

Exploring the Semantics of Communication Acts

(Revised version of the invited paper delivered orally to the Achen LAP Workshop in 2000 but not in written form)

Ronald Stamper

29 Dartmouth Park Road, London, NW5 1SU.

(Staffordshire University, UK-EPSRC SEDITA Project with Reading University)

rstamper@blueyonder.co.uk

Abstract

The meanings of communication acts depend on what they are about. From a new ontology that recognises only realities constructed by responsible agents, the paper illustrates how this idea gives rise to a concept: that of *ontological dependency*, which leads to a canonical schema (SNF) for a given action- or problem-domain. Applied to communication acts, this augments the well-developed Language-Action Paradigm. The SNF approach is contrasted with the State Transition Diagram (STD) protocol approach. It highlights the distinction between the momentary speech act (such as a request) and the social construct (also a request) that persists to govern the conduct of participants in the conversation. The SNF compels one to examine how conversations connect parts of reality that exist at different times, thus exposing the machinery for the construction and use of our social reality. Some aspects of conversations are examined in a legal context to suggest further research.

Keywords: ontology, actualism, semantics, speech acts, canonical schema, time, responsibility.

1 Introduction

We don't just talk, we use our speech act to talk *about* things that do not necessarily appear on the surface of our conversations that, nonetheless, make use of the participants' assumed knowledge of those things. "Have you stopped beating your wife?" makes lots of assumptions but succeeds in suggesting what they are. "I'm inviting you to a party on Saturday!" keeps most of its assumptions hidden. It would help to have a method of analysis that brings the assumptions to the surface thereby clarifying the meanings of those speech acts.

2 Meaning and Reality

For a word, a number, an expression, formula, sentence or other sign to have a meaning in the world of practical affairs, it must stand for something that does or could exist. In the terminology of Charles Morris, it must either *denote* something that does exist, or *signify* something that could exist. Signs function in both these ways by directing the interpreter's attention to whatever they stand for, even indicating how they might be reached. If that goal is attainable because the thing exists then the sign denotes it; otherwise, the sign can only signify by providing the person able to interpret it with criteria that would successfully terminate the search should it come into existence. "Venice" and "The LAP Community" both denote, but not "unicorn" and "King Lear" that do succeed in signifying because we have enough tests to check that we have actually encountered either of them. "Vamadadoodle" neither denotes

nor signifies, having been conjured up this moment without my having the least intention of using the word again.

When we converse, we frequently do so in order to create understanding about things that no longer exist (history) or that could be brought into existence (proposals, plans and designs) or that are too distant for the interpreter to reach. All this construction of conceptual worlds depends upon those engaged in the conversation having a large enough vocabulary whose meanings have already been established by their successfully denoting particular things. Helping us to span the gulf of time between the actual experience of particular things and the potential experience of things we wish to signify are the universal terms we employ. We have met John and Mary and Bill; knowing them as instances of the universal 'person' enables us to talk meaningfully of Napoleon as another person; if we have also visited France and number a few emperors or other autocrats among our acquaintance, we shall be even better able to know what 'Napoleon' signifies when history has supplied us with a more detailed description.

Clearly we cannot say much about meaning without first making rather basic assumptions about the nature of reality. Most people working on information systems embrace the objectivist assumption that reality consists of a vast population of individuals capable of being identified by names or codes, if we choose to assign them. For anyone working on speech acts and conversations, this ontological assumption has the unfortunate effect of deflecting their attention from one of the most interesting aspects of human use of signs: their role in constructing what we understand to be the real world.

3 Another Ontology: Actualism

The expression 'an ontology' is now widely used to mean roughly the same as 'a data model', which has the unfortunate effect of deflecting attention from the pre-IT-era meaning of a metaphysical assumption about the nature of reality. The danger of this semantic slippage becomes evident when one moves from investigating how to manipulate data by computer and begins to investigate the connections between those data and the things they stand for in the real world. The research reported here made this shift of focus in an attempt to model organised human behaviour in terms of the social norms that govern the people involved.

Beginning to search for a formalism to represent social norms, using legal norms as our experimental material, it soon became evident that the objectivist ontology does not work because we actually construct our social world – everything from tax regimes to football games – and, moreover, we use speech acts to do that. That began the loosening of our adherence to the objectivism that we had learned at school and that seemed so compellingly natural. However, trying to build a theory (embedded in a language for representing norms) that regarded the world as objectively given except for socially constructed additions (as Searle accepts in his book *The Construction of Social Reality*) proved difficult and gave one the feeling of offending against the principle of Occam's Razor. How to achieve unification and avoid having to accept simultaneously two different ontological principles only became clear on reading James Gibson's Theory of Affordances.

Gibson, in studying the psychology of perception, found that the classical theory, based on objectivism, which assumed an organism viewed a given reality through the 'window' of its senses to register perceptions of given individuals, failed to explain perception. He proposed, instead, that an organism, bombarded with external and internal signals from the events it encountered and its own actions, had to discover within that flux useful, stable patterns. To survive and avoid discomfort, the organism needs to know at any time what patterns of *behaviour* the outer-and inner-world **afford** it. Perception in these terms is the recognition of *invariant repertoires of behaviour* that Gibson appropriately called **affordances**.

