

Donnellan's Blocks

By John Perry

In Joseph Almog and Paulo Leonardi (eds.), *The Philosophy of Keith Donnellan*. Oxford University Press, 2012

§1. DONNELLAN'S THEORY OF BLOCKS

The name 'Jacob Horn' comes from the book, *The Horn Papers* (Horn 1945), published and apparently written by William Horn of Topeka, Kansas. He presented this book to the public as if it were the newly discovered diary of an eighteenth century American, his great, great, great grandfather. At first the diary was taken to be authentic, but scholars eventually concluded that it was a hoax. I'll assume that the scholars were right, and Jacob Horn does not exist and never did. The example is due to Keith Donnellan, in "Speaking of Nothing" (1974).

Now consider these statements,

(1) Jacob Horn does not exist.

(2) Jacob Horn exists.

(3) Jacob Horn was an important person in Colonial America.

(1) is true. (2) and (3) are false.

I take statements to be assertive utterances of declarative sentences.¹

Utterances are acts, concrete events that happen once. The numbered examples in this paper are used somewhat inconsistently but not I think incoherently.

Primarily, they are supposed to stand for particular statements, assertive utterances, made by me for illustrative purposes in writing this paper.

Secondarily, they are used for the declarative sentences used in these utterances,

¹ This is a somewhat special use of 'statement', not intended to capture all the ways we ordinarily use the term. 'Statement' is often used for something various assertive utterances have in common, those that "make the same statement".

as in, "Consider an utterance of (2) by someone who believed in Jacob Horn was real." I don't think this inconsistent usage will be particularly confusing.

Utterances then are concrete acts that occur at times and have speakers. Utterances typically have content, which we usually report by ascribing contentful actions to the speaker at the time: Smith said yesterday that Obama was born in Hawaii; Elwood asked whether Hawaii was a state; Patricia denied that Obama was born in Kenya. The acts in question are a statement, a questioning, and a denying. The contents of statements are propositions, the propositions that the person making the statement asserted, i.e., what the person said. So the proposition that Obama was born in Hawaii is what Smith said, and the content of his statement.

I take propositions (and other contents) to be abstract objects that encode the truth-conditions or other relevant satisfaction conditions of utterances. I don't rely on any particular theory of propositions in this essay, but I do appeal to singular propositions, that is, propositions with respect to which we can speak of an object or objects as constituents. A theory that does not support this way of talking won't do.

According to the *referentialist* theory of proper names, an utterance of "Donnellan once taught at Cornell," expresses a singular proposition with Donnellan and Cornell as constituents. This proposition will be true in any possible situation in which Donnellan once taught at Cornell, whatever he is called, and Cornell is called, in that situation. The last remark ignored issues of tense and time, and for the most part I will continue to do so.

Utterances and sentence tokens are both particulars, but they should be distinguished. Sentence tokens are typically produced by the speaker in uttering; that is, she produces a burst of sound, or marks on a piece of paper. Sentence types have meaning, in accord with the lexical conventions and syntactic rules of the languages in which they occur. I call 'Jacob Horn' an 'empty name'. However, being empty is an attribute only of certain uses or utterances of 'Jacob Horn'; the name itself no doubt has many non-empty uses for perfectly real Jacob Horns. I

use 'what is said', at least for now, as an equivalent of 'the proposition expressed'. I assume that the somewhat philosophical locution, 'what is said' stands for just what the speaker of the utterance says.

Empty names like 'Jacob Horn' present what seems to be a basic problem for referentialism about names. To repeat, referentialism about proper names is the view that utterances of them contribute the thing they designate to the content of the statements of which they are a part, to what is said by the speakers of such utterances, to the singular propositions expressed by such utterances. Since 'Jacob Horn' designates nothing, it has nothing to contribute. We don't seem to have an account of which propositions (1)--(3) express. The problem seems especially acute with respect to (1) since it is true, and so presumably not radically defective in any way. How can it be true, if it doesn't express a true proposition?

Like Donnellan, I will defend the view that (1) is true, and (2) and (3) are false, in spite of the fact that there seem to be no singular propositions available for them to express, and that this can be explained within a recognizably referentialist theory of proper names. I will claim that utterances can be true or false, in spite of not expressing such propositions, in virtue of meeting or not meeting their *network* truth-conditions. And I will eventually adjust the account of "what is said" to accommodate this. The key concepts in developing these claims are *network content* and *blocks*.

Both of these concepts are based on ideas in Donnellan's essay, in which he is primarily concerned with showing that a referentialist theory of the sort he espouses can explain the truth of (1) and the falsity of (2)² Donnellan had criticized the descriptive theory of names, and suggested historical chains of reference as an alternative way of accounting for the link between the bearer of

² The idea of an historical or communicative or informational or causal chain, linked by intentions to continue using the name in the same way, stretching between a name or a use of a name and its referent, also comes up in discussion of names by Peter Geach (1969), Saul Kripke (1980) and Gareth Evans (1984). See also David Kaplan (1969) for related ideas.

the name and uses of the name, in "Proper Names and Identifying Descriptions" (1979) In "Speaking of Nothing" he appeals to these historical chains to give truth-conditions for (1) and (2).

Donnellan provides the following rule for determining the truth-value of a statement of the form 'N does not exist' where N is a proper name:

If N is a proper name that has been used in predicative statements with the intention to refer to some individual, then 'N does not exist' is true if and only if the history of those uses ends in a block.

Donnellan explains a block as follows:

When the historical explanation of the use of a name (with the intention to refer) ends . . . with events that preclude any referent being identified, I will call it a 'block' in the history (23).

Think of the history of a use (of a name) as a stream of events that leads up to that use. In determining the referent of a use of 'Jacob Horn', we go upstream, back in time; we are looking for an object that plays the right role to be the referent. Sometimes we don't make it back to such an object. There is an event that blocks us. In the case of 'Jacob Horn', the event was William Horn writing a fictional diary, *The Horn Papers*. That's where the stream begins. No participant in this event plays the right role to be the referent of 'Jacob Horn'. Hence, (1) is true.

