

This is a repository copy of *When the donkey lost its fleas: Persistence, minimal situations, and embedded quantifiers*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/44120/>

Version: Submitted Version

Article:

Zweig, E. orcid.org/0000-0002-3535-7259 (2007) *When the donkey lost its fleas: Persistence, minimal situations, and embedded quantifiers*. *Natural Language Semantics*. pp. 283-296. ISSN: 0925-854X

<https://doi.org/10.1007/s11050-006-9003-6>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

When the donkey lost its fleas: Persistence, minimal situations, and embedded quantifiers^{*}

Eytan Zweig
New York University

October 3, 2006

Abstract. This paper revisits the question of whether propositions in situation semantics must be persistent Kratzer (1989). It shows that ignoring persistence causes empirical problems to theories which use quantification over minimal situations as a solution for donkey anaphora Elbourne (2005), while at the same time modifying these theories to incorporate persistence makes them incompatible with the use of situations for contextual restriction Kratzer (2004).

1. Introduction

Kratzer (1989) introduces a framework for situation semantics that was taken as a starting point by a substantial body of later work. One property of this framework is that what is true of a small situation must remain true of larger situations that it is a part of. This is known as **persistence**. Kratzer's argumentation for this condition, however, is primarily of a conceptual nature. This led most of the work which adopted her framework to overlook this condition, and neglect to incorporate it into their theories.

In this paper, I will return to the issue of persistence, with several goals in mind. First and foremost, I aim to show that the persistence condition is not just motivated on conceptual grounds, but it is justified empirically. While doing so, I will also explore some of the requirements that are necessary for a proposition to be persistent. Finally, I shall discuss the consequences of persistence to different lines of research in situation semantics. Specifically, I will show that theories of donkey anaphora that require quantification over minimally small situations are in conflict with Kratzer's (2004) theory of contextual restriction, as the latter requires that quantification involve large situations in order to ensure persistence.

^{*} I would like to thank Anna Szabolcsi, Paul Elbourne, Chris Barker, Zoltan Szabo, François Recanati, Lena Baunaz, Jon Brennan, Andrea Cattaneo, Tom Leu, Lisa Levinson, Liina Pylkkänen, Laura Rimell, Oana Savescu-Ciucivara, and Jason Shaw for all their useful discussion and criticism in different stages of this paper's development. A version of this paper has been presented in *Sinn und Bedeutung X* in Berlin. Both the audience and the conference reviewers were very helpful with their comments.

2. Persistent Propositions

Kratzer (1989) introduces a situation semantics which relies heavily on the part-whole relationship of situations¹. Situations, according to this framework, are groupings of entities, their properties, and relations between them. Reference to situations is handled through situation variables, which can be quantified over just like other variables. Much of the power of this framework is derived from the fact that situations in this system are partially ordered by the sub-situation operator \leq . If $s \leq s'$, then s' may contain at least one entity, property, or relation that s does not. There is a maximal element to this ordering - the possible world, which, naturally, includes all the entities, properties, and relations that exist in that world. For brevity, I shall call a situation s' an **extension** of a situation s iff $s \leq s'$ and $s \neq s'$.

In this system, a proposition is defined as a set of situations, such that a proposition p is true in a situation s if $s \in p$. Nothing said so far prevents a proposition from being true in a situation s , but false in some extensions of it. For example, take the proposition p which is expressed in (1):

- (1) There are no living kings.

(1) is, under a straightforward analysis of its meaning, true of a situation s_1 that includes only an individual x and the fact that x is alive. However, there may be a larger situation s_2 that includes x , the fact that he lives, and the fact that he is a king. (1) is not true of s_2 . But note that $s_1 \leq s_2$.

As mentioned above, Kratzer (1989) takes the view that this is an unwelcome result. She suggests that a condition be added such that all natural-language propositions be **persistent**, following the definition below:².

- (2) A **persistent proposition** is a proposition of which it is true that, for every s such that $s \in p$, for every s' such that $s \leq s'$ it holds that $s' \in p$.