Thus, the only world that any organism can know is one that it constructs by discovering the useful affordances and perceiving them when they become available. Thus a rabbit, to survive, must recognise when it can use the environment for hiding or for running fast; similarly the air force pilots Gibson studied had to recognise the various patterns of *towards* as they approached an operational landing strip. We also perceive objects as invariant repertoires of behaviour: what do a cup and a hammer allow one to do? Correct perception of the world around depends on having ample experience of these affordances – watch a small child with a hammer or with a cup testing out all that they can do!

By way of illustration, consider the perception of *towards* by taking the simple case of walking towards a vertical wall. You will experience the kinaesthetic signals from your legs and the rest of your body's motion in conjunction with many other visual and auditory signals; among the invariant features, you will notice a stationary point on the wall around which the visual texture of the wall expands at speeds also invariant over the circles centred on the stationary point; moreover, those speeds increase as an invariant mathematical function of the radius and the kinaesthetically registered speed of motion; if you are not sighted, you would use the invariants in the changing echoes of gravel crunching beneath your feet or you tapping stick. Gibson's landing strips make the notion of *towards* more complex to analyse!

A few points to note: Gibson's Theory of Affordances applies to the material world. Everything falls within the category of affordances. Our temptation to draw upon syntactic categories for analysis purposes is probably misleading, as we have seen affordances labelled with verbs (to hide), nouns (hammer) and prepositions (towards). Similarly, the favourite categories for conceptual modelling – entities, attributes and relations – become redundant. Many organisms will have their abilities to perceive important affordances hardwired as a result of evolution. Nevertheless, the organism – either as individual or genotype – must in a biological sense bear responsibility for the world they construct by selecting these invariant affordances. Also note that the agent *realises* an affordance – that is has available that repertoire of behaviour, whether hiding, hammering or going towards – for a finite period of time.

The obvious step towards our goal of formalising social norms was to extend Gibson's idea by adding to the physical affordances the social affordances that we define using norms. These invariant patterns of behaviour that we permit through their cultural evolution or by means of law- or rule-making function socially in much the same way as Gibson's affordances in creating the known physical reality. Just consider the examples of the person pointing to the empty chair in a restaurant asking,

“Is someone sitting there?” or the law of copyright that creates something that one can buy or sell just like a table: in both cases we are confronting socially defined, acceptable repertoires of behaviour. Certainly the idea of responsibility applies in the case of legal norms though less obviously in the case of cultural norms.

A species may acquire through the normal processes of evolution a hardwired ability to perceive many important affordances; similarly, humans (and to some extent a few other species) establish most of our affordances and hand them over to succeeding generations by cultural mechanisms. So the reality we know – though apparently solid and objective and not of our own making – has been created by generations of our ancestors. Hence Society must play a key role in any analysis based on this new ontology. A kind of biologically enforced responsibility for its constructed reality is imposed on every species when its members face the tests of evolution; similarly, it seems reasonable to charge Society with responsibility for its culturally evolved norms: a culture that does not work in a practical sense may deliver misery, even extinction to a species, including our own. (I have the impression that growth-orientated economics¹ on a finite Earth might be risking the ultimate punishment!) There is a further basic role for Society in our analysis: a simple organism lives only in the here-and-now with access to a past, a future or distant reality strictly limited by its capacity to remember or signify using signs. Society hugely extends our meagre individual capacities to create realities that are not here and now, by means of its languages and its acts of speech and conversation.

This ontology that accepts the common sense assumptions that

1. the only world we can know depends upon there being a knowing agent; and that
2. knowledge depends on the actions of the knowing agent.

It may be labelled appropriately ‘**actualism**’ to emphasise its defining characteristic of having to do things to create the reality we know. It’s no good hanging around and looking hoping that objectively it’s there!

4 Ontological Dependency Schemas

We now apply these philosophical notions to the construction of an instrument that we can use for the analysis of conversations. The instrument takes the form of a kind of logical language constructed so that not only can we say things about the world but also establish the meanings of such utterances. Classical logics provide hints of what we need. Predicate logic can be treated as a pure formalism but we disregard that approach simply because it evades the key issue of how to connect statements to the reality they stand for. The objectivist position that informs predicate logic assumes the existence of given individuals that have names (John, Mary, Fido) and an ability to assemble them into sets that can also be given the names of predicates (boy, girl, person) or arranged into ordered pairs that can then be assembled into sets (likes, owns). With these devices we can make composite statements about the membership of individuals in various sets such as:

boy (John) & girl (Mary) & likes (John, Mary) & ~ & owns (Mary, Fido)

¹ Driven by our pathological money systems – see the chapter on the semantics of money in Stamper (forthcoming).