Donnellan's rule is *reflexive*. The truth-condition for a statement containing an utterance of an empty name is a condition on the utterance *itself*. Donnellan's account is an instance of providing what I call a *reflexive truth-condition* for an utterance. That is, the truth-conditions are given in terms of the utterance *itself*, rather than the subject matter of the statement. The propositions that encode these truth-conditions I call *reflexive contents*. Donnellan's strategy is thus congenial to, and foreshadows, what I call the *reflexive-referential theory* (Perry 2001, Perry forthcoming). In this essay I try to use that theory to extend Donnellan's ideas to provide an account not only of the truth-conditions of (1)

and (2), but also an account of their content, what one says with (1) and (2), and what one says generally with statement like (3) that involve empty names, and what the participants in a debate about existence, like the scholars who poured over *The Horn Papers*, are saying to one another.³

§2. THE REFLEXIVE-REFERENTIAL THEORY

The basic idea of this theory is that statements have several levels of truth-conditions, depending on what we take as fixed. According to this theory, is that utterances have truth conditions *relative to* a set of facts we take as fixed or given about the utterance. The truth-conditions are what *else* the world has to be like, for the utterance to be true, *in addition to* what we take as fixed. The reflexive content of Donnellan's statement, "I used to teach at Cornell," *given* that the speaker is using English, is that the speaker of *that very utterance* taught at the place to which he refers with 'Cornell' before the time of *that very utterance*. The content is a singular proposition about the utterance *itself*, hence 'reflexive content'. Truth-conditions are reflexive so long as the truth-conditions put conditions on the utterance itself. If we add to what is given, that the speaker uses 'Cornell' to refer to the university in Ithaca (as opposed, say, to the college in Iowa), then the truth conditions are that the speaker of *that very utterance* taught at Cornell University, a singular proposition about the utterance and Cornell. These truth-conditions are still reflexive, since they impose conditions on the utterance itself. Thus, strictly speaking, I shouldn't talk about *the* reflexive contents.

However, I will continue to do so, having in mind the truth conditions given only the language and the meanings it associates with the words and syntactic structures, but not the facts about the basic context (speaker, time and location), the extended context (what the speaker is attending to, and his intentions relevant to the use of demonstratives, anaphora and other such issues) and the reference of

³ A version of this paper was presented in March 2008 at the Donnellan Conference in Bologna. That version was based on the account of the "no-reference" problem provided in the first edition of *Reference and Reflexivity* (2001). This account will be substantially changed in the second edition (forthcoming), and the present essay is based on this revised account.

names. I ignore issues of lexical and syntactic ambiguity, although the theory is useful in dealing with them.

If we add to what is given the fact that the speaker is Keith Donnellan and the time 2010, we get what I call the *referential content*. Given all of that, truth requires *in addition to what is given* that Keith Donnellan taught at Cornell before 2010. This proposition is not reflexive; it does not impose conditions on the utterance; it would be true, even if the utterance did not exist.

The reflexive content of an utterance is not what we would think of as what the speaker said. In this case, it is a proposition about Donnellan's utterance, not about Donnellan. We would take what he said to be the referential content of his utterance, a proposition about him but not about his utterance, that *he* used to teach at Cornell. However, the two propositions, the reflexive content of the utterance and its referential content, will have the same truth-value in the actual world, although not in all counterfactual circumstances, in some of which the utterance itself will not have occurred.

I claim that various problems that are thought to attach to referentialist theories of names can be solved if we take reflexive content into account, that is, if we generalize Donnellan's strategy. So construed, referentialism maintains that the proposition expressed by an utterance of "I used to teach at Cornell," spoken by Donnellan, and an utterance of "Donnellan used to teach at Cornell," spoken by anyone who is using 'Donnellan' to refer to Keith Donnellan, both express the same proposition; what is said is the same. But it allows that other levels of content can be used to explain differences in "cognitive significance". For example, the reflexive contents of "I am I", and "I am Donnellan" differ, even if both are spoken by Donnellan. The first is true if its speaker is its speaker; the second is true if its speaker is the person its speaker uses 'Donnellan' to refer to. As I explain below, it is due to this difference that the second is a reasonable way for Donnellan to introduce himself to someone, but the first is not.

Reflexive content gives us a start on explaining empty names in a way that can incorporate Donnellan's account. The reflexive content of (1) is that there is no

one that *its* speaker uses 'Jacob Horn' to refer to', because the historical chain of uses ends in a block. The reflexive content of (2) is that there is such a person; of (3) that there is such a person and he was an important person in colonial America. This provides the utterances with some sort of content, and explains why (1) is true and (2) and (3) are not.

But this account isn't satisfactory as it stands. For example, the reflexive truth-conditions of each utterance of "Jacob Horn was an important person in colonial America" will be different, because the reflexive content of an utterance are conditions on the utterance itself; that is, a singular proposition about that very utterance. But it seems that people who utter this sentence agree about something, and others who deny it disagree with that same thing. This thing, that some scholars, taken in by the Horn papers, thought was true, and others who were more skeptical thought was false, doesn't seem to be the reflexive content of any particular utterance of the sentence. If there were a Jacob Horn, a singular proposition about him would serve this purpose. Lacking such referential content, we still seem to need something more intersubjectively available than reflexive content, something which different utterances of (1), (2) and (3) can share. Networks --- something like Donnellan's historical chains --- are the natural place to find it.

§3. NETWORKS: A SIMPLE THEORY

NAMES

On my view of names, many people, places, and other things have the same name. I know several David Kaplans. I know only one Keith Donnellan, but a Google search shows that there are others. There are probably many Jacob Horns that do exist, in addition to the one we seem to be talking about, who doesn't. The forty-first and forty-third presidents of the United States share the name 'George Bush'. I share the name 'John' with an enormous number of other people. Thoughtful people have defended other views. Where I see shared names, David Kaplan (1990) would say that we really have shared vocables; the

different David Kaplans have different names, because names are individuated historically; the different names sound the same and are spelled the same, a bit like 'pole' meaning stick and 'pole' meaning axis, the first coming from Latin, the second from Greek. This makes names a lot like individual constants, which has its attractions. Epistemologically, it trades the difficulties of knowing which thing a name is being used to refer to, for difficulties in knowing which name is being used. The theory developed below could be altered to accommodate Kaplan's view of names. But I'll stick with the view that seems more natural to me.

