



UTRECHT UNIVERSITY

The productivity of Dutch diminutives

A thesis presented for the degree of Master of Arts in Linguistics

Rosita Willemijna van Tuijl

5490383

Supervised by prof. dr. Peter Coopmans
Second reader: dr. Luisa Meroni

July 2020

Contents

1	Introduction	4
2	Background	6
2.1	Dutch Diminutives	6
2.2	Previous research	9
2.3	The Tolerance Principle	16
3	Corpus Research	20
3.1	Adult data	20
3.1.1	Gathering diminutives	20
3.1.2	Determining allomorphs	22
3.1.3	Feature selection	25
3.1.4	500 most frequent	26
3.2	Investigating individual children	27
3.2.1	Yang’s method	28
3.2.2	Application to Dutch diminutives	29
3.2.3	Problems with Dutch diminutives	30
4	WUG-test	33
4.1	Participants	33
4.2	Stimuli	34
4.3	Procedure	34
4.4	Analysis	36
4.5	Productions	39
4.6	Discussion	41
5	Analysis of the Dutch Diminutive system with the Tolerance Principle	42
5.1	Analysis	42
5.2	Discussion	53
6	Conclusion	57
	References	58
	Bibliography	58
	Acknowledgements	61
	Appendix	62

List of Figures

1	Example of a stimulus used in a PowerPoint set-up.	35
2	Boxplot for the real and nonce diminutive production task. The scores are grouped by age.	37
3	Boxplot for the real and nonce diminutive production task where the scores are grouped by suffix.	37
4	Child production data from the WUG-test: suffixes given to words with obstruent endings.	39
5	Child production data from the WUG-test: Suffixes given to words with sonorant endings and long final vowels. ‘P’ and ‘T’ below the bars indicates that the <i>target</i> answers were the <i>-pje</i> and <i>-tje</i> suffix respectively. . .	39
6	Child production data from the WUG-test: suffixes given to words with sonorant endings, short final vowels and penultimate stress.	40
7	Child production data from the WUG-test: suffixes given to words with sonorant endings, short final vowels and no penultimate stress.	40
8	Hypothesis 1	43
9	Hypothesis 2.1	45
10	Hypothesis 2.2	45
11	Hypothesis 3.1	46
12	Hypothesis 3.2	47
13	Hypothesis 4	48
14	Hypothesis 5	52

List of Tables

1	Type and token frequencies of the diminutive allomorphs from all corpora currently available on PaQu (except the CHILDES corpus)	24
2	Computational subregularities of the total of the combined corpora gathered from PaQu.	26
3	Hand annotated types of the subregularities of the Child directed speech corpus (from CHILDES in PaQu).	26
4	Hand annotated subregularities of the 500 most common nouns of the Subtlex database.	27
5	Nonce-words used in the WUG-test. Obstr.: Obstruent ending, Son.: Sonorant ending, Long: Long final vowel, Short: Short final vowel, Penult: Penultimate stress, No Penult: no penultimate stress.	34

6	Real words used in the WUG-test. Obstr.: Obstruent ending, Son.: Sonorant ending, Long: Long final vowel, Short: Short final vowel, Penult: Penultimate stress, No Penult: no penultimate stress.	35
7	Mixed effects model predicting performance based on age.	38
8	Mixed effects model predicting performance based on target suffix.	38
9	Hand annotated subregularities of the Child directed speech (from CHILDES from PaQu).	50
10	Hand Annotated subregularities of the 500 most common nouns in the Subtlex database.	51

List of examples

1	Non-noun words that can receive a diminutive suffix	6
2	Stems with obstruent endings get the <i>-je</i> allomorph	8
3	Stems with sonorant endings and long final vowels get the homorganic allomorph	8
4	Stems with sonorant endings, a short final vowel and penultimate stress get the homorganic allomorph	8
5	Stems with sonorant endings, a short final vowel and no penultimate stress get the <i>-etje</i> allomorph	8
6	Examples of the diminutive rules according to Kooij and Van Oostendorp (2003)	9
7	Examples of the diminutive rules according to Booij (1999)	9
8	Examples of the diminutive rules according to Huber (2005)	10
9	The Tolerance Principle	16
10	Application of the TP	16
11	R is productive if $T(N, e) < T(N, N)$	17
12	MMP	18
13	Simplified Dutch diminutive XPath query	21
14	Simplified Dutch diminutive XPath query	21
15	Problematic categorisations	22
16	Examples of the application of the match function	23

Abstract

This study investigates the Dutch diminutive paradigm, with its five suffixal allomorphs (*-je*, *-tje*, *-kje*, *-pje*, *-etje*), in relation to the Tolerance Principle of dr. Charles Yang. The Tolerance Principle lets us determine if a system can be productive or not. By studying the frequency of the allomorphs and the environments in which they occur, we can determine if a productive rule system can be found for the Dutch diminutives. In order to collect the necessary data I conducted a corpus study and an experimental study (a WUG-test) on Dutch diminutives. Applying the Tolerance Principle to the resulting data, determined that a productive system can be found for the Dutch diminutive paradigm.

1 Introduction

In Dutch, diminutives are used often and in many different contexts. Their most obvious use is to indicate that something or someone is of a small size. But using a diminutive does not necessarily indicate that something is small. For example, saying *Doe mij maar een biertje* ('I'll take a small beer') in no way means you are ordering a small beer, just a beer. Diminutives can also be used as a way to express affection, insecurity, condescension and in many other specific situations.

My focus, however, will not be on the many uses of the Dutch diminutive, but on the different forms the diminutive suffix can take. The Dutch diminutive is formed by adding a suffix to the noun. Not all nouns, however, receive the same suffix. The paradigm consists of five phonological variations, or allomorphs, that are attached to the stem: *-je*, *-tje*, *-kje*, *-pje* and *-etje*.

In this study I will focus on the underlying rule system of the Dutch diminutive paradigm. When do we use which suffix in order to diminutivise a noun? We can describe the paradigm as we perceive it, but this is not necessarily the way the system is represented cognitively.

The paradigm and its underlying rule system have been investigated by many researchers. Over the years more and more specific descriptions of the construction have been given. Many researchers have proposed different theories as to how the diminutive suffixes are represented cognitively by adults (Booij (1999); Huber (2005); Kooij and Van Oostendorp (2003); Trommelen (1984); Van de Weijer (2002)). Some have performed computational studies to address the learnability of the paradigm (Daelemans, Berck and Gillis

(1997)). Others have studied the diminutive productions of adults and children to further our understanding (Boersma (2018); Den Os and Harder (1987); Gillis (1997); Peelaerts (2008); Snow, Smith and Hoefnagel-Hohle (1980)).

In this study I will attempt to apply the Tolerance Principle developed by Yang (2016) to the Dutch diminutive system. The Tolerance Principle is a calculation that can be used to decide if a rule can be productive or not. I will try to find an underlying rule system that is productive according to this principle. The research question is: *Can a productive rule system be found for the Dutch diminutive system with the Tolerance Principle?* An answer to this question could possibly serve as further support for the Tolerance Principle and provide new insights in the domain of Dutch diminutives.

The principle needs frequency data of the paradigm in order to make a calculation. So, in this study I have gathered frequency data of the Dutch diminutive paradigm and applied the Tolerance Principle to it. In my analysis of this data, it is shown that the Tolerance Principle can indeed be used to find a productive rule system for the Dutch diminutives

In section 2, I start with some theoretical background: a description of the Dutch diminutive paradigm, a discussion of previous research on Dutch diminutives and the Tolerance Principle as presented by Yang. Section 3 describes the corpus studies I have performed in order to gather data to which I could apply the Tolerance Principle. In section 4 I present the results of a WUG-test I conducted to gather child productions of diminutives to support my analysis of the diminutive system with production data. Section 5 presents my analysis of the Dutch diminutive system with the Tolerance principle and discusses the results. Finally, section 6 is the conclusion to this study.

2 Background

Before I go into the theory behind the Tolerance Principle and its application it is important to discuss the paradigm we are going to apply it to: the Dutch diminutive paradigm. This section will give an overview of the Dutch diminutive system, a description of previous research on the subject of Dutch diminutives and finally a description and example of the Tolerance Principle of Yang (2016).

2.1 Dutch Diminutives

In the Dutch language diminutives are used to refer to an object as being of little or small format. But aside from this they can, on the one hand, also be used to express affection or a positive feeling towards something. For example, *vriend* is Dutch for ‘friend’, while *vriendje* is the Dutch diminutive form for friend, but also the word for boyfriend. On the other hand, diminutives can also be used as derogatory forms of speech. For example, among adults, calling someone *meneertje* (‘sir-DIM’) or *mevrouwetje* (‘lady-DIM’) is usually meant to be antagonising or condescending.

In Dutch, diminutives are formed by attaching a suffix to the stem of a word, where the stem is the base form of a lexeme. As Souman and Gillis (2007) point out, quite a few lexical categories can be made into a diminutive. Nouns can get the diminutive suffix, but so can adjectives, proper nouns, verbs, numerals, pronouns, prepositions and even short phrases (see (1) for an example of each). When a word gets a diminutive suffix, it is assigned neuter gender no matter the category or gender of the stem.

- (1) Non-noun words that can receive a diminutive suffix

	Stem	English	Diminutive	English
Adjective	<i>zout</i>	‘salty’	<i>zoutje</i>	‘salty candy’
Proper N.	<i>Roos</i>	‘Rose’	<i>Roosje</i>	‘little Rose’
Verb	<i>kopen</i>	‘to buy’	<i>koopje</i>	‘bargain’
Numeral	<i>tien</i>	‘ten’	<i>tientje</i>	‘tenner’
Pronoun	<i>iets</i>	‘something’	<i>ietsje</i>	‘a little bit’
Preposition	<i>om</i>	‘around’	<i>ommetje</i>	‘short walk’
Phrase	<i>tussen door</i>	‘in between’	<i>tussendoortje</i>	‘snack’

Not in all cases does a diminutivised word get the same suffix. There is an allomorphy that consists of five different suffixal forms. Namely, *-je*,

-tje, *-kje*, *-pje* and *-etje* (respectively [-jə], [-tjə], [-kjə], [-pjə] and [-ətjə]). Which allomorph is used is determined by the phonological features of the stem. There are three phonological features that determine which allomorph is used. They are: (i) the sonority of the ending of the word, (ii) the length of the final vowel and (iii) stress pattern.

(i) In phonetic contexts sonority refers to the airflow in the vocal tract. Obstruent sounds are sounds that are formed by obstructing the airflow, like the consonants [p, t, f, s]. Sonorants, in contrast, are formed with little to no obstruction, like the consonants [m, j, l, n, ŋ, r] and also vowels.

(ii) Vowel length indicates the difference between short and long vowels. Dutch has a set of five short vowels [ɪ, ɛ, ə, ʏ, ʌ], seven long vowels [i, y, u, e, ø, o, a] and a few diphthongs [ɛi, œy, ou] (Booij (1999)). From now on, in this study, the long vowels and diphthongs will be taken together as long vowels.

(iii) Finally, the stress pattern of a word is also important in determining which allomorph a word should take. The choice between allomorphs depends on whether a word has a stressed penultimate syllable or not. A word like *wō-ning* ('home') has penultimate stress while, any monosyllabic word or words like *wa-gōn* ('wagon'), which has ultimate stress, does not.

With these stem features it can be determined which stem is diminutivised with which allomorph. In (2) to (5) some examples are given for all the allomorphs. As indicated above the examples, words with obstruent endings get the *-je* suffix, words with sonorant endings and a long final vowel get the homorganic allomorphs and words with sonorant endings, short final vowels and penultimate or no penultimate stress get the *-etje* or the homorganic allomorphs respectively.

The examples in (3) and (4) get the homorganic allomorph. When two consonants are homorganic it means that they are pronounced in the same place of articulation. The allomorphs *-kje* and *-pje* are formed through place assimilation. The first consonants of *-kje* and *-pje* replace the 't' of the *-tje* suffix when they share place of articulation with the final consonant of the noun. In case of *-pje*, the 'p' replaces the 't' when the final consonant of the noun is an 'm': like *raampje* ('little window'), with the stem *raam*. In case of *-kje*, the 'k' replaces the 't' when the noun ends with 'ŋ': like *koninkje* ('little king'), with the stem *koning* (Boersma, Rispens, Weerman and Baker (2019)). This place assimilation process is why the *-tje*, *-kje* and *-pje* suffixes are lumped together as the homorganic suffix.

This description of *-tje*, *-kje* and *-pje* as the homorganic allomorph makes

it seem like *-tje* is the default suffix from which *-pje* and *-kje* are derived. As will be discussed in the previous research section below, not all researchers agree with *-tje* as the default rule of the paradigm. This matter will also be discussed further in the analysis (section 5).

In (4) the example *zalm* is indicated as having penultimate stress, contrary to what spelling may suggest. This is because in stems like these (also *olm*, *scherm*, *berm* etc.) a schwa is inserted between the liquid and the nasal in pronunciation, making something like *zaləm* (Gillis (1997), p. 2). The word becomes bisyllabic and can thus have penultimate stress.

Also in (4) the diminutivised form of example *woning* has lost the /g/: *woninkje*. In Dutch the spelling ‘ng’ stands for the velar nasal ‘ŋ’. However, when the ‘g’ is removed here, the left-over ‘n’ is still pronounced as ‘ŋ’. The only change is to the written form of the word.

- (2) Stems with obstruent endings get the *-je* allomorph

<i>boek</i>	‘book’	<i>boekje</i>
<i>drop</i>	‘licorice’	<i>dropje</i>
<i>roos</i>	‘rose’	<i>roosje</i>
<i>taart</i>	‘cake’	<i>taartje</i>

- (3) Stems with sonorant endings and long final vowels get the homorganic allomorph

<i>deur</i>	‘door’	<i>deurtje</i>
<i>zoen</i>	‘kiss’	<i>zoentje</i>
<i>auto</i>	‘car’	<i>autootje</i>
<i>boom</i>	‘tree’	<i>boompje</i>

- (4) Stems with sonorant endings, a short final vowel and penultimate stress get the homorganic allomorph

<i>rechter</i>	‘judge’	<i>rechtertje</i>
<i>woning</i>	‘residence’	<i>woninkje</i>
<i>zalm</i>	‘salmon’	<i>zalmpje</i>
<i>bodem</i>	‘bottom’	<i>bodempje</i>

- (5) Stems with sonorant endings, a short final vowel and no penultimate stress get the *-etje* allomorph

<i>zin</i>	‘sentence’	<i>zinnetje</i>
<i>wagon</i>	‘wagon’	<i>wagonnetje</i>
<i>tekening</i>	‘drawing’	<i>tekeningetje</i>
<i>vlam</i>	‘flame’	<i>vlammetje</i>

2.2 Previous research

In previous research there have been various theories on the formation of the different allomorphs of the Dutch diminutive paradigm. Some researchers have argued that the Dutch diminutive has a default form - a default suffix - from which the other allomorphs are derived. The two most popular candidates for this default suffix are *-tje* and *-je*, which is not surprising as these are the two suffixes that occur with the highest frequency.