¹ While Kratzer's later work revises key aspects of this framework (see Kratzer, 2002 and Kratzer, 2005), the issues discussed in this paper remain constant throughout the various iterations of this system, both those authored by Kratzer herself and those adapted by other authors

² Terminology due Barwise and Perry (1983). It is important to distinguish this use of *persistent* from the use of the same term in Barwise and Cooper (1981), where it is used to denote "right upwards monotone".

With this condition in place, then, in the world described above, s_1 cannot be a member of the proposition expressed by (1), due to the existence of s_2 .

It is important to note that Kratzer does not enforce this condition by somehow filtering out non-persistent propositions. Rather, she provides denotations for quantifiers that encode persistence. For example, instead of the non-persistent denotation for *every* provided in (3), she suggests (4)³:

- (3) **Non-persistent quantification:**
 $\llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s. \text{ For all } x_{\langle e \rangle}: \text{ if } f(\lambda s.x)(s) = 1, g(\lambda s.x)(s) = 1$
- (4) **Persistent quantification:**
 $\llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s. \text{ For all } x_{\langle e \rangle}: \text{ if } f(\lambda s.x)(w) = 1, f(\lambda s.x)(s) = 1 \text{ and } g(\lambda s.x)(s) = 1$

The difference between the two quantifiers is as follows: in (3), the quantifier is restricted to entities which have property f in s , and it predicates of them that they also have property g . In (4), the quantifier is restricted to all the f s in the world, and it states that they have that property in s , as well as g . Thus, a proposition only holds of situations that include all the f s in w , and in which all of them are also g s. Both these properties will hold of every larger situation⁴.

While writing persistence into the determiner denotation ensures that all sentences end up denoting persistent propositions, it also complicates these denotations. Since Kratzer does not provide empirical justification for doing so, most of the literature following her work chose to use the simpler, non-persistent denotations⁵. The next section will examine one such theory, and show why this choice leads to empirical problems.

³ The denotations given below differ from Kratzer's in their notation, as I use the same formalism as Elbourne (2005). Nonetheless, the ideas are the same, with one major simplification: Kratzer (1989) deals with some distinctions which go beyond the scope of this paper, such as the distinction between propositions that are true accidentally and propositions that are true by some inherent fact about the nature of the world. I will ignore such distinctions here.

⁴ This is actually not entirely correct. Take the sentence *Every professor owns an even number of hats* - there can be a situation s that includes all the professors, and each of them has an even number of hats in that situation, but there's a situation s' in which one professor has an additional hat. I will ignore this issue in the discussion that follows, since it will not carry over to the quantifier denotations that use minimal situations.

⁵ For a discussion of persistence in non-Kratzarian situation theory, see Cooper (1991)

3. Minimal situations and donkey anaphora

3.1. THE HEIM/ELBOURNE SOLUTION FOR DONKEY ANAPHORA

One recent promising use of situation semantics has been to solve a problem that arises in the resolution of donkey anaphora. This line of research was first suggested by Heim (1990), and elaborated in work such as Elbourne (2005), von Steinhilber (2004, 1995) and Büring (2004). In the following discussion I shall make reference directly only to Elbourne's theory; however, a similar point could be made with respect to the other implementations.

Situation semantics become necessary because of an apparent problem for the E-type analysis (Evans, 1977, 1980) of donkey anaphora. The E-type analysis is one of the most attractive explanations of this phenomenon. In it, the donkey pronoun is taken to have semantics similar to a definite description, such that (5) is interpreted as (6):

(5) Every farmer who owns a donkey beats it.

(6) Every farmer who owns a donkey beats [the donkey].

However, there is a major problem with this solution: definite descriptions require a unique referent. Such a referent does not seem to be available in donkey anaphora; (5) can clearly be true in a context that contains multiple donkeys (and in fact, if there was only a single donkey, it would be hard to imagine (5) used with felicity).

The Heim/Elbourne solution relies on the insight that even if there is more than one donkey involved in the overall world, situation semantics allows us to access sub-situations of that world that contain only one donkey. Thus, it is possible to make use of those situations to ensure unique referents for the donkey pronouns.

