IMAGINE NO POSSESSION: JOHN LENNON IN THE CONSTRUCT STATE

GABI DANON
Bar-Ilan University

1 Introduction

The Hebrew construct state (CS) has received considerable interest in the linguistic literature of the last few decades. Yet despite all the attention that this construction has received, there are still numerous poorly understood empirical facts regarding its basic distribution. One such fact, first explicitly discussed in the literature only recently (Rothstein, 2012, 2017), is the ungrammaticality of proper names (PNs) in CSs like the following:

(1)   a.   *xatul   sara*
        cat.M.SG   Sara
        Intended: 'Sara's cat'

    b.   *mexonit   gabi*
        car.F.SG   Gabi
        Intended: 'Gabi's car'

If the CS, like other genitive constructions, is basically the surface realization of a syntactic mechanism for DP recursion, one would not expect DPs headed by proper names (PNs) to be excluded from the embedded position. This apparent ban on PN embedding in CSs is also puzzling and unexpected from the point of view of most previous theoretical analyses of the semantics of CSs. Several analyses, such as Dobrovie-Sorin (2000, 2003), Heller (2002), and Doron and Meir (2013), argue that the head of a CS is interpreted as a function of type $< e, e >$. As such, PNs would be expected to be among the most prototypical DPs to which such a function could apply. Hence, the ungrammaticality of examples like those in (1) is unexpected not only on syntactic grounds, but also on semantic ones.

The facts are actually more complex. In contrast to the examples above, where PN embedding is ungrammatical, the embedding of a PN in a CS is perfectly grammatical in (2).

(2) a. *miškafey   jon lenon*
         glasses.M   John Lennon
         'John Lennon glasses'

     b. *oyvey   stalin*
         enemies.M   Stalin
         'Stalin's enemies'

What is at stake, hence, is both a descriptive and a theoretical problem. Descriptively, we should be able to characterize the conditions in which a PN can, and those in which it cannot, be embedded in a CS. And theoretically, we should be able to explain *why* a PN is blocked in the latter case.

This paper is devoted to clarifying the empirical status of PNs in CSs and the conditions under which such embeddings are (un)grammatical. We argue that, in fact, there is no ban on PNs in CSs at all, and that the observed effects are epiphenomenal to several other constraints. Most importantly, it is claimed that the CS in contemporary Modern Hebrew (MH) [1] is highly constrained in its productivity, not only when the embedded phrase is a PN, but also when other referential DPs are involved. This contrasts with *modificational* CSs, where the embedded nominal is interpreted as a kind modifier; such constructs are highly productive. Hence, going beyond the observation in works such as Doron and Meir (2013), that the head noun in a CS must be relational, this paper argues that one of the two major subtypes of this construction is subject to further lexical idiosyncrasies, and that a novel CS with an embedded referential DP is ungrammatical unless the head noun is lexically licensed as a CS head. Additionally, we argue that (true) possessive relations are not productive in contemporary Hebrew CSs. The contrast between PN embedding and the embedding of other definite noun phrases is then reduced to the availability of a modificational reading of the CS, where a non-modificational reading is not lexically licensed.

2 Background: Hebrew Genitives

2.1 The Genitive Alternation

The CS is one of three genitive constructions in Hebrew; the other two are the *šel*-genitive (SG), often referred to as *free genitive*, and the double genitive (DG); see e.g., Berman (1978), Ritter (1988, 1991), Ravid and Shlesinger (1995), and others. The three genitives are illustrated below:

(3) a. *tmunat   ha-yeled*
         picture.F.SG   the-boy
         'the picture of the boy'   (CS)

     b. *ha-tmuna   šel   ha-yeled*
         the-picture.F.SG   of   the-boy
         'the picture of the boy'   (SG)

     c. *tmunat-   o   šel   ha-yeled*
         picture.F.SG -POSS.M.SG   of   the-boy

---

[1] Modern Hebrew (MH) is of course subject to variation like any other language. Where it is important to distinguish between the language as it is used by most speakers today, and more archaic uses of the language (which would still be considered MH), we use the term 'contemporary MH'. Unless stated otherwise, the shorter term 'Hebrew' is used in this paper as equivalent to 'MH' without being specific about the contemporary/archaic distinction.

‘the picture of the boy’ (DG)

The factors that govern the alternation between these three genitive constructions are still, to a large extent, an open question: While conventional wisdom states that register is the dominant factor in native speakers' choice of genitive type, this seems to be an oversimplification which glosses over many important aspects of their distribution. Crucially, a purely register-based account would most likely rely on the assumption that the semantics of the three constructions is essentially the same, and that whatever is grammatical in one of the constructions is also grammatical in the others. While much of the syntactic literature has indeed (implicitly or explicitly) assumed that the differences are only in form and configurational structure, it is not entirely clear to what extent this assumption is indeed true. In the literature on the semantics of Hebrew genitives, it has in fact been repeatedly claimed that these constructions differ in their compositional semantics. Even though this paper deals only with the CS, it aims to contribute to this debate by discussing several constraints on CS formation which prevent free alternation due to the fact that one of the three constructions is ruled out by the grammar.

2.2   Preliminary Corpus Data

At a basic descriptive level, the three genitive constructions are far from equal in their distribution. According to the small-scale corpus study in Ravid and Shlesinger (1995), the CS is by far the most frequent, followed by SG, with DG accounting for only a small percentage of all genitive tokens. The proportion of CS out of all genitives in their study ranges from around 50-60% in spoken language to 80-90% in textbooks. Note that these numbers already pose a serious problem to a register-only account of the alternation, as they show that even in spoken MH the CS is the dominant type of genitive. The question is hence: What characterizes the cases where the CS is nevertheless avoided?

These numbers are in agreement with our own corpus data. In the current study, a random sample of around 130,000 genitive tokens was extracted automatically from the 2013 Hebrew Wikipedia corpus (Goldberg 2014), using the dependency parsed treebank produced by Goldberg's parser. Given that this treebank was produced automatically, the data extracted from it should be viewed with caution, and the discussion of this data in the current paper indeed assumes a certain (non-negligible) proportion of misanalyses. Nevertheless, manual inspection suggests that genitive identification by this parser is fairly accurate, and the large size of the sample should minimize the effect of misclassifications of specific tokens.

Table 1 summarizes the overall distribution of the 3 genitive types in this corpus sample.

GENITIVE TYPE	TOKENS
CS	105,580 (80.4%)
SG	17,205 (13.1%)
DG	8,543 (6.5%)
Total	131,328 (100%)

Table 1. Overall distribution of genitive constructions in the Hebrew Wikipedia corpus sample. CS = Construct State; SG = *šel* Genitive; DG = Double Genitive.

The distribution, however, differs when the embedded nominal is headed by a PN, compared to cases when the head of the embedded phrase is a common noun (CN). In the case of the former, we see a lower proportion of CS, a much higher proportion of DG, and a slightly higher proportion of SG, as shown in Table 2 below (where 'N2' stands for the lexical head of the embedded nominal).

