More than Islands of Regularity: An Investigation of the Sensitivity of Morphological Decomposition to Higher-level Linguistic Properties

Roberto Petrosino, PhD

University of Connecticut, 2020

Abstract

A wealth of psycholinguistic evidence has shown that words, before being visually recognized, decompose into smaller units which seem to correspond to morphemes (Rastle et al., 2000, 2004). Such a procedure of *morphological decomposition* seems to occur in words that are made of more than one morpheme, independently of whether they are semantically transparent (e.g., *driver*, which means "someone who drives") or not (e.g., *brother*, which does not mean "someone who broths"). Decomposition, however, does not seem to occur in words that contain a root plus an additional, non-morphemic string (e.g., *brothel*, where *el* is not an English suffix). Current models of visual word processing assume the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS, whereby decomposition obligatorily relies on *islands of regularity* – namely, statistical orthographic regularities across letter strings, so that a string such as *er* is detected as a single morpho-orthographic unit, but *el* is not (Rastle and Davis, 2008). As such, decomposition is also argued to occur at early stages of word processing and *before* contact with the lexicon — namely, before semantic information is accessed, so that decomposition occurs regardless of semantic transparency. This dissertation builds on these findings and primarily elicits the visual masked priming response to test the sensitivity of decomposition to any of the following higher-level linguistic properties (in addition to islands of regularities): phono-orthographic syllabification (along with whole-word lexicality, though indirectly), whole-word frequency, phonological conditioning, and syntactic selectional restrictions. By doing so, this dissertation additionally provides insights regarding the unfolding of linguistic properties during visual word processing. The results reported here suggest that decomposition is affected by whole-word lexicality and whole-word frequency. Given that these two properties are generally considered to be accessed at late stages of processing, these findings seem to be arguing against current models. A model of early visual decomposition is proposed that is able to account for these findings without necessarily challenging the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS. In this model, decomposition accesses whole-word lexicality and frequency through a multi-step mechanism that first generates multiple possible morpho-orthographic decomposition patterns of the visual stimulus and then evaluates them in parallel in order to choose the optimal candidate for activation.

More than Islands of Regularity: An Investigation of the Sensitivity of Morphological Decomposition to Higher-level Linguistic Properties

Roberto Petrosino

B.A., Università degli Studi di Napoli Federico II, 2009
M.A., Università degli Studi di Siena, 2012
M.A., University of Connecticut, 2017

A Dissertation
Submitted in Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy
at the
University of Connecticut

2020

APPROVAL PAGE

Doctor of Philosophy Dissertation

More than Islands of Regularity: An Investigation of the Sensitivity of Morphological Decomposition to Higher-level Linguistic Properties

presented by

Roberto Petrosino, B.A., M.A., M.A.

Major Advisor: ____
Andrea Calabrese

Major Advisor: ____
Jon Sprouse

Associate Advisor: ____
Jonathan D. Bobaljik

University of Connecticut

2020

Acknowledgements

Two and half years ago, I decided to grow my hair out. It all began as a fun thing to do. Then, it ended up becoming a daily symbol and reminder of my commitment to the project reported in this dissertation, and, more generally, to the linguistic endeavor. When you read acknowledgments in others' dissertations as a graduate student, you really do not understand what all the fuss is about. In hindsight, I now realize that it is the *process* that makes it "interesting." The process of accepting that you knew nothing before and will probably not know enough after. The process of wavering over quitting at each and every step of the way. The process of seeing progress even in the tiny daily result. My hair finally reached the length of roughly 10 inches when I wrote the conclusions and decided to cut it. It symbolically represents the length of this process and this dissertation is its final, much longed-for outcome.

The journey that this dissertation represents would have been so much different without my advisors. Thank you, Andrea, for guiding me through the intricacies of phonological and morphological theory – and Salentino cuisine, of course. Thank you, Jon, for your painstaking assistance, patience, and care you showed throughout the whole process, right from the start, and even in the darkest hours. Thank you, Jonathan, for all discussions on morphological theory we have had and the climbing days we've spent together. I also wish to thank all the graduate students and professors of the Linguistics and the Psychology departments I was so lucky to engage with.

The last few years have been a roller-coaster of ever-changing plans. I am grateful to all of the people who had to deal with me throughout these messy times – COVID-19, warts, and all. First and foremost, I could not have done this without the silent, unconditional support from my family. Grazie di continuare ad avere fiducia cieca in me e nelle mie aspirazioni. Questo lo dedico a voi.

Claudia & Simon, thank you for your dedication to pulling me out of the cage I had built myself in for the past year. Grazie per la pazienza, per ogni gita dell'ultimo minuto, per ogni sabato mattina dedicato alle pulizie domestiche in quelli che sono stati – se posso dire – i "nostri" 40 mq, e per tutte le "solite insalate" che abbiamo preparato insieme. Grazie anche ai nostri compagni di mirobolanti avventure canine: Gina, Luca e Michela.

A huge thank you to my friends in the US: Chantale, Christos, Dabney, Emma, Fabiana, Ivana, Neda, Vanessa, Sabine, Sandra, and Steven. I cherish every home-made pizza, every Thai food, every 5-minute and 2-hour tomato sauce we have made and eaten together. Grazie a voi, amici della vita, per essere presenti, fisicamente e col cuore, soprattutto nelle avversità. Sapete chi siete, vi voglio bene.

Athletics have been an important part in my life and helped me find resilience and reset. I am proud that I have been part of two great Crossfit communities for the last few years. The Columbia Crossfit Community (Columbia, CT) first welcomed me into this world. Thank you to the coaches: Pat and Andrea Stone, Sam Gallmeyer, Shelley Jones, and all the fellow athletes (no list could actually give justice to this amazing bunch of people) for your guidance and friendship. The Crossfit Due Torri Community (Bologna, Italy) has been my go-to stress-relief spot for the whole past year. Grazie ai coaches e amici: Chiara, Carlét, Gube, Nau, Mattia e Elisa; e a tutti i membri di questa splendida community – in particolare i pazzi del "*7am party*": Almu, Andrea "Chef", Enrico, Klajdi, Farid, Franci, Melo e Vero "Genny". Grazie per aver supportato ogni mattina me e le mie urla durante i workout. Siete speciali.

Throughout the last few years, you all have made my life better, every day. I owe you big time.

Contents

List of figures

List of tables

catholic, *n.*
A person addicted to cats.
retired, *adj.*
Tired again.
fasten, *v.*
Make or become fast.

— *Decomposition matters*. And you may never realize it.

Chapter 1.

Morphological decomposition: what, when, how

1.1. Decomposition and the lexicon

The past few decades have yielded psycholinguistic research that has suggested that before being recognized, words are decomposed into smaller units that seem to correspond to morphemes (among others, Rastle et al., 2000, 2004; Fiorentino and Poeppel, 2007; Solomyak and Marantz, 2010; Lewis et al., 2011). This procedure, called MORPHOLOGICAL DECOMPOSITION, has been found to adhere to the following three-way pattern: (i) it occurs in words that are made of more than one morpheme (morphologically transparent words; e.g., *driver→{drive}-{er}*, meaning "someone who drives"); (ii) it occurs in words that appear to be made of more than one morpheme, but are actually monomorphemic (morphologically opaque words; e.g., *brother→{broth}-{er}*, even if it does not mean "someone who broths"); (iii) it does not occur in words that contain a root plus additional phonemes that do not form a morpheme (e.g., *brothel↛{broth}-{el}*, where *el* is not an English suffix). This pattern of results suggests that the morpho-orthographic form of morphemes (that is, the phonological realization and the orthographic sequence of letter strings associated with a given morpheme) drives decomposition (since *-er* elicits decomposition, but *-el* does not), whereas semantics does not seem to matter (since *-er* elicits decomposition even in morphologically opaque, monomorphemic words like *brother*). This dissertation builds on these findings and asks, in addition to the morpho-orthographic form of morphemes, what other linguistic properties affect decomposition (Figure 1.1, dotted section of the rectangle).

(1) EMPIRICAL QUESTION ON DECOMPOSITION
In addition to the morpho-orthographic forms of the morphemes, what linguistic properties, if any, affect decomposition?

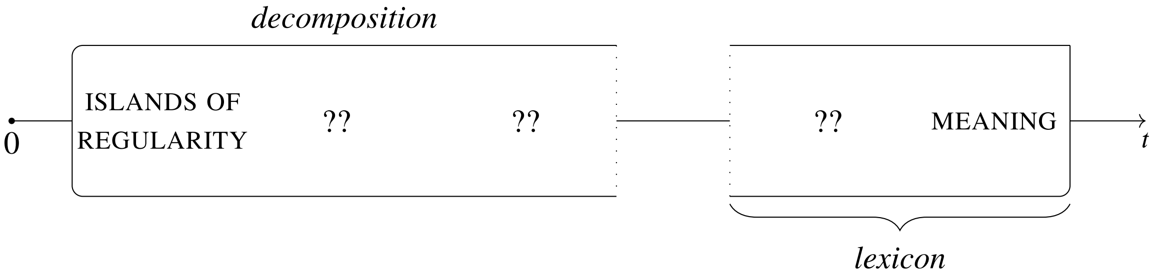

Figure 1.1.: Time-course of the unfolding of word information in word processing (initial version). After being presented (0), a word is processed though a series of stages unfolding over time t. At each stage, a given set of properties is assumed to be available for processing purposes. *Decomposition* reportedly accesses the form of the constituent morphemes of a word via an orthography-based segmentation procedure, before contacting the *lexicon* (i.e., before accessing, for example, meaning). The '??' are provisional placeholders for the properties that will be tested in the upcoming chapters.

In asking (1), this dissertation indirectly provides insights regarding the long-standing debate on the organization of lexical properties during word recognition. In the traditional view, all lexical properties are contained in a single place – the *mental lexicon*. Under this view, lexical access is all-or-none: if one property is accessed, all of other properties are accessed too. However, this prediction is challenged by the psycholinguistic evidence showing that decomposition is sensitive to the orthographic form of morphemes, but not to their semantic properties (see sec. 1.3 for further details).

(2) MODULAR MENTAL LEXICON
All lexical and linguistic properties associated to morphemes are stored in one place – the *mental lexicon*.

To accommodate (2), psycholinguistic models of lexical access have argued that decomposition is not sensitive to the actual stored morpho-orthographic form of morphemes; rather, it is sensitive to statistical regularities in letter strings that are part of the reader's knowledge ("morpho-orthographic units," also known as "islands of regularity": Rastle and Davis, 2008; see sec. 1.4.2). For example, the string $ e r $ is very likely to form a morpheme and therefore triggers decomposition; the string $ e l $ is, instead, unlikely to form a morpheme and therefore does not trigger decomposition. This perspective is still compatible with the modular view of the mental lexicon (2) and, crucially, leads to the prediction (3), whereby decomposition occurs before any undeniably lexical property could be accessed ("contact with the lexicon"). As of today, (3) is commonly accepted, and also corroborated by recent neuromagnetic evidence (among others, Fruchter and Marantz, 2015).

(3) MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS
Morphological decomposition relies on morpho-orthographic regularities and occurs before any contact with lexicon.

Admittedly, under (3), the question in (1) is theoretically vacuous; if decomposition occurs before any contact with the lexicon, the set of linguistic properties potentially affecting it must then be rather small, if not empty. However, as we will see in sec. 1.4.2, current models of decomposition do not strictly comply with (2). These models, more or less explicitly, seem to actually adopt a *distributed*, rather than modular, view of the mental lexicon, whereby availability of linguistic properties of morphemes is spread out across different stages of the recognition process. Therefore, the distributed view of the mental lexicon, as formalized in (4), validates the investigation reported here, while still assuming segmentation relies on the aforementioned orthographic regularities.

(4) DISTRIBUTED MENTAL LEXICON
All lexical and linguistic properties associated to morphemes are accessed at different stages of the word recognition process.

In this sense, while primarily addressing the question (1), this dissertation aims to (i) situate each property tested in models of lexical access, and, indirectly, (ii) situate decomposition within word processing. To this end, we* will test the sensitivity of decomposition to a representative subset of higher-level linguistic properties. We will primarily use the masked priming design, which has been the primary experimental tool to probe morphological decomposition. We primarily focus on decomposition in the visual modality (Chapters 2-5). Our results seem to be suggestive of a decomposition procedure that is far more complex than previously thought. On one hand, in contrast with what is currently assumed, the segmentation procedure does not (at least solely) rely on morpho-orthographic regularities (see Chapter 4). On the other hand, the recognition of decomposed morphemes seems to be sensitive to the lexicality status of the whole-word visual string (i.e., whether or not it forms a real word; see Chapter 2) and whole-word frequency (i.e., *dominance*; see Chapter 3), but not to syntactic affixal restrictions (see Chapter 5). We take stock of the whole investigation in Chapter 6, where we discuss the theoretical implications of the results obtained. In particular, these results seem to challenge the modular view (2) in favor a distributed view (4), in which some properties may be accessed at early stages of processing (thus affecting decomposition), and some other properties may be only available at later stages. We claim that such a seemingly clear divide is due to the time needed for the different properties to be properly accessed and computed. A competition-based model of decomposition is therefore proposed, in which multiple decomposition pattern candidates of the same visual stimulus are evaluated in parallel in order to identify the most likely pattern to decompose the stimulus.

This chapter is meant to give a concise, but hopefully thorough introduction to the literature on visual morphological decomposition and, more in general, visual word processing. Section 1.2 reviews the features of the priming design and the mechanisms proposed to account for it. Section 1.3 summarizes the main results from the priming literature on visual morphological

---

**Disclaimer.* This dissertation reports an experimental investigation that, while being completely envisioned, designed, and fulfilled by the PhD candidate, has benefited from countless discussions and exchanges with the members of advisory committee. The usage of *we* should therefore not be taken as pretentious, rather as a way to acknowledge the importance of collaborative work in the present endeavor. Nonetheless, the PhD candidate assumes full responsibility for any error present in these pages.

decomposition. Section 1.4.2 reviews four major models of lexical access; these models will be used to interpret the results of the experiments reported in the following chapters. Section 1.5 describes the general experimental methodology adopted in the dissertation. Finally, section 1.6 briefly outlines the following chapters.

1.2. The priming design

This dissertation primarily uses the lexical decision priming paradigm. Generally speaking, the priming paradigm consists of two stimuli being presented one after the other; the first stimulus is called *prime*, the second is called *target*. The experimental task always concerns the target stimulus; in this dissertation, the main task will always be a lexical decision task, in which subjects are asked to decide whether the target stimulus is an existing word by pressing a response button box. It has been shown that recognition of the target stimulus is faster when it is related to the prime stimulus to some degree: e.g., the target word *cat* is recognized as a word faster when it is preceded by a semantically related prime word like *kitten* than when it is preceded by an unrelated prime word like *hand*.

Priming experiments typically have two main conditions. A RELATED CONDITION, in which the prime and the target are related along some dimension (e.g., they may be identical: *cat-CAT*; orthographically similar: *cap-CAT*; or morphologically related: *cats-CAT*) and an UNRELATED CONDITION (baseline condition), in which the prime and the target are completely different and, crucially, unrelated, items (e.g., *hand-CAT*).[1] There are two main designs in which the related and unrelated conditions may be arranged. The basic design is the *Latin Square* design. Latin Square designs involve choosing the whole set of target stimuli that will be presented to subjects (e.g., words), splitting them in the number of conditions to be tested (e.g., related and unrelated condition), and finally rotating the groups of targets through the conditions to create different *lists*. Subjects are presented with one list only, so that they see exactly the same targets (which may be preceded by a different prime type), regardless of the list they were assigned. This kind of design has the advantage of naturally controlling for stimulus properties (e.g., words may be controlled for frequency, orthographic/morphological/phonological length, neighborhood density, etc.).[2] On the other hand, Latin square designs are not often used in lexical decision tasks because they are only possible when all target stimuli can be paired with all prime stimuli across conditions. This is rare, and mostly happens when the primes are pseudo-words. The standard design that is typically used instead involves creating each related condition separately (e.g., morphological condition: *driver-DRIVE*) and then arranging the stimuli across two different lists, such that, in each list, half of the targets of each condition are preceded by the corresponding related primes and the other half are preceded

---

[1] In visual priming experiments, targets are commonly presented in UPPERCASE and primes in lowercase. This is to ensure that lexical decisions are not influenced by merely graphical similarity between the prime and the target. For the sake of clarity, I will maintain this typographical difference between primes and target throughout the entire dissertation.

