Working Draft — for comments + criticisms. Please do not quote.

PRONOUNS AS PARAPHRASES

Terry Parsons

I INTRODUCTION

The purpose of this paper is to investigate a certain proposal about pronouns, namely, that certain of them ought to be generated from underlying definite NP's. For example, that the *it* in:

every man who owns a donkey beats it

ought to be generated from

the donkey he owns.

The proposal, when developed, has the virtue that it answers some outstanding questions about pronouns. Its major defect is that it is probably incorrect. My hope is that there are interesting things to learn from criticizing it, even if it fails to survive the criticism.

The proposal will be formulated roughly within the framework of Montague's system of "The Proper Treatment of Quantification in Ordinary English" (hereafter PTQ). I've tried to present the proposal in such a way that people who are unfamiliar with the framework will be able to follow it, at least in broad outline.

What is important about the PTQ framework for my purposes is that it provides a system of rules which mechanically associate English sentences with sentences of a fancied-up predicate calculus (I'll need hardly any of the fancy parts, so I'll just call it, somewhat inaccurately, "the predicate calculus"). Montague himself has a particular way of then giving a semantics for the predicate calculus translations of English sentences, and thus indirectly for the English sentences themselves. But I won't discuss his techniques here since there are many different views about how to give such a semantics, and for my purposes it simply won't be necessary for me to choose among them.

The system of PTQ is a generative one; it has rules which generate pairs of English sentences and predicate calculus sentences. But it's well-known in the computer science literature how to convert this system, anyway, into a recognition system, one that takes arbitrary strings as inputs and says whether they are generable and what possible translations they have. I think that the proposal I will make won't alter that fact.

There is one other way in which the system of PTQ is more general than is sometimes thought. Although Montague's syntactic rules are formulated in a way that is alien to the general custom in linguistics, it can be reformulated as a transformational grammar which has exactly the same outputs, both syntactic and semantic.

II PRONOUNS: THE PROBLEM

PTQ treats pronouns of English as bound variables of predicate logic. And there is a certain amount of prima facie evidence for this approach. For example, suppose that you "symbolize" the sentence "Some boy dates Mary" in the predicate calculus; you would write something like this:

$(\exists x)(Bx \mathbin{\&} xDm)$.

If asked to justify such a symbolization you could respond by "reading" the formula in English, in such a way that what you say is recognizable as a paraphrase of "Some boy dates Mary", namely:

$$(\exists x)(\qquad Bx \qquad \& \qquad xDm \qquad)$$

something is such that it is a boy and it dates Mary

Notice that when you come to the bound variables (outside the quantifier) you pronounce them as pronouns.

On the other hand, if you have a sentence of English which contains an explicit pronoun, the pronoun will often give rise to a bound variable. For example, if you symbolize "Every man who dates a woman that <u>he</u> likes is happy" you would produce something like:

$(x)(Mx \mathbin{\&} (\exists y)(Wy \mathbin{\&} xLy) \rightarrow Hx)$.

The pronoun "he" gives rise to the indicated bound variable "x" (if the original sentence contained "Sam" instead of "he" the translation would contain "s" instead of "x" in the indicated place).

The system of PTQ makes systematic use of this idea; whenever it generates a sentence which has a pronoun with an antecedant it simultaneously produces a predicate calculus translation which contains a bound variable that "comes from" the pronoun in an identifiable way. And the whole system works so well that in some cases it is almost startling. For example, the sentence:

a dog chased every cat

is ambiguous; it can mean either that there was a certain dog, and it chased every cat, or it can mean that every cat got chased by some dog or other. In PTQ these two readings are correlated with two different ways to generate the sentence; these two ways are embodied in the following two "analysis trees", along with the predicate calculus sentences correlated with the trees:

a dog chased every cat — $(\exists x_0)(Dx_0 \& (x_1)(Cx_1 \rightarrow x_0CHx_1))$

- a dog
  - dog
- he_0 chased every cat
  - every cat
    - cat
  - he_0 chased him_1
    - he_0
    - chase him_1
      - chase
      - he_1

a dog chased every cat — $(x_1)(Cx_1 \rightarrow (\exists x_0)(Dx_0 \& x_0CHx_1))$

- every cat
  - cat
- a dog chased him_1
  - a dog
    - dog
  - he_0 chased him_1
    - he_0
    - chase him_1
      - chase
      - he_1

Now a closely related sentence is *not* ambiguous; the sentence:

a dog chased every cat that provoked it

has only a reading corresponding to the former one above. And intuitively this is because of the presence of the pronoun in the object. Well, that's exactly what happens in PTQ; the correct reading is generated as follows (the system of PTQ uses only "unreduced" relative clauses; I'll stick to this, but in giving examples I'll alter them to their more idiomatic forms):

