Distributional Models of Meaning

Question 1 Oedipus Rex

We consider a sentence space S spanned by the set $\{T, F\}$ together with the usual map $not: S \to S$. (*)

Our noun space will be spanned by the noun "Thebes" (Th) together with the set of characters:

Oedipus (O)

Laius (L)

Jocasta (J)

As in any classical tragedy, we will need an oracle (no analogy with a quantum oracle intended...) which we will model as a verb taking a sentence as input:

(*) we also assume $and: S \otimes S \to S$

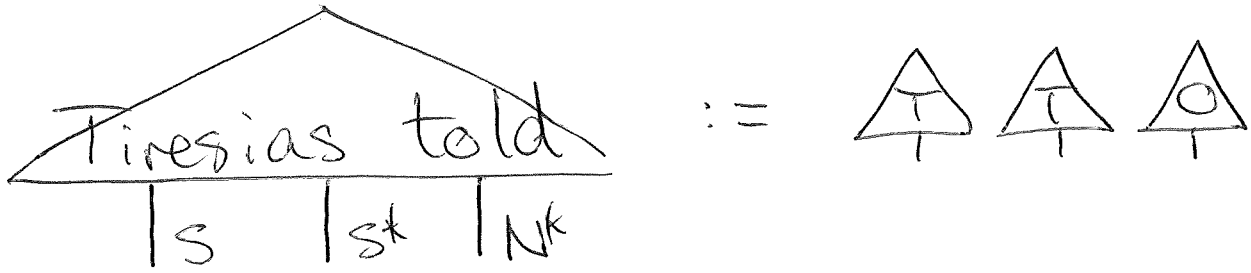

So that we can say the sentence:

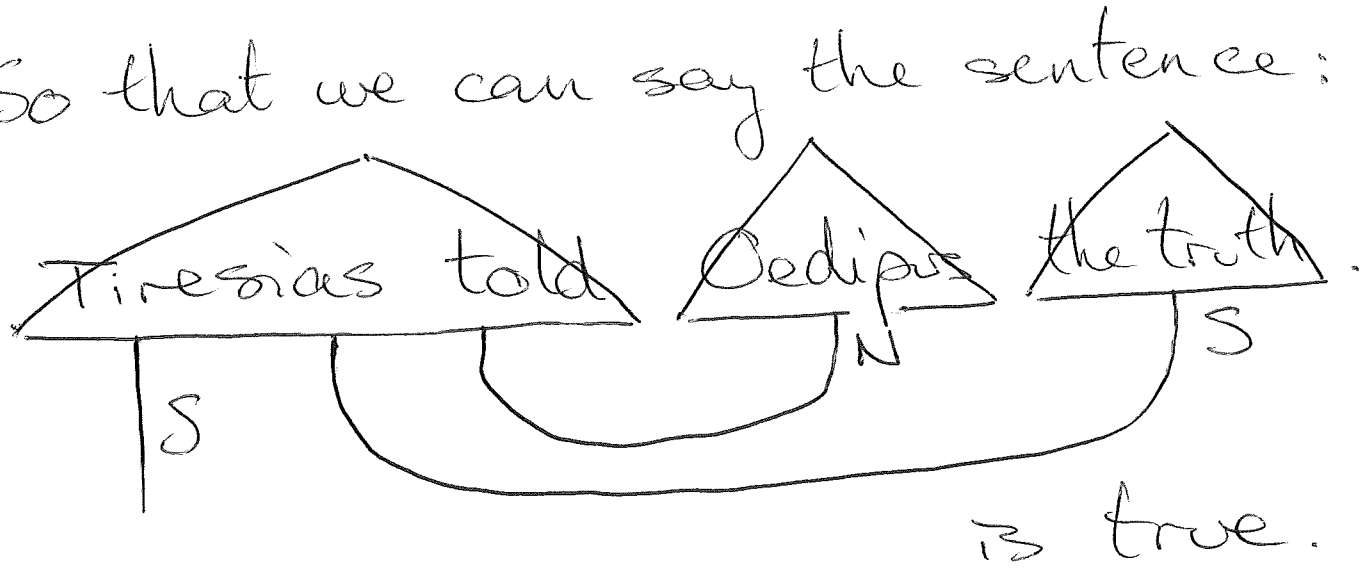

Now we will need the following intransitive verbs:

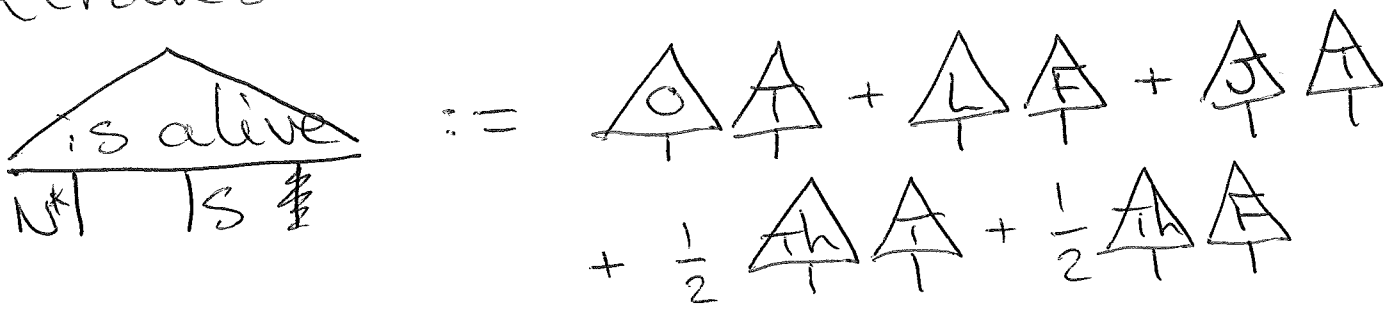

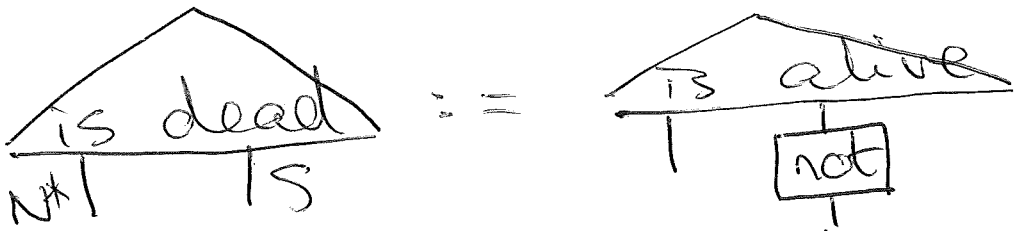

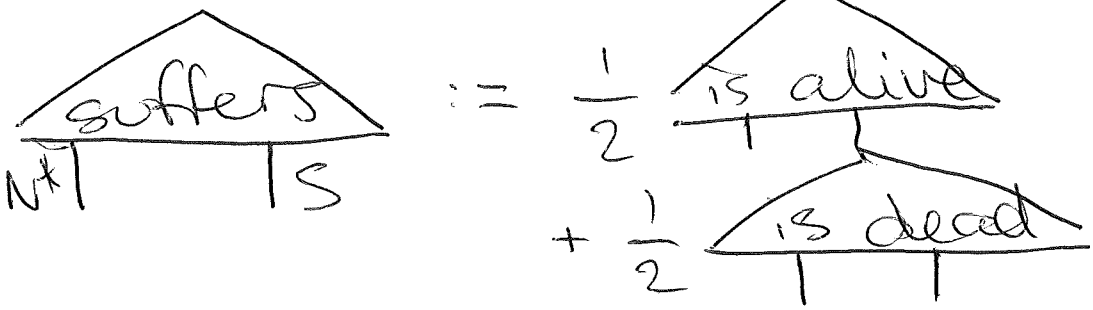

together with the following transitive:

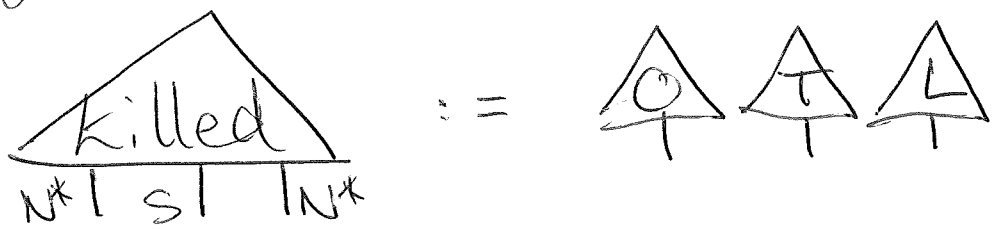

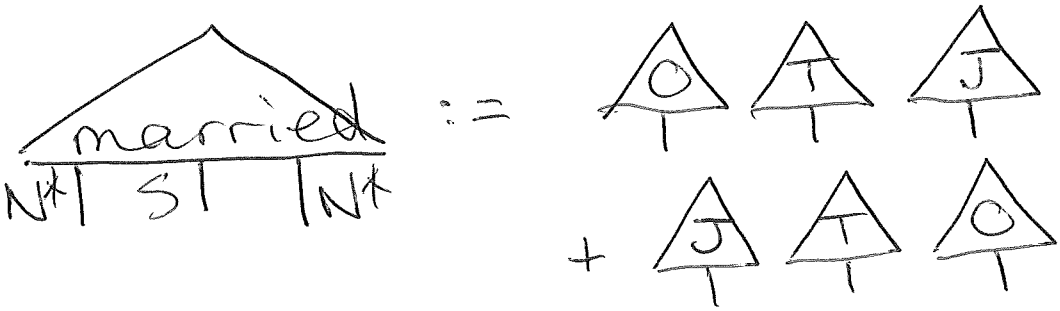

Furthermore, we will make use of the following nouns:

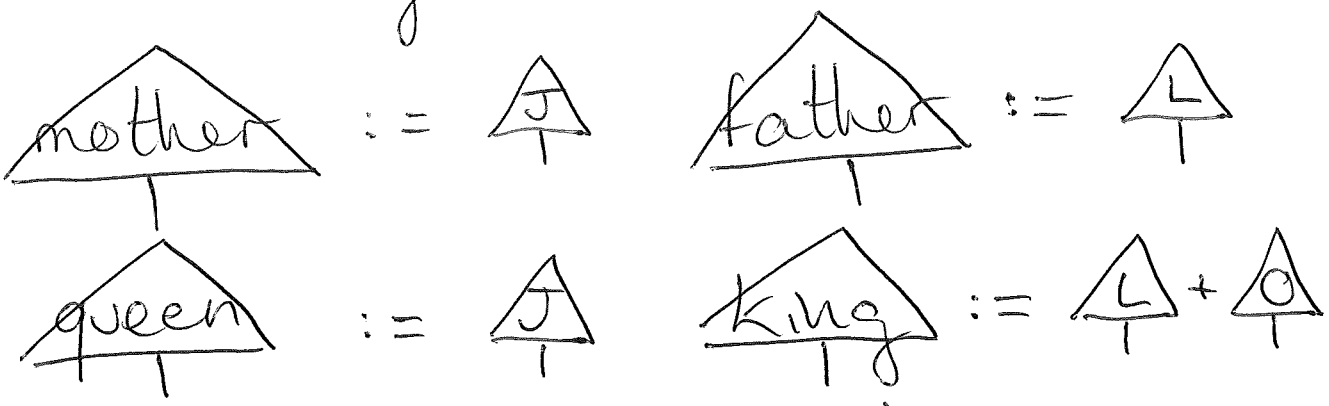

as well as two adjectives:

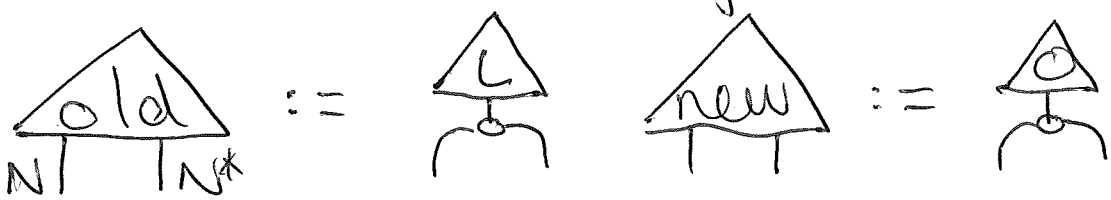

where o is the classical structure induced by the basis of our noun space.