All that needs to be done is to take care to only refer to situations small enough to contain exactly one donkey. For this purpose, instead of making reference to just any situations within the denotation of the quantifiers, instead they are taken to quantify over **minimal situations**. A minimal situation such that p holds is a situation $s \in p$ such that there is no situation $s' \in p$ such that $s' \leq s$ ⁶.

For example, the following is Elbourne's denotation for *every*:

(7) **Minimal quantification:**

$\llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s_1$. For all $x_{\langle e \rangle}$: for each minimal situation s_2 such that $s_2 \leq s_1$ and $f(\lambda s.x)(s_2) = 1$, there is a

⁶ The use of minimal situations is not novel to Heim (1990) but comes from Berman (1987). Heim, however, was the first to use minimal situations within the context of an E-type solution.

situation s_3 such that $s_3 \leq s_1$ and s_3 is a minimal situation such that $s_2 \leq s_3$ and $g(\lambda s.x)(s_3) = 1$

Paraphrased informally, *every* quantifies not over individuals that have a certain property (the NP restriction), but over sub-situations of its argument situation that contain only the individual and said property. For each of these situations, *every* claims that it is possible to extend it in such a way that a second property (the VP denotation) holds true of the individual⁷.

By adding this quantifier denotation to the E-type story, (5) can be informally paraphrased as (8):

- (8) Every situation can be divided up in such a way that for every sub-situation that involves a farmer, a donkey he owns, and nothing else, there is a situation that involves the farmer, the donkey, the ownership, and the fact that the farmer beats the unique donkey in that situation.

An example of such a situation can be seen in figure 1, where for each of the minimal farmer/donkey situations s_1 , s_2 , and s_3 there is an extension s_1' , s_2' , and s_3' respectively wherein the farmer beats the donkey.




Figure 1. Example of Elbournian situations

At first blush, this solves the problem, as, by virtue of being minimal, the minimal situation will never contain more than the single donkey necessary to make the subject have the property of being a farmer who owns a donkey. This donkey makes a good unique referent

⁷ It's worth noting that this approach is not trivially generalized to cardinal quantifiers such as *three*, as it is not clear how to count minimal situations. For interesting discussion on this topic, the reader is referred to von Stechow (2005) and Kratzer (2002, 2006).

(within the context of the situation) for the definite description to pick up. Thus, the E-type reference problem seems to be solved⁸.

3.2. THE PROBLEM

The preceding discussion, however, contains a henceforth unstated assumption. Namely that, whenever donkey anaphora occurs, an appropriate minimal situation that will provide a unique referent is available. Unfortunately, this is not always the case.

3.2.1. *The donkey that lost its fleas*

For example, take a world in which there are three farmers (A,B,C), each of which owns a donkey. Farmers A and B each take good care of their respective donkeys, grooming them daily. As a result, their donkeys have no fleas. Farmer C, however, does not groom his donkey, which has many fleas. This situation is shown in figure 2 below.




Figure 2. The donkey that lost its fleas

It is pretty uncontroversial that sentence (9) is true in this context (ignoring causality for the sake of simplicity):

- (9) Every farmer who owns a donkey which has no fleas grooms it.

But applying the minimal situation analysis as given above to this sentence, (9) is false in this scenario.

This can be seen by closely examining figure 2. There are three situations that involve farmers/donkey pairs (s_1 , s_2 , and s_3). Note especially s_3 , which involves farmer C, his donkey, the owning relationship between them, but no fleas, nor possession relations between the fleas and the donkey. s_3 conforms to the requirements of being a minimal situation that contains a farmer who owns a donkey which has no fleas.

⁸ There are further issues to be addressed as to what happens when a single farmer owns more than one donkey and similar cases. I refer the reader to Elbourne (2005) for detailed discussion.

Due to the denotation of *every*, every such minimal situation needs to have an extension wherein the farmer in question (farmer C) grooms the donkey. However, there is no situation that satisfies that requirement, and thus the sentence is false.

3.2.2. *The donkey hiding out of the situation's reach*

A second manifestation of this problem can be seen in the following sentence:

- (10) Every man who owns a farm beats every donkey in it.