Names are assigned to people and things by permissive conventions. When my wife and I named our first son 'Jim', we established a convention that permits people to refer to him with that name; it doesn't preclude anyone from referring to other Jims with the same name. When people use 'Jim' to refer to Jim, they are *exploiting* the convention Frenchie and I established. On a simple view, we could identify such conventions with pairs of names and things named. But this won't work for Jacob Horn. We only have a name, not a thing named. We can imagine a second, independent hoax using the same name. It seems like we have two conventions, each established without benefit of a thing named. The simple view won't work. But on Donnellan's account, we have two historical chains, leading back to different blocks. The two conventions involving the same name can be distinguished by the networks involved.

COCO-REFERENCE

I will say that names *exploit* conventions, and naming conventions are *supported by* networks, even when the networks end in a block. I think of networks as having utterances that refer or purport to refer as nodes. Think of the universe of utterances that are suited to refer to things, laid out in space and time, from the beginning of language up to the present. A network will be a possibly branching path through a subset of such utterances. Paradigmatically, the path will start with an object, what I'll call the *origin* of the network. Around the time a place or object is discovered for the first time, or comes into existence --- a person is born, a tree sprouts, a house is built --- people start referring to it. When my son Jim

was born in 1963, his mother and I and the obstetrician had already been referring to him for some time, but not by name, since in those days one did not know whether one had a boy or girl until birth. Once he was born, people started referring to him in a variety of ways, "that baby", "your baby", and so forth. By evening we had established the permissive convention that he could be called 'Jim', and people began to refer to him with that name. We spread the word by telephone, so people began referring to him who had never seen him. Their references to him were dependent on ours. So we have a variety of utterances standing to one another in two important relations. All these references were to the same individual, Jim; they co-referred. In addition, between some pairs of later and earlier utterances there is a relation of *intentional co-reference*; the later reference was intended to co-refer with the earlier one. Finally, there is a relation that I call *conditional co-reference* or *coco-reference*. First we'll look at intentional co-reference, and then at coco-reference.

When I told my mother that Jim had been born, she began to use the name 'Jim' with the intention of co-referring, that is, of referring to the same baby I had referred to with that name. When she told my father, he started referring to Jim using his name with the intention of referring to the same baby my mother had referred to.

People who see a thing or are otherwise aware of it fairly directly, will often have an intention to refer to the object they are aware of, and when appropriate, to co-refer with other utterances in virtue of doing that. Suppose I am looking at Jim, and say to the nurse "Isn't he the cutest baby you have ever seen?" She replies, as sincerely as she can, "He certainly is." She intends to refer to Jim. According to the rather relaxed and implicit standards of intention I shall employ, however, she also intends to co-refer with my utterance. The intentions involved in intentionally co-referring needn't be full-blown conscious ones; it is a practice that comes naturally to those who know how to refer and carry on conversations, most of whom have no explicit concept of reference or co-

reference. We knew how to refer, and to co-refer, long before philosophers started worrying about names and producing theories of what we were doing.

As word spreads to those who are only aware of an object as a result of hearing or reading about it, the pattern changes. At least initially my mother intended to refer to Jim *by* co-referring with me. My father intended to refer to Jim *by* co-referring with my mother. Whether the co-referring is the primary intention (one intends to refer by co-referring), or a derived one (one intends to co-refer by referring), it suffices to get a network going, as the chains of co-reference spread out in time and space.

Suppose however that my wife's pregnancy and Jim's birth were all part of an elaborate hoax my wife and I perpetrated to get our parents to quit harassing us, at least for a while, about their need for grandchildren. In this case there would have been no Jim, no reference to Jim, and no pairs of utterances that co-referred to Jim. Still, the causal relation between my utterance and my mother's, and her utterance and my father's, would have been basically the same as they really were, as would be the intentions involved in their utterances. My mother would have intended to refer to the same child I had referred to, and if I had referred to a child she would have succeeded.

Coco-reference is short for "conditional intentional co-reference". This suggests that coco-reference is a species or special case of co-reference, but that's not quite right. A later utterance *co-refers* with an earlier one, if both utterances refer, and refer to the same thing. A later utterance conditionally co-refers, or *coco-refers*, with an earlier one, if conditions are such that, the later utterance will refer if the earlier one does, and if it refers it will refer to the same thing as the earlier one. So there are cases of coco-reference that are not cases of reference, and so not cases of co-reference. We know how to make utterances that coco-refer, and we do so intentionally, but not in a way that requires that we have the concepts of reference, co-reference, or coco-reference.

To make my assumptions clear, I think there is a relation between utterances of the sort apt for referring, which insures that they will co-refer, if

either of them refers. This relation does not derive from referential relations the two utterances have to the same object. An analogy is that wheels attached to the same axle will rotate at the same speed, if they rotate at all. Call rotating at the same speed, 'co-rotating'. Wheels can co-rotate, even if they are not attached to the same axle, and they can be attached to the same axle, and not co-rotate, because neither of them is rotating at all. Call 'coco-rotating' the state of being connected so that, if either rotates, both will, and at the same speed. Connecting wheels to the same axle is the simplest way of establishing the relation of coco-rotating. Being connected by an axle does not insure that either wheel rotates, or ever has or ever will, only that if one does, both will, and at the same speed.

Just as an axle makes it the case that if one wheel rotates, the other one will too, and at the same speed, the relation of coco-referring makes it the case that if the earlier utterance refers, the later one will too, and to the same thing. So coco-referring isn't a special case of co-referring; it is rather being related, due to perception, memory and intention, in such a way that if reference takes place, it will be co-reference.

I've said what connection is involved in the case of coco-rotating (or at least identified a simple one that suffices): being attached to the same axle. In the case of coco-referring, there are various forms the connection may take. The relation may arise in three ways, which I'll call 'dependent', 'co-dependent' and 'convergent'.