Finally, we will encode basic facts about our story in the following tensor for the verb "is the":

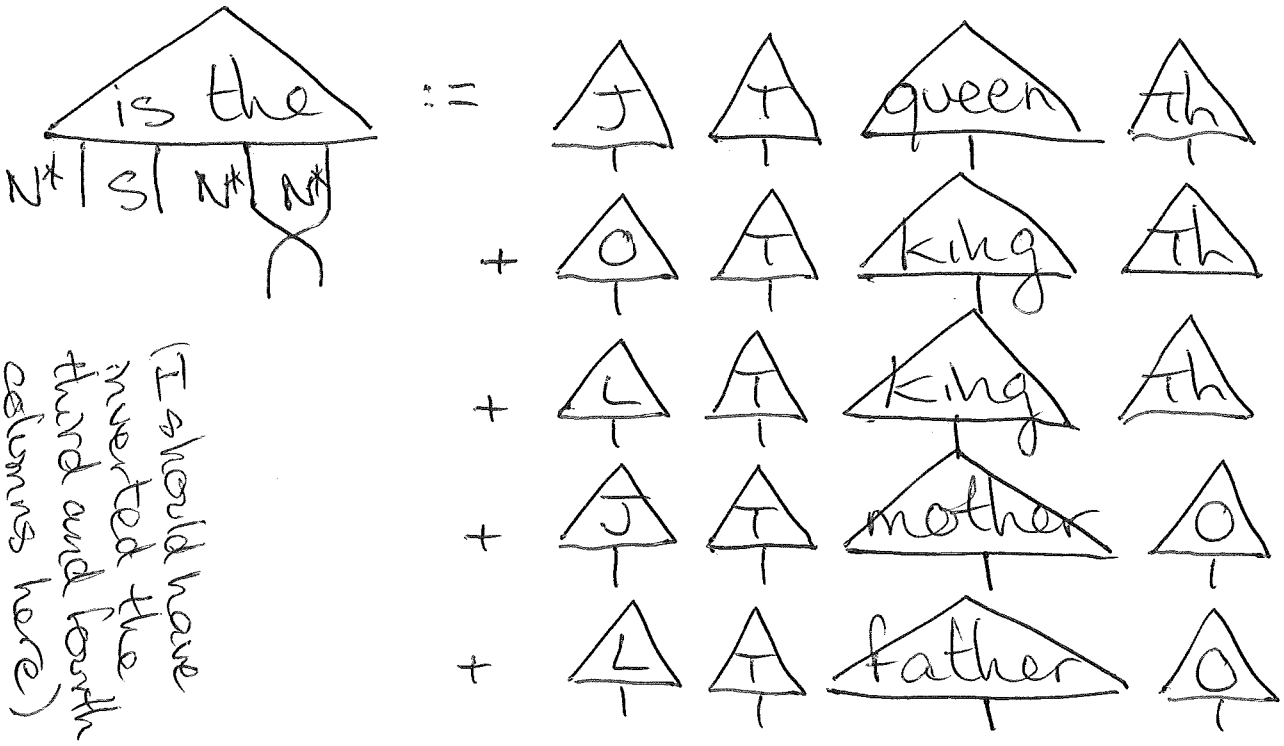

So that, for example, we can say:

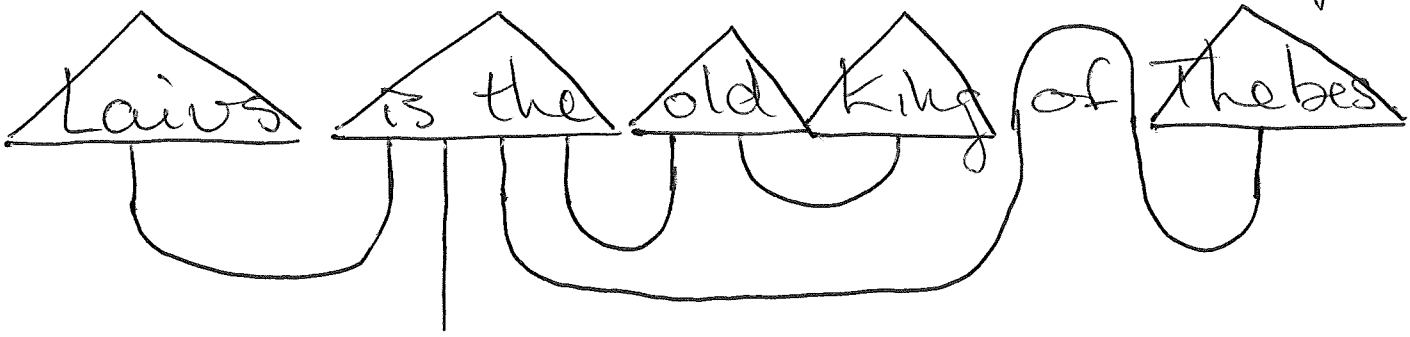

reduces to true. Indeed, we have:

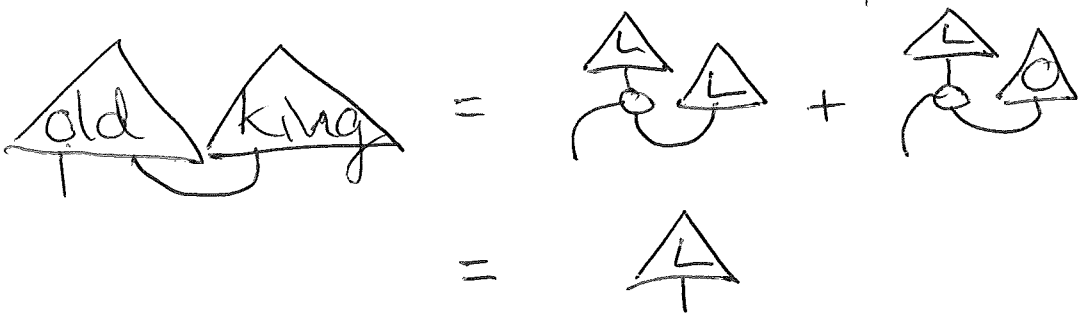

and

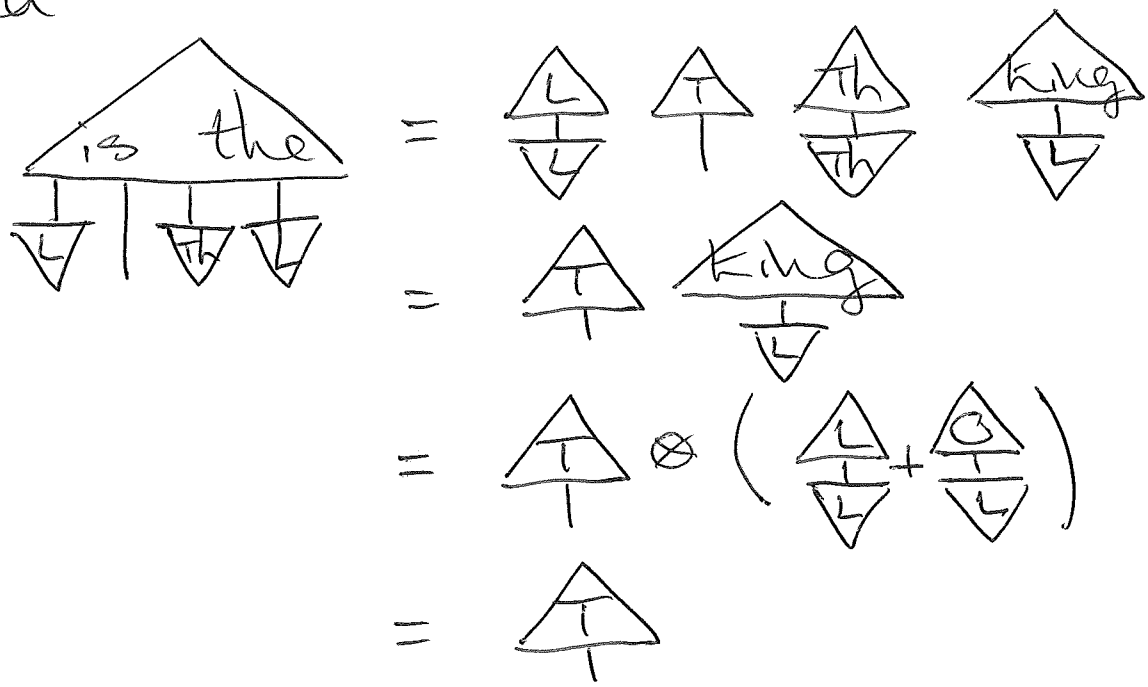

Note that in order to make the language of the story more natural, we introduce a mecanism for pronouns and reference as follows:

1) we introduce a new type $\overline{N}$ for noun references together with a map $\square : N \to N\overline{N}$ depicted as

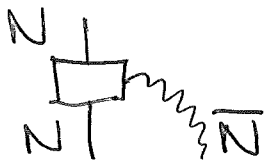

2) we define the pronoun "He":

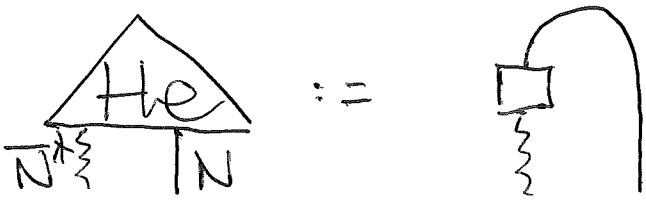

3) we interpret both the pronoun and the map □ as the spider o, using squiggly wires to distinguish a noun and its reference.

