

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/8198436>

Morphological Family Size in a Morphologically Rich Language: The Case of Finnish Compared With Dutch and Hebrew

Article in *Journal of Experimental Psychology Learning Memory and Cognition* · November 2004

DOI: 10.1037/0278-7393.30.6.1271 · Source: PubMed

CITATIONS

105

READS

429

5 authors, including:



Fermín Moscoso del Prado
University of California, Santa Barbara

25 PUBLICATIONS 1,049 CITATIONS

[SEE PROFILE](#)



Raymond Bertram
University of Turku

63 PUBLICATIONS 3,197 CITATIONS

[SEE PROFILE](#)



Tuomo Häikiö
University of Turku

35 PUBLICATIONS 679 CITATIONS

[SEE PROFILE](#)



Robert Schreuder
Radboud University

147 PUBLICATIONS 7,815 CITATIONS

[SEE PROFILE](#)

Morphological family size in a morphologically rich language: The case of Finnish compared to
Dutch and Hebrew.

Morphological Family Size in Finnish

Fermín Moscoso del Prado Martín^{1,2}, Raymond Bertram³, Tuomo Häikiö³, Robert Schreuder²,
& R. Harald Baayen²

¹ Medical Research Council – Cognition and Brain Sciences Unit, Cambridge, U.K.

² University of Nijmegen & Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands

³ University of Turku, Finland

Address all correspondence to:

Fermín Moscoso del Prado Martín

MRC–Cognition and Brain Sciences Unit

15 Chaucer Road

CB2 2EF Cambridge

United Kingdom

e-mail: fermin.moscoso-del-prado-martin@mrc-cbu.cam.ac.uk

Abstract

Finnish has a very productive morphology in which a stem can give rise to several thousand words. This study presents a visual lexical decision experiment addressing the processing consequences of the huge productivity of Finnish morphology. We observed that, in Finnish, words with larger morphological families elicited shorter response latencies. However, in contrast to Dutch and Hebrew, it is not the complete morphological family of a complex Finnish word that codetermines response latencies, but only the subset of words directly derived from complex word itself. Comparisons with parallel experiments using translation equivalents in Dutch and Hebrew showed substantial cross-language predictivity of family size between Finnish and Dutch, but not between Finnish and Hebrew, reflecting the different ways in which the Hebrew and Finnish morphological systems contribute to the semantic organization of concepts in the mental lexicon.

Introduction

In languages such as English and Dutch, stems differ in their productivity. Some stems give rise to a great many complex words. For instance, in English, the stem *man* appears in nearly 200 complex words. Other stems hardly ever give rise to complex words, e.g., the noun *scythe*, which only has its corresponding verb (*to scythe*) as morphological relative. Previous research has shown that the morphological family size of a stem, defined as the number of different complex words in which the stem appears as a constituent, is a robust predictor of response latencies in tasks such as visual lexical decision, auditory lexical decision, and subjective familiarity rating. Words with a larger morphological family size elicit shorter response latencies and higher subjective familiarity scores than do words with smaller family sizes matched for frequency (Schreuder & Baayen, 1997).

The effect of family size is present when measures of word form such as orthographic neighborhood size and bigram frequency are controlled for (Schreuder & Baayen, 1997). Family size counts are highly correlated with measures of morpheme frequency (Baayen, Tweedie, & Schreuder, 2002; Schreuder & Baayen, 1997). Both factorial studies (De Jong, Schreuder, & Baayen, 2000; Schreuder & Baayen, 1997) as well as regression studies (Baayen et al., 2002) have shown that family size effects can be observed independent from the effects of morpheme frequency. The family size effect is present independently of morpho-phonological inconsistency (De Jong et al., 2000), and remains a significant predictor when age of acquisition is partialled out (De Jong, 2002; see also Moscoso del Prado Martín, 2003).

The effect of morphological family size is well-established for Germanic languages (Dutch: Schreuder & Baayen, 1997, Bertram, Baayen, & Schreuder, 2000, De Jong et al., 2000; English: De Jong, Feldman, Schreuder, Pastizzo, & Baayen, 2002; German: Lüdeling & De Jong, 2002). Recently, an effect of morphological family size has also been established for a Semitic language,

Hebrew (Moscoso del Prado Martín, Deutsch, Frost, Schreuder, De Jong, & Baayen, 2003). In this language, morphological family size is defined in terms of the number of words that share a given consonantal root. The morphological family size in Hebrew ranges between 1 and 30, and is therefore much more restricted than English family sizes (range 1 to 200) and Dutch family sizes (range 1 to 550). Even though morphological families tend to be small in Hebrew, morphological family size emerged as a reliable predictor of response latencies independently of word frequency.