Anaphora is a dependent way of coco-referring. You say, "Robert lost his job," and I reply, "He must be very worried". I use the pronoun 'he' as an anaphor for your use of 'Robert'; I intend for to "pick up" the reference of your use of 'Robert'; my utterance only refers in virtue of your utterance referring. In intending to co-refer, I implicitly intend to coco-refer, that is, to refer to the same object that you do, if there is such an object. My utterance coco-refers because the mechanism of reference for my utterance is the same as for your utterance, plus the additional anaphoric step. Anaphoric relations reflect the intentions of the

speaker. Such intentions may be conveyed, particularly within a given utterance, by the use of suitable syntactic forms.

I could also reply, "Robert must be very angry." This is co-dependent coco-referring; the reference of my use of 'Robert' does not depend on your use of that name; I could have said the same thing even if you had not spoken. But I intend to exploit the same convention that you do, and hence employ the same mechanism of reference. Again I implicitly conditionally refer, in that I intend that my utterance will have the same referent as yours, if yours has a referent.

In convergent coco-reference the reference of the second utterance does not really depend on that of the first, nor does it exploit the same convention, but the later speaker intends and expects, because of various auxiliary beliefs, that his utterance will refer if the earlier one does, and to the same object. You say "San Sebastian is a nice place" and I reply "Yes, I have been to San Sebastian many times." That's co-dependent co-referring: I employ the same name and exploit the same convention. But I can also coco-refer using a different form of reference that doesn't depend on yours. Perhaps I reply, "Yes, you are correct, Donostia is quite beautiful," intending for you to learn a new name for the city in virtue of recognizing my intention to coco-refer. My intention is based on an auxiliary belief, that Donostia and San Sebastian are the same city. Or perhaps we are flying over the Basque country, and I point out the window and say, "You are right, that city is quite beautiful," believing that the city I point to is San Sebastian. My use of the demonstrative 'that city' is intended to coco-refer to your use of 'San Sebastian,' but it is not anaphoric on it, and does not exploit the convention associated with 'San Sebastian'. These are cases of convergent coco-referring.

One can intend to coco-refer with an earlier utterance, but fail to do so, because one's auxiliary beliefs are false. I say, referring to David Kaplan the doctor who teaches in Stanford's medical school, "David Kaplan lives in Los Altos." My wife, referring to David Kaplan the logician, asks, "When did David Kaplan move to Los Altos?" She intended to exploit the familiar convention for using the name 'David Kaplan' that would ordinarily be relevant for a

conversation between us, and thought that by doing so she would be employing the same convention I did, and hence coco-referring. She succeeded in referring to the logician David Kaplan, but not in coco-referring or co-referring with me.⁴ Along the same lines, I might have been wrong about 'Donostia' being another name for San Sebastian, or I might have been wrong in thinking that the city we saw from the plane was San Sebastian.

One can intend to coco-refer with more than one earlier utterance. Several of us have been talking about President Obama; you continue the conversation by saying, "I hope he gets re-elected." You intend to co-refer with several of our utterances. This raises another possibility of mistake. Suppose someone overhears the conversation between my wife and I of the last paragraph, and says, "I don't think he has moved to Los Altos. He still teaches at UCLA." He intends to dependently co-refer with my utterance of 'David Kaplan' and my wife's utterance of the same name. But he can't succeed in doing both, since we refer to different David Kaplans, exploiting different conventions.

When the later speaker thinks the object in question does not exist, or suspects it does not, the conditional intention is likely to be more conscious and explicit. Suppose Anissa is talking about Santa Claus to her younger brother Everett. Everett believes in Santa Claus; Anissa isn't sure whether or not he exists. Everett asks, "Does Santa Claus live right on the North Pole, or just nearby?" Anissa replies, "No one really knows exactly where he lives". Her use of 'he' is a case of dependent coco-reference. If she had replied, "No one knows exactly where Santa Claus lives," her utterance of 'Santa Claus' would have been a case of co-dependent coco-reference.

Her intention is that her utterance and Everett's will have a certain relationship so that if Everett's utterance refers to something, hers will refer to the

⁴ Alternatively, she might have intended to refer to the logician by coco-referring, and succeeded in co-referring with me, in referring to the doctor, and in saying something false about him, but failed to refer to the logician. Or, most likely, her intentions have no clear structure, and the case is indeterminate.

same thing. In the first, anaphoric, case, the dependence of her utterance on his guarantees this relationship. In the second case, her utterance is sufficiently similar to his, so that the facts sufficient to make a certain object the referent of his utterance would suffice to make the same object the referent of hers. If, contrary to her suspicions, there is someone suitably related to Everett's utterance of 'Santa Claus' to be its referent, that same person will be the referent of her utterance of the same name. And, as we saw in the San Sebastian case, she doesn't have to use the same name. She could have said "No one knows where Kris Kringle lives," with the intention of conveying to Everett that Santa Claus has another name, too. She believes that if the facts are such that 'Santa Claus' refers, then they are such that 'Kris Kringle' refers, and refers to the same person.

Even if Anissa has become a complete non-believer, she can still coco-refer with Everett. She can carry on the conversation; she intends her utterance to fit with and be responsive to Everett's in the way that intentionally co-referring utterances are, even though she is under no illusions that reference is taking place. In each case Anissa coco-refers with the Everett's attempts to refer to Santa Claus. She intends for her utterance to meet all of the conditions of co-reference with Everett's except the condition of reference.

An utterance can coco-refer with several earlier utterances. Everett and Anissa are conversing, after their mother told him about Santa Claus. In answering Everett's questions, Anissa is coco-referring with utterances of Everett and utterances of her mother. And an utterance can be coco-referred to by any number of later utterances. So coco-referring is both a many-one and a one-many relation; hence, coco-reference chains can branch in both directions.

COCO-NETWORKS AND CONVENTIONS

The relation of coco-reference induces the networks I have in mind; the networks can exist even if there is no origin, no reference, and no co-reference. Thus the historical chain that Donnellan appeals to in his explanation of the truth of (1) can be considered such a network.

I'll say that an utterance is a *coco-descendant* of an earlier utterance if there is a chain of coco-referring utterances from it to the earlier one. Elwood tells you that Jack drank the last beer. You leave me a note saying "Jack drank the last beer." I tell Marsha, "Jack should be whipped." She tells Fred, "Jack's becoming unpopular because of all the beer he drinks." Your note coco-refers with Elwood's utterance. My utterance coco-refers to your note, and Marsha's utterance coco-refers to mine. Marsha's utterance doesn't coco-refer with Elwood's utterance, but it is a coco-descendant of it, because of the chain of coco-references. None of this depends on Jack existing; I could have made Jack up in order to blame him for drinking missing beers that which I actually had taken.