We are now in a position to ruin the suspense of our story (there can be no suspense in tragedy) with the first sentence of Oedipus Rex:

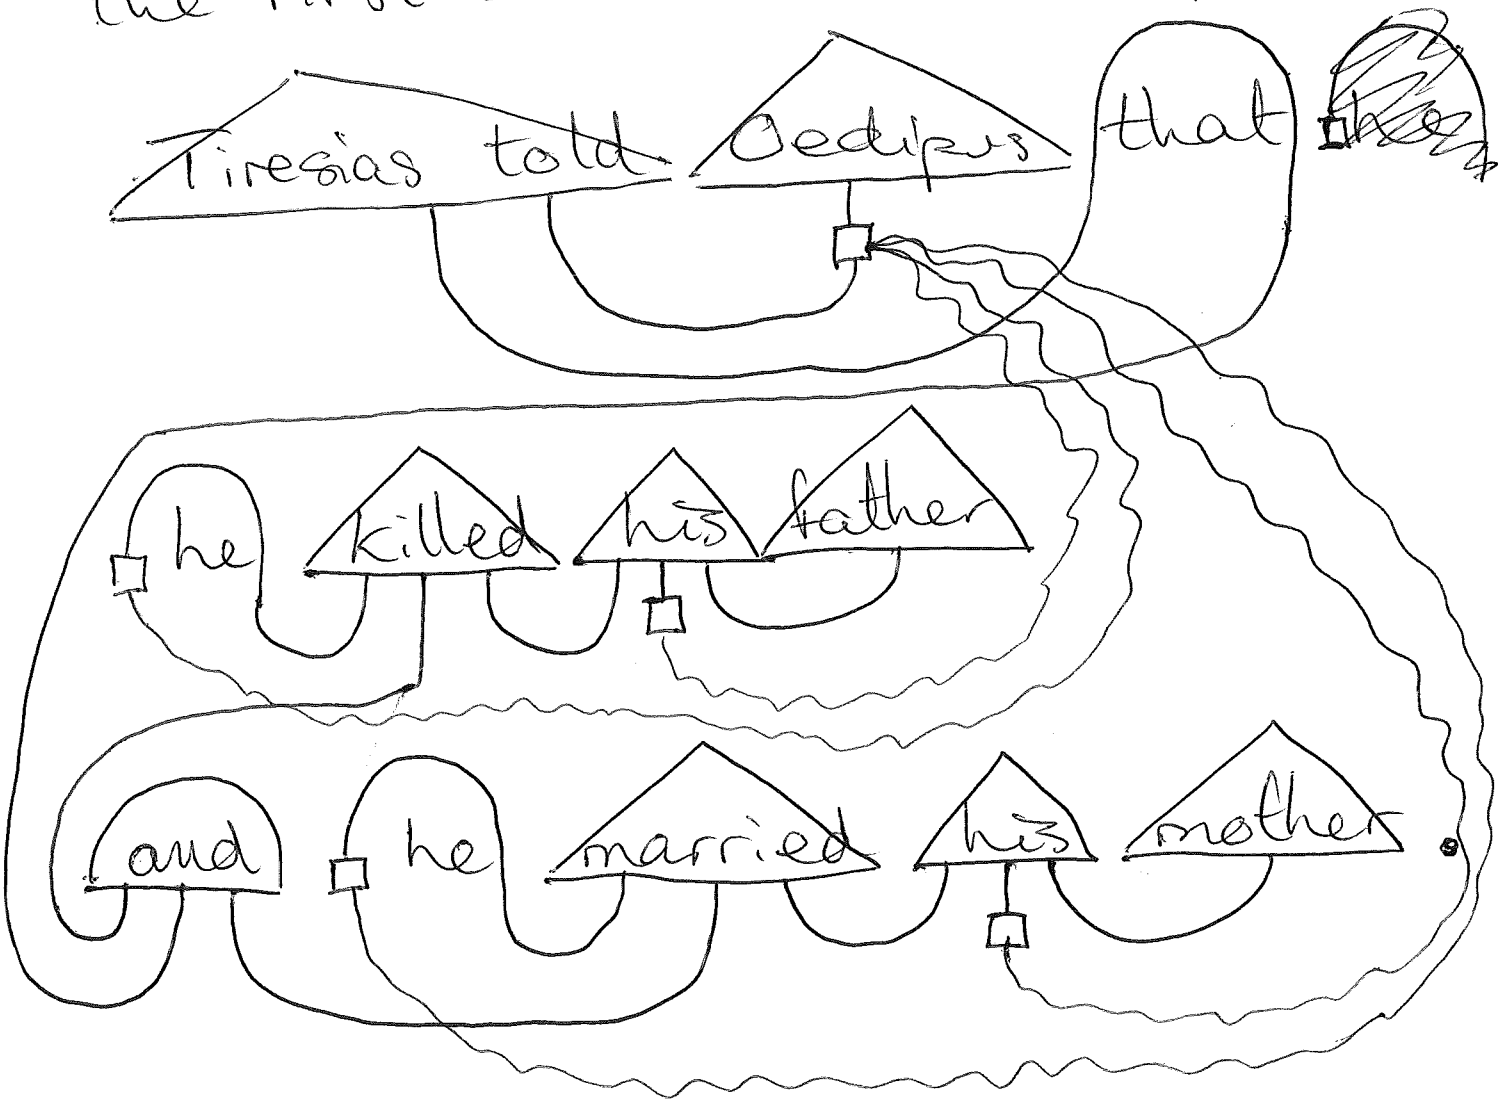

where we define the possessive pronoun:

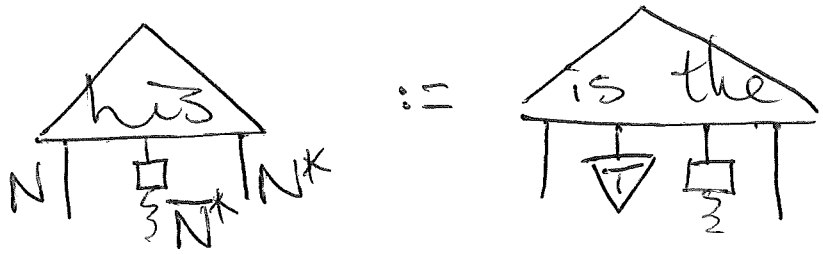

this can be seen more clearly with the help of an example:

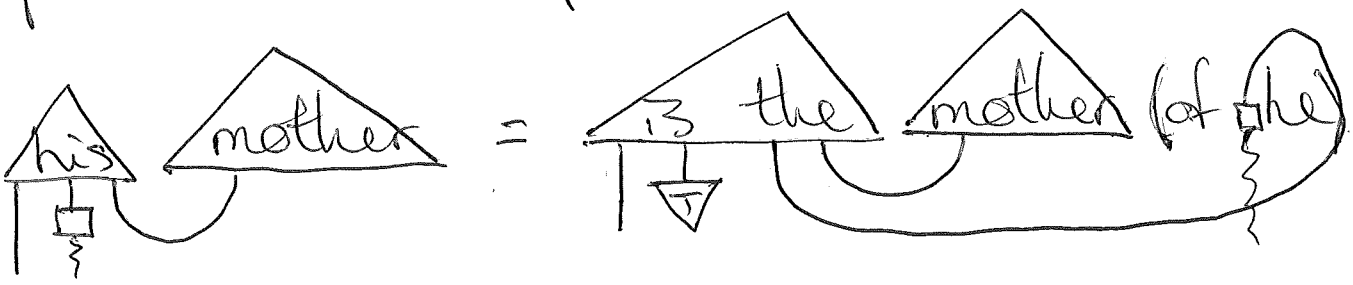

Using the fact that Oedipus is a copyable state of our classical structure, we can simplify the sentence by "resolving the references" of the pronouns:

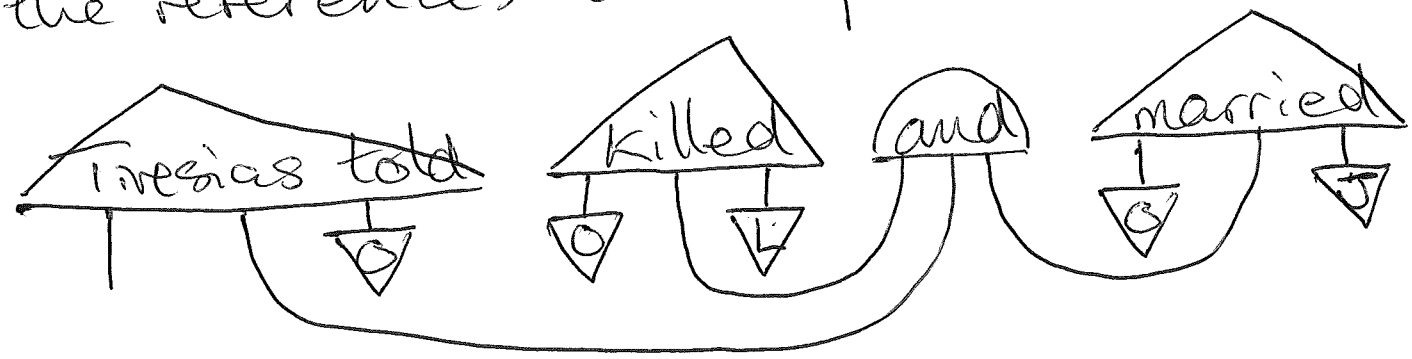

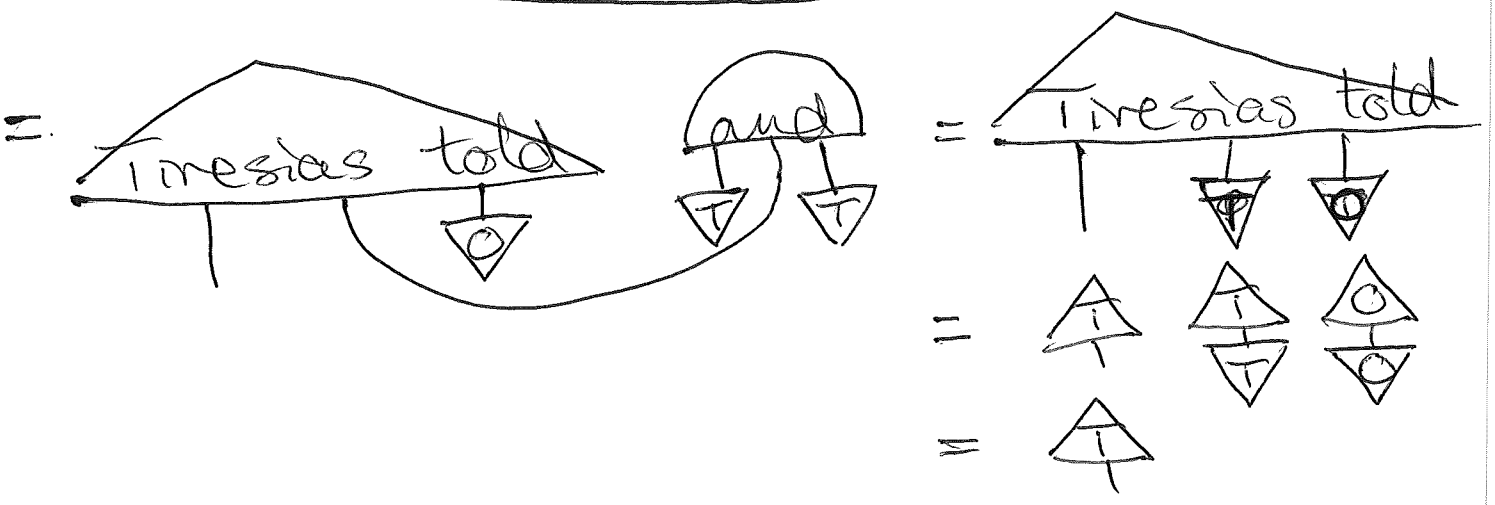

Now, using the inner product, we compare the meanings of the two sentences:

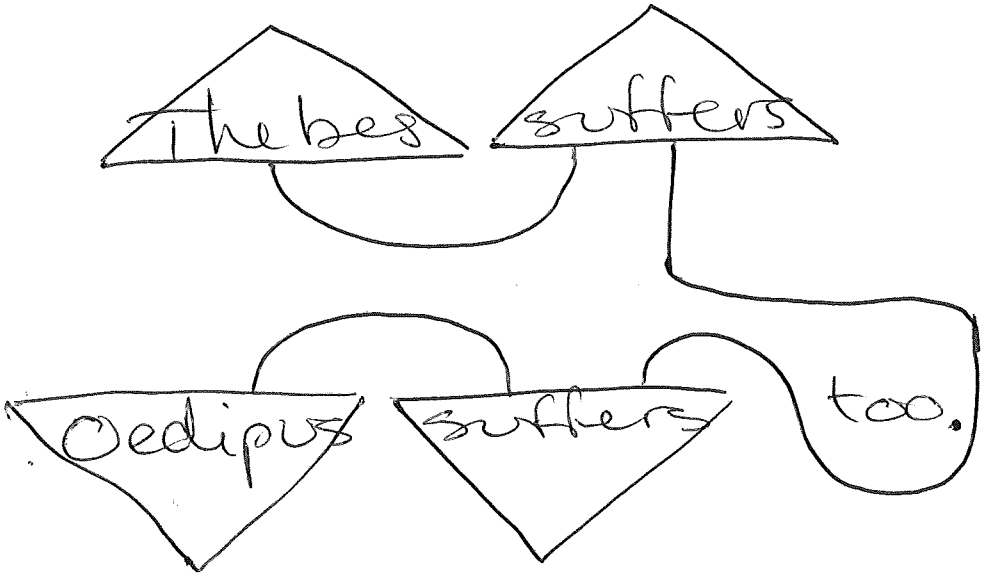

$=$

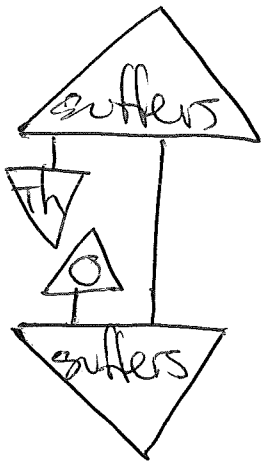

$=$

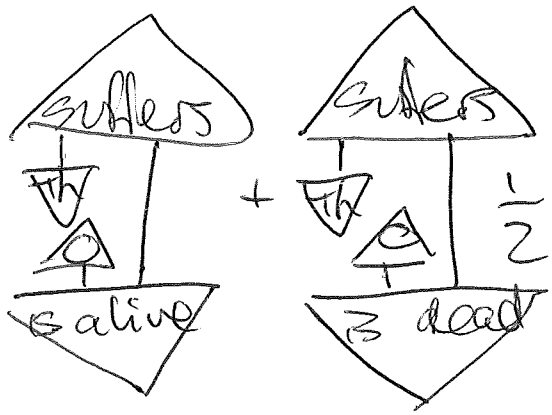

$=$

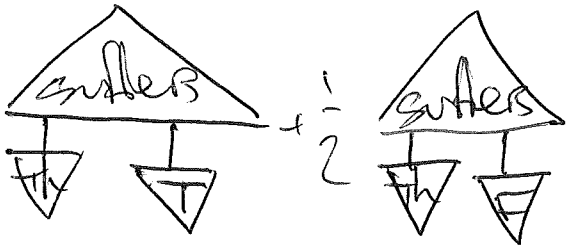

$$= \frac{1}{2}\left(\frac{1}{2}+\frac{1}{2}\right)+\frac{1}{2}\left(\frac{1}{2}+\frac{1}{2}\right)$$

$$= 1\ !$$

We can conclude that Oedipus and Thebes suffer as much. $\square$

As a conclusion, we can sum up our story as the reduction of a question to an answer that Oedipus is looking for:

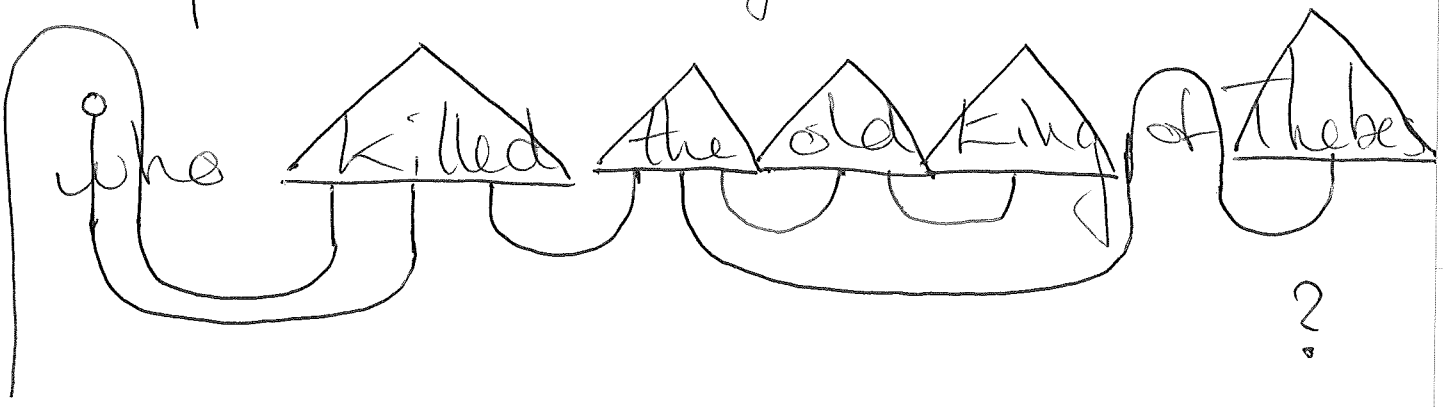

=

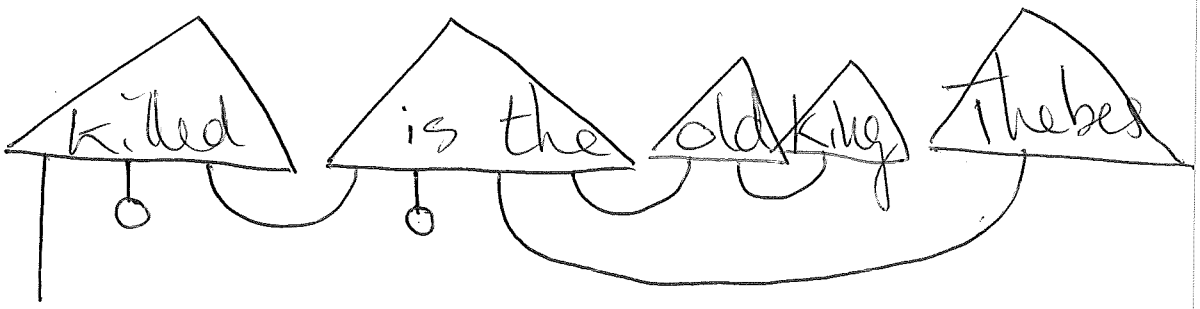

=

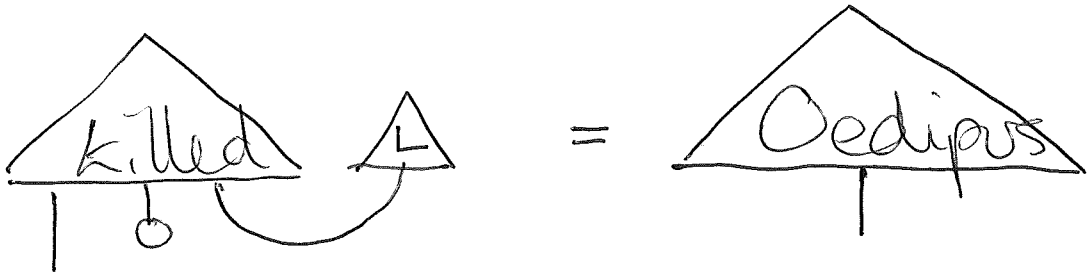

=

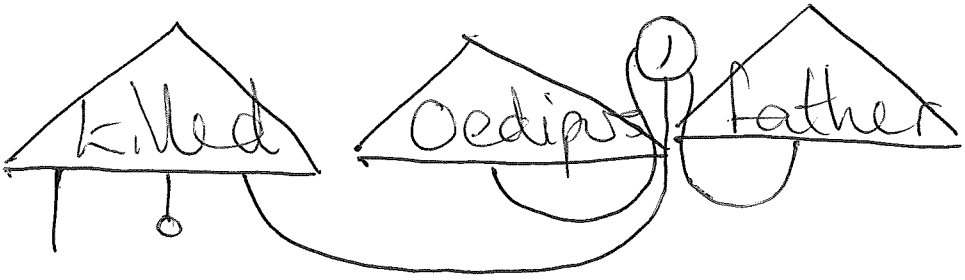

(where we can encode the possessive:

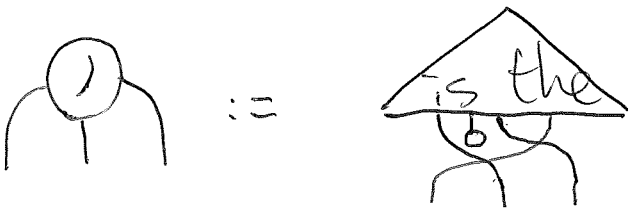

)

Question 2 Waiting for Godot

We will tell the story of Beckett's "Waiting for Godot" by generalizing our model of meaning along two dimensions:

- introduce a basic notion of time by defining meanings for the adverbs "yesterday", "today" and "tomorrow".

- generalize our notion of truth to model the meaning of "maybe".