The family size effect is semantic in nature (Bertram et al., 2000; De Jong, 2002). Recent evidence supporting this conclusion has been obtained in Hebrew as well as for Dutch-English bilinguals. The Hebrew family size effect has a specific property that is particular to the Hebrew root, namely, that for words with homonymic roots the semantically related family members lead to facilitation while the semantically unrelated family members give rise to inhibition. In Dutch, such an effect has not been observed for homonymic stems (De Jong, 2002). However, a similar effect has been observed for interlingual homographs for Dutch-English bilinguals (Dijkstra, Moscoso del Prado Martín, Schulpen, Schreuder, & Baayen, 2003). Interlingual homographs are non-cognate words with identical spelling but different meanings across two languages. For instance, *angel* refers to a celestial being in English, and to a sting of a bee or wasp in Dutch. When Dutch bilinguals performed a Dutch simple visual lexical decision task with interlingual homographs as target words, the number of Dutch family members of the interlingual homographs is negatively correlated with response latencies (facilitation) while the number of English family members is positively correlated with the response latencies (inhibition). In contrast, when participants perform English simple lexical decision, the same homographs elicited response latencies that correlated positively with the Dutch family counts and negatively with the English counts. Since interlingual homographs such as *angel* have different meanings in the two languages, the

opposite effects of family size observed for the Dutch and English family sizes supports the hypothesis that the morphological family size arises at semantic levels of lexical processing.

It is important to realize that the family size measure explores semantic relations between sets of words (see, e.g., Moscoso del Prado Martín, Kostić, & Baayen, in press), while virtually all other studies known to us focus on semantic relations between pairs of words, such as synonymy, hyponymy, and hyperonymy, or associative measures obtained from ratings (e.g., McRae, DeSa, & Seidenberg, 1997). In fact, the family size effect turns out to be a much stronger predictor than such measures, see, e.g., the reanalysis of the data of McRae et al. provided by De Jong, Schreuder, and Baayen (2003). This greater predictivity is not so surprising once it is realized that a great many semantic relations in the lexicon are expressed morphologically. Furthermore, semantic categorization judgements and latencies are codetermined by morphological family size (De Jong, 2002). Finally, an information-theoretic account of the family size effect (in which the token frequencies of the family members are also considered) can be found in Moscoso del Prado Martín et al. (in press). This study shows that the family size count is a very good estimate of the informational complexity of morphological paradigms.

In this study, we report an experiment addressing the possible existence of a family size in Finnish. Finnish belongs to the Finno-Ugric language family, and is well known for its rich and complex morphology. It combines a complex inflectional system with a great many cases with productive derivation involving rampant stem allomorphy and very productive compounding. In Finnish, a stem such as *työ*, 'work', has roughly 7000 family members, including *työntekijä*, 'employee', *työehtosopimus*, 'wage rate treaty', *työstökone*, 'machine tool', *työläs*, 'laborious', and *työväenluokka*, 'working class'. Obviously, most Finnish stems have smaller morphological families, but many are very sizeable anyway with family sizes of some two hundred words or more.

While the Hebrew study established that family size effects generalize from Germanic concatenative morphology to Semitic non-concatenative morphology, the present study investigates whether family size effects also exist in a language with agglutinative¹ morphology like Finnish. It is far from evident that this would be the case. Just as the word frequency effect, the family size effect is logarithmic in nature. Robust effects are typically observed in the range of 0–40 family members, after which we generally have a floor effect. Given the large families counted for Finnish stems, no effect of family size might be observed due to an overall floor effect. As we will show below, this prediction is partially correct, requiring a more limited family size definition for complex words.

In a previous study, Moscoso del Prado Martín et al. (2003) performed two lexical decision experiments in Hebrew and Dutch. The materials for their Dutch experiment were the Dutch translations of the words that were used in the Hebrew experiment. Their results showed that Hebrew response latencies can be predicted from the Dutch family sizes of the corresponding translation equivalents even after Hebrew frequency and Hebrew family size are partialled out, and vice versa. This indicates that there is substantial similarity in semantic lexical organization in Dutch and Hebrew, even though these languages are typologically fundamentally different.

A second question addressed in the current study is whether a similar cross-language predictivity might be observed for Finnish and Dutch translation equivalents, and for Finnish and Hebrew translation equivalents. The patterns of cross-language predictivity have important implications for the degree of isomorphy in semantic organization across languages with typologically different morphological systems.

The purpose of the present study can therefore be summarized as follows. First, by attesting the role of a family size effect in Finnish, we provide further validation for this measure. Second,

by means of cross-linguistic comparisons, we further illustrate the potential of this measure as a research tool for investigating semantic organization in the mental lexicon.

The following visual lexical decision experiment addresses the questions raised above. It is designed along the lines of the Hebrew and Dutch experiments reported by Moscoso del Prado Martín et al. (2003), and makes use of translation equivalents of the Hebrew and Dutch words used in that study.

Experiment

Method

Participants. Twenty-six undergraduate students of the University of Turku participated in the experiment. All were native speakers of Finnish and had normal or corrected-to-normal vision.

