

Language Dynamics and Language Instantiation

Richmond H. Thomason
Philosophy Department
University of Michigan
rthomaso@umich.edu

Version of: September 14, 2020

0. Preface for September 15, 2020 presentation to the University of Michigan Philosophy Faculty Working Group

This paper is a working document that I expect will never appear in print, though I may recycle parts of it.

It's a prospectus for a project that I hope to be working on for an extended period of time. Although some of the ideas in it are over 35 years old, and some that are more recent have appeared in print in some form, there are things that I'm not satisfied with and things that I don't understand well enough. For rhetorical and stylistic reasons, I have pretended more confidence than I actually have for some parts of the project description.

1. Introduction

The project described in this paper has some things in common with the many applications of dynamic logic to semantics and pragmatics that have appeared over the last twenty-five years or more, and that continue to appear. But the aims and focus are different from most of that work. I'm interested in the philosophical and foundational opportunities offered by the dynamic vision. In applying to human languages an interpretive framework that was developed for computer programs, we make it possible to formulate an account of language and language use that is more inclusive, and that unifies human languages not only with programming languages but with other phenomena that in some sense are linguistic, such as musical scores and the genetic code.

2. Embodied languages contrasted with human languages

A computer program is a very explicit and often quite complex imperative, aimed at an agent which, if it is functioning correctly, will slavishly follow the instructions of the program. Its interpretation, then, involves a mathematical model of the interpreting agent (and perhaps, of parts of its environment), of the actions the agent can perform, and of the changes of state induced by these actions.

The architecture of a digital computer enables an "agent loop" consisting of repetitions

of the following action.

*execute current instruction;*¹

For a digital computer, instructions guiding the agent cycle are strings consisting of 0's and 1's of some fixed uniform length, this length being a power of 2. The device on which this paper was written is a 64-bit computer.

To qualify as a user of a programming language, a computer must be able to compile high-level programs into machine language. Both stages of the compilation—the high level programming language and machine language—count as languages. But machine language has the special property that it is incorporated into the device: it is at once a symbolic structure and part of the electronic state of a device. In its latter capacity, as a matter of engineering and electronic causality, it shapes the dynamics of device operation.

This sort of symbolic-mechanical duality predates the electronic age. The keys of a music box and the cards of a player piano are at once musical scores and parts of a device that performs the score. The gears of a cuckoo clock are at once instructions for the cuckoo's hourly behavior, and part of the device that executes the behavior. The moral is that sometimes, language can be embodied and causal.

But at first glance, engineered language usage is nothing like human language usage. A digital computer understands machine language because of the agent loop that is installed as part of its operating dynamics. It understands a programming language \mathcal{L} if it has a machine language compiler that renders programs in \mathcal{L} into machine language. So understanding \mathcal{L} is determined by the electronic properties of the composed system consisting of the high-level program, the compiler, and the device. And it uses \mathcal{L} by executing programs in \mathcal{L} . A human understands \mathcal{L} if she has learned and assimilated its syntax and semantics, and using \mathcal{L} is (mostly) a matter of deploying what she has learned in communication.

The project of this paper will have to account for such differences and to articulate similarities that are not apparent at first glance.

3. David Lewis on \mathcal{L} -speaking populations

David's Lewis' 1967 Ph.D. dissertation had to do with conventions in general, and linguistic conventions in particular. Immediately afterwards, Lewis interacted for a short time with Richard Montague and Barbara Partee at UCLA, and so was able to participate in the origins of formal semantics. [Lewis, 1975] relates these two themes, posing the question "What makes a language, conceived of as a syntax and an association between syntactic structures and model-theoretic meanings, the language spoken by a population?"

His answer (p. 5) is:

A language \mathcal{L} is used by a population P if and only if there prevails in P a convention of truthfulness and trust in \mathcal{L} , sustained by an interest in communication.

¹For the loop to function properly, agent states must include a pointer to the current instruction. This pointer must be updated with each instruction execution.

This brief statement is accompanied by pages of elaboration and systematic replies to various objections; the details aren't important for our purposes.

Lewis' account is an analysis with wriggle room—there are several places (e.g., the use of 'prevails') where things are intentionally left vague. As usual, Lewis presents a compelling vision of a problem and its solution. And the moving parts of the vision—most especially his theory of convention—illuminate important aspects of the social dimension of language.

3.1. Framing the issue

Nevertheless, I believe that Lewis's framing of the question is problematic, as well as the idea that it can be answered by an analysis. Linguists are scientists; among other things, they produce theories of syntax and semantics, and Lewis' \mathcal{L} s are linguistic theories. To ask what makes \mathcal{L} the language of a population is to ask what makes a linguist's syntactic and semantic grammar a correct theory of its subject matter. This is the sort of question that philosophers of science are equipped to answer. But Lewis' paper created a tradition of "metasemantics," pursued by metaphysicians and philosophers of language, not philosophers of science. It is part of an unhealthy philosophical habit of treating linguistics not as an independent science, but as a sort of adjunct of philosophy, and as such to be addressed in philosophical terms.

