

PRODUCER INTERPRETATIONS OF THE ENGLISH PRE-NOMINAL GENITIVE

Per Anker Jensen, University of Southern Denmark

paj@sitkom.sdu.dk

Carl Vikner, Copenhagen Business School

cv.id@cbs.dk

Abstract

This paper presents a theory of the pre-nominal genitive in English based on the claims that it is argument-only and polymorphic, on the one hand allowing a control interpretation delivered by 's itself, and, on the other hand, three 'constructional interpretations': an inherently relational, a part-whole, and a producer interpretation delivered by the head noun of the genitive construction. The paper addresses four types of evidence seeming to raise problems for the claim that producer interpretations are semantic, i.e. available as default interpretations, rather than pragmatic. We argue that in spite of this evidence, the hypothesis can be upheld. Finally, independent evidence is adduced concerning the semantics of the Hebrew Construct State Construction, which allows precisely the range of semantic interpretations we have proposed for the English pre-nominal genitive, including the producer interpretation, while not allowing pragmatic interpretations.

1 Introduction

We start by giving a very rough outline of our theory of the semantics of the pre-nominal genitive in English, in order to show how the producer interpretation fits into that theory (cf. Vikner & Jensen 2002). Then we turn to the subject matter proper of this paper, which is split up into two parts. The first part deals with a number of problems connected with positing the producer interpretation as a semantic reading of the pre-nominal genitive. The second part deals with our proposed solutions to those problems. Before concluding we briefly mention an interesting fact about the Hebrew Construct State Construction, which relates to the producer interpretation.

2 A Theory of the Semantics of Pre-nominal Genitive Phrases

Our point of departure is that all English pre-nominal genitive constructions crucially involve a relation which is not explicitly expressed (we call this relation the 'genitive relation'), and that they allow for a great variety of relational interpretations. Our principal aim is to identify the range of relations available to the pre-nominal genitive, and we focus on the problem of determining which relations are made available, and where they come from.

Our theory is an 'argument-only hypothesis' in that it holds that the first NP of a genitive construction (by some called the 'possessor phrase') always delivers an argument to the genitive relation. Further, it is polymorphic in that it ascribes two semantic representations to the genitive element 's. Importantly, it is not a theory of the salience of particular interpretations in particular situations, but a theory which predicts possible interpretations only.

2.1 Methodological Remarks on Semantic and Pragmatic Interpretations

We distinguish between two kinds of interpretation of a phrase, and in so doing we follow an idea proposed by Briscoe, Copestake and Boguraev (1990: 42-43; cf. also Pustejovsky 1991: 429-31), who outline a “default theory” which introduces a distinction between “default” and “non-default” interpretations. A default interpretation of *John began a novel* is ‘John began to read a novel’, whereas ‘John began to translate a novel’ would be a non-default interpretation made possible by pragmatic inference in a marked, informationally rich context (Briscoe, Copestake, Boguraev 1990:43). We call these two kinds of interpretations “semantic” and “pragmatic”, respectively.

The *semantic interpretations* of a phrase are “privileged” in the sense that the information needed to determine them is part of the lexical semantics of the elements of the phrase, and thus, in principle, it should be possible to extract this information from the relevant lexical entries, provided the lexical description given in the lexicon is rich enough. So, for instance, in connection with the example *John began a novel*, it is assumed that the lexical entry for *novel* contains the information that a novel is designed for reading, but not the information that a novel is designed for translating. Thus, there are fairly strong restrictions on the amount of information to be included in the lexicon, a point we shall elaborate on in a little while. The determination of *pragmatic interpretations*, on the other hand, cannot be carried out on the basis of lexical knowledge alone, but depends essentially on pragmatic knowledge and discourse knowledge. For instance, the interpretation ‘John began to translate a novel’ of the example just cited, presupposes that the hearer knows that John is a translator (or has available some similar piece of information about John or about the actual situation).