According to the minimal situation analysis as given above, this is a tautology.

This is because the restriction of the quantifier requires that the quantification be over minimal situations in which a man own a farm. These situations obviously do not include any donkeys, as none are mentioned in the quantifier's restriction. But every such situation has many extensions which have nothing to do with donkeys or beatings. Let us take the scenario given in figure 3 below. In such a scenario, the sentence intuitively comes out as false, as not all the men beat all the donkeys in their farm. Specifically, the man in s_3 fails to beat one of the two donkeys in the farm he owns.




Figure 3. The donkey hiding out of the situation's reach

But note that s_3 has an extension s_3' wherein there is only one donkey, which is beaten by the man in question. Thus, the sentence is true. Even if the man in question was actually beating none of the donkeys in his farm, then the sentence would still be true, since the minimal situation including him and his farm would have many arbitrary extensions (say, the one including the man, the farm, the owning relationship, and the man's hat) that contain no donkeys and

thus the farmer beats every donkey in them by default. Thus, (10) is always going to be true⁹.

3.2.3. *What went wrong*

There is a clear intuitive notion of what is wrong in these examples. In (9), The minimal situation that includes farmer C and his donkey includes no fleas; yet it feels like it should not count as a minimal situation *of a farmer who owns a donkey with no fleas*, as the donkey in question does have fleas outside this situation. In (10), it does not feel sufficient that for every man/farm pair there is an arbitrary extension in which all the donkeys in that extension are beaten. Rather, it seems that the man should beat every donkey in an extension includes all the donkeys in the farm.

It is here that persistence is needed.

In (9), what is necessary is to quantify over minimal situations that involve a donkey with no fleas, **and** are not sub-situations of a situation for which said donkey has fleas. In (10), it is required that the man beat every donkey in the farm in the situation in question, and that there be no extension of that situation in which the farmer doesn't beat every donkey in the farm.

Thus, it can be seen that ignoring persistence creates problems for Elbourne's framework. The obvious way to correct these problems is to reintroduce persistence into the equation.

4. Persistence - consequences and implementation

In the previous section, I found some problems for the Heim/Elbourne analysis of donkey anaphora and suggested that modifying their theory to ensure persistence will solve these problems. In this section I shall demonstrate this.

4.1. PERSISTENCE AND MONOTONICITY

Not all determiners need to have persistence explicitly written into their denotations. Those that denote quantifiers that are upwards monotone on both arguments are, in fact, persistent by default.

⁹ This ignores the possibility that *every* has an existence presuppositions. If such a presupposition is reintroduced, then (10) will no longer be a tautology. However, this does not solve the problem, as the sentence will only require that the man beats *at least one* donkey in his farm to be true, so the scenario in figure 3 would still erroneously validate the sentence.

To see why monotonicity matters, it is helpful to look at the denotation of a quantifier that does not have persistence written in, such as the denotation of *every* given in (3), repeated below as (11):

$$(11) \quad \llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s. \text{ For all } x_{\langle e \rangle}: \text{ if } f(\lambda s.x)(s) = 1, g(\lambda s.x)(s) = 1$$

The quantifier is restricted to entities x that have property f in a situation s . Because the sub-situation relation \leq is upwards monotone, then, assuming that f does not in itself contain any downwards entailing operators, if something has the property f in s it has the property f in every s' such that $s' \leq s$. In other words, the set of xs that have property f in s is a subset of the set of xs that have the property f in s' .

Thus, going from a situation to an extension of it in essence replaces the domain argument of the quantifier by a superset of it. This is always safe if the determiner is upwards monotone in its restriction, but not if it is downwards or non-monotone in that argument. Parallel reasoning applies to the nuclear scope of the determiner. This means that if a determiner is upwards monotone in both arguments, nothing needs to be added for it to provide persistent quantification¹⁰.