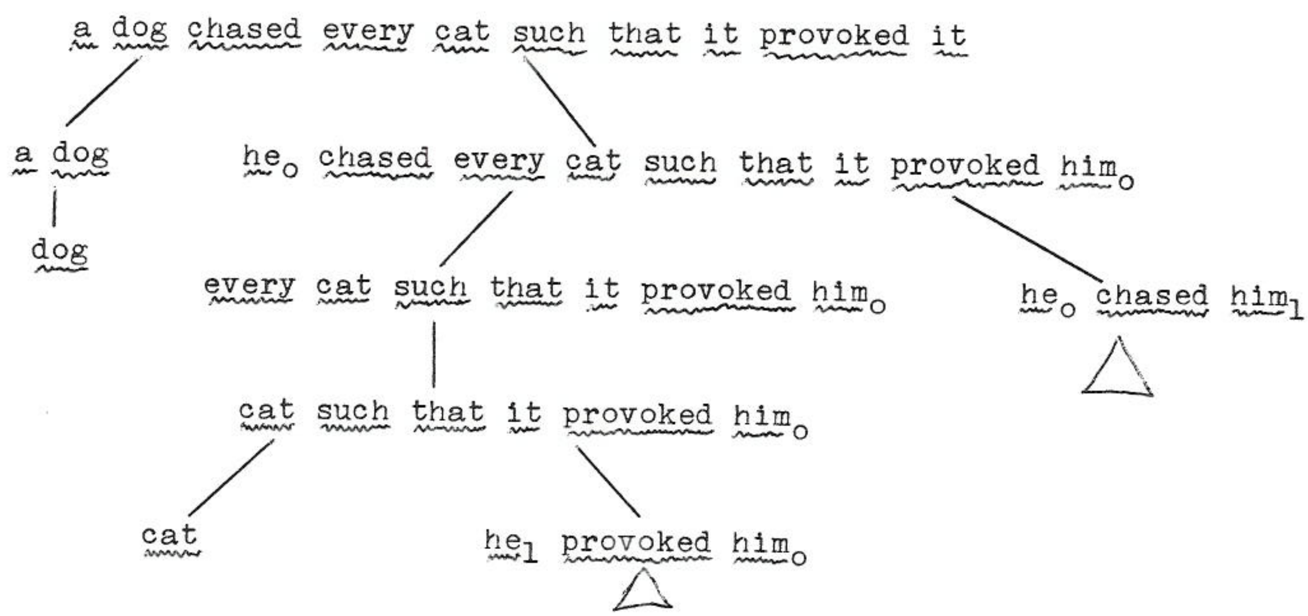

If we try to generate the sentence "the other way" it can't be done. Briefly, we can't have:

every cat such that it provoked it

hanging off the top node, because there is no way to generate this phrase all by itself that will have the second *it* be anaphorically related to *a dog*. And if we have:

every cat such that it provoked him$_0$

hanging off the top node there will be no way to turn the *him$_0$* into *it*.

Unfortunately, there are problems. Consider the sentence:

every man who owns a donkey that he likes beats it.

This sentence cannot be generated at all in PTQ; and the mechanism that prevents it from being generable is the same as the mechanism that so nicely disambiguates the sentence cited earlier. The first problem that I want to solve is: how can we include sentences of the sort just given, without giving up the advantages that the PTQ system already has?

What is to be done? Well, if you ask someone what the *it* means in the "ungenerable" sentence cited above there is a ready answer: it means *the donkey that he owns that he likes*. This paraphrase of the pronoun eliminates one pronoun at the expense of introducing two more, but it turns out that these do not cause problems of the sort raised by the *it*.

The idea I want to investigate is whether the problematic pronouns might be short for paraphrases of an appropriate sort. The best way I know of to investigate such an idea is to formulate some specific rules that embody the idea, and see what they yield. Specifically, since this is a generative system, the rules would produce the pronouns from their paraphrases. For example, given a derivation of the sentence:

every man who owns a donkey that he likes beats the donkey that he likes that he owns

(along with a symbolization of the sentence) the rules would generate the sentence at the top of this page (along with the same symbolization). That is the task of the next section.

III THE PROPOSAL

I begin by adding to PTQ a device for indicating and controlling relations between pronouns and their antecedants. The device employs indexes, which are used in closed* analysis trees as follows:

(1) Every basic noun (CN or NP), excluding pronouns, gets a non-negative integer as its index when it enters the tree; this index appears as a superscript on the noun.

(2a) Derived CN's (i.e. CN's modified by relative clauses) receive the index of their head CN, and (2b) when a derived CN is formed by the operation $F_{3,n}$ the n must be the same as the index.

(3a) Derived NP's (i.e. of the form Det+CN) receive the index of their constituent CN, and (3b) when the quantification operation $F_{10,n}$ is used with an NP the n must be the same as the index of the NP.

(4) Whenever a subscripted pronoun loses its subscript it gains an index which is the same as the lost subscript.

An example of an indexed analysis tree is the following (the numbers in circles indicate applications of the conditions 1-4 above):

---

*A <u>closed</u> analysis tree is defined to be an analysis tree whose top node is an S and which contains no subscripted pronouns in its top node.