In order to do this, we will construct a "Custom Hypergraph Category" in the sense of [Marsden and Genovese], defining a custom quantale internal to our time presheaf.
(We will restrict ourselves to the full subcategory of presheaves over finite sets.)

We define $\mathcal{T}$ as the functor category of presheaves over the following finite category:

$$\begin{tikzcd}
 & \bullet \arrow[dl, "\text{yesterday}"'] \arrow[d, "\text{today}"] \arrow[dr, "\text{tomorrow}"] & \\
\bullet & \bullet & \bullet
\end{tikzcd}$$

An object of $\mathcal{T}$ can be seen as a set T together with projection functions to the three time components.

As a presheaf category, $\mathcal{T}$ is naturally a topos and we can construct the category $\mathrm{Rel}_{\mathcal{T}}(Q)$ with the symmetric monoidal structure given by the finite products of $\mathcal{T}$, and the compact closure generated by the graph of the diagonal comonoids.

Q is the internal quantale generated by the following truth lattice:

False $\longrightarrow$ Maybe $\longrightarrow$ True

along each time component, where union and intersection are defined component wise, e.g.

$$(F, M, T) \vee (T, M, F) = (T, M, T)$$

A morphism $X \to Y$ in $\mathrm{Rel}_T(Q)$ will be an arrow $X \times Y \to Q$ in T, which we can construct explicitly as a matrix of triples (or a triple of matrices as will be convenient) taking values in true, false and maybe.

We pick Q to be our sentence space, and define an adverb to be a map $Q \to Q$ in $\mathrm{Rel}_T(Q)$. For simplicity, all the adverbs we will consider arise as the graphs of morphisms in T, using the functor defined in Prop. 5 of [Custom Hypergraph Categories]

Hence, it is sufficient to define a triple of functions acting on each time component. We first define "maybe" component-wise:

$$\text{maybe}(T) := F$$
$$\text{maybe}(M) := T$$
$$\text{maybe}(F) := F$$

as well as an extended "not" map:

$$\text{not}(T) := F$$
$$\text{not}(M) := M$$
$$\text{not}(F) := T$$

Note that we have $\text{maybe} = \text{maybe} \circ \text{not}$.

Then we define our three time adverbs as simply discarding the irrelevant time components (setting them to T):

$$\text{yesterday}(x, y, z) := (x, T, T)$$
$$\text{today}(x, y, z) := (T, y, T)$$
$$\text{tomorrow}(x, y, z) := (T, T, z)$$

So that for a sentence s, we have

that today(s) is true iff s is true in the today component, and maybe(not(today(s))) is true iff s is maybe in the today component.

So far we have defined the semantics of our adverbs, but this is only half of what a full-blown model of natural language should specify. Clearly, we want to be able to say that the phrase:

maybe not tomorrow

is grammatical, whereas:

today not maybe

is nonsensical.

In order to do this, we need to modify our pregroup to take adverbial phrases into account.

First, we will assume that time adverbs always come to the end of a sentence, and that a sentence can ~~only~~ contain at most one of them.

To model this, we introduce a new element S_T in our pregroup, together with the reduction $S \to S_T$ (which encodes the fact that a sentence can have no time adverb.)

Time adverbs have type $S^r S_T$ so that

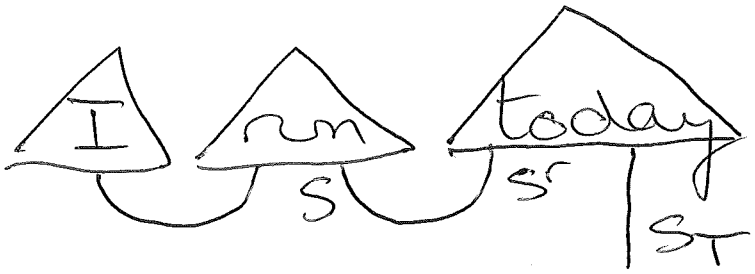

reduces to a well-typed sentence, but

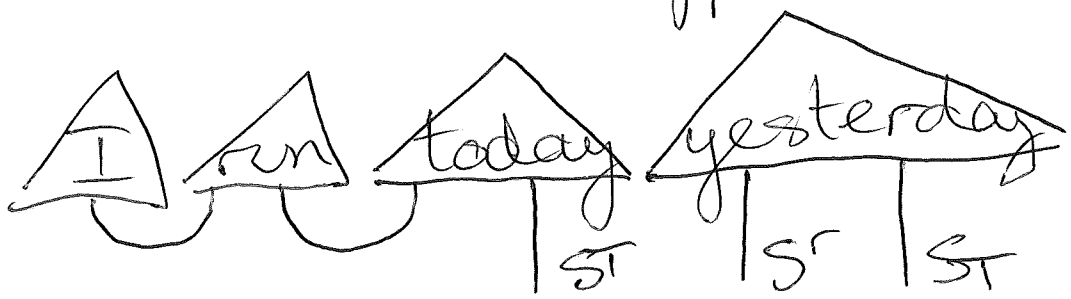

does not.

Now, the syntax of maybe and not is given by the basic rule that the only valid combinations are:

- no adverb (positive sentence)
- "not"
- "maybe"
- "maybe not"

We introduce the types S_+ and S_- for positive and negative sentences respectively, together with the reductions:

$$S_+ \longrightarrow S_- \longrightarrow S$$

so that "not" has type $S_+^r S_-$
and "maybe" has type $S_-^r S$

(The use of auxiliaries in English complicates this picture slightly, but this need not concern us here.)

To conclude the description of our new model of meaning, we follow the general recipe:

1) We generate freely a pregroup from the following poset:

$$\begin{array}{c} N \\ S_+ \to S_- \to S \to S_T \end{array}$$

which will serve as our "grammar category".

2) We construct $\mathrm{Rel}_p(Q)$ as our "semantics category".

3) Bingo! We can define the story of "Waiting for Godot" as a monoidal functor from grammar to semantics.

We define our noun space N to be the set of characters:

Vladimir (V)
Estragon (E)
Godot (G)

with identity functions for the three time components (characters do not change over time.)

We define the ~~in~~transitive verb:

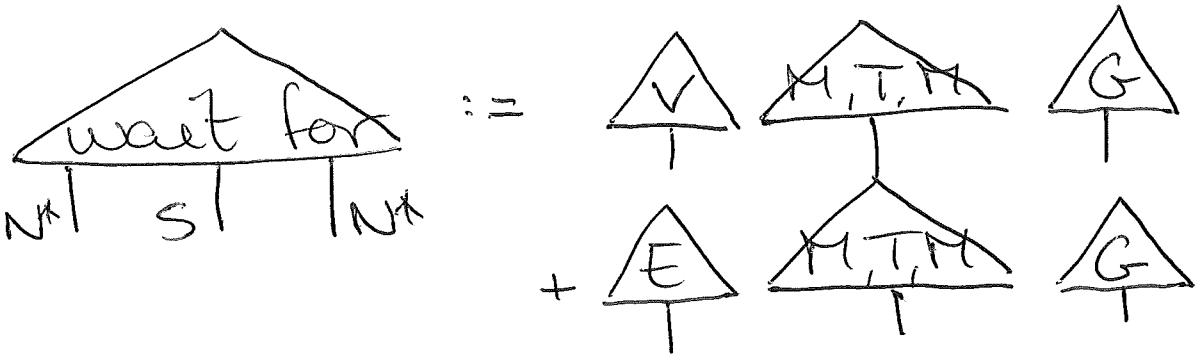

as well as the two intransitives:

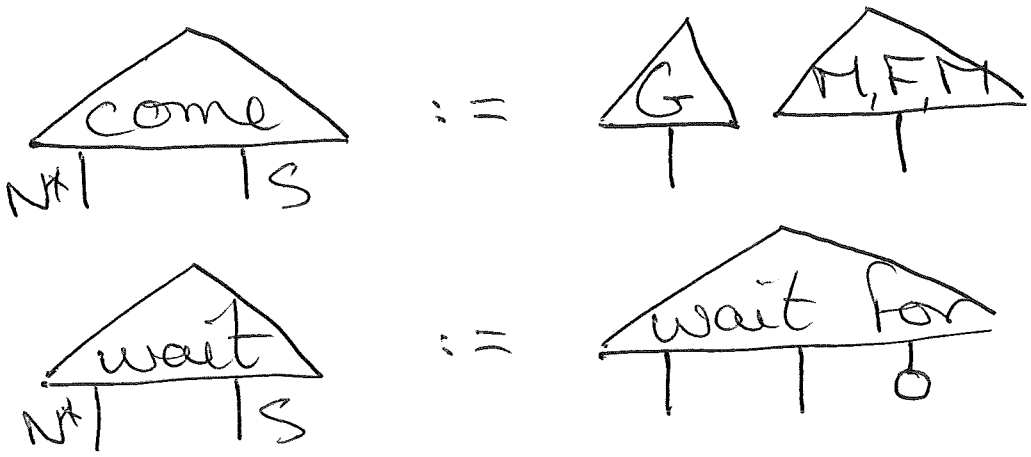

where o is the classical structure induced by our noun space, hence we can interpret "wait" as "wait for someone".

The first sentence of our story will be simple:

Vladimir and Estragon wait.

However, we will need to make use of subtle coordination, as introduced

in "Coordination in Categorical Compositional Semantics" [Kartsaklis 2016]

First, note that we have the following pregroup reduction:

$$N \xrightarrow{\eta^l \otimes N} SS^l N = S(N^r S)^l$$

which means intuitively that a noun can be seen either as an atomic type N, or as a compound type waiting for an intransitive verb on the right to yield a sentence. Graphically, this simply amounts to tensoring our atomic noun with a sentence cap.