[2] It should be noted that the actual extent to which lexical information of words may actually affect priming effects has not yet been shown. This is though common practice in the field, as a way to remove any potential confounding effect resulting from it.

by the corresponding unrelated primes. This way, each list consists of several self-contained sub-designs; in each sub-design, the targets are either related to the prime (related condition: *driver-DRIVE*) or unrelated to the prime (unrelated condition: *lovely-DRIVE*). While this kind of design offers more freedom (because it allows for counterbalancing across word-word pairs, unlike Latin Square designs), it has the drawback of not naturally controlling for the lexical properties of the stimuli. This means that primes and targets of each condition must be controlled for any given lexical information (e.g., frequency or length) across all conditions, in order to guarantee similar distributions (but see fn. 2).

Once response times (RTs) to target recognition have been recorded, they are averaged together and used to calculate priming effects for each condition. *Priming effects* are calculated as the difference between the averaged RT to the related condition and the averaged RT to the unrelated (i.e., baseline) condition. According to the most common practice in priming studies, the sign of the difference is then reversed for clarity.

(5) $priming = -(RT_{related} - RT_{unrelated})$

If positive, the priming magnitude calculated in (5) expresses the amount of *facilitation* (expressed in ms) that the presentation of a related prime exerted on the recognition of the correspondent target. If negative, the priming magnitude calculated in (5) expresses, instead, the amount of *inhibition* (expressed in ms) that the presentation of a related prime exerted on the recognition of the correspondent target.

Morphological decomposition has been mainly investigated by adopting a sub-type of the visual priming design. This design is called *masked* priming and was pioneered by Forster and Davis (1984). In this design, a *mask* (that is, a series of hashmarks, #####) is presented for 500 ms before the prime stimulus, which is usually presented for 30-70 ms. At such a short duration, primes are not consciously recognizable to subjects. The advantage of the masked priming technique is two-fold: (i) it prevents facilitation effects due to strategic processes triggered by episodic memory and (ii) it elicits them before full lexical access of the prime word is complete. The target immediately follows the prime and usually stays on the screen until the subject performs the task on the target (e.g., lexical decision). As already discussed in sec. 1.1, we will follow the literature on morphological decomposition and use the masked priming design as the primary tool to explore the our research questions. In doing this, we consider the masked priming design as the right tool to probe the linguistic sophistication of early decomposition, under the assumption that prime masking prevents the prime stimulus from being further processed after target presentation. As we will see in Chapter 6, this methodological assumption about the experimental design used is crucial for the results to be properly interpreted.

1.2.1. The mechanics of priming

There are two main interpretations of priming. In the *retrospective interpretation* of priming, priming is seen as arising *after* target presentation, as a function of the speed in which the target can be integrated with the prime (Neely, 1991; Forster, 1981). In the *prospective interpretation* of priming, priming is seen as arising *before* target presentation. In this dissertation, we adopt the prospective interpretation of priming, as it has been widely implemented in most models of lexical access (e.g., the multiple read-out model, MROM: Grainger and Jacobs, 1996; the dual-route cascaded model, DRC: Coltheart et al., 2001).

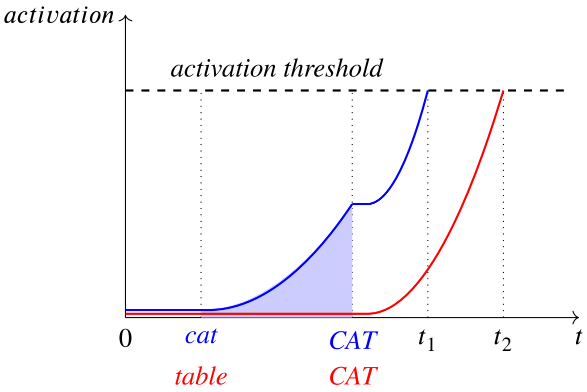

Figure 1.2.: Priming as explained in a spreading activation model. The lines represent the level of activation of the cortical representation of the target word *CAT* over time. When a prime is presented, the activation level of its representation starts increasing. If the prime word is related to the following target word (blue line, for which the word *cat* is presented as both prime and target), the activation threshold of the target word is reached faster, as recognition of the target word occurs at t_1. The blue-shaded area represents the magnitude of the head start given to target recognition by the related prime. If the prime word is related to the following target word (red line, for which the prime word is *table* and the target word is *CAT*), the activation threshold of the target word cannot benefit from the level of activation of the representation of the prime; therefore, recognition of the target word occurs at t_2.

The most prominent way of explaining priming involves activation of cortical representations as in *spreading activation models* (e.g., interactive activation models: McClelland and Rumelhart, 1981). In these models, computation occurs through activation of nodes (also called units), which are the analogs of brain neuronal cells. Node *activation* refers to the action potential spike that the node may generate; node activation typically spreads to all other nodes connected to it. The activation value of a given node is calculated via an activation function, which sums the activation values of the nodes connected to it. It is generally assumed that node activation is contingent on whether the activation value of the node reaches a fixed *activation threshold* (see Kriesel, 2007, for an accessible introduction to neural networks). Priming can, therefore, be explained as result of the activation of the node corresponding to the target prior to its actual presentation. The cortical representation of a prime stimulus is connected, or overlaps with, the cortical representation of the related target stimulus, so that when the prime is presented, it also activates (at least part of) the representation of the target (for a review, see Forster et al., 2003; Forster, 1999; Taft, 1994). Prior presentation of a related prime therefore helps the target reach the activation threshold; this results in facilitating its recognition (Figure 1.2). As we will see, the activation mechanism is extensively used to describe node activation in the models of lexical access we are considering in this dissertation (see sec. 1.4.2).

1.3. Morphological decomposition as revealed in visual masked priming: a review

In this section, I briefly summarize the main results on morphological decomposition as reported in the visual priming literature. The review below represents the bulk of the findings we will be referring to throughout the dissertation. Table 1.1 at page 9 is meant to serve as a summary of the relevant results the reader can go back to whenever needed.

1.3.1. Root priming

Priming experiments on morphological processing have primarily focused on root priming, in which the prime and target words share the same root (e.g., *driver-DRIVE*). The studies by Rastle et al. (2000) and Rastle et al. (2004) were among the first reporting root priming in a masked design. Rastle et al. (2000) looked at priming effects in four different conditions: in the morphological condition, the prime derives from the target (e.g., *adapter-ADAPT*); in the semantic condition, the prime and the target are semantically related (e.g., *cello-VIOLIN*); in the identity condition, the prime and the target are the same word (e.g., *church-CHURCH*); finally, in the orthographic condition, the prime shares the first letters with the target (e.g., *electrode-ELECT*). The four conditions were tested across three different prime-target SOA (Stimulus Onset Asynchrony, namely the time between the presentation of the prime and the target) durations. At the SOAs of 42 and 72 ms, pairs in both the morphological and identity conditions showed priming effects, wheres pairs in the semantic and the orthographic condition did not. At the SOA of 230 ms, pairs in the identity, morphological, and semantic conditions showed priming effects, while pairs in the orthographic condition did not. The fact that unlike semantic priming, morphological priming arises even at short SOAs has been interpreted as confirmation of the hypothesis that morphological decomposition occurs subliminally (namely, even at short SOAs, in which the prime duration is so fast that the presentation of the prime before the target cannot be acknowledged by subjects) and cannot be induced by semantic relatedness (since semantically related, but morphologically unrelated pairs such as *cello-VIOLIN* did not show priming at short SOAs). Building on these results, Rastle et al. (2004) looked at subliminal priming elicited in the following three conditions: the morphologically transparent condition, in which prime and target words were morphologically, semantically and orthographically related (*alarming-ALARM*); the morphologically opaque condition, in which prime and target words looked like the former was a derived form of the latter, but were not semantically related (*brother-BROTH*); finally, the orthographic condition, in which prime and target words were only orthographically related so that the target word was contained in the prime word (*brothel-BROTH*; henceforth, ‘pseudo-affixed condition’). With a 42ms-long SOA, priming effects were found in the transparent and opaque conditions, and not for orthographic condition. Taken together with Rastle et al. (2000)’s results, these results further suggest that decomposition occurs only in presence of real morphemes, since *brother* primes *BROTH* as much as *alarming* primes *ALARM* (even though *brother* and *broth* are not semantically related to one another at the syncronic level), but *brothel* does not prime *BROTH* (possibly because *el* is not a real English morpheme).

The studies that were subsequently published can all be seen as aiming to further understand the mechanism of decomposition. Morris et al. (2011) compared the priming effects in three conditions: a morphological condition (*flexible-FLEX*); a pseudo-morphological condition, in which the pseudo-word prime consisted of the corresponding monomorphemic target suffixed with a real, but syntactically illicit morpheme (e.g., *flexity-FLEX*); and a pseudo-affixed condition, in which the pseudo-word prime consisted of the corresponding monomorphemic target suffixed with a phonologically licit, non-suffixal ending (e.g., *flexire-FLEX*). All three conditions were found to elicit priming effects whose magnitude did not vary across conditions. These results are at odds with the results reported in Rastle et al. (2000, 2004); in particular, the pseudo-affixed condition eliciting priming is at issue with the argument that decomposition occurs only when the prime contains extant morphemes. We will get back to this apparent contradiction below (sec. 1.3.3).

1.3.2. Affix priming

Only a few studies have looked at affix priming. Dominguez et al. (2010) tested priming effects in response to Spanish prefixes by comparing two conditions: a prefixed condition, in which prime and target shared the same prefix (e.g., *infeliz-INCAPAZ* ‘unhappy-incapable’) and a syllabic condition, in which prime words had the first syllable being phonologically and orthographically identical to the prefix of the corresponding target (*industria-INCAPAZ* ‘industry-incapable’). In line with Rastle et al. (2004), Dominguez et al. (2010) found significant priming effects for both the prefixed and the syllabic condition only at the short SOA of 33 ms (experiment 1a); the effects to the syllabic condition gradually decreased as the SOA increased, with the effects to the prefixed condition not varying throughout (experiment 1b, c). In English, Chateau et al. (2002) found priming effects for both transparently prefixed (*impatient-IMMOBILE*) and opaquely prefixed (*imitate-IMMOBILE*) pairs, but no priming effects for orthographic pairs sharing the initial letters (*pursuit-PURPLE*). Moreover, transparent pairs elicited priming regardless of whether the prefix was high or low in spelling-meaning consistency (for example, im- is high consistency, whereas *con-* is low consistency), whereas opaque pairs elicited priming only when affixed with high-consistency prefixes (e.g., *context-CONFORM* did not prime).

As for suffix priming, Andoni Duñabeitia et al. (2008) reported a series of Spanish experiments comparing a morphological condition, in which prime and target stimuli shared the same suffix and an orthographic condition in which prime and target stimuli shared the same last three letters. In the morphological condition, the target words (e.g., *IGUALDAD* ‘equality’) were preceded by primes being the sole corresponding suffixes (*dad*; experiment 1), the suffixes along with non-alphabetic symbols (*%%%%dad*; experiment 2), or words sharing the same suffix (e.g., *brevedad* brevity; experiment 3). Similarly, in the orthographic condition, the target words (e.g., *CERTAMEN* contest) were preceded by the primes that were only the corresponding pseudo-suffixes (*men*; experiment 1), the pseudo-suffixes attached to non-

PRIMING TYPE	CONDITION	EXAMPLE	PRIME LEXICALITY	PRIMING?	REFERENCE
ROOT PRIMING	transparent	*driver-DRIVE*	W	Y	Rastle et al. 2000, 2004
	opaque	*brother-BROTH*	W	Y	Rastle et al. 2004
	root-containing	*brothel-BROTH*	W	N	
	root+suffix	*flexity-FLEX*	NW	Y	Morris et al. 2011
	root+non-suffixal ending	*flexire-FLEX*	NW	Y	
PREFIX PRIMING	Pseudo-prefixed	*industria-INCAPAZ*	W	Y	Dominguez et al. 2010
	prefix+rooted	*infeliz-INCAPAZ*	W	Y	
	transparent	*impatient-IMMOBILE*	W	Y	Chateau et al. 2012
	opaque	*imitate-IMMOBILE*	W	Y	
	orthographic	*pursuit-PURPLE*	W	N	
SUFFIX PRIMING	suffix only	*dad-IGUALDAD*	NW	Y	Duñabeitia et al. 2008
	pseudo-suffix only	*men-CERTAMEN*	NW	N	
	transparent	*brevidad-IGUALDAD*	W	Y	
	orthographic	*volumen-CERTAMEN*	W	N	
	suffix	*sheeter-TEACHER*	NW	Y	Crepaldi et al. 2015
	word-ending	*poolel-BARREL*	NW	N	
COMPOUND PRIMING	non-head constituent	*flagpole-FLAG*	W	Y	Fiorentino & Fund-Reznicek 2009
		hallmark-HALL	W	Y	
		plankton-PLAN	W	N	
	head-final constituent	*classroom-CLASS*	W	Y	
		honeymoon-MOON	W	Y	
		battalion-LION	W	N	
	pseudo-word+word	*slegrack-RACK*	NW	Y	Fiorentino et al. 2015
	word+word	*drugrack-RACK*	NW	Y	

Table 1.1.: Summary of results of visual priming studies.

alphabetic symbols (i.e., *%%%%men*; experiment 2), or words sharing the same ending (e.g., *volumen*; experiment 3). In all three masked experiments, priming effects were found for all morphological conditions only, regardless of whether the suffixes were presented stand-alone (experiment 1), together with non-alphabetic symbols (experiment 2), or embedded in a real bimorphemic word (experiment 3). This suggests that decomposition occurs in the presence not only of alphabetic strings (e.g., brevidad is decomposed into *breve-dad*), but also of non-alphabetic strings (i.e., *%%%%dad* is decomposed into *%%%%-dad*). Finally, Crepaldi et al. (2016) is the only study we know of that looked at suffix priming in English. The RTs to pairs in which pseudo-word primes shared the same suffix with derived targets (suffix condition; *sheeter-TEACHER*) were found to be faster than the RTs to pairs in which pseudo-word primes were affixed with a different suffix (suffix control condition; *sheetal-TEACHER*) and the RTs to pairs in which pseudo-word primes were affixed with a non-suffixal ending (control condition; *sheetub-TEACHER*). On the other hand, the RT to pairs in which pseudo-word primes shared the same non-suffixal ending with monomorphemic targets (word-ending condition; *poolel-BARREL*) were found not different from the RTs to pairs in which pseudo-word primes had a different suffix (word-ending control condition; *poolic-BARREL*) or from the RT to pairs in which pseudo-word primes has a different non-suffixal ending (control condition; *poolut-BARREL*). This suggests that morphological priming may be elicited only when a given letter cluster constitute a real morpheme.

1.3.3. Compound priming

A substantial body of studies suggests that decomposition occurs in compounds too. In the priming literature, both compound-initial and compound-final constituents were reported to be primed by their respective compounds. Fiorentino and Fund-Reznicek (2009) looked at priming effects to both head and non-head compound constituents in three conditions: the transparent, opaque, and orthographic conditions. In experiment 1, targets were non-head constituents of both transparent and opaque compounds (e.g., transparent condition: *flagpole-FLAG*; opaque condition: *hallmark-HALL*). In experiment 2, targets were head constituents of the both transparent and opaque compounds: e.g., transparent condition: *classroom-ROOM*; opaque condition: *honeymoon-MOON*. In both experiments, the orthographic condition consisted of a pseudo-compound prime and the target contained therein (experiment 1: *plankton-PLAN*; experiment 2: *battalion-LION*.) In both experiments, the transparent and opaque conditions elicited similar priming effects and the orthographic condition did not.