In general, you can learn a lot from case studies of what makes specific scientific hypotheses correct, but—and I stand ready to be corrected by the experts if I'm wrong—I suspect that most philosophers of science wouldn't expect a one-sentence analysis to capture the relation between a comprehensive scientific theory and its subject matter, even with pages of explanation. Like other scientific theories, linguistic theories are complex, and linguists bring many types of evidence and criteria to bear in supporting or attacking these theories.

The question "What makes a physical process the execution of a computation?" is a close analogy to Lewis' \mathcal{L} -population problem. Since the initial and not very helpful examination of this question in an appendix to [Putnam, 1988], this topic has been debated by a relatively small group of philosophers. What is the outcome? I'd put it like this: it is a mistake to expect a simple answer to this question, even though we have a relatively simple theory of computation and well understood examples of engineered computing mechanisms. See [Milkowski, 2013, Piccinini, 2015] for overviews.

None of this is a direct argument against Lewis' proposal. His answer may be among the best available, if you begin with his question and his expectations of what the answer should look like. And maybe you can be happy with this answer, if you think that the empirical practices of linguists are somehow irrelevant.

3.2. Simplicity versus plausibility

Lewis started with Montague's semantic framework. This suggested a strategy connecting \mathcal{L} to a population P by way of the truth value of a sentence in a possible world. Taking semantic dynamics seriously, as I do here, leads to an \mathcal{L} - P connection based on the change

of state induced by a sentence in a given state of a dynamic system.

The dynamic picture is causal, not static, and affords three possible ways in which sentence uses can be connected with their effects: (1) the natural causal order, (2) conventions, and (3) pragmatic inferences based on intention recognition. Combinations of these sources are not ruled out.

Lewis' approach has simplicity going for it. At the same time, it makes it difficult for him to do justice to some features of human languages. He is aware of this, and this accounts for much of what he has to say in the objections & replies portions of [Lewis, 1975]. Whether or not you accept his account or the dynamic one depends on how you weigh the tradeoff between simplicity and plausibility.

Nondeclaratives illustrate the point. Lewis takes imperatives to be true or false: "Mow the lawn!" is equivalent to "I am telling you to mow the lawn." This promotes simplicity, but looks implausible. And there are linguistic arguments against the equivalence. Of course, although here his view is on the defensive, Lewis has ways of explaining away the apparent implausibility and can sustain a defense.

From the dynamic standpoint, and limiting things to human languages only, Lewis' picture of language instantiation is uncomfortable in five main ways: (1) the exclusive concentration on declaratives and static semantics; (2) emphasis on truth—an evaluative feature of some illocutionary acts—rather than on the acts themselves; (3) emphasis on conformity to conventions at the expense of ingrained individual psychology, (4) failure to allow for natural, nonconventional components of language use and (5) neglect of the role of spontaneity in language use.

I'll elaborate briefly on these five.

(1) One of the earliest and most successful applications of dynamic semantics to natural language concerned a theory of indefinite noun phrases as discourse referent introducers. Standard examples of cruelty to animals like

- 1a. If a farmer owns a donkey he beats it.
- 1b. Pedro owns a donkey. He beats it.

draw attention to several things. (i) Utterances of declarative sentences often are associated with nondeclarative effects; in particular, they can introduce a discourse referent. Concentrating on truth and trust alone misses this, as well as failing to do justice to many discourse effects that Lewis himself lists in [Lewis, 1979]. (ii) Single sentences may not be the basic unit of semantic analysis; some semantic effects are multisentential.

(2) Lewis' approach contrasts with J.L. Austin's insight that utterances are associated with various illocutionary acts, that different acts are evaluated in different characteristic ways, and that truth is one among many ways of evaluating some illocutionary acts. Austin's picture does more justice to the diversity of language usage.

(3) A speaker of a language—especially a native speaker—has absorbed the language at a profound psychological level that is not properly explained merely in terms of assent to coordination strategies. A large part of associating an utterance with its conversational contribution is managed by automatic, ingrained mechanisms, such as parsing, lexical access,

and disambiguation. It stretches Lewis' account of convention to or even beyond the breaking point to count these mechanisms as conventional. Whether or not a parsing mechanism can be explained as one solution among others to a game-theoretic coordination problem, deploying the mechanism is not a matter of choosing the mechanism because of the expectation that others share it. It is not chosen at all—it is ingrained in the first years of life as part of acquiring a first language.

(4) Some interpretive effects may well be a matter of natural causation, rather than of convention. The best candidates include prosodic effects, such as intonation contours, emphasis, and contrastive stress. Suppose you and your friend are photographing a bison in Yellowstone Park. Your friend's attention is drawn to the side, and at that point the bison lowers its head and begins to move, faster and faster, in your direction. When you yell "It's charging!" much of the effect on your friend is due to the characteristic intonation. Prosodic effects like this are pervasive and seem to be more or less directly connected to the human emotional system. A conventional account of these effects isn't at all plausible.