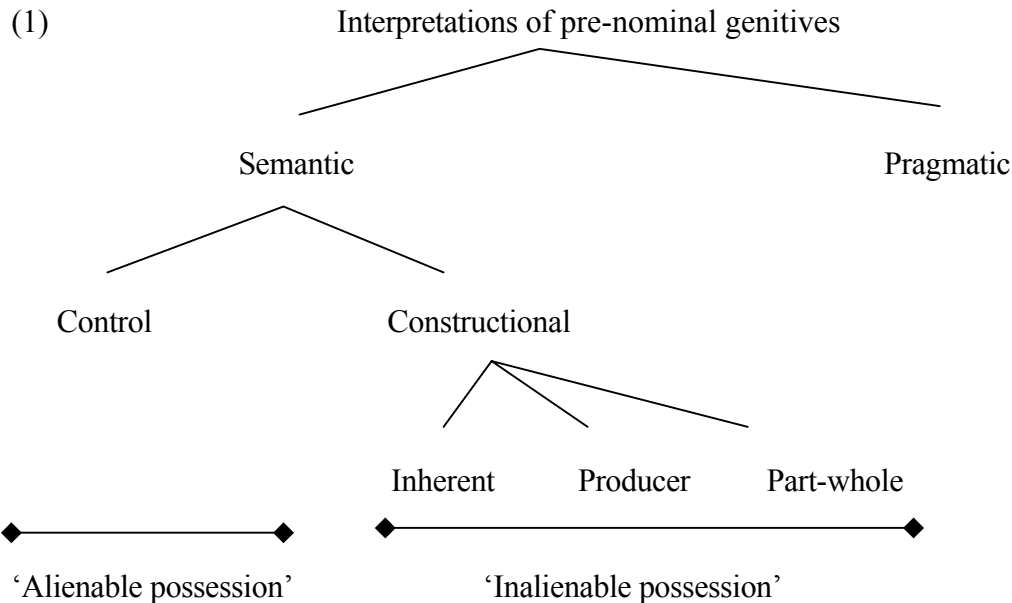
Thus, semantic interpretations only require knowledge of the meaning of the relevant *words*, whereas pragmatic interpretations not only require knowledge of the meaning of the relevant *words*, they also require knowledge of the *individuals* referred to by each of the two nominal expressions in the pre-nominal genitive construction. This being so, it is evident that there may be infinitely many pragmatic interpretations of a given phrase, whereas there are very few semantic ones.

2.2 Semantic Interpretations

We distinguish four types of semantic interpretations, which we call the ‘control relation’, the ‘inherent relation’, the ‘producer relation’, and the ‘part-whole relation’, respectively.

We do not think that these semantic relations originate from the same linguistic source. The control interpretation seems to stem from the genitive element itself. For instance, in an example like *Ann’s car* control comes neither from *Ann* nor from *car*, so ‘s seems to be the most plausible lexical source. The other three semantic interpretations differ radically from the control relation in that they all depend on the meaning of the head noun of the genitive construction. That is, in these three cases the genitive relation is picked up from another part of the genitive construction. Therefore, we bring together these three kinds of interpretation under the term ‘constructional interpretations’, cf. Borschev & Partee (2000: 179,192).

In overview, then, the types of interpretations we distinguish, are the ones shown in (1):



The four semantic interpretations are not mutually exclusive, and hence a genitive construction may be several ways ambiguous. Thus, for instance, an example such as *the farmer’s picture* is three ways semantically ambiguous between a control interpretation meaning ‘the picture that the farmer has at his disposal’, an inherently relational interpretation meaning ‘the picture that depicts the farmer’, and finally the producer interpretation meaning ‘the picture that the farmer has painted’.

2.3 Modelling Lexical Semantic Knowledge

Our semantic analysis draws heavily on information assumed to be encoded in the lexicon. When organizing this information, we follow James Pustejovsky’s theory of lexical structure (cf. Pustejovsky 1991, 1995). His theory assumes four levels of lexical representation. We concern ourselves only with two of these: Argument structure and Qualia structure. In (2) we have given some sample lexical entries. In a lexical entry x is a distinguished variable representing the object denoted by the lexical item in question, cf. Pustejovsky (1991:427).

