looking ER, Niam, Relational or other orthodox data models. Unfortunately, these students are given the impression that they are discovering the semantics of the domain. Genuine semantic analysis is far more intellectually demanding and costly an exercise. But the effort is worthwhile. Huge savings can be obtained if we share a solid and accumulating corpus of re-usable semantics from which to begin the design

of each new system. Finally, as we have found in practice, the uniformity of this analysis (every affordance has the same attributes) makes it far easier for users who have no technical interest in computers. They can understand the ontology charts quite easily and use them to formulate precise information requirements to suit their needs.

A caveat: although I am confident that we are approaching a canonical semantic form, more theoretical work and empirical testing are called for.

Building a computer system to support Japan Wines Inc.

The ontology charts introduced above are rather easy to implement directly. However, an interpreter is of little value unless one can construct high quality semantic schemas. Operations on data held under bad schemas with nonsensical time constraints will only generate semantic garbage. On the other hand, good semantics can be profitably implemented with any systems software.

To read a chart and build a computer system to interpret it, one must remember that the arcs are directed from left to right and represent constraints on existence *not* causal dependencies. This applies both to the universals shown in the figure as well as to the particular instances of them. Any element on the right can only exist if relevant individuals exist on the whole of the stem, the lattice leading back to the root Society at the root. The computer system we have built to interpret this schema (called a Semantic Temporal Data Base, STDB) automatically checks the temporal consistency of new data and draws the logical conclusions from the finishing of any element by finishing also its ontological dependants. These time constraints are enforced by the meta-model, whereas conventional methods require one to supply them in additional documentation.

Behind every element, whether universal or particular, there stands a list of attributes, including the start and finish times and authorities. This information enables us, if necessary, to trace the responsible agents who can give operational definitions of the meanings of the elements and account for the data we hold. Unlike the Core model, which depends on some taken-for-granted, shared understanding about the conceptions in the inaccessible minds of unspecified persons, this schema pins responsibility for meanings on identifiable individuals or groups who can be required to account for their observations, choices and actions. It also imposes the exacting semantic discipline of knowing the temporal limits on the existence of each element. A maxim anyone can adopt is: If you cannot account for the start and finish of the existence of things you name, you cannot be sure about the meanings of you terminology.

The STDB, of course, can operate on these start and finish times and authorities. Powerful temporal operators such as *while*, *orwhile*, *whilenot*, *before*, *after* and so on can be invoked to retrieve and manipulate data and formulate norms. The authorities may be norms, which the computer can interpret as program modules (eg: rules to compute the stock parameters or to assign a part of a stock-holding to a particular contract). But they may be agents with whom the computer can communicate to obtain data or decisions (individuals may be appointed to consign goods or accept orders, for example). It would be far more difficult to drive a computer interpreter from the Core model in chapter 5.

Is the Social View better than the Core View?

I offered some criteria for judging any alternative view with its theory and modelling methods. Let us check that it measures up to them.

1) Does it deal adequately with the relevant IS concepts identified in the Broad View, including meaning, intention, responsibility and society?

Yes. It provides rigorous, operational definitions of all these concepts and methods of accounting for them in the models created using the new approach.

2) Does it handle the multiple meanings of key IS concepts such as information, communication, meaning and knowledge?

Yes, but there is not space to examine all these concepts in detail in this paper. On pp. 141-146, the Report notes the different meanings of “meaning”, “communication” and “information” that apply on each of the semiotic levels. These were demonstrated in ontological analyses that I provided for information systems concepts that belong on each level. A published version of these ontology charts you can find in Stamper 1996. All these definitions depend on each semiotic level confining our attention to a different repertoire of behaviour that signs make available to us. “Knowledge” is defined quite precisely in terms of norms, which are then used to give formal precision to what we understand by society.

3) Does its structure of definitions lead back to operationally sound primitive notions so that its propositions can be tested empirically?

Yes! The primitive notions can all be experienced directly and in the public domain where empirical tests can be applied. They include persons, Society, reproducible repertoires of behaviour called ‘affordances’, including the complex ones, norms, and signs. Norms are the regularities in behaviour that we can observe in the conduct of the norm subjects or elicit from them as the rules they endeavour to follow, often specified as laws. Signs are

affordances that can be demonstrated to stand for other affordances in relationship to which people can behave appropriately given only the sign-tokens representing them.

4) *Does it lend itself to a formal modelling with such advantage as ensuring systems are logically complete and internally consistent?*

Probably yes! The new method employs the novel logical concept of 'ontological dependency' and a syntax based on a wff that forces us to recognise the agent responsible for any knowledge. The ontological structures enforce a high degree of consistency on any model. I conjecture that Norma, the new formalism has the potential to become a logic of social action, better suited than the predicate logic used in the Core model for representing social systems. Whereas predicate logic has had the benefit of the research work of a vast academic industry over almost a century, Norma is embryonic. I suspect that the mathematical methods needed to develop it fully will be found in category theory. In a practical sense, Norma is formal enough for generating directly the computer system needed to support the system being modelled.