4.2. QUANTIFIER MONOTONICITY VS. SENTENTIAL ENTAILMENT

It is worth noting that it is the monotonicity of the quantifiers that matters, rather than the entailment properties of any particular sentence. For example, note that for (9), the quantifier *no fleas* is embedded in the restriction of the quantifier *every farmer*. This means that the argument slots of *no fleas* are actually an upwards entailing environment, as can be seen from the following inference pattern:

- (12) Every farmer who owns a donkey which has no fleas grooms it.
 (13) a. \nRightarrow Every farmer who owns a donkey which has no red fleas grooms it.
 b. \Rightarrow Every farmer who owns a donkey which has no parasites grooms it.

Based on this information, one could be led to expect that there should be no persistence problems associated with the arguments of *no*. But,

¹⁰ It is worth noting that determiners are not the only possible source of downwards monotonicity; negation, negative adverbials, and “negative” predicates such as *lose* all are downwards monotone with respect to some of their arguments. Whether or not these elements include quantification over situations in their denotations is a question beyond the scope of this paper. If they do, however, then the conclusions arrived here will apply to them as well.

as shown in section 3.2.1, that is incorrect. The reason is that while entailment is calculated by the sentence as a whole, persistence must be ensured in embedded propositions as well as matrix ones. (9) can be paraphrased as the follows:

- (14) Every x of which it holds that x **is a farmer that owns a donkey that has no fleas** is such that x grooms the relevant donkey.

For the whole sentence to express a persistent proposition, the bolded proposition must itself be persistent for each x . If it is not, then going from a situation to an extension of it may alter the domain of the matrix quantifiers, by changing whether individual farmers fall under the restriction or not. This is the nature of the problem in example (9).

Thus, the nature of the embedded quantifier is relevant, even if ultimately its arguments end up being an upwards entailment environment. This shows that the decision in Kratzer (1989) to include the persistence condition in the denotation of (non-upwards monotone) quantifiers is the correct way to handle persistence, and I will follow suit.

4.3. IMPLEMENTING PERSISTENCE

Since failures of persistence arise when a proposition that was true in a small situation fails to be true in a larger one, the best way to prevent this is to check that the proposition holds in as large a situation as possible. This is a potential problem, as the Heim/Elbourne solution for donkey anaphora relies on the presupposition that minimal situations give unique referents. Can persistence be implemented in a way that satisfies both demands?

In fact, there is no need to look beyond what was already discussed to find an implementation that makes this possible. The persistent quantification in Kratzer (1989) adds a condition that the individuals quantified must satisfy the restriction of a quantifier in the largest situation available (i.e., the entire world) **in addition** to the situation quantified over. This denotation allows checking persistence against the maximal situation w , while at the same time the actual quantification remains on truly minimal situations. Thus, the best of both worlds has apparently been achieved, at least as far as using situations to account for donkey anaphora. Adding such a condition to Elbourne's *every* results in the following:

- (15) **Persistent minimal quantification:**
 $\llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s_1$. For every $x_{\langle e \rangle}$: if $f(\lambda s.x)(w)$

$= 1$, then $f(\lambda s.x)(s_1) = 1$ and for every minimal situation s_2 such that $s_2 \leq s_1$ and $f(\lambda s.x)(s_2) = 1$, there is a situation s_3 such that $s_3 \leq s_1$ and s_3 is a minimal situation such that $s_2 \leq s_3$ and $g(\lambda s.x)(s_3) = 1$

This denotation of *every* (and a similarly modified denotation for *no*) would avoid both of the problems for Elbourne's system. In the case of the donkey that lost its fleas, the reasoning is simple: farmer C is not a farmer who owns a donkey with no fleas in w , and thus does not fall under the domain of quantification. The other problem is a bit more complex: the matrix *every* quantifies over all the men in w that own a farm, and for each minimal situation that includes such a pairing, it states that there is an extension wherein every donkey in the farm is beaten. So far, the persistence makes no difference. But the embedded *every* now quantifies over every entity in w that is a donkey in the farm in the relevant minimal situation, rather than just those donkeys that are present in an arbitrary situation. Thus, no donkeys can escape notice.

But this denotation is only possible under the assumption that reference to w in a determiner denotation is unproblematic. In the following section, it shall be shown that this does not fit comfortably with other recent uses of situation semantics.