(2) *farmer*

Argument structure: $\lambda x[\text{farmer}'(x)]$

Qualia structure: ...

poem

Argument structure: $\lambda x[\text{poem}'(x)]$

Qualia structure:

TELIC: $\lambda x[\lambda y[\text{read}'(x)(y)]]$

AGENTIVE: $\lambda x[\lambda y[\text{compose}'(x)(y)]]$

cake

Argument structure: $\lambda x[\text{cake}'(x)]$

Qualia structure:

TELIC: $\lambda x[\lambda y[\text{eat}'(x)(y)]]$

AGENTIVE: $\lambda x[\lambda y[\text{bake}'(x)(y)]]$

car

Argument structure: $\lambda x[\text{car}'(x)]$

Qualia structure:

TELIC: $\lambda x[\lambda y[\text{drive}'(x)(y)]]$

AGENTIVE: $\lambda x[\lambda y[\text{manufacture}'(x)(y)]]$

Every semantic predicate used in the lexicon is associated with a set of selectional restrictions on its arguments, which may filter out some unacceptable interpretations of a given phrase. We shall return to this point later.

2.4 Genitive 's

As for the genitive 's, our hypothesis proposes the two representations for the meaning of 's shown in (3.a) and (3.b):

(3) a. *Constructional interpretations*

$$[G\ s] = \lambda P[\lambda R[\lambda P[\lambda P(\lambda u[\exists x[\forall y[R(u)(y) \leftrightarrow y = x] \& P(x)]]]]]]$$

$\langle\langle e,t \rangle, t \rangle, \langle\langle e, \langle e,t \rangle \rangle, \langle\langle e,t \rangle, t \rangle \rangle$

b. *Control interpretation*

$$[G\ s] = \lambda P[\lambda Q[\lambda P[\lambda P(\lambda u[\exists x[\square y[[\text{control}'(u)(y) \& Q(u)] \leftrightarrow y = x] \&$$

$$P(x)]]]]]$$

$\langle\langle e,t \rangle, t \rangle, \langle\langle e,t \rangle, \langle\langle e,t \rangle, t \rangle \rangle$

The representation in (3.a) is designed to take care of all the constructional interpretations regardless of whether the actual head noun of the full genitive construction is relational or not. As it stands, the representation in (3.a) accepts only relational nominals, so this raises the question of how to get at the information in the Qualia roles. In order to achieve this, we need a set of meaning shifting operators. The technical details of this have been worked out in (Vikner & Jensen 2002).

The representation in (3.b) takes care of the control interpretation. The most important differences between these two representations is that while (3.a) shows that a relation is picked up from the head nominal, corresponding to the R -variable of (3.a), in (3.b) we can see that the control'-predicate is already present as a constant in the lexical entry for 's, and ready to receive its arguments. Furthermore, the representation in (3.b) shows - by the variable Q - that in this case the head nominal has to be a one-place predicate.

2.5 Summary of the Theory

The predictions which our analysis makes about the possible semantic interpretations of pre-nominal genitive constructions, are summarized in the table in (4):

(4) *Possible semantic interpretations of pre-nominal genitive constructions*

N ₂		N ₁	Interpretation
Relational		Meets selectional restrictions of N ₂ -relation	Inherent relation
Relational or Non- relational	Marked as possible part in Constitutive role of N ₂	Meets restrictions on whole in Constitutive role of N ₂	Part-whole relation
	Artefact	Meets selectional restrictions on agent position in N ₂ 's agentive role	Producer relation
	“Controllable object”	Animate	Control relation

3 Producer Interpretations

We now address the subject proper of this paper, that is, whether the producer interpretation is a semantic interpretation – as we claim - and not just one out of infinitely many pragmatic ones. As mentioned earlier, the methodological basis of our theory is the availability of default readings of linguistic constructions. Thus, since the producer relation seems to come readily into play as a default reading in examples such as (5),

- (5) *Ann's picture*
Ann's poem
Ann's cake

This – to our minds - is a strong indication that the producer relation is indeed semantic. A closer look at the kinds of examples allowing the producer interpretation leads naturally to the hypothesis that artefact-denoting nouns in general make the producer interpretation available when occurring as head nouns of genitive constructions.

