Processing morphological ambiguity: An experimental investigation of Russian numerical phrases

The Harvard community has made this article openly available. Please share how this access benefits you. Your story matters

Processing morphological ambiguity: An experimental investigation of Russian numerical phrases[1]

Ming Xiang[a], Boris Haarizanov[b], Maria Polinsky[c], and Ekaterina Kravtchenko[b]

University of Chicago, University of California at Santa Cruz[b], and Harvard University[c]

Abstract

Russian nouns in nominative and accusative numerical expressions appear in three different forms, depending on the numeral: nominative singular with the numeral 1, genitive singular with the paucal numerals 2-4, and genitive plural with all other numerals. Results from an acceptability judgment task and a self-paced reading task on Russian case/number marking provide support for a theory stating that the suffix used with paucal nouns is morphologically ambiguous. The ambiguity resolution process involving this suffix leads to extra processing cost, compared to the unambiguous suffixes in other numeral contexts (the number 1, and the numbers 5+). This would account for the additional processing time observed with the paucal nouns. The status of the form occurring with the paucal numerals has long been a challenging issue in Russian linguistics, and the new results add to the growing body of literature that makes use of experimental methods to address issues of linguistic theory and analysis.

1. Introduction

In Russian numerical phrases, the case of the noun co-occuring with the numeral varies depending on the numeral. A noun used with the numeral 1 has to appear in the nominative singular; with numerals 5 and higher, one finds the genitive plural, and the genitive singular is used with 2, 3 and 4 (so-called "paucal numerals"):

(1) a.   odin   kruglyj   stol   propal

          one   round.NOM.SG   table.NOM.SG disappeared.SG

     b.   tri   kruglyx   stol-a   propal-i

          three round.GEN.PL   table-GEN.SG disappeared-PL

     c.   sem´   kruglyx   stol-ov   propal-i

          seven   round.GEN.PL   table-GEN.PL disappeared-PL

          'One/three/seven round tables disappeared.'

In (1b), semantically speaking, the sentence refers to a quantity of the noun that is greater than one; unsurprisingly, both the modifying adjective and the matrix verb are marked as plural. What is surprising, though, is that the head noun itself is marked as genitive singular (GEN.SG.). As noted above, this pattern only appears with the paucal numerals two, three and four.

There are two major theoretical approaches to this morphosyntactic phenomenon, which on the surface does not seem to be an instantiation of a regular agreement pattern. One approach treats the unexpected marking as a case of syncretism due to accidental homophony or underspecification. Under this approach, the same phonological form spells out both the genuine genitive singular ending and the paucal suffix (which we will discuss in section 2). An alternative approach considers the surface phonological form to be mapped unambiguously to a single

underlying morpheme, genitive singular; under this approach, the observed syncretism is due to the underlying featural identity of the genitive singular ending and the paucal suffix. We review these different views in section 2.

In this paper, we address the unusual genitive singular suffix from an experimental point of view. Using acceptability rating and self-paced-reading paradigms, we examine native speakers' sensitivity to nominal forms in numerical contexts. Up to now, researchers have not studied processing profiles of the three different numerical environments. The way these environments are processed, however, has a bearing on our models of the interaction between morphology and phonology, and of the mapping from underlying features to surface phonological strings. An investigation of the processing of numerical phrases will also help us distinguish between competing theoretical accounts of Russian case/number morphology.

The paper is organized as follows. In section 2, we review the basic facts of Russian case and number morphology, and introduce the relevant theoretical proposals. Section 3 lays out the logic of our study and the experimental predictions made by different theoretical proposals. Section 4 presents our experimental tasks and their results, and section 5 is a general discussion of the results.

2. Case and number morphology in Russian

In this section, we present the background on Russian case and number morphology, and survey the existing approaches to the puzzling agreement marking pattern presented in (1).

2.1. The puzzle of the genitive singular

Case and number in Russian are realized morphologically on the noun as a synthetic suffix (which also contains declension class information).[2] In the current discussion, our main focus is

on nouns in a numerical context - namely, nouns that are preceded by a numeral, with or without an intervening adjective. The numerical phrase has the reading of a precise quantity.[3] Based on the surface marking of the noun in the numerical phrase, such phrases can be divided into three groups. When the numeral is 1, the noun following it typically appears in the (nominative) singular (with the exception of *pluralia tantum,* and some other cases which are not relevant here—see Corbett 1983, Babby 1973, Hahm 2006, Wechsler and Zlatic 2003), and the numeral shows gender agreement with the noun. When the numeral is a number from 2 to 4, the following noun acquires a suffix which is phonologically (and potentially underlyingly) identical to the genitive singular morpheme used in other contexts[4]. The numeral 2 (but not 3 or 4) shows gender agreement with the noun. For the numerals 5 and above, the following noun is marked as genitive plural, and there is no agreement in gender. The following examples illustrate these three patterns, for each gender (feminine, masculine, and neuter):

(2)   odin-Ø   mal´čik-Ø/ odn-a   devočk-a/   odnojablok-o

       one-MASC   boy-NOM.SG/one-FEM girl-NOM.SG/ one-NEUT   apple-NOM.SG

       'one boy, one girl, one apple'

(3)   tri   mal´čik-a/ tri   devočk-i/   tri   jablok-a

       three   boy-GEN.SG/   three girl-GEN.SG/   three   apple-GEN.SG

       'three boys, three girls, three apples'

(4)   šest´ mal´čik-ov/   šest´   devoček/   šest´   jablok

       six   boy-GEN.PL/   six   girl.GEN.PL/   six   apple.GEN.PL

       'six boys, six girls, six apples'

Numerical phrases starting with 2, 3, or 4, which we have referred to as paucal, are typically treated as plural for purposes of subject-verb agreement, as shown below (see also Ionin and Matushansky 2006).

(5)   dva   mal´čik-a   spjat/*spit

       two   boy-GEN.SG   sleep.PL/sleep.SG

       'Two boys are asleep.'

On the surface, the pattern in (3) is puzzling. Even though "three boys" is conceptually plural, and is able to trigger plural agreement on the verb, as in (5), the head noun "boy" is marked as genitive singular. When a numeral is not present, the genitive singular morpheme is only used to mark nouns that are indeed singular. The following examples show that genitive singular markings are only present for singular nouns, whereas plural nouns have a different suffix (e.g., GEN.PL.).

(6)   a. genitive of negation (complement of a negative existential predicate)

          net/ne okazalos´   mal´čik-a/mal´čik-ov

          be.NEG.PRES/not appeared.PAST.SG.NEUT   boy-GEN.SG/boy-GEN.PL

          'There is no boy./No boy appeared.'//'There are no boys./No boys appeared.'

       b. prepositional genitive

          {do, u, bez, za}   mal´čik-a/mal´čik-ov

          to, by, without, instead of   boy-GEN.SG/ boy-GEN.PL

          'to, by, without, instead of the boy(s)'

c. possessive genitive

sobaka mal´čik-a/mal´čik-ov

dog   boy-GEN.SG/boy-GEN.PL

'the boy's/boys' dog'

Next, the use of the genitive singular with the numerals 2-4 (and the genitive plural with higher numerals) is only visible when the numerical phrase is in the nominative or accusative (non-oblique positions):

(7) a. priexali   tri   mal´čik-a

arrived.PL   [three.**NOM** boy-GEN.SG]

'Three boys arrived.'

b. ja kupil   tri   apel´sin-a

1SG bought   [three.ACC   orange-GEN.SG]

'I bought three oranges.'

In this paper, we will only concern ourselves with contexts like (7a) where the numerical phrases are nominative. In what follows, we will refer to the numerals 2-4 as *paucal* numerals,[5] and to the underlying morpheme phonologically realized as the genitive singular.