GENITIVE TYPE	N2: PN	N2: CN
CS	15,189 (**69.7%**)	90,391 (**82.5%**)
SG	3,420 (15.7%)	13,785 (12.6%)
DG	3,176 (14.6%)	5,367 (4.9%)
Total	21,785 (100%)	109,543 (100%)

Table 2. Distribution of genitive constructions in the Hebrew Wikipedia corpus sample: Embedded proper name versus common name. PN = Proper Name; CN = Common Name; CS = Construct State; SG = *šel* Genitive; DG = Double Genitive.

Corpus data thus support the observation that the CS is not as common with PNs as it is otherwise; but at the same time, they also show that the CS is in fact still by far the most frequent type of genitive even with PNs. Our goal is thus to characterize the environments in which a PN is indeed blocked from appearing in a CS, without overgeneralizing this to block all instances of PN embedding in a CS.

2.3 CS as a Function

With these distribution patterns as a baseline, we can now ask whether there are any cases where there is systematically *no* alternation, or an alternation that involves only 2 of the 3 genitives. One major factor that has been claimed to constrain the availability of the CS is the lexical semantics of the head noun, and more specifically, whether or not this noun is relational. Rosén (1957) was perhaps the first modern linguist to argue explicitly that when the head noun is relational, the CS, as opposed to the SG, is strongly favored; and that nouns that are

ambiguous between a relational and a sortal reading receive only the former in the CS.[2] Later generative analyses that propose formal encodings of this intuition can be found in the works of Dobrovie-Sorin (2000, 2003), Heller (2002), and Doron and Meir (2013). According to these authors, the head of a CS must denote a function from individuals to individuals, and hence a CS is predicted to be ruled out with sortal nouns for which no such functional semantics is available. For instance, Doron and Meir (2013) cite the example of the noun *iša*, which may mean either 'woman' (sortal/non-relational) or 'wife' (relational). In a CS like (4) below, however, only the relational reading is available:

(4) *ešet   ha-cayar*
     wife   the-painter.M.SG
     'the painter's wife'   (adapted from Doron and Meir, 2013:(51))

Heller (2002) extends the functional semantics with the claim that possession is also a lexical relation (following Vikner and Jensen, 2002), and hence possessive CSs would not pose a counterexample to the proposed semantics of CSs.

2.4 Modificational CS versus Referential CS

One major exception to this formalization of the semantics of CSs, however, would be CSs in which the embedded nominal is interpreted as a non-referential kind modifier. Borer (2009) distinguishes between R-constructs (referential constructs; henceforth R-CS), where the embedded nominal is a referential DP, and M-constructs (modificational constructs; henceforth M-CS), where the embedded nominal is an NP which functions as a modifier.[3,4] The functional semantics discussed above might indeed be the correct analysis in cases like (5), which involve a referential embedded DP and are hence R-CS:

(5) a. *horey*   [*ha-yeled   še-   šavar   et   ha-xalon*]
         parent.M.PL   the-boy.M.SG   that   broke.3M.SG   ACC   the-window
         'the parents of the boy who broke the window'

     b. *menahel*   [*ha-xanut   ha-hi*]
         manager.M.SG   the-store.F.SG   the-DEM.F.SG
         'the manager of that store'

In contrast, the embedded nominal in the following M-CS examples is non-referential (possibly a bare NP rather than a full DP), and it bears the interpretation of a kind modifier:

(6) a. *mic*   [*perot   tropiyim*]
         juice.M.SG   fruit.M.PL   tropical.M.PL
         'tropical fruit juice'

     b. *texnay*   [*tanurim   ve-   mekarerim*]
         technician.M.SG   oven.M.PL   and-   refrigerator.M.PL
         'oven and refrigerator technician'

---

[2] Rosén (1957) did not use these terms, nor did he attempt to offer any formalization of these observations. Nevertheless, his claims and generalizations are surprisingly similar to those found in theoretical generative works half a century later.

[3] For a discussion of genitives as arguments and as modifiers, see Partee and Borschev (2003).

[4] Some works refer to these as unlexicalized compounds (see e.g., Berman and Ravid, 1986; Clark and Berman, 1987; Ravid and Shlesinger, 1995). For a discussion of the distinction between compounds and the CS, see Borer (1988, 2009).

As shown by Borer, M-CSs (M-constructs) bear major similarities to compounds, but they are nevertheless distinct from (lexicalized) compounds, and should be analyzed as fully compositional syntactic constructions. For instance, as seen in examples (6a-b) above, the embedded nominal in an M-CS may be syntactically complex, and its interpretation is fully compositional and transparent.

In a language like English, the equivalent of an M-CS would often have a surface form that distinguishes it from 'full' genitives: The modificational *peacock feather*, for instance, is overtly distinguishable from the genitive *a peacock's feather*. In Hebrew, however, the two have the same surface form, which renders many cases potentially ambiguous; hence, *nocat tavas* (lit. 'feather peacock') could be translated into either of the above English phrases, where it denotes either a kind of feather (i.e., the M-CS reading) or a feather belonging to a specific peacock (namely, the R-CS reading). As pointed out by Heller (2002), the M-CS reading is strongly favored when a bare noun is embedded, whereas an R-CS reading can be forced by adding the postnominal specificity marker *exad* (lit. 'one'):

(7) a. *nocat*   *tavas*
         feather.F.SG   peacock.M.SG
         'peacock feather' (preferred reading) / 'a peacock's feather'

     b. *nocat*   *tavas*   *exad*
         feather.F.SG   peacock.M.SG   one.M.SG
         'a (specific) peacock's feather'

2.4   Register: M-CS versus R-CS

Given this distinction between M-CS and R-CS, it is important to note what seems to be a fairly strong correlation with the use of CS in different registers. As mentioned above, it is sometimes assumed that the CS is used mostly in high register Hebrew, while informal language overwhelmingly favors the SG. One major shortcoming of this view is that it fails to note that the *modificational* CS is in fact extremely common in all registers, including fairly informal and colloquial MH. It has also been shown that this kind of CS is acquired and used spontaneously by relatively young children (Clark and Berman, 1987; Berman, 1988).[5] Taking this observation into account, a more accurate claim would be that *R-CS* use is *sometimes* seen as characteristic of formal or archaic MH. However, even this is not a categorical generalization, but rather an overall tendency, with numerous exceptions which often seem unpredictable. This observation will play a central role in the analysis proposed in this paper (§4).

2.5   PNs in M-CS

Looking once again at PNs, one might at first wonder whether the subclass of *modificational* CSs is even relevant when the embedded nominal is a PN. Since PNs are usually considered to be among the most prototypical referential DPs, we might expect them to occur in R-CS but not in M-CS. But in fact, there exists a class of M-CS in which the modificational nominal is a PN. These are illustrated in (8) below:

(8) a. *miškafey   jon   lenon*
         glasses.M   John   Lennon
         'John Lennon glasses' (but not: 'John Lennon's glasses')

---

[5] The literature on the acquisition of Hebrew genitives does not usually make explicit reference to the M-CS/R-CS distinction, but it seems that the majority of examples cited in these works are of the former type.

b. *kol*   *donald*   *dak*
   voice.M.SG   Donald   Duck
   'Donald Duck voice' (but not: 'Donald Duck's voice')

Some additional examples include *ragley maradona* 'Maradona legs'; *tisroket elvis* 'Elvis haircut'; *bdixat david levi* 'David Levy joke'; *xiyux jek nikolson* 'Jack Nicholson smile'; etc.