The apparent contradiction between the results reported in Rastle et al. (2004) and Morris et al. (2011; see sec.1.3.1) was first acknowledged in Fiorentino et al. (2015). Recall that in previous visual priming studies decomposition seems to occur in two apparently contradicting circumstances. It occurs in response to primes sharing a morphologically transparent (e.g., *alarming*) or opaque (e.g., *brother*) relationship with their targets Rastle et al. (2004, *ALARM, BROTH*). It also occurs in response to primes that contain the corresponding target (*FLEX*), irregardless of whether the primes are licit affixed word (*flexible*), illicit affixed pseudo-words (*flexity*), or pseudo-words ending with a non-suffixal letter cluster (*flexire:* Morris et al., 2011). Fiorentino et al. (2015) used two kinds of novel compounds to prime the corresponding compound-final constituent: compounds made of two extant English words (e.g, *drugrack-*

RACK) and compounds made of an English pseudo-word and an extant English word (e.g., *slegrack-RACK*). The results showed priming effects in both conditions, in line with Morris et al. (2011), and challenging Rastle et al. (2004)'s argument whereby decomposition occurs only for extant morphemes. The next chapter will address this contradiction and try to find a way of solving the contradiction while maintaining the hypothesis in (3).

1.4. From letters to words: Models of visual word recognition

While exploring the linguistic sophistication of decomposition, the dissertation indirectly engages with (i) orthographic processing, i.e. the process whereby orthographic stimuli (i.e., letters) are decoded by the parser; and (ii) lexical access, i.e., the process whereby the linguistic and lexical properties of words are retrieved. The section is divided in two subsections. In the first subsection, we describe the time-course of the orthographic word stimulus and briefly review the currently well-accepted model of orthographic processing; no other model will be taken into consideration, since the debate on orthographic processing is irrelevant to the purposes of this dissertation. In the second subsection, we take a deeper look at how decomposition lexical access is believed to occur in four different models; as being specifically relevant to our purposes, these models will be worked through in each of the following chapters.

1.4.1. Orthographic processing

When a visual word is presented, it is recognized (i.e., it is *read*) within a fraction of a second, regardless of the potential variations in position, size, case, and font (henceforth, *orthographic-feature invariance*). As being invented less than 6,000 years ago, reading (also known as visual word recognition) unlikely taps into a specific neuro-cognitive mechanism implemented in the brain. Rather, it is far more likely that reading recruits (or "recycles"; Dehaene and Cohen, 2007) general visual object identification neural mechanisms, which eventually trigger the development of letter-specific detectors due to the extensive training we are exposed to since we reach school age (Grainger, 2018). In particular, it has been shown that a localized region of the left occipito-temporal sulcus, close to the fusiform gyrus – commonly labeled as the Visual Word Form Area (VWFA) – is particularly sensitive to letter identification (Cohen et al., 2000).

In this dissertation, we assume the anatomically-implemented neurobiological framework proposed by (Dehaene et al., 2005), which takes advantage of the large amount of evidence about the organization of the ventral occipito-temporal route for visual object recognition in the human brain. While other models argue that letter feature-invariance is achieved in a single step, Dehaene et al. (2005)'s hierarchical model of orthographic recognition argues instead that visual receptive fields gradually widen their visual scope, aligning with the increase in the degree of complexity and abstractness of the properties neuronal populations that are gradually sensitive to. At lower levels, only basic features are detected (e.g., local contrasts). At the intermediate level, letter contours and (case-specific) shapes are used to detect a letter. At the higher level of analysis, the letter just detected is matched with the corresponding abstract (i.e., feature-invariant) letter representation. Once letters are abstractly identified, the system starts grouping letters together to form bigram sequences. Bigram sequences are argued to be the best trade-off letter sequence that allows the system to be flexible about letter location and order (both within the sequence and with respect to the flanking letters and sequences), as compared to smaller (monogram) or bigger (e.g., trigram) sequences. Finally, bigrams are analyzed together in combination, so that they may be identified as morpho-orthographic units (namely, orthographic units that corresponds to the abstract morpheme units).[3]

1.4.2. Models of lexical access and decomposition

While aiming to answer the question in (1), this dissertation indirectly provides insights on the process of *lexical access*, whereby linguistic and lexical properties are organized and retrieved during word recognition. As such, this dissertation engages with the current models of lexical access. As already discussed in sec. 1.1, the models of lexical access that have been proposed vary in the ontology of the mental lexicon. In some models, the mental lexicon is purely modular, in the sense that all linguistic properties (e.g., the morpho-orthographic form, alternations, morpho-syntactic properties, meaning) are stored in a single place (e.g., Hauk et al., 2006; Pulvermüller et al., 2001; Sereno et al., 1998). In other models, the mental lexicon is distributed throughout word processing, so that properties may gradually unfold over time, as word processing occurs (e.g., Dehaene et al., 2005; Vinckier et al., 2007). In this dissertation, we consider the following four models of lexical access: the bin model (Forster, 1999), the race model (Schreuder and Baayen, 1995), the full decomposition model (Taft, 1994, 2004), and the morpho-orthographic model (Crepaldi et al., 2010). On one hand, all four models seem to share a fairly distributed view of the mental lexicon. It is also important to reiterate that all of these models share the starting assumption that *segmentation* (i.e., the acquisition process of morpho-orthographic units) is distinct from *decomposition* (i.e., the on-line process of unit identification). On the other hand, they differ from one another in the mechanisms and, therefore, in the time-course of lexical retrieval. In the subsections below, we give a detailed description of the four models. The structure of each subsection is the same. In the first paragraph, we outline the mechanisms. In the second paragraph, we examine how lexical properties are distributed across stages. Finally, in the third paragraph, we exemplify how the model works through Rastle et al. (2004)'s critical results (see sec. 1.3.1). As dealing with *visual* word recognition, all of these models will be referred to when discussing the results of the visual experiments reported in the chapters 2-5 below. In doing this, we indirectly assume that different modality-specific strategies may be recruited in word processing (for a review, Grainger, 2018).[4]

---

[3] To ensure clarity, in this dissertation we will use square brackets (e.g., [...]) to refer to the phonological forms of morphemes, and dollar signs (e.g., $ . . . $) to refer to the orthographic forms of morphemes.

[4] The literature on modality-specific lexical access is very vast and touches on several related topics, among which are phonological processing (Vitevitch et al., 1999; Vitevitch and Luce, 1998, 1999), syntactic processing (Dikker et al., 2009), and connectionist modeling (e.g., interactive activation models: McClelland and Rumelhart, 1981; parallel distributed models: Seidenberg and McClelland, 1989; the cohort model: Marslen-Wilson and Welsh, 1978; Taft and Hambly, 1986; the TRACE model: McClelland and Elman, 1986).

A. BIN MODEL

Forster (1999) argues that orthographically similar words are grouped together in *bins*; each bin is identified with a specific hash code. When the stimulus is presented, the parser first converts the orthographic strings into a hash code indicating the bin where the correct entry is likely to be located. Recognition is then carried out through two stages. In the first stage – the entry-opening stage (originally proposed by Forster and Davis, 1984) – each entry within the same bin goes through a fast, frequency-ordered comparison with the stimulus and is assigned a goodness-of-fit score. This is a rather fast, low-level operation that sorts entries into subcategories depending on the degree of the matching comparison: perfect matches involve entries that are no different from the stimulus, close matches involve entries that are a bit different from the stimulus, no-matches involve entries that are completely different from the the stimulus. Entries that are either perfect or close matches with the stimulus are flagged as potential candidates. Once an entry has been flagged, it is "opened," in the sense that the information therein may be retrieved. The whole process operates in parallel for all entries and all identified bins. As soon as an entry is opened, the second stage – the verification stage – begins; at this stage, the comparison between the stimulus and the opened entries is much slower and detailed, and goes through each candidate in sequence. In this stage, if a candidate entry does not match with the stimulus, it is rejected, so that the verification of the next candidate begins. When a candidate entry is found to match the stimulus, the entry is selected as the correct one and all other entries are immediately closed even if they have not yet been evaluated. Priming is argued to arise when the entry for the target word has already been opened due to previous prime presentation. In this sense, priming is seen as *savings effects*. The prime masking procedure is assumed to block later stages of processing of the prime word. The model was proposed to explain the fact that the magnitude of identity priming is the same for high- and low-frequency words (e.g., *nation-NATION* and *plague-PLAGUE*, respectively; Forster and Davis, 1984; Forster et al., 2003). In the bin model, orthographically similar words are ordered in the same bin according to their frequency. Since priming is felt only after the lexical entry has been retrieved, the fact that a given word is low- or high-frequency does not interact with the priming magnitude (see Figure 1.3).

The bin model does not assume the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3), or any other specific constraint on decomposition. In this model, lexical properties are divvied up between the two stages. In the entry-opening stage, the orthographic form and frequency of entries is accessed. During this stage, the comparison mechanism evaluates the orthographic compatibility of frequency-ordered entries with the stimulus. Only entries that are close or perfect matches with the stimulus are opened and move on to the verification stage. In the verification stage, opened entries are further accessed, and more detailed information (thus, morphological alternations, syntactic category, meaning, ...) is therefore accessed.

Notice that the bin model does not include any specific mechanism of morphological decomposition, which is seen as resulting from the orthography-based entry-opening process described above. For this reason, the bin model does not seem able to fully account for Rastle et al. (2004)'s results. In particular, the absence of priming effects in the orthographic condition (*brothel-BROTH*) is rather problematic in the bin model. Since the comparison mechanism in the entry-opening stage is purely orthography-based, a prime word should always trigger opening of the entry of the related target, as long as they are orthographically related to one

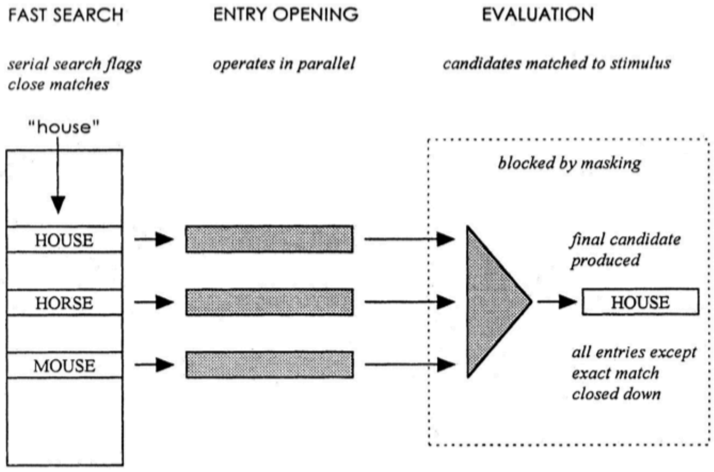

Figure 1.3.: The entry-opening model of lexical access (adapted from Forster (1998)).

another. Therefore, *brothel* is wrongly expected to trigger priming onto *BROTH* as much as *driver* and *brother* trigger priming onto *DRIVE* and *BROTH*, respectively.

B. RACE MODEL

The race model belongs to the type of models that are usually referred to as dual-route models. Dual-route models of lexical access argue that words are either recognized via decomposition (the parsing route) or as whole units (the storage route). In the original model proposed by Pinker (1991), words are recognized via either route depending on whether they were irregular (e.g., *fell*) or regular (e.g., *walked*), respectively. In Schreuder and Baayen (1995)'s race model, words go through both routes in parallel (Figure 1.4). In the storage route, the whole-form representation of words is linked to the lemma node, which activates the semantic and syntactic representation. In the parsing route, words are first decomposed into affixes and stems (segmentation stage). Then, the features of the activated constituents (e.g., it is checked that each affix is attached to the legitimate root category) are checked (licensing stage). Finally, the constituent morphemes are interpreted together (composition stage). Notice that the race model is somewhat unclear regarding the strategies the segmentation stage relies on when operating with words. The two routes race against one another, so that whichever route reaches the recognition point first "wins the race," in the sense that it is responsible for recognition. The immediate prediction in the race model dictates that words with high surface frequency are recognized via the storage route rather than the parsing route, since the time required for high-frequency words to be searched through in the storage route is less than the time required for the constituent morphemes of the same words to be decomposed and recognized in the parsing route. Conversely, words with low surface frequency and/or high stem frequency are recognized via the parsing route rather than the storage route, since the time required for the constituent morphemes of these words to be decomposed and recognized is less than the the time required for the whole-words to be searched through in the storage route. A number of behavioral studies have confirmed this prediction. For example, in traditional (e.g., non-priming) lexical decision tasks, recognition of plural forms such as *worlds* and *windows* has been reported to be depend on the surface frequencies of the singular and the plural forms of the same noun (*dominance*; among others, Baayen et al., 1997, 2007, 2002). Plural forms such as *windows*, which is more frequent than the singular form *window*, are visually recognized faster than plural forms such as *worlds*, which less frequent than the singular form *world.* The asymmetrical effect reported in these traditional lexical decision studies can be easily explained in the race model. When a singular-dominant plural form (*worlds*) is presented, the parsing route wins the race because the low-frequency plural form is decomposed via the parsing route faster than it is looked up via the storage route. When a plural-dominant plural form (*windows*) is presented, the storage route wins the race because the high-frequency plural form is looked up via the storage route faster than it is decomposed via the parsing route.

Albeit not explicitly, the race model assumes the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). In this model, lexical properties are distributed throughout the whole system. In the storage route, a pure, frequency-based lexical search of the full-form stimulus occurs. This means that, as soon as a full-form entry is looked up in the storage route, the full bundle of lexical and linguistic properties associated to it is accessed too. In the parsing route, the licensing stage accesses the syntactic information of the decomposed morphemes. No further information is provided regarding the properties accessed in other stages (decomposition stage, composition stage).

The race model is however unable to explain Rastle et al. (2004)'s results. This is because the model does not provide an explicit mechanism that could explain the differential effects between the transparent and opaque conditions, and the orthographic condition. The decomposition procedure of the parsing route is expected to apply blindly and send the segmented strings to the licensing stage. At this stage, any lexical entry matching with a segmented string should activate and therefore trigger priming onto the related target. This means that in the race model, root-containing monomorphemic words such as *brothel* should prime the respective target roots contained therein, and therefore decompose as much as morphologically transparent and opaque words do. After a transparent or opaque prime (*driver, brother*) is presented, it goes through both the storage and the parsing routes. In the storage route, the whole-form representation that matches the visual stimulus starts being searched through the lexicon. While the lexical search runs in the storage route, in the parsing route the visual stimulus is first segmented (segmentation stage; *driver*→driv-er, *brother*→broth-er). Segmentation also ensures that the segmented units match with the form of existing morphemes, which are therefore activated ({*drive*}, {*broth*}, {*er*}). By the time the related target is presented (*DRIVE, BROTH*), neither route has reached completion, but the lexical entry corresponding to the target stimulus is already activated; hence, priming arises. Similarly, after an orthographic prime (*brothel*) is presented, it goes through both the storage and the parsing routes. In the storage route, the whole-form representation that matches the visual stimulus starts being searched through the lexicon. While the lexical search runs in the storage route, in the parsing route the visual stimulus is first segmented as (*broth*→broth-el); one of the segments is then matched with the corresponding lexical entry in the licensing stage ({*broth*}). By the time the related target is presented (*BROTH*), both routes do not reach completion, but the root entry is activated and

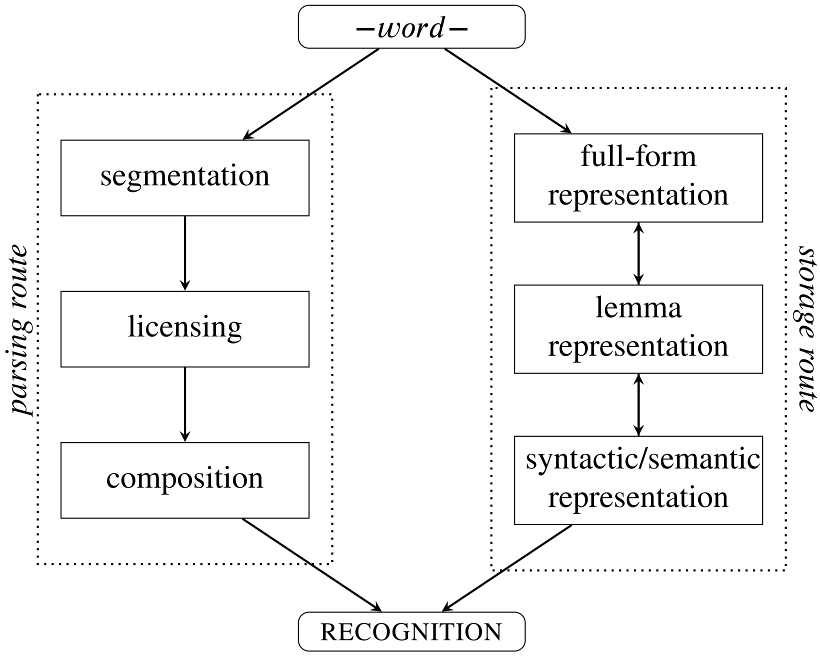

Figure 1.4.: Graphic representation of Schreuder and Baayen (1995)'s dual-route race model.

priming should arise.