Notice that we can deduce the truth of the whole expression from the truth-values of all its simple predicate expression. Moreover, time receives no mention but we can introduce it by assuming there exist whole populations of individual, named instants or intervals. Clearly we are all expected to be looking at the same objectively given reality composed of independent creatures and times for whose existence no one has any responsibility; the truth-values provide the connection between the logical expressions and this given reality, but we have no hint of how those values can be tested. Predicate logic conforms quite naturally with the underlying, tacit philosophical assumptions of objectivism.

Actualism starts from the position that, to say anything about any knowable reality, we must find the responsible knowing agent and his or her selected invariant repertoire of behaviour (affordance) that the agent selects from the enveloping flux of experience. So our sentences have the shape:

Agent affordance

These sentences don't assert any truth but say that a certain repertoire of behaviour exists in the world that the agent realises, or brings into operation.

That agent – perhaps holding a hammer – is a modified form of agent, now capable of realising repertoires of behaviour unavailable without the tool. Thus

(Agent affordance) affordance
(Stamper hammer) heft

Another elaboration arises when an agent realises simultaneously two affordances, which together provide further extensions to its repertoire of behaviours:

Agent (affordance **while** affordance) affordance
Stamper (hammer **while** nail) hit

Thus hitting requires two *antecedent* affordances, hammer and nail and hit is the **ontological dependent**. In certain cases this building of realisation upon realisation may continue recursively. Note that affordances have no existence without an agent to experience them; they have to be realised, that is made real by the agent: hence the term 'realisation'. Mention of an affordance in the context of a schema always implies a realised affordance.

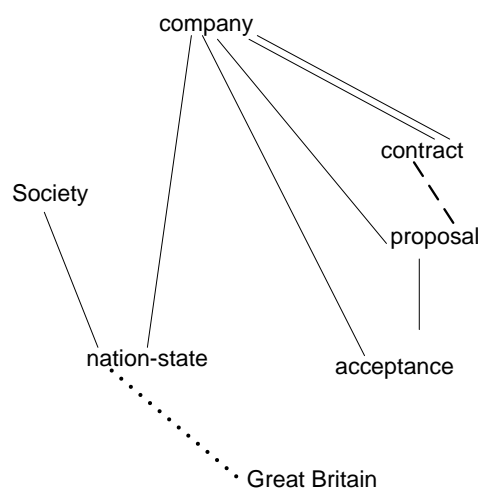


Figure 1: Ontological Dependency Schema for a Contract

With these concepts we have the elements of a new logical formalism². But before we begin to apply it to Language Action, we must take note of some important features. First, note the implied time constraints imposed by the ontological dependency relationships and certain logical consequences. If any one of the antecedents of an affordance goes out of existence, the agent cannot realise it. So, if someone takes away either the hammer or the nail, I cannot do any hitting. These powerful temporal constraints have some even stranger implications that we shall return to shortly.

Illustrating this with making a contract, an example that involves a speech act, will emphasise this point and draw attention to some others and, simultaneously introduce the graphical conventions. The schema in Figure 1 has nodes that label *universal* affordances such as nation-state (in lowercase) except for Great Britain (with uppercase capitals), which is a *particular* instance of a nation with a jurisdiction that determines the existence of certain companies. Note that universals are not sets but are also affordances that capture the ability to recognise any particular providing certain repertoires of behaviour. A *dotted line* links the particular to its universal. The solid lines, with a left-right progression represent the ontological dependencies. The only other exceptional arc, a *line of dashes* connects proposal to one of its antecedents, contract.

Obviously the contract cannot exist when someone is just proposing it, so the broken line indicates that the antecedent is not the contract itself but a *sign* that stands for the particular contract. It may look as though this proposal is a speech act but that could not be the case as its existence would be momentary which would make the acceptance impossible because its existence depends upon that of the proposal. Clearly we have encountered a lurking ambiguity: a company certainly does perform the appropriate, *momentary speech act* which then brings into existence a *persistent attitude* that may well go by the same name. We could add the speech act to the schema, indicating the difference – I suggest by adding an exclamation mark – thus: proposal! Notice that the interpretation of the illocution by an appropriate perlocutionary act will have brought into existence an item of ‘furniture’ in the social world: a complex cluster of attitudes implying an obligation on the speaker, a permission on the addressee and various expectations about further communications. The proposal will exist until it finishes with a withdrawal or as a result of other actions by the parties involved.

² It would be unwise to call it a logic because its properties are different from classical logics in a number of fundamental respects.