5. Persistence and contextual restriction

One property of persistent quantification as discussed so far is that it is global; every quantifier in some sense quantifies over the whole world.

If nothing further is said, this leads to strange-looking predictions. Take the following sentence, for example:

(16) Every tree is laden with wonderful apples.

By global persistence, (16) would only be true if every tree in the entire world is laden with wonderful apples. Kratzer (1989) solves this by appealing to contextual domain restriction to fill in additional descriptive material. According to her, (16) really should be given a reading along the lines of the following:

(17) Every tree [in my orchard] is laden with wonderful apples.

This is an intuitively appealing notion, as it is a well-established fact that contextual restriction must come into play in exactly these sentences anyway. However, the viability of this option depends heavily on the way in which contextual restriction is implemented. While Kratzer

(1989) does not provide an actual theory of contextual restriction, she is clear that this must be done by an additional mechanism rather than the situations themselves, explicitly rejecting the theory of contextual restriction provided in Barwise and Perry (1983) because it relies on non-persistent propositions.

5.1. CONTEXTUAL RESTRICTION VIA TOPIC SITUATIONS

In contrast to her earlier position, Kratzer (2004) proposes that contextual restriction should be accounted for not by adding descriptive material to the sentence, but rather by applying the proposition in question to a *topic situation*, which contains only the contextually relevant entities.

According to Kratzer, utterances in context represent an *Austinian proposition* (after Austin, 1950) - that is, a pairing of a topic situation and a proposition $\langle s, p \rangle$. An assertion operator ASSERT is responsible for applying the topic situation as a situation argument for the proposition (i.e., the one required by the λs of the highest scope operator)

$$(18) \quad \llbracket \text{ASSERT} \rrbracket(\langle s, p \rangle) = p(s)$$

Since every embedded operator is passed a situation variable by the next higher operator which is a sub-situation of the situation parameter of that operator, this ensures that all quantifiers are restricted to elements of the topic situation.

Put differently, this system relies on the principle that each operator only has access to the situation that the operator above gives it, and can only pass down parts of that situation to lower operators. This, indeed, recaptures one of the intuitive uses of situations; they are used in order to talk about just part of the world¹¹.

This principle would be nullified if direct reference to w is allowed, such as used above to ensure persistence. Doing so allows a quantifier to see information that was not strictly passed down to it by a higher operator. For example, imagine the following scenario: yesterday, a semantics exam was graded. Exactly one student got a B; surprisingly, she did so without making any actual errors, but just by failing to answer questions in a satisfactory manner. (19) is true in this scenario:

¹¹ Note that Kratzer (2004) does not specifically rule out an additional mechanism for contextual restriction. In fact, she argues that such a mechanism must exist for restrictions that are based on cultural conventions. But for the purposes of this paper, what is important is that normal contextual restriction, i.e. the kind that determines the relevant apples for the use of *every apples* in (16), is handled via topic situations.

(19) Some student who made no errors got a B.

(19) requires the existence of a student who made no mistakes in the relevant context - i.e., on her semantics exam. It will not be falsified if that same student made an error in her phonology exam.

However, if persistence is checked relative to the world, then the error on the phonology exam will be enough to remove the student from the domain of quantification (for there are errors in w which she made), thus falsifying the sentence¹².

5.1.1. *Local persistence*

Accepting the theory of contextual restriction in Kratzer (2004), then, means that a different way of implementing persistence is necessary: one wherein persistence is local to the situation which the quantifier received as an argument.

Note that, if minimal situations are ignored, local persistence actually comes for free in Kratzer (1989). The denotation of *every* given in (3) (repeated below as (20)) is only problematic as far as persistence is concerned because the situation argument a given sentence takes would often be smaller than the entire world, which was the global level up to which persistence was desired. If, following Kratzer (2004), this situation variable is taken to always reflect the contextual domain wherein persistence needs to hold, (3) (repeated as (20)) will suffice.

$$(20) \quad \llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s. \text{ For all } x_{\langle e \rangle}: \text{ if } f(\lambda s.x)(s) = 1, g(\lambda s.x)(s) = 1$$

In the Heim/Elbourne system, however, things are not so simple. The first problem is that having the property specified in the restriction is only checked in a minimal situation, not in the actual contextual situation. This can be solved with a minimal modification of (15), replacing

¹² One possible objection at this point is that the sentence is true even in the larger situation because of implicit anaphora within it; in other words, that the sentence really means:

(i) Some student who made no errors IN AN EXAM got a B IN THAT EXAM.