(5) Consider the first conversational turn of the example from Shaw's *Candida* in Section 12. The stage directions have Morell opening a letter and uttering a "comic groan of despair." Proserpine then says "Another lecture?" and her meaning is glossed as "Are you annoyed about another lecture invitation?"

Proserpine is able to achieve this effect with just two words because she is familiar with how the morning's post is likely to evoke Morell's peculiar sense of humor—familiarity which, because the two have worked together for a considerable time, she can take to be part of the common ground. The reasoning that allows the two to converge on the full illocutionary act is circumstantial, dependent on intention recognition, and not at all conventional. If we take expert recognition and production of illocutionary acts to be the fundamental skill exhibited by speakers of a language, linguistic conventions are only part of the story, even if they are a fundamental part.

4. Language and convention

If we wish to take seriously the idea that programming languages are languages, we will have to rethink convention and its role in language instantiation. The account of convention in [Lewis, 1969][p. 78] is based, naturally enough, on patterns of behavior that solve a group problem, that are arbitrary in that there are other equally good solutions, and that are based on knowledge and expectations in a population. Populations adopt conventions as a matter of consent and can, and sometimes do, choose to change them. On the other hand, even though computer engineers and programming language designers make many arbitrary choices, the devices that use the language do not, and there is no straightforward way of adopting Lewis' account to this case.

The genetic code is even more problematic. Here there are no engineers or language designers, and the choices might not be arbitrary in any plausible sense. Apparently, only one genetic code has emerged on this planet, and—being part of our fundamental makeup—it is certainly not arbitrary for us or other earthly creatures. And whatever reasons we have

for or against deciding that the genetic code is a language have nothing to do with whether it is arbitrary or the only choice offered by nature.

It is tempting to contrast convention with nature. But if we want to situate convention in a general account of language and language use, we seem to have two choices. (1) We could allow some conventions to be a matter of chemistry or biology and to have been engineered or to have evolved, perhaps even without there being any genuine alternatives. Or (2) we could try to relate languages to populations of users using some more general notion, treating convention as one form of language fixation. Elements of Lewis' account can be generalized, and the idea of natural conventions isn't entirely incoherent. Although, for instance, choice and preferences are part of Lewis' game theoretical account, it isn't essential that these should be human choices and human preferences.

I won't pursue this matter further here.

5. Imperatives, declaratives, and illocutionary acts

The fundamental linguistic unit of dynamic logic is imperative, so the fundamental semantic unit is the repertoire of actions available to some agent. Lewis' fundamental linguistic unit is declarative, and the semantic primitive is truth in a possible world. But dynamic logic is a generalization of truth-conditional logic, not a radical departure. Truth enters into dynamics by way of conditional imperatives. The antecedents of conditional imperatives are declarative, and the dynamic logic replicates the familiar static logic in its interpretation of these antecedents.

So maybe something similar can be contrived in framing the philosophical issues. Perhaps we can make imperatives and actions fundamental for human languages, and capture much of what Lewis has to say—for instance, his ideas about conventions of truthfulness and trust—as insights about the declarative subsystem of a fundamentally dynamic human language.

But what are these actions? There is really only one choice: they must be illocutionary acts or speech acts.²

This raises an initial difficulty: not all illocutionary acts are imperative. But perhaps the greatest obstacle to unifying natural and engineered imperatives and the illocutionary acts of human languages is automaticity. A computer is designed to obey the imperatives of its machine language: unless it's broken; obedience is built into its causal structure. Obviously, natural language imperatives are not like this. Overcoming both difficulties depends on realizing that linguistic imperatives—speaker directives concerning hearer actions—are not the right analogy. They are not the fundamental constructs of the dynamics of human languages.

The relevant states for our purposes are *states of a conversation*. The relevant actions are changes in these states, and the relevant imperatives are directives about changes in

²I will use the former term, although it's the less common one, because I want to emphasize connections between things I will say and themes that I find in J.L. Austin. I myself see no point in distinguishing speech acts from illocutionary acts.

conversational state.

6. Presupposition, common ground, and the conversational record

The idea that conversations as such track certain features, and that linguistic acts serve to modify these features, goes back to Stalnaker’s work on pragmatic presupposition in [Stalnaker, 1973] and can also be found in [Fillmore, 1997]. It is elaborated in [Lewis, 1979]; [Thomason, 1990] seeks to turn it into a foundation for pragmatic theory.

Stalnaker limits his attention to assertion, which updates a data structure tracking what is understood to have been established at any point in a conversation. Fillmore is mainly concerned with reference points anchoring words like ‘come’ and ‘go’. Lewis is more general, discussing a long list of components of “conversational score.” Like Lewis, I want a multidimensional record, with linguistic tests for the features that it incorporates.

[Clark and Marshall, 1981] introduce the term “common ground.” [Thomason, 1990] proposes an alternative term, “the conversational record,” and, like Clark and Marshall, insists on its commonality, in a technical sense corresponding to the way that “common knowledge” is used in game theory, the theory of communications protocols, and Lewis’ account of linguistic conventions.