C. FULL-DECOMPOSITION MODEL

In contrast with dual-route models in (B), single-route models argue that all words are recognized through a single path. In particular, Taft (1994, 2004)'s full-decomposition model argues that all words are decomposed before being recognized. The model involves a three-stage process. In the first stage (decomposition stage), words are decomposed into their putative constituent morphemes (also referred to as "form codes"). This procedure relies on orthographic regularities provided by the morpho-orthographic transitional probability of bi- and trigrams (Seidenberg, 1987). Bigram and trigram frequencies are high within stems and affixes, but low across morpheme boundaries. Formation of morpho-orthographic units – also called "islands of regularity (Rastle and Davis, 2008) – is triggered by the placement of a morpheme boundary, which occurs whenever a bigram has a lower transitional probability than the flanking bigrams (the "trough pattern"; see 4.2.5 for a detailed description of the algorithm). For example, a bigram cluster such as $ e r $ has a higher probability than the orthographically similar cluster $ e l $; therefore, the former string is more likely to identified as an island of regularity than the latter string. Although it has never actually been made explicit, low-probability strings are assumed to affect the whole segmentation pattern, including high-probability strings (Rastle et al., 2004): in a word like *brothel*, the low-probability string $ e l $ also affects the high-probability string $ b r o t h $. In the second stage (lookup or lemma stage), the lexical entries of each morpheme (stem and eventual affixes) are accessed. While the decomposition stage is assumed to be costless, the lookup stage is assumed to be affected by stem and/or affix frequency. In particular, the decomposition and lookup stages are claimed to be responsible for the base frequency effect, whereby words with high stem frequency are recognized faster than words with low stem frequency (Taft, 1979). This is also true for bound stems; for example, a word (e.g., *vent*) that can be the bound stem of a higher frequency word (e.g., *invent, prevent*) is recognized as a word slower than a frequency-matched word (e.g., *coin*) that cannot be a bound stem at all (Taft and Forster, 1975). Finally, in the third stage (recombination stage), the morphemes are then recombined so that the semantic interpretation of the whole form could be accessed. This stage is claimed to be responsible for the reverse base frequency effect, whereby low base-frequency words (*fangs*) were recognized faster than high base-frequency words (*moons*) when presented together with root-containing pseudo-words (*mirths*; Taft, 2004). This is because, when the constituent morphemes are put back together at the recombination stage, the advantage of activating a high-frequency stem (*moon*) is counterbalanced by the relative advantage of activating a low base-frequency word (*fangs*). Recent work has also suggested that any RT-taxing potential cost correlating with morphological complexity may be attributed to the recombination stage (as confirmed by Lehtonen et al., 2006, 2007; but see Zweig and Pylkkänen, 2008). The full decomposition model has also been corroborated by neurophysiological evidence (among others, Stockall and Marantz, 2006; Solomyak and Marantz, 2010; Lewis et al., 2011; Fruchter and Marantz, 2015).

The full-decomposition model assumes the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). Compared to the race model, the full-decomposition model provides a clearer by-stage distribution of the properties available to the parser. In the decomposition stage, only islands of regularity are accessed and, thus, no contact with the lexicon occurs. In the lookup stage, the putative morphemes are looked up in the lexicon; it is, however, unclear what this means in terms of lexical properties being accessed at this stage. Finally, in the recombination stage, when the morphemes are put back together, the full bundle of lexical information associated with the full form is expected to be retrieved.

In the full-decomposition model, priming is driven by activation of the lexical entry of the prime word in the lookup stage. Under this assumption, this model explains Rastle et al. (2004)'s results as follows (see Figure 1.5). After a transparent or opaque prime (*driver, brother*) is presented, it is first decomposed on the basis of morpho-orthographic regularities (decomposition stage; *driver*→driv-er, *brother*→broth-er). The units are then looked up and matched with the corresponding lexical representations, which are therefore activated (lookup stage; {*drive*}, {*broth*}, {*er*}). By the time the related target is presented (*DRIVE, BROTH*), the prime does not go through the recombination stage, but the lexical unit corresponding to the target stimulus is already activated; hence, priming arises. After an orthographic prime (*brothel*) is presented, $broth$ is identified as a legitimate morpho-orthographic unit, while el is not. The latter string prevents the segmented string $broth$ from being sent to the lookup stage. When the related target is presented (*BROTH*), target recognition does not benefit from the activated lexical entry, and priming does not arise. As it is, the full-decomposition model does not provide a mechanistically clear explanation for the segmented string $broth$ not being activated, thus leaving inhibition of $broth$ essentially unaccounted for. If orthographic processing occurs from left to right (e.g., Rumelhart and McClelland, 1982), it is hard to explain how, in a word like *brothel*, the string el is able to prevent the previous string $broth$ from activating. As we will see, this is a problem that is also common to the morpho-orthographic model (D).

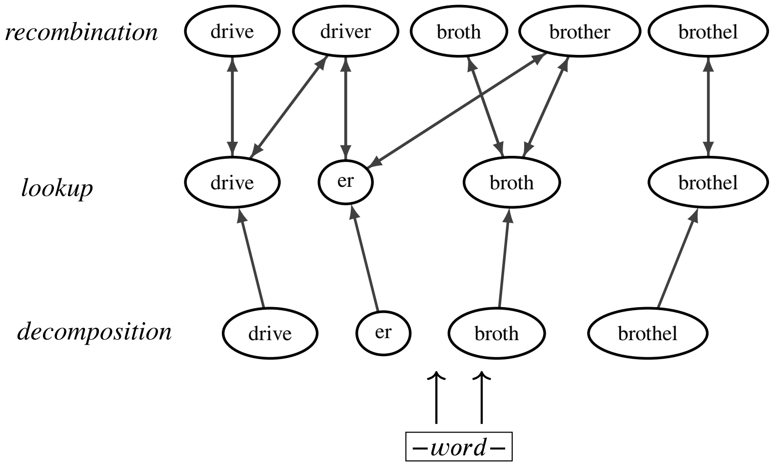

Figure 1.5.: Graphic representation of the full-decomposition model to explain the results reported in Rastle et al. (2004). Cf. Figure 1.6.

D. MORPHO-ORTHOGRAPHIC MODEL

Crepaldi et al. (2010)'s morpho-orthographic model is a connectionist-oriented development of Taft's. This model involves three levels (or stages) of computation: morpho-orthographic segmentation, orthographic lexicon, and lemma level. The *nodes* in a given level are connected to the nodes of the next level (see Figure 1.6). When a word is presented, it is first decomposed into constituent units (i.e., morphemes) in the morpho-segmentation level. The morpho-orthographic segmentation level in the morpho-orthographic model is similar to the decomposition stage in the full-decomposition model in that they both rely on Rastle and Davis (2008)'s morpho-orthographic islands of regularity (see above for further details). Additionally, the morpho-orthographic model assumes a "decomposability constraint" whereby only words that are entirely decomposable into morphemes are decomposed. The decomposability constraint is nothing but another way to implement the effect that low-probability strings have onto decomposition patterns (see paragraph above). As Crepaldi et al. (2010) also acknowledge, such a constraint is hard to implement computationally, especially under the common assumption that orthographic decoding occurs in a serial, left-to-right fashion. The activated morpho-orthographic nodes then activate the corresponding nodes in the a orthographic lexicon. At this level, the orthographic forms of bare roots, and inflected and derived forms are stored (along the lines of Coltheart et al., 2001). The activated nodes in both the morpho-orthographic segmentation level and the orthographic lexicon are assumed to trigger priming effects. The activated nodes in the orthographic lexicon finally feeds into the lemma level, which contains the actual lexical representations. As compared to the full-decomposition model, the morpho-orthographic model does not include a recombination level, and therefore assumes word recognition to be solely contingent on activation of the corresponding full-form nodes at the lemma level. Both the orthographic lexicon and the lemma level contain (a) bare root forms (e.g., *cat*) and (b) inflected and derived forms (e.g., *fell, falls, weakness*). However, neither level contains inflectional/derivational affixes (e.g., *-s, -ness*) as independent nodes. First, inflectional affixes are not contained in the orthographic lexicon and in the lemma level in order to ensure that regular, but ungrammatical forms such as **falled* may be distinguished from irregular, but grammatical forms such as *fell*. Second, derivational affixes are not contained in the orthographic lexicon or in the lemma level in order to explain the similar priming effects reported for morphologically transparent pairs (e.g., *darkness-DARK*) and morphologically opaque pairs (e.g., *brother-BROTH*). If derivational affixes presented as independent nodes in both levels, *darkness* should activate the nodes {*dark*} and {*ness*} at both levels of computation; whereas, *corner* should activate {*corn*}, {*er*}, and {*corner*} in the orthographic lexicon, and {*corner*} in the lemma level. Therefore, the differential activation of nodes across both levels between the two conditions should result in an greater priming magnitude for *darkness-DARK* than for *corner-CORN*. The presence of fullly derived forms in both levels ensures the same amount of activation energy for both morphologically transparent and opaque prime words.[5]

The morpho-orthographic model assumes the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). In this model, the distribution of lexical and linguistic properties is similar to the full-decomposition model, with the crucial difference being in the addition of the orthographic lexicon feeding the lemma level. At the morpho-orthographic segmentation level, decomposition relies on islands of regularity and within the decomposability constraint. In the orthographic lexicon, the *actual* orthographic forms of lexical entries are stored. Here, bare roots, and fully inflected and derived forms are included; this ensures that ungrammatical regular and irregular forms (e.g., **falled*) not be activated at this level. At the lemma level, activation of the lexical entry triggers full access to all of the properties associated to it.

The morpho-orthographic model explains Rastle et al. (2004)'s results as follows. When a morphologically transparent or opaque prime word is presented (*driver, brother*), the morpho-orthographic segmentation activates the nodes {*drive*}, {*broth*}, and {*er*}, which then feed into the compatible orthographic nodes in the orthographic lexicon. When the related target is presented (*DRIVE, BROTH*), the matched orthographic nodes are already activated, thus triggering priming. When an orthographic prime word is presented (*brothel*), the morpho-orthographic segmentation cannot segment it as $ b r o t h - e l $, which violates the decomposability constraint because the segment $ e l $ does not match with an existing morpheme. Therefore the segment $ b r o t h $ is not sent to the orthographic lexicon to be activated (whereas the non-decomposed segment $ b r o t h e l $ may be activated instead). Thus, when the related target is presented (*BROTH*), no priming arises.

Taking stock

It is important to underline that each of the four models above fails, in one way or another, to provide a explicit mechanism to explain Rastle et al. (2004)'s results. In the bin model (A), the comparison mechanism behind the entry opening process is solely based on orthographic relatedness, and therefore is unable to explain the differential effects reported between morphological (transparent and opaque) conditions and purely orthographic conditions. Similarly, in the race model (B), the mechanism within the segmentation stage has never been fully fleshed

---

[5] Some theories of morphology also highlight the difference between derivational and inflectional morphology, in that only the former lead to formation of independent lexical entries (among others, Aronoff, 1976; Stump, 2001). However, there is little-to-no consensus among theoretical morphologists on this topic; see, for example, Embick (2015) for a different perspective.

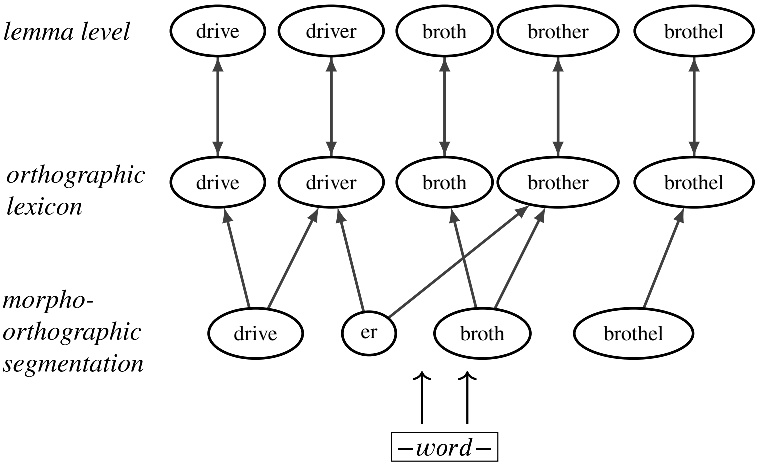

Figure 1.6.: Graphic representation of the morpho-orthographic model (adapted from Crepaldi et al., 2010 to explain the results reported in Rastle et al. (2004).

out, therefore it is not clear how the model can explain the differential effects. Finally, the full-decomposition model (C) and the morpho-orthographic model (D) attempt to explain the effects by arguing that segmented strings with an low orthographic probability (e.g., \$ e l \$) affect the other segmented strings with a higher orthographic probability (e.g., \$ b r o t h \$); however, they do not clearly specify the underlying mechanism. As it turns out, explaining these results is not easy to implement mechanistically especially because visual letter analysis is generally assumed to feed into the word recognition system in a serial fashion (e.g., Rumelhart and McClelland, 1982). In these terms, when *brothel* is presented, its constituent letter would be orthographically decoded serially, from left to right, which would, in turn, lead to recognition of the string \$ b r o t h \$ and therefore elicit priming onto the target *BROTH*. In the next chapters, we will interpret our results within the four models above, with the ultimate goal of overcoming the implementational complications thereby and therefore filling the gaps with missing explicit mechanisms. As we will see, the competion-based model of decomposition proposed in Chapter 6 seem to overcome these implementational issues by allowing the decomposition procedure to generate and evaluate multiple potential candidates in parallel.

1.5. Experimental methodology

As already said above, in this dissertation we primarily make use of the visual masked priming design. We mainly arrange conditions in self-contained subdesigns, in which half of the targets are paired up with related primes and the other are paired up with unrelated primes; we take advantage of the Latin-Square arrangement whenever possible. For the sake of clarity, the different designs used (i.e., the rating design) will be described in the relevant sections.

Analysis of the data will be thoroughly described in the dedicated section of each experiment in each chapter. Statistical analyses will involve calculation of the best-fit model by using both linear mixed-effect regression (henceforth, LMER) and the Bayes Factor (henceforth, BF) statistical algorithms for all conditions of each experiment. The LMER algorithm is a statistical linear model that tries to explain the distribution of a given dependent variable (in our case, RTs) as an effect of a set of predictors (or independent variables; in our case, relatedness of the prime word to the target word). It is called *mixed* because it allows us to take into account the distribution of our data as explained by both fixed effect factors (namely, our independent variables) and random effect factors (namely, factors whose levels may not be replicated exactly and do not exhaust the full range of possible level values). We calculate the LMER estimates using the restricted maximum likelihood (REML) criterion implemented in the `lmerTest` R-package (Kuznetsova et al., 2017). P-values were estimated using the Satterthwaite approximation of degrees of freedom. We adopt the usual α-criterion (α=.05) and interpret any resulting p-value$<\ \alpha$ as a case in which we can conclude either (i) the null hypothesis is incorrect or (ii) a rare event occurred (*Fisher's disjunction*). We also want to keep in mind that all null-hypothesis approaches only provide $p(data|H_0)$ (where the pipe symbol, |, is to be read as "given that") in an explicit way and make inferences from that about which competing hypothesis to choose.