We can stipulate that referential utterances (that is, utterances that are suited to refer, that "purport" to refer, whether or not they actually do) coco-refer to themselves. Coco-reference is then reflexive but neither transitive nor symmetrical. Being a coco-descendant of is reflexive and transitive but not symmetrical, and is many-one and one-many. Being a *coco-ancestor* is the inverse of being a coco-descendant. Utterances u and u' are *coco-connected* if there is a sequence of utterances, starting with u and ending with u' , each of which is a coco-descendant or coco-ancestor of the previous one. Being coco-connected is, (subject to qualifications below) an equivalence relation: reflexive, symmetrical, and transitive.

Consider the universe of utterances. Pick one utterance to my son Jim. It, together with all of its ancestors and descendants, is a coco-network. It is likely that there are a number of coco-networks with Jim as their origin. There is certainly a large network, which began with early references to him by my wife and I and probably includes most references to him since then. But suppose, as I assumed at the time, that everyone who walked by the hospital nursery noticed the unusually cute red-haired baby, and told their friends and family about him. These folks discussed the cute baby for a while, and then forgot about him. Each of these incidents would have started a coco-referring network with Jim as the origin that was not coco-related to the dominant network. Although Jim is not as

cute as he once was, this process no doubt continues. Someone walks by an auditorium where Jim is giving a talk. They ask the person next to them, "do you know who he is?" That person responds, "No, I have no idea who he is." The latter person's use of 'he' coco-refers with the first person's use of 'he'. We have a short coco-network with Jim as its origin. If they drop the matter and never find out whom they saw, this network will never attach with the dominant one. It will remain detached even if they later meet Jim and come to know him well, as long as they don't recognize him as the man they saw giving the talk.

We can think of the more typical network having a trunk with roots and branches. Roots begin with references that are not intended to co-refer; perhaps when Jim was born my wife and the doctor and the nurse all referred to him in various ways, quite independently. But as they continued to admire him, their references coco-referred to each other's earlier references; the roots formed a trunk. This trunk has extended for many years. The branches are more or less local subnetworks that start off from the trunk and continue until the end of a conversation; they may be parts of larger subnetworks that rejoin the trunk in a later conversation --- a bit unlike most real trees.

Also unlike real trees, a root can start independently of the trunk, and give rise to a branch, that later joins the trunk. I see a new person at the department meeting, and ask you, "Was she appointed while I was on leave?" Perhaps you don't know either, and pose the same question to the person sitting next to you. After a while someone answers the question with a name: "That's Alice Fritchey, our new logician." The root begun by my question joins the trunk, and my original reference to Alice Fritchey comes to share a network with thousands of references to Alice most likely going back to the day she was born, or before. That evening I tell my wife about Alice Fritchey and some of the interesting things I noticed about her; she tells her friends the next day at lunch; this is a branch, that may continue growing for years without ever feeding back into the main trunk.

Of course, a key fact about coco-networks for our purposes is that they need not begin with a reference to a real object; they can begin as in the Jacob Horn case or the case of Jack a couple of paragraphs above, with an act of invention. In such cases we have a block; none of the utterances in the network refer, although they do coco-refer. In such cases, we talk about "referring to the same person," and "having the same person in mind." But there is no reference, and no identity of referents; it is a loose way of talking about coco-referring.

This is our basic picture of networks. It is an idealization. The complications come the fact that coco-referring networks, along which reference flows, are also networks along which information flows, and such networks often involve misidentifications. Misidentifications often create what I call a *messes*. Messes occur when one referential utterance ends up being parts of two (or more) coco-networks, because of confused intentions, usually rooted in confused beliefs or misperceptions. Messes mean that coco-connectedness as defined is at best a near-equivalence relation. I discuss messes at length in *Reference and Reflexivity*, but will ignore them here.

A convention for name N is supported by a network, if there is a practice along the network to use N to coco-refer. A use of a name that exploits a convention refers to the origin of the network that supports the convention, if it has one; otherwise the convention and the use are empty.

A practice is inaugurated when a name is used as part of a network. It may be explicitly introduced for this purpose, by reference to an existing, salient object. Perhaps Alice's mother says, "Let's call her 'Alice'." Or perhaps she just starts using the name to refer to Alice. The practice is continued when others use the name as a part of the Alice-network. The permissive convention to call Alice 'Alice' is supported as long as the practice continues.

The network contemplated above supported a convention for the use of 'Jack', and the Jacob Horn network supports a convention for the use of 'Jacob Horn'. The conventions permit one to coco-refer with other acts of reference on the same network. If there is an origin, then by coco-referring, one will refer to it.

In this case, we have a permissive convention to refer to the origin with the name. Recall that by 'permissive', I mean to emphasize that the convention doesn't mean that one cannot use the name to refer to other individuals, or to coco-refer on other networks. So the convention that permits me to use 'Jack' for the beer-drinking guy I invented, doesn't preclude me from using 'Jack' to refer to other people I know by that name, or to coco-refer with people talking about Jack Frost, Jack Sprat, or the Jack found in the beanstalk.

A network can support more than one convention for referring to an individual in this sense. The network that leads back to the seventh president of the U.S. supports conventions for calling him 'Andrew Jackson' and 'Old Hickory'. New naming conventions can come into being with only indirect connections to the named object. Perhaps an inept student asks a question about Jackson, calling him 'Old Hickuppy'. Soon everyone in class calls Jackson 'Old Hickuppy'; the students have inaugurated a new naming convention.

The same network might support two conventions for calling the same person by the same name. Suppose the first American Pope was called 'Gary' by his parents, but was orphaned at an early age, and renamed 'Ellsworth' by his adoptive parents. Much later, when he becomes Pope, he chooses a new name, as is the practice, and quite by accident picks 'Gary', and becomes Pope Gary I. This rather unlikely possibility makes the point that when a network supports a convention, it is really a matter of some parts of the network supporting the convention. In this case the early part of the network, involving utterances by Gary's original parents and their friends and relatives support the first 'Gary' convention, and utterances by the Pope, his Cardinals, and the wider world after he becomes Pope support the second 'Gary' convention. I'll mostly ignore this complication.