In the examples above, the embedded PN serves as a kind modifier. For instance, (8a) denotes a kind of glasses; (8b) denotes a kind of voice; etc. Interestingly, the R-CS interpretation is *not* available in these cases: (8a) cannot be interpreted as a pair of glasses that belonged to John Lennon himself, (8b) does not mean *Donald Duck's voice*, etc.[6] Thus, despite the general referential properties of PNs, embedding them in a CS sometimes blocks the referential reading, leaving only a modificational reading.[7]

We thus have a somewhat surprising distribution, where a PN cannot get its typical referential reading, but only a modificational one. This observation points to a possible generalization regarding the degree of productivity of R-CSs, as opposed to M-CSs. This is discussed in the next section.

3 Hypothesis 1: Only M-CS

3.1 Hypothesis: Only M-CS Is Productive

The possibility of embedding PNs in M-CSs but not in R-CSs, as discussed above (see §2.5), suggests a straightforward hypothesis, which we consider but eventually reject:

Hypothesis 1: R-CS is no longer a productive construction in MH.

Such a hypothesis would fit naturally with the observation made earlier regarding CSs and register: If R-CS were productive in earlier stages of the language, but is in the process of losing its productivity, we would indeed expect it to be associated with formal, literary, or archaic language use more than with informal usage. M-CS, on the other hand, is still fully productive, and hence not associated with any particular register. Furthermore, such a hypothesis would mean that frozen instances of R-CS taken from, e.g., Biblical Hebrew, such as examples (9a-b) below, do not pose a problem. The claim is not that R-CS is ungrammatical, but merely that it may not be systematically used for the formation of novel, *unstored* phrases.

(9) a. *ešet*   *lot*
        wife   Lot
        'Lot's wife'

    b. *sulam*   *ya'akov*
        ladder.M.SG   Jacob
        'Jacob's ladder'

In testing Hypothesis 1, we should thus be careful not to use such stored examples as evidence, but rather focus only on unstored ones.

One immediate objection to the hypothesis above is that it is much stronger than simply excluding productive uses of PNs in R-CS, and perhaps too strong. What this hypothesis entails is that *all* grammatical cases of R-CS, whether or not their embedded DP is a PN, are stored.

---

[6] For similar observations regarding Swedish compounds with embedded PNs, see Koptjevskaja-Tamm (2009).

[7] These kinds of modificational constructs with embedded PNs are productive and are not limited to a few frozen examples, but the formation of novel CSs of this type requires a highly salient and familiar property to be associated with the PN. Hence, they are common either with names of famous individuals, or in contexts where there is substantial shared knowledge to license such 'linguistic caricatures'.

Testing this with non-PN definites, however, can be somewhat challenging due to the phenomenon known as *definiteness spreading* (Borer, 1999; Dobrovie-Sorin, 2000, 2003; Alexiadou, 2005; Danon, 2001, 2008, 2010), whereby definiteness marked on the embedded nominal of a CS might in fact be interpreted on the CS as a whole; as a result, not every embedded nominal *marked* as definite is really a referential DP.[8] This is illustrated in the following example:

> (10)   *tmunat*   *ha-yeled*
>   picture.F.SG   the-boy
>   '[the boy]'s picture/the [boy's picture]'

Due to definiteness spreading, the definite article in *ha-yeled* is not necessarily interpreted in this location: (10) could, in principle, be interpreted as either 'the picture of the boy' or 'the picture of a boy', where the latter could be argued to be, in fact, an M-CS (i.e., 'boy' is a restrictive modifier for 'picture'). The empirical question is hence whether (10) also has the 'real' referential reading. Speakers' intuitive judgments are that it does. This can be supported by the fact that pronominal reference to the embedded nominal is indeed possible (11a), which, as argued by Borer (2009), is possible only with R-CS; as well as by the fact that it is possible to make the embedded phrase unambiguously referential by using a demonstrative (11b):

> (11) a. *tmunat*   *ha-yeled*   *ve-*   *ktovt-*   *o*
>   picture.F.SG   the-boy   and-   address-POSS.3M.SG
>   'the boy's picture and his address'
>
>   b. *tmunat*   *ha-yeled*   *ha-ze*
>   picture.F.SG   the-boy   the-this.M.SG
>   'this boy's picture'

Hence, there is reason to think that the proposed hypothesis is too strong, and that it incorrectly blocks not only the ungrammatical R-CSs, but also grammatical ones.

3.2   Productive Non-Modificational CS

Much of the discussion in this paper revolves around the issue of productivity, rather than of the grammaticality of isolated examples. Hence, it is important to establish the existence of entire *classes* of genitives for which the R-CS is *systematically* available. Many of these are fully grammatical not only with embedded nominals headed by common nouns, but also with PNs (Rothstein, 2012, 2017). The following is a non-comprehensive list of such cases:

- Derived nominals with thematic arguments: *sgirat ha-mis'ada* 'the closing of the restaurant', *nicul ha-macav* 'the exploitation of the situation', *kibuš mosul* 'the conquest of Mosul';

- Locative relations with embedded geographical PNs: *cfon london* 'the north of London', *merkaz sfarad* 'the center of Spain';

- Group membership: *tošav tel aviv* 'a resident of Tel Aviv', *ezrax carfat* 'a citizen of France', *dayarey ha-binyan* 'the inhabitants of the building';

- Genitives headed by a variety of other relational nouns: *sof/txilat ha-hofa'a* 'the end/ beginning of the show', *manhig/nesi/mefaked ha-irgun* 'the leader/president/

---

[8] It is sometimes assumed that in a definite-marked CS, both the embedded nominal and the CS as a whole are interpreted as definite. See Danon (2008, 2010) for the alternative claim that definiteness must be interpreted on either one of the two nominals; as well as the discussion of definiteness of an M-CS in Borer (2009).

commander of the organization'

It thus seems that Hypothesis 1 should indeed by weakened in order not to rule out these productive cases.

3.3   Rothstein (2012, 2017)

Alternatively, it might be argued that the basic intuition behind Hypothesis 1 could be maintained with a more sophisticated formal characterization of the semantics of the CS. Indeed, Rothstein (2012, 2017) argues that, with few exceptions, there is no real R-CS in the sense of a construction whose head is interpreted as a function of type $< e, e >$. Instead, she proposes that the embedded nominal in a CS must be an NP (rather than a DP), and that it must receive a predicational interpretation. This includes not only the straightforward cases of M-CS, but also CS with a definite-marked embedded nominal, which Rothstein argues can also be predicational. For CS headed by relational nouns, Rothstein (2017) proposes an incorporation-based analysis that makes it possible to maintain the hypothesis that the embedded XP is predicative, even in what looks like a prototypical R-CS. Proper names, however, are mostly excluded under this analysis due to this combined syntactic/semantic constraint on the embedded XP.