Materials. The materials of these studies are the translation equivalents of the Hebrew and Dutch words used in the experiments reported in Moscoso del Prado Martín et al. (2003). As our point of departure, we took the 162 Dutch words from their Experiment 2, and translated them into Finnish. The translations were done using a Dutch-Finnish Dictionary (Suomi-Hollanti-Suomi taskusanakirja, Porvoo: WSOY, 1992), and they were extensively validated by the second and third author. When a word had different possible translations into Finnish with different meanings, we included all translation possibilities in the Experiment. Four of the original Dutch words could only be translated into Finnish using multi-word utterances, and were excluded from the experiment as well. In this way we obtained a set of 167 Finnish words. Of these words, 71 were morphologically simple, 81 were derived words, and 15 were compounds. Within the derived words there are only three that contain a suffix for which a base frequency effect has ever been reported in the literature (Bertram, Baayen, Schreuder, Laine & Hyönä, 2000; Järvikivi,

Bertram, & Niemi, 2003; Vannest, Bertram, Järvikivi, & Niemi, 2002).

Frequency counts for these words are based on the unpublished computerized Turun Sanomat Finnish newspaper corpus of 22.7 million word forms accessed with the help of the WordMill database program of Laine and Virtanen (1999). Morphological family size counts were also based on this database, with each of the potential family members evaluated by the third author, in some occasions aided by a dictionary (Nykysuomen sanakirja, Porvoo: WSOY, 1978). Each of these words was paired with a pseudo-word whose phonotactics did not violate the phonology of Finnish. The pseudo-words were derived from the experimental target words by changing 2 to 7 characters. Monomorphemic words were predominantly changed into pseudo-words without any morphological structure (e.g., *jalka* 'leg' was transformed into *solka*, *varas* 'thief' became *turas*), occasionally we pseudo-words created with a real stem but no suffix (e.g., *vaalia* 'to take care' became *puulia*, which contains the stem *puu* 'tree'). For derived words, only the stem was altered into a pseudostem, but the suffix remained intact (e.g., *hävytön* 'shameless' became *selytön* in which *-tön* corresponds to '-less'). For compound words, sometimes the first or second constituent was replaced by an alternative existing constituent (e.g., first constituent: *lampunvarjostin* 'lampshade' was transformed into *lennonvormostin* in which *lennon* means 'flight' (in genitive); second constituent: *nenärengas* 'nose ring' became *nypäkangas* in which *kangas* means 'textile'), sometimes the whole word was changed so that no sublexical morphological structure was present anymore (e.g., *itsemurha* 'self murder, suicide' became *istekorha*).

Twenty practice trials, ten words and ten pseudo-words were run before the actual experiment. We constructed three different permutations and their corresponding reversed versions of the original word list for counterbalancing. Table 1 provides a summary of the distributional properties of the data set.

Procedure. Participants were tested in noise-attenuated experimental rooms. They were asked to decide as quickly and accurately as possible whether the letter string appearing on the computer screen was a real Finnish word. Following a pause after the test trials, the experiment was run with two further pauses, dividing the experiment into three blocks, each containing one third of the materials. Items were preceded by a fixation mark in the middle of the screen for 500 ms. After 500 ms., the stimulus appeared at the same position. Stimuli were presented in white lowercase 12 point Helvetica letters on a dark background and they remained on the screen for 1500 ms. The maximum time span allowed for a response was 2000 ms. from stimulus onset.

Results and Discussion

All participants in this experiment performed with an error rate of less than 15%. One item elicited errors for more than 30% of the participants, and was thus excluded from the analyses. Additionally, we excluded four items that elicited response latencies of more than two and a half standard deviations above or below the mean.

INSERT TABLE 1 AROUND HERE

Table 1 provides the medians, means, standard deviations, and ranges for the frequency, family size, and word length counts for this data set, and the average response latencies in the experiment after excluding the four outliers. In addition, it lists details about the error scores. As the analyses of the error data revealed the same pattern of results as the reaction times, separate analyses of the errors will not be reported.

A multilevel regression model (Pinheiro & Bates, 2000, Baayen et al., 2002, a more powerful extension of the technique described by Lorch & Myers, 1990) fit to the dataset, with log response latency as dependent variable and log frequency, log family size, and word length as independent

variables revealed a facilitatory main effect for word frequency ($F(1, 3625) = 521.86, p < 0.0001$), an inhibitory main effect of word length ($F(1, 3625) = 137.66, p < 0.0001$, after partialling out the effect of frequency), and a facilitatory main effect of family size ($F(1, 3625) = 24.62, p < 0.0001$), after partialling out the effects of frequency and word length. We also observed a significant interaction between word length and word frequency ($F(1, 3625) = 89.21, p < 0.0001$), after partialling out the main effects: Longer words elicited longer response latencies, but only for lower-frequency words.