Since Russian is characterized by extensive case syncretism, we used only masculine nouns in the experimental design, described in sections 3 and 4. For those readers that are not familiar with Russian declensional patterns, we present here the basic declensional types of Russian masculine nouns. Note that the form of the genitive plural has several possible exponents, of which we used only *-ov* and *–ej* (this was done to make sure all the nouns had a non-null genitive plural ending).

Table 1: Main declension patterns of masculine nouns in Russian

	Animate				Inanimate			
	SG		PL		SG		PL	
	Subtype I	Subtype II	Subtype I	Subtype II	Subtype I	Subtype II	Subtype I	Subtype II
Nom.	mal´čik-**Ø** 'boy'	stroitel´-**Ø** 'builder'	mal´čik-**i**	stroitel-i	apel´sin-**Ø** 'orange'	gvozd´-**Ø** 'nail'	apel´sin-**y**	gvozd-**i**
Gen.	mal´čik-**a**	stroitel´-**a**	mal´čik-**ov**	stroitel´-**ej**	apel´sin-**a**	gvozd´-**a**	apel´sin-**ov**	gvozd´-**ej**
Acc.	mal´čik-**a**	stroitel´-**a**	mal´čik-**ov**	stroitel´-**ej**	apel´sin-**Ø**	gvozd´-**Ø**	apel´sin-**y**	gvozd-**i**

2.2. Main theoretical approaches to the "paucal suffix"

There are two ways to approach the surface genitive singular marking in paucal contexts. One possibility is that the paucal suffix is phonologically realized identically to the genitive singular because it is, in fact, the same morpheme as the genitive singular suffix, specified for the case and number features [gen] and [+sg], respectively (e.g. Pesetsky 2007, 2010). Such an approach would treat these nouns as genitive singular, and the genitive singular suffix on these nouns as the same formal element that appears in the contexts of the prepositional and possessive genitives, and the genitive of negation.

The other possibility is to treat the suffix on nouns following paucal numerals as distinct, but syncretic with the genitive singular suffix, having the same phonological realization (e.g. Bailyn and Nevins 2008). Under such a view, the paucal suffix and the genitive singular suffix are formally distinct feature bundles, which happen to be realized by the same surface-ambiguous phonological string.

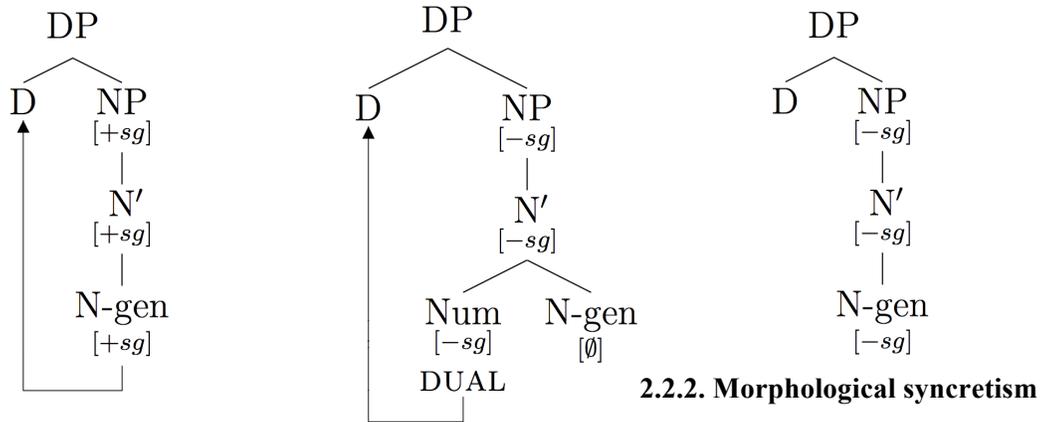Both conceptions of the paucal suffix have been proposed in the literature, as discussed below.

2.2.1. "What you see is what you get"

Pesetsky (2007, 2010) argues that genitive case is the default case on all nouns: i.e., Russian nouns enter the syntactic derivation already containing a default "gen" case feature. In addition, paucal nouns themselves are not directly marked by any number feature, and by default appear as singular in morphology. Crucially for our purposes here, this would mean that the surface suffix GEN.SG. reflects the underlying featural composition. We sketched the basic idea below (and in diagram (8) as well), which derives the correct case and number features for different numerical phrases. Paucal nouns look like singular on the surface because the [–sg] number feature is not realized as a bound morpheme attached to the nouns themselves; instead, the free-standing paucal numerals are treated as number features (dual, trial, quadral) that attach to the nouns (which are numberless) to pluralize the whole DP. The paucal numerals undergo a Num to D movement, and get the corresponding nominative case feature from D. These derivations lead to a final representation of "paucal$_{NOM}$ + adj$_{Nom. Pl}$ + N$_{Gen}$". For "1" and "5+" contexts, number features are realized as [+sg] and [+pl] pre-syntactically as bound morphemes on nouns, and numerals are real numerals rather than free-standing number features. For the numeral "1", nouns go through N to D movement, and hence inherit the nominative case from D, which overwrites the original genitive case on the noun, and results in a final NOM.SG. noun. For "5+" numerals, there is no N to D movement, and nouns stay as genitive (i.e., GEN.PL.).

(8) a. numeral 1   b. paucal numerals   c. numerals 5+

DP
D — NP [+sg]
N′ [+sg]
N-gen [+sg]

DP
D — NP [−sg]
N′ [−sg]
Num [−sg] — N-gen [∅]
DUAL

DP
D — NP [−sg]
N′ [−sg]
N-gen [−sg]

2.2.2. Morphological syncretism

Another family of accounts treats the paucal suffix as formally distinct from the true genitive singular. Bailyn and Nevins (2008) suggest that it has the following feature specification: [+nom, +paucal]. The [+paucal] number feature can be represented as [-singular, +augmented], and this easily accounts for the plural agreement observed on modifying adjectives and the main verb. One piece of evidence Bailyn and Nevins cite to support the [+paucal] feature analysis is that with some nouns in the paucal numeral contexts, the [+nom,+paucal] morpheme actually has a phonological form distinct from that of the genitive singular morpheme:

(9)   a.   bez   šága/* šagá

           without step-GEN.SG   'without a step'

      b.   tri   šag-á/*šág-a

           three   step-PAUCAL/*-GEN.SG   'three steps'

Consistent with this proposal, which suggests a paucal feature on the noun, Ionin and Matushansky (2004, 2006) also suggest that the paucal suffix only expresses a paucal case form, i.e. it only contains a case feature, and is different from the genitive. Under this type of account, Russian has a three-way distinction in marking its number morphology: singular, paucal, and plural, making it more similar to Polish, Czech or Serbo-Croatian (Franks 1994, 1995, 1998, 2002). The paucal suffix on nouns following paucal numerals is determined by the underlying

10

paucal feature on the noun. This paucal suffix happens to be syncretic with the suffix of genitive singular.

In a different version of the account which treats the paucal suffix as underlyingly distinct from the genitive singular suffix, Zaliznjak (1967) offers the following view: the paucal suffix is a special count form, which is outside of the nominal case paradigm. This puts this form on par with such fossilized Russian forms as *večerom* 'in the evening' (historically, the instrumental singular form of *večer*), *domoj* 'homeward' (historically the dative singular form of *dom* 'house, home'), or *bez sprosu* 'without permission.' These forms have the appearance of case forms, but they are not part of the general case paradigm of a respective noun. According to this view, nouns that bear the paucal suffix are simply count forms, which makes them essentially caseless, in terms of morphological case, which further means that they are neither nominative nor genitive. As far as their number is concerned, Zaliznjak suggests that they may be underspecified for number, which makes them compatible with numerals.