Noting that there are, however, classes of CS which do allow PN embedding, Rothstein proposes specialized explanations for the possibility of PN embedding in nominalizations, which are claimed to be licensed thematically; as well as for geographical names, which are argued to denote locative predicates. Additionally, PNs in 'author CS' like the following example are argued to have only a modificational (predicative) reading:

(12)   *širey*   *le'a*   goldberg
         poems.M   Lea   Goldberg
         'Lea Goldberg's poems'

According to Rothstein's analysis the embedded PN in this example denotes a property (roughly, 'authored by Lea Goldberg'), and hence the grammaticality of such examples does not constitute a counterexample to her analysis.

3.4   Problems with Rothstein's Analysis

In what follows we briefly discuss several potential problems that arise from the attempt to reduce all cases of CS to predicative modification.

3.4.1   Quantified XPs

Rothstein's predicational analysis predicts that the embedded phrase cannot be quantified, and indeed Rothstein provides several examples where a quantifier is impossible. This generalization, however, is too strong; the following examples from the Wikipedia corpus involve perfectly grammatical quantified genitives XPs:

(13) a.   *mefaked*   *kol*   *tayeset*
          commander.M.SG   every   squadron.F.SG
          'the commander of every squadron'

     b.   *matrat*   *kol*   *hesder*
          purpose.F.SG   every   arrangement.M.SG
          'the purpose of every arrangement'

c. *gova*   *kol*   *migdal*
    height.M.SG   every   tower.M.SG
    'the height of every tower'

It is thus clear that the semantics of the CS should not exclude *any* quantification of the embedded nominal. It should also be noted that the examples above are headed by functional nouns. If this is the typical environment where such quantified CSs are possible, then the analysis should be sensitive to the properties of the head noun and not only to those of the embedded XP.

3.4.2   Author CS: Both M-CS and R-CS

The second objection to a unified modificational/predicative analysis of the embedded nominal in a CS is that an author CS, like the one mentioned in (12), is in fact ambiguous, where only one of the two interpretations seems to have the properties that would be expected on the basis of Rothstein's analysis. To show this, we start by revisiting the properties of the definiteness spreading phenomenon observed in CSs. As Hebrew displays morphosyntactic sensitivity to definiteness in terms of definiteness agreement on attributive adjectives, and in terms of differential object marking, definiteness spreading can be diagnosed without having to rely on interpretation. The basic generalization is that a CS displays morphosyntactic definiteness iff its embedded nominal is definite (see e.g., Danon, 2008, 2010, and references cited there). Hence, in the following example, the CS must be preceded by the accusative object marker *et* (which is used only with definite objects); and an adjectival modifier of the entire CS must show definiteness agreement:

(14)   *ha-katav*   *ri'ayen*   *(et)*   *manhig*   *ha-miflaga*   *(ha-)zaken.*
        the-reporter.M   interviewed.3M.SG   ACC.   leader.M   the-party.F   the-old.M.SG
        'The reporter interviewed the old party leader.'

One exception, however, involves M-CS where the embedded nominal is a PN, like the ones discussed in §2.5. Example (15a) shows that a PN on its own, unsurprisingly, functions as a definite. (15b), on the other hand, shows that definiteness does *not* spread to the entire CS in a modificational structure:[9]

(15) a. *ba-* *sirton*   *ro'im*   *(et)*   *jon lenon*   *(ha-)cair.*
         in.the-video.M   see.PRES.PL   ACC.   John Lennon   the-young.M.SG
         'The video shows the young John Lennon.'

     b. *hi*   *hexlita*   *liknot*   *(*et)*   *miškafey*   *jon lenon*   *(*ha-)xadašim.*
         she   decided.3F.SG   buy.INF   ACC.   glasses.M   John Lennon   the-new.M.SG
         'She decided to buy new John Lennon glasses.'

With this background, we now see that (12) is in fact ambiguous between the modificational reading, where no definiteness spreading takes place, and the R-CS reading where it does, as demonstrated in (16) below.[10]

(16)   *hi*   *ohevet*   *(et)*   *širey*   *le'a*   *goldberg*   *(ha-)mukdamim.*
        she   loves.F   ACC.   poem.M.PL   Lea   Goldberg   the-early.M.PL

---

[9] In the pragmatically less plausible R-CS reading of 'John Lennon's glasses', definiteness spreading does occur, and hence both *et* and *ha-* would be required for that reading.

[10] Rothstein (2017) provides a similar example where definiteness does spread, and notes that when preceded by a cardinal numeral an indefinite construal is also acceptable. As (16) illustrates, we do not think that a cardinal is necessary for this option to exist, and we assume that there is true, systematic, ambiguity in such constructs.

'She loves the early poems of Lea Goldberg/early Lea Goldberg poems.'

The optionality of *et* and *ha-* correlates with an interpretational difference, where the indefinite CS (lacking *et* and *ha-*) gives rise to a kind reading ('Lea Goldberg poems'), and the definite CS gives rise to a specific object reading ('the poems of Lea Goldberg').[11] This strongly suggests that the semantic interpretation proposed by Rothstein covers only one subtype of CS rather than all constructs.

3.4.3 Geographical Names

Similar arguments can also be applied to the case of geographical PNs embedded in CSs. Rothstein (2017) proposes that these are in fact predicative NPs, and hence they fit into the overall semantics of the CS. The question is whether a CS with an embedded geographical PN gives rise to a single kind of structure (and associated interpretation). As the following examples show, definiteness spreading is either obligatory or optional in this case, where spreading correlates with an interpretational difference:

(17) a. *ani   mexapes   *(et)   merkaz   london   *(ha-)amiti.*
I   seek.PRES.SG   ACC.   center.M.SG London   the-real.M.SG
'I'm looking for the real center of London.'

    b. *ha-itona'im   kiblu   (et)   degel   mauritanya   (*ha-)xadaš.*
the-reporters.M   received.3M.PL   ACC.   flag.M   Mauritania   the-new.M.SG
'The reporters received the new flag of Mauritania/a new Mauritania flag.'

The CS in (17a) is headed by a strictly functional noun, and shows the definiteness spreading and interpretation expected from an R-CS with an embedded definite. (17b), on the other hand, allows either the functional reading ('flag-of') or a modificational reading (a flag of a certain kind), with definiteness spreading taking place in the former case and not in the latter. This once again suggests that there is more to CS than what a *uniform* predicational analysis would predict.[12] In §4 we propose a totally different kind of explanation for why geographical PNs pattern differently from human-denoting PNs.