These results document, for the first time, the presence of a morphological family size effect in Finnish. As in English, German, and Dutch, and as in Hebrew, words with larger families give rise to shorter response latencies than words with smaller families. The presence of a morphological family size effect in three genetically unrelated language families, Indo-European, Hamo-Semitic, and Finno-Ugric, shows that, across typologically very different morphological systems, the organization of related words in morphological paradigms (i.e., the set of all the words that share a given morphological constituent) is an important factor in lexical processing.

Thus far, it would seem that the possibility we considered in the introduction, namely, that the large family sizes of Finnish compared to English or Dutch would lead to a floor effect, is not borne out. However, consider a selection of the members of the morphological family of *kirja*, 'book' in Finnish:

<i>kirja</i>	book
<i>väitöskirja</i>	dissertation
<i>muistikirja</i>	notebook
<i>päiväkirja</i>	diary, notebook
<i>romaanikirjallisuus</i>	novel literature

<i>aikakauskirja</i>	journal
<i>kirjasto</i>	library
<i>lainakirjasto</i>	public library
<i>kirjastonhoitaja</i>	librarian
<i>kirjoitus</i>	writing
<i>kirjoitusjärjestelmä</i>	writing system
<i>kirjepaino</i>	paper weight
<i>kirjailija</i>	author
<i>kirjailijantoiminta</i>	authorship
<i>asiakirja</i>	document
<i>kirjailla</i>	embroider
<i>kirjoittaa</i>	write
<i>kirje</i>	letter (a written communication)
<i>kirjain</i>	letter (the symbol)
<i>kirjeenkantaja</i>	postman
<i>kirjeenvaihtaja</i>	correspondent
<i>kirjeenvaihtotoveri</i>	pen-pal
<i>kirjoittautua</i>	register
<i>kirjoituskone</i>	typewriter

Note that while there is a family member that has a translation in English that contains the stem *book* (notebook), all other family members require translations with quite different stems in En-

glish, ranging from *author* to *library* and from *register* to *dissertation*. Note furthermore that some family members form semantically cohesive clusters, such as the words for *library*, *librarian*, and *public library*. This suggests the possibility that the family size effect in Finnish might be carried predominantly or perhaps even exclusively by the semantically more closely related family members.

One way of obtaining an objective and replicable way of defining the notion of being more closely related semantically, is to make a distinction between the family members of a word that are its direct descendants (its dominated family) and the other family members (its non-dominated family). Figure 1 illustrates the distinction between the dominated and non-dominated family size for the Finnish family of *työläinen* ('worker'). The dominated family size of *työläinen* consists of the words that are shown in bold in the figure. Its non-dominated family size consists of the remaining words. Note that the dominated family members are in general more closely related in meaning to each other than is the case for the non-dominated family members. This leads to the hypothesis that in Finnish, the morphological family size might be carried predominantly or perhaps exclusively by the dominated family size.

INSERT FIGURE 1 AROUND HERE

In order to test this hypothesis, we selected the 83 complex Finnish words in our dataset. (We excluded the monomorphemic words from the analysis, as for monomorphemic words the family size as a whole is identical to the dominated family size, the non-dominated family size being an empty set.) For these complex words, we determined the dominated and non-dominated family size. We then carried out a regression analysis, with log word frequency, word length, log dominated family size, and log non-dominated family size as independent variables, and log

response latencies as the dependent variable. A multilevel regression analysis revealed a highly significant effect for the dominated family size ($F(1, 2127) = 20.25, p < 0.0001$) and no effect whatsoever for the non-dominated family size ($F < 1$). In fact, it turns out that the total family size is not a good predictor for the complex words in our data. This shows that adding the non-dominated family members to the family size count for complex words in Finnish amounts to adding so much noise that the effect of the true predictor, the dominated family size, is completely masked.

The non-existence of a family size effect for the non-dominated family is partly in line with the intuition outlined in the introduction that with the large family sizes of Finnish the family size effect might be reduced due to a floor effect. However, restriction of the effect to the dominated family suggests that the degree of semantic relatedness in the family might be the key determinant rather than size as such. To gain further insight into the weight of these two factors, the magnitude of the family on the one hand, and its semantic cohesion on the other, we reanalysed the Dutch analogue of the present experiment reported in Moscoso del Prado et al. (2003), in which the translation equivalents of the Finnish words studied in the present paper were analysed. From their Dutch items, we selected the 59 words that were morphologically complex. A multilevel regression model revealed significant effects of both dominated and non-dominated family size, although the beta weight for the dominated family size ($\hat{\beta} = -0.085$, standard error = 0.017, $t(2018) = -4.875, p < 0.0001$) was more than twice as large as the beta weight of the non-dominated family size ($\hat{\beta} = -0.030$, standard error = 0.011, $t(2018) = -2.790, p = 0.0053$).