Kooij and Van Oostendorp (2003) are among the researchers that have argued for *-tje* as the default form. They chose this form because the *-tje* allomorph occurs after stems that end in a glide or a vowel (see (6a) and (6b)). They assume that this is the environment in which the suffix would need to be changed the least. According to them, the ‘t’ of *-tje* would be omitted after obstruents to form the *-je* suffix (see (6c)). The *-pje* and *-kje* allomorphs would be formed through place assimilation of the final consonant of the stem and the ‘t’ of *-tje* (see (6d) and (6e)). Finally, the *-etje* allomorph is construed through schwa insertion (see (6f)).

- (6) Examples of the diminutive rules according to Kooij and Van Oostendorp (2003)
- a. leeuw + tje → leeuwtje (‘lion’)
 - b. dino + tje → dinootje (‘dinosaur’)
 - c. vos + tje → vosje (‘fox’)
 - d. zalm + tje → zalmpje (‘salmon’)
 - e. haring + tje → harinkje (‘herring’)
 - f. lam + tje → lam + e + tje → lammetje (‘lamb’)

Booij (1999) also chose the *-tje* suffix as the default form. His theory was similar to that of Kooij and Van Oostendorp (2003), but instead of place assimilation for the *-pje* and *-kje* allomorphs he argued that in these cases the ‘t’ was deleted and the ‘m’ and ‘p’ were inserted as homorganic stops (see (7a) and (7b)).

- (7) Examples of the diminutive rules according to Booij (1999)
- a. zalm + tje → zalm + je → zalm + pje → zalmpje (‘salmon’)
 - b. haring + tje → haring + je → haring + kje → harinkje (‘herring’)

A researcher with a different approach is Van de Weijer (2002). The theoretical background of Van de Weijer’s proposal is Optimality theory. In

this framework constraints and rules are assumed to be universal and thus not specific to certain languages or subgroups within that language. This is why he must dismiss the specific morpho-phonological rules mentioned above. According to Van de Weijer (2002), these analyses would need too many morpho-phonological rules that would be specific to the domain of Dutch diminutives, like the deletion of the ‘t’ or the insertion of the schwa. Instead, Van de Weijer uses universal constraints to arrive at the different allomorphs, still using *-tje* as the default.

On the other hand, there are studies that support the *-je* suffix as the default form, like Huber (2005). His claim is that the underlying form *-je* is turned into ‘something palatal’ whenever it is intervocalic (in the *-etje* cases) or when it is preceded by a stem with a sonorant ending.

As is visible in example (8a), after a stem with a sonorant ending a ‘t’ is added to turn the suffix into *something palatal*. The *-pje* and *-kje* allomorphs are formed by stems that receive an epenthetic stop and the *-je* suffix (see (8c) and (8d)). And according to Huber, because the diminutive suffix only attaches to a metrically appropriate base, the plural stem of *lam* is used instead of the singular (because the vowel in *lam* is short). After pluralising the stem it has become intervocalic, so a ‘t’ is added (see (8e)). Huber thus states that, in essence, the allomorphy is simply between *-je* after obstruents and its strengthened form *-tje* after sonorants.

(8) Examples of the diminutive rules according to Huber (2005)

- a. dino + je → dino + tje → dinootje
- b. vos + je → vosje
- c. zalm + je → zalm + p + je → zalmpje
- d. haring + je → haring + k + je → harinkje
- e. lam + je → lamme + je → lamme + tje → lammetje

One of the most cited works on the subject of Dutch diminutives is the study by Trommelen (1984). This work gives a very elaborate description of the Dutch diminutive paradigm. In this study she argues for *-tje* as the default form of the allomorphy. One of her central claims is that the selection of the allomorph used can be decided entirely based on the final syllable. This is in contrast with the description given by most researchers, because it excludes from the equation whether or not the penultimate syllable is stressed.

Trommelen argued that ə is actually bimoraic, with a null element as the first mora, making it equal to a long vowel. Consequently, words with schwa as their final vowel are excluded from the set of words with short final vowels. Trommelen argues that the rule for the *-etje* allomorph would then be *-etje only after short final vowels*. The schwa would be counted among the long vowels and these words would receive the homorganic suffix. In this analysis the stress of the penultimate syllable is not needed.

Daelemans et al. (1997) investigated Trommelen's claim by simulating the acquisition of the Dutch diminutive paradigm with a data mining method. They used the C4.5 algorithm (Quinlan (1993)) to generate a decision tree to determine what suffix a noun should get. C4.5 is an algorithm that maximizes information gain: the feature that provides the most information about which suffixes nouns should get is chosen to be highest in the decision tree.

Daelemans et al. (1997) tested the claim made by Trommelen (1984) by generating decision trees based on different selections of the noun features: onset, nucleus, coda and stress pattern of the syllables. They made a model that built a decision tree based only on the final syllable of the noun, as was proposed by Trommelen to provide sufficient information, and a model that used the features of all syllables. The model with only the features of the final syllable performed well, with only 2.3% of errors. It made most of its mistakes on the nouns that should get the *-etje* suffix. The model that received the features of all available syllables performed better overall and better on the *-etje* suffix. However, the difference was minimal, as the later had an error rate of 1.9%.

In the model that only takes the features of the last syllable into account the *-kje* after stems ending in *-ng* is overregularised and thus incorrectly gives some nouns ending in *-ŋ* the *-kje* suffix. For example, a word like *ring* ('ring'), that ends with *-ŋ* would be diminutivised as **rinkje*. However, as this word has no penultimate syllable and thus no penultimate stress, it should receive the *-etje* suffix (*ringetje*). This mistake makes sense because this model is not allowed to take the stress on 'possible' previous syllables into account. So, although the model with only the features of the final syllable gets very far, it still needs information about the penultimate syllable to get *-etje* right.

Aside from this theoretical and computational research quite a few experimental studies have also been conducted. The first and earliest I will discuss here is a study by Snow et al. (1980). Among some other Dutch morphological rules they studied the acquisition of the Dutch diminutives with a WUG-test (Berko (1958)). They looked at the production of the diminutives

by Dutch native children of 7 and 12 years old and by various age groups of second language learners (5 - 10, 12 - 18, adult). Their results showed that the order of acquisition of the various allomorphs seemed to correspond with the type frequency of that particular allomorph. The higher the type frequency of the allomorph, the earlier it is learned. This was the same for both first and second language learners.

Another experimental study that looked at the acquisition of Dutch diminutives by means of a WUG-test is a study by Den Os and Harder (1987). They investigated Dutch native children between 4 and 12 years old and a group of adults as a control group. They had the participants produce the diminutive forms of real Dutch nouns and nonce words. They also found a clear order in the acquisition of the different allomorphs. Their results show that the *-je* allomorph is not fully acquired until the children are already 7 years old. At this age they score over 80% correct for the first time. The *-je* allomorph is then followed by the *-tje* and the *-pje* allomorphs at age 9. The *-kje* is not acquired until the age of 12 years old. Finally, the *-etje* allomorph did not seem to be acquired completely, even by the oldest children tested. What is worth noting here is that even the adults made mistakes with the nonce words that should receive the *-etje* suffix.

Den Os and Harder (1987) also looked at the wrong answers given by the children. The children made few to no mistakes with the *-je* suffix. When they replaced the *-tje* and *-pje* suffixes they mostly used the *-etje* suffix, e.g. **stoeletje* and **boomtje*, which should be *stoeltje* ('chair-DIM') and *boompje* ('tree-DIM') respectively. Finally, *-kje* and *-etje* were replaced the most and in these cases the preferred replacement was the *-tje* suffix, e.g. **koningtje* or **zontje*, which should be *koninkje* and *zonnetje* respectively.

Gillis (1997) studied the diminutive acquisition of a single child from the age of 1;5 to 2;5. The analysed data consisted of bi-weekly recordings of natural conversations between the child and her mother. The child only started producing diminutives from the age of 1;7. Up to when she was 2 years old she used diminutives sporadically and from then on her use became more regular. She started out by only using the *-je* and *-tje* suffixes, which Gillis (1997) attributed to the suffix distribution in her mother's speech (60% *je*, 30% *tje*, 10% other).

In order to determine the extent to which the child had acquired the semantic relation between the diminutive suffix and the smallness of the objects, Gillis (1997) compared the child's use of diminutive forms with the child's use of the lemma forms. The comparison showed that relatively few

words were encountered in both diminutive and lemma forms. This suggests that had not figured out the pragmatic use of diminutives yet. Some additional evidence for this suggestion is that she also used the adjectives *klein* ('small') and *groot* ('big') along with the diminutives.

Just like Gillis (1997), Peelaerts (2008) also mentions a period where children do produce many diminutives but do not seem to understand the connection between the diminutive form and smallness yet. At this stage they use diminutives excessively. After this stage, the distribution (of diminutives vs stem forms) normalises into adult-like behaviour. At this point, when they start to understand the meaning and regularity of the diminutives they also derive incorrect stems from diminutives like **meis* from *meisje* ('girl-DIM'). *Meisje* is a lexicalised word, it is not formed by any stem+suffix. This indicates that the child now understands that the suffix adds some meaning to a stem.

To further investigate the diminutive production of children, Peelaerts (2008) conducted a WUG-test at a Dutch primary school, with children aged 4 to 11. Differently from the results of Den Os and Harder (1987), even the youngest children (4-year-olds) already performed very well on the *-je* and the *-tje* allomorphs (slightly better on *-je*).

In the rare cases that *-je* was replaced by another allomorph the children mostly chose *-tje* or *-etje* (e.g. *boekje* ('book') as **boeketje* or **boeketje*). If *-tje* was replaced the choice was most often *-etje* (e.g. *oortje* ('oor') as **ooretje*), but sometimes also *-pje* or even *-je*. If *-etje* was replaced, the children mostly chose *-tje* (e.g. *sterretje* ('star') as **stertje*), but sometimes also *-pje* and *-kje*. When *-pje* and *-kje* were replaced, the chosen suffix was also most often *-tje* (e.g. *koninkje* ('king') as **koningtje*) and to a lesser extent *-etje*.

Finally, one of the most recent studies in the field of Dutch diminutives was carried out by Boersma (2018). Her dissertation consists of multiple studies related to the acquisition of Dutch diminutives. She investigated the Dutch diminutive system in adult speakers of Dutch and she performed a judgement and production test with 5 to 10 year old Dutch children. She studied the results of these judgement and production tests in relation to the linguistic development of the children tested (Boersma, Rispens, Weerman and Baker (2018)) and in relation to the frequency effects and morpho-phonological characteristics of the diminutives themselves (Boersma et al. (2019)).

In order to investigate the Dutch diminutive system in adult speakers, Boersma performed a WUG-test with real and nonce words. She found that

the accuracy of the adults was at ceiling level for the real nouns, but the participants seemed to struggle with the *-etje* suffix for the nonce words. When the target suffix was *-etje*, the participants produced the right stem+suffix combination only 43% of the time. Instead of *-etje* they chose either *-tje*, *-kje* or *-pje*, in accordance with place assimilation. So for the nonce words *gol*, *ming* and *vom*, which should all receive the *-etje* suffix, many participants answered with **goltje*, **minkje* and **vompje*.

However, not only did most participants not follow the expected pattern (namely *-etje*), they were also inconsistent in their approach of this category of words. Some participants consistently chose *-etje* in the right circumstances (5%), some consistently chose *-tje*, *-kje* or *-pje* in accordance with the place assimilation rules (23%), but for the most part they used both of these approaches inconsistently (72%). For example, some participants would first diminutivise *ming* as *mingetje*, using the correct strategy, and when diminutivising the word *ting* they would answer **tinkje*, using another strategy on a word that is almost exactly the same.

According to Boersma, these results suggest that the nouns that are diminutivised with *-etje* are lexicalised forms, because it seems like the adults do not have a set strategy for diminutivising this set of words. She argues that the *-etje* rule is a hard rule to acquire, because the subset to which it applies is more constrained: only for sonorant endings, short final vowels and no penultimate stress. While the other subsets to which the other allomorphic variants apply only have final consonant and final vowel as constraints. Another reason for the difficulty of the acquisition of the *-etje* suffix is its appearance as an exception to other rules. For example, a noun like *brug* ('bridge') should receive the *-je* suffix, because it has an obstruent ending. However, the correct diminutivised form is *bruggetje*. According to Boersma (2018), these two factors cause the *-etje* suffix to be harder to acquire.

Aside from adults, Boersma also tested 5 to 10 year old children on their perception and production of diminutives. First she studied the children's performance in relation to their processing skills as measured by other tests (Boersma et al. (2018)). The idea behind this study is that, in order to accurately process the morpho-phonological characteristics of the diminutives, the children need to have good phonological processing skills. This is why, aside from the testing on diminutives, the participants performed a number of tests to determine, for example, the level of their phonological awareness or the size of their vocabulary. As was expected, the results showed that phonological processing skills contribute to both processing and producing

diminutive forms.

Finally, Boersma et al. (2019) also studied the performance of the 5 to 10-year-olds in relation to the morpho-phonological characteristics of the diminutives and various frequency effects. At the beginning of this study, Boersma et al. (2019) theorise which suffixes should be the hardest to acquire according to phonological constraints and frequency effects. They argue that the *-etje* suffix should be acquired last, because the subset of nouns after which it appears is the most specific. According to frequency, the *-kje* and *-pje* allomorphs should be acquired last, because these suffixes have the lowest type (and token) frequencies (this will be shown in more detail in section 3).