§4. NO-REFERENCE RECONSIDERED

Let's return to our problematic statements about Jacob Horn:

- (1) Jacob Horn does not exist.

(2) Jacob Horn exists.

(3) Jacob Horn was an important person in Colonial America.

We can extend Donnellan's strategy, using these concepts, to provide reflexive contents for these utterances.

(Px 1) That there is a speaker S, a convention C and a network N such that S exploits C with the utterance of 'Jacob Horn' in (1), and N supports C, but there is no x such that x is the origin of N.

(Px 2) That there is an individual x, a speaker S, a convention C and a network N such that S exploits C with the utterance of 'Jacob Horn' in (1) and x is the origin of N.

(Px 3) That there is an individual x a speaker S, a convention C and a network N such that S exploits C with the utterance of 'Jacob Horn' in (1), N supports C, x is the origin of N, and x was an important person in Colonial America.

These reflexive contents seem to get the truth-values of (1)--(3) correct. But two problems remain. First, although we have an explanation of why (2) and (3) are not true, they seem not only to fail to be true, but to be false. In this way they differ from, say, a Lewis Carroll-inspired utterance of, 'Borogroves are mimsy.' Such an utterance doesn't really *say* anything; it is not true, but also, at least by my lights, doesn't even manage to be false. Second, it's not clear that we have a plausible candidate for what was at various times said and denied, believed and disbelieved by people who read *The Horn Papers*, and lots of other people they talked to, many of whom may not have read the book, or even known of its existence. To deal with these problems, we need to recognize levels of content between reflexive and referential content.

The reflexive contents listed get at what *else* the world has to be like, for the utterances to be true, *given* that they are English sentences and the words have their usual meaning. As I said above by *the* reflexive content I have in mind what

else has to be the case, for an utterance to be true, with the language and meaning fixed. As we instantiate the existential quantifiers on the objects that actually play the indicated roles, taking more as fixed and given, thus change what else has to be true. This may still be *a* reflexive content, as in the example above, where we fixed the referent of 'Cornell' but didn't fix the speaker. Similarly, if we fix Elwood Fritchey as speaker of (3), then the conditions that remain are conditions on Fritchey and his utterance:

That there is an individual *x*, a convention *C* and a network *N* such that Fritchey exploits *C* with the utterance of 'Jacob Horn' in (3), *N* supports *C*, *x* is the origin of *N*, and *x* was an important person in Colonial America.

These are still reflexive. But if we fix *C*-JH as convention the Fritchey exploits, we obtain a condition that is no longer bound to the utterance or Fritchey, but only to the convention:

That there is an individual *x*, and a network *N* such that *N* supports *C*-JH, *x* is the origin of *N*, and *x* was an important person in Colonial America.

This is not reflexive, but it still not referential. I call this the *convention-bound* content of (3).

Finally, if we fix *N*-JH as the network that supports *C*-JH, we obtain the 'network-bound content' or *network content* of (3):

That there is an individual *x* who is the origin of *N*-JH, and *x* was an important person in Colonial America

Again, this is not reflexive, but also not referential. This is a proposition about *N*-JH. The network content will be the same for any utterance of (3) supported by that network, including those that do not use the name 'Jacob Horn', but some other way of coco-referring to utterances along that network. It is promising, therefore, as a content that can serve as what people agree disagree about when they utter and deny (3)

§5. INCREMENTAL CONTENTS AND FALSITY

The three contents just considered are species of what I call *incremental contents*. In the first case, where we fixed the speaker but not the convention, the content was both incremental and reflexive. In the latter cases, the convention-bound content and the networks content, what we fixed was sufficient to replace occurrences of the utterance itself as a constituent of the level of content determined, in favor of objects in the world beyond the utterance, although not sufficient to get us all the way to referential content. Referential content is incremental content, but not all incremental content is referential.

In planning and understanding utterances all kinds of incremental but not referential content may be relevant. To return to an earlier example, suppose, Keith Donnellan introduces himself to someone at a party by saying "I am Keith Donnellan". The referential content is the same as if he had said, "I am I" or "Keith Donnellan is Keith Donnellan". But is not the referential content that Donnellan plans on conveying to his interlocutor. Given Donnellan's modest nature, he probably simply intends that this person learn his name, that is, comes to believe that the person he is talking to is named 'Keith Donnellan'. His implicit plan is that the person will grasp the reflexive content of his remark, in virtue of knowing English:

This utterance I hear is true iff there is a convention for the name 'Keith Donnellan', and a network, such that the speaker exploits the convention, the network supports the convention, the network has an origin, and that origin is the speaker.

The person will realize that the speaker is the person he is talking to, and so infer:

This utterance I hear is true iff there is a convention for the name 'Keith Donnellan', and a network, such that the person I am talking to exploits the convention, the network supports the convention, the network has an origin, and that origin is the person I am talking to.

From which she will infer something she might express with:

This person is named 'Keith Donnellan',

and, using the name she has just learned,

This person is Keith Donnellan

If the person were a philosopher or a linguist, she probably would already know of Keith Donnellan, and, perhaps given the nature of the party and the fact that 'Keith Donnellan' is not a terribly common name, could infer that the person she is talking to wrote the famous articles she had read in graduate school

The key to her learning Donnellan's name, and possibly figuring out that she was talking to someone whose articles she had read, is not the referential content of his utterance, but the content with the speaker fixed as the person she is talking to. It is this content that leads her to the knowledge Donnellan wanted to convey. This would not be a content of "I am I" nor of "Keith Donnellan is Keith Donnellan," and that is the reason "I am Keith Donnellan" would be a better way for Donnellan to introduce himself.