This result suggests that the dominated family size is the prime carrier of the family size, but that the non-dominated family size may also have some predictive power, at least in Dutch. This is probably due to the relatively small sizes (at least compared to Finnish) of morphological

families in Dutch. Within these small families, there is enough semantic similarity between the non-dominated and the dominated family members to allow a non-dominated family size effect to emerge. In Finnish, by contrast, the range of meanings covered by the non-dominated family is too broad, leading to semantic neighborhoods that are too sparsely populated to give rise to a measurable family size effect in the response latencies.

At this point, it should be made explicit that we do not claim that the distinction between the dominated and the non-dominated family is an absolute distinction for Finnish. To the contrary, we believe that closely related non-dominated family members will also contribute to the family size effect. However, we leave it to further research to establish principled ways in which the contributing non-dominated family members might be ascertained.

Summing up, the crucial contribution of the present experiment to our knowledge of the family size effect in human cognition is that by examining the family size effect in a highly productive agglutinative language such as Finnish, the semantic nature of the effect is clarified in more detail. If the family size effect were just a form effect, the distinction between the dominated and non-dominated family should not have been relevant, contrary to fact. This shows that the family size effect depends on the combination of shared morphological form and shared semantics. When the condition of semantic overlap is not met, as for most non-dominated family members in Finnish, those family members no longer contribute to the effect.

Cross-language Analyses

As mentioned in the introduction, Moscoso del Prado Martín et al. (2003) observed that Hebrew response latencies can be predicted from the Dutch family sizes of the corresponding translation equivalents even after Hebrew frequency and Hebrew family size have been partialled out

first, and vice versa. This result is indicative of substantial similarity in semantic lexical organization in Dutch and Hebrew, even though these languages are fundamentally different typologically. We now turn to investigate whether a similar cross-language predictivity might be observed for Finnish and Dutch translation equivalents, and also for Finnish and Hebrew translation equivalents. This will allow us to obtain insight in the extent of cross-language predictivity across typologically unrelated languages and its implications for the degree of isomorphy in semantic organization across radically different morphological systems.

For this cross-language multiple regression analysis, we selected those items that elicited less than 30% errors in the three experiments in Hebrew, Dutch, and Finnish. In this way, we obtained a total of 131 items, each with three response latencies. For each word in each of the three languages, we added as predictors length (in letters), word frequency, and morphological family size in that language. The key question of interest is whether length, frequency, and family size of, e.g., Dutch, predict response latencies in Finnish, even after the effects of Finnish frequency, Finnish word length, and Finnish family size, have been partialled out first.

Table 2 summarizes the results obtained for the 6 pairwise comparisons (Hebrew to Dutch, Hebrew to Finnish, Dutch to Hebrew, Dutch to Finnish, Finnish to Hebrew, and Finnish to Dutch). When predicting from language A to language B, we took the best multilevel regression model fitted to the data from language B as point of departure. The columns of Table 2 list the language for which the response latencies are predicted. The rows of Table 2 list the language from which an additional predictor (frequency, length, or family size) is taken. Each F-statistic and associated p-value corresponds to a separate analysis including the within-language variables and one additional predictor from another language. (Including more than one additional predictor at a time would have led to a serious collinearity problem.) For the details of the within-language

regression models for Hebrew and Dutch, the reader is referred to Moscoso del Prado Martín et al. (2003). For the Finnish data, the within-language model incorporates the effects of word frequency, length in letters, and family size: the full family size for the monomorphemic words, and the dominated family size for the complex words. The third row of Table 2 shows that Finnish frequency is an excellent predictor of Dutch response latencies, after having partialled out the effect of Dutch length, frequency, and family size. Finnish family size likewise emerged as a highly significant predictor, and even Finnish length turned out to have some predictive value.

INSERT TABLE 2 AROUND HERE

What Table 2 shows is that frequency is an excellent additional predictor in five out of six cases. The only instance in which frequency fails to have additional predictivity is when Finnish frequency is used to predict Hebrew reaction times. Note that, in turns of stem productivity, the typological distance is greatest between Hebrew and Finnish, with Dutch taking an intermediate position. Family size emerges alongside with word frequency as a remarkable explanatory variable in four out of six cases. The two cases where family size fails as a cross-language predictor is from Finnish family size to Hebrew reaction times and from Hebrew family size to Finnish reaction times. Again, cross-language predictivity breaks down where the typological difference in morphological structure and stem productivity is greatest. Finally, even word length shows some cross-language predictivity. The only language pair for which word length is predictive in both directions is Finnish and Dutch. The small differences in word length in Hebrew seem not to be predictive for Dutch but predictive for Finnish. Conversely, the big differences in word lengths in Finnish emerge as predictive for Dutch but not for Hebrew.