3. Processing Russian numerical phrases

The basic logic of our study is as follows: if the paucal suffix is featurally distinct from the genitive singular suffix, the corresponding phonological form is then an ambiguous input to the parser; namely, after encountering a noun that bears an exponent that can be mapped either to the paucal morpheme or to the genitive singular one, the parser needs to decide which morphological representation it is indexing. On the other hand, if no phonological syncretism is involved and the paucal and genitive singular suffixes are in fact the same morpheme, then the problem of ambiguity resolution does not arise. These two parsing possibilities, which arise from the two potential morphological representations of the suffix that appears on paucal nouns, are associated with different processing costs.

Before discussing the specific predictions that the two proposals make, we want to lay out our experimental assumptions with respect to sentence parsing and ambiguity resolution. First, we make the standard assumption that sentence processing is highly incremental, and the parser not only builds structures that satisfy the current input, but also actively predicts the upcoming structures based on the currently available evidence (e.g., Crocker et al. 2000, Ferreira 1996). This means that when speakers process the numerals, they should have already projected a DP structure that contains the numeral and other necessary pieces, such as the NP structure and the N-head. Second, with respect to ambiguity resolution, we assume that multiple underlying representations of an ambiguous input string on the surface are activated simultaneously, and the final disambiguation and selection of the appropriate representation is affected by the relative frequency of each representation and also the relevant context.

Although ambiguity resolution of syncretic inflectional affixes has not been discussed much in the literature, we can borrow some insight from the literature on lexical ambiguity resolution. Using a cross-modal priming paradigm, Swinney (1979) found that the possible meanings of an ambiguous word were all activated even in contexts that favor one of the interpretations. For instance, the noun "bug" is ambiguous between the meaning "an insect" and "a surveillance device." In a context that strongly biased toward one of the meanings, both meanings were activated. The same sort of effect was also found with partial lexical input that is temporarily ambiguous. Zwitserlood (1989) looked at the possible candidates the parser generates at different points of an auditorily presented word, which may cause temporary ambiguity. For example, in the word *captain*, the initial segment *cap* is consistent with either *captain* or *capital*. Before the audio stream becomes unambiguous, people activate both possible candidates, even if the context may be biased toward one of the two meanings. In addition to these findings, previous studies on lexical syntactic ambiguity have also shown that context does not inhibit the lexical

activation of multiple candidates of a syntactically ambiguous noun. Noun-verb homographs such as "park" are syntactically ambiguous between a noun and a verb interpretation. In a clearly syntactically biased context such as "John went to the park," it is very clear that the word after a determiner has to be a noun. In contexts like this, it seems that syntactic cues should absolutely block the lexical activation of the inappropriate interpretation; i.e. syntactic context should disambiguate the ambiguous word immediately. Somewhat surprisingly, contrary to this intuition, studies on N-V homographs have shown the opposite trend: despite the syntactic contexts that should clearly disambiguate the upcoming word, multiple senses of an ambiguous word are still activated (Tanenhaus, Leiman and Seidenberg 1979; Federmeier, Segal, Lombrozo and Kutas 2000; Lee and Federmeier 2009). For instance, in an ERP study, Lee and Federmeier (2009) found that in a sentential context, an NV homograph such as "watch" elicited a bigger sustained frontal negativity compared to an unambiguous word, even when context has clearly disambiguated the interpretation of the NV homograph word - such as "to watch," or "the watch." The extra processing cost from NV homographs, compared to the unambiguous control, suggests that both interpretations of the homograph were being processed.

All of the studies mentioned above looked at ambiguity resolution on open class words. It is known that close class words and inflectional morphemes are processed with a distinct time course from the open class ones. The time course difference could mean that functional morphemes and open class words are accessed through completely different routes. However, for our purposes here, a more critical question is the effect of context on functional morpheme ambiguities, i.e., whether ambiguities arise at all when contexts are biased towards only one of the possible representations. As far as we know, there is no discussion about this in the literature. While keeping in mind the processing difference between open and close classes, we want to use the current study as an opportunity to extend previous studies on lexical semantic ambiguity
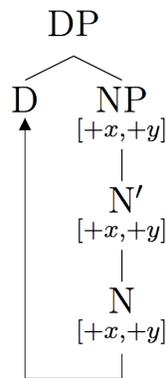
resolution to the functional morpheme domain. If contexts constrain ambiguity resolution in the same manner for lexical and functional categories, we can make different predictions with respect to the processing cost involved in processing the paucal suffix.

Consider first the account that treats the paucal suffix as featurally distinct from the genitive singular suffix. Recall that there are different proposals regarding the exact feature specification of the paucal suffix, but for current purposes, let's call it [+x, +y], where x is case, and y is number. Under this account, while processing the numerals 2, 3, or 4, speakers should have established a DP structure containing a numeral, and should also be anticipating a noun head which carries the features [+x,+y]. The structure that is constructed at this point is demonstrated in (10a). Furthermore, native Russian speakers also know that the [+x,+y] features will be realized phonologically in the same way that [+gen,+sg] is realized. When subjects get to the noun carrying a paucal suffix, the phonological form of the paucal suffix is perceived as ambiguous – it may represent a [+gen,+sg] morpheme, or a [+x,+y] morpheme. Both representations will be generated at the noun, and the contextual information (such as the fact that a paucal numerical phrase is being processed) facilitates the selection of the appropriate one. Note that, based on the discussion about lexical ambiguity resolution above, it is expected that this contextual information will not lead to disambiguation before the head noun is reached. The process of ambiguity resolution - namely, the processing load associated with the generation of multiple underlying candidate-representations, and the selection of the appropriate one - should impose a larger processing load on the parser.
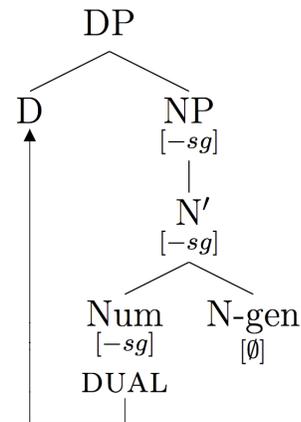
On the other hand, an account that treats the paucal suffix as largely identical to an underlying genitive singular form, such as Pesetsky's (2007, 2010), does not predict an increased processing load associated with ambiguity resolution in paucal contexts. At the numeral, an incremental parser again establishes a DP structure that contains a numeral and a noun head.

According to Pesetsky, this noun head is [+gen], and unspecified for number. The structure constructed at this point is shown in (10b). The suffix found on the head noun, in this case, does not present the parser with any kind of ambiguity, given that the phonological form of the suffix maps to a single underlying representation: [+gen]. Without the complications induced by ambiguous input, the parser should be able to incorporate the noun straightforwardly into the previously established DP structure (10b).

(10a)

DP
├─ D
└─ NP [+x,+y]
   └─ N′ [+x,+y]
      └─ N [+x,+y]

(10b)

DP
├─ D
└─ NP [−sg]
   └─ N′ [−sg]
      ├─ Num [−sg] DUAL
      └─ N-gen [∅]

In addition, the structure in (10b) involves some extra steps (additional syntactic operations) that could potentially bear on the processing output. In particular, since these syntactic operations are carried out at different positions within a DP (see (8) and the explanation about the derivation therein), they may induce extra processing cost, and such processing load will also occur at a different point in time. First, the Num-to-D movement is only present for paucal numerical phrases. At the numeral position, there is no extra syntactic operation involving the other two types of numerical phrases. This suggests that only for numerals in the paucal context, but not for numerals in the other two contexts, there should be some extra processing associated with head movements. In addition, at the noun position, only the "1" nouns undergo N-

to-D movement. Therefore we would expect some extra processing cost for the "1" context at the noun, but not the other two contexts.