5 Semantic Normal Form

The schema illustrated here has some hidden features that arise from the philosophical position that underlies the formalism. It might help to picture the graph as a schema for a database. Indeed, we do use it as a schema for a Semantic Temporal DataBase (STDB) which can be imagined as assigning to every node a relation with exactly the same attribute structure. Each of these tuples functions as a surrogate for an affordance realised between its specified start and finish. Thus:

Table 1: Attributes of the surrogate for a realised affordance. One may picture this as a relational table, one for every universal and particular, such as this proposal by a certain company:

Attribute	Value	Meaning
Surrogate number	12345	proposal
Modifier (e.g “-“ for a sign)		
Its universal	34567	attitude
Antecedent-1	56789	company [buyer]
[Antecedent-2	76543	“contract”
Start	20050520	“chronometer”
Finish		(not yet)
Authority for start	65432	proposal!
Authority for finish		(not yet)

The second column contains the recorded surrogate numbers of the items labelled in the third column. The labels in English do not appear in the surrogates, which stand for realised affordances; labels are attached in any desired language or jargon as other surrogates. Note that times never exist, as in predicate logic, as independent individuals; they only exist as signs for the something or other that coexisted with a start or a finish. In this illustration we have provided a date which names the state of some indicated chronometer that would have to be provided as the default source of times for the application.

From this schema structure, we can see that it imposes a very strict constraint: every realised affordance can have only one or a maximum of two antecedents. This imposes a canonical structure and ensures that different investigators will produce the same schema if they are modelling the same situation. Differences, will reveal a) differences of scope or b) error on the part of one, other or both investigators or c) they will model different real situations that are described in the same, ambiguous terminology. We can provide no formal proof of this canonical property: the justification for the claim is empirical.

The importance of this canonical or Semantic Normal Form lies in its virtual elimination of arbitrariness in the analysis. A schema does not represent one analyst’s opinion, which may be set against alternative proposed by other analysts, as is the case with almost all other schema modelling methods. (Any one interested in the practical consequences of the SNF should see Stamper and Ades, 2004, which shows how it leads to massive cost reductions in systems development, support and maintenance.)

For the creation of a branch of science that can generate a cumulative body of, a canonical form is essential. If two physicists represent the same situation in apparently mathematical terms, then each model should be capable of being transformed by formal operations into the other. They may differ if one provides more information than the other, if one or both is in error, or if they are actually describing different things; their can be no claiming that ‘my conceptualisation is just as valid as yours’ in the Post Modernist style of thinking!

6 The Instructive Nature of the SNF

Applying this method of analysis to a huge range of practical problems has demonstrated its power to uncover ambiguities, omissions, misconceptions and other modelling failures. We feel confident that it could throw light on at least some aspects of speech acts and conversations.

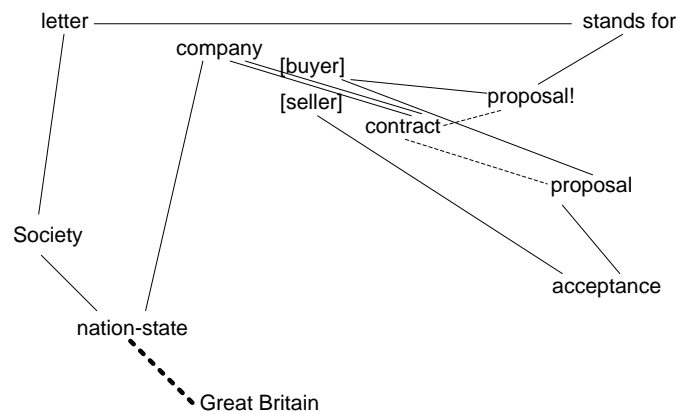


Figure 2: More detailed and revised Ontological Dependency Schema for a Contract

Let us first make a note of the few point that have already emerged from the introduction above and build further on the earlier analysis.

Figure 2 contains some more details. The schema now includes the role names of the participants to the contract; this draws attention to a defect in the original version, which showed the contract between two firms being proposed by some firm that need not necessarily be one of those in the contract – it may be the buyer’s agent – and the proposal being accepted by possibly a fourth firm – perhaps the seller’s agent. If we want those elaborations, we need only introduce the relationship shown in Figure 3 and use the agent to perform the speech act that creates the proposal.