The results showed that the children improved in applying all the suffixes, except for *-etje* and *-kje*. In case of the real nouns, their performance on *-etje* and *-kje* stayed about the same in the different age groups. Their overall accuracy on the nonce words was a little lower than on the real words. Interestingly, in case of the nonce words, the performance of the children worsened for *-etje* as they got older. The older the children become, the more adult-like they become in their use of the *-etje* suffix. This result further suggests that the nouns diminutivised by *-etje* are lexicalised and not produced by a productive rule.

Boersma et al. (2019) also notice the difference between the acquisition of *-pje* and *-kje*. While both have comparable type frequencies, *-pje* is acquired significantly faster than *-kje*. Boersma et al. (2019) take this as evidence that the order of acquisition is not entirely due to the frequency of the various allomorphs.

In the final chapter of her (2018) dissertation Boersma brings together the results of word-specific characteristics of the diminutives and the linguistic skills of the children. She concludes that while frequency and vocabulary size have an effect on the accuracy scores of the children, the phonological skills of the children and the difficulty of the phonological features seem to be ‘the more direct causes of children’s protracted acquisition of some allomorphs compared to others’ (Boersma (2018), p. 147).

In this subsection (2.2), I have discussed a number of points that will be of importance throughout the rest of this thesis. First of all, there is a still ongoing debate about the default form of the diminutives in Dutch. Most researchers argue for either *-je* or *-tje* as the default, but a definitive default rule of the underlying system has not been found yet. Secondly, the agreed upon order of acquisition of the five allomorphs is: *-je* < *-tje* < *-pje* < *-kje* <

etje. And finally, the allomorph with which children seem to struggle most while acquiring the diminutive suffixes is the *-etje* allomorph. Even adults seem to have trouble applying this allomorph to nonce words correctly.

2.3 The Tolerance Principle

Yang (2016) proposed the Tolerance Principle as a model for the way we distinguish the core of paradigms from the periphery. In other words, it allows us to find the general rules that define a paradigm and the exceptions to these rules. When acquiring language we optimize our model of a paradigm by looking for a pattern, a rule. The Tolerance Principle is a way to determine if any of these rules can be accepted as a productive rule, or not. The Tolerance Principle states that a productive rule can only have a limited amount of exceptions. The amount of exceptions has to stay below a threshold, which is determined by the amount of lexical items in the paradigm. If the exceptions do not stay below the threshold, the learner has to revise her rule. The mathematical equation used to determine this threshold is shown in (9).

- (9) The Tolerance Principle :
 If R is a productive rule applicable to N candidates, then the following relation holds between N and e, the number of exceptions that could but do not follow R:

$$e \leq \theta_N \text{ where } \theta_N := \frac{N}{\ln N}$$

From Yang (2016), p. 61-64

The Principle comes with a step by step routine for its application (see (10)). First a rule *R* is obtained, this rule has a certain structural description and has a specific set of lexical items it applies to. The next steps are counting the number of lexical items that fit this structural description (*N*) and counting the items within this set that are exceptions to the rule (*e*). With this data the threshold can be calculated and the learner can decide if the rule is productive or not.

- (10) Application of the TP :
1. Obtain a rule R
 2. Count the number of lexical items to which R applies, N

3. Count the number of exceptions to R (a subset of N), e
4. Compare e and θ_N

By following these steps the learner could be in a constant state of testing hypotheses and accepting or rejecting them based on the Tolerance Principle (TP). By constantly applying this principle to input while acquiring language it can be determined which rules are most productive to use and which lexical items are better off being lexicalised as exceptions to the rule.

As stated above, the exceptions to a rule have to stay below the calculated threshold for the rule to be productive. The underlying equation that leads to the derivation in (9) actually compares the time complexity of two systems: a system with the productive rule (with a number of exceptions) and a system where *all* items are lexicalised. If the amount of exceptions to the rule stays below the calculated threshold, the system *with* the rule is more optimal than lexicalising the whole system. So, when we compare e to θ_N we actually compare the time complexity of two language models. This comparison is shown in (11). In this equation $T(x, y)$ calculates the expected time complexity of a rule over x lexical items with y exceptions.

$$(11) \quad R \text{ is productive if } T(N, e) < T(N, N) \text{ ; otherwise } R \text{ is unproductive.}$$

From Yang (2016), p. 61-64

Yang incorporates a traditional approach to handling exceptions during language processing in the form of the Elsewhere Condition (Kiparsky (1973)). The Elsewhere Condition poses that exceptions are processed as items to be dealt with by more specific rules before the application of the general rule is considered. This process can be visualised a list of exceptions that need to be checked before the general rule can be applied. So, a longer list of exceptions makes the overall application of the rule slower. If there are too many exceptions, the model with the rule becomes less optimal in time complexity than lexicalising all the items.

This list of exceptions is ordered according to frequency of occurrence. This means that the words that occur often will be at the top of the list and can be retrieved faster. Clahsen, Eisenbeiss and Sonnenstuhl-Henning (1997) investigated this difference in processing between regulars (items that follow a rule) and irregulars (items that are exceptions to this rule). They conducted lexical decision tasks with German participles and plural markers to determine if the items were processed faster if they had a higher frequency.

They found that irregulars with a higher frequency were processed faster than irregulars with a lower frequency, but this difference was not found for regulars. The processing times for regulars were about the same, regardless of frequency effects. These results are in line with the Elsewhere Condition: irregulars are listed according to frequency, regulars are processed by one rule.

Now that we have discussed the theory behind the Tolerance Principle, let's move on to its application. In his book, Yang (2016) gives various examples of the application of the TP to corpus data, including English past tense, English stress, German plural markers and many others. I will use the case of the German plural markers to give an example of the TP in action here, because it is very similar to the Dutch diminutive system. As will become apparent later, this paradigm is similar to the Dutch diminutive paradigm in that there is no single overwhelming default rule (like there *is* for i.e. English past tense verbs).

German nouns are made into plural forms by adding one of five suffixes: *-(e)n*, *-s*, *-e*, *-er* and *-∅*. Which suffix is chosen depends on the gender and the phonological form of the noun. None of the suffixes occur often enough to be the default pattern according to the Tolerance Principle. No matter which suffix is chosen, there would always be too many exceptions to allow the rule to be productive.

There are, however, signs of productive rules in the productions of the speakers and learners. Learners overregularise some of the more frequent suffixes (*-e* and *-(e)n*). Even the least frequent suffix (*-s*) shows signs of productivity, as it is often applied to words that are new to the language. When studying the full paradigm, some structural generalisations can be made. So how can the TP find them? In order to achieve this Yang proposes the Maximize Productivity Principle (12). This principle encourages the learner to look for different rules if a hypothesis over the whole set N does not work out. This holds particularly for rules that divide N into *subsets*.

(12) Pursue rules that maximize productivity (MMP).

From Yang (2016), p. 72

So, what happens when no rule can be found that is productive over the whole set of German nouns? The learner will go on a search for subsets that might have a productive rule. In the case of German plural markers, a feature that can be used to make some subdivisions is gender.

First, Yang takes the feminine nouns as a subgroup. In this class most of the nouns are pluralised by adding the -(e)n suffix. The exceptions to this rule stay below the threshold, so the rule can be productive. The other side of this partition is the class of non-feminine nouns. In this class there is, again, no suffix that could become the default rule. Again, we apply the Maximize Productivity Principle (MPP) and look for subdivisions within this class.

Within the non-feminine class Yang finds some phonological regularities that take the \emptyset and -n suffix. If these items are excluded, more than 70% of the remaining set of non-feminine nouns takes -e as a suffix. Unfortunately, this is still not enough to make the set productive. But, a productive rule can be found if we just select the masculine nouns from this set. At this point, what remains of the non-feminine set is the class of neuter nouns. If we make a division between monosyllabic words with back vowels and the rest of the set, this class also turns out to be productive.

Finally, we arrive at the -s suffix. This suffix is divided over all three genders. So the -s suffix emerges as the most general rule, as it has no restrictions on gender. This indicates that -s must be the default rule. It also makes the rule perfect for new words that enter the language, as these often can not be placed in a gender category yet. A new word like *iPhone*, that entered the German language from English, does not have a firm gender marking yet. This results in the loanword getting the -s suffix: the most general rule without restrictions on gender.

In conclusion, even for a paradigm that seemed to be impossibly irregular by looking at the whole, making subdivisions in the data allows for the discovery of productive rules. This same method of rule finding will be applicable to the set of Dutch diminutives, as I will show in section 5.

3 Corpus Research

In this section I will discuss two kinds of corpus research I carried out in order to further investigate Dutch diminutives in relation to the Tolerance Principle. In section 3.1, I present a study where I extract diminutives from a large corpus of Dutch native adult data. In section 3.2, I apply a method presented by Yang to study the production data of individual children.

3.1 Adult data

In order to make any analysis of the Dutch diminutive system with the help of the Tolerance Principle we need frequency data. We need to know how often and in which environments (features of the stem) the various suffixes occur. In section 3.1.1, I describe how I gathered a corpus of Dutch diminutives. Section 3.1.2 explains how I determined which suffix each diminutive had and in section 3.1.3 it is explained how I divided this corpus based on the relevant features of the stems (sonority, vowel length and stress).

3.1.1 Gathering diminutives

The first step to get the frequencies of the different allomorphs is gathering a collection of Dutch diminutives from a large corpus. To collect these diminutives I used two online platforms: GrETEL and PaQu. GrETEL is used to make a detailed XPath query (an expression used to select specific nodes in a treebank) for diminutives that can be used in PaQu (further explanation below) to select all the diminutives from available Dutch corpora.

GrETEL (Augustinus, Schuurman, Vandeghinste and Van Eynde (2014)) is a corpus tool that works with syntactically annotated data. It can be used to search for specific linguistic patterns in corpora. In GrETEL one can investigate any linguistic pattern by specifying the linguistic features of this pattern in an XPath query. However, one can also use the tool to make an XPath query. By giving the tool an example of the pattern and giving a few more specifications, one can extract a very specific XPath query. With this query one can search Dutch treebanks for the pattern in question.

I used GrETEL to form an XPath query specifically for diminutives. I gave the tool the simple diminutive *boekje* ('book-DIM') as an example and checked if it was parsed correctly (as a noun, diminutive, singular etc.). Then

I asked GrETEL for a detailed query, which can be tested on Dutch corpora to see if it returns what I need. Example (13) shows the resulting XPath:

- (13) Simplified Dutch diminutive XPath query :
nountype = "common" & pos tag = "noun" & degree = "diminutive"
& number = "singular" & gender = "neuter" & case="standard"¹

This query was a little more detailed than was necessary, because it also included a specification for singular nouns (number = "singular") and a specification for common nouns (nountype = "common") that excludes proper nouns. As I would like to include both plural and proper diminutives I removed these specifications (see example (14)):

- (14) Simplified Dutch diminutive XPath query :
pos tag = "noun" & degree = "diminutive" & gender = "neuter" &
case = "standard"²

I used the XPath query in example (14) to select diminutives from large corpora with another online tool: PaQu. PaQu (Odijk (2015)) is an online tool that allows one to search through Dutch corpora (and even allows for the upload your own corpus) with an XPath query. PaQu first scans the selected corpus and finds all the instances of the nodes the XPath points to. In this case, all diminutives. Then, PaQu allows one to download this data in the form of a list of all occurring diminutives in the corpus and a count of their appearances. One can ask for all sorts of additional information, like the stem or gender of the word, or the previous and following words. In this search I had PaQu return the diminutives themselves and the stems of the words. The stemmer that PaQu uses to get the stems is FROG(Bosch, Busser, Canisius and Daelemans (2007)). It will be explained below why the stems were necessary.

With this tool I was able to collect diminutives from the various large corpora available on PaQu (the corpora that were included are listed in Appendix A).

All the diminutives gathered from the different corpora resulted in in 913,959 tokens and 16,889 types to be investigated.

¹Original XPath query: `//node[@ntype="soort" and @pt="n" and @graad="dim" and @getal="ev" and @genus="onz" and @naamval="stan"]`

²Original Xpath query: `//node[@pt="n" and @graad="dim" and @genus="onz" and @naamval="stan"]`

3.1.2 Determining allomorphs

In this section, I will explain how it can be determined computationally which allomorph a noun receives and why this process is not as straightforward as it may seem. When determining which suffix a noun gets, the simplest approach would be to compare the end of the diminutive form with the respective suffixes and check if they match (e.g. the end of *boekje* matches *-je*, so it receives the *-je* suffix). However, due to the similarity of the different suffixes this approach would often result in false positives.

For example, if I built an algorithm that started by comparing all the diminutive endings with *-je*, all would return a match, because all the other allomorphs also contain *-je* (*-tje*, *-kje*, *-pje*, *-etje*). A solution could be to start out with the longer suffixes, so none of those will be miscategorised as having the *-je* suffix, and leave *-je* as the default. After all the longer allomorphs had their turn, the remaining set of words would automatically be the set of words getting the *-je* suffix. However, as is shown in (15), this would also result in incorrect categorisations.

(15) Problematic categorisations :

1. A word like *boekje* ('book-DIM'), stem *book* would match with '*-kje*', while it should be '*-je*'.
2. A word like *hertje* ('deer-DIM') would be categorised as '*-tje*'. However, as the stem is *hert* it should be categorised as '*-je*'.
3. A word like *dametje* ('lady-DIM'), stem *dame*, would be categorised as '*-etje*', while it actually has the '*-tje*' allomorph.
4. A word like *aap* ('monkey-DIM'), stem *aap* would match '*-pje*', while it should be '*-je*'.

The examples in (15) show that a simple search that matches the diminutive with a suffix is not feasible. In order to categorise the diminutives correctly, more information is needed. Luckily, there is additional data because, as mentioned earlier, we also have the stems of the diminutives.

Now that we have all the necessary data about the diminutives, we can start counting the suffixes. This process can be divided into three steps: 1) cleaning the data, 2) separating stems and suffixes and 3) selecting the appropriate suffix.

In the first step the gathered data is cleaned up. No matter how great a stemmer FROG is, there are still some mistakes left in the data. Some diminutives have not been stemmed correctly, some of the selected diminutives

are not really diminutives and some words still have some punctuation that needs to be removed.