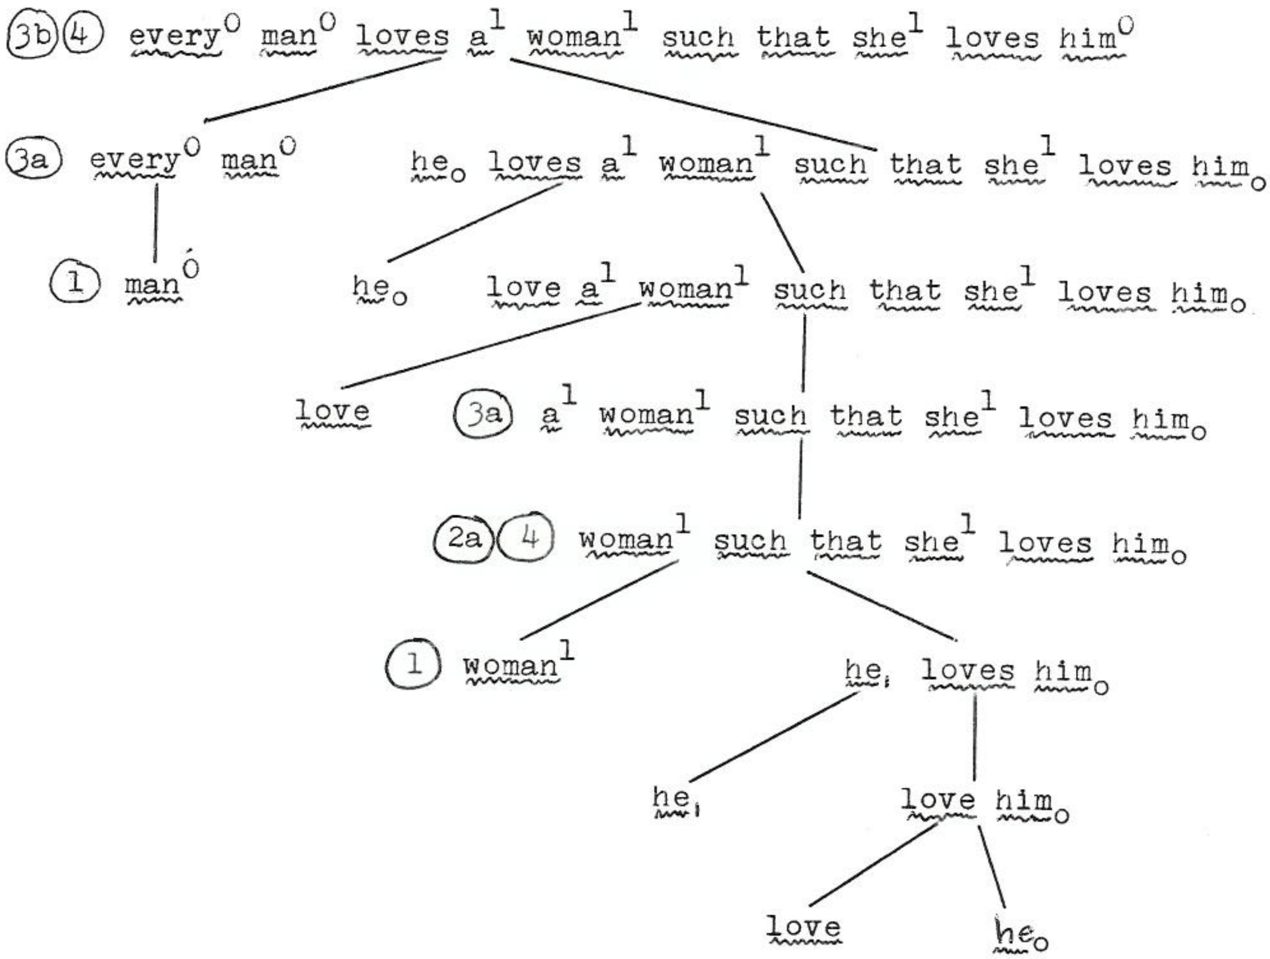

The rule for generating pronouns from paraphrases now is:

If: (1) there is an NP-subtree (the ANTECEDANT) of the form: $\underset{\triangle}{\gamma^k \alpha^k}$ where γ = every or the or a,

and: (2) the smallest S-subtree containing the antecedant subtree is of the form $\beta_1 \underset{\triangle}{\gamma^k \alpha^k} \beta_2$ (where β_1, β_2 may be null)

and: (3) there are one or more NP-subtrees (the SOURCES) of the form:

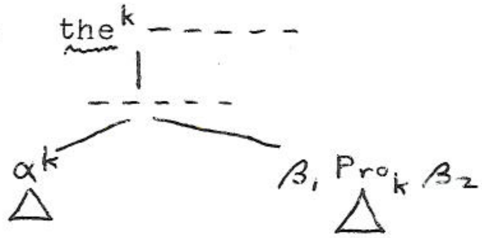

which do not both precede and command the antecedant, then: (4) each source subtree may be replaced by Pro^k (where Pro is chosen to be the same gender as α), and likewise all higher occurences of strings that are exactly like *the*^{k} ----- , or which differ from it only by containing Pro^j instead of Pro_j for some j.