I’m happy with “conversational score” and “common ground,” but in this paper will mostly stick to my own terminology: “CR” or conversational record. Whatever you call it, the conversational record provides a way to articulate the commonalities between human languages and languages that secure their effects electronically, mechanically, or chemically.

7. Illocution and Perlocution

The distinction between illocutionary and perlocutionary acts is introduced in [Austin, 1962], and is meant to mark the difference between a linguistic act *per se* and its downstream causal effects. Austin recognizes that his readers may find the distinction problematic, and that often it is difficult to pry an action apart from its effects. But he thinks it marks a real and useful difference. And it’s important for Austin that illocutionary acts are constituted by conventions—including, of course, linguistic conventions.

We obtain something close to the spirit of Austin’s distinction if we say that illocutionary actions are effects on the CR. If I argue that capitalism is evil, I am adding to the CR an argument for a conclusion. If I convince you that capitalism is evil, I have achieved an effect on you that may be caused by what was added to the CR, but which is no part itself of the CR. (Unless, of course, you say “I’m convinced.”) Adding something to what is assumed for the sake of the conversation is not the same as believing it, or even pretending to believe it. And what is added to the CR gets there for lack of objection, not because it has carried conviction.

This account of illocutionary acts has substantive and methodological consequences—especially if, as I do, you combine it with the idea that the CR is confined to common or

mutual information.

8. Commonality

For propositional attitudes, commonality for a group \mathcal{G} of attitude-holders is a matter of arbitrary iterations of the attitude. Most of those who formalize commonality, like [Lewis, 1969], have *common knowledge* in mind. A proposition p is common knowledge for \mathcal{G} if for all members x_1, x_2, \dots, x_n of \mathcal{G} , x_1 knows that x_2 knows that $\dots x_n$ knows p .

I don't believe that knowledge is the right attitude for the declarative or assertional component of the CR.³ In general, standing attitudes are not appropriate for practical purposes—agents create *ad hoc* attitudes appropriate for these purposes.⁴ Among these attitudes is *C-supposition*, or supposition for the sake of the conversation. A specific attitude, tailored to the conversation, and subject to a supposed mutuality constraint, helps to keep private information out of participants' views of what can be considered common ground in a conversation. Anyone who has wondered what they can take for granted in framing a conversational contribution has felt this constraint at work.

Various authors have expressed skepticism about the possibility of obtaining common knowledge, as well, possibly, as common versions of other attitudes, such as common belief. [Thomason, 2000] describes a way in which commonality—or something very like it—can be achieved using defaults about attitude transfer. For example, when two neighbors in a small town meet, each will assume that whatever she C-supposes will be C-supposed by her conversational partner. And she will be careful only to C-suppose things that were made available in ways that would be available to a neighbor. These assumptions are defeasible, and of course will sometimes be wrong. (So this mechanism for achieving commonality is more appropriate for attitudes that are more like belief than knowledge.) We have all experienced conversations where a discrepancy in coordination was discovered, and repairs have had to be made.

Not all that is common in a conversation is propositional. Salient questions can be part of the CR, as well as salient individuals. As far as I know, no one has developed an account of commonality for either of these types, and this is something that needs to be done. In Section 12, below, I label common individuals with variables $D-x$, which are to be construed as discourse referents: common labels presumed to be anchored to individuals.

The idea of a discourse referent originates, as far as I can tell, in [Karttunen, 1967]. I use a notation for discourse referents based loosely on that of Discourse Representation Theory, [Kamp and Reyle, 1993a, Kamp and Reyle, 1993b]. Throughout this paper there is alignment with and indebtedness to Hans Kamp's approach to discourse.

³As a corollary, I reject the idea that successful assertions create conditions for knowledge, and that a participant in a conversation has grounds for believing whatever has been successfully asserted in a conversation. There are in fact many counterexamples; conversations are not always as serious as many philosophers pretend.

⁴For more about this, see [Thomason, 2013].

9. What must be included in the CR

David Lewis' list of components of conversational score is fairly large, but there is, I believe, an ulterior motive in [Lewis, 1979]. He wishes to use context-sensitivity to refute skepticism, and softens the reader up for the cases that serve this purpose by beginning with a number of less controversial examples. But if conversational score tracks illocutionary acts, we have linguistic tests for what should go on the scoreboard, and can keep from overpopulating it.⁵

Treating illocutionary acts as CR changes creates a methodology for structuring the CR. If stating is an illocutionary act, then *established for the sake of the conversation* must be a component. If promises are illocutionary acts, then public commitments need to be a component. If questions are illocutionary acts, then salient issues (or *questions under discussion*, QUDs), need to be a component. If imperatives are illocutionary acts, then the goals explicitly attached to an agent in the course of the conversation need to be a component.⁶

This leads to a large and open-ended picture of the CR. But conversations themselves serve many purposes and must keep track of many things, so in this respect the picture matches its target.

10. Microinstitutions

J.L. Austin drew impressive insights from the idea that illocutionary acts are actions, and that considerations that apply generally to the successful performance and the less-than-successful misperformance of actions apply to them. Likewise, Paul Grice reaped important conclusions from the idea that communicative transactions are interactions between agents with intentions and beliefs.