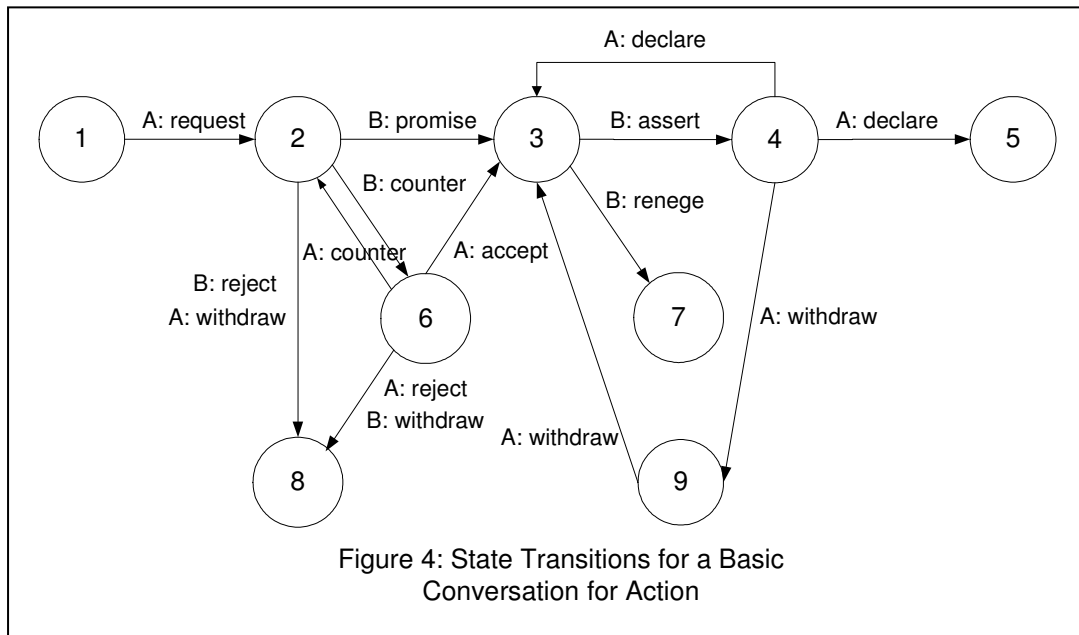


Figure 3: An additional relationship

The analysis should present all these details unambiguously. Of course the surrogates for the particular contracts, proposals and acceptances would contain the exact details of the companies involved. However, a generic model of the conversation to form a

contract benefits from introducing the roles of the participants. This is especially valuable when analysing the semantics of illocutions, as further examples will demonstrate.

Let us make some comparisons between the SNF method of representing conversations and the state transition diagram (STD) for which I quote the figure that has the honoured position as a principal inspiration for the Language-Action research programme (Winograd and Flores, 1986). It helps to have more than one way to present a given 'story' but we should check that any new notation must, at least, equal the expressive power of the STD.



When reading the SNF schema, one should remember that every node, or realised affordance, has its own surrogate consisting of the attributes listed in Table 1, above; most nodes in a schema represent universals, each of them having a population of particulars when treated as database schema. The numbered nodes of the STD in Figure 4 implicitly represent the attitudes created by the interpretation of the illocutionary acts but the SNF requires them to have appropriate names instead of numbers. The SNF also requires the surrogates to include the relevant times, which we often need when specifying the norms governing the conversation. The STD labels the participants (A and B) and their permitted actions on the arcs; these are handled in more detail in the authority attributes in the surrogates.

Notice the uniform modular structure of the SNF: universals and particulars all have the same surrogate structure. Thus, the particular, Great Britain, had its start with the Act of Union between England and Scotland but the start of a universal is often difficult to pin down: contract for example began its existence as a common sense, cultural product in the dim past but for commercial purposes we may need to deal with the notion of a contract in a precise legal fashion, in which case the authority for the universal will specify the precise law of a definite jurisdiction that brought it into existence. The STD model would have difficulty handling this aspect of authority in conversations. What about the authority for starting the particular contract?

Subject to the relevant contract law, the contract will start with the acceptance of the proposal for the contract. A reference in the surrogate to the acceptance brings with it all the relevant information in a manner that deserves attention because it illustrates well some of the expressive power of the SNF. An important geometric feature associated with every node of the SNF is its *branch*, which consists of all the selected affordance's antecedent nodes leading all the way back to Society, the root. We have already noted that the existence of that node depends on the coexistence of its entire branch. Hence, when we provide a reference to a node, we automatically provide links to all the necessary contextual information embedded in the branch. In the case of the authority for the start of the contract, this includes the acceptance and all its times and authorities, the proposal and its details, including the details of the contract with its universal that supplies the reference to the actual law governing the formation of this contract. This power to assemble the relevant information to determine that the contract is correctly formed stems from the discipline imposed by the ontological dependencies; moreover, the SNF achieves this richness of expression in a quite simple, readable form of the branch of the triggering affordance.

In Figure 2 we have already introduced the exclamation mark to differentiate between the illocutionary act, *proposal!*, and the attitude, *proposal*, which the perlocutionary act generates. So, we may also note that the authority for the start of the contract will also lead us to the exact illocutionary acts required.

Incidentally, the use of *proposal!* and *proposal* without confusion draws attention to the point that the syntactic categories we choose for the labels of the nodes signify very little. We can use the same label for more than one node because the location of the node in the SNF structure will remove any ambiguity. The SNF does not represent a linguistic structure but the structure of repertoires of behaviour: the affordances. In addition, if we choose, either to facilitate translation or to reduce misunderstanding, we may label each node in several different languages, registers or jargons or even with a description or video clip

The analysis in Figure 1 pays no attention to the physical sign-tokens employed to communicate or record the intentions involved; for many purposes, we are not concerned with the signalling aspects of the conversation. However, to show that the SNF can readily handle this feature, we have introduced it into Figure 2 where we have called the physical sign-token employed a *letter*. This document used to place an order on the supplier stands for the illocutionary act expressing that intention. When the *letter* physically coexists with the buyer's intention, the stands for relationship is one of denotation and at other times – when contemplating the order or recording it – the relationship is one of signification. Obviously, the intentions expressed by illocutionary acts and the attitudes that people adopt, exist in the minds of the members of the community involved, even without a persistent form of physical expression; some common forms of contract only involve transitory sign-tokens; a passenger normally make that kind of contracting conversation with a taxi driver.