For example, a word like *zonnebankkleurtje* which is the diminutive of ‘tanning bed colour’, is not stemmed correctly by FROG. These kind of words can be compounded in Dutch, but hardly ever occur in text. The stemmer just returns the full diminutive as the stem in these cases. Another example is a word like *marine* (‘navy’), which was incorrectly selected by PaQu as a diminutive. These incorrectly stemmed and incorrectly selected words were removed from the data. These removals made up a large percentage of the types (80%) and tokens (19%). This high percentage can be explained by the large amount of compound words in Dutch. 93% of these removals were incorrectly stemmed words, mostly compound words, and only 7% were incorrectly selected words, like *marine*. It is clear that FROG’s stemmer can be improved, but that is outside the scope of this study. Additionally, the remaining data was cleaned up a bit by removing some irrelevant symbols that complicate the categorisation process (slashes, points, apostrophes etc.).

After this cleaning process there is still quite a sizeable list of diminutives left to investigate (3572 types and 589981 tokens). The second step is dividing the diminutives into stem and suffix. This is done by comparing the diminutive with the stem. For example, if the stem is *boek* (‘book’) and the diminutive is *boekje* (‘book-DIM’), the suffix is found by searching *boekje* for *boek* and returning what is left over. A few examples of this process are presented in (16).

For words ending in -ng the final -g is removed from the stem used in the comparison, as is made visible in (16d, e). This is because in nouns that receive the *-kje* suffix the final -g is deleted when the suffix is added. If the -g were not removed before comparison, the algorithm would return nothing in these cases.

(16) Examples of the application of the match function :

- a. $\text{match}(\textit{boek}, \textit{boekje}) \rightarrow \textit{je}$
- b. $\text{match}(\textit{hert}, \textit{hertje}) \rightarrow \textit{je}$
- c. $\text{match}(\textit{boom}, \textit{boompje}) \rightarrow \textit{pje}$
- d. $\text{match}(\textit{konin}, \textit{koninkje}) \rightarrow \textit{kje}$
- e. $\text{match}(\textit{rin}, \textit{ringetje}) \rightarrow \textit{getje}$
- f. $\text{match}(\textit{dame}, \textit{dametje}) \rightarrow \textit{tje}$

Unfortunately, in this step of the process a small percentage of the diminutives is lost from the analysis. Some of the words are correctly stemmed, but are discarded as incorrect because of the selection process described above. For example, a word like *scheepje* (‘ship-DIM’) has the stem *ship* (‘ship’). This is the correct stem, but when comparing *ship* and *scheepje* as described above, the algorithm will not be able to find *ship* in *scheepje* and will discard the diminutive. About 2% of the types and 1% tokens is lost here, leaving 3245 types and 722645 tokens.

Finally, the diminutives are categorised into one of the five categories based on the suffix that was extracted. Overall this is a simple process. The left-over suffix can then be compared with the allomorphs from longest to shortest (*-etje*, *-kje/-pje/tje*, *-je*), and categorised accordingly. For example, in (16e) the suffix of *ringetje* is determined to be *-getje* (the *-g* is still attached to prevent words receiving the *-kje* suffix from being excluded, as explained above). When comparing this suffix with *-etje*, it matches, so it is correctly categorised as having the *-etje* suffix.

In this way almost all diminutives can be matched with a suffix and the frequencies of the allomorphs can be counted. In table 1 the diminutive frequencies from all corpora available on PaQu are displayed.

Allomorph	Type	Token
<i>-je</i>	1824 (56%)	511973
<i>-tje</i>	964 (30%)	155114
<i>-kje</i>	113 (3%)	625
<i>-pje</i>	82 (3%)	10436
<i>-etje</i>	234 (7%)	44362
rest	28 (0.1%)	135
total	3245 (100%)	722645

Table 1: Type and token frequencies of the diminutive allomorphs from all corpora currently available on PaQu (except the CHILDES corpus)

Here, the diminutives categorised into the ‘rest’ category are words like *manneke* (‘little man’) or *stukske* (‘little piece’). These are diminutives used in dialects in some Western and Southern regions of Holland.

Table 1 presents the type and token frequencies. Here the type frequencies are the amounts of different diminutives with the same suffixal allomorph and the token frequencies are the total amounts of occurrences of the respective

allomorphs.

Appendix A presents the frequencies of the separate corpora.

3.1.3 Feature selection

As described in the previous section, for each diminutive I have determined the actual allomorph that was used to diminutivise the noun. Now, we need to divide the nouns into categories based on the features described in section 2.1.

As was described in section 2.1, the Dutch diminutive system is based on three phonological features: sonority, length of the final vowel and the stress pattern of the nouns. By dividing the gathered noun stems into subsets based on these features we can see the frequency of the five allomorphs within these subsets. For example, according to the paradigm described in 2.1, all stems with obstruent endings should get the *-je* allomorph as a suffix (like *boekje*, ‘book-DIM’). If we select all the stems with obstruent endings from the corpus, we can see how many of these words actually take the *-je* suffix and if there are any exceptions. In other words, it will now become clear how many diminutives actually follow the rules, and which rules can be productive according to the TP.

In this section, I will explain how the stems were categorised based on the features: sonority of the ending, length of the final vowel and stress pattern.

For the sonority feature, each stem was sorted as either obstruent or sonorant. Each stem was checked on its final character. If this character was an obstruent, the stem was sorted as obstruent, otherwise it was sorted as sonorant (except for words ending in *-ng*).

Next, each stem was sorted as either having a long or a short final vowel. In Dutch long vowels can be formed in three ways: a double vowel like in *boom* (‘tree’), a single open vowel at the end of a syllable (can be followed by a glide) like in *vlo* (‘flea’) or *zwaluw* (‘swallow’) and finally a diphthong like in *vrouw* (‘woman’). Exceptions were made for the sequences ‘*ia*’ and ‘*io*’, which occur in words like *maniak* (‘maniac’) and *spion* (‘spy’) and are pronounced with a short final vowel. The selections on the basis of both sonority and vowel length are displayed in table 2. These frequencies were gathered computationally based on the combined corpora from PaQu.

The final feature is the stress pattern of the stem. Specifically, if the word has penultimate stress or not. This feature seems fairly straightforward but is hard to select computationally. In order to achieve this, a measure of the stress

Suffix	Obstruent		Sonorant		Long		Short	
	Type	Token	Type	Token	Type	Token	Type	Token
<i>-je</i>	1824	511973	0	0	0	0	0	0
<i>-tje</i>	9	118	955	154996	587	115883	368	39113
<i>-kje</i>	0	0	113	625	0	0	113	625
<i>-pje</i>	0	0	82	10436	42	3750	40	6686
<i>-etje</i>	13	3021	221	41341	1	683	220	40658
rest	13	27	15	108	9	72	6	36
total	1859	515139	1386	207506	639	120388	747	87118

Table 2: Computational subregularities of the total of the combined corpora gathered from PaQu.

pattern of the word is needed. Getting this measurement was beyond the scope of this study. In order to still get data on the stress pattern, I selected a smaller corpus of child directed speech (extracted from the CHILDES corpus in PaQu) and manually annotated it. In table 3 the frequencies for these hand annotated features are displayed.

Suffix	Total	Obs	Son	Long	Short	- Pen	+ Pen
<i>-je</i>	391	391	0	0	0	0	0
<i>-tje</i>	225	0	225	127	98	1	97
<i>-kje</i>	2	0	2	0	2	0	2
<i>-pje</i>	20	0	20	12	8	0	8
<i>-etje</i>	70	8	62	0	62	62	0
total	708	399	309	139	170	63	107

Table 3: Hand annotated types of the subregularities of the Child directed speech corpus (from CHILDES in PaQu).

3.1.4 500 most frequent

Instead of only studying diminutives I also decided to look at the most frequently occurring nouns. I gathered the 500 most frequent nouns from the SUBTLEX-NL database (Keuleers, Brysbaert and New (2010)). The SUBTLEX-NL database is based on 44 million words from film and television subtitles. I selected the 500 most frequent nouns with an online search

in their database. Nouns that could not be diminutivised, like *vrede* (‘peace’) and *geweld* (‘violence’), were deleted.

I took these 500 most frequently occurring nouns and determined the suffix they should get manually. Additionally, I determined computationally if these words had a sonorant or obstruent ending and if their final vowel was long or short. Finally, I hand annotated all nouns to include the stress patterns and the diminutive suffix they should receive. The results of these divisions are shown in table 4.

Suffix	Total	Obs	Son	Long	Short	- Pen	+ Pen
<i>-je</i>	247	247	0	0	0	0	0
<i>-tje</i>	177	0	177	109	68	1	67
<i>-kje</i>	15	0	15	0	15	0	15
<i>-pje</i>	18	0	18	12	6	0	6
<i>-etje</i>	43	5	38	0	38	35	3
total	500	252	248	121	127	36	91

Table 4: Hand annotated subregularities of the 500 most common nouns of the Subtlex database.

For the analysis I will make in section 5, using the Tolerance Principle, we will need these collected frequencies. Actually, using only one of the corpora, either the CDS data or the 500 most frequent nouns, would do to make the analysis. In this study, I will be using *both* to support the reasoning in my analysis, as will be shown in section 5.

3.2 Investigating individual children

In this section I will describe a method Yang proposed to look at the acquisition of productive rules within individual children. This is one of the methods that might be of use to me to further investigate the productivity of Dutch diminutives. Yang (2016) discusses this method in relation to the acquisition of the English past tense system. The Dutch diminutive system, however, differs from the English past tense in that there is no obvious default rule, like there is the ‘-d’ rule for English past tense. This section will first describe Yang’s application of the method to the English past tense and then present some details surrounding the difficulty in applying it to the Dutch diminutive paradigm.

3.2.1 Yang's method

According to the Tolerance Principle (TP), in order to acquire a productive rule, the regulars within the set considered need to be overwhelmingly present. Yang shows that in the case of English past tense the most frequently occurring verbs are actually irregulars. The most frequently used verbs are often acquired first, so a large part of a child's early vocabulary will consist of irregulars. This means that the presence of the regulars is not immediately overwhelming. It takes the child a while to acquire more and more verbs before the set of regulars (ending with '-d') she knows becomes large enough to be productive.

This point in time - when the child learns enough regular verbs to make the rule productive - is marked by the point where the child starts overregularising this rule. As Yang states, 'the very first instance of overregularisation unambiguously marks the emergence of productivity' (Yang (2016) p. 87). This is the point where the child can recognize the rule as productive and starts overusing it, because she has not figured out all the exceptions yet. Even irregular verbs she previously appeared to pronounce correctly can now be overregularised with the '-d' rule.

For example, looking at the data analysed by Marcus et al. (1992), a child can produce the correct irregular form of 'win' as 'won' when she is around 2 years old and then go through a period where she overregularises this to 'winned', before returning to the irregular form 'won' (same for broke/breaked, came/comed etc.).

Consequently, if this is the moment when the rule becomes productive, this is the point in time when the amount of regular verbs the child knows should be enough to account for the irregulars as exceptions. If we had a list of the verbs known to the child at precisely this point, we could test this hypothesis. According to Yang, this calculation should be done with the effective vocabulary of the child (and not based on a child directed speech corpus).

Yang exemplifies this approach with the English past tense. He investigated the CHILDES (MacWhinney (2000)) data of some individual English children. He gathered the effective vocabulary of these children in relation to their usage of the past tense. If a verb occurred in any of its conjugations it was added to the child's vocabulary. One of the children started overregularising the English past tense '-d' rule at the age of 2;11. In the recordings before this the child used irregular verbs in their correct forms. The overreg-

ularisation only started after the recording at the age of 2;11. So according to Yang's assumption, this is the point where the child has acquired enough regular verbs to allow for the emergence of a productive rule.

As was mentioned above, this productive rule can be tested by extracting all the verb lemmas the child used up until that point. The extraction from CHILDES resulted in a list of 300 verbs, of which 57 were irregular. Applying the TP results in $\theta_{300} = 53$ allowed exceptions. While this is very close, 57 irregulars still cross the threshold of 53 exceptions.

To explain this discrepancy, Yang argues we should take into account the extracted vocabulary is not the full vocabulary of the child. The extracted vocabulary is a subset of the child's current vocabulary. The recordings are not likely to have captured all the words she knows, both regulars and irregulars are probably undersampled. Additionally, due to the irregular verbs being more frequent, the regular verbs are probably undersampled to a greater extent than the irregulars. Correcting for this undersampling would change the ratio (irregular/regular) to be slightly more in favour of the regulars. This would make the real threshold a bit higher than 53, and also high enough to account for the undersampled irregulars.

This example shows how the TP can be investigated on an individual basis. The method also clearly illustrates how easy it is for children to have slightly varying developmental tracks. Just a few extra lemmas in a vocabulary can make the difference between a rule becoming productive or not. The speed with which rules can be acquired depends at least partially on the individual child's rate of vocabulary acquisition.

3.2.2 Application to Dutch diminutives

As I discussed earlier, the Dutch diminutive system is more like the German plural markers, in its productivity, rather than like the English past tense. There is no obvious default rule that can account for all the exceptions like there is for the English past tense. But there are certainly productive rules, it's just that a child will have to look for subregularities in order to find them.

Is it still possible to apply the method Yang used on the English past tense in such a case? We can at least make an attempt. One of the obvious rules a child might find is the rule: *-je after nouns with obstruent endings* (as we will see in the analysis section). What we would expect for the acquisition of this rule is that the child starts overregularising the *-je* allomorph to exceptions within the set of nouns with obstruent endings. There are a few nouns with

obstruent endings that get the *-etje* suffix. An example of this is *bruggetje*, with the lemma *brug* ('bridge'). If the child starts overregularising the *-je* suffix to such words it would indicate that she has acquired the productive rule of the nouns ending in obstruents receiving the *-je* suffix.

After finding the instances where a child overregularises, we would need the vocabulary of the child at that moment to see if the rule that seems to emerge can be productive. The next section briefly describes how the components necessary to apply the method were gathered and why applying this method to the Dutch diminutives does not, unfortunately, return the desired results.

3.2.3 Problems with Dutch diminutives

In order to apply Yang's method I studied the children in the Groningen corpus, available on CHILDES. This corpus contains longitudinal data from seven Dutch children. The recordings were made when the children were between the ages 1;5 and 3;7 years old. I analysed the transcribed data by extracting all diminutives (every word ending with *-je* or *-jes*) produced by one of the children and checking if the word was diminutivised correctly. This resulted in a few incorrectly diminutivised words per child.

The next step was to gather the vocabularies of the children at the time of the different recordings. Comparable to the methods used to extract the diminutives from various large corpora as described in section 3.1, I made an Xpath query for nouns with GrETEL (Augustinus et al. (2014)) and ran this Xpath query on the CHILDES corpus on PaQu (Odijk (2015)). This resulted in a list of nouns uttered by adults and children in the CHILDES corpus. From these results I selected the utterances from children in the Groningen corpus and organised them by file (the file name also indicates the age of the child at the time of recording).