In the experiment we will present in the next section, we used a self-paced-reading task to look at the processing cost related to nouns with the paucal suffix. Specifically, in the paucal context we examined the processing cost difference between grammatically marked nouns and ungrammatically marked nouns (e.g. a noun in a paucal context carrying a NOM.SG suffix). For the other two contexts, the numeral 1, and numerals 5 and higher (5+) we also examined the processing cost difference between the grammatical and ungrammatical forms. Since in these other two contexts, the grammatical form is unambiguous (i.e., NOM.SG and GEN.PL respectively), we expect to see a clear advantage of processing grammatical forms in these contexts. Namely, the ungrammatical forms should involve a bigger processing cost, and thus elicit a much longer reading time than the grammatical ones. These two contexts then, serve as our control conditions. If we see a clear contrast between grammatical and ungrammatical forms with regard to reading time in the contexts involving the numerals 1 or 5 +, but not in the paucal context, this would indicate that the processing of nouns in the paucal context is special. Specifically, if the paucal suffix is indeed underlyingly distinct from the genetive singular, then nouns with the paucal suffix may be more costly to process than the unambiguous grammatical forms in the other contexts. This would result in the absence of a significant difference between processing grammatical and ungrammatical forms for paucal nouns. This processing cost would not be predicted, however, by a proposal which claims that genuine genitive singular morphemes appear in paucal contexts.

4. Experiment

Below, we present the results from one experiment that included two experimental tasks. In one task, we collected native speakers' acceptability judgment data in an offline task. This allowed us to establish the basic judgments for the relevant Russian forms. Using the same group of subjects, we also collected their online response to (un)grammatical forms in a self-paced reading task. Self-paced reading is a widely used paradigm in psycholinguistics. Processing cost is indexed by the reading time on each region in a sentence. With this task we may expect to see divergence of the paucal contexts from the other two environments with respect to the processing difference between grammatical and ungrammatical forms.

4.1. Materials

The two tasks shared the same set of materials. There were 12 conditions, using a 3 x 4 design based on 2 independent variables. The first variable is the type of numerical context. Subjects read sentences that either have the number 1, the paucal numbers (2 to 4), or numbers above 4 (i.e. 5 or higher)[6]. The second variable we manipulated involved the degree of the featural match between the number-case marking on the noun, and the grammatical expectation in a given numerical context. For instance, following the number 1 in its nominative form, the grammatical form was a nominative singular noun. Therefore, a noun that is indeed marked as NOM.SG would fully match the grammatical expectations, with regard to both number and case features. However, a noun that is marked as NOM.PL would only match the grammatical expectation for case marking, and not the number feature. Similarly, a noun marked as genitive singular (here, by GEN.SG. we mean the phonological form normally associated with genitive singular in non-numerical contexts) only matches the grammatical version in its number feature, and not in the

17

case feature. Finally, a noun marked GEN.PL. does not match any feature of the grammatical version.

For each of the three numerical contexts, the nouns following the numeral varied in 4 possible ways, namely: full match, case match only, number match only, and no feature match. In order to avoid a parsing strategy that maps a numeral-noun sequence into a counting routine, we inserted an adjective between the numeral and the noun. We always used the grammatical form of the adjective. Similarly, for the matrix verb of the sentence, we only used the grammatical form, which matched the numerical context. In Table 2, we give an example of the experimental stimuli.

Table 2: Example Stimuli

PP	Numeral	Adjective	Noun	PP	Verb	PP
V xore in choir	odin one.NOM	malen'kij little.NOM.SG	mal'čik (boy.NOM.SG) *mal'čika (boy.GEN.SG) *mal'čiki (boy.NOM.PL) *mal'čikov (boy.GEN.PL)	v očkax in glasses	stojal stood.SG	vperedi vsex before all
	tri three.NOM	malen'kix little.PL.GEN	*mal'čik (boy.NOM.SG.) mal'čika (boy.GEN.SG) *mal'čiki (boy.NOM.PL) *mal'čikov (boy.GEN.PL)		stojali stood.PL	
	pjat' five.NOM	malen'kix little.PL.GEN	*mal'čik (boy.NOM.SG *mal'čika (boy.GEN.SG) *mal'čiki (boy.NOM.PL) mal'čikov (boy.GEN.PL)		stojali stood.PL	
'In the choir, one/three/five little boys in glasses stood in front of everybody.'						

We used a total of 60 experimental items, each appearing in the 12 conditions described above. These sentences were distributed into 12 lists, using a Latin square design, such that only one condition from each item appeared in any given list. Within the experimental items, the grammatical/ungrammatical ratio was 1:3. In addition, there were 108 filler sentences, which

included some ungrammatical sentences (their ungrammaticality was not related to numerical contexts). The overall grammatical/ungrammatical ratio in the entire stimuli set was 7:5. Our critical stimuli included 29 animates and 31 inanimates. In order to minimize case syncretism, we used only masculine nouns, given that they show the different cases required by the numerals most clearly (e.g. most feminine nouns exhibit identical nominative plural and genitive singular forms). We constructed different types of filler items in order to sufficiently disguise the critical items. Among the 108 filler sentences, 24 involve reflexive binding and 24 involve NPI licensing, and the remaining 60 were constructed randomly.

4.2. Participants and procedure

37 native Russian speakers participated in our study. 9 were tested in Moscow, and 28 were tested in the Boston area. All of them were adult speakers (age range from 19 to 45 years, mean=28 years, SD=6.6 years; 27 females, 9 males). Those participants who were tested in the US had been outside Russia for, on average, 4.5 years (range from less than 1 to 14 years, SD=3.9 years). For the same group tested in US, two left Russia as teenagers, and the rest left Russia as adults (age of immigration: range 13-40 years, mean=24years, SD=6 years). All subjects were reimbursed for their participation. They completed the self-paced reading task first. Sentences were presented using the Linger Software package (Rohde 2003) on a PC. Participants pressed the space bar in order to continue reading each sentence in a word-by-word fashion. One third of the sentences were followed by a yes-no question for the participants to answer; the question concerned the content of the immediately preceding sentence. Results from their accuracy data show that the average accuracy was 89%, with every participant having an accuracy rate of at least 80%. No one was excluded from the data analysis in the result section below.

After they finished the self-paced reading task, the subjects were asked to complete an acceptability-judgment task. 35 of the 37 people above finished the acceptability-judgment task. All of the stimuli in the judgment task, including the experimental items and fillers, were the same as in the self-paced reading (except that there were no yes-no questions in the rating task), but participants may have seen different item conditions in their reading and judgment tasks. This is because we made the 12 judgment lists using a Latin Square Design and distributed them evenly across our subjects; for the self-paced reading task, the data acquisition program we used, Linger, created the 12 lists automatically, and assigned to each subject a list in a random manner (but also evenly distributed across all subjects). Participants were tested individually, in a quiet room. Some were given a pencil and a printed copy of the sentences, and some filled out a spreadsheet version on the computer. The task instruction was to rate each sentence on a 1 to 7 scale, where 1 represented a completely unacceptable sentence, and 7 represented a fully acceptable one. The participants were instructed to make the judgments based on their intuitions, rather than any prescriptive rules acquired in classroom settings.

4.3. Rating task results

The average rating results are presented in figure 1.