To coordinate Vladimir and Estragon, we need first to copy the verb wait with the following spider:

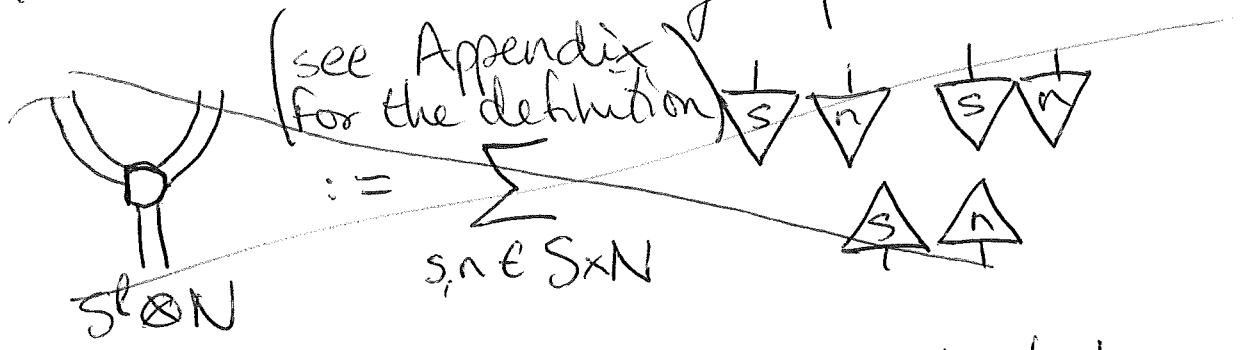

then we merge the two subsentences with the spider of the sentence space (which will behave as and)

Hence, we can now define the coordinator and as the product:

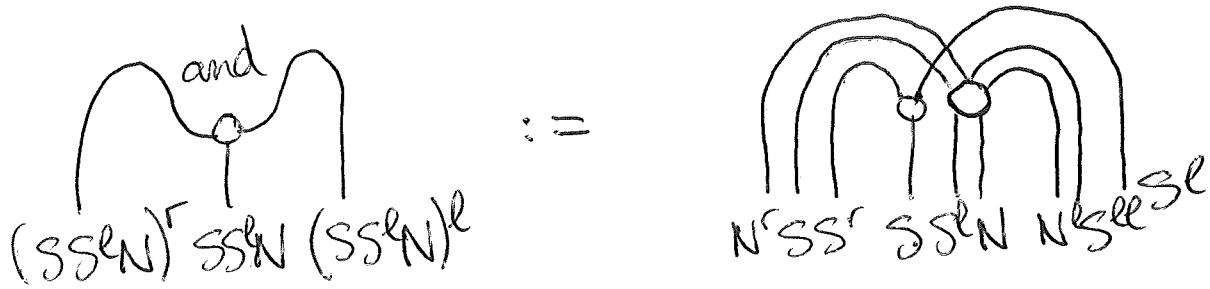

which allows us to compute the meaning of "Vladimir and Estragon" and how it interacts with "wait":

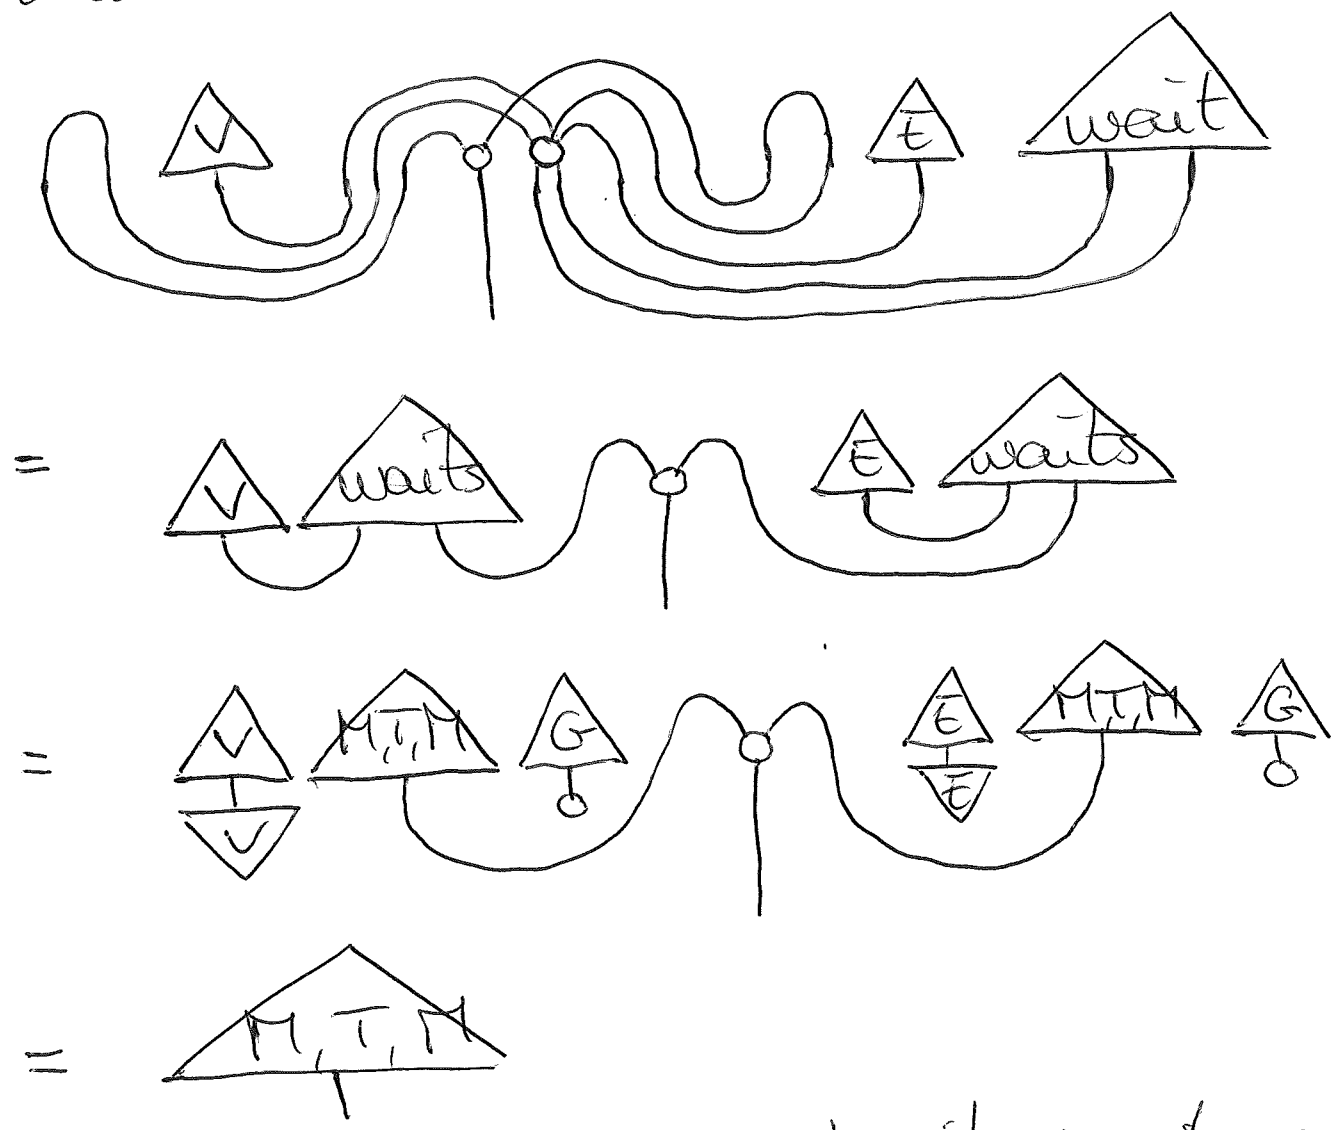

It is true today, maybe it was true yesterday, it may be true tomorrow.

The rest of our story will be told in the form of questions and answers. Following [Hamblin 1973] "Questions in Montague English", we will take the meaning of a question to be the set of its answers.

Following [Lambek 2008] we can turn a verb in declarative form into its question form by "bending the subject wire around", e.g.:

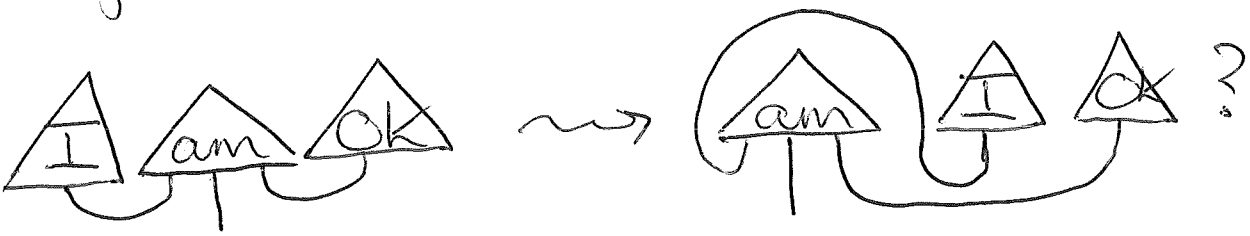

Then, a wh-word for a question of type X will have type XX^rS^r, e.g.

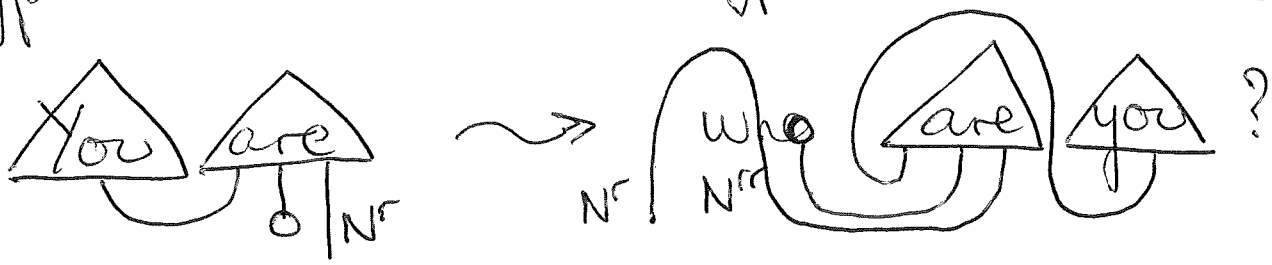

where we discard the sentence part of the verb to keep only the answer of the question.

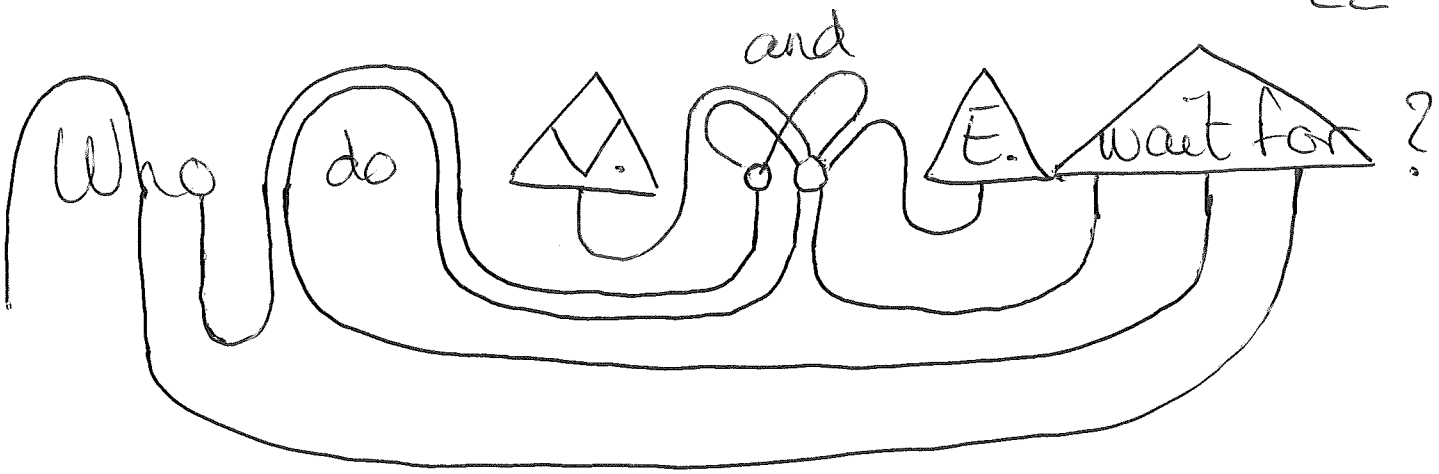
and
Who
do
V.
E.
wait for
?