However, after gathering all the necessary resources, applying Yang's method to the Dutch diminutives turned out to be more difficult than I expected. Here, I will explain the problems I ran into by presenting the differences between the English past tense system and the Dutch diminutive paradigm.

The differences between the two systems make it hard to apply the method Yang uses on the English past tense to the Dutch diminutives. The English past tense system is a system with one overwhelming productive rule. This makes it easier to apply the method to this system, because there is

really only one type of overregularisation we have to watch out for.

The biggest problem, however, is the difference in the distribution of the regulars and irregulars when the children start acquiring the system. As I mentioned above, Yang explained that many of the most frequently used verbs are actually irregulars. Because of their frequent use these irregulars are acquired early. This means that many of the first verbs a child acquires will actually be irregular verbs. The child essentially starts out with a skewed perception of the system.

This skewed perception creates a situation where the child has to acquire more regular verbs before the rule can become productive. This in turn creates a very specific turning point where the child starts overregularising the irregulars, as a result of the threshold being crossed.

When acquiring the Dutch diminutive system the children do not start out with a skewed perception of the paradigm. There is only a small amount of exceptions and the majority of the first words a child learns are regulars. For example, this is the case if we consider the first split in the data, between nouns ending in obstruents and sonorants, and look at the first rule a child would acquire: nouns with obstruent endings get the *-je* suffix. Here we might then find some overregularisations of this rule to the irregulars. One of the children I investigated started saying **slabje* instead of *slabbetje* ('bib-DIM') at the age of 2;06. But if we then check the set of nouns with obstruent endings in the vocabulary of the child, we observe that this rule has been productive from the beginning, or at least from the moment of the first recording (at the age of 1;10). There simply are not enough exceptions to ever have the child doubt the productiveness of this rule.

This means that when a child shows signs of overregularisation of the rule, I can argue that the rule is productive in the child's grammar, but not that it became productive when the child starts overregularising. Because the rule was already productive before the child showed signs of overregularisation.

Moreover, the number of incorrectly diminutivised (possibly overregularised) nouns is not large enough to make a meaningful analysis. Because of these insoluble problems, I can simply not take this analysis any further than this.

This is, fortunately, not the end of my efforts to analyse the Dutch diminutive paradigm with the Tolerance Principle. The method described above was used to look specifically at individual children, but we can still study the paradigm in general. In section 5, I will analyse the Dutch diminutives with the child directed speech data and the 500 most frequent nouns corpus

(described in section 3.1). But first I will look at a different approach to getting a better grip on children's productions of the diminutive and its various forms.

4 WUG-test

A WUG-test (Berko (1958)) was performed to get diminutive production data from Dutch children between the ages of 4 and 12 years old. Based on the experiments described in section 2.2 it may be assumed that children of these ages are still developing their diminutive paradigms. Between these ages children are still looking for the most efficient way to process diminutives, that is to say, they are looking for productive rules. In this period they might go through stages when they overregularise nouns in certain categories. For example, as they learn the rule: "nouns with obstruent endings get the *-je* suffix", they could start overregularising words like *biggetje* to **bigje* ('piglet').

Collecting their productions of diminutivised Dutch words and (Dutch) nonce words can give us insight into the developmental process they go through. The questions this experiment aimed to answer are: 1) Do Dutch children (between the ages of 4 to 12) make mistakes when diminutivising real and nonce words? 2) If so, what mistakes do they make? 3) Do they use the wrong suffix? 4) If so, which suffix?

4.1 Participants

This experiment was focussed on children of 4 to 12 years old. In total 14 children between these ages participated in the experiment. Most participants were recruited by contacting the parents via their primary school. Some others were also acquired through personal connections. All participants were contacted via email. The parents conducted the experiment with their child themselves and sent the results to the experimenter via email. The experiment had to be conducted this way because of circumstances surrounding COVID-19.

The low number of participants was also caused by the COVID-19 pandemic. Many of the schools that were contacted politely refused, because they did not want to burden the parents (most of whom were caring for their children while working from home) any further. Details about the necessarily adapted method and other challenges these circumstances caused are presented in section 4.3 and 4.5.

4.2 Stimuli

As discussed in section 2.1, in the Dutch diminutive system nouns can be specified into four categories: those with 1) obstruent endings, 2) sonorant endings, long final vowel, 3) sonorant endings, short final vowel, penultimate stress and finally those with 4) sonorant endings, short final vowel and no penultimate stress.

For each category three words were chosen in the development of nonce stimuli for the WUG-test. The nonce words were generated with Wuggy (Keuleers and Brysbaert (2010)), a multilingual pseudo-word generator. By selecting the intended language and giving the generator a word, this programme returns different versions of the real word, provided that all conform to the phonotactic constraints of the language.

After gathering some nonce words for each category with Wuggy, I conducted a small survey among friends and family to determine if the resulting words ‘could be’ Dutch and if so, which ones looked the most Dutch. Participants were asked to score the words on a scale of 1 to 5, where 1 was ‘non-Dutch looking’ and 5 was ‘Dutch looking’. Some Lithuanian and Icelandic words were included in the survey as control items. At least 70% of the participants scored the Lithuanian and Icelandic words as non-Dutch. All the words that were eventually selected to be used in the WUG-test (see table 5) were scored at least 4 by 70% of the participants.

Obstr.	Son./Long/	Son./Short/Penult	Son./Short/No Penult
Spaaf	Kroem	Fander	Geng
Wuig	Zaar	Mering	Lem
Woek	Kreen	Zelm	Termon

Table 5: Nonce-words used in the WUG-test. Obstr.: Obstruent ending, Son.: Sonorant ending, Long: Long final vowel, Short: Short final vowel, Penult: Penultimate stress, No Penult: no penultimate stress.

4.3 Procedure

The study was reviewed and approved by the ethics review board of the Faculty Ethics assessment Committee of the Faculty of Humanities (FEtC-H), Utrecht University. Schools were approached and asked whether they wanted to participate in the study. Parents were then contacted and received


Obstr.	Son./Long/	Son./Short/Penult	Son./Short/No Penult
Boef	Veer	Leerling	Ring
Big	Boom	Helm	Lucifer
Pop		Viking	
Fles			

Table 6: Real words used in the WUG-test. Obstr.: Obstruent ending, Son.: Sonorant ending, Long: Long final vowel, Short: Short final vowel, Penult: Penultimate stress, No Penult: no penultimate stress.

an information letter with an active consent form to be signed if the parents and child were willing to participate.

Because of the circumstances surrounding COVID-19 it was not possible to go to the schools and conduct the WUG-test in person. To solve this problem, I made a PowerPoint with audio-recordings and pictures to prompt the children into producing the diminutive forms. This PowerPoint was sent to the parents of children in the right age groups. The parents were asked to make an audio recording of the experiment and send it to the experimenter.

In the PowerPoint the children were introduced to the character Roos, voiced by me. Roos asked the children to speak loudly and clearly before she introduced them to a second character, Leo, voiced by Sjoerd Eilander (a male fellow student in the Linguistic Master Programme at Utrecht University). Roos introduced Leo as a student of Dutch and asked the children if they would like to help Leo learn what to call small things.



(a) *Dit is een wuij!*
‘This is a wuij’

(b) *Hoe noem je een kleine wuij?*
‘What do you call a small wuij?’

Figure 1: Example of a stimulus used in a PowerPoint set-up.

After the introductions, the children were given three examples of how to help Leo. Two examples of real items and one example of a nonce item

were presented to them. After these practice items, the actual experiment started. For each item, Roos pointed out what the object on the slide was called (see figure 1a). On the following slide Roos asked Leo what the small version of the object would be called and Leo gave a negative response that made it clear he did not know the answer. This response was varied between the Dutch versions of phrases like 'I don't know' and 'Can you help me?' to avoid monotony. After that, the child was meant to give the right answer and was prompted slightly by flickering the image of the small version of the object in question (see figure 1b). The images to go with the nonce words that were used in the PowerPoint were drawn by me.

There were 20 test items in total, of which 12 were nonce words and 8 were real words. As shown in table 6, the number of real words used was actually 11. This is because 3 of these items were used to bring some variety in the experiment. Leo gave the right answer for *fles*, *boom* and *viking* after 5, 10 and 15 items respectively. This was done as it might discourage the children if Leo never appeared to know the right answer and to break the monotony of the test.

Two Dutch native speakers listened to all the sound recordings and scored the answers independently of each other. Both of the raters were trained linguists. After discussing all differences, we reached an agreement on all productions.

4.4 Analysis

In order to analyse the data a mixed effects logistic regression model was used. Mixed effects logistic regression is used to model binary outcome and these type of models can control for variation between participants and items in the form of random effects. All tests were performed in R (R Core Development Team (2013)) with the lme4 package (Bates, Maechler, Bolker and Walker (2015)). In the final analysis one child was excluded because he did not understand the goal of the task.

The participants were divided into 3 age-groups: 4 to 6, 6 to 9 and 9 to 12 years old. The mean scores per age group are displayed in figure 2. Here the scores are the correct (target) answers the children gave on a scale of 0 to 1, where 0 is incorrect and 1 is perfectly on target. For example, in the graph of the real words, the age group 9 to 12 scores 0.85. This means that, on average, the 9- to 12-year-olds gave the target answer 85 percent of the time.




Figure 2: Boxplot for the real and nonce diminutive production task. The scores are grouped by age.

Two separate mixed effects logistic regression models were looked at. In both the dependent variable was whether the child gave a correct or incorrect answer. In one of the models the independent variable (fixed effect) was age and in the other it was the allomorphs. In both models the participant numbers and items were included as random intercepts to account for random variation between participants and test-items.³

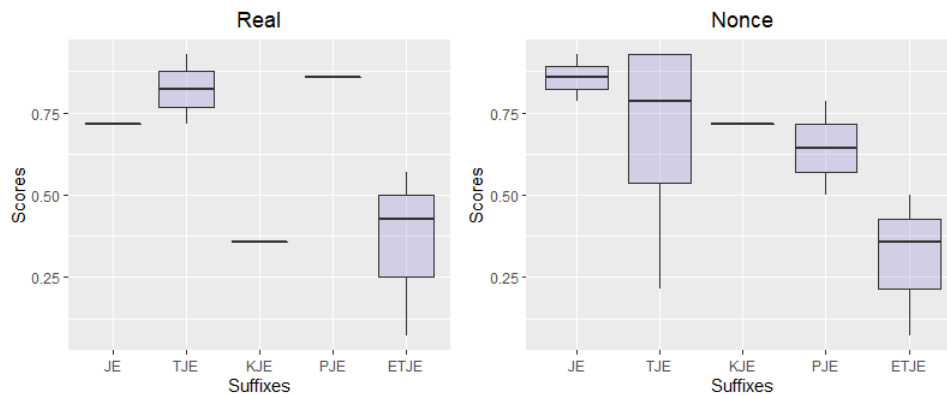


Figure 3: Boxplot for the real and nonce diminutive production task where the scores are grouped by suffix.

The model with age as the independent variable and figure 2 show that

³Model 1: $\text{answer} \sim \text{age} + (1|\text{item}) + (1|\text{participant})$,
 Model 2: $\text{answer} \sim \text{allomorph} + (1|\text{item}) + (1|\text{participant})$.

the children perform better, give more on-target answers, as they grow older. The results of this model are visible in table 7.

	Estimate	SE	z value	p
(Intercept) 4 - 6	0.35	0.48	0.73	0.46
6 - 9	1.10	0.36	3.02	< 0.01
9 - 12	1.34	0.43	3.09	< 0.01

Table 7: Mixed effects model predicting performance based on age.

The model with the suffixes as the independent variable shows that the children struggle the most with the *-etje* and *-kje* suffixes. For all three other allomorphs (*-je*, *-tje*, *-pje*) the answers correctly significantly more often. This can be derived from the results in table 8 and the mean scores per suffix in figure 3.

	Estimate	SE	z value	p
(Intercept) ETJE	- 0.78	0.58	-1.34	0.18
JE	3.25	0.98	3.30	< 0.001
KJE	1.17	1.11	1.05	0.29
PJE	2.29	1.00	2.27	< 0.05
TJE	2.66	0.88	3.04	< 0.01

Table 8: Mixed effects model predicting performance based on target suffix.

Based on the results of previous studies (Boersma (2018), Peelaerts (2008)) it was expected that the *-etje* suffix would be acquired later than the other suffixes. The poor performance on the *-kje* suffix may be explained by the relatively low frequency of the *-kje* suffix in the input of the children.

4.5 Productions

Here, I will discuss some of the productions in more detail. The children make very few mistakes with the nouns with obstruent endings (see figure 4) nor with the nouns with sonorant endings and long vowels (see figure 5). If they made mistakes it was mostly because they did not correctly hear the noun from the audio recording in the presentation.




Figure 4: Child production data from the WUG-test: suffixes given to words with obstruent endings.




Figure 5: Child production data from the WUG-test: Suffixes given to words with sonorant endings and long final vowels. ‘P’ and ‘T’ below the bars indicates that the target answers were the *-pje* and *-tje* suffix respectively.

With the sonorant ending nouns, with short final vowels and penultimate stress (6), the children make surprisingly few mistakes as well. I expected more mistakes here because the children have to take three features (sonority, vowel length, stress) into account to learn this rule. All children use the homorganic suffix and use the right allomorph depending on the stem. For example, *helmpje* in the case of the real nouns, and *zelpje* or *merinkje* for the nonce words.

Finally, with the sonorant ending nouns with short final vowels and no penultimate stress (see figure 7), which should receive the *-etje* suffix, the children make the most mistakes. With the word *termon* all children in the youngest age group use the *-tje* suffix instead of the *-etje* suffix. The middle age group performs slightly better and in the final age group almost half of the children still uses the wrong suffix. For the nonce word *geng* the answers in the lower age groups are divided between *-tje* and *-kje*. Only in the final age group do we find some participants that use the *-etje* suffix. And in case of *lem*, the youngest children mostly use *-pje*, the middle age group is divided between *-pje* and *-etje* and the older children all use *-etje*.




Figure 6: Child production data from the WUG-test: suffixes given to words with sonorant endings, short final vowels and penultimate stress.