Imagine that we are driving. You are at the wheel. Suddenly, noticing the erratic behavior of the car you are about to pass, I shout, "That driver is drunk!" My plan is that you will grasp the reflexive content of my remark in virtue of knowing English, that you will identify me as the speaker, that you will realize that I am using 'that car' to refer to the driver of a car I am attending to, that you will follow my gaze to identify the person in question, and you will think, of a man you perceive, "That man is drunk," and take suitable precautions. The key to my plan is the demonstrative, "that driver". If I had recognized the driver as, say, our friend Clem, who is drunk most of the time, I could have conveyed the same referential content by saying, "Clem is drunk." But then my plan wouldn't have worked. Rather than taking defensive measures, you might have simply said, "So what else is new?" My plan relied on your grasping the incremental

content of my utterance, given that I was the speaker and I was using "that driver" to refer to someone I perceived.

I suggest that the difference between an utterance simply failing to be true, and being false, is that in the latter case the utterance provides an incremental content, a claim about the world beyond its mere reflexive content. Our Lewis Carroll inspired utterance, "Borogroves are mimsy," does not satisfy that condition. Its lack of truth has to do with deficiencies in the utterance itself, rather than in an incremental claim about the rest of the world that the utterance makes. It isn't true, but it isn't false either. It doesn't say anything. From now on, I will reserve the term 'statement' for assertive utterances that make such an incremental claim about the rest of the world. However assertively I utter "Borogroves are mimsy" or "Borogroves are not mimsy," I have not made a statement. But statements involving empty names that exploit a convention governed by a network, however short, do make statements, even if the network is blocked.

This means that an utterance can be a statement, and be false, even if it lacks referential content, so long as it has network content. Statement (2) has the network content that the Jacob Horn network has an origin. It not only fails to be true, it succeeds in being false.

Network content gets at what else the world has to be like, for an utterance to be true, given meaning and context and the facts about which networks support the naming conventions that are exploited. In "no-reference" cases, the relation between utterances of having reference governed by the same network, and having the same network content, will do part of the job ordinarily done by having the same referent, and the same referential content.

§6. BLOCKING AND PROMOTING

In a paradigm case of saying, we will have a speaker, a time, and a proposition. But there is more structure involved than is provided by three items. Saying is a kind of action, a kind of doing. In action, the *by* and *way of* relations are crucial.

What did Oswald do? He killed Kennedy, *by* shooting the rifle, which he did *by* pulling the trigger. Pulling the trigger of a loaded rifle is a *way of* shooting, and shooting a rifle accurately aimed at the head of a person is a *way of* killing them. And, typically, one can do the same thing in different ways, relying on different circumstances, and one can do different things in the same way, in different circumstances.

This structure, with respect to saying, is marked with the distinction between direct and indirect discourse. Donnellan *said that* he taught at Cornell, *by* saying, "I taught at Cornell". I said the same thing, *by* saying "Donnellan taught at Cornell". Sydney Shoemaker could do the same thing, while walking on campus, *by* saying "Donnellan taught here." But when he says "I taught at Cornell" he doesn't say the same true thing Donnellan says with those words.

To develop an account of appropriate answers to the question "what did he do?" a two-step procedure seems promising. First we consider all of the movements made and results accomplished: Oswald moved his finger, fired the rifle, shot Kennedy, killed Kennedy, killed the President, revealed his location in the School Book Depository, and changed the course of history. Then we could consider which of these possible answers is appropriate in a given conversational situation. Importantly, the various things he did have different properties, which may be relevant to different inquiries. Shooting the rifle, and thereby killing the President, are things Oswald did intentionally. The latter was his goal, the former his means. Revealing his location was probably something he did knowingly, but it wasn't part of his goal, or his means; he didn't do it on purpose. Perhaps the *default* answer to "What did he do?" is his goal, the intentional act his movement was directed towards accomplishing. If he didn't succeed, we typically revert to what he *tried* to do.

I think our concept of saying and what is said needs to be approached in a similar way. Referential contents pick out an extremely important aspect of most cases of saying. They are typically the focus of agreement and disagreement, and capture what people have the goal of conveying, although often with further

goals. Conversations may focus on referential contents, which their participants get at quite different ways: "He taught at Cornell," "I taught at Cornell", "You taught at Cornell," "Donnellan taught at Cornell". We pay attention to what is said, in the sense of referential content, not how it is said. Referential content is the *default*. This is what the great texts of referentialism draw our attention to, as part of weaning us from descriptive theories of names. It is why I sometimes call referential content 'official content', with that authority of these texts in mind. But it is not the only thing that can count as what is said.

I suggest that when we have a block, as in the Jacob Horn case, alternative incremental contents may be *promoted* to the status of what is said. When scholars discuss whether there was a Jacob Horn, what they disagree about is whether the Jacob Horn network ends in a block, or has an origin. Blocks promote alternative contents to the status of the subject matter that is the focus of agreement and disagreement.

However, I think that it is not only blocks of the type Donnellan contemplated that can do this. Not only the *lack* of referential content, but also the *irrelevance* of referential content, can block the path. We stop, and seek alternative contents to be our focus, even when there is a referential content, when it cannot supply an explanation of the relevance of the statement in question.

Return to the simple case in which Donnellan introduces himself, by saying "I am Keith Donnellan." The referential content of his remark is the same as it would be if he had said "I am I" or "Keith Donnellan is Keith Donnellan". But as we saw, this referential content does not capture what Donnellan meant to convey. His interlocutor could felicitously say "He said his name was 'Keith Donnellan'", but only a philosopher biting some bullet is likely to say, "He said that Keith Donnellan was Keith Donnellan." The triviality and irrelevance of the referential content leads us to choose alternative contents, in answering the question, "What did he say?"

In such cases, we may not expect the questions, "What did he say?" or "What was said?" or "What are they arguing about?" to have a unique acceptable answer. The structure, or possible structure, of available contents may have to be considered to understand what is going on. This is often the case with conversations about existence.

§7. TALKING ABOUT EXISTENCE

I will assume that existence is a condition that can be reasonably defined in terms of identity and the existential quantifier:

A exists iff there is an object x , such that $A = x$.

So existence is a property everything has, or at least a condition everything meets. Since George W. Bush is a real person,

That George W. Bush exists

is a perfectly good proposition, which we can take to be what someone says if they say, "George W. Bush exists". But what if someone says,

(4) George W. Bush does not exist.