A key point to notice in the above discussion is that the SNF analysis of a conversation to create a contract has many standard features that would be shared by numerous other types of conversation, such as the illocution leading to an attitude plus the physical token to stand for it plus the assignment of responsibilities etc. The

representation of these reusable standards hugely reduces the scale and complexity of the graphical model. Just as the labelled node implies all the information in its surrogate, we could also, for instance, dispense with SNF references to the speech act and other standard features implied by that attitude; all these affordances and their ontological dependencies follow as logical implications. The succinct graphical record would need some supplementary norms forming part of the standard. Let us make use of this simplification while recasting the STD in the form of an SNF. We omit both the illocutionary acts and the their physical representations, assuming that the relevant attitudes entail their existence.

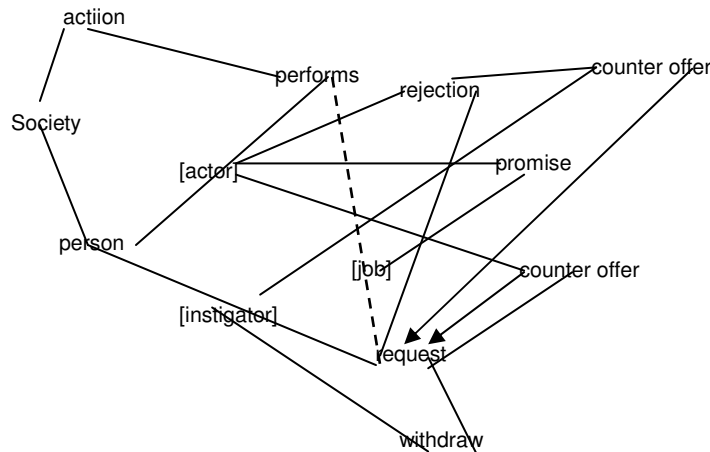


Figure 5: SNF for the ‘basic conversation for action’ – reaching agreement

Figure 5 contains the left-hand side of the STD in SNF form. Person A becomes the instigator and B the actor. The instigator takes the first step by making a request to the actor to perform a certain job but the request may be withdrawn. The actor can promise to do that, reject the request or make a counter offer; if the actor rejects the request, the instigator can also make a counter offer. Between nodes 2 and 6 in the STD counter offer may pile up upon counter offer until the last is withdrawn or rejected, or accepted and a promise made to act as proposed.

At this point an important distinction between the STD and the SNF model forces itself on our attention: the STD does not *necessitate* the recording of the history of the negotiations (of course, a suitable computer implementation can do that) but the SNF necessarily incorporates a complete historical record. In the SNF, we could also pile up counter offer upon counter offer in the graph but that would be both inelegant and unnecessary. The choice of terminology in the STD implies a fundamental distinction between the request and the counter offers while, in fact, all of them are proposals for the actor to do a job. The SNF shows that these are essentially the same by means of the graphic convention of the specific-to-generic arrows from each counter offer to the request, which tell us that the counter offer inherits all the affordances of the original request.

Clearly it would be more appropriate to call the request a proposal – a more neutral term – and the other a counter proposal – but because the SNF structure removes the potential

ambiguity, this is not important. Wherever a proposal stands – however it is labelled – it allows any of the three courses of action: withdraw, reject or accept. Note, however that the schema by itself does not incorporate the constraints governing who may take these actions, so these would have to be included in the semantic norms associated with the illocutions.

Finally, notice the small anomaly that the STD takes us to node 3 either directly from 2 with a promise or via 6 with an acceptance; perhaps more appropriately we could treat 3 as the state where the promise exists reached in both cases by accepting the standing proposal. This promise, unlike the contract in Figure 1, does not accommodate a compensating action by the instigator as the contract in the SNF can do so readily. Of course, the envisaged promise does not have the legal status associated with a contract; however they have an important feature in common in that the promise is linked to the request: it is not any old promise because the instigator is involved in ways – not accounted for here – that entail some responsibility in the subsequent performance and its effects.

At this point we should develop an SNF for the right hand part of the STD dealing with the performance of the job, establishing its completion and the instigator accepting or rejecting the result. We shall only do a small part of this analysis.