=

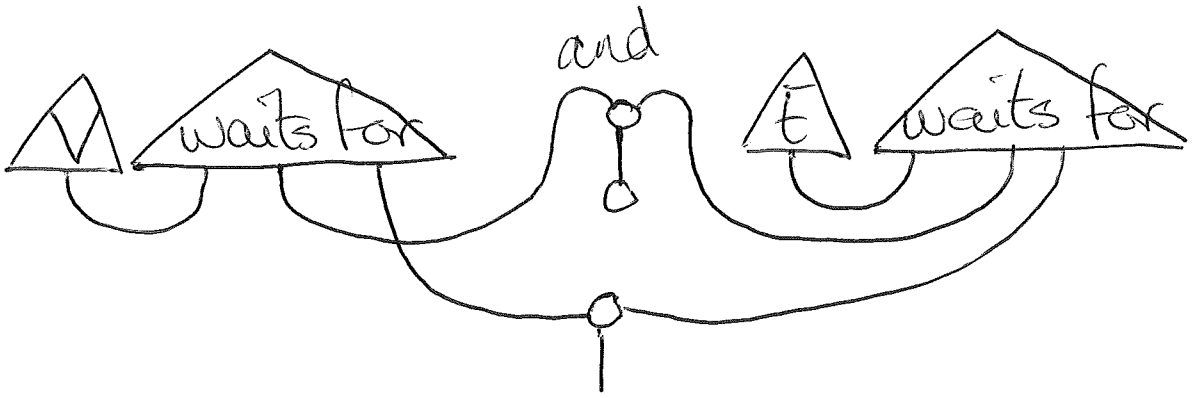
and
V
waits for
E
waits for

=

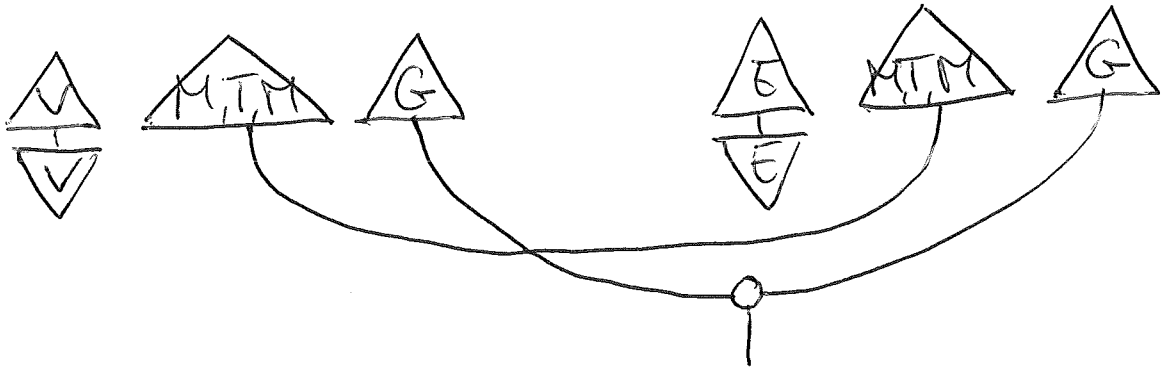
V
V
M,T,M
G
E
E
M,T,M
G

=

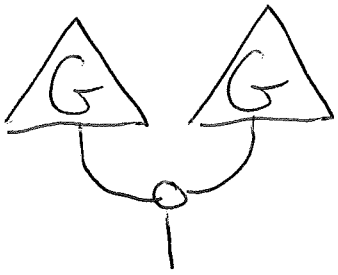
G
G

=


Godot

We now look at a more complex question where the desired answer is an intransitive verb:

=

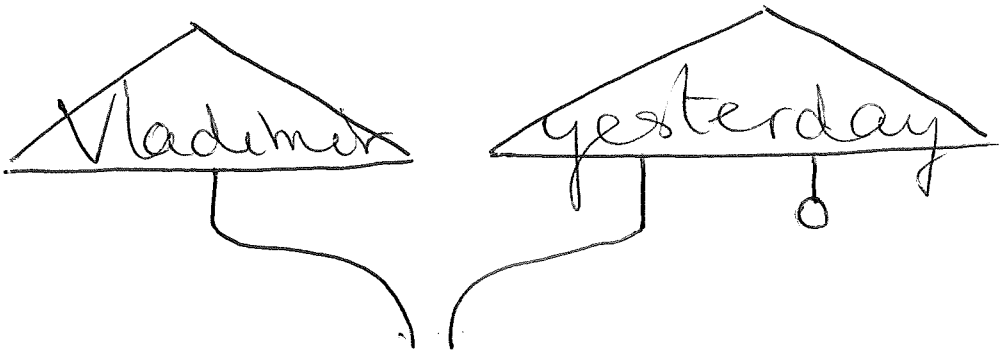

Hence, an answer to the question would be a verb X such that "Vladimir X yesterday" is true.
We can take the inner product with a potential answer to get a scalar:

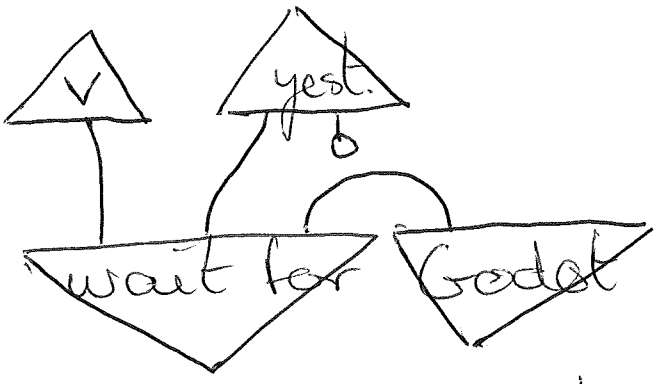

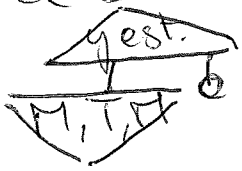

= Maybe (*)

Indeed, this is absurd theatre afterall.

(*) in the past dimension, true in both future and present.)

In order to model "when" questions, we need to use the type S^{S_T} for time adverbs:

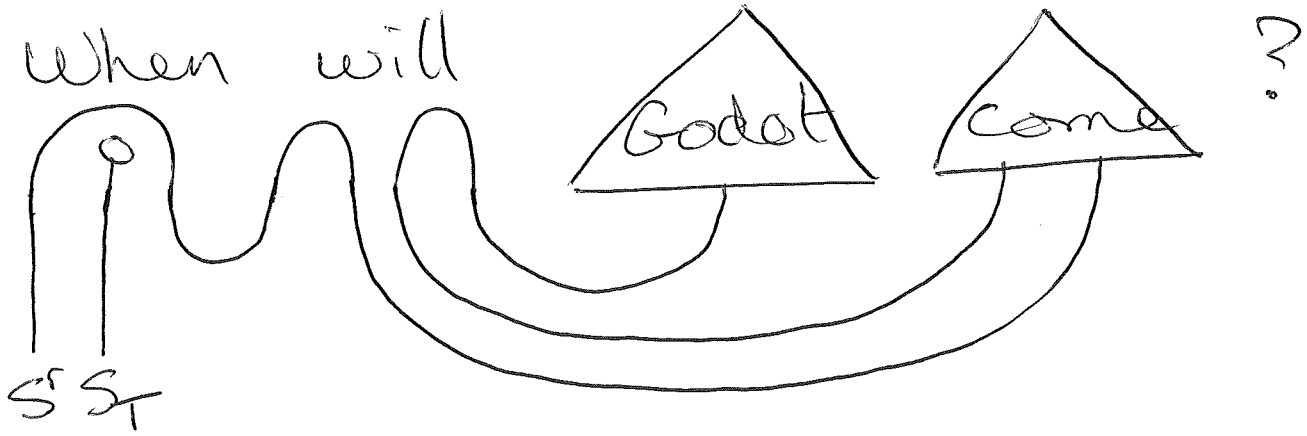

The messenger comes into the play and delivers Vladimir and Estragon a letter from Godot with the answer:

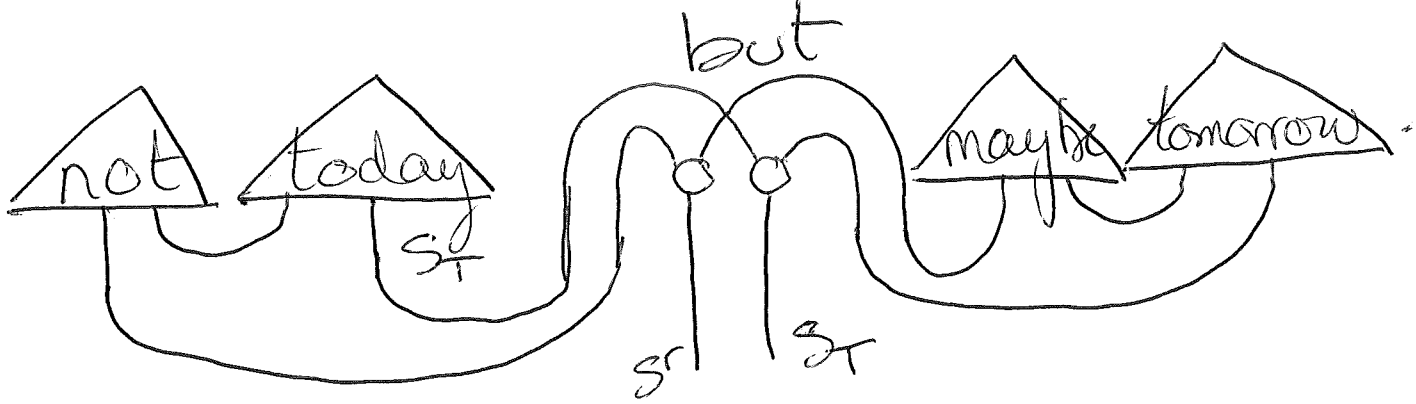

Taking the inner product we get:

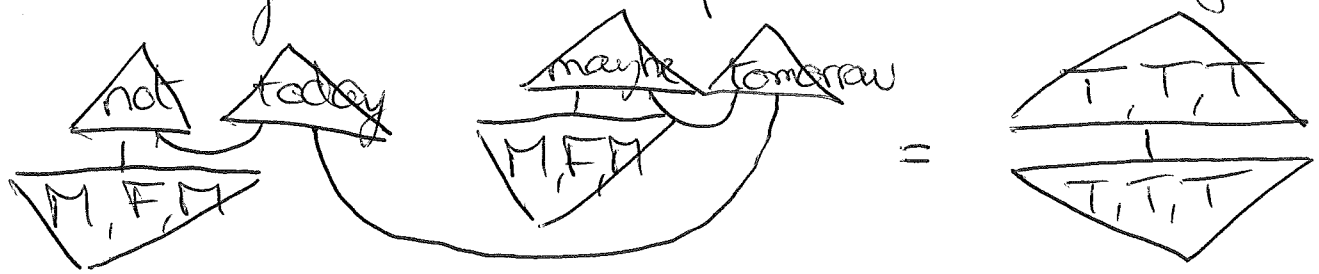

= True
(at all times)