Bayesian approaches use Bayes Theorem (Bayes, 1763) to provide $p(H|data)$, where H can be any of the hypotheses considered (see also Kruschke, 2011, for an accessible introductory textbook). In particular, Bayes Factors (BFs) are the ratio of the likelihood of the data under one hypothesis (i.e., the null hypothesis, $p(data|H_0)$) to the likelihood of the data under the other hypothesis (the alternative hypotheses, $p(data|H_1)$; Wagenmakers, 2007). In other words, they calculate the strength of the evidence for one hypothesis over another. In our analyses, we will calculate BFs by using the R-package `BayesFactor` (Morey, 2018).

(6) a. $BF_{1,0} = \frac{p(data|H_1)}{p(data|H_0)}$ b. $BF_{0,1} = \frac{p(data|H_0)}{p(data|H_1)}$

Notice that the formulae above vary in which hypothesis is in the numerator and which hypothesis is in the denominator. Mathematically, (6b) is the inverse of (6a). BFs calculated as in (6a) express how much more likely the data is under the alternative hypothesis H_1 compared to the null hypothesis H_0; conversely, BF-values calculated as in (6b) express how much more likely the data is under the null hypothesis H_0 compared to the alternative hypothesis H_1. We are going to use both BF-values in our statistical analyses, depending on the purpose. We will typically calculate $BF_{1,0}$'s when we want to know the likelihood of the data under the alternative hypothesis H_1; conversely, we will typically calculate $BF_{0,1}$'s when we want to know the likelihood of the data under the null hypothesis H_0. As odds ratios, the BF-values are always between 0 and ∞. A BF of 1 indicates that the data is equally likely under either hypothesis and therefore suggests that the data collected is not able to provide enough evidence to prefer one hypothesis over the other. A BF greater than 1 indicates that the data is more likely under the hypothesis in the numerator than under the hypothesis in the denominator. A BF smaller than 1 indicates that the data is more likely under the hypothesis in the denominator than under the hypothesis in the numerator. Since there are no explicit rules for interpreting BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in Table 1.2 below. In compliance with these guidelines, we adopt the conventional threshold of 3 in the BF analyses to decide on competing hypotheses.

$BF_{1,0}$		$BF_{0,1}$	
BF-VALUE	INTERPRETATION	BF-VALUE	INTERPRETATION
0 to 0.01	extreme for H_0	0 to 0.01	extreme for H_1
0.01 to 0.1	strong for H_0	0.01 to 0.1	strong for H_1
0.1 to 0.33	substantial for H_0	0.1 to 0.33	substantial for H_1
0.33 to 1	anectodal for H_0	0.33 to 1	anectodal for H_1
1 to 3	anectodal for H_1	1 to 3	anectodal for H_0
3 to 10	substantial for H_1	3 to 10	substantial for H_0
10 to 100	strong for H_1	3 to 10	strong for H_0
100 to ∞	extreme for H_1	100 to ∞	extreme for H_0

Table 1.2.: Jeffreys (1961)'s guidelines for interpreting Bayes Factors.

Similarly, we also calculate pairwise comparisons between the effects in each condition of a given experiment. For each comparison, we only report the Dunn-corrected p-value, as well as the corresponding uncorrected BF-value. The general purpose of reporting both LMER and BF analyses is to guarantee *continuity* with the previous statistical methodologies and *reliability* in th interpretation of the results obtained. On one hand, LMER analyses ensure that our results can be compared to the results reported in the literature. On the other hand, BF analyses provide a clear measure of the evidence for one hypothesis over the other. For these reasons, in the event that p-values calculated in the LMER analysis look in conflict with BFs calculated in the BF analysis, our interpretation of results will mainly refer to latter.

1.6. Outline of the dissertation

CHAPTER 2 asks whether syllabification drives decomposition at early stages of processing. We examine this hypothesis because it may explain why words that contain a root but not a second morpheme do not decompose (*brothel*↛broth-el; "*brothel non-effect*"; Rastle et al., 2004), but non-words that contain a root but not a second morpheme do (*slegrack*→sleg-rack; *flexire*→flex-ire; "*slegrack/flexire effect*"; Fiorentino et al., 2015; Morris et al., 2011). One possibility is that the slegrack/flexire effect results from the syllabic chunking operation (*flexire* → FLEX.IRE → flex-ire; where the dot '.' signals a loose syllabic boundary). The brothel non-effect would then be a consequence of averaging the priming response to words with a syllabic, non-affixal ending (*brothel* → BROTH.EL → broth-el) together with the non-priming response to words with a non-syllabic, non-affixal ending (i.e., *against* → AGAINST ↛ again-st). In experiment 1, masked priming was elicited in two different orthographically related, but morphologically unrelated conditions: (i) a syllabic condition, in which the prime word contained the corresponding target word, plus a syllabic unit (e.g., *ban.jo-BAN*), and (ii) a non-syllabic condition, in which the prime word contained the corresponding target word, plus a consonantal unit (e.g., *starch-STAR*, where the underlined string signals the extra consonantal unit). Results showed that both conditions did not elicit priming, thus suggesting that syllabification does not drive morphological decomposition (at least in the visual modality). These results potentially imply that decomposition is affected by the lexicality of the stimulus. In words, decomposition occurs only if they can be transparently broken down into morphemes; in non-words, decomposition occurs as long as they contain at least one morpheme.

In the same chapter, we also include results for the same experiment calculated on data collected on *PsychoJS*, a newly introduced javascript-based program that allows for data collection online through subjects' browsers (Peirce et al., 2019). We compare the results obtained online with the results obtained in the more traditional environment (i.e., in a quiet lab room) as a way to evaluate the use of the online environment to conduct the priming experiments presented in the following chapters.

CHAPTER 3 asks whether decomposition is sensitive to *dominance*. This term refers to the asymmetrical ratio between frequencies of the forms in the same nominal paradigm, For example, a word such as *windows* is plural-dominant because the surface frequency of the plural form (*windows*) is higher than the surface frequency of the singular form (*window*); a word such as *worlds* is instead singular-dominant, because the surface frequency of the plural form (*worlds*) is lower than the surface frequency of the singular form (*world*). Lexical decision studies have reported that plural-dominant plural forms are recognized faster than singular-dominant plural forms. The hypothesis that dominance potentially impinges on decomposition was tested in experiment 2, as a way to understand whether word frequency is accessed during decomposition or, rather, at later stages of processing. The masked priming response was tested in two different morphological conditions: (i) the sgdom-sgdom condition, where the pairs contained singular-dominant plural forms as both primes and targets (e.g., *worlds-HEAVENS*), and (ii) the pldom-sgdom condition, where the pairs contained a plural-dominant plural prime and a singular-dominant plural target (e.g., *windows-GODS*). Results showed that priming occurred in the sgdom-sgdom condition, but not in the pldom-sgdom condition, thus suggesting that dominance asymmetries (and therefore word frequency) impinge on decomposition.

CHAPTER 4 asks whether decomposition is sensitive to *phonologically-conditioned morpho-phonological alternations*. These alternations involve different exponents that realize the *same* morpho-syntactic feature; for example, the words *impossible* and *intolerant* present the same prefix *in-*, which is realized as *in* or *im*, depending on the place of articulation of the following root-initial segment. Morpho-phonological alternations are fully predictable via language-general phonological operations (in our example, an operation of place assimilation). Experiment 3 elicited priming of each alternant in all *prime-TARGET* combinations. Unfortunately, this experiment encountered a series of technical problems and we were unable to interpret the results obtained in terms of the two competing hypotheses at issue. Rather, our results shed light on a more fundamental issue in decomposition: namely, the factual inability of the current orthography-based algorithm to account for morpheme segmentation.

CHAPTER 5 asks whether decomposition is sensitive to syntactic violations to affixation of bimorphemic words. Across languages, affixes are endowed with specific syntactic restrictions concerning which syntactic categories they may attach to. For example, the suffix *-able* attaches to verbs and yields adjectives (e.g., *detectable*, but **blissable*), whereas the suffix *-ness* attaches to adjectives and yields nouns (e.g., *weakness*, but not **blissness*). Before addressing this question, we identified two possibly interacting factors. First, we explored the extent to which the zero-morpheme may license syntactically illicit forms and make them grammatical. Experiment 4 revealed that subjects judge *able-, ity-, ment-,* and *ness*-suffixed illicit forms as grammatical. This suggests that the questions being asked in the chapter may not be explored without ruling in the possibility that zero-derivation may be at play in decomposition procedures of syntactically illicit forms. Second, we explored the extent to which decomposition may selectively occur for some morphemes, but not others, in line with what the experiments in previous chapters seem to suggest. Experiment 5 elicited the masked priming response to syntactically licit words being presented as primes of their corresponding target roots (e.g. *detectable-DETECT*); the results seem to ensure that all four affixes decompose. Finally, experiment 6 elicited the masked priming response to both syntactically licit and illicit prime stimuli (e.g., *blissity-BLISS* vs. *blissful-BLISS*). In the light of the two previous experiments, our results cannot be used to make a conclusive argument regarding the question being asked in the chapter. On the other hand, the results seem to suggest that syntactically illicit forms still decompose regardless of violations to the syntactic restrictions of the corresponding affix.

CHAPTER 6 takes stock of all the results reported here, tries to (at least partially) answer the question in (1), and addresses the issue about the unfolding of properties during early stages of visual word processing. In particular, we developed a model in which decomposition generates multiple decomposition pattern candidates (including the whole-form string) that compete with one another on-line for activation. This allows for whole-word properties such as frequency and (possibly) lexicality to impact decomposition, while leaving other more abstract properties (e.g., syntactic affixal restrictions and meaning) for later stages.

We finally lay out a broader research program designed to further explore the theoretical and methodological issues raised in this dissertation.

Chapter 2.

Morphological decomposition and *syllabification*

2.1. Introduction

In Chapter 1, we saw that morphological decomposition is commonly argued to occur before contact with the lexicon, as stated by the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). Though substantiated by a large body of evidence in the visual masked priming literature (see sec. 1.3 for a review), the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3) hinges on only two crucial results. First, priming arises in both morphologically transparent (*alarming-ALARM*) and opaque (*brother-BROTH*) pairs, regardless of their semantic relatedness (*alarm* and *alarming* are indeed morphologically and semantically related to one another, but *broth* and *brother* are not; Rastle et al., 2004). Second, priming does not arise in semantically related word pairs (*cello-VIOLIN*; Rastle et al., 2000). Since it occurs before contact with the lexicon, decomposition is therefore believed to rely on the orthographic transitional probability of letter clusters (namely, the probability that two, three, or more letters are found together; see sec. 1.4.2). As such, it hinges on statistical regularities, while implying no contact with the lexicon whatsoever. Letter clusters with a low orthographic transitional probability are flagged as potential morpheme boundaries, whereas letter clusters with a high orthographic transitional probability are flagged as morpho-orthographic units ("islands of regularity"; Rastle and Davis, 2008). As mentioned in sec. 1.4.2, the islands of regularity approach wrongly predicts that monomorphemic words (e.g., *brothel*) prime the semantically and morphologically unrelated word embedded in it (*BROTH*). If decomposition exclusively relied on islands of regularity, it would identify the string $\$ b r o t h \$$ as a morpho-orthographic unit and therefore trigger priming onto the target *BROTH*. The absence of priming in pairs like *brothel-BROTH* therefore suggests that morpho-orthographic decomposition occurs only when the whole visual string may be fully decomposed. Thus, morphologically transparent and opaque words such as *alarming, brother* decompose because they are fully decomposable: *alarm, broth, -ing,* and *-er* are all morphemes of English; monomorphemic words such as *brothel* do not decompose because they are not fully decomposable: *broth* is a morpheme, but *-el* is not. This pattern of effects, which we refer to as the '*brothel non-effect*', has turned out particularly hard to implement within models of decomposition (see sec. 1.4.2). The bin and the race models provide orthographic-based mechanisms that are therefore unable to explain the brothel non-effect described above. The full-decomposition and the morpho-orthographic models only provide a descriptive explanation of the brothel non-effect: i.e., *brothel* does not decompose because *-el* is not an English morpheme and inhibits decomposition ("decomposability constraint"). However, the lack of mechanistic explanation for such a constraint is mainly due to the fact that connectionist models make use of the activation mechanism (see sec. 1.2.1), which, however, appears to be insufficient to explain the brothel non-effect. It is indeed unclear how orthographic transitional probabilities may allow *brother* to decompose and *brothel* to not. This is particularly true in models of orthographic processing in which letter identification and integration occurs serially, from left to right: how could the non-morphemic string $ e l $ inhibit activation of the preceding morphemic string $ b r o t h $, so that the corresponding lexical entry does not contribute to priming onto the target word *BROTH*? Failure in explaining the brothel non-effect seems inevitable, and suggests that decomposition does make contact with the lexicon at least to some extent, in contrast with the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3).

The hypothesis that early decomposition may make at least partial contact with the lexicon seems to be further confirmed by more recent masked priming studies looking at both affixal and compound morphology. As we have seen in sec. 1.3, Fiorentino et al. (2015) reported that pseudo-compounds consisting of a pseudo-word and a real word (*slegrack*) prime their target constituent (*RACK*) as much as lexicalized compounds (e.g., *flagpole*) do (*POLE*). Similarly, Morris et al. (2011) reported that pseudo-suffixed non-words (*flexire*) prime the correspondent target root (*FLEX*), similarly to suffixed words (*flexible*). Given the wealth of evidence showing that compound decomposition is akin to affixal decomposition (among others, Fiorentino and Poeppel, 2007; Fiorentino and Fund-Reznicek, 2009), we conflate these two sets of results and refer to them as the '*slegrack/flexire effect*.' These results seem to contradict the brothel non-effect, and therefore the "decomposability constraint" assumed in the full-decomposition and morpho-orthographic models of decomposition. The *brothel non-effect* suggests that decomposition does not occur in presence of a pseudo-morphemic string (e.g., *-el* as in *brothel*), in compliance with the "decomposability constraint". The *slegrack/flexire effect* suggests instead that decomposition occurs for non-words like *slegrack, flexire* regardless of the pseudo-morphemic pieces *sleg-, -ire*, thus circumventing the "decomposability constraint." The most obvious way to explain the two effects would be to make reference to the different lexicality status of the prime stimuli tested: *brothel* is a word and does not prime *BROTH*, since *-el* is not a morpheme; *slegrack* and *flexire* are non-words, and they prime *RACK* and *FLEX* respectively, regardless of the fact that *sleg-, -ire* are not morphemes. However, accessing the lexicality status of a given visual stimulus essentially entails contacting the lexicon, which would lead us to challenge the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3) again.