[Figure 1] Acceptability Judgment

Acceptability Judgment
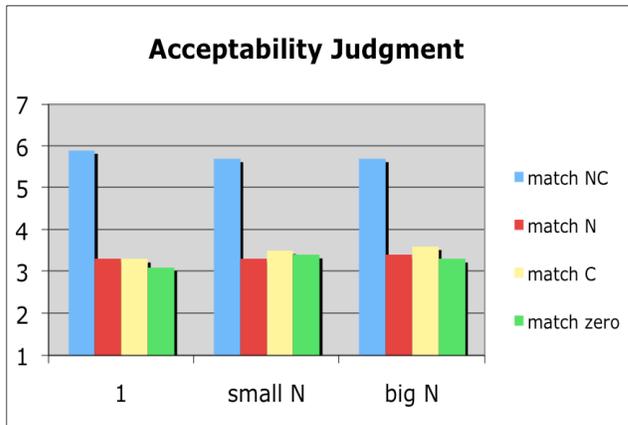
A 3 by 4 ANOVA shows no main effect of numerical context, but does show a main effect of feature match (F_1 (3, 102)=81, p_1<.001; F_2(3, 177)=220, p_2<.001). The interaction is significant by subject only (F_1(6, 204)=2.9, p_1<.01; F_2(6, 354)=1.2, p_2>.3). A separate analysis for each numerical context shows that for each numerical context, the grammatical condition is rated higher than the ungrammatical conditions, and there is no difference among the three ungrammatical conditions.

The offline acceptability-judgment task confirmed the acceptability of the grammatical noun forms following the three different types of numerals. Namely, NOM.SG, GEN.SG, and GEN.PL are indeed the acceptable surface case-number markings on nouns in these numerical contexts.

4.4. Self-paced reading task results

4.4.1. Data analysis

Self-paced RT results were examined word-by-word. Because of the large number of conditions in our experiment, rather than presenting RTs for each word of the sentence, in the case of each condition, we will discuss only the three regions of interest around the critical noun: the word before the critical N, the critical N itself, and the spill over word after the critical N. RTs above

2000ms were considered outliers and removed from data analysis, which affected 3% of the data. For the rest of the RT data, mixed-effect models were carried out for each region of interest, with subjects and items as random effects, and numerical context and feature match as fixed effects. This analysis allows simultaneous generalization to the population of participants and items, and can avoid potential spurious effects arising from the traditional ANOVA done on group averaged data (Baayen 2008). Analyses were carried out using R, an open-source statistical computing software (R Development Core Team 2008), and in particular the lme4 package for linear mixed-effect models (Bates et al. 2008).

4.4.2. Results

The numeral and the adjective before the critical N

Analyses were carried out for the numeral and the adjective regions before the critical noun. For each region, mixed effect models revealed no effect of either the numerical context or the morphological marking on the N, nor an interaction between the two. Planned pair comparisons also did not find any reliable difference between conditions.

The critical N and the spill over word

In Table 3 we present the RT and the standard derivation (SD) for each condition at the critical N and the one word after the critical N (spill over region).

Table 3: Mean reading times (standard derivations) for the critical N and the spill over word

	Critical N			N+1		
	1	5+	paucal	1	5+	paucal
match.both	465(155)	489(146)	501(177)	425(121)	414(115)	430(110)
match.case	540(219)*	537(183)	521(212)	496(140)**	473(124)**	455(115)
match.num	546(208)**	538(198)*	507(180)	504(149)**	469(120)*	465(104)
match.zero	576(233)***	517(168)	507(191)	512(155)***	513(144)***	441(109)

Significant codes (***.001, **.01, *.05) reflect p-values for mixed-effects models with the match-both condition as intercept.

The above table only presented mixed model effects with the "match.both" conditions as intercept, which should be read as comparisons between the "match.both" condition and the other three conditions (i.e. "match.case", etc.). We also carried out comparisons between the three ungrammatical conditions, and found no significant effects anywhere. All of our significant effects have medium to large effect sizes (0.4< Cohen's d <0.8), except for the "match.num" condition above in the 5+ context (effect size <0.3). All of our insignificant effects have small effect sizes (<0.3) (Cohen 1992).

We see from Table 3 that people showed different degrees of sensitivity to ungrammatical suffixes, depending on the numeral context that they were reading. When the numeral was 1, speakers' RT showed sensitivity to grammaticality as early as at the critical noun. The grammatical form was read faster than all of the other (ungrammatical) forms. There was no difference between the ungrammatical forms, and this effect continued to the spill over region. When the numeral was 5 or higher, there was no difference between the fully grammatical form and the fully ungrammatical form (zero-match) at the critical noun. However, the partially matched ungrammatical conditions (with the correct number feature or correct case form) approached significant difference from the grammatical condition (with a small effect size). At the spill over region, the 5+ context showed the same kind of grammaticality effect as the number

1 context—namely, the grammatical form was read faster than all other ungrammatical forms, and there was no difference among the three ungrammatical forms. With the paucal numerals, we did not find any robust difference between the grammatical condition and the ungrammatical conditions, either at the critical noun, or in the spill over position.

To summarize, the basic findings of the self-paced reading study are as follows. First, the paucal numeral context is very different from the other two numerical contexts, in the sense that speakers did not show any advantage when processing grammatical forms in this context, even though in the offline judgment task there was a clear preference for the grammatical form. Second, the number 1 context triggered responses to grammaticality at the critical noun, whereas the 5+ context showed a delayed effect at the spill over region. Third, whenever we see sensitivity to grammaticality, it appears to be categorical. If speakers clearly distinguished the fully grammatical and fully ungrammatical forms, the "intermediate" ungrammatical forms with partially matching features were treated exactly like the fully ungrammatical ones.

5. General discussion

5.1. Processing difference between the paucal context and other numerical contexts

The offline acceptability judgment task showed that native speakers are fully aware of the correct and incorrect forms that nouns may take in different numerical contexts. Across all numerical contexts, speakers showed very consistent performance, i.e. ungrammatical forms are rated significantly lower than grammatical forms. There is no difference among different numerical contexts.

However, the online data revealed crucial differences among different numerals: case/number violations in "1" and "5+" contexts evoked much longer reading time as compared

to the grammatical conditions, presumably reflecting speakers' detection of the ungrammaticality. Dissimilarly, in the paucal context the grammatical and ungrammatical conditions showed comparable reading times.

Why did speakers fail to show an online grammaticality effect for the pacual context when they clearly showed a grammaticality judgment difference in a rating task? We suggest that the key to accounting for these two observations is a "double life" of the phonological form associated with the GEN.SG suffix. If it is ambiguous - i.e., it is the phonological exponent of both [+gen, +sg] and a different morphological feature bundle (the paucal suffix) - the ambiguity resolution process at the suffix will slow down the reading time on the noun. Crucially, the extra processing time is only required for the grammatical form that carries an ambiguous phonological form, which decreases the RT difference between the grammatical and ungrammatical forms for the paucal nouns. For the other two numerical contexts, i.e. "1" and "5+", the grammatical forms are not ambiguous, and the morphological processing of a suffix is straightforward in these cases. For instance the NOM.SG suffix in the "1" context is unambiguously mapped to a [+nom,+sg] feature complex, which fits nicely with the processing prediction generated before the noun. On the other hand, an unexpected ungrammatical form in the same context leads to a much longer RT because of the "surprisal" effect on the parser, and also in some cases a reanalysis process of the parser, which tries to come up with a coherent parse.