Appendix: Definition of a verb spider

The aim would be to define a spider which can copy the verb "wait". However, the way we defined our verb spider is just the product of a noun spider and a sentence spider, we cannot hope that it copies an entangled verb, indeed we have:

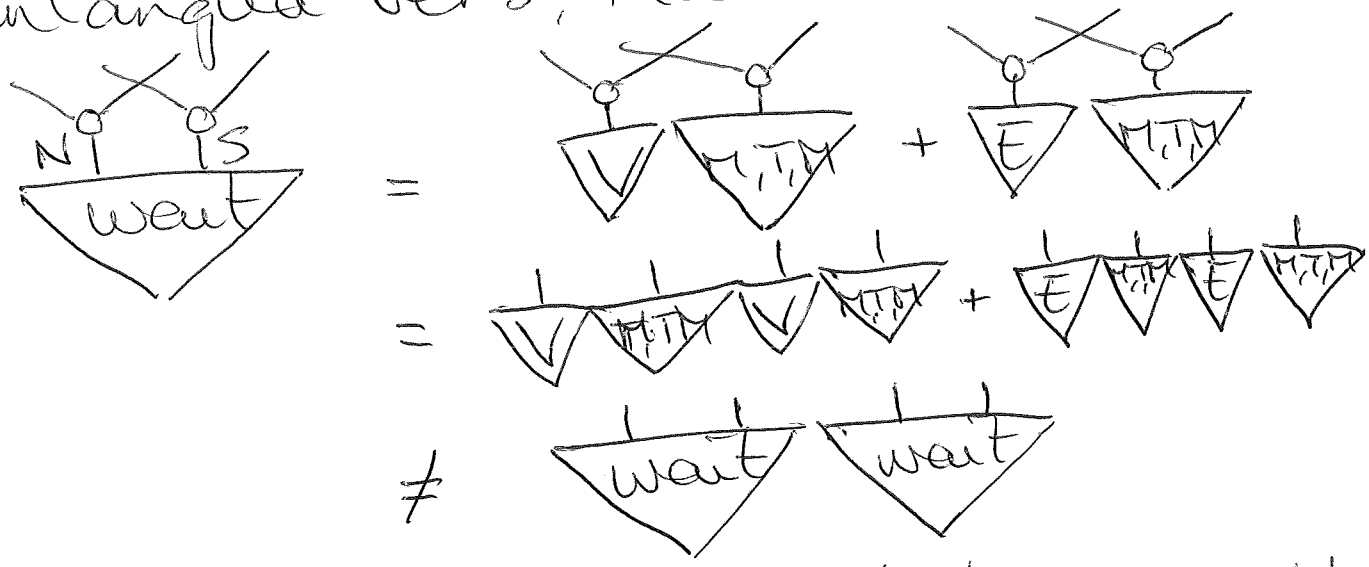

If we were working in Vect, we could easily define an ONB on $N \otimes S$ which includes the vector for wait. However, it is not obvious how to do this inside $Rel_\pi(Q)$. Hence, for the sake of our example, we will define the verb spider in the most ad hoc way possible: just specify that it copies "wait".

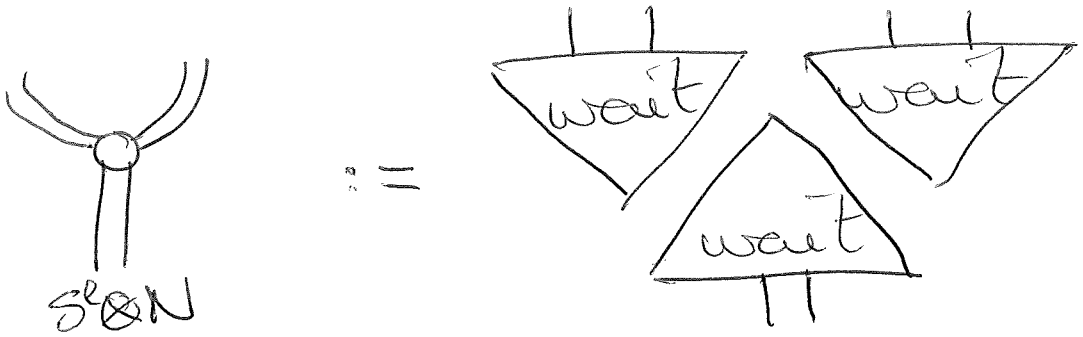

Then, because of how we setup our definition of "wait", we have the nice property that:

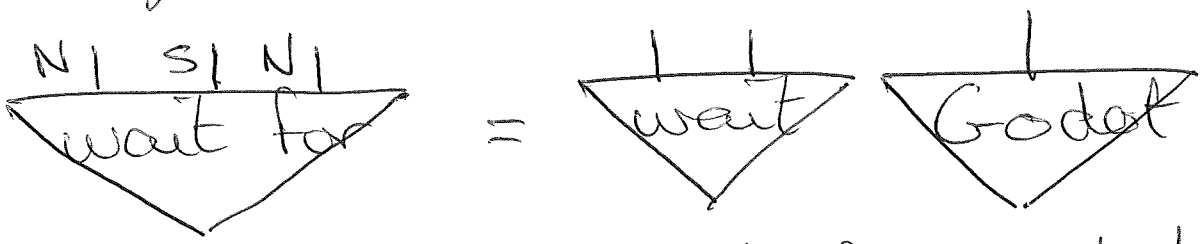

(indeed, one only ever waits for Godot) which implies that

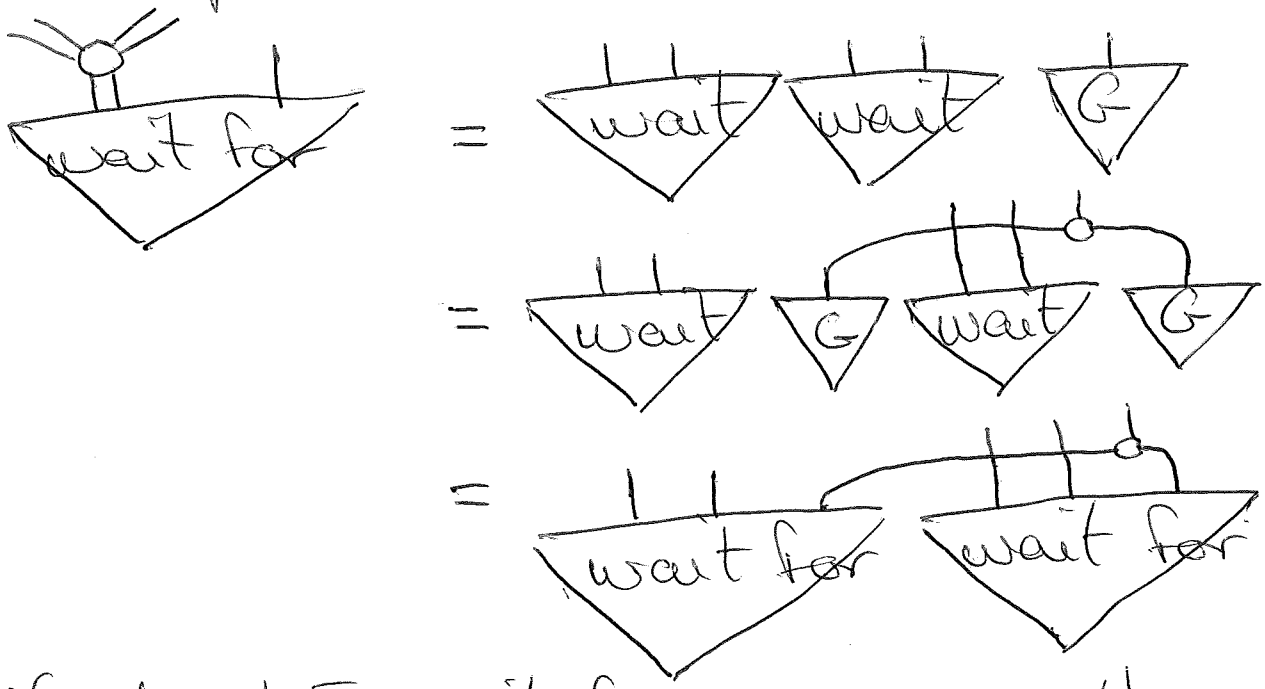

i.e. if V and E wait for someone, then they individually wait for the same person.

We use this property in the analysis of "Who do V. and E. wait for?", but it feels like it would require further investigation.

One way to look at this property is to see it as a naturality condition on our spiders, indeed we would like that:

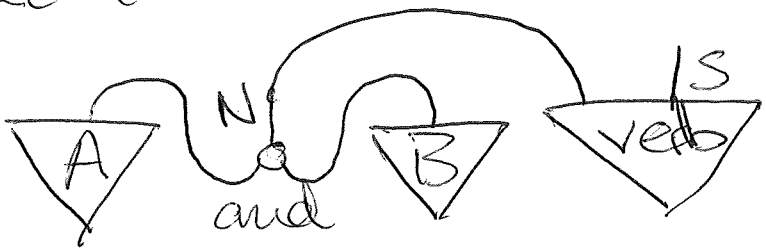

=

which is just a naturality square between the two spiders:

$$\begin{tikzcd} N \otimes N \arrow[r, "\text{verb} \otimes \text{verb}"] \arrow[d, "\text{and}_N"] \arrow[dr, phantom, "//"] & S \otimes S \arrow[d, "\text{and}_S"] \\ N \arrow[r, "\text{verb}"] & S \end{tikzcd}$$

Spiders in Vect depend on a basis, hence we cannot hope for naturality in Vect-based models of meaning.

Conclusion

To conclude this mini-project, we sum up the work we presented

Question 1

- We instantiated a compositional vector space model to tell the story of Sophocles' Oedipus Rex.
- We introduced a mechanism for modelling noun reference and the meaning of personal pronouns like "he", as well as possessives like "his mother".

Question 2

- We generalised a Rel-based model along two dimensions, modelling a basic notion of time with adverbs like "today" and "tomorrow", as well as a notion of uncertainty with "maybe".
- We introduced a basic notion of question and answer inside our model, and used this to tell the story of waiting for Godot.

- We experimented with complex patterns of coordination in natural language sentences, expanding on current research in the field and playing with the limits of the framework

Overall, I can say I am very pleased with the outcomes of this exam and this reading course more generally.

It allowed me both to explore a lively field of research, and to come to the point where I can start asking meaningful questions within that field!

:)