It doesn't seem reasonable, in ordinary circumstances, to suppose that they intend to convey the proposition

That there is no object x such that $x = \text{George W. Bush}$

since if they were correct, there would be no such proposition. I think the ordinary use of this would be to express the belief that the George W. Bush network has no origin. It's the sort of thing that might be said by someone who feared a cadre of Republican industrialists who have control of the media had created an illusion based on old pictures of George Bush senior and skilled impersonators, in order to get an otherwise unelectable Dick Cheney into a position of power. So we might suppose that the self-defeating nature of the attempt to express the referential content is a sort of a block, which promotes the

false network content, that the George W. Bush network has no origin, to the level of what is said.

But suppose George W. Bush overheard (4) and protested,

(5) I do exist.

He would not seem to be making a remark about his network, but about himself. Shouldn't we suppose that "George W. Bush does not exist" and "I do exist," said by George W. Bush, directly contradict one another?

We need to bring in network content to understand the conversation. What is Bush's communicative plan? He intends to refer to the speaker of (5) in virtue of using 'I', and so to refer to himself. But he also intends to contradict (4). He wants to do this, because he takes himself to be the origin of the network that supports the use of 'George W. Bush' in (4); he takes himself to coco-referring with the use of his name in (4), and so referring to the George W. Bush whose existence is under discussion. The reference of his use of 'I' does not derive its reference from the earlier use of 'George W. Bush'. But the remark is unmotivated if he is not coco-referring with that use. He intends to use 'I' in the usual way, and thereby refer to himself, but he also intends, by doing that, to coco-refer with the utterance of 'George W. Bush' in (4). To get at Bush's no doubt somewhat implicit and unconscious plan for refuting (4), we need to bring in the whole structure of his plan.

My suggestion is that with existence statements, the default shifts to network content. Network content gets at the issue in question between those who assert something like (1) or (4), and those who disagree with them. Those who are asserting existence, rather than denying it, will think that by expressing the network content, they also predicate the property of existence to a real object, that is, they think that their remark will have referential content. Even when reference is not via a network, and has a referential content, in that, as with (5), there will typically be a salient network that is needed to understand the conversation of which it is a part. But not always. If George W. Bush, in an

uncharacteristically philosophical mood, simply uttered (5) to himself, there would be no relevant network content, and we would have to override that default.⁵

“George W. Bush might not have existed,” seems to have an unproblematic referential reading. The possibility would be witnessed by a variety of worlds in which George W. was never born. Among others, these would be worlds in which the senior George Bush and Barbara Bush did not have this particular child, worlds in which the Bush clan never got started, worlds in which life never evolved on earth. This can be reconciled with a generally network reading of existence statements, however, if we suppose it comes to:

There is an origin for the GWB network, but there might not have been. This is the sort of thing a former believer in the conspiracy theory of George W. Bush might say, in claiming that his former belief, although not true, wasn't completely absurd.

Someone who had believed in Jacob Horn, but became convinced otherwise, might say “Jacob Horn might have existed,” with similar motives. But, referentially interpreted, the remark seems clearly false. Since there is no Jacob Horn, there are no possible worlds in which Jacob Horn exists. What he means to claim, and what might be true, is that the Jacob Horn network, based on the evidence about it that was accessible to him, might have had an origin.

Assessment of counterfactual possibilities played a large role in motivating referentialism. We distinguish what the world would have to be like for an *utterance* to be true, from what it would have to be like for *what was said by the utterance* to be true. Given this distinction, we usually see that what the counterfactual possibilities have in common is the object referred to, rather than names for it, descriptions of it, or networks leading back to it. If we keep the link between what is said and counterfactual possibilities intact, but keep an open

⁵ I doubt, however, that a satisfactory account of such a Cartesian thoughts can get by simply with referential content; other forms of reflexive and incremental content would be required. But I won't go into these issues here.

mind about what those counterfactual possibilities are, it seems to me that the story is different when we consider existence statements. The counterfactual possibilities relevant to inquiry, the ones investigators seek to eliminate or establish as actual, do involve networks. The possibilities that the scholars investigating the reality of Jacob Horn tried to establish or eliminate were by and large possibilities in which the Jacob Horn network had an origin. The issue they were investigating, and used (1) and (2) to express their opinions about, was an issue about network content.

§8. CONCLUSION

Donnellan's theory of blocks provided a plausible account of the truth-conditions of affirmations and denials of existence that fit with his rejection of the descriptive theory, and his historical account of how name are linked with their bearers. He did not claim that it provided us with an account of what is said by affirmations and denials of existence, nor by other statements that involve empty names. I have argued that his ideas can be developed in ways that provide such an account.

REFERENCES

- Donnellan, Keith, 1966. Reference and Definite Descriptions *The Philosophical Review*, LXXV: 281--304.
- Donnellan, Keith, 1970. Proper Names and Identifying Descriptions. *Synthese*, 21: 335--358.
- Donnellan, Keith, 1974. Speaking of Nothing. *Philosophical Review*. LXXXIII: 3--31.
- Evans, Gareth, 1973. The Causal Theory of Names. *Supplementary Proceedings of the Aristotelian Society*, 47: 187—20
- Geach, Peter, 1969. The Perils of Pauline. *Review of Metaphysics*, 23, 287--300.
- Horn, William, ed., 1945. *The Horn Papers: Early Western Movement the Monongahela and Upper Ohio*. Scottsdale: The Herald Press.
- Israel, David, John Perry and Syun Tutiya, 1993. Executions, Motivations and Accomplishments. *The Philosophical Review*: 515--40.

Kaplan, David, 1969. Quantifying In. In *Words and Objections* eds. Donald Davidson and Jaakko Hintikka. Dordrecht: Reidel Publishing Company: 206--42.

Kaplan, David, 1990. Words. *Proceedings of the Aristotelian Society Supplementary Volume*, 64: 93--119.

Korta, Kepa and John Perry, forthcoming. *Critical Pragmatics*. Cambridge: Cambridge University Press.

Kripke, Saul, 1980. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.

Perry, John, 2001. *Reference and Reflexivity*. Stanford: CSLI Publications.

Perry, John, forthcoming. *Reference and Reflexivity*, 2nd edition Stanford: CSLI Publications.