Under this analysis, the surprising similarity in RTs between grammatical and ungrammatical forms of paucal nouns is a consequence of the extra processing effort involving the grammatical paucal nouns. In other words, the difficulty of processing a paucal noun blurred the grammaticality effect in this context. One reviewer pointed out that this analysis should predict a difference between the grammatical paucal form on the one hand, and the other two grammatical forms on the other hand, such that the grammatical paucal form has longer RT than

the other grammatical forms. We did a post hoc test to compare the three grammatical forms, and found no difference among them. The same post hoc analysis also found no reliable difference for each ungrammatical condition across different numerical contexts, with the only exception being that at the spill over region, the "match.zero" condition in the paucal context is read faster than the "match.zero" condition in the "1" context or the "5+" context (p<.01, all other p's>.05). The failure to find a reliable difference between the three grammatical forms could be due to a few factors. In particular, the suffixes on these forms could show other differences with regard to their frequency, length, markedness, etc. There could also be some inherent processing differences at the level of morphological access of these suffixes. For instance, in the discussion below (5.2.), we note that previous literature has found that the NOM.SG suffix seems to have processing advantage over other kinds of suffixes in lexical decision tasks. Variation on these dimensions could have interacted with the ambiguity resolution process, and washed away any ambiguity effect, especially if the actual effect size is small when only comparing grammatical forms. We also do not think much conclusion about ambiguity resolution can be drawn from the difference between the "match.zero" condition in the paucal context, and the other two contexts. This is a relatively late effect (at the spill over region), and it doesn't generalize to other ungrammatical conditions (i.e. "match.case" or "match.num" conditions). In addition, the RTs on this condition (i.e. NOM.SG in the paucal context) could have potentially been affected by what speakers perceive as the underlying grammatical form in this context (see discussion below in this section).

If the phonological form of the suffix on the grammatical paucal noun is ambiguous, an account which treats the paucal suffix as identical to the "genitive singular" suffix (as in Pesetsky 2007, 2010) becomes untenable, given that this account would predict that paucal phrases should show the same grammaticality effect as the other two numerical phrases. Moreover, we have

noted in section 3 that the account in Pesetsky (2007, 2010) evokes extra steps of derivations, with head-movement in particular, at a different time point of the processing of a DP. This may predict a higher processing cost for the numeral word in the paucal context, and for the noun in the number 1 context. Neither of these predictions was borne out. At the numeral and the following adjective position, all conditions have comparable RTs; nouns in the number 1 context are also not more costly than nouns in other contexts. However, we also want to point out that a lack of the effect associated with head-movements is not a direct argument against the analysis in Pesetsky (2007, 2010). Under the assumption of very active incremental parsing, a full-blown DP structure (without the actual lexical content) could be established with the minimal information of a numeral word. In principle, then, different syntactic operations could all be carried out at the numeral word, on just the terminal nodes of a syntactic structure, before any of the actual lexical input. In this sense, head movements themselves do not directly map onto the time course predictions in processing.

To summarize, we take the grammaticality effect, which is only present in "1" and "5+" contexts, but not in the paucal context, to be our key evidence in distinguishing the two types of analyses of the GEN.SG. marking on paucal nouns. Our results are problematic for any analysis that assumes an identity relation between the surface GEN.SG. form and the underlying representation; instead, we argue that the current result is most compatible with proposals that treat the GEN.SG. marking on paucal nouns as merely syncretic to some other underlying representations. As for the next question, i.e., what exactly this underlying representation could be, there are different implementations, some of which we mentioned in section 2.2.2. The experimental results we observed do not clearly distinguish these possibilities at this point. However, we want to point out one interesting observation that might bear on this discussion. As we mentioned earlier in this section, if we look at the comparison across different numerical

contexts, the "match.zero" condition in the paucal context (i.e. the NOM.PL suffix) is read faster than the "match.zero" condition in the other two numerical contexts. One possible way to explain this is that NOM.PL is closer to the actual underlying form. It is relatively easier to process than other "match.zero" ungrammatical conditions, because it actually fits in with what speakers know about the underlying form. This remains only bold stipulation before we carry out further investigations, but if it is on the right track, it is consistent with the analysis in Bailyn and Nevins (2008), which treats the paucal suffix as [+Nom, +paucal].

A critical assumption in our argument concerns ambiguity resolution. Psycholinguistic research on ambiguity resolution has focused more on syntactic and semantic rather than on morphological ambiguities, and we have attempted in this paper to draw a parallel between the processing of ambiguous functional suffixes and previous work on open-class words. More future research is certainly called for to further investigate the processing of morphological affixes.

5.2. Advantage in the "1" context

Our study also found processing differences between singular and plural contexts. Comparing the context where the numeral is "1," and the context where the numeral is a "5+" number, we see that in the former context, speakers are faster in detecting morphological ill-formedness. The grammaticality effect was observed at the critical noun for the singular number context, but was delayed to the spill over region for the 5+ number context.

One possibility is that morphologically speaking, nouns following "1" are easier to process than nouns following "5+" numbers. First of all, following the numeral "1," nouns carry the NOM.SG suffix, whereas following "5+" numerals, nouns are marked as GEN.PL. The NOM.SG suffix has consistently shown processing advantages over other case forms in a number of lexical decision studies of languages closely related to Russian, such as Serbo-Croatian (Lukatela et al.

1987, Feldman et al. 1987) and Serbian (Katz et al. 1995). In these studies, the lexical decision time on nouns with NOM.SG suffixes is significantly faster than nouns with other suffixes. It is still an open question as to why NOM.SG would have a processing advantage over other case forms (for possible explainations, see Katz et al. 1995 and Lukatela et al. 1987), but this could potentially drive the current observation that people are faster at detecting ungrammaticality related to NOM.SG forms.

The second difference between the "1" and "5+" context is the singular/plural distinction. There is evidence at least for English showing that morphological markedness could play a role when processing singular and plural nouns. Wagers et al. (2009) reported that when a noun is separated from its mismatched number-specific demonstrative (e.g., *this…apples* and *these…apple*), increasing the distance between the demonstrative and the noun only slows down the retrieval of the demonstrative for the singular noun conditions, but not for the plural ones. Wagers et al. argued that at least in English, the plural feature on nouns is marked, and the singular feature is unmarked. Because of this difference in markedness, when people process the number-specific demonstrative *these*, an expectation for [+plural] nouns is generated, and carried forward in the plural context, but no expectation is generated for [+singular] nouns in the singular context (*this*). As a result, less processing cost is needed to retrieve the plural demonstrative at the noun, since the [+plural] feature has always stayed in the focal attention of the working memory. It is conceivable that the difference we observed in the current study could also be due to the difference between the NOM.SG and GEN.PL in morphological markedness.[7] In the case of the masculine nouns, the NOM.SG form is the least marked form in the case/number paradigm, as its exponent is commonly a morphologically null form, cf. the form *malʹčik-Ø* 'boy.' On the other hand, gen.pl has three different morphological realizations, depending on the stem (cf. Zaliznjak 1967, 1977, Jakobson 1984). People could react faster to the morphological violations when they

error-check a less complex form, or a form that has few potential exponents. It may be that it simply takes more time, and more processing effort, to detect an error on a more morphologically marked form.

Another alternative account, suggested by one anonymous reviewer, has to do with the pre-nominal adjectives in the current set of testing stimuli. Taking the sample item in Table 2 as an example, for our testing items, at the adjective position, only the adjectives in the "1" context unambiguously signal the morphological form of the following noun - i.e., an adjective in NOM.SG indicates that the following noun is also NOM.SG. For the "5+" context (and for that matter, the paucal context as well), pre-nominal adjectives themselves do not unambiguously determine the form of the following noun: i.e., an adjective marked GEN.PL could potentially be followed by either a GEN.SG or a GEN.PL noun, depending on the actual numeral preceding the adjective. Since the pre-nominal adjective is intervening between the numeral and the noun, there could be a "agreement interference" effect in the "5+," but not in the "1" context. That is to say, in the "5+" context, although a GEN.SG. marked noun is unacceptable, it is still locally coherent with a GEN.PL adjective. This seemingly well-formed local relationship between the adjective and the noun could have interrupted people's rejection of the GEN.SG noun in the "5+" context. This is reminiscent of the well-known "agreement attraction" effect found in cases like "the key to the cabinets are…", where people fail to detect the problematic agreement between the matrix subject and the predicate (i.e. *the key….are…), due to a locally coherent agreement (i.e. …cabinets are…)—see Wagers et al. 2009. This explanation, although likely, is not entirely supported by the current data. Crucially, this account would predict that among all three ungrammatical conditions for the "5+" context, the "match.case" only condition, i.e. GEN.SG marked nouns, should lead to longer RT than the other two ungrammatical conditions, given that interference from previous adjectives would only affect the GEN.SG nouns, but not the other ungrammatical

30

forms. In contrast to this prediction, we observed no reliable difference among the three ungrammatical conditions.