While accounting for this covert anaphora is an interesting problem in its own right, it does not allow the sentence to apply to a wider situation. To see this, instead of the scenario above imagine that the only student who made no errors in semantics got an A in that exam. At the same time, she also made no errors in phonology, yet got a B in phonology regardless. If I were to utter (19) while I was discussing the grading of the semantics exam, the sentence would be false. Thus, the anaphora effect on its own is insufficient as a means to rescue persistence checked relative to the world.

the reference to w with reference to *every*'s situation parameter s_1 , as follows:

- (21) **Locally persistent minimal quantification:**
 $\llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s_1$. For every $x_{\langle e \rangle}$: if $f(\lambda s.x)(s_1) = 1$, then for every minimal situation s_2 such that $s_2 \leq s_1$ and $f(\lambda s.x)(s_2) = 1$, there is a situation s_3 such that $s_2 \leq s_3$ and $g(\lambda s.x)(s_3) = 1$

(21) can handle the problem of the disappearing fleas as well as (15) can. Simply put, it is not sufficient that a minimal situation can be found that contains a farmer, his donkey, and no fleas, it is also necessary that he has no fleas in the context situation. This is all that is necessary to get the correct reading for that sentence.

However, there is a second problem. Unlike in the simple case of (3), in the minimal situation-based theory embedded quantifiers no longer have access to everything in the topic situation, but only have access to what is in the situation passed down to them from the higher quantifier, as desired. This, unfortunately, reintroduces the other problem. To see this, let's return to (10), repeated as (22):

- (22) Every farmer who owns a farm beats every donkey in it.

As before, the minimal situation (call it s_{farm}) in which a farmer x owns a farm contains no donkeys. Now take an arbitrary extension (s_{farm+}) of that situation, such that s_{farm+} contains no donkeys. By the definition of the quantifier, it is now necessary to check whether *beats every donkey in it* is true of x in s_{farm+} . This involves passing s_{farm+} as the situation parameter of the embedded quantifier *every*. This is the largest situation which the persistence condition of *every* can see. But there are no donkeys in the farm in s_{farm+} . Thus, the persistence condition is toothless in this scenario.

Thus, domain restriction that relies on situations variables being passed down from one operator to the next prevents using persistence to solve the problem of elements hiding outside minimal situations.

5.1.2. Possible alternatives

Other methods of using situations for domain restriction may not suffer from this problem:

One possible solution is to claim that the topic situation is always available for direct reference in a discourse. Thus, it is possible to use the definition in (15), simply replacing the reference to w with s_{topic} :

- (23) **Locally persistent minimal quantification (alternative):**
 $\llbracket \text{every} \rrbracket = \lambda f_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda g_{\langle \langle se \rangle, \langle st \rangle \rangle} \lambda s_1$. For every $x_{\langle e \rangle}$:
 if $f(\lambda s.x)(s_{topic}) = 1$, then $f(\lambda s.x)(s_1) = 1$ and for every minimal situation s_2 such that $s_2 \leq s_1$ and $f(\lambda s.x)(s_2) = 1$, there is a situation s_3 such that $s_3 \leq s_1$ and s_3 is a minimal situation such that $s_2 \leq s_3$ and $g(\lambda s.x)(s_3) = 1$

Another possibility, raised by Recanati (2004), is that topic situations are not used to saturate a situation argument slot, but rather are added as a form of semantic enrichment. Such a system would differ enough from Kratzer (2004) that the results above would not necessarily hold for it (though other problems may well rise, based on the exact implementation).