Figure 7: Child production data from the WUG-test: suffixes given to words with sonorant endings, short final vowels and no penultimate stress.

These results indicate that the younger children mostly apply the homorganic allomorphs instead of the *-etje* suffix. As they get older they get better at using *-etje*, but they are still not always on target.

4.6 Discussion

I designed this way of conducting the WUG-test out of necessity. Because of time constraints it was not possible to test this new method extensively and it was hard to know what kind of results to expect. Here I'll discuss some limitations and confounds of this method based on the responses received in this study.

After receiving a few responses to the experiment, I was pleasantly surprised with the quality of the recordings. The productions of the children were perfectly audible and they responded very well to the questions asked by Roos and Leo.

Unfortunately, this way of testing also has some drawbacks. Sometimes the children found it hard to cooperate in the experiment, because they knew the characters in the PowerPoint could not really hear their response. Additionally, by having the parents conduct the experiment themselves it was impossible to control the environment in which the children took part in the experiment. The children might have been distracted by siblings or pets. Their parents might even have unconsciously helped them with the answers.

Although I was pleasantly surprised with the recordings that were sent in regarding quality and audibility, there were some recordings where the PowerPoint seemed to malfunction. Sometimes the audio-recordings of Roos and Leo would not play automatically, or multiple recordings would play at the same time. None of these complications created too many problems, but they did cause confusion among participants and raters. These problems might have been caused by a difference in the version of PowerPoint that was used to play the presentation.

Finally, the biggest drawback was the low response to an experiment that is conducted this way. We asked the parents to conduct the experiment themselves instead of just reading and signing a form (as would normally be the case for an in-person experiment conducted at school). In order to gather data, parents all have to respond individually and many will forget or simply choose not to participate. In order to gather a reasonable amount of data an very large number of people need to be contacted.

5 Analysis of the Dutch Diminutive system with the Tolerance Principle

In this section I will show that the Dutch diminutive system can be successfully analysed with the Tolerance Principle. After this analysis the results are discussed in relation to the theoretical background and the experimental study (WUG-test).

5.1 Analysis

This section will show how the Tolerance Principle can be used to account for the Dutch diminutive paradigm. Similar to the argumentation in Yang’s analysis of the German plural markers, I will divide the Dutch diminutive paradigm up into subregularities and show that these subregularities are acceptable according to the Tolerance Principle. I will use the results of the WUG-test I conducted to support the argumentation.⁴

For each potential rule I will select the total number of nouns in the (sub)set the rule should apply to from the table(s) specified below. This number is N in the formulation of the Tolerance Principle in 9. With this equation we can calculate the threshold, the maximum amount of exceptions, for that particular rule. By checking how many exceptions there actually are in the table, we can find out if the rule is acceptable or not. If the Tolerance Principle can explain the Dutch diminutive paradigm, we should ideally end up with a rule system that accounts for the formation of all diminutive forms.

In order to find these rules, a child has to test rules and discard or accept them. Let’s have a hypothetical child consider different hypotheses regarding the diminutive paradigm and test them with the TP. What is a rule she might start out with?

As a starting point we look at the frequencies of the five different allomorphs (*-je*, *-tje*, *-kje*, *-pje*, *-etje*) in the child directed speech (CDS) corpus that was extracted from the CHILDES corpus. We will use this corpus instead of the full corpus (of all corpora available on PaQu), because these

⁴In addition to the results from my own WUG-test I have also received the data from a WUG-test that was previously performed by Boersma et al. (2019). As the data they collected was very relevant to this research I contacted the researchers and Tiffany Boersma graciously allowed me to use their results. My analysis of their results is available in Appendix C for comparison to my own results.

Suffix	Total	Obs	Son	Long	Short	- Pen	+ Pen
<i>-je</i>	391	391	0	0	0	0	0
<i>-tje</i>	225	0	225	127	98	1	97
<i>-kje</i>	2	0	2	0	2	0	2
<i>-pje</i>	20	0	20	12	8	0	8
<i>-etje</i>	70	8	62	0	62	62	0
total	708	399	309	139	170	63	107

Table 3: Hand annotated subregularities of the Child directed speech (from CHILDES from PaQu).

frequencies are closest available to the actual input a child would get. The suffix frequencies of the CDS corpus are displayed in table 3, repeated here.

Aside from the CDS corpus I will also use the corpus of the 500 most frequent nouns I extracted from the Subtlex database. This corpus is similar to the CDS corpus in the distribution of the different allomorphs. Because of this similarity, applying the Tolerance Principle *mostly* yielded the same results. I will only expand on the data of the Subtlex corpus if they differ from those of the CDS corpus. For ease of comparison, the results of all calculations are displayed in Appendix B.

Our child starts by looking for a main productive suffix. If only one of the five allomorphs followed a productive rule, all other allomorphs would be exceptions to this rule. The productive allomorph should have such a high frequency that it could account for all the other allomorphs as exceptions. The allomorph with the highest frequency is *-je*. So, a rule the child might consider is the "everything is *-je*" rule, where the default is always *-je* and the other allomorphs are exceptions (see figure 8).




Figure 8: Hypothesis 1

Taking the total amount of nouns (708, in table 3) and applying the TP results in a threshold of $\theta_{708} = 107$ allowed exceptions. This number is far too low to account for the 317 (i.e. $708 - 309$) exceptions. As *-je* is the most frequently occurring suffix, none of the other allomorphs can result from one productive rule over the whole system either. This shows that the

Dutch diminutive paradigm can not just consist of one productive rule with exceptions. One productive rule for the whole system can not be accepted as a valid rule by the child. As is shown in figure 8 unacceptable rules are red, while acceptable rules will be green (as will be shown in later figures).

According to the Maximize Productivity Principle (MMP, in example 12) the child will ‘pursue the rules that maximize productivity’. Instead of settling on a system that has an overwhelming amount of exceptions, the child will search for other productive rules, ‘especially rules that divide the data into subsets’ (Yang (2016) p. 73).

Let’s begin our search for subsets with the partition between words with sonorant and obstruent endings. As is shown in table 3, words with an obstruent ending take *-je* as their suffix almost all of the time. There are only a few cases where a noun ending in an obstruent takes *-etje* as its suffix. These exceptions are words like *trap* (‘stairs’) and *weg* (‘road’) which are diminutivised as *trappetje* and *weggetje* respectively.




Figure 4: Child production data from the WUG-test: suffixes given to words with obstruent endings.

Applying the TP to the subset of words with obstruent endings results in a threshold of $\theta_{399} = 66$. The actual amount of exceptions to the rule within the set of nouns with obstruent endings is 8, so this is a perfectly acceptable rule. Figure 4, repeated here, shows that even the youngest tested children already made few mistakes in diminutivising words with obstruent endings, in both the real and the nonce task, and they only got better as they got older. This supports the idea that the split between obstruent and sonorant endings is the first learners find and that the *-je* suffix is thus acquired relatively early.

As the following suffix with the highest number of occurrences is *-tje*, let’s




Figure 9: Hypothesis 2.1

suppose that the child has a hypothesis space as illustrated in figure 9. As calculated above, the obstruent side of this tree makes for a good rule. The sonorant side, however, needs a little work. Looking at table 3, the amount of exceptions the *-tje* rule can tolerate is $\theta_{309} = 53$. As the number of actual exceptions would be 84 (2+20+62), this is not a sustainable rule.

Before completely throwing this tree out, the child might consider another hypothesis (figure 10). In the case of a homorganic rule, which takes the allomorphs *-tje*, *-kje* and *-pje* together as one rule, the only exceptions would be words taking the *-je* and *-etje* allomorphs. In this scenario the amount of exceptions would be 62 instead of 84. However, as this number is still larger than the allowed amount (53), this rule is not sustainable either.⁵




Figure 10: Hypothesis 2.2

Again, following the MPP, the child searches for a break in the data, a subregularity. She may find one in the length of the final vowel. As is shown in figure 11, a hypothesis could be made that breaks the sonorant set into two subsets: one consisting of words with long final vowels and one with short final vowels.

⁵However, according to the Subtlex data this rule *is* tolerable. Calculated with the Subtlex data the threshold is $\theta_{248} = 44$ and the number of exceptions is 38. I will ignore this for now, because the CDS data is closer to what our child would actually receive as input, but I will come back to this later in the analysis.




Figure 11: Hypothesis 3.1

In table 3 the separation between the long and short final vowels (within the sonorant set) is displayed. For the long vowels the threshold calculated by the TP is $\theta_{139} = 28$. Depending on their development regarding the homorganic rule, children could hypothesise a productive *-tje* rule or the homorganic rule. In both cases the amount of exceptions, 12 for the *-tje* rule and 0 for the homorganic rule, would stay below the threshold. This indicates that either *-tje* or homorganic can be productive rules for the words with sonorant endings and long final vowels.

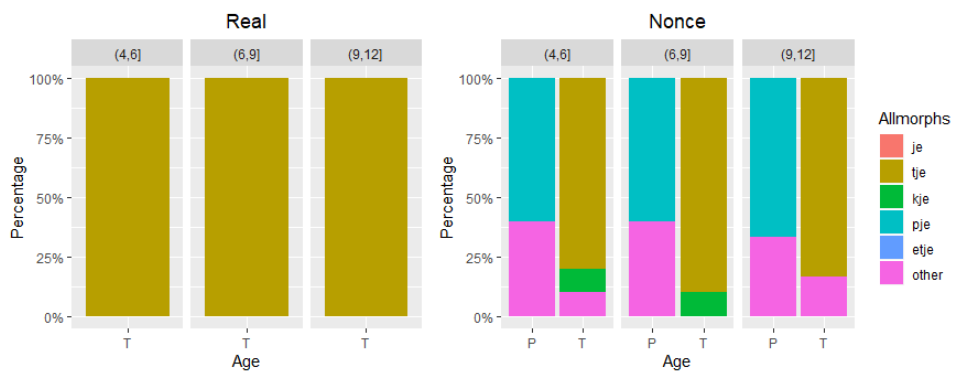


Figure 5: Child production data from the WUG-test: Suffixes given to words with sonorant endings and long final vowels. ‘P’ and ‘T’ below the bars indicates that the *target* answers were the *-pje* and *-tje* suffix respectively.

Figure 5, repeated here, shows the production data from the WUG-test on nouns with sonorant endings and long final vowels. In these graphs the bars have been separated based on their target suffix. The ‘P’ bars and ‘T’ bars represent the answers to nouns for which the target answers were *-pje*

and *-tje* respectively. There is no bar for the *-kje* suffix, because no words with long final vowels and velar nasals were included in the experiment (and these are very rare if not non-existent in Dutch). In both the real and the nonce task the children chose the right homorganic variant most of the time. The children showed improvement with age in the nonce task and performed perfectly in the real task. This data indicates that the children do not have trouble acquiring this rule.

As for the words with sonorant and short final vowels, the child could hypothesise that words with these features get the *-etje* allomorph (figure 11), or the *-tje/homorganic* allomorphs (figure 12).

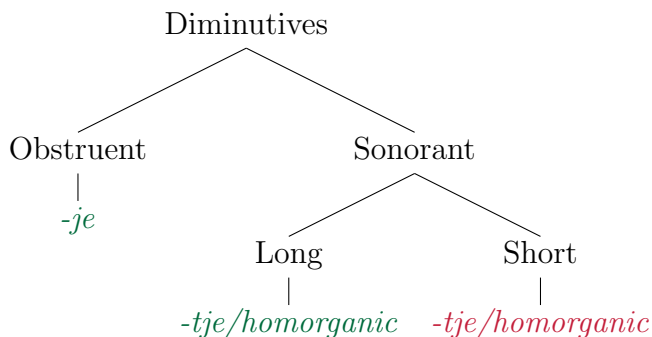


Figure 12: Hypothesis 3.2

Both of these options would allow for $\theta_{170} = 33$ exceptions, as calculated with the TP for the set of words with sonorant endings and short final vowels. The *-etje* rule would result in 108 exceptions, the homorganic rule in 62 exceptions and the *-tje* rule in 72 exceptions. All of these exceed the allowed amount. Once again the child will need to find another solution.

As displayed in figure 13, the child could come up with one final split in the data. A split between words with and without penultimate stress. For the words with penultimate stress (sonorant ending and short final vowel), the rule would either be the homorganic or the *-tje* rule again.

Looking at the frequencies in table 3, the subregularity of words with penultimate stress allows for $\theta_{107} = 22$ exceptions. Depending on whether we take the *-tje* rule or the homorganic rule, the number of exceptions is either 10 or 0. In both cases the rule would be productive and could be accepted by the child.

Again the bars in figure 6, repeated here, are separated based on which suffix the included words should get. Here the children performed surpris-

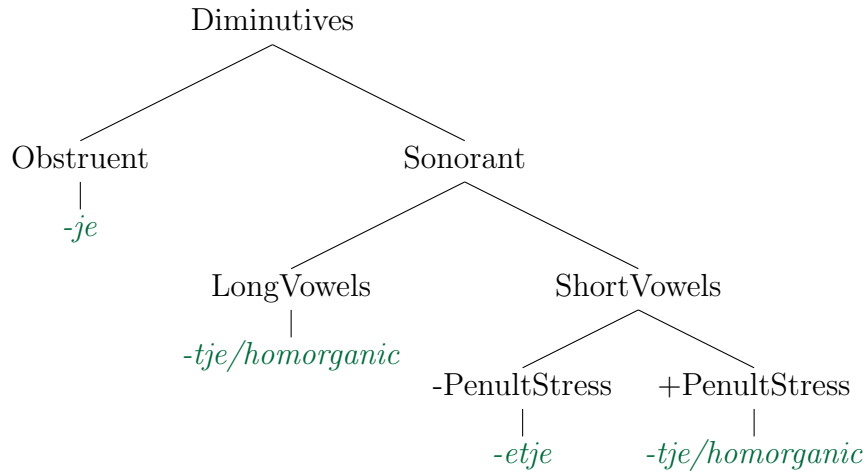


Figure 13: Hypothesis 4

ingly well. Surprisingly because this rule would be acquired relatively late and requires the children to take more features into account (sonority, vowel length *and* stress). The children performed well from the start and only get better with age.




Figure 6: Child production data from the WUG-test: suffixes given to words with sonorant endings, short final vowels and penultimate stress.

Finally, for the words without penultimate stress the rule would be *-etje*. The total amount of words without penultimate stress (and a short and sonorant ending) results in a threshold of $\theta_{63} = 15$. As there is only one word

that does not take the *-etje* suffix, the exceptions definitely stay below the allowed amount.