5.3. Gradience effect

Recall that in the experimental stimuli, each grammatical form is compared with three types of ungrammatical forms. in the ungrammatical form, either only one case or number feature is incorrect, or both features are incorrect. This was designed to test whether violations of different case or number features, and the number of feature violations, would lead to different degrees of grammaticality violation, such that some ungrammatical forms would be perceived as "more ungrammatical" than others, or that more feature violations would lead to a higher degree of unacceptability.

Gradient acceptability has been reported for word recognition in a sentential context. In studies addressing these contexts, word recognition is facilitated because the semantic features of the target word have been activated in advance, through prediction from sentential context. But crucially for our topic, the facilitation effect is not only found for the target word that best fits the discourse context, but also for words that share semantic features with the best-fit word, even if these semantically related words, strictly speaking, do not fit in the sentential context (Kleiman 1980; Kutas and Hillyard 1984; Federmeier and Kutas 1999). In other words, one often finds gradient effects among the expected target, an unexpected but semantically related target, and the unexpected and also unrelated target. For instance, in an ERP study, Federmeier and Kutas (1999) reported that in a context where people are expecting the target word *palms* in the sequence *They wanted to make the hotel look more like a tropical resort. So along the driveway, they planted rows of palms*, the incorrect word *pines* induced a smaller N400 compared to another incorrect

word *tulips*. Federmeier and Kutas suggested that this is due to more semantic feature overlap between *palms* and *pines* (e.g. they are both from the same semantic category of trees).

It is not clear from the literature, however, whether the processing of grammatical features also presents a pattern of graded acceptability. Taking agreement processing as an example, one can imagine two possibilities. The parser could signal absolute ungrammaticality as long as one feature value is mismatched. On the other hand, the parser could also keep track of each individual feature, and signal different degrees of ungrammaticality when different features are violated, or when different numbers of features are violated. Cross-linguistic studies so far in the literature have yielded support for both possibilities (Lukatela et al. 1987 for Serbo-Croatian; de Vincenzi 1999 for Italian; Barber and Carreiras 2005 for Spanish; Nevins et al. 2007 for Hindi).

Our results did not show clear signs of gradient grammaticality. At the critical word, only the number 1 condition showed the basic grammaticality effect between the grammatical form and the completely ungrammatical form (in which both case and number features are incorrect). In this context, there is no difference between the violation of a single feature (case or number), and the violation of both features. There is also no difference depending on whether it is a case feature or a number feature that is violated. At the spill over region, both the number 1 and the 5+ number context showed the basic grammaticality effect. In the number 1 context, the number of features that were violated had no effect on the reading time. In the 5+ number context, however, there was a slight difference in the features. Although all three ungrammatical forms were read significantly longer than the grammatical form, the condition that had the incorrect case feature but correct number feature (NOM.PL) was faster (only approaching significance) than the two-feature violation condition (469 vs. 513ms, p=.056; there is no difference between the case violation and the number violation conditions). Whether this numerical trend indicates gradient

32

grammaticality is not clear. Note that this is a delayed effect at the spill over region, and that more importantly, this effect only appeared in the 5+ number context. It is possible that this effect is specific to the processing of the NOM.PL form in that particular context.

We want to leave open the question of whether the lack of gradience in Russian is a language-specific or more general property. We want to point out that Russian inflectional forms are mostly synthetic. It is hard to separate different features on a single morpheme, and fast online processing in such contexts may facilitate categorical decisions. If this explanation is on the right track, we can predict a more likely gradient effect in those languages that have agglutinative morphology. Of course, in order to test this last notion, more cross-linguistic research on morphological processing is needed.

6. Conclusions

This paper presented experiments designed to understand the structure of Russian numerical expressions, and in particular, the underlying representation of the noun forms occurring with the numerals 2-4 (paucals). We showed that although speakers are sensitive to grammaticality violations in the offline judgment task, their online reading time only showed comparable sensitivity for "1" and "5+" numerals, but not for paucals. We argue that an account which extends previous findings on lexical ambiguity resolution to suffix ambiguity resolution can best capture the current results, and that our results also assist in distinguishing different groups of morpho-syntactic analyses. These results add to the growing body of literature that makes use of online processing methodologies to address theoretical linguistic issues.

References

Baayen, R.H., Davidson, D.H., and Bates, D.M. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390-412.

Baayen, R.H. 2008. *Analyzing linguistic data: A practical introduction to statistics*. Cambridge, UK: Cambridge University Press.

Babby, L. 1973. The deep structure of adjectives and participles in Russian. *Language*, 49, 349-360.

Bailyn, J. F., and A. Nevins. 2008. Russian genitive plurals are impostors. In A. Bachrach and A. Nevins (eds.). *Inflectional identity*, 237-270. Oxford: Oxford University Press.

Barber, H. and Carreiras, M. 2005. Grammatical Gender and Number Agreement in Spanish: An ERP Comparison. *Journal of Cognitive Neuroscience 17:1*,137-153.

Bates, D.M., Maechler, M., and Dai, B. 2008. lme4: Linear mixed-effect models using S4 classes. R package version 0.999375-1. [Software]. Vienna, Austria: R Foundation for Statistical Computing.

Billings, L. 1995. *Approximation in Russian and the single-word constraint*. Ph.D. Dissertation, Princeton University.

Cohen, J. 1992. A power primer. *Psychological Bulletin* **112**: 155–159.

Corbett, G. 1983. *Hierarchies, targets and controllers: Agreement patterns in Slavic*. London: Croom Helm.

Corbett, G. 1991. *Gender*. Cambridge: Cambridge University Press.

Corbett, G. 2000. *Number*. Cambridge: Cambridge University Press.

Crocker, M., M. Pickering, and C. Clifton, eds. 2000. *Architectures and mechanisms in sentence comprehension.* Cambridge: Cambridge University Press.

de Vincenzi, M. 1999. Differences between the morphology of gender and number: Evidence from establishing coreferences. *Journal of Psychology Research, 28:5*, 537-553.

Dikker, S., Rabagliati, H., and L. Pylkkänen. 2009. Sensitivity to syntax in visual cortex. *Cognition,* 110(3), 293-321.

Dikker, S., Rabagliati, H., Farmer, T. A., and L. Pylkkänen. 2010. Early occipital sensitivity to syntactic category is based on form typicality. *Psychological science: A journal of the American Psychological Society*/APS, 21, 629-34.

Federmeier, K. D. and Kutas, M. 1999. A rose by any other name: Long-term memory structure and sentence processing. *Journal of Memory and Language*, 41: 469-495.

Federmeier, K. D., Segal, J. B., Lombrozo, T., and Kutas, M. 2000. Brain responses to nouns, verbs, and class ambiguous words in context. Brain, 123, 2552-2566.

Feldman, L. B., and Fowler, C. A. 1987. The inflected noun system in Serbo-Croatian: Lexical representation of morphological structure. *Memory and Cognition, 15*, 1-12.

Ferreira, V. 1996. Is it better to give than to donate? Syntactic flexibility in sentence production. *Journal of Memory and Language*, 35, 724-755.

Franks, S. 1994. Parametric properties of numeral phrases in Slavic. *Natural Language and Linguistic Theory*, 12, 570-649.

Franks, S. 1995. *Parameters of Slavic morphosyntax.* Oxford: Oxford University Press.