I will say that a tree formed by zero, one, or more applications of this rule is well-formed if it does not contain any two distinct occurences of NP's which are not pronouns and which have the same index. In any well-formed tree the antecedant of any indexed pronoun is defined to be the NP which is not a pronoun and which has the same index as the pronoun.

At this point we badly need an illustration. I will begin by showing how:

every man who owns a donkey beats it

can be derived from a tree whose semantics is that of:

every man who owns a donkey beats the donkey he owns.

On the next page I give the source tree. The antecedant subtree, the smallest S-subtree containing it, and the source subtree are all circled. The phrase that will be turned into a pronoun has been enclosed in square brackets. In this tree, β_1 is *he*$_0$ *owns*, and β_2 is null.

EXAMPLE # 1

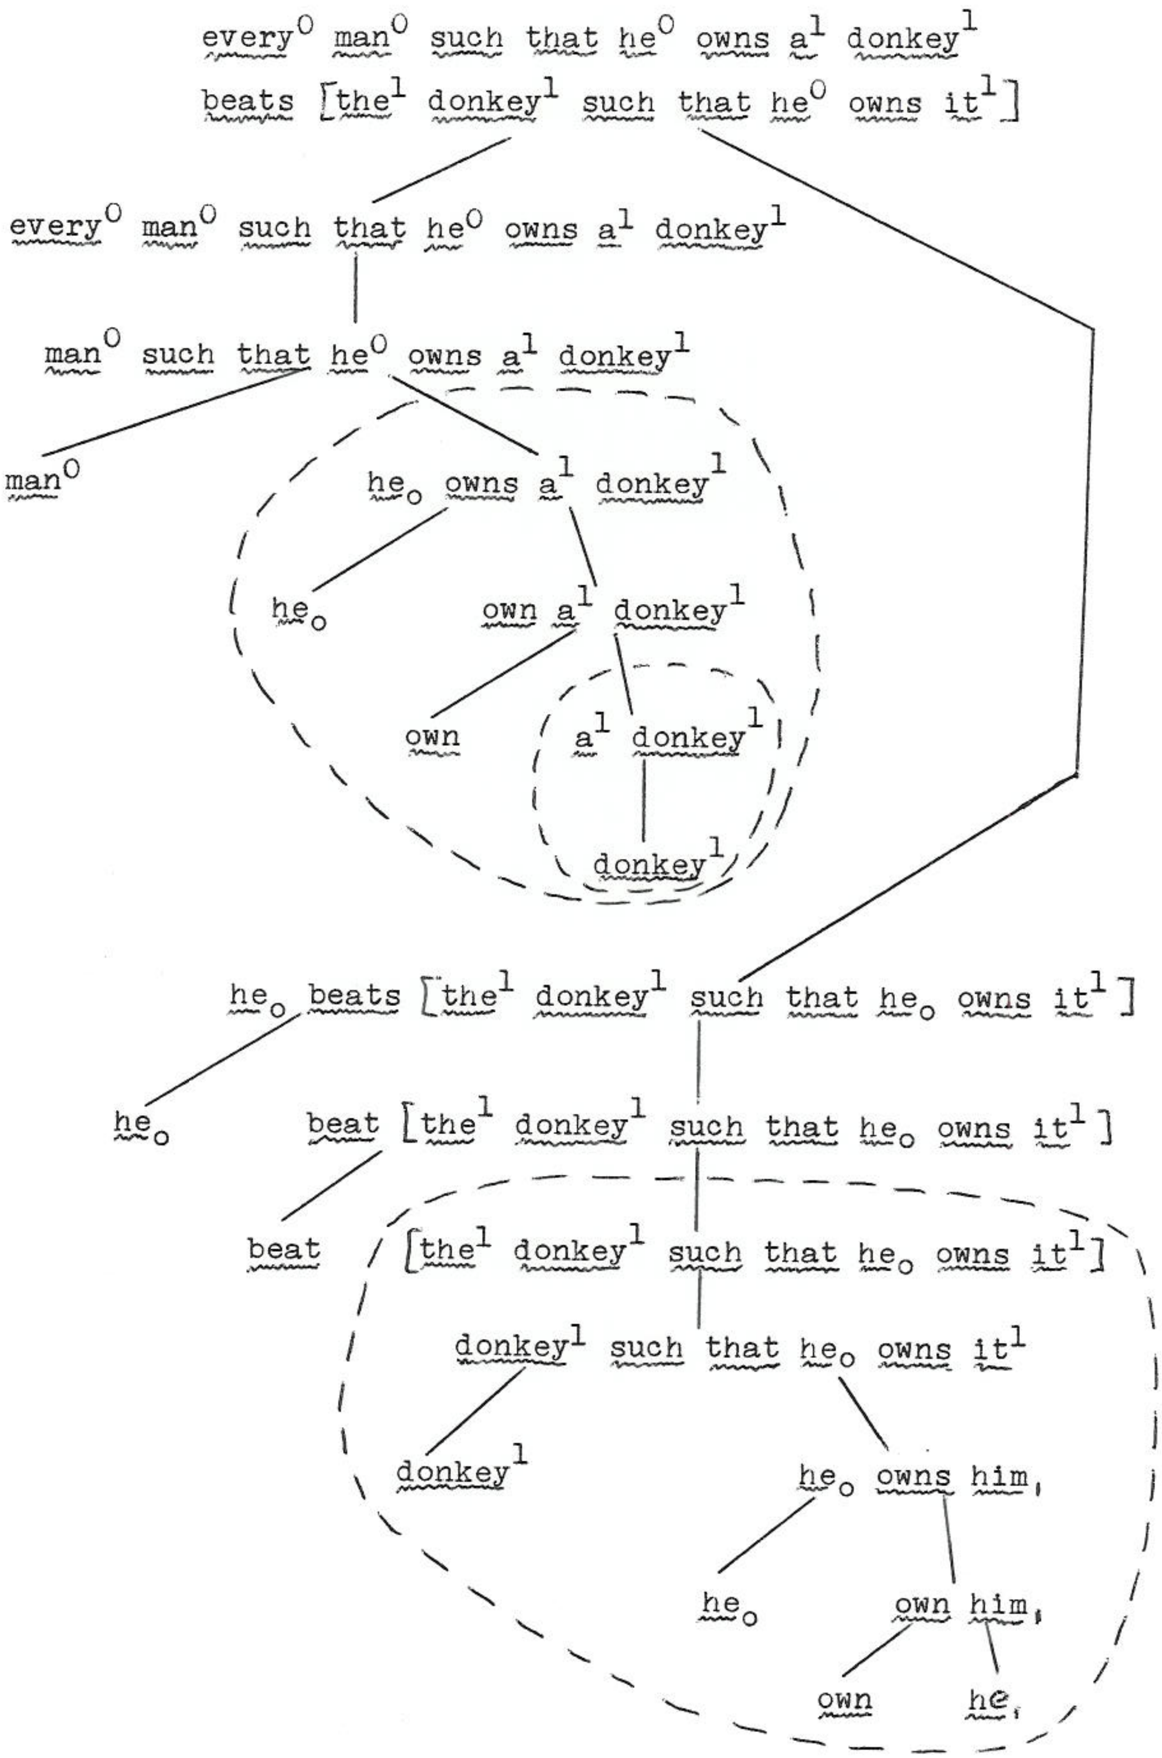

every^0 man^0 such that he^0 owns a^1 donkey^1 beats [the^1 donkey^1 such that he^0 owns it^1]
every^0 man^0 such that he^0 owns a^1 donkey^1
man^0 such that he^0 owns a^1 donkey^1
man^0
he$_0$ owns a^1 donkey^1
he$_0$
own a^1 donkey^1
own
a^1 donkey^1
donkey^1
he$_0$ beats [the^1 donkey^1 such that he$_0$ owns it^1]
he$_0$
beat [the^1 donkey^1 such that he$_0$ owns it^1]
beat
[the^1 donkey^1 such that he$_0$ owns it^1]
donkey^1 such that he$_0$ owns it^1
donkey^1
he$_0$ owns him$_1$
he$_0$
own him$_1$
own
he$_1$

Conditions (1)-(3) apply, and the rule converts the tree into the well-formed:

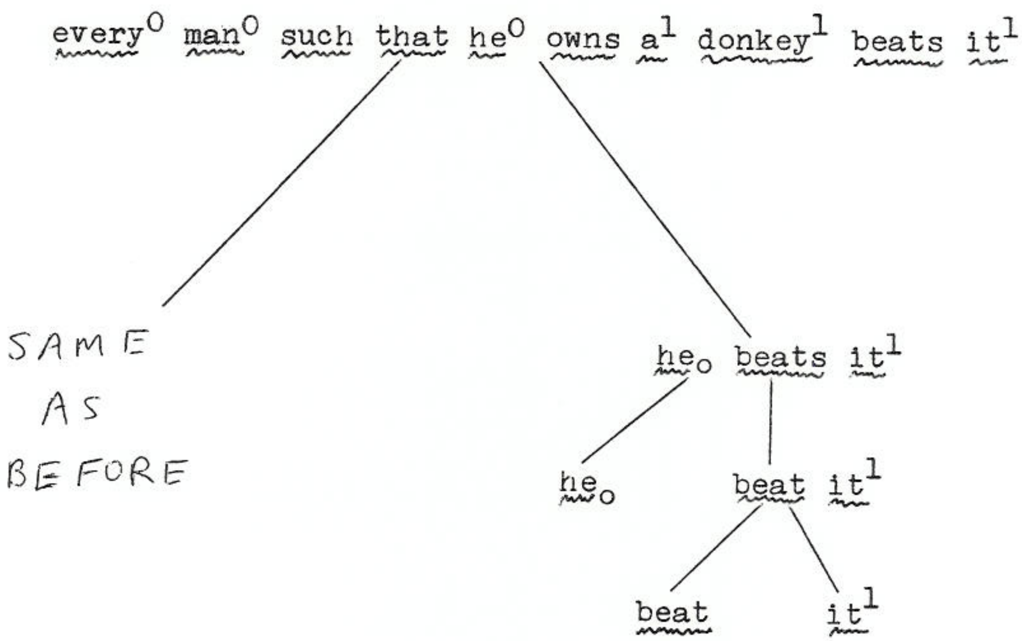

I will now state some further consequences of the rule, but without displaying the actual trees; they will be found in the appendix. First, the rule easily generates:

every man who owns a donkey that he likes beats it

from:

every man who owns a donkey that he likes beats the donkey that he likes that he owns. (EXAMPLE 2)

But a problem arises in trying to generate:

every man who marries a woman who owns a donkey that he likes beats it.

It turns out that this could be generated from:

every man who marries a woman who owns a donkey that he likes beats the donkey that he likes that he owns (EXAMPLE 3)

if we allowed ourselves to start with non-closed analysis trees, but this would get the meaning wrong, which is why we disallowed such trees at the outset. The trick is to get the pronoun *she*^{1} in the place where the *he*$_{1}$ is above, and this can be done by using our special rule to generate the *she*.^{1} But then we need to be able to treat this *she*^{1} on a par with other anaphoric pronouns for purposes of *re*applying our rule to get the *it*. And this can be done by relaxing the input conditions for our rule as follows: the β_1 and β_2 in the top node of the S-subtree must be just like those in the source subtree *except* that the latter may contain Pro^{j} in one or more places in which the former contains Pro_{j}. This relaxation allows us to generate the sentence we want from:

> *every man who marries a woman who owns a donkey that he likes beats the donkey that he likes that [the woman who owns a donkey that he likes whom he marries] owns.* (EXAMPLE 4)

(Try reading this the first time through putting *she* in place of the bracketed NP; the next time it'll be easier to process). We can now also generate a sentence that is superficially the same, but which has (in my experience) resisted treatment by other means; we get:

> *every man who marries a woman who owns a donkey that she likes beats it*

from:

> *every man who marries a woman who owns a donkey that she likes beats the donkey that [the woman who owns a donkey that she likes whom he marries] likes that [the woman who owns a donkey that she likes whom he marries] owns.* (EXAMPLE 5)

(If this proposal for treating pronouns were both correct and psychologically real it would be apparent by now just *why* we have pronouns in English).

Now consider the sentence:

every woman who doubts that a man respects her resents him

The relative clause of this sentence is potentially ambiguous; it can mean either:

woman who doubts that any man respects her

or:

woman who doubts that a certain man respects her.

In the given sentence the anaphoric *him* at the end of the sentence resolves this ambiguity in favor of the latter reading. And the pronoun rule handles this reading nicely, generating the sentence from:

every woman who doubts that a man respects her resents the man such that she doubts that he respects her. (EXAMPLE 6A)

Unfortunately the rule also produces a reading corresponding to the incorrect construal of the relative clause; it produces the sentence from:

every woman who doubts that a man respects her resents the man who respects her. (EXAMPLE 6B)

The moral: this rule, like so many others, overgenerates. A defect to be kept in mind.

Lastly, the rule *can* be made to apply to Bach-Peters sentences. These sentences violate the general rule that no pronoun may both precede and command its antecedant. Most sentences which violate this rule aren't good English, but certain of them seem fairly natural. In fact if we relax this rule in condition (3) then we can generate, e.g.:

a pilot who saw it downed a MIG that chased him

from:

a pilot who saw the MIG that chased him that he downed downed a MIG that chased him. (EXAMPLE 7)

IV A SPECULATION ABOUT DISCOURSE

If there is any hope for the proposal sketched above then it might be extended to include discourse anaphora as follows. First, we define a *discourse tree* to be a structure of the form:

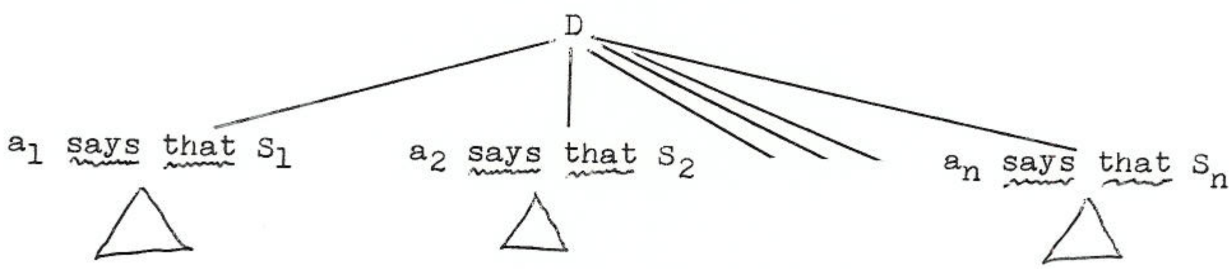

where each subtree below the D is a closed analysis tree of PTQ, and where the a_i's need not be distinct. The idea is that such a discourse tree is realized phonetically by some speakers (representing the a_i's) pronouncing the S_i's in sequence. (Actually this would require some adjustments for tense, for indexicals such as "I", etc.). Now we allow the pronoun generating rule to apply when the antecedant is in a <u>previous</u> tree. Then we could generate, for example:

a_1: "A man died."

a_2: "He left a will"

from:

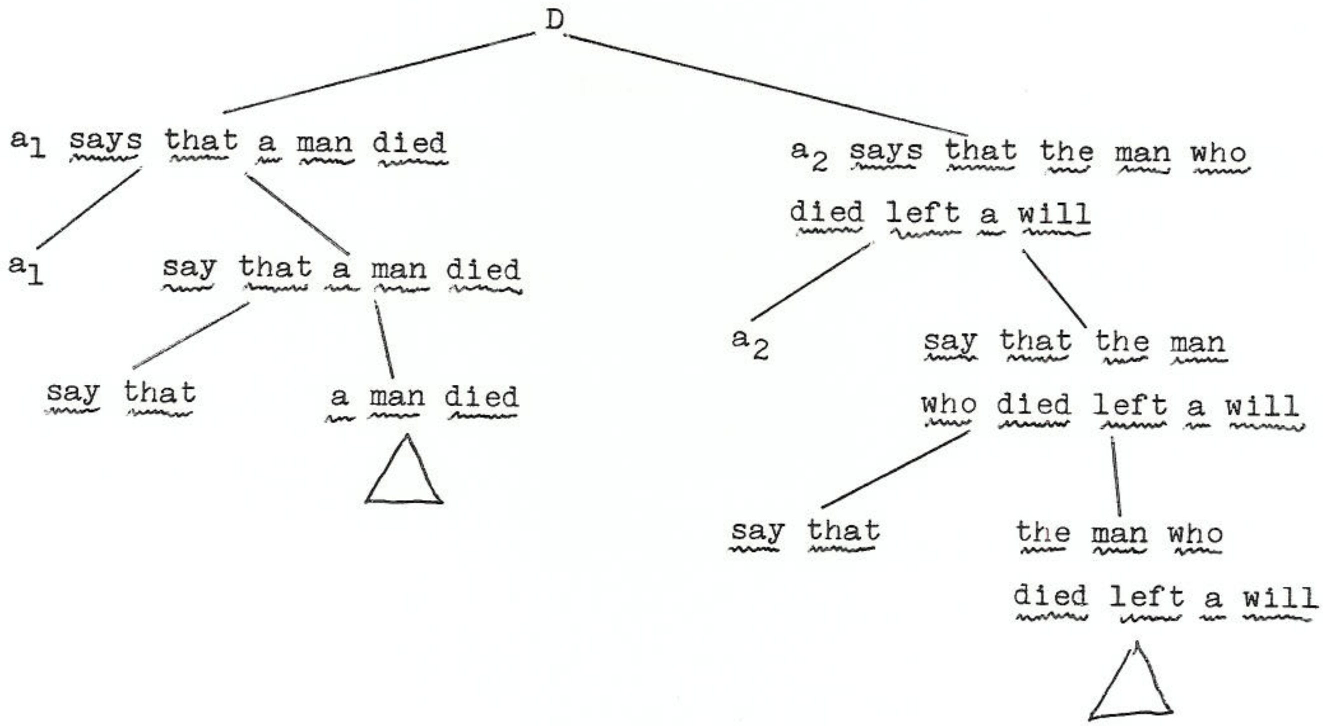

Similarly, we could get:

a_1: "A man jumped"

a_2: "He didn't jump (he was pushed)"

from:

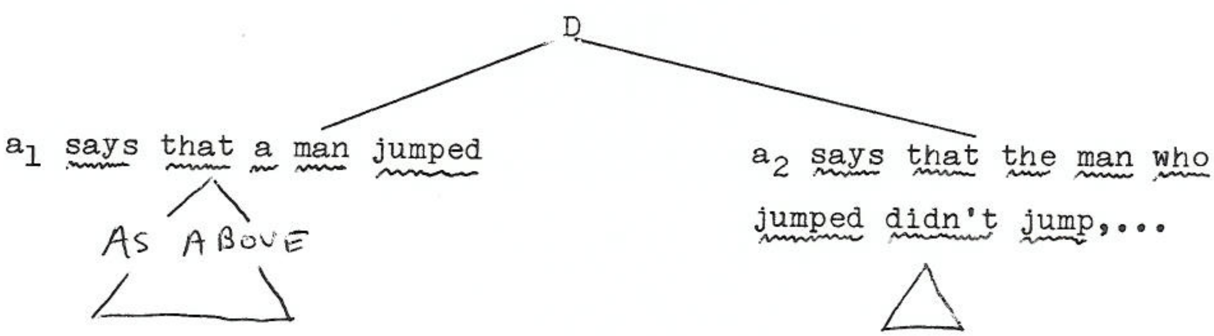

which would be pragmatically bizarre, and rejected as the right interpretation for that reason. Fortunately we could also get it from:

D

a_1 says that a man jumped

a man

a_1 says that he_o jumped

a_2 says that the man such that a_1 says that he jumped didn't jump,...