In contradicting one another, the brothel non-effect and slegrack/flexire effect seem to argue for a decomposition procedure that has (at least partial) access to the lexicon – thus, against the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3), which a large body of behavioral and neurophysiological results seem instead to support (among others, Stockall and Marantz, 2006; Solomyak and Marantz, 2009; Lewis et al., 2011; Fruchter and Marantz, 2015). In this chapter, we therefore try to traverse through the impasse by entertaining a hypothesis that, if true, could explain the two effects, while still maintaining the MORPHO-ORTHOGRAPHIC DE-

COMPOSITION HYPOTHESIS (3). Our testing hypothesis is that visual decomposition could rely on a phono-orthographic procedure of syllabification, which breaks down the visual stimulus into syllable-like chunks (similar to Taft, 2003's BOSS units). Occurring at early stages, this procedure is rather crude and clusters together letters from left to right. This procedure first identifies an orthographic nucleus (that is, a vocoid string) and groups with it an orthographic coda (that is, a consonantal cluster string) according to the phono-orthographic syllabic constraints of the language. Under this hypothesis, the brothel and the slegrack/flexire effects are explained as follows. On one hand, the slegrack/flexire effect is a direct result of the syllabification procedure just described. Stimuli such as *slegrack, flexire* are loosely syllabified as $sleg.rack$, $flex.ire$ (where the dot "." signals a loose syllabic boundary). These syllabic chunks then activate the lexical entries {*rack*} and {*flex*}, which in turn trigger priming onto the target words *RACK, FLEX*. On the other hand, the brothel non-effect is a confound resulting from the diverse syllabic nature of the word pairs tested in the same condition. In the literature on morphological priming, the brothel non-effect has been reported in orthographic conditions in which the primes consisted of the primed root followed by a syllabic (i.e., *brothel-BROTH*) or by a consonantal (*against-AGAIN*) non-suffixal ending (as also suggested by Taft and Nguyen-Hoan, 2010). Thus, under the syllabification hypothesis we are entertaining here, a word with a syllabic, non-suffixal ending like *brothel* is loosely syllabified as $broth.el$; the syllabic chunk $broth$ activates the lexical entry {*broth*}, which in turn triggers priming onto the target word (*BROTH*). Conversely, a word with a non-syllabic, non-suffixal ending like *against* is loosely syllabified as $against$, which activates the lexical entry {*against*}, thus inhibiting/suppressing priming onto the target word (*AGAIN*). When averaged together as part of the same condition, the opposite priming magnitudes cancel each other out, thus leading to the brothel non-effect. The syllabification-driven decomposition procedure seems like a promising solution for the impasse we are in, as it would fill three needs with one deed: (i) it explains the difference betwe the brothel and slegrack/flexire effects by providing a reasonable mechanistic explanation that (ii) is relatively easy to implement in models of lexical access, (iii) without running afoul of the commonly-accepted MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3).

Experiment 1 below was designed to test this hypothesis. In this experiment, we elicited priming effects to the transparent and opaque conditions, similarly to Rastle et al. (2004); additionally, we tested priming effects to two different orthographic conditions: a *syllabic condition*, in which primes were words made of a monosyllabic real word and a syllabic non-suffixal ending (e.g., *can.vas-CAN*), and a *non-syllabic condition*, in which the primes were words made of a real word and a consonantal non-suffixal ending (e.g., *starch-STAR*). Sec. 2.2 reports the version of experiment 1 that was collected in a quiet, dark room in our lab. Our results disconfirm the hypothesis that decomposition relies on syllabification, and further suggest the need for a more abstract model of decomposition. Sec. 2.3 reports the version of experiment 1 that was instead collected on-line through the newly-developed javascript-based program *PsychoJS*. This was done with two goals in mind. First, it allowed us to make a direct comparison between the online and the in-lab environments. The section also reports a thorough comparison of the data collected in the two different environments (in-lab vs. online), so as to validate the online method against the in-lab standards. Second, it provided a way to confirm the results obtained in the in-lab environment, while benefiting from a larger sample and therefore a higher statistical power. Regardless of their distributional differences, the two versions of experiment 1 show a similar pattern of results. Sec. 2.4 takes stock of the results reported in previous sections and offers further remarks on a mechanism of decomposition that is able to account for the brothel and the slegrack/flexire effects.

2.2. Experiment 1 – syllabicity/root priming (in-lab)

2.2.1. Materials

One-hundred and sixty pairs were selected from the Worldlex corpus (WL: Gimenes and New, 2016), 32 in each of the following five conditions.

1. Pairs in the identity condition included monomorphemic words that were presented both as prime and target (e.g., *fuss-FUSS*).
2. Pairs in the *transparent condition* carried a semantically and morphologically transparent relationship (e.g., *boneless-BONE*).
3. Pairs in the *opaque condition* carried an apparent morphological relationship, but not a semantic one (e.g., *belly-BELL*).

In lieu of Rastle et al. (2004)'s orthographic condition (*brothel-BROTH*), we constructed two different orthographic conditions.

4. Pairs in the *syllabic condition* consisted of prime words (e.g., *banjo*) that were made of the corresponding target word (*BAN*) plus an additional syllabic pseudo-suffix (*-jo*).
5. Pairs in the *non-syllabic condition* consisted of prime words (*starch*) that were made of the corresponding target word (*STAR*) and an additional consonantal, non-syllabic pseudo-suffix (*-ch*).

In compiling the lists of these two conditions, we were particularly careful so that both of them adhered to the constraints listed below. Since syllabification (namely, the process whereby words are chunked into syllabic units) is part of phonology, these constraints guaranteed that prime and target words consistently carried the same orthographic and phonological information, without mismatches (e.g., different spellings in homophonic pairs) or conflicts (e.g., different pronunciations in homographic pairs).

i. the TARGET was phonologically and orthographically identical to the left most portion of the prime (e.g., *heaven-HEAVE* was excluded because of the root vowel change)

ii. the TARGET did not contain extra letters with respect to the corresponding prime word (e.g., *sort-SORE* was excluded because of the silent e present in the target but absent in the prime);

iii. the prime word could not possibly be segmented differently from the way the TARGET is extracted (e.g., *restore-REST* was excluded since *restore* may be possibly paired up with *REST*, *STORE*, or *TORE* as targets);

iv. the TARGET was not an abbreviation or a proper name (e.g., *caroline-CAROL* was excluded because both prime and target words can be proper names);

v. the pseudo-suffix in the prime was not a real suffix from either the orthographic and phonological point of view (e.g., *covenant-COVEN* was excluded since *covenant* may possibly segmented as affixed with the *-ant* suffix; similarly, *brothel-BROTH* was excluded since the pseudo-suffix *-el* in brothel is phonologically identical to the suffix *-al*);

vi. prime and TARGET pairs were not morphologically, semantically and/or (pseudo-)etymologically related (e.g., *arcade-ARC* was excluded);

vii. the pseudo-suffix did not consist of more than one syllable (e.g., *armada-ARM* was excluded since *armada* is trisyllabic);

viii. the TARGET was not an inflected word form (e.g., *rancho-RAN* was excluded since *ran* is the past tense form of the verb *run*);

ix. in the syllabic condition, the pseudo-suffix was a legitimate syllable (e.g, *galaxy-GALA* was excluded since [ksi] is not a legitimate syllabic unit in English).

To ensure best comparability across conditions, none of the primes in the transparent, opaque, or syllabic conditions were monosyllabic bimorphemic words (e.g., a word like *dipped* was excluded); each of these conditions had only bimorphemic prime word with more than two syllables.

Primes and targets across the five conditions were matched as closely as possible on frequency and length. In addition to the frequency values reported in the WL corpus, we ran the same calculations against the HAL frequency values (Hyperspace Analogue to Language; Lund and Burgess, 1996) reported in the English Lexicon Project corpus (ELP: Balota et al., 2007).[6] The mean values across the three conditions and the ANOVA results are shown in Table 2.1 below. Because of the substantial differences between the non-syllabic condition and the other conditions, we were unable to control for orthographic length. In line with previous studies (Rastle et al., 2000, 2004), we therefore added orthographic length as a covariate in the data analysis. Items of the identity condition were matched in length and frequency with the targets of the other five conditions.

Thirty-two unrelated prime words were selected for all target words; these words were orthographically, morphologically and semantically unrelated to the corresponding targets and were matched as closely as possible on frequency and length (WL: $t(159)$=-0.45, p=.65, $BF_{0,1}$= 10.26; HAL: $t(159)$=-0.16, p=.87, $BF_{0,1}$=11.10; length: $t(159)$=0.03, p=.97, $BF_{0,1}$=11.34.

---

[6] Notice that BF-values for the HAL frequency were smaller than the BF-values for the WL frequency. This was due to the fact that a few words of our lists were not present in the HAL corpus and were therefore removed from the BF analysis.

Property		Conditions					STATISTICS
		identity	*transparent*	*opaque*	*syllabic*	*non-syllabic*	
PRIMES	WL frequency (log10)	–	2.40	2.85	2.08	2.18	$F(3,124)=1.74$, $p=.16$ $BF_{0,1} = 3.28$
	HAL frequency (log10)	–	7.25	7.83	7.19	7.23	$F(3,121)=1.29$, $p=.27$ $BF_{0,1} = 0.87$
	length	–	6.69	6.19	6.75	(4.78)	$F(3,90)=2.2$, $p=.10$ $BF_{0,1} = 2.25$
TARGETS	WL frequency (log10)	9.22	9.39	8.79	9.75	9.22	$F(4,155)=1.53$, $p=.2$ $BF_{0,1} = 5.23$
	HAL frequency (log10)	4.15	4.27	3.69	4.56	3.93	$F(4,155)=1.52$, $p=.19$ $BF_{1,0} = 2.61$
	length	3.5	3.56	3.78	3.47	3.41	$F(4,155)=1.24$, $p=.29$ $BF_{0,1} = 7.96$

Table 2.1.: Experiment 1 – syllabicity/root priming. Summary of the lexical properties of the stimuli used. The mean length of the non-syllabic condition is in parentheses because it was not included in the statistical calculations (see sec.2.2.1).

Property	Conditions				
	identity	*transparent*	*opaque*	*syllabic*	*non-syllabic*
log10 WL	$t(31)=0.19$, $p=.84$ $BF_{0,1}=5.2$	$t(31)=-0.7$, $p=.48$ $BF_{0,1}=4.23$	$t(31)=0.11$, $p=.9$ $BF_{0,1}=5.26$	$t(31)=-0.21$, $p=.83$ $BF_{0,1}=5.18$	$t(31)=0.06$, $p=.5$ $BF_{0,1}=4.29$
log10 HAL	$t(31)=0.04$, $p=.96$ $BF_{0,1}=5.29$	$t(31)=-0.46$, $p=.64$ $BF_{0,1}=4.8$	$t(31)=-0.003$, $p=.99$ $BF_{0,1}=5.29$	$t(31)=0.008$, $p=.99$ $BF_{0,1}=5.06$	$t(31)=-0.11$, $p=.9$ $BF_{0,1}=5.26$
length	$t(31)=0$, $p=1$ $BF_{0,1}=5.29$	$t(31)=0.09$, $p=.92$ $BF_{0,1}=5.27$	$t(31)=0$, $p=1$ $BF_{0,1}=5.29$	$t(31)=0$, $p=1$ $BF_{0,1}=5.27$	$t(31)=0$, $p=.1$ $BF_{0,1}=5.29$

Table 2.2.: Experiment 1 – syllabicity/root priming. Within-condition statistical results of the lexical properties of related and unrelated primes across the five conditions.

By-condition statistical results are reported below in Table 2.2). All unrelated primes were suffixed (bimorphemic) words. To further guarantee exact matching, all stimuli (targets, related, and unrelated primes) were also matched in the number of syllables.

As a way to assess semantic relatedness in prime-TARGET pairs, we retrieved pairwise Latent Semantic Analysis (LSA) values for each prime-TARGET pair of all conditions; these values express the semantic similarity/relatedness between two given words (`http://lsa.colorado.edu`; Landauer and Dumais, 1997). Pairwise comparisons (see Table 2.3) show

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	$p < .0001$	$BF_{1,0}$>1000
transparent	*syllabic*	$p < .0001$	$BF_{1,0}$>1000
transparent	*non-syllabic*	$p < .0001$	$BF_{1,0}$>1000
opaque	*syllabic*	$p = 1$	$BF_{1,0}$=0.3
opaque	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.27
syllabic	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.38

Table 2.3.: Experiment 1 – syllabicity/root priming. Pairwise comparisons of the LSA coefficients across conditions.

that the LSA values for the opaque (mean: 0.07) did not vary from the LSA values for the syllabic (mean: 0.06) and non-syllabic conditions (mean: 0.08); whereas the LSA values for the transparent condition (mean: 0.43) were different from each of the other conditions. These results were expected because of the additional semantic relatedness between prime and target words in the transparent condition (which was not present in the other conditions).

Thirty-two pairs of unrelated words were added to the final item set to reduce the prime-TARGET relatedness proportion to 40%. The filler targets were matched on frequency (WL: $F(5, 186)$=1.40, p=.22, $BF_{0,1}$=7.81; HAL: $F(5,184)$=1.82, p=.1, $BF_{0,1}$=3.60) and length (F(5, 186)=1.07, p=.37, $BF_{0,1}$=13.89) to the targets of the experimental conditions; they were preceded by unrelated suffixed word primes. All word stimuli used in the experiment can be found in the Appendix I.

An additional set of 192 pseudo-word targets were chosen from English Lexicon Project corpus. All pseudo-word targets complied with English spelling and phonotactics, and were generated by changing one or two letters in a corresponding target word. All pseudo-word targets were matched to the word targets used in the experiment in length ($F(6, 377)$=.9, p=.49, $BF_{0,1}$=39.06). They were preceded by unrelated suffixed word primes.

All target stimuli were counterbalanced and divided at random into two different lists of equal number of pairs (384 in total). In each list, half of the target words were preceded by a related prime and half by an unrelated prime word. Each participant was assigned either list, so that they saw each target stimulus exactly once.

2.2.2. Participants & procedure

Eighty-four subjects participated in experiment 1 (68 females, 16 males; mean age: 19.25, s.d.: 0.82). All subjects were undergraduate students at the University of Connecticut and native speakers of American English. None of the subjects reported atypical vision and/or other reading-related neurological impairment (e.g., dyslexia). All subjects were compensated in the form of course credit.

Stimulus presentation and data recording were performed with a python script run through the PsychoPy software (Peirce, 2009). Subjects were tested individually in a dimly lit, quiet testing room. They were asked to read the capitalized letter strings on the display and decide as quickly and accurately as possible whether or not each string is a word. To do that, they used a mechanical keyboard; mechanical keyboards are accurate devices for proper recording of response time. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case for 33 ms; the target word immediately followed in uppercase, and remain on the screen until a response is made. The order of pairs was chosen randomly across participants. Participants were given 20 practice pairs before the actual experiment began. A total of 384 pairs were presented to each participant. During the experiment, participants were also given the possibility to take two brief breaks.

2.2.3. Predictions

Keeping the predictions for the identity, transparent, and opaque conditions constant (in line with all the previous literature), two possible sets of predictions are possible (see Figure 2.1).

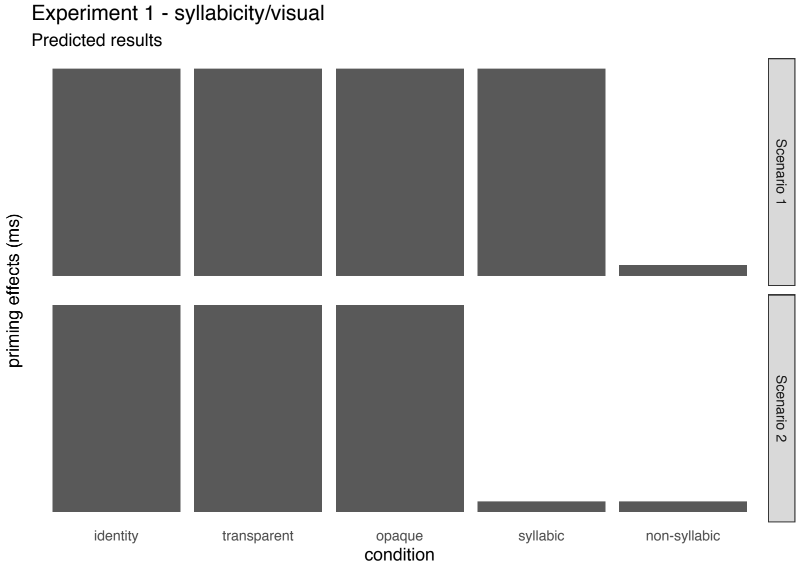

Figure 2.1.: Predicted results for experiment 1 - syllabicity/root priming.