INSERT FIGURE 2 APPROXIMATELY HERE

Figure 2 summarizes the family size effects for the three languages by means of non-parametric regression lines. Note that the effect of family size is greatest for lower family sizes across all three languages, and that it levels off for greater family counts. The point of inflection is different for the three languages. For Hebrew, it is around 7.4 (e^2), for Dutch, it is around 20.1 (e^3), and for Finnish, it is around 148 (e^5). Although, arguing from Dutch, the large families of Finnish might have given rise to an early floor effect, we actually observe a floor effect only at a much larger family size.

General Discussion

The questions addressed in this study were, first, whether the family size effect might be observed in Finnish, and second, to what extent Finnish might participate in the cross-language predictivity of family size observed for Hebrew and Dutch. As to the first question, a visual lexical decision experiment revealed that, as in Germanic languages such as Dutch, English, and German, and as in Hebrew (Semitic), the morphological family size is also relevant for lexical processing in Finnish, a Finno-Ugric language. This finding provides further evidence for the cross-linguistic generality of the family size effect.

Earlier studies (De Jong, 2002; Moscoso del Prado Martín et al., 2003) established that the observed effect of the morphological family size probably arises at the level of semantic processing. These studies also established that semantic similarity shared between the family members is crucial for the effect to emerge.² Inspection of morphological families in Finnish, however, suggests that the larger families as a whole are semantically fairly diverse. To obtain further insight into the role of semantic similarity, we introduced the notion of the dominated versus the non-dominant family size for complex words. The dominated family size (consisting of the semantically more

similar morphological descendants of a complex word) turned out to be the crucial predictor for Finnish. A reanalysis of Dutch data showed both dominant and non-dominant family size to be relevant in this language. Given that morphological families in Dutch are both smaller and semantically more cohesive, we argued that this result supported the hypothesis that the family size effect crucially depends on semantic similarity. The operationalization of semantic similarity in terms of dominated versus non-dominated family size is a first objective and replicable operationalization for differentiating between clusters of semantically related words. We leave it to future research to develop more fine-grained operationalizations of semantic relatedness within morphological families.

Bates et al. (2003) studied response latencies in picture naming across a broad range of languages. They observed that picture naming latencies in one language could be predicted from the frequency and word length counts in another language. They interpreted these results as arguing in favor of a substantial semantic component to the word frequency effect. Following the line of research developed by Bates et al. (2003) for the cross-linguistic predictivity of frequency in picture naming, and the cross-linguistic predictivity of frequency and family size in Moscoso del Prado Martín et al. (2003), we investigated the cross-language predictivity of frequency and family size across Finnish, Dutch, and Hebrew.³ We observed substantial cross-language predictivity for frequency across the three languages, and more limited cross-language predictivity for word length. This suggests that there is considerable similarity in concept frequency in these languages, and that Zipf's observation that more frequent words tend to be shorter holds to some extent even across unrelated languages. Following Bates et al. (2003), we interpret these results as another indication of a substantial semantic component to the word frequency effect.

The most important cross-linguistic finding, however, is that the cross-language predictivity

of family size is absent when the distance between the morphological systems, as reflected in the degree of stem productivity, becomes very large. Finnish and Hebrew, the languages with the greatest and the smallest stem productivity, showed no additional predictivity for family size once the within-language measures (frequency, length, and family size) have been taken into account. This lack of predictivity contrasts markedly with the significant predictivity of family size from Hebrew to Dutch and vice versa. This suggests to us that there is a higher degree of overlap between the semantic organization in the mental lexicon of morphologically related words in Hebrew and Dutch, and in Finnish and Dutch, than there is for Finnish and Hebrew. Given that we have thus far only investigated three language families, and only very few languages within these families, this line of explanation remains necessarily tentative, and requires further research.

Although the cross-language predictivity of family size suggests that there may be considerable overlap in semantic organization, in the sense that words in dense morphological neighborhoods tend to have translation equivalents that also have dense morphological neighborhoods, the absence of such predictivity for Finnish and Hebrew suggests that there are limits to this cross-language predictivity. To understand why these limits might arise, consider, for instance, the consequences of the different degrees of productivity of compounding in Finnish, Dutch, and Hebrew. In Finnish, compounding is extremely productive, in Dutch, it is productive, and in Hebrew, it is marginally productive at best. Thus, complex concepts expressed by compounds in Finnish will have lexical (instead of phrasal) counterparts in Dutch relatively often, but very seldom in Hebrew. In Hebrew, many Finnish words will require phrasal translations. Consequently, the patterns of lexical co-activation in Finnish will resemble the coactivation patterns of their translation equivalents to a much larger degree in Dutch than in Hebrew. If, as has been argued by De Jong et al. (2003), the co-activation of the morphological family members indeed co-determines the

semantic percept of a word, then the present results support the Whorfian view of language, according to which language co-determines thought (see, e.g., Boroditsky, 2001). For languages with similar morphologies, the morphology might guide thought along similar paths, thereby giving rise to considerable cross-language predictivity of family size. When morphological systems are very different, as for Hebrew and Finnish, the well-worn paths along which morphology might lead thought become notably different, as witnessed by the breakdown of the cross-linguistic predictivity of family size for these languages.