In a similar vein, it can be helpful to situate conversations in a more general class of social phenomena and, in particular, to think of them as *microinstitutions*. Microinstitutions are social phenomena in which the participants have role-based expectations and responsibilities as well as shared attitudes, and which create and maintain shared information.

I've introduced awkward terminology for what I have in mind for two reasons: to point out the relationship to the larger social structures we usually call institutions and to downplay the well-worn analogy between conversations and games,⁷ and for lack of a better term. Games also are microinstitutions, but so are waltzes danced by a couple, musical performances, say by a string quartet, cello master classes, committee meetings, family dinners, business lunches, courtroom trials, religious rituals, and—two minimal cases—Hume's example of two

⁵If the illocutionary criterion is necessary as well as sufficient, Lewis' crucial example of vagueness is more problematic than the others,

⁶For the QUD, see [Roberts, 1996, Ginzburg, 2012]. For imperatives, see [Portner, 2009]. Portner talks about "to-do lists" and I talk about goals. These are not the same as actions, but in simple cases actions can be derived from goals by a trivial exercise in planning.

⁷With just this analogy to go on, rather than a wider range of examples, we are likely to be misled about conversations, which (unlike many games) are not zero-sum, are often improvised rather than rule-governed, and which rely heavily on accommodation.

men in a rowboat and David Lewis' example of a truck driver backing into a parking space with the help of a signaller.

All of these involve coordination in certain shared activities and mutual attunement to the intentions and expectations of others. They differ in regimentation and spontaneity, in the need for conventionalized means of interaction, and in the suite of attitudes that the participants need to track in order to keep up to date. The rowing example is merely a matter of maintaining rhythm and equilibrium—if we imagine the rowers are side by side with one oar each, the goal is to manage the oars as if just one person were rowing with an oar in each hand. If there are many rows of oarsmen and the rowers can't sense each other's motions, a coxwain and signalling conventions will be needed, and the crew will need a certain amount of training. Courtroom trials occupy another end of the spectrum. There are specific roles, with explicit rights and responsibilities, as well as elaborate rules about the agenda. In many cases there is even an official record, with rules about what must and cannot go into it.

Conversations of all kinds clearly fit into this category. It is hard to imagine a genuine conversation taking place without exploiting some conventions for communication, although some simple communicative interactions can do this. Although conversations differ greatly in their purposes, formality, and complexity they all involve mutual attunement to the intentions and expectations of others, and the need to track at least declarative, as well most likely as imperative and interrogative information.

Each of these dimensions of the CR can be modeled by invoking a data structure that is presumed to be common, that is changed by the actions of the participants and perhaps also by its own dynamics, and that is tracked by the participants as it changes.

Microinstitutions are a special case of what Allen Newell called *problem spaces* [Newell, 1992]; they provide a framework for reasoning about action, by making goals, action types, and expectations available. What differentiates them from single-agent problem spaces is a division between public and private information, and attention to interpersonal activities. Like any problem space, conversational microinstitutions provide the structure that shapes and informs the associated activity.

At this point we can (finally) return to the automaticity problem, raised in Section 5.

11. Linguistic acts and automaticity

Although Austin doesn't put it this way, he believes, in part because of his assimilation of illocutionary acts to actions in general, that there is a (qualified) automaticity to illocutionary acts. He says (paraphrasing, [Austin, 1962][pp. 5–6] that to utter 'I bet you sixpence that it will rain' in the appropriate circumstances is to make a bet. I myself would want to treat this as a case of contingent identity and say that (again, in the appropriate circumstances) the locutionary act is the same as the illocutionary act. But Austin is not explicit about this.

In [Austin, 1962][Lecture 2] he goes into more detail about appropriate circumstances,

classifying inappropriate circumstances into *misfires*—things upon which the performance of the act depends—and *abuses*—cases where it is performed, but performed poorly in some way. And he goes on to subclassify these two categories.

It isn't safe or even quite fair to attribute views to a philosopher based on a work that was never thought through entirely or prepared by the author for publication. So I will merely say that Austin might have wished to conclude that the performance of a locutionary act in appropriate circumstances is inseparable from the performance of the illocutionary act, and that the connection between the two (whether or not, as I want to say, it's identity) is a matter of convention.⁸

But in that case, the connection between saying (for instance) 'I bet you sixpence that it will rain' in appropriate circumstances and offering to make a bet is just as immediate as the connection between an instruction in machine language (in the appropriate circumstances) and its performance. Even though in the one case the connection is social and conventional and the other is engineered and electronic.

Unfortunately, things are more complicated than this, for several reasons.

I identify illocutionary acts with changes in the CR. So if I say that these changes are automatic and conventional in appropriate circumstances, I have to say that the performance of an illocutionary act in appropriate circumstances automatically amounts to an update of the CR, and that this is a matter of convention. Neither of these is right.

Automaticity can fail because of the commonality of the CR, discussed in Section 8. A change can't be added to the CR unless the parties to the conversation can assume common recognition of the change. Whatever a speaker says, if the hearer says "Huh?" the conversational update is suspended.