3.5 Intermediate Summary

To summarize, it seems that trying to derive a ban on PN embedding from a general unavailability of an interpretation strategy for CS with an individual-denoting DP is too extreme and leads to a serious undergeneration problem, where a wide range of grammatical constructs are falsely predicted to be impossible, and many that are predicted to have only a single (predicational) analysis have in fact a referential one, either in addition to, or instead of, a predicational one. We hence reject this line of analysis, and turn instead to address the question of what restricts or constrains the distribution of R-CSs. In other words, we assume that true R-CSs do in fact exist alongside M-CSs, but hypothesize that the former are not as freely available as one might expect.

---

[11] It seems that while all speakers easily accept the definite/object reading, some judge the modificational/kind reading as somewhat borderline. See also Rothstein (2017).

[12] Rothstein (2017) discusses at length the patterns of definiteness spreading and interpretation associated with an embedded geographical PN, but the discussion revolves almost entirely around the group membership noun *tošav* ('resident'), which triggers what may be described as a (weak) anti-uniqueness presupposition (i.e., it is normally presupposed that a city has more than one resident). This kind of noun is exceptional when it comes to definiteness spreading (Danon, 2008, 2010), and it is left open to what extent the same analysis can be extended to CSs headed by functional nouns, like the ones in (17).

4 Hypothesis 2: Reduced Productivity of R-CS

4.1 Relational Nouns: A Necessary and Sufficient Condition?

As mentioned earlier, R-CSs have been claimed by multiple authors to be restricted to head nouns that are relational or functional, as well as to possessive relations. This generalization blocks, for instance, the possibility of an R-CS involving a contextual relation. Heller (2002) illustrates this with the example *horey ha-psixologit* ('the psychologist's parents'), which only has the *lexically relational* reading in which the psychologist is the daughter, and cannot have a reading in which someone else's parents are *contextually* related to the psychologist. In this section, however, we argue that this only imposes a necessary condition, but not a sufficient one, for R-CS formation. Specifically, we show that not every relational noun is equally acceptable as a CS head.

4.2 Lexical Gaps and Idiosyncrasies

If being headed by a lexically relational noun was a necessary and sufficient condition for R-CS formation, one would expect sets of relational nouns that have similar meanings to pattern similarly in their acceptability as CS heads. This prediction is not borne out. Consider for example the nouns *tmuna* ('picture') and *ciyur* ('painting'), both of which can be claimed to be relational, with the depicted entity serving as their lexically selected argument. It is thus surprising to notice the contrast in acceptability between the often-cited example in (18a) and its parallel in (18b):

(18) a. *tmunat*   *ha-xamaniyot*
          picture.F.SG   the-sunflowers.F
          'the picture of the sunflowers'
      b. ??*ciyur*   *ha-xamaniyot*
          painting.M.SG   the-sunflowers.F
          'the painting of the sunflowers'

Similarly, Doron and Meir (2013) illustrate the relational noun restriction using the example in (4), repeated below in (19), which is headed by *iša* ('wife'). This CS allows only the relational reading of the head noun despite the fact that outside the CS context, this noun is ambiguous between the relational reading of 'wife' and the sortal reading of 'woman':

(19)   *ešet*   *ha-cayar*
          wife   the-painter.M.SG
          'the painter's wife'   (adapted from Doron and Meir, 2013:(51))

But this example does not generalize to other head nouns. Replacing the head noun with the Hebrew words for *girlfriend* or *lover*, for instance, would render the CS totally ungrammatical:

(20)   **xavrat/me'ahevet*   *ha-cayar*
            girlfriend/lover.F.SG   the-painter.M.SG
          Intended: 'the painter's girlfriend/lover'

Similarly, the parallel CS with the genders reversed is ungrammatical (or at least highly degraded):[13]

---

[13] Most other kinship nouns are similarly degraded as CS heads in contemporary MH; overall, it seems that CS headed by a kinship noun is only acceptable in highly formal or archaic Hebrew, or in fixed contexts such as *axey ha-kala* (lit. 'siblings the-bride') often used in wedding contexts.

(21)   ??/*ba'al   ha-cayeret
             husband   the-painter.F.SG
             Intended: 'the (female) painter's husband'

Furthermore, it is well known that the word *ba'al* is in fact ambiguous between the meanings of 'husband' and 'owner', both of which are clearly relational; and while the former reading seems to be impossible as a CS head (21), the latter is perfectly acceptable (22):

(22)   *ba'al*   *ha-toyota*   *ha-levana*
             owner.F.SG   the-Toyota.F.SG   the-white.F.SG
             'the owner of the white Toyota'

In short, it seems that there are no simple semantic generalizations that can capture the exact set of nouns that can head an R-CS. At most, semantics imposes an upper bound on the availability of R-CSs.

4.3   Storage and Frequency Effects

Another aspect of the idiosyncratic and not fully productive nature of R-CS formation is the existence of storage effects. As noted in §3.1, older stages of the language seem to have allowed the use of a much broader range of nouns as CS heads than does contemporary MH. As a consequence, educated speakers readily accept as grammatical familiar phrases from Biblical Hebrew and other well-known sources, such as the CSs cited in (9). But when trying to generalize such examples with alternative embedded nominals, it seems that there is a continuum in which acceptability is correlated with the likelihood of finding such constructs in high register texts. Similarly, when trying to generalize using different (but semantically similar) head nouns, the acceptability of the resulting CSs seems to correlate with the use of the head noun in high register Hebrew. These claims are somewhat speculative, and to the best of our knowledge, no existing works provide empirical evidence for such correlations. We thus leave this to be confirmed in future research.

4.4   Possession: How Productive Is It?

Going beyond lexically relational nouns, one more kind of relation that has been assumed to be possible in CSs is that of possession. This turns out to be one of the major areas in which the limited productivity of R-CSs can be demonstrated. If the R-CS were a fully productive genitive construction, possession should perhaps be the most obvious place in which this productivity would be visible, as possession, often taken to be the prototypical relation encoded by genitive constructions crosslinguistically, is rarely restricted lexically. The literature on Hebrew CS, however, provides almost contradictory descriptions of the status of possessive CS. On the one hand, individual examples of possessive CS are cited in numerous works; Heller (2002) further makes this explicit by arguing that CSs can indeed encode possession and that this is in line with the general semantics of the CS (Vikner and Jensen, 2002). Some other works, on the other hand, note that the CS is disfavored as a means of expressing possessive relations, with the SG being the dominant way of expressing possession in contemporary Hebrew (Rosén, 1957; Ravid and Shlesinger, 1995).

Crucially, if possession were indeed a productive relation in CSs, we would expect *every* noun denoting a concrete object to be able to combine with *any* human/animate possessor DP to form a possessive CS. The isolated examples cited in the literature do not provide evidence that *all* such combinations are acceptable, only that some are (and, in fact, it seems that the same head nouns are cited repeatedly, such as *bayit* 'house', *xulca* 'shirt'). We should thus

check for productivity by attempting to construct arbitrary possessive CSs with pairs of nouns that were not specifically selected on the basis of their ability to give rise to a valid construct.