3.1 Problems for the Hypothesis

In this section we take up some evidence which seems not to support the hypothesis that the producer relation is semantic. There are four problems, which we shall mention in turn, and in section 3.2 we present our proposals for dealing with these problems.

Problem 1

There are numerous examples where artefact-denoting nouns appear as head nouns of genitive constructions, but where the producer interpretation is not so readily available as one would expect if our hypothesis were correct, cf. examples such as (6):

- (6) *Ann's knife*
Ann's cigarette
Ann's beer
Ann's car

In all of these cases the control interpretation is clearly salient, and it is fairly hard to get the producer interpretation when these examples occur out of context.

Problem 2

An empirical investigation we have conducted comprising a total of 2333 examples gleaned from one fictional and one non-fictional text¹, shows a markedly smaller number of producer interpretations than of the three other semantic interpretations as is evident from the table in (7):

- (7) *Count of interpretations of pre-nominal genitive constructions*

Total		
Inherent	1666	71.4 %
Control	313	13.4 %
Part-whole	282	12.1 %
Producer	57	2.4 %
Pragmatic	9	0.4 %
Doubt	6	0.3 %
Total	2333	100 %

Problem 3

A third factor which might lead one to suspect that the producer interpretation is pragmatic rather than semantic, is the fact that usually genitive constructions with semantic readings correspond to a *have*-construction preserving the relation, cf. (8):

- (8) *Ann's sister* → *Ann has a sister* Relational Head Noun (sister-relation)
Ann's car → *Ann has a car* Non-Relational Head Noun (Control-relation)
The horse's tail → *The horse has a tail* Non-Relational Head Noun (Part-Whole-relation)
Ann's cake → *Ann has a cake* Non-Relational Head Noun (Control-relation. NB! The producer interpretation of *Ann's cake* is not possible with *Ann has a cake*)

¹ Jesse Byock: *Viking Age Iceland*. Penguin Books. 2001, pp. 1-203.

Margaret Drabble: *A Summer Bird-Cage*. (1963). Reprinted in Penguin Books 1967.

The important point here is that, since quite generally it is not possible to get pragmatic interpretations with *have*-constructions, the facts of (8) apparently favour treating the producer interpretation as pragmatic.

Problem 4

Finally, it should be mentioned that many of the examples we have found with the producer interpretation, exhibit contextual support, which does not prove that they are pragmatic, but which makes it somewhat harder to claim that the producer interpretation is purely semantic.

3.2 Proposed Solutions to the Problems

In spite of the apparent counterevidence presented in the four problems mentioned above, we would like to maintain that the producer interpretation is in fact semantic and made available by the lexical semantics of the head noun. In this section we explain how we think this position is tenable.

Re Problems 1 and 2

The reason why the producer interpretation (in examples such as those in (6)) is often not readily available and therefore not very frequent, we think, must be ascribed to the technological development, which has made individual production of most artefacts obsolete or rather uncommon. The reason why we think so, is that, if one considers different kinds of artefacts that are still usually produced by individuals, the producer reading of a genitive construction is available without contextual support. So, what we can observe, is a subset of artefact-denoting nouns which, when they appear in pre-nominal genitive constructions, do license a producer interpretation. The relevant subset of artefact-denoting nouns which license a producer interpretation, includes those enumerated in (9). Possibly, there are others, but perhaps not many others. We have not investigated this in detail yet:

(9) *Subclasses of nouns denoting artefacts commonly produced by individuals*

Informational content objects

Ann's words,
her theory, idea, plan
Joe's paper, computer program

Cooking objects

John's salad,
his pancakes

Social occasions

John's party, banquet, picnic
Babette's feast

Works of art and handicraft

Sue's watercolour,
her embroidery

Objects created in play

the children's snowman,
their sand castles

Artefacts created by animals
the beavers' dam.