Acknowledgements

The authors wish to thank Judith Kroll, Maryellen MacDonald, Sandy Pollatsek, and one anonymous reviewer for helpful comments and suggestions on a previous version of this manuscript. This study has been supported by the Dutch Research Council (NWO) through a PIONIER grant to the last author. During the last stages of preparation of the manuscript, the first author received additional support from the Medical Research Council (U.K.) and the European Community, under the “Information Society Technologies Programme” (IST-2001-35282).

Footnotes

1. Agglutinative languages are languages with an especially rich concatenative morphology in which grammatical functions tend to be expressed by separate affixes. For instance, the Finnish word *Taloissanikinko* translates into the English sentence '(Do you mean) in my houses, too?' and consists of the morphemes *talo-i-ssa-ni-kin-ko*.

2. In this study, we have used word length as a means for assessing the effect of family size while controlling for an important variable relating to word form. Other measures, such as neighborhood size and orthographic bigram frequency, were not included in our analyses, for two reasons. First of all, Schreuder & Baayen (1997) showed that the family size effect is not confounded with neither neighborhood size nor with orthographic bigram frequency. Second, adding such measures to the regression models leads to a very large increase in collinearity, with the condition number (Belsley, 1991; Belsley, Kuh, & Welsch, 1980) increasing from 12 to 79 when just bigram frequency is added. With such high collinearity the coefficients of the regression model become unstable. When a regression model with bigram frequency as an additional predictor is nevertheless fitted to the response latencies of the Dutch experiment, the effects of frequency, word length, and family size remain highly significant.

3. The cross-linguistic comparisons in this study are based on experimental data from highly-educated participants with at least a working knowledge of English. Thus, it might be possible that some of the Dutch-Hebrew and Dutch-Finnish cross-linguistic predictivity is due to this shared knowledge. More precisely, it is possible that the knowledge of English induces a more 'Dutch-like' (since Dutch and English have comparable morphologies) representation of the cor-

responding concepts in the Hebrew and Finnish mental lexicon. Note that such an explanation is perfectly in line with our main argument, that the morphology of the languages one knows does indeed shape ones semantic representations.

References

- Baayen, R. H., Tweedie, F. J. and Schreuder, R.: 2002, The subjects as a simple random effect fallacy: Subject variability and morphological family effects in the mental lexicon, *Brain and Language* **81**, 55–65.
- Bates, E., D’Amico, S., Jacobsen, T., Szekely, A., Andonova, E., Devescovi, A., Herron, D., Lu, C.-C., Pechmann, T., Pléh, C., Wicha, N., Federmeier, K., Gerdjikova, I., Gutiérrez, G., Hung, D., Hsu, J., Iyer, G., Kohnert, K., Mehotcheva, T., Orozco-Figueroa, A., Tzeng, A. and Tzeng, O.: 2003, Timed picture naming in seven languages, *Psychonomic Bulletin and Review* **10**(2), 344–380.
- Belsley, D. A.: 1991, *Conditioning Diagnostics: Collinearity and Weak Data in Regression*, Wiley, New York.
- Belsley, D. A., Kuh, E. and Welsch, R. E.: 1980, *Regression Diagnostics. Identifying Influential Data and sources of Collinearity*, Wiley Series in Probability and Mathematical Statistics, Wiley, New York.
- Bertram, R., Baayen, R. H. and Schreuder, R.: 2000, Effects of family size for complex words, *Journal of Memory and Language* **42**, 390–405.
- Bertram, R., Laine, M., Baayen, R. H., Schreuder, R. and Hyönä, J.: 1999, Affixal homonymy triggers full-form storage even with inflected words, even in a morphologically rich language, *Cognition* **74**, B13–B25.
- Boroditsky, L.: 2001, Does language shape thought? English and Mandarin speakers’ conceptions of time, *Cognitive Psychology* **43**(1), 1–22.

- Cleveland, W. S.: 1979, Robust locally weighted regression and smoothing scatterplots, *Journal of the American Statistical Association* **74**, 829–836.
- De Jong, N. H.: 2002, *Morphological Families in the Mental Lexicon*, MPI Series in Psycholinguistics, Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands.
- De Jong, N. H., Feldman, L. B., Schreuder, R., Pastizzo, M. and Baayen, R. H.: 2002, The processing and representation of Dutch and English compounds: Peripheral morphological, and central orthographic effects, *Brain and Language* **81**, 555–567.
- De Jong, N. H., Schreuder, R. and Baayen, R. H.: 2000, The morphological family size effect and morphology, *Language and Cognitive Processes* **15**, 329–365.
- De Jong, N. H., Schreuder, R. and Baayen, R. H.: in press, Morphological resonance in the mental lexicon, in R. H. Baayen and R. Schreuder (eds), *Morphological structure in language processing*, Mouton de Gruyter, Berlin, pp. 65–88.
- Dijkstra, T., Moscoso del Prado Martín, F., Schulpen, B., Schreuder, R. and Baayen, R.: 2003, Family size effects in bilinguals, *Manuscript submitted for publication, University of Nijmegen*.
- Järvikivi, J., Bertram, R. and Niemi, J.: 2003, Affixal salience and the processing of derivational morphology: The role of suffix allomorphy., *Manuscript submitted for publication, University of Turku*.
- Laine, M. and Virtanen, P.: 1999, *WordMill Lexical Search Program*, Center for Cognitive Neuroscience, University of Turku, Finland.
- Lorch, R. F. and Myers, J. L.: 1990, Regression analyses of repeated measures data in cognitive research, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **16**, 149–157.