Should we find priming effects in the syllabic condition (scenario 1), we may suggest that early decomposition relies on orthographic and/or phonological syllabification. If true, this scenario would suggest that decomposition relies on phono-orthographic (syllable-based) information, in addition to morpho-orthographic information. Additionally, it would solve the contradiction between the brothel non-effect (*brothel* not priming *BROTH* and *brother* priming *BROTH*; Rastle et al., 2004) and the slegrack/flexire effect (*slegrack* priming *RACK* and *flexire* priming *FLEX*; Fiorentino et al., 2015; Morris et al., 2011), as it suggests that the brothel non-effect merely results from the averaging between the responses to orthographically related pairs with a syllabic ending (e.g., *brothel-BROTH*) and orthographically related pairs that ended with a sub-syllabic and non-suffixal ending (e.g., *against-AGAIN*). Should we instead find no priming effect for both the syllabic and non-syllabic conditions (scenario 2), the result would then align with the previous literature on morphological priming and confirm that morphological decomposition is not driven by phono-orthographic syllabification. Furthermore, scenario 2 would suggest that the contradiction between the brothel non-effect and the slegrack/flexire effect might be affected by the different lexicality statuses of the prime word (*brothel* is a word, but *slegrack* and *flexire* are not), thus potentially challenging the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3) assumed by current models of decomposition. We are not taking into consideration the remaining two possible sets of predictions (that is, priming effects arising in both the syllabic and non-syllabic conditions or priming effects in the non-syllabic condition only) because they seem unlikely given the previous results reported in the literature.

2.2.4. Results

Response times (RTs) were measured from target onset and cleaned of outliers as follows. First, we calculated by-subject error rates for words and pseudo-words separately. Since the means of the two distributions did not vary significantly ($t(145.52)=1.52$, $p=.12$; $BF_{0,1}=2.03$), we calculated by-subject overall error rates (that is, including words and pseudo-words) and removed all subjects whose error scores were higher than 20%. Second, items were excluded from the analysis if their overall error rate was higher than 30%. Incorrect responses and fillers (words and pseudo-words) were excluded from analysis. Finally, RTs were first log10-transformed to guarantee near-Gaussian distribution as suggested by Baayen (2008); then, individual log10 RTs were excluded from the analysis if they were more than 2.5 standard deviations away from the by-subject and overall log10 mean RT. Outlier rejection resulted in excluding a total of 361 points (3.62% of the dataset). A total of 9,598 datapoints were included in the analysis.

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
identity	550	530	20	0.38
transparent	558	540	18	0.36
opaque	571	558	15	0.28
syllabic	567	561	8	0.15
non-syllabic	556	546	10	0.21

Table 2.4.: Experiment 1 – syllabicity/root priming. Mean RTs and Cohen's d (effect sizes, ES) across conditions.

Figure 2.2 and Table 2.4 above show that the identity, transparent, and opaque conditions elicited priming effects, in line with the previous literature. The priming magnitudes and effect

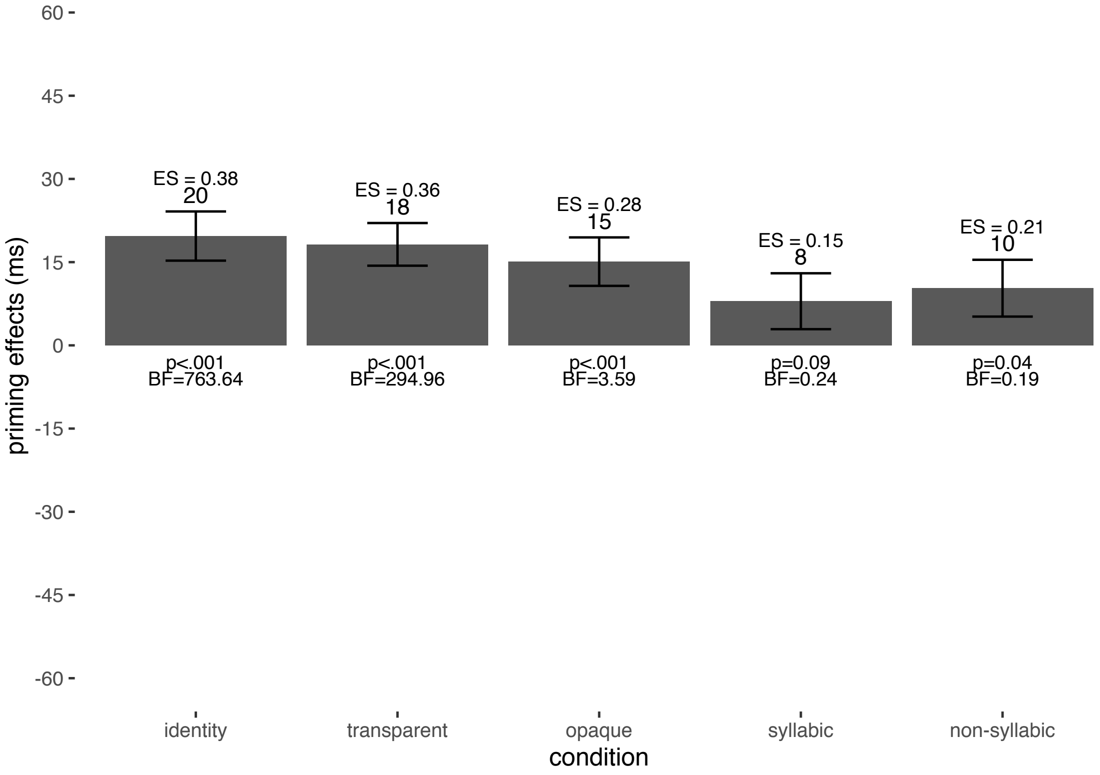

Figure 2.2.: Experiment 1 – syllabicity/root priming. Summary of the priming effects. The numbers over the bars are the mean priming magnitudes and Cohen's *d* (effect sizes, ES). The number below the bar are the *p*- and the $BF_{1,0}$-values resulting from the statistical analyses.

sizes for the syllabic and non-syllabic conditions are instead lower. We constructed a series of linear mixed-effect regression (LMER) models for each sub-design (that is, identity, transparent, opaque, syllabic, and non-syllabic; Baayen, 2008; Barr et al., 2013). Each model had *raw RT* (in ms) as the dependent variable, RELATEDNESS (2 levels: related vs. unrelated) as fixed factor, and SUBJECT and ITEM as random factors (intercept only). Estimates of each parameter of the models were determined using the restricted maximum likelihood (REML) criterion using the `lmerTest` R-package (Kuznetsova et al., 2017). P-values were estimated using the Satterthwaite approximation of degrees of freedom. We also estimated Bayes Factors for each subdesign using the R package `BayesFactor` (Morey, 2018). Recall from sec. 1.5 that Bayes factors (BFs) express the ratio of the likelihood of the data under one hypothesis (say, H_1, the alternative hypothesis) to the likelihood of the data under the other hypothesis (H_0, the null hypothesis). We calculated $BF_{1,0}$s (that is, the ratio of the probability of the data under H_1 to the probability of the date under H_0) similar to the ANOVA calculations above. For the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.5. The details of the LMER and the BF analyses are reported in Table 2.5 below. Priming effects were found to be significant in the following conditions: identity, transparent, and opaque (*p*s<.05; $BF_{1,0}$s>3). Priming effects in the non-syllabic condition reached significant in the LMER analysis, but were instead found to substantially support the null hypothesis in the

BF analysis ($BF_{1,0}$<0.33). Priming effects in the syllabic condition did not reach significance in both analyses (*p*=.09; $BF_{1,0}$s<0.33).

CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$ interpretation
identity	28.94	<.0001	763.63	extreme for H_1
transparent	24.02	<.0001	294.95	extreme for H_1
opaque	11.98	<.0005	3.59	substantial for H_1
syllabic	2.83	.09	0.24	substantial for H_0
non-syllabic	4.37	.04	0.19	substantial for $H_0 H_1$

Table 2.5.: Experiment 1 - syllabicity/root priming. Summary of the statistical results.

Pairwise comparisons of priming magnitudes were also performed across conditions. Table 2.6 reports both Dunn-corrected *p*-values and uncorrected $BF_{1,0}$-values for each combination. None of the comparisons were found to be significant.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	$p = 1$	$BF_{1,0}$=0.18
transparent	*syllabic*	$p = 1$	$BF_{1,0}$=0.2
transparent	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.35
transparent	*identity*	$p = 1$	$BF_{1,0}$=0.59
opaque	*syllabic*	$p = 1$	$BF_{1,0}$=0.30
opaque	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.43
opaque	*identity*	$p = 1$	$BF_{1,0}$=0.2
syllabic	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.18
syllabic	*identity*	$p = .7$	$BF_{1,0}$=0.71
non-syllabic	*identity*	$p = 1$	$BF_{1,0}$=0.43

Table 2.6.: Experiment 1 – syllabicity/root priming. Pairwise comparisons of the priming effects across conditions.

2.3. Experiment 1 – syllabicity/root priming (on-line)

The results reported above refer to data that was collected in an in-lab environment. We reran experiment 1 online primarily as a way to validate the online environment against the in-lab standards. The ultimate goal of this section is to show that, regardless of the evident distributional differences between the two environments, the data collected online for priming experiments is actually comparable to the data collected in lab for priming experiments. This allows us to recruit large samples of participants faster. For this reason, all the visual experiments reported in chapters 3-5 were run online. The reader that is not interested in the purely methodologically issues connected to the environment in which the experiments reported in the following chapters were run may skip this section.

In 2018, the same developers of the free-source python-based program *PsychoPy* started developing *PsychoJS* (Peirce et al., 2019). The intent was to develop an online program that could be easily accessible to the growing scientific community using PsychoPy in their labs, in order to boost collaboration among scientists all over the world (for example, by script-sharing, which may ensure replicability of results) and increase sample size (which may decrease of Type-II errors). Since the in-lab version of experiment 1 reported above was coded in PsychoPy and PsychoJS is very similar to PsychoPy, we took the chance to conduct two replications of experiment 1 on the PsychoJS platform as a way to compare the results across the two different environments and ultimately evaluate for the first time, to our knowledge, whether PsychoJS can be used to conduct priming studies online. The design of the online replications was coded in javascript as required by PsychoJS, but it was as close as possible to the in-lab version of the experiment. We ran two different online versions of the same script. In one version of the script (henceforth, *online ms version*), we set the prime duration in milliseconds; this was initially done to ensure compatibility across different monitors (with different refresh rates) that subjects might have been using during the experiment. In the other version of the script (henceforth, *online frames version*), we set the prime duration in numbers of frames to test the reliability of the javascript screen clock. The report below follows a similar outline to the report for experiment 1 (sec. 2.2), but includes an in-depth description of the pipeline adopted, the issues encountered, the data distributions of the two online versions as compared to the data distribution of the in-lab version, and the statistical analysis.

2.3.1. Materials

The materials used were the same as described for the in-lab version of the experiment (experiment 1; sec. 2.2.1). We did not include the 32 word-word and 32 pseudo-word-word filler pairs in order to to reduce the overall duration of the experiment of around 20%.

2.3.2. Participants & procedure

All subjects recruited were native speakers of American English. Eighty-four subjects participated in the experiment in the *in-lab frames version* (68 females, 16 males; mean age: 19.25, s.d.: 0.82). They were undergraduate students at the University of Connecticut and received study credits for their participation. One hundred and forty subjects participated in the *online ms version*(55 females, 84 males; mean age: 36.43, s.d.: 11.07). One hundred and thirty-eight subjects participated in the online experiment in the *online frames version* (49 females, 89 males; mean age: 35.06, s.d.: 11.51). Subjects participating in the online versions were recruited through Amazon Mechanical Turk and received monetary compensation for their participation.

In the *in-lab frames version*, stimulus presentation and data recording were performed with a python script run through the PsychoPy software (Peirce, 2009). Subjects were tested individually in a dimly lit, quiet testing room. They were asked to read the capitalized letter strings on the display and decide as quickly and accurately as possible whether or not each string is a word. To do that, they used a mechanical keyboard that ensured reliable response times. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case for 2 frames (which is roughly 34 ms in 60Hz-monitors). To monitor prime duration throughout the two runs, the code also reported prime duration both in frames and in milliseconds for each trial. The target word immediately followed in uppercase and remain on the screen until a response is made. The order of pairs was chosen randomly across participants. Participants were given 20 practice pairs before the actual experiment began. A total of 384 pairs were presented to each participant. During the experiment, participants were also given the possibility to take two brief breaks.

In the two online versions, stimulus presentation and data recording were performed with a python script run through the PsychoJS program (Peirce et al., 2019). Subjects were asked to find a comfortable and quiet room where they could take the experiment without getting distracted. They were also asked to turn off the computer's sound and their mobile devices during the experiments. The task was the same as for the in-lab versions. Response times were recorded through subjects' keyboard. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case either for 34 ms (*online ms version*) or 2 frames (*online frames version*). The target word immediately followed in uppercase and remained on the screen until a response was made. The order of pairs was chosen randomly across participants. Participants were given 10 practice pairs before the actual experiment began. During the experiment, participants were also given the possibility to take 7 brief breaks. To detect bots, subjects were also asked to answer three open-ended questions immediately before a break, during which subjects were reminded of the main task.

2.3.3. Aims

The goal of this experiment was to ensure that the online program PsychoJS could be used to conduct priming effects. Therefore, we compared the results of the two online versions of experiment 1 between each other, as well as with the results of the in-lab version of experiment 1. If the results turn out to be similar between the three versions, we can be sure that utilizing the online program PsychoJS may give us a reliable platform to test our questions on. Additionally, we also want to compare the actual prime durations in the two online versions, in order to determine which prime duration setting (i.e., in milliseconds or in number of frames) is able to present prime stimuli more reliably at the wanted duration.

2.3.4. Data analysis

For each of the two online versions of experiment 1, preliminary trimming was performed to remove trials in which the duration of the prime was longer or shorter than 34±8 ms. After this, we followed the same pipeline used in the in-lab experiment. First, subjects and words were removed if their error rate was above 20% and 30%, respectively. Incorrect responses (both words and pseudo-words) were then excluded from analysis. RTs were then log-transformed to guarantee near-Gaussian distribution as suggested by (Baayen, 2008) and to reduce the high between-subject variability. Individual log RTs were excluded from the analysis if they were more than 2.5 standard deviations away from the by-subject and overall log mean RT. At this point, statistical analyses were performed on the trimmed data. In the subsections below, we thoroughly describe each of the aforementioned steps.

Prime duration

The scripts of the three versions of experiment 1 contained a clock object that timed the duration (in milliseconds) of each stimulus (forward mask, prime, and target) for each participant. The purpose of this was two-fold. First, it allowed us to make sure that all stimuli were presented at the wanted duration (regardless of the whether the duration was set in milliseconds or in number of frames). Second, it allowed us to compare the distributions of the two online versions of experiment 1 with the in-lab version of experiment 1, to assess reliability of the online platform in terms of stimulus presentation.

In the analysis, we primarily focused on the durations of the prime words, as it substantially impinges on target recognition (Forster et al., 2003), while we set aside the durations of the forward mask. The prime durations in the online environment fluctuated substantially, as compared to the prime durations in the in-lab environment (see Figure 2.3).