In fact, at just about any point an interlocutor may interrupt with an objection. At that point, either commonality breaks down (as sometimes happens in hostile conversations) or update is suspended and what will go on the CR must be negotiated.⁹ This negotiation will be a subordinate conversation, about how to agree on the CR of the main conversation.

I can say, however, that a locutionary contribution to a conversation is a bid to make a change in the CR, and that by default this bid will be accepted if there is no reason to think the bid wasn't recognized and accepted by the parties to the conversation, and (of course) if the circumstances are appropriate. And I see no problem in identifying the locution, the bid, and the corresponding update of the CR when all goes well and the utterance succeeds. It is fair to call this process automatic.

Indirect speech acts introduce a more worrisome difficulty, since they appear to sever the conventional connection between the locutionary and illocutionary acts. Linguistic convention would make "Do I look like an idiot?" a bid to open a QUD, whereas it is intended, and normally would be taken to perform quite a different change.

Some of these cases, like "Can you pass the salt," can be handled by allowing that a

⁸A number of passages in [Austin, 1962] support this last point.

⁹The subdialogs that ensue after an objection or a call for clarification or some other recognized conversational mishap are sometimes called *repair subdialogs*.

locutionary act can perform multiple illocutionary acts, and saying that as well as compositional conventions associated with the syntax and semantic rules of a language, there may be idiomatic conventions. We could say, for instance, that ‘Do I look like an idiot?’ does open a QUD, but in the same breath it answers it negatively. Many, maybe most indirect speech acts are like this.

But surely there can be genuinely spontaneous and idiosyncratic indirect speech acts. Maybe “Please, don’t tell me that you think I’m an idiot,” used to say “I’m no idiot,” is an example. Such examples would show that the connection between a locutionary act and the corresponding illocutionary act is not always conventional.

Conversational implicatures raise the same problem. Sometimes they serve to update the CR, although some things that Grice counts as conversational implicatures—insinuations, for instance—do not. Because conversational implicatures are not conventional, we must admit that not all updates to the CR are conventional. Nor must they be automatic—occasionally, implicatures have to be worked out deliberately.¹⁰

So I have to say that some locutionary acts are directly and conventionally associated with a corresponding illocutionary act, and some are indirectly and pragmatically connected with the associated act.

Maybe I can also say that conventional associations between locutionary and illocutionary acts are in some sense not only normal but fundamental. It’s easy to imagine, and not much harder to find, examples of long discourses—literal, practically oriented conversations, for instance—where every utterance invokes a conventional update. Furthermore, if Grice is right in thinking that conversational implicatures must be grounded in terms of reconstructed reasoning that depends, among other things, on the literal or “favored sense” of the locution, we have an argument that conventional updates are fundamental.

There are genres—poetry is one—where the automaticity of interpretation fails. But in most cases we encounter in day-to-day life and even in literature, it is straightforward to reconstruct the CR, and automaticity prevails. I offer an example in Section 12 that I believe supports this point, although it raises many questions about pragmatic interpretation that would have to be addressed in a more complete account. So humans can and do use language in a way that calls for deliberate and reflective interpretation; but most cases we encounter are effortless and automatic, and in this respect human language use is quite like the engineered case of computer languages.

This is a fairly complicated picture, threatening to become still more complicated. And it’s incomplete. In particular, it lacks an appropriate account of locutionary acts, something Austin doesn’t say much about. I won’t attempt that here.

¹⁰Nevertheless, most conversational implicatures are inferred automatically, not deliberately. One of Grice’s tests for conversational implicatures is deniability; but an automatic inference can be deniable.

12. An example

I'll use the opening of Act I of George Bernard Shaw's play *Candida* to illustrate how ideas about the CR apply to a conversation. Shaw is notoriously generous with stage directions; this is helpful for our purposes, because it provides details about the props that an author thinks are required to make sense of a conversation. And this can provide clues to the evolving CR.

The characters in the play incorporate these environmental features into the CR and use them in their conversational reasoning. Indeed, in her first turn Proserpine treats Morell's groan as part of the common ground. In the following transcript, I've omitted some of the stage directions; even with lacunae they are quite long enough. As in all conversations, some information that is not explicit in the utterance itself is added to the CR by pragmatic reasoning. This is incorporated in the annotation.

The CR in this conversation tracks five sorts of information:

- (1) Declarative information (DCG for "declarative common ground"), the things that are commonly supposed for the sake of the conversation.
- (2) Common referents, e.g. 'D- x_2 '.
- (3) The question-under-discussion list, "QUD".
- (4) The salient goal (if one of the purposes of the conversation is to form a group plan, the salient goal is a component of the planning process).
- (5) Participant goals (these are commitments assigned to individual participants in the course of the conversation).

Stage directions:

A fine October morning in the north east suburbs of London . . . The sun is shining cheerfully; there is no fog; and though the smoke effectually prevents anything, whether faces and hands or bricks and mortar, from looking fresh and clean, it is not hanging heavily enough to trouble a Londoner.