6. Conclusion

In this paper, I explored the notion of persistence and have shown that the form in which it is implemented has crucial consequences for the applications of situation semantics in linguistics. While it is debatable whether a global implementation of persistence, wherein it holds all the way up to the entire world, is necessary, some form of persistence is needed to avoid information “leakage” between embedded quantifiers and the quantifiers scoping over them. One example of such a problem, which arises in the system of Elbourne (2005), was described.

Attempting to solve this problems taught us more about the nature of persistence and how it interacts with minimal situations. Along the way it was shown that implementing a persistent minimal situations approach to donkeys is impossible if the contextual restriction method proposed in Kratzer (2004) is also used. This incompatibility goes unnoticed if persistence is ignored.

Situation semantics are a powerful tool because they allow access to, and quantification over, parts of worlds, thus allowing information to be partitioned into chunks of the correct granularity for various semantic application. However, this tool arrives with a cost, since without care it is easy to lose important information along the way, arriving at incorrect results. Persistence is important in that it controls the flow of information between different levels of embedding in a sentence. While, with sufficient care, it may be possible to write a situation semantics (for a situation semantics fragment) that does not use persistence, simply ignoring it will likely result in hidden flaws in the system.

References

- Austin, J. L.: 1950, 'Truth'. In: *Philosophical Papers (3rd edition)*. Oxford: Oxford University press, pp. 117–133.
- Barwise, J. and R. Cooper: 1981, 'Generalized quantifiers and natural language'. *Linguistics and Philosophy* **4**, 159–219.
- Barwise, J. and J. Perry: 1983, *Situations and attitudes*. Stanford, California: CSLI Publications.
- Berman, S. R.: 1987, 'Situation-based semantics for adverbs of quantification'. In: J. Blevins and A. Vainikka (eds.): *University of Massachusetts occasional papers 12*. Amherst: GLSA.
- Büring, D.: 2004, 'Crossover situations'. *Natural Language Semantics* **12**, 23–62.
- Cooper, R. P.: 1991, 'Persistence and structural determination'. In: J. Barwise, J. M. Gawron, G. Plotkin, and S. Tutiya (eds.): *Situation Theory and its Applications, Volume 2*, Vol. 26 of *CSLI Lecture Notes Series*. Stanford, California: CSLI Publications, pp. 295–309.
- Elbourne, P.: 2005, *Situations and individuals*. Cambridge, Mass.: MIT Press.
- Evans, G.: 1977, 'Pronouns, quantifiers and relative clauses'. *Canadian Journal of Philosophy* **7**, 467–536.
- Evans, G.: 1980, 'Pronouns'. *Linguistic Inquiry* **11**, 337–362.
- von Fintel, K.: 1995, 'A minimal theory of adverbial quantification'. In: B. H. Partee and H. Kamp (eds.): *Context Dependence in the Analysis of Linguistic Meaning: Proceedings of the Workshops in Prague, February 1995, Bad Teinach, May 1995*. IMS Stuttgart Working Papers, pp. 159–193.
- von Fintel, K.: 2004, 'Restrictions on quantifier domains'. Ph.D. thesis, U. Mass, Amherst.
- von Fintel, K.: 2005, 'How to Count Situations (Notes Towards a User's Manual)'. ms, MIT.
- Heim, I.: 1990, 'E-type pronouns and donkey anaphora'. *Linguistics and philosophy* **13**, 137–177.
- Kratzer, A.: 1989, 'An investigation of lumps of thought'. *Linguistics and philosophy* **12**, 607–653.
- Kratzer, A.: 2002, 'Facts: particulars or information units'. *Linguistics and philosophy* **25**, 655–670.
- Kratzer, A.: 2004, 'Covert domain quantifier restrictions'. U. Mass Ms.
- Kratzer, A.: 2005, 'Constraining premise sets for counterfactuals'. *Journal of semantics* **22**, 153–158.
- Kratzer, A.: 2006, 'Situations in Natural Language Semantics'. Ms, U. Mass, Amherst. To appear in Stanford encyclopedia of philosophy.
- Recanati, F.: 2004, *Literal meaning*. Cambridge, England: Cambridge University Press.