Consider for instance a sample of 5 concrete nouns (the set of possessees), given in (23a), and a sample of 5 animate denoting nouns (the set of possessors), provided in (23b):

(23) a. *tapuz* 'orange', *iparon* 'pencil', *kumkum* 'kettle', *kadur* 'ball', *ec* 'tree';

b. *xatul* 'cat', *nehag* 'driver', *balšan* 'linguist', *tuki* 'parrot', *yalda* 'girl'

Note the following properties of the selected sets of nouns:

- None of these involves what could be considered lexical possession; an orange, for instance, is not inherently conceived as being owned by a possessor (in contrast to nouns like *house* or *shirt*, for which a strong association with a possessor is assumed).

- None of the possible combinations from these sets is likely to be stored or memorized as a high frequency possessive noun phrase.

- The English translations of these can be freely combined to form perfectly grammatical possessives, as seen in (24) below.

(24) a. *the cat's orange*
b. *the driver's pencil*
c. *the girl's tree*

For the sake of illustration, we provide only a handful of the full 5x5 set of possible combinations in (25). Crucially, for all the constructs below, grammaticality is borderline at best:[14]

(25) a. *tapuz   *ha-xatul*
orange.M.SG   the-cat.M.SG
Intended: 'the cat's orange'

b. ??*iparon*   *ha-nehag*
pencil.M.SG   the-driver.M.SG
Intended: 'the driver's pencil'

c. *ec   *ha-yalda*
tree.M.SG   the-girl
Intended: 'the girl's tree'

Note, however, that some of these are acceptable with a *modificational* (M-CS) reading. If, for instance, there is a special kind of pencil which is a 'driver's pencil', example (25b) would be grammatical with such a reading; this should not be confused with a true possessive R-CS. Modification of the embedded nominal with a restrictive relative clause might help disambiguate this and avoid the modificational reading:

(26)   *iparon   *ha-nehag*   *še-*   *hisi'a*   *oti*   *etmol*
pencil.M.SG   the-driver.M.SG   that   drove.M.SG   me   yesterday
Intended: 'the pencil that belongs to the driver who drove me yesterday'

---

[14] Speakers do not unanimously judge these as ungrammatical, and some native speakers seem to accept all of these to a certain extent. However, our informal observation is that even these speakers hesitate somewhat regarding such examples, whereas many cases of non-possessive R-CSs with relational head nouns are judged as grammatical without any doubt or hesitation by most speakers. Furthermore, even speakers who do accept such possessive CSs clearly associate this exclusively with high register MH. Pending systematic experimental testing, we thus conclude that these do not show the full degree of grammaticality that one would expect from a fully productive construction.

We thus conclude that possessive CSs are no longer productive in contemporary MH. In light of this, it is not surprising that many of the examples of possession cited in the literature on CSs involve *lexically* possessed nouns such as *bayit* ('house, home') or items of clothing, which are not representative of free, unlexicalized possession.

The lack of productivity of possessive CSs has, of course, consequences for the PN issue: Since PNs are often names of people, animals or organizations, many cases of CSs with an embedded PN might be such that the only non-contextual interpretation that could be considered is possession, which we now see is ruled out independently of whether the possessor is a PN. Indeed, some of the examples in Rothstein (2012) fall exactly into this category, and we can see that those would be ungrammatical (or at best marginal) not only with PNs but even with common nouns:

(27) a. *et*   *ya'el/ha-balšanit*
        pen.M.SG   Yael/the- linguist.F.SG
        Intended: 'Yael's/the linguist's pen'

   b. *sefer*   *fred/ha-balšan*
        book.M.SG   Fred/ the-linguist.M.SG
        Intended: 'Fred's/the linguist's book'   (adapted from Rothstein, 2012:(32))

Once we have excluded possessives from the discussion, we can now go back to the list of cases where we *do* find productive R-CSs and reconsider their status. In particular, we can now consider the possibility that geographical PNs, argument-taking nominalizations, and 'author' CS are not *exceptions* to a general (direct or indirect) ban on PN embedding in a CS, which require special explanations, but the opposite: These are the cases that require no special explanation, whereas the ungrammatical embeddings are the ones whose ungrammaticality requires an explanation. The main component of the proposed explanation is that the head noun of an R-CS must be *lexically licensed*, and such licensing is (productively) possible only with a subset of lexically relational nouns.

4.5   Corpus Data Revisited

The claim that possessive CSs are not productive in contemporary Hebrew has testable predictions regarding corpus data. At the most basic level, the prediction is that individual inspection of genitives in a corpus would not yield any examples of possessive CSs other than potentially frozen or lexicalized ones. Such a prediction, however, is somewhat hard to establish in a robust manner, as the lack (or scarcity) of examples in a given sample provides only weak evidence for the grammatical status of the construction. A better approach would be to show that the proportion of genitives that are in the CS (as opposed to the other 2 types of genitives) is significantly lower when a possession relation is involved than otherwise; in other words, we aim to show an association between possession and the choice of genitive type.

To establish this, we consider data from the Wikipedia corpus again. In order to be able to test for association within a fairly large sample, manual identification of possession relations is not a feasible option. Instead, we aimed to test a weaker generalization by using animacy (or, more precisely, the distinction between human- and non-human-denoting nouns) as a proxy, given that animacy is strongly correlated with possession, at least in one direction, namely: Non-human nouns are very unlikely to serve as possessors. If possession, which only occurs with animates, is incompatible with the CS, then there is a subset of the human cases which are categorically excluded from appearing in the CS, and hence the prediction is that we would find a significantly lower percentage of CSs with human nouns.

To test this prediction, a list of human-denoting nouns was extracted from the corpus, and subsequently all genitives were classified according to whether or not the head noun of the

embedded nominal is on this list. It should be noted that this process necessarily leads to some classification errors, for at least two reasons:

1. The list of human-denoting nouns is not exhaustive and does not cover *all* such nouns in the corpus; as a result, the set of 'non-human' nouns (= the complement set of the ones identified as human-denoting) includes a certain percentage of human ones as well.

2. Lexical ambiguities and parsing errors might lead to errors in both directions.

As a result, we should treat the numbers below cautiously, and, at a minimum, aim for a much higher level of significance than if the data were manually annotated.

The following table shows the distribution of the 3 genitive types when the embedded nominal is headed by a CN.

GENITIVE TYPE	N2: NON-HUMAN	N2: HUMAN	TOTAL
CS	85,733 (**83.5%**)	4,658 (**67.3%**)	90,391
SG	12,566 (12.2%)	1,219 (17.6%)	13,785
DG	4,322 (4.2%)	1,045 (15.1%)	5,367
Total	102,621 (100%)	6,922 (100%)	109,543

Table 3. Distribution of genitive constructions with embedded common nouns in the Hebrew Wikipedia corpus sample: [±human]. N2 = the lexical head of the embedded nominal; CS = Construct State; SG = *šel* Genitive; DG = Double Genitive.