As is shown in figure 7, repeated here, the children had a harder time with this rule. As this rule requires the learner to take three different phonological features into account, it is also one of the hardest rules to acquire (although they did fine with the nouns with penultimate stress). The graph for the production data of the real words shows that the younger children struggle with the rule, but do acquire it by the age of 9. In the graph with the nonce data it is visible that the youngest children almost never use *-etje*. As they get older, they do use it more often, but still only about 50% of the time.




Figure 7: Child production data from the WUG-test: suffixes given to words with sonorant endings, short final vowels and no penultimate stress.

As was discussed in the previous research section, Boersma (2018) notes that even adults are not always accurate in giving the *-etje* suffix to nonce words (with sonorant endings, long last syllables and no penultimate stress). An adult diminutive production test indicates that they often use the homorganic rule instead of *-etje*. This would indicate that the *-etje* rule is not as productive as the other rules, even for adults.

So, if the TP indicates that the *-etje* rule can be learned, why does it seem as if neither children nor adults have a productive *-etje* rule? Even though these rules are learnable according to TP, the productions of speakers suggest that the underlying rule system is structured differently. Let's backtrack to the first split in the data, between obstruent and sonorant endings, and consider some other rules.

Figure 9, repeated here, shows the hypothesis space after the split between obstruent and sonorant endings. As was discussed before, the child can not




Figure 9: Hypothesis 2.1

allow the *-tje* rule as a productive rule for the nouns with sonorant endings. There are too many exceptions. But if we find some other subsets within the set of nouns with sonorant endings, perhaps the *-tje* rule can be allowed.

Table 9 shows a new distribution of the child directed speech data gathered from PaQu. After the split between sonorant and obstruent, the child could hypothesise that there is a rule that says "if a noun ends in *-m*, the suffix added is *-pje*" and another rule "if a noun ends in *-ŋ*, the suffix added is *-etje*". Both of these rules are acceptable. The respective thresholds are $\theta_{28} = 8$, with 8 exceptions and $\theta_7 = 3$, with 2 exceptions.

Suffix	Total	Obs	Son	-ŋ	-m	Son. without -ŋ and -m ⁶
<i>-je</i>	391	391	0	0	0	0
<i>-tje</i>	225	0	225	0	0	225
<i>-kje</i>	2	0	2	2	0	2 - 2 = 0
<i>-pje</i>	20	0	20	0	20	20 - 20 = 0
<i>-etje</i>	70	8	62	5	8	62 - 5 - 8 = 49
total	708	399	309	7	28	309 - 2 - 20 - 5 - 8 = 274

Table 9: Hand annotated subregularities of the Child directed speech (from CHILDES from PaQu).

In the final column of table 9 both of these sets are excluded from the remaining nouns with sonorant endings. If we recalculate the acceptability of the *-tje* rule with this dataset, the threshold is $\theta_{274} = 48$. The amount of exceptions is 49. While this amount of exceptions is still too high, it's very close to being tolerable.

What I am proposing here is that within the set of nouns with sonorant endings the nouns ending with both *-m* and *-ŋ* have their own productive

⁶liquidae, glides, vowels, other nasals

Suffix	Total	Obs	Son	- η	- m	Son. without - η and - m
- je	247	247	0	0	0	0
- tje	177	0	177	0	0	177
- kje	15	0	15	15	0	15 - 15 = 0
- pje	18	0	18	0	18	18 - 18 = 0
- $etje$	43	5	38	5	2	38 - 5 - 2 = 31
total	500	252	248	20	20	248 - 15 - 18 - 5 - 2 = 208

Table 10: Hand Annotated subregularities of the 500 most common nouns in the Subtlex database.

rules. If these subsets, of nouns ending with $-m$ and $-\eta$, are excluded from the set of sonorant ending nouns, the $-tje$ rule suddenly comes very close to being acceptable as a rule over the rest of the sonorant set (liquidae, glides, vowels etc.). And in another dataset, the $-tje$ rule *is* acceptable, as we will see below.

Here, we are finally taking a look at the Subtlex data (see table 10). With this data we find the same acceptable rule for the nouns ending with $-m$: "if a noun ends in $-m$, the suffix added is $-pje$ ". In this case the threshold would be $\theta_{20} = 6$ and the number of exceptions is 2. However, we find a slightly different rule for words ending with $-\eta$. Namely, "if a noun ends in $-\eta$, the suffix added is $-kje$ ". The threshold is $\theta_{20} = 6$ and the number of exceptions is 5.

The difference between these two rules can be explained by the fact that children see very few nouns that get the $-kje$ suffix (there are only two in the whole CHILDES corpus). So, our child could consider the rule that nouns ending in $-\eta$ get the $-etje$ suffix, with the nouns ending in $-kje$ as exceptions. But as she gets more input, she would realise it is actually the other way around.

Finally, let's also calculate the acceptability of the $-tje$ rule (also excluding the sets of nouns ending in $-m$ and $-\eta$). The threshold is $\theta_{208} = 38$ and there are only 31 exceptions. So, in the Subtlex corpus the rule can actually be accepted! This final system is displayed in figure 14

Coming back to my question, "why does it seem like $-etje$ is not a productive rule?" we can now determine that these calculations indicate that it is not a productive rule, but that all nouns diminutivised with $-etje$ are exceptions.




Figure 14: Hypothesis 5

As was mentioned in the previous research section (2.2), the WUG-test performed with adults by Boersma (2018) indicated that most adults do not follow the expected pattern when it comes to the *-etje* suffix. When the nonce nouns should be receiving the *-etje* suffix, some participants consistently chose *-etje* in the right circumstances (5%), some consistently chose *-tje*, *-kje* or *-pje* in accordance with the place assimilation rules (23%), but for the most part they actually used both these approaches inconsistently (72%).

This distribution of strategies could indicate that not all adults have the same underlying system when it comes to diminutives. The adults that consistently chose *-etje* in the right circumstances seem to have acquired a system like hypothesis 4 (see figure 13). The adults that consistently use place assimilation could have a system like hypothesis 5 (see figure 14). That leaves about 72% of the people that use both strategies inconsistently.

This last group of people, that uses both strategies inconsistently, might be explained by the different ways children and adults acquire language. As mentioned in the previous research section, Boersma (2018) found that as children get older they become more like adults in their nonce substitutions for the *-etje* suffix. This suggests children gradually become more adult-like in their way of learning.

A study by Schuler (2017) shows that while children follow the Tolerance Principle while acquiring language, adults seem to apply probability matching: adults seem to match the distribution of what they see in their input. When children grow older, they become more and more like adults in their way of learning. So, a possible explanation for their inconsistent use of the *-etje* suffix is that it is acquired so late in their development that by the time children acquire it, they already lean more towards probability matching than the TP. And they thus match the distribution of the *-etje* suffixes

that they find in their input.

5.2 Discussion

In this discussion section I will discuss the results of my study in relation to the theoretical background. I will also highlight some limitations of this study, some things I would have done differently in hindsight and some possible directions for further research.

In section 2.2 I discussed various accounts that argue for either *-je* or *-tje* as the default suffix. The researchers who proposed these accounts argued their points by presenting morphophonological rules that show how the other suffixes can be derived from the default rule. The *-je* and *-tje* allomorphs are most often chosen as defaults, because they are the two suffixes with the highest frequency. However, if we follow a strict line of reasoning laid out by the TP, neither can be the default rule.

As was shown in the analysis above, I tested the *-je* suffix as a default for the whole set of diminutives. Although it is the suffix with the highest frequency, this frequency is still not high enough to account for all the other suffixes as exceptions. This means that all suffixes with lower frequencies (among which *-tje*) certainly can not approach productivity.

Because none of the five allomorphs can result from a productive rule with the other four as its exceptions, a learner needs to go looking for a way to split the data (in accordance with the MMP). A split the learner may find is one based on sonority. The set of nouns can be divided into two subsets, one with stems ending with obstruents and one with stems ending with sonorants. In the obstruent set the default suffix is *-je* and in the sonorant set the default suffix is *-tje* (this split is shown in figure 14). This suggests that *-je* and *-tje* are actually both productive rules, within their own subset.

The analysis also shows that a rule system can indeed be found for the Dutch diminutives, that is to say, a rule system can be found in which all exceptions are tolerable. We even found a rule system that would explain why the *-etje* suffix seems to be unproductive (according to Boersma (2018) and in line with my own WUG-test results). Figure 14 shows the rule system in which most of the nouns that receive *-etje* as a suffix are in fact exceptions to the *-tje* rule ("if a noun ends with a sonorant, the suffix is *-tje*"). As *-etje* is not a productive suffix in this system, it explains why children and adults struggle with applying the *-etje* suffix to nonce words.

Although I am very happy with the overall results of this study, there

are some limitations and things I would have done differently in hindsight. There clearly was one obvious limitation to my study that influenced the way I had to conduct the WUG-test. Because of the outbreak of COVID-19, it was not possible to go to a primary school and conduct the tests in person. In order to still be able to conduct the experiment, I made a version that could be performed at home with the help of the parents.

A downside of this method of testing was the low response rate. Many parents were taking care of their children while working from home during the quarantine. So, I think that having to administer the experiment themselves took up too much of most parents' time.

The quality of the audio recordings I did receive were great. The answers the children gave were audible and most children understood what was asked of them. In some of the recordings there seemed to be some difficulties with the PowerPoint presentation that was used to conduct the experiment. I believe these issues were due to the differences in the versions of PowerPoint that were used.

In hindsight, it would have been better to record a video of the Powerpoint presentation and publish this on a video streaming site, like youtube. Sending the parents a link to the youtube video would have been an easier way to conduct the experiment. Most people know how to use this site and it would have been less likely to cause technical difficulties.

Another complication of the PowerPoint presentation was that the form of all the test items was not clear to all the participating children. For example, a few children heard *kroen* instead of *kroem*. As the parents did not know the real forms of the test items either, they could not correct these kind of mistakes. A solution would be to provide only the parents with a list of the test items, so they could then correct the pronunciation in these cases.

Another limitation of this study is related to the corpus research section. As was discussed in section 3.1.3, the features obstruent/sonorant endings and final vowel length were investigated computationally. A limitation of this study was that the last feature, the stress pattern of the words, could not be investigated computationally. Some smaller sets (CDS, Subtlex) were annotated for stress manually. However, in future research, it could be interesting to look at larger datasets. In that case annotating the stress patterns manually would be very time consuming, so a way to computationally determine the stress pattern would be beneficial. A possible source for a measure of these stress patterns could be an online dictionary that includes stress patterns.

Finally, I want to discuss some possibilities for further research. In the section on the Tolerance Principle (2.3), I discussed the Elsewhere Condition. Briefly recapping, it states that exceptions are processed before rules. The exceptions can be visualised as a list that needs to be checked before any regular rule is applied. This list is ordered by frequency: the most frequent exceptions are at the top of the list. In other words, for irregular words the mental lexicon is accessed, while regular inflection needs morphological decomposition.

I briefly discussed a study by Clahsen et al. (1997), cited by Yang, that supports this theory. They performed certain lexical decision tasks to determine if regulars and irregulars are processed differently in the case of German plural markers. They took two suffixes, one that follows a productive rule (-s) and one that does not (-er). For each of these suffixes they took high and low frequency nouns for their lexical decision task. The results showed that while the difference in frequency does not matter for the regulars (-s), it does for the irregulars (-er). The irregulars with a high frequency were processed significantly faster than the irregulars with a low frequency.

These findings indicate that there is a difference between the way regulars and irregulars are processed. Their results are confirmed by a replication study from Sonnenstuhl and Huth (2002). These researchers also performed a lexical decision task and found the same difference between regulars and irregulars. This method could also be useful in determining if there is a rule for the *-etje* suffix.

In the analysis I discussed the view that there might be no productive rule for *-etje*. A lexical decision task could show if there is a difference in processing time for high frequency *-etje* nouns and low frequency *-etje* nouns. The results of such a study could indicate if nouns diminutivised with *-etje* are indeed exceptions.

Another interesting line of research can be found in statistical learning experiments. Newport (2019) and Schuler (2017) describe statistical learning experiments testing the Tolerance Principle. They had both adults and children learn noun plurals of an artificial language in two conditions. In the first condition they were presented with 9 nouns, 5 of which followed a rule when pluralised, the other 4 words were the exceptions (5R4E) that followed no rule. In the second condition only 3 of the nouns followed the rule, and 6 were exceptions (3R6E). After presenting the participants with the stimuli (each noun was presented multiple times in a sentence, both in singular and plural form), a WUG-test was used to assess whether learners had formed a

productive rule.

The expectation was that the rule could only be learned in the 5R4E condition. In this condition the exceptions fall below the threshold ($\theta_9 = 4$), not so in the 3R6E condition.

The results showed that the children learned the productive rule in the 5R4E condition and applied it categorically (they used the productive rule 100% of the time), but that in the 3R6E condition the children did not acquire the rule (they only use it about 17% of the time). The adults seemed to apply probability matching in both conditions. They did apply the rule more often in the 5R4E condition, but not categorically like the children. This suggests there is a difference between the way adults and children acquire a new language.

Schuler (2017) suggests that this difference may be due to the ‘less is more’ theory (proposed by Newport (1990)). Children might acquire rules more easily precisely because their processing power is limited and their vocabulary is still small. Because of their limited processing resources, they need to be as efficient as possible while acquiring language. This could lead to the categorical use of productive rules and consequently the overregularisations we observe.

So far these statistical learning experiments (testing the TP) have been designed to test if the participants can learn one productive rule in the system. As for the Dutch diminutive system it would be interesting to design a statistical learning experiment where the learners would need to find sub-regularities within the data in order to learn the paradigm.

6 Conclusion

In this study I have gathered frequency data of the Dutch diminutive suffixes. I applied the Tolerance Principle to this data and found that a productive rule system can be found with the TP.

This finally answers the research question of this thesis: *Can a productive rule system be found for the Dutch diminutive system with the Tolerance Principle?* Not only did applying the TP result in a productive rule system, the system that was found possibly explains the unproductive *-etje* suffix.

More research is necessary to investigate the system that was found as the underlying rule system of the Dutch diminutives. I hope the results of this study can provide insights for future research in the domain of Dutch diminutives.