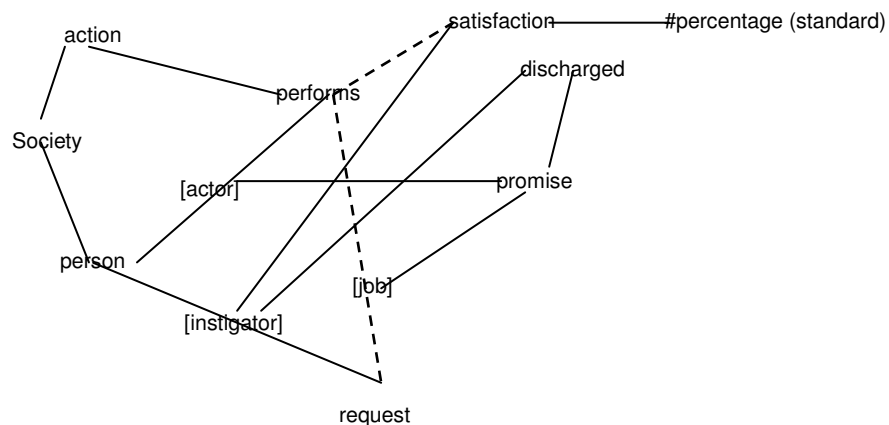


Figure 6: SNF for the 'basic conversation for action' – performing the action

This side of the analysis raises yet more interesting issues because it deals with knowing that the work has been done and making judgements about it, which are complex matters. Many of the complications may be put to one side if we remember that the SNF expresses only the aspects of the domain that we want to handle explicitly, leaving much to be known through the informal system. For example, the instigator can observe what the actor has done without any explicit reporting activity, leaving her to pass judgment on whether to discharge actor from the promise, which is all we have achieved by adding discharge to the analysis in Figure 6. In fact, it would be crazy to attempt to put into a schema all the possible patterns of discovery, communication, judgement, attitude formation and so forth that take place in the informal, tacit world of human thought and informal behaviour. So, when reading an SNF schema one must carry that real background in mind. Our analysis need only concern the parts we need to make explicit.

The action will have to be complete before the instigator can make a complete judgement, so that will have to be based on a sign standing for the performance. The

affordance of satisfaction in the Figure can apply to the performance at any time to represent the instigator's judgment; she will register some degree of satisfaction, which is also included in the Figure as a #percentage against some standard. (We shall only mention but not examine the kinds of affordances called 'determiners', indicated by the #-prefix, which are akin to measurements in this paper. For more see Stamper (forthcoming) or Stamper 1973 Chapter 9 on Measurement.) The instigator, on the basis of that judgement may at any time discharge the promise or report their judgement to the actor, and could also make partial payments for the work etc.

This small amount of tinkering will have to suffice to keep the paper within bounds. We could look into various possible jobs of making a thing or establishing a condition based around the activities of *beginning* or *ending* the existence of something, or ending its beginning (= *preventing* it) or ending its ending (= *sustaining* it). We could examine the use of reports on the work or its progress coming from divers sources. We could examine how different must be the notion of withdrawal before a promise is made and after the performance has started. We shall leave that to the interested reader.

Instead I want to suggest that these further explorations could be fruitful because the SNF model allows one to enter into the semantic domain that the STD model does not explore. Indeed the STD does not make such analysis easy whereas the SNF has the virtue (encumbrance, you may say) of compelling one to say things in some detail concerning *what the conversation is about*. When we do that, we can begin to express most of the semantic constraints on speech acts that Searle and Vanderveken discuss in their 1985 book, *Foundations of Illocutionary Logic*.

7 One further instructive case

So much of the semantics of illocutions is bound up with what the conversation is about, rather than the protocol of the permitted turns that participants may take in their discourse, that a rigorous and succinct method of handling this detail could open up a fruitful field of research. The results of such endeavour would be of immense practical value as reusable chunks of schema that recur in the analysis and design of e-government, e-business computer-based systems, not to mention the general theory of speech acts and conversations.

The example comes from the legal domain, which I have chosen to indicate both the research value of legal norms for the study of language for action, and the potential for adding practical value through the design of software packaged solutions to common business problems. It concerns the making of a will.