We see that the quantitative prediction is borne out, with the percentage of CSs being much lower for human-denoting nouns than for other nouns. Note that due to the fact that the 'non-human' class contains some human nouns as well, as noted above, the actual contrast would probably be even sharper if this 'noise' were removed. We thus have confirmation for the prediction of the hypothesis that the CS is incompatible with possession. Furthermore, since the CS is nevertheless compatible with a variety of other relations that could involve animates, we still find a large proportion of CSs in the human group.

With this in hand, we may now return to PN embedding and ask to what extent this covers the apparent ban on PNs. Table 4 shows the distribution of genitives when the embedded phrase is a PN.

GENITIVE TYPE	N2: NON-HUMAN	N2: HUMAN	TOTAL
CS	13,830 (**76.8%**)	1,359 (**38.9%**)	15,189
SG	2,493 (13.9%)	638 (18.2%)	3,131
DG	1,676 (9.3%)	1,500 (42.9%)	3,176
Total	17,999 (100%)	3,497 (100%)	21,496

Table 4. Distribution of genitive constructions with embedded proper names in the Hebrew Wikipedia corpus sample: [±human]. N2 = the lexical head of the embedded nominal; CS = Construct State; SG = *šel* Genitive; DG = Double Genitive.

We see that indeed, when the PN is human, the percentage of CSs is radically lower than with other PNs. This supports the hypothesis that what is behind the apparent ban on PN embedding in CSs is in fact a ban on the expression of specific relations in the CS, rather than a categorical ban on PNs.

Finally, we must address the question of why the contrast between human and non-human genitives is much sharper with PNs than with CNs. At this point, we return to the distinction between M-CSs and R-CSs, and note that M-CSs are much more frequent with CNs than with PNs, which, as noted in §2.5 above, require a highly salient property to be felicitous in this construction. The numbers in Table 4 show the contrast for PNs, with which M-CSs are fairly uncommon, and hence we could speculate that these numbers are close to the numbers that we would get if only R-CSs were considered.[15] The CN numbers in Table 3, on the other hand, show a much weaker animacy effect due to the fact that the R-CS counts are highly 'diluted' in this case by numerous M-CS instances, where animacy does not constrain the formation of CSs. As an illustration, consider the following three constructs:

(28) a.   *simlat*   *yalda*
            dress.F.SG   girl
            'a girl dress/dress for girls'

      b.   *simlat*   *ha-yalda*
            dress.F.SG   the-girl.M.SG
            'the girl dress/dress for girls' / 'the girl's dress'

      c.   ??*simlat*   *sara*
             dress.F.SG   Sara
             'a Sara dress' (but not: 'Sara's dress')

In (28a), the M-CS reading is straightforward ('a girl dress/dress for girls'). In (28b), despite the definiteness marking on the embedded nominal it is still possible to get the same M-CS reading due to definiteness spreading, which renders the entire CS definite without the embedded nominal having to be definite or referential. (28c), on the other hand, cannot get a modificational reading unless used in a context where speakers have shared knowledge about a person named Sara whose dresses have a salient characteristic property[16]. Hence, this contrast in the availability of the M-CS analysis between (28b) and (28c) could be enough to account for the contrast between the numbers in Table 3 and Table 4.

What is still left as an open question, though, is why (28c) is not as acceptable as (28b) on its R-CS reading. Note that not all R-CSs display such contrasts:

(29) a.   *oyvey*   *ha-nasi/stalin*
            enemies.M   the- president.M.SG/Stalin
             'the president's/Stalin's enemies'

      b.   *ce'eca'ey*   *ha-melex/ferdinand*
            descendants.M   the- king.SG/Ferdinand
             'the king's/Ferdinand's enemies'

---

[15] The actual numbers would have probably shown an even sharper contrast if lexicalized constructs were excluded. Many names of places in MH, which are very common in the Wikipedia corpus, have the form of a CS where the embedded nominal is a PN: *kfar yehošu'a* (lit. 'Joshua's village'), *giv'at brener* (lit. 'Brenner's hill'), etc.

[16] As pointed out by Noa Brandel (p.c.), another kind of reading, which is quite productive, is the appositive reading in which the PN is given to a specific model/design, and the lexical head of the CS explicitly states the superset (in this example, *simlat sara* might be a dress design named 'Sara', and the full CS might be used in a catalog). We leave it as an open question to what extent the proposed analysis of M-CS with PNs can be generalized to these uses.

Note that the head nouns in (29) are unambiguously relational, as opposed to the head noun in (28), which is a sortal noun with an optional relational reading ('the dress worn by x'). We leave it for further research to determine to what extent this observation can be generalized, and whether or not the borderline status of (28c) can be reduced to the fact that its (non-possessive) reading is based on the (weak) relational reading of an otherwise sortal noun.

4.6   Intermediate Summary: Semi-Productivity

To summarize, there is evidence that the CS is not as productive as often thought. In particular, cases of R-CSs which involve a possession relation are mostly ruled out. This already accounts for the borderline (or ungrammatical) status of a significant number of the examples of embedded PNs cited in the literature. We now proceed to combine the various observations that we've seen into a single claim regarding the productivity of the CS.

5   Conclusions: M-CSs and Lexically Licensed R-CSs

5.1   R-CS as a Lexically Licensed Construction

Our conclusion partially overlaps with the claims of Heller (2002), and partially contradicts them: Like Heller, we argue that only lexical relations are possible in R-CSs, but we diverge from her by arguing that the behavior of possessives shows that they should *not* be analyzed as lexical relations. We thus arrive at a more constrained version of the hypothesis that an R-CS must be headed by a relational noun.

Beyond the lack of possessive relations, a second conclusion is that even for lexically relational nouns, the CS is only partially productive. Hence, having a relational noun as the construct head is a necessary condition but not a sufficient one. We thus propose the following:

> **Lexical constraint on R-CSs:** An R-CS must be headed by a lexically relational noun with a (stored) CS lexical entry.

This constraint thus reduces the ability to form novel R-CSs to the lexical status of the head noun, where general grammatical principles provide only an upper bound for this construction. This state of affairs reflects a process of language change, where a construction which was fully productive in earlier stages of the language is still available to speakers but with somewhat unpredictable levels of acceptability.

M-CSs, on the other hand, seem to be mostly unconstrained and fully productive, as long as a suitable modification relation is available. It is thus somewhat ironic that while the majority of the generative literature on Hebrew genitives has focused on R-CSs (most notably nominalizations), it is in fact the M-CS which is much more frequent and productive in actual usage than the R-CS. The contrast in productivity between these two subtypes argued for in this paper suggests that the difference between them is not merely a minor typological issue, but rather a more substantial difference in their grammatical status.

Going back to the status of PNs in a CS, we now conclude that their ungrammaticality is merely an epiphenomenon of the CS's shift towards a modificational construction. Possession, possibly the most prototypical genitive relation when the embedded nominal is animate, is no longer productive in CSs, and this already rules out the use of CSs in a substantial number of genitives with an embedded human PN. Additionally, relational nouns which might in principle take human PNs as their (non-possessor) argument are not systematically grammatical as heads of novel (unstored) R-CSs. This leaves us with M-CS as the only fully productive paradigm, but when it comes to PNs, this type of CS is constrained by the fact that using PNs modificationally is a marked operation which usually depends on having a salient property

associated with the PN. We thus have an account not only of the tendency of PNs in a CS to be degraded, but also of the fact that there is no categorical ban against such embeddings.