. . . The best view of Victoria Park is from the front window of St. Dominic's Parsonage . . . In [the drawing room], the only sitting-room that can be spared from the children and the family meals, Reverend James Mavor Morell does his work. He is sitting in a strong round backed revolving chair at the right hand end of a long table, which stands across the window . . . At the opposite end of the table, adjoining it, is a little table; only half the width of the other, with a typewriter on it. His typist is sitting at this machine, with her back to the window. The large table is littered with pamphlets, journals, letters, nests of drawers, an office diary, postage scales and the like. A spare chair for visitors having business with the parson is in the middle, turned to his end. Within reach of his hand is a stationery case, and a cabinet photograph in a frame. Behind him the right hand wall, recessed above the fireplace, is fitted with bookshelves. . . . Opposite him on the left, near the typewriter, is the door. Further down the room, opposite the fireplace, a bookcase stands on a cellaret, with a sofa near it. There is a generous fire burning; and the hearth, with a comfortable armchair

and a japanned flower painted coal scuttle at one side, a miniature chair for a boy or girl on the other, a nicely varnished wooden mantelpiece ... Altogether the room is the room of a good housekeeper, vanquished, as far as the table is concerned, by an untidy man, but elsewhere mistress of the situation. ...

The Reverend James Mavor Morell is a Christian Socialist clergyman of the Church of England, and an active member of the Guild of St. Matthew and the Christian Social Union. ... He has a healthy complexion, a good forehead, with the brows somewhat blunt, and the eyes bright and eager, a mouth resolute, but not particularly well cut, and a substantial nose, with the mobile, spreading nostrils of the dramatic orator, but, like all his features, void of subtlety.

The typist, Miss Proserpine Garnett, is a brisk little woman of about 30, of the lower middle class, neatly but cheaply dressed in a black merino skirt and a blouse, rather pert and quick of speech, and not very civil in her manner, but sensitive and affectionate. She is clattering away busily at her machine whilst Morell opens the last of his morning's letters. He realizes its contents with a comic groan of despair.

Dialog, annotated for changes to the CR:

«Mutually recognized emotional states of participants may be part of the CR. So before anything is said, we have: **[[Morell finds something he has read annoying, and is amused that he is annoyed]]** is added to the DCG. »

PROSERPINE. Another lecture?

«**[[Are you annoyed about another lecture invitation?]]** becomes the QUD.

[[What is the nature of the invitation]] is added as a secondary QUD. »

MORELL. Yes. The Hoxton Freedom Group want me to address them on Sunday morning [great emphasis on "Sunday" ...]. What are they?

«The primary QUD is answered positively and removed from the QUD list.

The secondary QUD is answered and removed from the QUD list, which is now empty.

[[Morell is annoyed about another invitation $D-x_1$]] and **[[$D-x_1$ is an invitation to address The Hoxton Freedom Group $D-x_2$ on Sunday morning.]]** are added to the DCG.

[[Who are $D-x_2$?]] becomes the QUD. »

PROSERPINE. Communist Anarchists, I think.

«**[[$D-x_2$ are Communist anarchists]]** is added to the DCG.

Optionally, it may be added that the evidence for this is somewhat weak.

(Notice, this is a natural case in which a claim to knowledge doesn't accompany an assertion.)

The QUD list is emptied. »

MORELL. Just like Anarchists not to know that they can't have a parson on Sunday! Tell them to come to church if they want to hear me: it will do them good. Say I can only come on Mondays and Thursdays. Have you the diary there?

«A facetious piece of rhetoric is ignored.
 «Turn down $D-x_1$ » is added to Proserpine’s goals.
 «Tell $D-x_2$ that Morell can only come on Mondays and Thursdays» is added to Proserpine’s goals.
 A diary $D-x_3$ is made salient.
 The QUD becomes «Does Proserpine have $D-x_3$ handy?»
 A salient goal «Find an open date to lecture to $D-x_2$ » is added. »
 PROSERPINE (taking up the diary). Yes.
 ««Proserpine has $D-x_3$ » is added to the DCG.
 Progress has been made on the salient goal of finding an open date.
 The QUD list is emptied. »
 MORELL. Have I any lecture on for next Monday?
 ««Does Morell have a lecture on next Monday?» becomes the QUD. »

⋮

13. Conclusion

This paper was advertised as an attempt to unify three topics: human language, computer programming languages, and the genetic code. The unifying scheme is this.

The *Dynamic language scheme* (DL scheme):

- (1) There is a compositional symbolic system producing a set of complete utterance types, a set of expressions belonging to each type, and a semantics for each expression.
- (2) There is a dynamic system—a discrete arrangement of states which can change over time—associated with the language.
- (3) Complete utterances can be mapped to state changes of the dynamic system. The mapping may be fully determined by the semantics of the utterance, or may, in the case of human languages, also depend on situational pragmatic reasoning.

I felt that no justification was required for the claim that computer programming languages conform to the DL scheme; this is the standard account accepted by computer scientists. The body of this paper was devoted to showing that looking at human languages in the same way can be useful and illuminating.