- Lüdeling, A. and De Jong, N. H.: 2002, German particle verbs and word-formation, in N. Dehé, R. Jackendoff, A. McIntyre and S. Urban (eds), *Verb-particle explorations*, Mouton de Gruyter, Berlin, pp. 315–333.
- McRae, K., DeSa, V. and Seidenberg, M. S.: 1997, On the nature and scope of featural representations of word meaning, *Journal of Experimental Psychology:General* **126**, 99–130.
- Moscoso del Prado Martín, F.: 2003, *Paradigmatic Effects in Morphological Processing: Computational and cross-linguistic experimental studies*, MPI Series in Psycholinguistics, Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands.
- Moscoso del Prado Martín, F., Deutsch, A., Frost, R., Schreuder, R., De Jong, N. H. and Baayen, R. H.: 2003, Changing places: A cross-language perspective on frequency and family size in Hebrew and Dutch, *Manuscript submitted for publication*, Max Planck Institute for Psycholinguistics.
- Moscoso del Prado Martín, F., Kostić, A. and Baayen, R. H.: to appear, Putting the bits together: An information theoretical perspective on morphological processing, *Cognition*.
- Pinheiro, J. C. and Bates, D. M.: 2000, *Mixed-effects models in S and S-PLUS*, Statistics and Computing, Springer, New York.
- Schreuder, R. and Baayen, R. H.: 1997, How complex simplex words can be, *Journal of Memory and Language* **37**, 118–139.
- Vannest, J., Bertram, R., Järvikivi, J. and Niemi, J.: 2002, Counterintuitive cross-linguistic differences: More morphological computation in english than in finnish, *Journal of Psycholinguistic Research* **3**, 38–106.

	median	mean	standard deviation	range
frequency	670	3,155	7,097	1–56,193
word length	7	7.2	2.3	3–14
family size	298	620	892	8–6,029
dominated family size	88	273	485	0–3,080
non-dominated family size	29	347	762	0–5,835
response latency	604 ms.	617 ms.	63 ms.	530–808 ms.
error rates	0.000	0.017	0.033	0.000–0.192

Table 1: Medians, means, standard deviations, and ranges for the different counts, response latencies, and error rates in Experiment 1, after excluding four outliers.

		Hebrew	Dutch	Finnish
	frequency	-	$F(1, 4603) = 7.39^{**}$	$F(1, 3263) = 20.57^{***}$
Hebrew	word length	-	$F(1, 4603) = 2.53$	$F(1, 3263) = 11.51^{**}$
	related family size	-	$F(1, 4603) = 20.22^{***}$	$F(1, 3263) = 1.38$
	frequency	$F(1, 3184) = 15.28^{***}$	-	$F(1, 3263) = 35.99^{***}$
Dutch	word length	$F(1, 3184) = 6.17^*$	-	$F(1, 3263) = 56.94^{***}$
	family size	$F(1, 3184) = 15.03^{***}$	-	$F(1, 3263) = 18.62^{***}$
	frequency	$F(1, 3184) = 2.20$	$F(1, 4603) = 15.98^{***}$	-
Finnish	word length	$F < 1$	$F(1, 4603) = 4.90^*$	-
	family size	$F(1, 3184) = 1.52$	$F(1, 4603) = 16.75^{***}$	-

Table 2: Cross language predictivity of word frequency, word length, and morphological family size between translation equivalents in Hebrew, Dutch, and Finnish, in sequential analyses of variance in multilevel regression analyses. Significance codes are: $*p < 0.0500$, $**p < 0.0050$, and $***p < 0.0005$. The columns specify the language for which the response latencies are predicted, the rows indicate the languages from which the independent variables are taken. The value listed in a given cell specifies the significance of the predictor listed in the row of the cell after partialing out the within-language effects of frequency, length, and family size.




Figure 1: The position of *työläinen* ('worker') in the family of *työ* ('work'). The items in bold represent the dominated family of *työläinen*, the rest of the items are the non-dominated family.




Figure 2: Response latency (RT) against family size for Hebrew, Dutch, and Finnish using non-parametric regression (Cleveland, 1979).