5.2   Consequences: Possession versus Lexical Relations

Going beyond the Hebrew CS, our analysis has broader consequences for the status of genitives in general. As opposed to Heller's (2002) claim, possession does not seem to have the same distribution as 'real' lexical relations. Hebrew thus does not support the view that possession is a lexical relation (Vikner and Jensen, 2002).

As to the essence of genitive constructions, there are two views that are sometimes (implicitly) reflected in the choice of terminology. According to one view, possession is the core relation of the genitive, and other relations are 'peripheral' extensions of an inherently-possessive construction (see e.g., the discussion in Nikolaeva and Spencer, 2010); this view is often reflected in the use of the terms 'genitive' and 'possessive' as being roughly equivalent and interchangeable. The alternative view sees 'genitive' as a purely structural or morphosyntactic notion, where the essence of a genitive is nominal recursion or embedding; while 'possessive' is a semantic concept, which is often expressed using the genitive but is nevertheless an entirely independent notion. In this paper we have argued that the CS is becoming a non-possessive genitive. Under the former view, which takes possession as the core of the genitive, this seems almost contradictory. We thus believe that the properties of the CS support the view that 'genitive' is a purely morphosyntactic notion which should not be confused with possession.

Why, then, does it often look like possession *is* the core of the genitive? We will merely speculate that there could, in principle, be three types of meaning associated with a nominal recursion construction:

- Lexical;

- Contextual;

- Systematic relations that apply to an entire class of nouns

Possession belongs to the third type: It is semantically predictable whenever the embedded nominal is animate. But this is not the only possibility for the third type of meaning, with kind modification being a competing option. In many languages, possession has 'won' this competition, with modification expressed otherwise. MH demonstrates, however, that a language may choose differently, especially if, like Hebrew, it has more than one genitive construction. Hebrew is thus in the process of restricting the CS to lexical relations and to modification, with contextual and possession relations being associated more with the *šel*-genitive. The fact that these are mere tendencies, rather than categorical mappings, must be kept in mind, as it is among the main reasons for why it is so hard to characterize the exact nature of the genitive alternation.

References

Alexiadou, Artemis. 2005. Possessors and (in)definiteness. *Lingua* 115.6:787–819.

Berman, Ruth A. 1978. *Modern Hebrew Structure*. Tel Aviv: University Publishing Projects.

Berman, Ruth A. 1988. Language knowledge and language use: Binominal constructions in Modern Hebrew. *General Linguistics* 28.4:261–285.

Berman, Ruth A., and Dorit Ravid. 1986. On the degree of lexicalization of *smixut* constructions. *Hebrew Computational Linguistics* 24:5–22.

Borer, Hagit. 1988. On the morphological parallelism between compounds and constructs. In *Yearbook of Morphology*, ed. Geert Booij, and Jaap van Marle, 45–65. Dordrecht: Foris.

Borer, Hagit. 1999. Deconstructing the construct. In *Beyond Principles and Parameters*, ed. Kyle Johnson, and Ian Roberts, 43–89. Dordrecht: Kluwer.

Borer, Hagit. 2009. Afro-Asiatic, Semitic: Hebrew [compounds]. In *The Oxford Handbook of Compounding*, ed. Rochelle Lieber, and Pavol Stekauer, 491–511. Oxford: Oxford University Press.

Clark, Eve V., and Ruth A. Berman. 1987. Types of linguistics knowledge: Interpreting and producing compound nouns. *Journal of Child Language* 14:547–567.

Danon, Gabi. 2001. Syntactic definiteness in the grammar of Modern Hebrew. *Linguistics* 39.6:1071–1116.

Danon, Gabi. 2008. Definiteness spreading in the Hebrew construct state. *Lingua* 118.7: 872–906.

Danon, Gabi. 2010. The definiteness feature at the syntax-semantics interface. In *Features: Perspectives on a Key Notion in Linguistics*, ed. Anna Kibort, and Greville G. Corbett, 143–65. Oxford: Oxford University Press.

Dobrovie-Sorin, Carmen. 2000. (In)definiteness spread: From Romanian genitives to Hebrew construct state nominals. In *Comparative Studies in Romanian Syntax*, ed. Virginia Motapanyane, 177–226. Oxford: Elsevier.

Dobrovie-Sorin, Carmen. 2003. From DPs to NPs: A bare phrase structure account of genitives. In *From NP to DP, Volume 2: The Expression of Possession in Noun Phrases*, ed. Martine Coene, and Yves D'hulst, 75–120. Amsterdam: John Benjamins.

Doron, Edit, and Irit Meir. 2013. Amount definites. *Recherches Linguistiques de Vincennes* 42:139–65.

Goldberg, Yoav. 2014. The Hebrew Wikipedia Dependency Parsed Corpus Ver.1.0. Bar-Ilan University.

Heller, Daphna. 2002. Possession as a lexical relation: Evidence from the Hebrew construct state. In *Proceedings of WCCFL 21*, ed. Line Mikkelsen, and Christopher Potts, 127–40. Somerville, MA: Cascadilla Press.

Koptjevskaja-Tamm, Maria. 2009. Proper-name nominal compounds in Swedish between syntax and lexicon. *Rivista di Linguistica* 21.1:119–148.

Nikolaeva, Irina, and Andrew Spencer. 2010. The possession-modification scale: A universal of nominal morphosyntax. Unpublished manuscript, SOAS/University of Essex.

Partee, Barbara H., and Vladimir Borschev. 2003. Genitives, relational nouns, and argument-modifier ambiguity. In *Modifying Adjuncts*, ed. Ewald Lang, Claudia Maienborn, and Cathrine Fabricius-Hansen, 67–112. Berlin: Mouton de Gruyter.

Ravid, Dorit, and Yitzhak Shlesinger. 1995. Factors in the selection of compound-types in spoken and written Hebrew. *Language Sciences* 17.2:147–79.

Ritter, Elizabeth. 1988. A head-movement approach to construct-state noun phrases. *Linguistics* 26:909–29.

Ritter, Elizabeth. 1991. Two functional categories in noun phrases: Evidence from Modern Hebrew. *Syntax and Semantics* 25:37–62.

Rosén, Haiim. 1957. *ivrit tova: iyunim be-taxbir* [*Good Hebrew: Studies in Syntax*]. Kiryat Sefer.

Rothstein, Susan. 2012. Reconsidering the construct state in Modern Hebrew. *Rivista di Linguistica* 24.2:227–66.

Rothstein, Susan. 2017. Proper names in Modern Hebrew construct phrases. *Folia Linguistica* 51.2:419–451.

Vikner, Carl, and Per Anker Jensen. 2002. A semantic analysis of the English genitive. Interaction of lexical and formal semantics. *Studia Linguistica* 56.2:191–226.