5) *Does it lead us to draw distinctions and investigate problems central to IS that would otherwise be overlooked by other methods and theories?*

Yes! Overwhelmingly, this is the special advantage of the new approach. Whereas the Core method does not distinguish computer games from systems for running social security or controlling nuclear power stations, the new approach does so by making the locus of responsibility explicit and by differentiating fiction from fact. Dungeons and dragons have totally different semantic standing from contributions and benefits or from reactor activity and control parameters. The Core View does not distinguish clearly beliefs, obligations and values nor draw attention, as the Social View does, to the problems of justifying the information we handle in our systems. Time has no special significance in the Core View whereas the Social View forces us to look at how we create it along with past and future worlds through the use of information. The Core View confines itself to the syntactic level of analysis whereas the Social View provides the means of exploring all semiotic levels, including those most important for users: the semantic, pragmatic and social levels.

6) *Does it support the continual accretion of testable knowledge of the information systems domain, as we would expect of any other modern science?*

Yes, at least in some areas. It treats the human systems of organisations, institutions and communities as the prime object of study. These have evolved for as long mankind; indeed human nature would not be possible without them. The formal models and theories of behaviour (norms and affordances) that we can create can be subjected to rigorous testing, not only of their internal consistency, but of their empirical validity. This takes us out of the non-scientific mode of methodological practice, in which information systems has been confined so far, into a scientific age where engineering is more than a kind of best practice, just as physics transformed the accumulating methodological knowledge of the blacksmith into mechanical engineering.

7) *Does it enable us to do the routine work on information systems conveniently and effectively?*

Yes, as the reworking of the Japan Wines Inc. case study indicates. Experience with real commercial applications confirms this. The new approach can drastically reduce the costs of system development, support and maintenance. Documentation shrinks by at least a factor of 10. The canonical substantive schema supports re-engineering processes. The apparent disadvantage is that the rigorous semantic analysis, imposed using the new approach is difficult compared with orthodox styles of modelling, but the resulting, rigorous solutions are re-usable. The new methods, by dealing rigorously and explicitly with the problems of meanings, intentions and social knowledge, greatly increase the likelihood that the users' actual needs will be met, and reduce the risk of the organisational failures that have plagued information systems work to date.

Industrial Testing

An essential test to apply to any framework of information systems concepts is to see how well it works in practice. In the case of the Core FRISCO framework, it has never been used to solve any industrial problem. Looking at the definitions and notations in Chapters 3 and 4, one suspects that it never will have an industrial use unless it is hidden in methods that are easier to read. But the Core View is not far removed from the many conventional systems analysis and specification methodologies in use today. In systems development work, these have not saved us from building a high proportion of systems that are organisational failures, however well they run as computer systems. Either way, FRISCO's Core View has not yet scored a high mark in industry.

The Social View has not yet been promulgated to a wide audience of systems developers. We have held back from doing so while the ideas were being tested. Those tests, of course, have included numerous industrial

applications. It is very difficult, unless you are already a leading consulting firm, to introduce radically new ideas to systems development from initial study to final implementation. Hence most of our tests have been in the analysis stage and, so far, only two have gone right through to the working system. Analysis based on the Social View forces attention on the business problems and almost invariably unarticulated problems were brought to the surface, clearly articulated and usually solved.

The two systems fully implemented on the basis of the Social View were successful. One of these was a software package being implemented in the conventional way. The documentation was vast and impenetrable to the professional experts whose ideas were being embodied in the package. When they adopted the concepts of the Social View explained in this paper, the user documentation shrank by about a factor of twenty and the professionals regained control of the work, extending the functionality of the package and advising on bad features. The structures produced were easily implemented using a multi-media, object-oriented database. The other system is of medium size and administers a university with rather complex, course-unit degrees. Ades 1999, describes some interesting aspects of this system. It has about 200 affordances in its ontological structures, that means, counting a minimum of four attributes per affordance, that it is equivalent to a system with at least 800 conventional elements. This system was created quickly and has been running for nearly 10 years with hardly any need for revision despite a rapidly evolving set of business requirements. One person deals with all the modifications, support and training. The rigour of the ontology structures explains the economy. They give us a semantic normal form (SNF) if the constraints are applied strictly. Interestingly, these constraints were relaxed at the outset, in the spirit of conventional systems development. The fact that standard relational normal forms were applied did not prevent the system being infected by a few “malignant features” – “malignant” in the sense that they could only be eliminated with the rewriting of application programs. The rest of the system conformed to the SNF and never required any program to be rewritten as requirements changed. Over the 10 year life of the system, the maintenance costs not arising from straightforward augmentation, all resulted from the struggle to eliminate these malignant features. We do hope that others will be interested in adopting the Social View that is set out in this paper.