It is not quite obvious what it means to say that the lacking availability of the producer interpretation in some cases is “due to the technological development.” One way to try to make sense of this claim in theoretical terms is to assume that over time the semantic predicates in the qualia description of certain nouns change, or new ones are added.²

For instance, in the case of the noun *car* we have proposed in (10) that the semantic predicate is *manufacture'*, the ontological type of whose producer argument, ‘manufacturing organization’, is supposed to represent something like a factory, company, firm or some such entity:

(10) *car*

Argument structure: $\lambda x[\text{car}'(x)]$

Qualia structure:

TELIC: $\lambda x[\lambda y[\text{drive}'(x)(y)]]$

AGENTIVE: $\lambda x[\lambda y[\text{manufacture}'(x)(y: \text{manufacturing organization})]]$

Unfortunately, this proposal still does not seem to help us get a producer interpretation, even when we have an NP like (11):

(11) *the factory's cars*

This example shows just another case of a salient control interpretation, and one still has to put in some effort to get the producer interpretation, which we claim to be there alongside the control interpretation. Probably, the reason why we still get the control interpretation as the most salient one, is that it is much more common for a factory to own cars than to produce them. That is, the normal relation between factories and cars is a control relation, not a producer relation.

If, on the other hand, we take a look at other objects, whose relation to companies or factories is not so commonly a control relation, it seems to us much easier to elicit a salient producer reading. Take for instance the examples in (12):

(12) *the factory's cloth*

the factory's tyres

the factory's valves

the factory's cigarettes

This lends support to our claim that, even though in many cases the producer interpretation of a pre-nominal genitive is not salient, it is certainly there as one of the four available semantic readings.

² We are not aware of any explicitly argued limit to the number of semantic predicates that may occur in a qualia role.

Re problem 3

The problem concerning paraphrasability by *have* is very tricky. But part of the explanation may be related to the fact that *have* seems to prefer static predicates to dynamic ones. So, whereas the semantic predicates involved in the examples in (13):

- (13) sister_of'
control'
part_of'

are all of a static nature, the predicates needed in order to get the producer interpretation have to be dynamic, as exemplified in (14):

- (14) build' (of houses)
manufacture' (of cars)
bake' (of cakes)
brew' (of beer)

Thus, the fact that we cannot get producer interpretations with *have*-constructions is probably due to its being a static verb requiring a static complement, and not to its inability to get pragmatic interpretations. This claim is supported by the fact that the producer relation is available with predicate genitives. Just like *have* the predicate genitive does not allow pragmatic interpretations. Compare, for instance, the pre-nominal genitive in (15) with the predicate genitive in (16):

- (15) *Ann's sky*

- (16) *That sky is Ann's*

Whereas (15) may have any number of pragmatic interpretations, like 'the sky Ann is always talking about', 'the sky Ann loves to paint', 'the sky Ann mentioned yesterday', etc., it is not possible to get such pragmatic interpretations with (16). But examples like (17):

- (17) *The cake over there is Ann's*
The biggest snowman is Fred's

show that the producer relation is readily available with the predicate genitive construction. So, this provides support of our proposal that the producer interpretation is semantic.

Re problem 4

As mentioned already, the fact that many of the examples we have found with a contextually supported producer interpretation is not proof that they are pragmatic genitives. If the producer interpretation were not semantic, we would have no explanation of the fact that in an example such as (18):

- (18) *The French believe they can, but one has only to read their books to mark ...*

the phrase *their books* is immediately interpretable as 'the books they have written'. And similarly with *their literature* in (19):

- (19) *Avoiding warfare, the Icelanders esteemed political flexibility and legal acumen, a cultural focus that is seen in their literature*

In neither of these examples does the surrounding context contain any information supporting the producer interpretation, and it is hardly likely that these instances of producer

interpretations derive from any common background knowledge of the French or the Icelanders. So, therefore, we posit that these must be semantic producer readings.