Franks, S. 1998. *Parameters of Slavic morphosyntax* revisited: A minimalist retrospective. In Z. Bošković et al. (eds.). *Formal Approaches to Slavic Linguistics 8: The Connecticut Meeting*, 134-165. Ann Arbor, MI: Michigan Slavic Publications.

Franks, S. 2002. A Jakobsonian feature-based analysis of the Slavic numeric quantifier genitive. *Journal of Slavic Linguistics*, 10, 141-181.

Friederici, A. D., Pfeifer, E., and A. Hahne. 1993. Event-related brain potentials during natural speech processing: effects of semantic, morphological, and syntactic violations. *Cognitive Brain Research,* 1, 183–192.

Friederici, A. D., Hahne, A., and A. Mecklinger. 1996. The temporal structure of syntactic parsing: early vs. late effects elicited by syntactic anomalies. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 22, 1219–1248

Hahm, H.-J. 2006. Number agreement in Russian predicates. Proceedings of the 13[th] International Conference on Head-Driven Phrase Structure Grammar, 132-154. Stanford:CSLI.

Ionin, T. and O. Matushansky. 2004. A singular plural. In B. Schmeiser, V. Chand, A. Kelleher and A. Rodriguez (eds.), *Proceedings of WCCFL 23*, 399-412. Stanford: CSLI.

Ionin, T. and O. Matushansky. 2006. The composition of complex cardinals. *Journal of Semantics,* 23, 315-360.

Jakobson, R. 1984. *Russian and Slavic Grammar: Studies, 1931-1981.* Berlin: Mouton de Gruyter.

Katz, L., Rexer, K., and Peter, M. 1995. Case Morphology and Thematic Role in Word Recognition. In L. B. Feldman (Ed.) *Morphological Aspects of Language Processing*. Lawrence Erlbaum Associates: Hillsdale, NJ. pp.79-107.

Kleiman, G. M. 1980. Sentence frame contexts and lexical decisions: Sentence-acceptability and word-relatedness effects. *Memory and Cognition,* **8,** 336 –344.

Kutas, M. and Hillyard, S. A. 1980. Reading senseless sentences: brain potentials reflect semantic incongruity. *Science* 207, 203–205.

Kutas, M. and Hillyard, S. A. 1984. Event-related brain potentials (ERPs) elicited by "novel" stimuli during sentence processing. In: R. Karrer, J. Cohen, and P. Tueting (Eds.). *Brain Information: Event-Related Potentials*. New York Academy of Sciences, Vol. 425, pp. 236-241.

Lau, E.F., Phillips, C., and D. Poeppel. 2008. A cortical network for semantics: (de)constructing the N400. *Nature Reviews Neuroscience*, 9, 920-933

Lau, E.F., Rozanova, K., and Phillips, C. 2007. Syntactic prediction and lexical surface frequency effects in sentence processing. *University of Maryland Working Papers in Linguistics*, 16, 163-200.

Lee, C. and Federmeier, K. D. 2009. Wave-ering: and ERP study of syntactic and semantic context effects on ambiguity resolution for noun/verb homographs. *Journal of Memory and Language*, 61: 538-555.

Lukatela, G., Carello, C., and Turvey, M. T. 1987. Lexical representation of regular and irregular inflected nouns. *Language and Cognitive Processes, 2,* 1 - 17.

Mel´čuk, I. 1985. *The surface syntax of Russian numerical expressions.* (Wiener slawistischer Almanach, 16.) Vienna: Institut fur Slavische Philologie.

Neville, H., Nicol, J., Barss, A., Forster, K. I., & Garrett, M. I. (1991). Syntactically-based sentence processing classes: evidence from event- related brain potentials. *Journal of Cognitive Neuroscience*, 3, 151–165.

Nevins, A., Dillon, B., Malhotra, S. and Phillips, C. 2007. The role of feature-number and feature-type in processing Hindi verb agreement violations. *Brain Research, 1164*, 81-94.

New, B., Brysbaert, M., Segui, J., Ferrand, L., and Rastle, K. 2004. The processing of singular and plural nouns in French and English. *Journal of Memory and Language*, 51, 568-585.

Pereltsvaig, A. 2006. Passing by cardinals: In support of Head Movement in nominals. In J. Lavine, et al. (eds.) *Formal Approaches to Slavic Linguistics 14: The Princeton Meeting*, 277-292. Ann Arbor, MI: Michigan Slavic Publications.

Pesetsky, D. 2007. Russian case morphology and the syntactic categories. Handout of the talk given at the Leipzig-Harvard Workshop on Morphology and Argument Encoding. http://web.mit.edu/linguistics/people/faculty/pesetsky/Pesetsky_Harvard_morphology_workshop_handout.pdf

Pesetsky, D. 2010. Russian Case Morphology and the Syntactic Categories. Ms. http://ling.auf.net/lingBuzz/001120

R Development Core Team. 2008. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. http://www.R-project.org

Rohde, D. 2003. Linger: a flexible platform for language processing experiments, version 2.94.

   http://tedlab.mit.edu/~dr/Linger/

Suprun, A. E. 1959. *O russkix čislitelʹnyx*. Frunze: Kirgizskij Gos. Universitet.

Swinney, D. 1979. Lexical access during sentence comprehension: (Re)consideration of context

   effects. *Journal of Verbal Learning and Verbal Behavior, 18*, 645–659.

Tanenhaus, M. K., Leiman, J. M., and Seidenberg, M. S. 1979. Evidence for multiple stages in

   the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and*

   Verbal Behavior, 18, 427–440.

Van Berkum, J. J., Brown, C. M., Zwitserlood, P., Kooijman, V., and Hagoort, P. 2005.

   Anticipating upcoming words in discourse: evidence from ERPs and reading times. *Journal*

   of experimental psychology. Learning, memory, and cognition, 31(3), 443-67.

Wagers, M., Lau, E., and Phillips, C. 2009. Agreement Attraction in Comprehension:

   Representations and Processes. *Journal of Memory and Language*, 61, 206-237.

Wechsler, S., and L. Zlatic. 2003. *The many faces of agreement.* Stanford:CSLI.

Zaliznjak, A. A. 1967. *Russkoe imennoe slovoizmenenie*. Moscow: Nauka.

Zaliznjak, A. A. 1977. *Grammatičeskij slovar' russkogo jazyka*. Moscow: Russkij jazyk.

Zwiterlood, P. 1989. The locus of the effects of sentential-semantic context in spokenword

   processing. *Cognition, 32*, 25–64.

The following abbreviations are used throughout the paper: ACC—accusative, FEM–feminine, GEN--genitive, MASC—masculine, NEG--negation, NEUT—neuter, NOM—nominative, PL--plural, PRES—present, SG—singular.

[2]   For details of Russian declensional classes, see Corbett 1983, 1991; Zaliznjak 1977.

[3]   We will not be concerned with approximative constructions where the numeral follows the noun, e.g., *štuk sem'* 'around/about seven items' (cf. Billings 1995, Mel´čuk 1985, Pereltsvaig 2006, Suprun 1959).

[4]   Deadjectival nouns (e.g., *životnoe* 'animal,' *buločnaja* 'bakery') always appear in the genitive plural with the numerals 2 and above, and we will not include or discuss these nouns here.

[5]   For a comprehensive discussion of paucal numerals and origins of paucal systems, see Corbett (2000).

[6]   Among the 60 sets of stimuli, the numerals in the 5+ condition include five (17 items), six (17), seven (12), eight (12), nine (1), and ten (1); the paucal numerals include two (21), three (24) and four (15).

[7]   We would like to emphasize that we are using "markedness" in a very narrow, specific sense, as defined by the number of possible exponents, and the availability of a null exponent.