Bibliography

- Augustinus, L., Schuurman, I., Vandeghinste, V. & Van Eynde, F. (2014). GrETEL: Searching for breadcrumbs in texts. *Centre for Computational Linguistics KU Leuven*. Retrieved 2nd July 2020, from <http://nederbooms.ccl.kuleuven.be/documentation/EducationalModule-v4b.pdf>
- Bates, D., Maechler, M., Bolker, B. & Walker, S. (2015). Fitting linear mixed effects models using lme4. *Journal of Statistical Software*, 67(1), 1-48.
- Berko, J. (1958). The child's learning of English morphology. *Word*, 14(2-3), 150-177.
- Boersma, T. (2018). *Variability in the acquisition of allomorphs*. PhD Dissertation, University of Amsterdam, Utrecht: LOT publications.
- Boersma, T., Rispens, J., Weerman, F. & Baker, A. (2018). The effects of phonological skills and vocabulary on morphophonological processing. *First language*, 38(2), 147-174.
- Boersma, T., Rispens, J., Weerman, F. & Baker, A. (2019). Acquiring diminutive allomorphs: taking item-specific characteristics into account. *Journal of child language*, 46(3), 567-593.
- Booij, G. (1999). *The phonology of Dutch (vol. 5)*. Oxford, United Kingdom: Oxford University Press.
- Bosch, A. V. D., Busser, B., Canisius, S. & Daelemans, W. (2007). An efficient memory-based morphosyntactic tagger and parser for Dutch. *LOT Occasional Series*, 7, 191-206.
- Clahsen, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and brain sciences*, 22(6), 991-1013.
- Clahsen, H., Eisenbeiss, S. & Sonnenstuhl-Henning, I. (1997). Morphological structure and the processing of inflected words. *Theoretical linguistics*, 23(3), 201-250.
- Daelemans, W., Berck, P. & Gillis, S. (1997). Data mining as a method for linguistic analysis: Dutch diminutives. *Folia Linguistica*, 31(1-2), 57-76.
- Den Os, E. & Harder, R. (1987). De verwerving van de regels voor meervouds- en verkleinwoordsvorming in het Nederlands. *De Nieuwe Taalgids*, 80, 240-250.
- Gillis, S. (1997). *The acquisition of diminutives in Dutch*. Antwerp, Belgium: University of Antwerp.

- Huber, D. (2005). Some notes about the Dutch diminutive suffix. *Government Phonology Workshop*, 2(4), 23-49.
- Keuleers, E. & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior research methods*, 42(3), 627-633.
- Keuleers, E., Brysbaert, M. & New, B. (2010). SUBTLEX-NL: A new frequency measure for Dutch words based on film subtitles. *Behavior Research Methods*, 42(3), 643-650.
- Kiparsky, P. (1973). Elsewhere in phonology. *Reproduced by Indiana University Linguistics Club.*, 93-106.
- Kooij, J. & Van Oostendorp, M. (2003). *Fonologie: uitnodiging tot de klankleer van het Nederlands*. Amsterdam, The Netherlands: Amsterdam University Press.
- MacWhinney, B. (2000). *The CHILDES Project: Tools for analyzing talk*. 3rd edition. Mahwah, USA: Lawrence Erlbaum Associates.
- Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., Xu, F. & Clahsen, H. (1992). Overregularization in language acquisition. *Monographs of the society for research in child development*, 57, i-178.
- Newport, E. L. (1990). Maturation constraints on language learning. *Cognitive science.*, 14(1), 11-28.
- Newport, E. L. (2019). Children and adults as language learners: Rules, variation, and maturational change. *Topics in cognitive science.*, 12(1), 153-169.
- Odiijk, J. (2015). Linguistic research with PaQu. *Computational Linguistics in The Netherlands*, 5, 3-14.
- Peelaerts, C. (2008). Van kleine boot tot scheepje: Experimenteel onderzoek naar de diminutievorming bij Vlaamse kinderen. *Unpublished Manuscript*.
- Quinlan, R. (1993). *Programs for machine learning*. San Francisco, USA: Kaufmann Publishers inc.
- R Core Development Team. (2013). R: A language and environment for statistical computing. Retrieved from <http://www.r-project.org/>
- Schuler, K. D. (2017). *The acquisition of productive rules in child and adult language learners*. PhD Dissertation, Georgetown University.
- Schuler, K. D., Yang, C. & Newport, E. L. (2016). Testing the Tolerance Principle: Children form productive rules when it is more computationally efficient to do so. *CogSci*, 38, 2321-2326.
- Snow, C. E., Smith, N. S. & Hoefnagel-Hohle, M. (1980). The acquisition of some Dutch morphological rules. *Journal of Child Language*, 7(3),

539-553.

- Sonnenstuhl, I. & Huth, A. (2002). Processing and representation of German plurals: A dual mechanism approach. *Brain and Language*, 81(1-3), 276-290.
- Souman, A. & Gillis, S. (2007). A longitudinal study of the acquisition of diminutives in Dutch. *Language acquisition and language disorders*, 43, 183-206.
- Trommelen, M. (1984). *The syllable in Dutch*. Berlin, Germany: Walter de Gruyter.
- Van de Weijer, J. (2002). An Optimality Theoretical analysis of the Dutch diminutive. *Linguistics in the Netherlands*, 19(1), 199-209.
- Yang, C. (2016). *The price of linguistic productivity - How children learn to break the rules of language*. Cambridge, Massachusetts: The MIT Press.

Acknowledgements

First of all, I want to thank previous researchers in the domain of Dutch diminutives and language acquisition. My thanks especially go out to professor Charles Yang. His book on the productivity of language was inspiring and motivated me a lot in the start of my master in linguistics.

Closer to home, I would like to thank dr. Tiffany Boersma and prof. dr. Judith Rispens and their colleagues for allowing me to use their data for my own research.

From the University of Utrecht, I would like to thank dr. Meaghan Fowlie for supervising my internship. She gave me the opportunity to start my own project that eventually turned into a thesis.

Of course, I also want to thank the supervisor of my thesis, prof. dr. Peter Coopmans. Two years ago I formed the idea for this thesis in the Foundations of Acquisition course he taught together with dr. Luisa Meroni, during which I wrote my first paper on the topic. Now that I'm finishing my master, with the same topic and the same supervisor, it feels like I've come full circle. Thank you for all your advice and last minute video calls!

I also have to thank my fellow linguistics students Romy van Drie and Sjoerd Eilander. Thank you Sjoerd for voicing an amazing Leo! Thank you Romy for your thoughts on my analysis! And thank you both for our daily meetings.

For being amazing last readers I also want to thank Yolana van der Horst and my sister Annabella van Tuijl. Thanks you for the helpful feedback on all my work and for correcting stupid mistakes I would have never spotted. And especially thanks to Annabella for listening to adorable audio recordings together and for hyping my project!

Last but not least, I want to thank my parents and sisters for always being there to listen to my complaints and excitement.

Appendix

A

Allomorph	2		3		4		5		6	
<i>-je</i>	235	1310	723	3244	465	1579	193	545	38	85
<i>-tje</i>	81	207	342	1154	213	505	61	117	20	26
<i>-kje</i>	0	0	3	6	0	0	0	0	0	0
<i>-pje</i>	1	1	20	40	9	12	10	14	0	0
<i>-etje</i>	37	90	101	411	64	216	41	90	17	28
rest	12	62	27	232	5	34	5	10	0	0
total	366	1670	1216	5087	756	2346	310	776	75	139

Table 11: Type and token frequencies of the diminutive allomorphs in the CHILDES corpus.

Allomorph	Type		Token
<i>-je</i>	54	(53%)	76
<i>-tje</i>	35	(34%)	45
<i>-kje</i>	1	(1%)	1
<i>-pje</i>	2	(2%)	2
<i>-etje</i>	5	(5%)	5
rest	5	(5%)	5
total	102	(100%)	134

Table 12: Type and token frequencies of the diminutive allomorphs in the Alpino corpus.

Allomorph	Type		Token
<i>-je</i>	428	(53%)	1628
<i>-tje</i>	251	(31%)	646
<i>-kje</i>	3	(0.4%)	7
<i>-pje</i>	19	(2%)	40
<i>-etje</i>	62	(8%)	191
rest	43	(5%)	54
total	806	(100%)	2566

Table 13: Type and token frequencies of the diminutive allomorphs in the CGN corpus.

Allomorph	Type	Token
<i>-je</i>	245 (61%)	549
<i>-tje</i>	119 (30%)	204
<i>-kje</i>	1 (0.2%)	1
<i>-pje</i>	6 (1%)	6
<i>-etje</i>	25 (6%)	47
rest	5 (1%)	10
total	401 (100%)	817

Table 14: Type and token frequencies of the diminutive allomorphs in the Lassy Klein corpus.

Allomorph	Type	Token
<i>-je</i>	1157 (56%)	76758
<i>-tje</i>	620 (30%)	22481
<i>-kje</i>	73 (4%)	195
<i>-pje</i>	57 (3%)	1388
<i>-etje</i>	164 (6%)	6562
rest	4 (8%)	11
total	2075 (100%)	107395

Table 15: Type and token frequencies of the diminutive allomorphs in the CLEF corpus.

Allomorph	Type	Token
<i>-je</i>	612 (58%)	21551
<i>-tje</i>	316 (30%)	5366
<i>-kje</i>	5 (0.5%)	10
<i>-pje</i>	26 (2%)	968
<i>-etje</i>	87 (8%)	1375
rest	1 (0.1%)	2
total	1047 (100%)	29272

Table 16: Type and token frequencies of the diminutive allomorphs in the Dutch Web Corpus corpus.

Allomorph	Type		Token
<i>-je</i>	278	(59%)	1045
<i>-tje</i>	136	(30%)	357
<i>-kje</i>	6	(1%)	6
<i>-pje</i>	11	(2%)	21
<i>-etje</i>	38	(8%)	103
rest	1	(0.2%)	3
total	470	(100%)	1535

Table 17: Type and token frequencies of the diminutive allomorphs in the Eindhoven corpus.

Allomorph	Type		Token
<i>-je</i>	1323	(56%)	240519
<i>-tje</i>	713	(30%)	72972
<i>-kje</i>	68	(3%)	233
<i>-pje</i>	58	(2%)	4046
<i>-etje</i>	188	(8%)	17668
rest	4	(0.2%)	105
total	2354	(100%)	335543

Table 18: Type and token frequencies of the diminutive allomorphs in the Groot Lassy: Krant corpus.

Allomorph	Type		Token
<i>-je</i>	851	(56%)	63795
<i>-tje</i>	487	(32%)	20487
<i>-kje</i>	15	(1%)	67
<i>-pje</i>	39	(3%)	1682
<i>-etje</i>	132	(9%)	6966
rest	2	(0%)	2
total	1526	(100%)	92999

Table 19: Type and token frequencies of the diminutive allomorphs in the Groot Lassy: Wiki corpus.

Allomorph	Type	Token
<i>-je</i>	911 (56%)	103350
<i>-tje</i>	530 (33%)	32049
<i>-kje</i>	15 (1%)	105
<i>-pje</i>	39 (2%)	2360
<i>-etje</i>	132 (8%)	10738
rest	1 (0%)	3
total	1628 (100%)	148605

Table 20: Type and token frequencies of the diminutive allomorphs in the NL-wiki: 2017 corpus.

Allomorph	Type	Token
<i>-je</i>	247 (64%)	2958
<i>-tje</i>	103 (27%)	526
<i>-kje</i>	0 (0%)	0
<i>-pje</i>	7 (2%)	134
<i>-etje</i>	28 (7%)	146
rest	1 (0.2%)	1
total	386 (100%)	3765

Table 21: Type and token frequencies of the diminutive allomorphs in the Wabliet corpus.

B

Rule	CDS θ	CDS e	Subtlex θ	Subtlex e
<i>Everything</i> \rightarrow <i>-je</i>	$\theta_{708} = 107$	317	$\theta_{500} = 80$	253
<i>Obstr.</i> \rightarrow <i>-je</i>	$\theta_{399} = 66$	8	$\theta_{252} = 45$	5
<i>Son.</i> \rightarrow <i>-tje</i>	$\theta_{309} = 53$	84	$\theta_{248} = 44$	71
<i>Son.</i> \rightarrow <i>homorganic</i>	$\theta_{309} = 53$	62	$\theta_{248} = 44$	38
<i>Son./Long</i> \rightarrow <i>-tje</i>	$\theta_{139} = 28$	12	$\theta_{121} = 25$	12
<i>Son./Long</i> \rightarrow <i>homorganic</i>	$\theta_{139} = 28$	0	$\theta_{121} = 25$	0
<i>Son./Short</i> \rightarrow <i>-tje</i>	$\theta_{170} = 33$	72	$\theta_{127} = 26$	59
<i>Son./Short</i> \rightarrow <i>-etje</i>	$\theta_{170} = 33$	108	$\theta_{127} = 26$	89
<i>Son./Short</i> \rightarrow <i>homorganic</i>	$\theta_{170} = 33$	62	$\theta_{127} = 26$	38
<i>Son./Short/+Penult</i> \rightarrow <i>-tje</i>	$\theta_{107} = 22$	10	$\theta_{91} = 20$	24
<i>Son./Short/+Penult</i> \rightarrow <i>homorganic</i>	$\theta_{107} = 22$	0	$\theta_{91} = 20$	3
<i>Son./Short/-Penult</i> \rightarrow <i>-etje</i>	$\theta_{63} = 15$	1	$\theta_{36} = 10$	1
<i>Ending with -m</i> \rightarrow <i>-pje</i>	$\theta_{28} = 8$	8	$\theta_{20} = 6$	2
<i>Ending with -ŋ</i> \rightarrow <i>-kje</i>	$\theta_7 = 3$	5	$\theta_{20} = 6$	5
<i>Son. without -m and -ŋ</i> \rightarrow <i>-tje</i>	$\theta_{274} = 48$	49	$\theta_{208} = 38$	31

Table 22: Comparison table of the Tolerance principle calculations with the Child directed speech and 500 most common Subtlex nouns data.

C




Figure 15: Child production data from Boersma (2018): suffixes given to words with obstruant endings.




Figure 16: Child production data from Boersma (2018): suffixes given to words with sonorant endings and long final vowels.




Figure 17: Child production data from Boersma (2018): suffixes given to words with sonorant endings, short final vowels and no penultimate stress.




Figure 18: Child production data from Boersma (2018): suffixes given to words with sonorant endings, short final vowels and penultimate stress.