4 The Hebrew Construct State Construction

Recent evidence from Hebrew (Heller 2002) can be taken to provide indirect support of our view of the producer relation as being semantic rather than pragmatic in that the Hebrew Construct State Construction, which never allows pragmatic interpretations, does allow the producer reading, and thus seems to follow our proposal for the semantic readings of the English pre-nominal genitive neatly.

We therefore think that our claim is plausible, that, even though it is rarely a salient interpretation, the producer interpretation is indeed a semantic interpretation, and all nouns denoting artefacts do contain an agentive role in their qualia structure.

5 Conclusion

We have presented a theory of the pre-nominal genitive in English based on the claims that it is argument-only and polymorphic, allowing a control interpretation delivered by 's itself on the one hand, and three constructional interpretations on the other, namely the inherently relational, the part-whole, and the producer interpretations, respectively. We have zoomed in on some evidence which seems to raise problems for our claim that producer interpretations are semantic rather than pragmatic. In particular, 1) the apparent absence of producer interpretations with many nouns denoting artefacts, 2) the relatively low incidence of producer interpretations in our empirical investigation, 3) the lack of *have*-paraphrases of producer-interpreted genitive constructions, 4) the frequent occurrence of contextual support for producer interpreted genitive constructions. We have argued that in spite of this evidence our hypothesis seems to be tenable: Problems 1 and 2 arise due to the technological development, which has made most artefacts industrially rather than individually produced. As to problem 3, *have*, due to its static semantics, is only possible with static semantic predicates, which explains why it does not allow the producer interpretation, which is dynamic. Concerning problem 4, the occurrence of frequent supporting contexts does not preclude that the producer interpretation is semantic. The important point is that there exist clear examples that require a producer interpretation without being contextually supported. Finally, recent independent evidence concerning the semantics of the Hebrew Construct State Construction shows that it allows precisely the range of semantic readings we have proposed for the English pre-nominal genitive, including the producer interpretation, while it allows no pragmatic interpretations.

References

- Barker, C. & Dowty, D. (1992): "Non-verbal Thematic Proto-Roles." Available from <http://ling.ucsd.edu/~barker/Research/barker.dowty.nominal.proto-roles.pdf>.
- Borschev, V. & Partee, B.H. (2000): "Genitives, Relational Nouns, and the Argument-Modifier Distinction." *ZAS Papers in Linguistics*, 177-201.
- Briscoe, T., Copestake, A., & Boguraev, B. (1990): "Enjoy the Paper: Lexical Semantics via Lexicology." In H. Karlgren (ed.): *Papers presented to the 13th International Conference on Computational Linguistics*, vol. 2, 42 - 47. Helsinki.
- Copestake, A. (1992): *The Representation of Lexical Semantic Information*. University of Sussex at Brighton, Cognitive Science Research Papers 280.

- Copestake, A., Sanfilippo, A., Briscoe, T., & de Paiva, V. (1993): "The ACQUILEX LKB: An Introduction." In Briscoe, T., de Paiva, V., Copestake, A. (eds.): *Inheritance, Defaults, and the Lexicon*: 148-163. Cambridge University Press.
- Heller, D. (2002): "Possession as a Lexical Relation: Evidence from the Hebrew Construct State." In L. Mikkelsen & C. Potts (eds.): *Proceedings of the WCCFL 21*. Somerville, MA: Cascadilla Press, 127-140.
- Jensen, P. A. & Vikner, C. (1996): "The Double Nature of the Verb *have*." *LAMBDA* 21, 25-37. Department of Computational Linguistics. Copenhagen Business School.
- Pustejovsky, J. (1991): "The Generative Lexicon." *Computational Linguistics*, 17, 409-441.
- Pustejovsky, J. (1993): "Type Coercion and Lexical Selection." *Semantics and the Lexicon*, ed. J. Pustejovsky, 73-94. Dordrecht: Kluwer Academic Publishers.
- Pustejovsky, J. (1995): *The Generative Lexicon*. Cambridge, Massachusetts: MIT Press.
- Vikner, C. & Jensen, P. A. (2002): "A Semantic Analysis of The English Genitive. Interaction of Lexical and Formal Semantics." *Studia Linguistica* 56,2, 191-226.