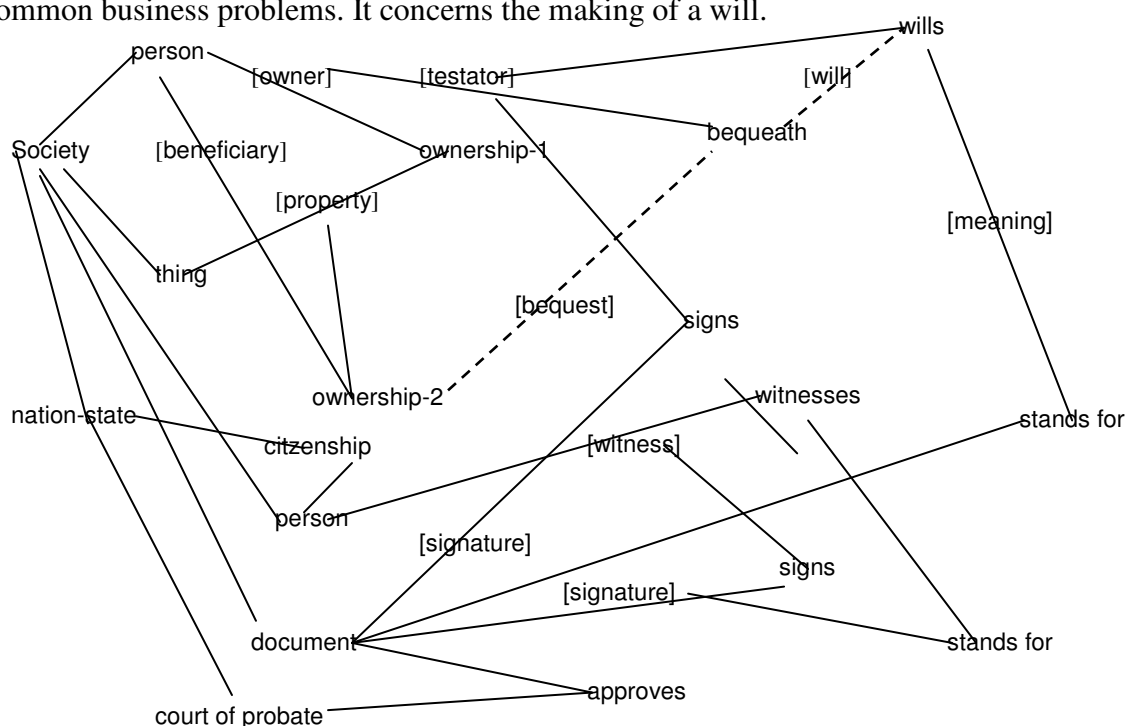


Figure 7: Speech acts for making a will and receiving probate onl

The example should serve as an exercise in reading an SNF, which skill underpins the much more difficult one of devising one. The reader should always regard the schema as an hypothesis about the social world capable of empirical testing. The strict constraints upon the ontological dependencies help one to make these tests but I am often surprised when people fail to spot violations of the simple rules – not least, I must say, when I notice my own failures. The value of the schema increases with testing.

The tests are simple: first, for each realised affordance, one must check coexistence of all antecedent realised affordances in the whole of its stem; secondly, one must be clear what authorities determine the realisation's start and finish. Remember that the *stem* is the lattice leading back from the affordance being examined to the root,

Society. Notice that the schema represents only what exists in the here-and-now; as one traces forward from the root, its scope narrows, raising the question “How do we link the present to anything that happened in the past or to what we might wish to do in future?” Anyone interested in the role of information in practical affairs, having to answer this fundamental question, should find it takes them to the heart of their problem domain where speech acts, conversations and talking do the essential work.

In the case of speech acts, they are always performed by an agent who obviously must exist at the time but using a sign (equivalent to the *propositional act*) that stands for whatever the speech is about. So one should always find a broken line from the illocution to the sign employed. In Figure 7 you will find no illocutionary acts, only the attitudes that their corresponding perlocutionary act generate. The nearest to the physical performance of an illocutionary act are the signings of the document by the testator and witnesses.

Taking the branch for wills we have the attitude established by the testator, which ceases to exist with his decease. What he want to happen is his will in the sense of the changes of ownership that should take place in future, so his ideas will perish but now shared with other people especially in the will, in the other sense as a document, persists.

The construction of the document, with due tests of the testator's soundness of mind, and the witnessing of his signature etc, is important for ensuring that the record of the illocution is well made. The illocution of willing something to happen does not meet its perlocutionary interpretation until the testator dies. At that point, Society or the relevant jurisdiction within it, as the agent to whom the illocution was addressed must interpret the expressed will as recorded. A court – in the UK the Family Division of the High Court – examines the will in order to check its validity as a document and the legality of the will as the intentions expressed therein. Only when proven, does the attitude labelled here bequeath start to exist. At that point the testator's roles are taken over by his executors who give effect to his wishes, but we do not analyse those stages.

Let us finish by noting the importance of avoiding the easy assumption that an illocutionary act will be followed promptly by the performance of the required perlocution. In this legal situation, Society, addressed by the testator, waits – perhaps decades – until the right moment and then checks many details of the making of the will, leaving the executors to check the factual details that obtain at his decease (did he still own the van Gogh bequeathed to his daughter when he died?). Our systems for automating speech act patterns have to take these considerations into account. The perlocutionary *activity* (‘act’ would suggest an oversimplification), as in this case, can be complex and extended, involving many agents and dependent on the subject matter of the speech acts and attitudes.

I hope that the methods introduced in this paper will interest some LAP colleagues in their application. I have had them in mind while writing *Analysis of Perceptions and Meanings for Information Systems Engineering*, a detailed account of the methods, which will be available to them in draft in 2005, in part to thank them for their stimulation.

References

- Searle, John R. and Daniel Vanderveken (1985), *Foundations of Illocutionary Logic*, Cambridge, Cambridge University Press.
- Stamper, Ronald, (forthcoming) *The Analysis of Perceptions and Meanings*
- Stamper, Ronald, 1973, Information, London, Batsford, New York, Wiley.
- Stamper, R and Y Ades, 2004, "Semantic Normal Form and System Quality, Proc. IEEE Conference on Requirements Engineering, Kyoto.
- Winograd, Terry. and Carlos F. Flores 1986, *Understanding Computers and Cognition*, Norwood, N.J., Ablex Publishing Co.