This requires that we interpret a_1 as saying of some particular man that he jumped, instead of just saying that some man or other jumped -- but this seems exactly right.

The technique could just as easily be expanded to include non-indicatives as well. E.g. if we allow "asks whether" instead of "says that" we can generate:

a_1: "A man jumped"

a_2: "Did he (really) jump?"

V WHAT'S WRONG WITH THE PROPOSAL?

First two fairly clear criticisms:

For one, I have only examined a few cases. I could easily have overlooked examples similar to the ones treated, but where the proposal works incorrectly or does not work at all. (One possible example of this is one of the "paycheck" sentences: "The man who gives his paycheck to his wife is wiser than the man who gives it to his mistress." The proposal does not generate "it" with the intended reading. This doesn't worry me too much because I think the "it" here is colloquial for "his", and this is an instance of a different phenomenon.)

Second, the framework of PTQ is in certain ways quite restricted, and so there is a vast array of cases involving pronouns that are not treated. For example, the fragment does not contain any plurals and does not contain any determiners other than *every*, *a* and *the*. So there is no treatment of:

most men who own donkeys beat them.

Such cases require further work.

But the central problem that I see is that perhaps each and every one of the examples treated in this paper is just plain wrong. Recall that the central goal of PTQ is semantic; the goal is to pair English sentences with their possible meanings. The proposal I have sketched does perhaps generate sentences containing pronouns from other sentences which do, in English, on the appropriate readings, mean the same as the sentences containing the pronouns. But it's not clear that the meanings assigned to these source sentences by PTQ lets them mean the same. This is because of the treatment of the word *the* in PTQ, which is patterned after that of Russell in "On Denoting".

Let's look at an example. I generated:

every man who owns a donkey beats it

from:

every man who owns a donkey beats the donkey he owns.

In PTQ this latter sentence means:

every man who owns a donkey owns exactly one donkey, and beats that donkey.

This doesn't sound right at all (all the other examples are peculiar in the same way).

I can think of two possible responses to this objection, but before giving them let me fill in a little background.

Suppose we ask what the right answer <u>should</u> be as to how to treat the above sentence. If we do this we will soon discover something pointed out by Barbara Partee: the data is <u>very</u> unclear. The unclarity has mostly to do with what we say about cases in which various men own more than one donkey. In a relatively untrustworthy experiment the following was "discovered":

If every man beats every donkey he owns, the sentence is true. If some man owns exactly one donkey and fails to beat it, the sentence is false.

In practically all other cases, speakers are either unable to make judgments about truth and falsity, or their judgments disagree with one another.

The issue can be made poignant by considering a different sentence with a closely related structure:

every man who has a son wills him all his money

(due to Partee). What does this sentence say about men who have more than one son? It somehow seems inappropriate to apply the sentence to them, yet it seems literally to be discussing them; it says "every man who has a son...", and any man with two or more sons is a man who *has* a son.

THE FIRST RESPONSE: One might suggest that the feeling of inappropriateness comes explicitely from the use of the pronoun. How would that work? Well, one purported meaning of "a" is "one", in the sense of "exactly one". For example, if I say "I just bought a car" or "I got a bicycle for Christmas" we would normally take the "a" to have the import of "exactly one". Usually this is thought to be a presupposition, implication, or implicature of the utterance rather than part of the content of what is said. But perhaps the use of a singular pronoun can make the import part of the official content.

The suggestion then is that "a" can mean either "at least one" or "exactly one". Normally it means the former, but certain grammatical constructions force the latter reading. The former reading is the "indefinite" one, and the latter is the "definite" one. And the rule for generating pronouns from paraphrases should be restricted so as to apply only in cases where the antecedant NP has its determiner marked "definite".

THE SECOND RESPONSE: Sometimes "the" doesn't mean "exactly one", but rather "at least one" or "every". It means "at least one" in:

everyone must pay the clerk five dollars

and it means "every" in:

you should always watch out for the other driver.

Or something like this.

So perhaps the treatment of pronouns as paraphrases is correct, but we have to tailor the meaning of "the" for the situation at hand. For example, in our sample sentence we need to read *the donkey he owns* as *every donkey he owns*.

This response would involve specifying some method for determining which reading of *the* is appropriate in a given paraphrase; I haven't carried this out.