The main gap in the present paper is the case of the genetic code; that is a work in progress, still in the initial phases and perhaps unworkable. It isn’t easy to reconstruct morphogenesis as a dynamic system, and in particular to find a way to characterize the appropriate system states. [Simon, 1962], a classic early contribution to the complex systems literature, provides a rationale for the analogy between ontogeny and program execution that is still plausible. On the other hand, there are reasons to think that the relation of genetic coding to morphogenesis may be significantly unlike the relation of a computer program to program execution in some ways; see, for instance, [Atlan and Koppel, 1990]. Nevertheless,

an appropriate formalization must be possible if molecular genetics is at all sound and correct, and only by trying to think through the details can we see which viewpoint is more correct.

Some readers may wonder why the “Language of Thought” wasn’t included as a fourth case. Well, the Language of Thought would certainly conform to the DL scheme. And in the engineered psychology of cognitive architectures this makes perfect sense and the details can be reconstructed from the implementations. But discussions of this topic for human psychology are speculative and belong more to the foundations of psychology than to psychology itself. If psychologists develop a way to confirm and disconfirm specific hypotheses about human LOT, I would be happy to include it as a case.

Bibliography

- [Atlan and Koppel, 1990] Henri Atlan and Moshe Koppel. The cellular computer DNA: Program or data. *Bulletin of Mathematical Biology*, 52(3):225–348, 1990.
- [Austin, 1962] John L. Austin. *How to Do Things With Words*. Oxford University Press, Oxford, 1962. Edited by J.O. Urmson and G.J. Warnock.
- [Clark and Marshall, 1981] Herbert H. Clark and Catherine R. Marshall. Definite reference and mutual knowledge. In Arivind Joshi, Bonnie Webber, and Ivan Sag, editors, *Elements of Discourse Understanding*, pages 10–63. Cambridge University Press, Cambridge, England, 1981.
- [Fillmore, 1997] Charles Fillmore. *Lectures on Deixis*. CSLI Publications, Stanford, California, 1997.
- [Ginzburg, 2012] Jonathan Ginzburg. *The Interactive Stance: Meaning in Conversation*. Oxford University Press, Oxford, 2012.
- [Kamp and Reyle, 1993a] Hans Kamp and Uwe Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics in Natural Language, Formal Logic and Discourse Representation Theory, Vol. 1*. Kluwer Academic Publishers, Dordrecht, 1993.
- [Kamp and Reyle, 1993b] Hans Kamp and Uwe Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics in Natural Language, Formal Logic and Discourse Representation Theory, Vol. 2*. Kluwer Academic Publishers, Dordrecht, 1993.
- [Karttunen, 1967] Lauri Karttunen. The identity of noun phrases. The Rand Corporation. Available at <https://www.rand.org/pubs/papers/P3854.html>, 1967.
- [Lewis, 1969] David K. Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, Massachusetts, 1969.
- [Lewis, 1975] David K. Lewis. Languages and language. In Keith Gunderson, editor, *Language, Mind, and Knowledge. Minnesota Studies in the Philosophy of Science, Vol. 7*, pages 3–35. University of Minnesota Press, Minneapolis, Minnesota, 1975.

- [Lewis, 1979] David K. Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359, 1979.
- [Milkowski, 2013] Martin Milkowski. *Explaining the Computational Mind*. MIT Press, Cambridge, Massachusetts, 2013.
- [Newell, 1992] Allen Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1992.
- [Piccinini, 2015] Gualtiero Piccinini. *Physical Computation: A Mechanistic Account*. Oxford University Press, Oxford, 2015.
- [Portner, 2009] Paul Portner. *Modality*. Oxford University Press, Oxford, 2009.
- [Putnam, 1988] Hilary Putnam. *Representation and Reality*. The MIT Press, Cambridge, Massachusetts, 1988.
- [Roberts, 1996] Craige Roberts. Information structure in discourse: Towards an integrated formal theory of pragmatics. Technical report, Linguistics Department, The Ohio State University, Columbus, Ohio, 1996. OSU Working Papers in Linguistics, volume 49, Jae-Hak Yoon and Andreas Kathol, editors.
- [Simon, 1962] Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.
- [Stalnaker, 1973] Robert C. Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457, 1973.
- [Thomason, 1990] Richmond Thomason. Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Philip R. Cohen, Jerry Morgan, and Martha Pollack, editors, *Intentions in Communication*, pages 326–363. MIT Press, Cambridge, Massachusetts, 1990.
- [Thomason, 2000] Richmond H. Thomason. Modeling the beliefs of other agents. In Jack Minker, editor, *Logic-Based Artificial Intelligence*, pages 375–473. Kluwer Academic Publishers, Dordrecht, 2000.
- [Thomason, 2013] Richmond H. Thomason. Loosening up agents and their propositional attitudes. In Franck Lihoreau and Manuel Rebuschi, editors, *Epistemology, Context, Formalism*, pages 167–184. Springer Verlag, Berlin, 2013.