If information systems is to become a science

I have deliberately courted controversy. That is because controversy, conducted in the right spirit, is essential for science to exist. I have also tried to be precise and clear enough to invite effective criticism, for without

criticism applied to theories that cannot escape it through slippery vagueness, there can be no science.

The views expressed here are at variance with nearly all the orthodox approaches to information systems practised in business at the turn of the century, though not entirely lacking academic support (Hirschheim *et al* 1995, for example). I should like these views to be given your time and open-minded attention and then subjected to your detailed critical scrutiny. Belief in a subjectivist position is only a belief, after all. Time and open-minded attention are essential because, all too often, I have had these ideas misinterpreted as some minor variation on an orthodox methodology. Perhaps the critical test that one has grasped the unusual point I am making is that one can appreciate the shift of ontology. One has to relinquish a belief in an objective reality (even one seen through the veil of an unspecified person's perceptions and conceptions) in favour of the radical belief, found in in FRISCO's Broad View, of a socially constructed reality, a world where there is no time but *now*, and no place but *here*. Only having made this leap of imagination can one see that a science of information systems has the exciting task of explaining how we construct this greater reality. Instead of taking information for granted as the transparent lens through which we see the 'real' reality, we begin to examine the lens and its rather magical 'optical' properties that depend on people expressing their intentions to create past, future and distant worlds. We do not take these processes for granted in an unrestrained relativism but we ask how our information systems can make sure our social constructions are justified, through the demand for, and exercise of responsibility by individuals and groups of people. This prospect of a true science, which FRISCO's Broad View opens for us, is a far cry from the old methodological information systems.

Even modern science is the product of information technology. It depends on open, critical discussion of clearly stated and testable facts and theories. Science could not come into being until the technical infrastructure was capable of reproducing factual and theoretical reports with sufficient accuracy and circulating them openly among a sufficiently large community of critical scholars. Criticism and openness are essential to the growth of scientific knowledge. Politics can easily intervene to prevent ideas circulating. In the past, churches and secular authorities often intervened and they continue to do so today wherever it is to their advantage. Institutional science and technology can also be politicised where it suits large corporations, governments and professional groupings. In a spirit that challenges any politics of information systems, I offer this extension of my dissenting reflections in the FRISCO Report.

References

- Ades, Yasser, 1999, "Semantic Normal Form: Compliance" in *Proc. Workshop on Organisational Semiotics*, Almelo October 12-14, University of Twente.
- Ernest, Paul, 1998, *Social Constructivism as a Philosophy of Mathematics*, New York, State University of NY.
- Falkenberg, Eckhard, et al, 1998, *A Framework of Information Systems Concepts*, International Federation for Information Processing, ISBN 3-901882-01-4 .
- Feyerabend Paul, 1975, *Against Method*, London, Verso.
- Gibson, James, J., 1979, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston
- Hersh, Reuben, 1997, *What is Mathematics Really?*, London, Vintage.
- Hirshheim, Rudy, Heinz Klein and Kalle Lyytinen, 1995, *Information Systems Development and Data Modelling: conceptual and philosophical foundations*; Cambridge University Press, Cambridge.
- Kitcher, Philip (1984), *The Nature of Mathematical Knowledge*, Oxford University Press.
- Lakatos, Imre and Alan Musgrave, 1970, *Criticism and the Growth of Knowledge*, Cambridge, CUP.
- Lakatos, Imré (1976), *Proofs and Refutations*, Cambridge, Cambridge University Press.
- Michaels, Claire F. and Claudia Carello, 1981, *Direct Perception*, Prentice-Hall, Englewood Cliffs, NJ.
- Popper, Sir Karl, 1959 (original 1935), *The Logic of Scientific Discovery*, Basic Books, New York.
- Popper, Sir Karl, 1963, *Conjectures and Refutations*, Routledge and Kegan Paul, London.
- Sokal, Alan and Jean Bricmont, 1998, *Intellectual Impostures*, London, Profile Books.
- Stamper, Ronald, 1973, *Information in Business and Administrative Systems*, Wiley, New York and Batsford, London.
- Stamper, Ronald, 1985, A Logic of Social Norms and Individual Affordances, in Gilbert, N. and G. Heath (eds), *Social Action and Artificial Intelligence*, Gower Press, Aldershot.
- Stamper, Ronald, 1989, *Information Management*, Enschede, University of Twente.
- Stamper, Ronald, 1996, Signs, Norms and Information Systems, in Holmqvist B. *et al* (eds), *Signs at Work*, De Gruyter, Berlin.