Figure 2.3 suggests that, despite the fact the distributions of the prime durations in the two online versions look similar to each other, they look different from the distribution of the prime duration in the *in-lab frames version*. To compare the distribution of the prime durations across the three versions, we ran a series of two-sample Kolmogorov-Smirnov tests (henceforth, K-S tests), one for each of three possible pairwise combinations of versions. The K-S test is a non-parametric test that compares the empirical cumulative distributions (ECDFs) of two given datasets by quantifying the distance D between the two distributions. ECDFs provide the probability estimates for a data point of a variable to be less than or equal to each of the other data points of the same variable. D measures the maximum vertical distance between the two ECDFs considered; the closer D is to 0 the more likely it is that the two samples were drawn from the same distribution. Estimation of the corresponding ps will not reported here, since it is highly inaccurate in presence of ties, as in our case. We found that D was larger when comparing the ECDFs of the prime duration distributions in the *in-lab frames version* and *online ms version* (D=0.95) and the *in-lab frames version* and *online frames version* (D=0.89) than when comparing the ECDFs of the prime duration distributions in the *online frames version* and *online ms version* (D=0.07), as also showed in Figure 2.4 below. This

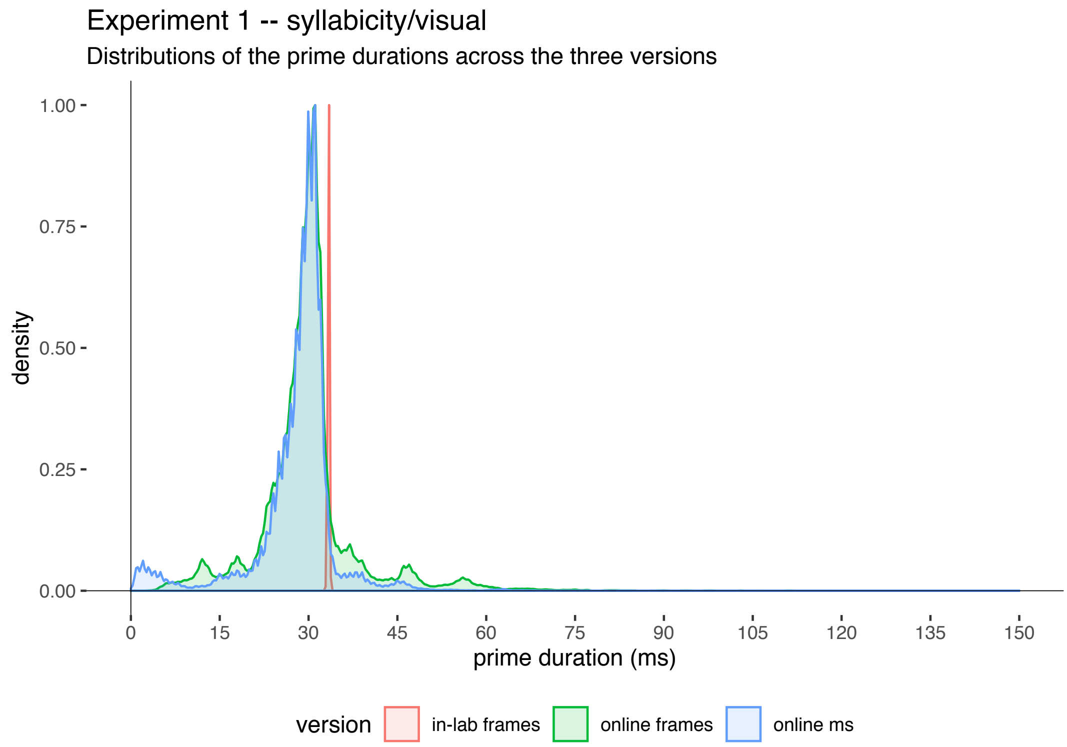

Figure 2.3.: Distributions on the prime duration across the three versions of experiment 1.

confirms that the prime duration was more reliable in the *in-lab frames version* than in either online version and that the distributions of the prime durations in the two online durations were similar to one another.

Figure 2.4 also shows that the fluctuations in the *in-lab frames version* were minimal, whereas the fluctuations in both online versions were substantial and forced us to apply a preliminary trimming of trials which had their prime lasting above and below the wanted duration (that is, 34 ms) ± 8 ms. The allowed deviation was set at 8 ms because it is half of the refresh rate (16 ms) of the most common monitors (60Hz; 1/60=16 ms). In the *online ms version*, 9,679 trials had their prime exceed the aforementioned threshold and were therefore removed (21.76% of the dataset). In the *online frames version*, 11,220 trials had their prime exceed the aforementioned threshold and were therefore removed (25.4% of the dataset).

Distributions

In all three versions, by-subject error rates did not vary for words and pseudo-words (*in-lab frames version*: $t(145.52)=1.52$, $p=.12$, $BF_{0,1}=2.03$); *online ms version*: $t(251.3)=-0.87$, $p=.38$, $BF_{0,1}=5.25$; *online frames version*: $t(256.75)=-0.76$, $p=.44$, $BF_{0,1}=5.71$). Subjects were therefore removed if their overall error rate was equal or higher than 20%. In the in-lab version of the experiment, 12.04% of the subjects (10/83) were removed; in the *online ms ver-*

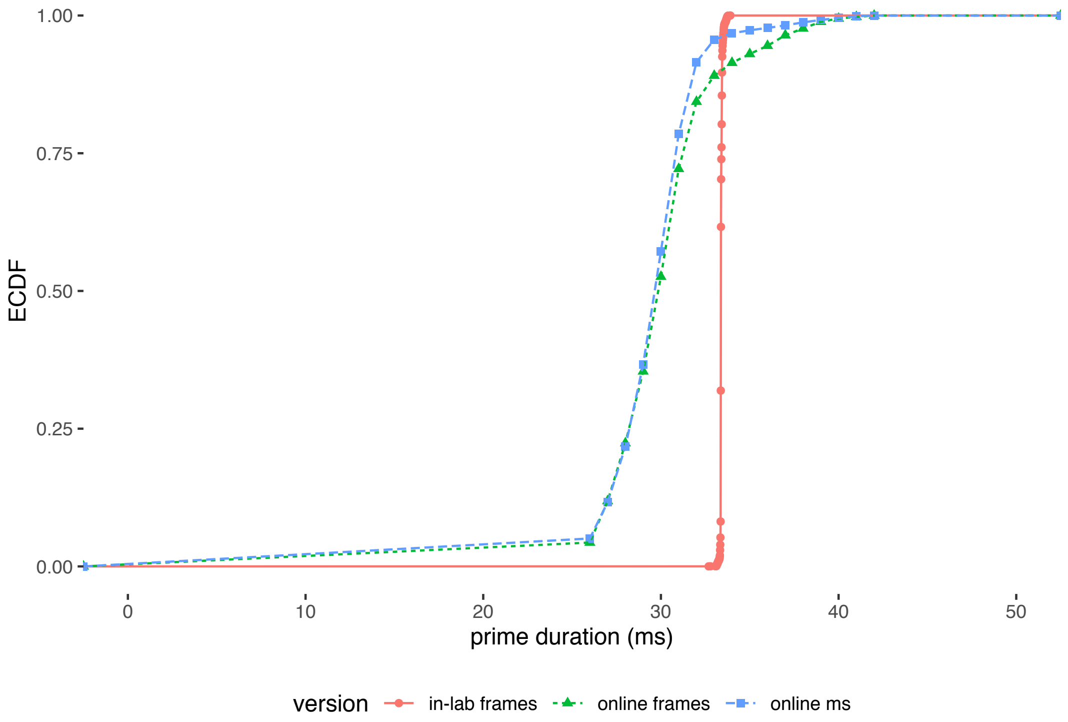

Figure 2.4.: Empirical Cumulative Density Function of the distribution of the prime duration across the three versions of experiment 1.

sion, 3.57% of the subjects (5/140) were removed; in the *online frames version*, 10% of the subjects (14/140) were removed. We compared the ECDFs of the distributions of the subject error percentages across the three versions (see Figure 2.5A). The results from the three pairwise, two-sample K-S tests revealed that the distributions were more similar between the two online versions (*online ms version* vs. *online frames version*: *D*=0.08) than between the *in-lab frames version* and either online version (*in-lab frames version* vs. *online ms version*: *D*=0.19; *in-lab frames version* vs. *online frames version*: *D*=0.22). Similarly, trials were removed that contained target words whose overall error rate was higher than 30%. In the in-lab version of the experiment, 8.3% of the words (16/192) were removed; in both the *online ms version* and the *online frames version*, 7.5% of the words (12/160) were removed. The following twelve words were removed in all of the three versions: *char, gob, gull, hag, hem, mar, opt, par, wee, whir, wit,* and *yaw*. These words were tested in the syllabic and non-syllabic conditions. The consistency in the high error rate for these words was probably due to the fact that these words are either fairly uncommon (e.g. *hem, mar, whir*) or rather used in multimorphemic forms (e.g., *charcoal*, *seagull*, *witty*). We compared the ECDFs of the distributions of the word error percentages across the three versions (see Figure 2.5B). The results from the three pairwise, two-sample K-S tests revealed that the distributions were similar in all combinations (*in-lab frames version* vs. *online ms version*: *D*=0.10; *in-lab frames version* vs. *online frames version*: *D*=0.16; *online ms version* vs. *online frames version*: *D*=0.1).

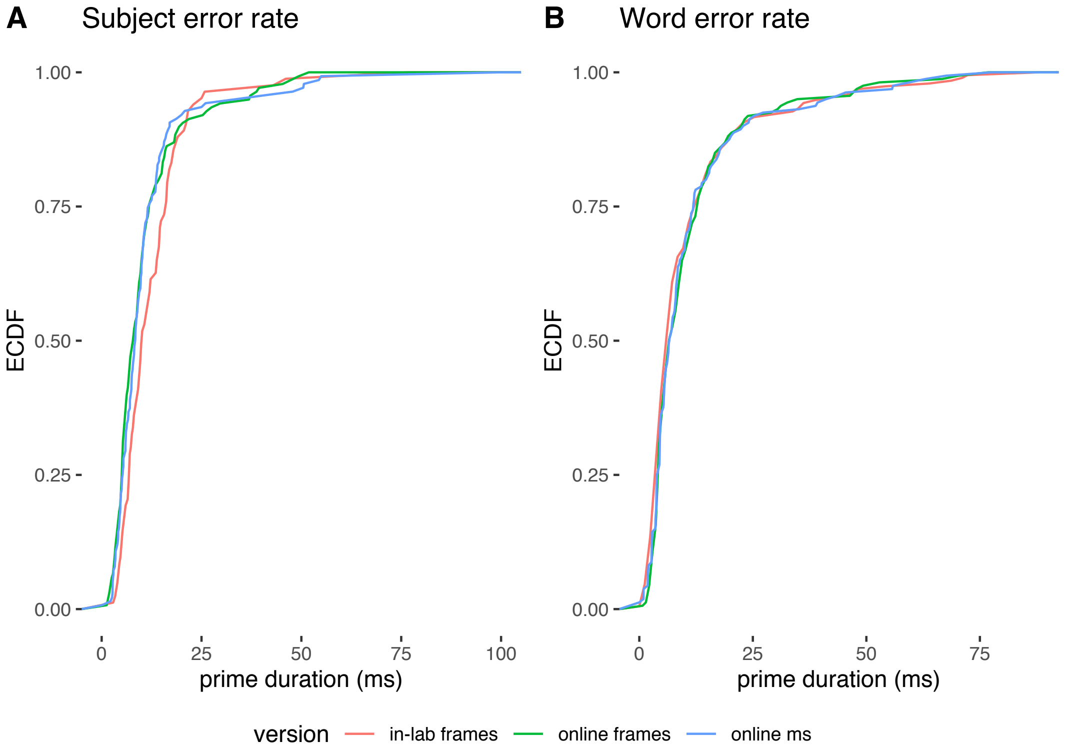

Figure 2.5.: Empirical Cumulative Density Function of the distribution of the subject error rate (A) and word error rate (B) across the three versions of experiment 1.

After removing trials by performance error, we compared sampled variance across the three versions of the experiments. The sample variance in the two online versions was much higher than the sample variance in the in-lab version of the experiment: it was 32,464.46 in the *in-lab frames version*, 190,681.8 in the *online ms version*, and 177,900.7 in the *online frames version*. Then, we removed trials whose log RTs were two standard deviations away from the by-subject and overall log RT means. The percentages of trimmed datapoints were similar across the three versions. In the in-lab version of the experiment, a total of 361 datapoints (3.62% of the dataset) were removed; a total of 9,598 datapoints were included in the analysis. In the *online ms version*, a total of 493 datapoints (3.72% of the dataset) were removed. A total of 12,001 datapoints were included in the analysis. In the *online frames version*, a total of 452 datapoints (3.73% of the dataset) were removed; a total of 12,770 datapoints were included in the analysis. Even after the RT trimming step, the sample variance in the two online versions was higher than the sample variance in the *in-lab frames version*. It was 11,636.24 in the *in-lab frames version*, 16,548.69 in the *online ms version*, and 19,254.13 in the *online frames version*. Distributions of the RTs across the three versions are plotted in Figure 2.6. We also ran three pairwise K-S tests to compare the ECDFs in all possible combinations. The RT distributions between the *in-lab frames version* and each of the two online versions differed (*in-lab frames version* vs. *online ms version*: D=0.28; *in-lab frames version* vs. *online frames version*: D=0.29) more than the RT distributions between the two online versions (*online frames version* vs. *online ms version*: D=0.02; Figure 2.7). These results suggest that the RT distributions in the three versions varied significantly, which was likely due to the high difference in the sample variances. Datasets with high variance are noisy, which leads to overfitting models that mistakenly include random error in the data. For this reason, *log RT* was used as the dependent variable in our statistical analyses of the two online versions (whereas *raw RT* was used as the dependent variable in the statistical analysis of the *in-lab frames version*). Log-transformation is a common way to reduce variability within the dataset and guarantee near-Gaussian (normal) distributions.

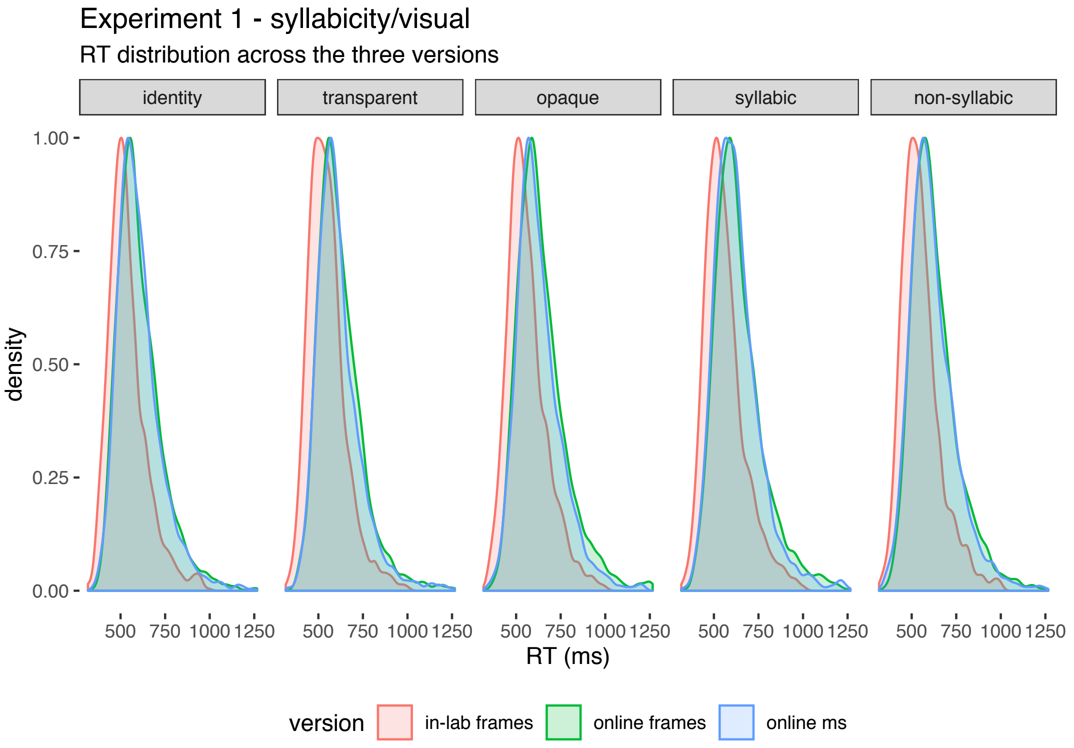

Figure 2.6.: Distributions on the RTs across the five conditions and the three versions of experiment 1.

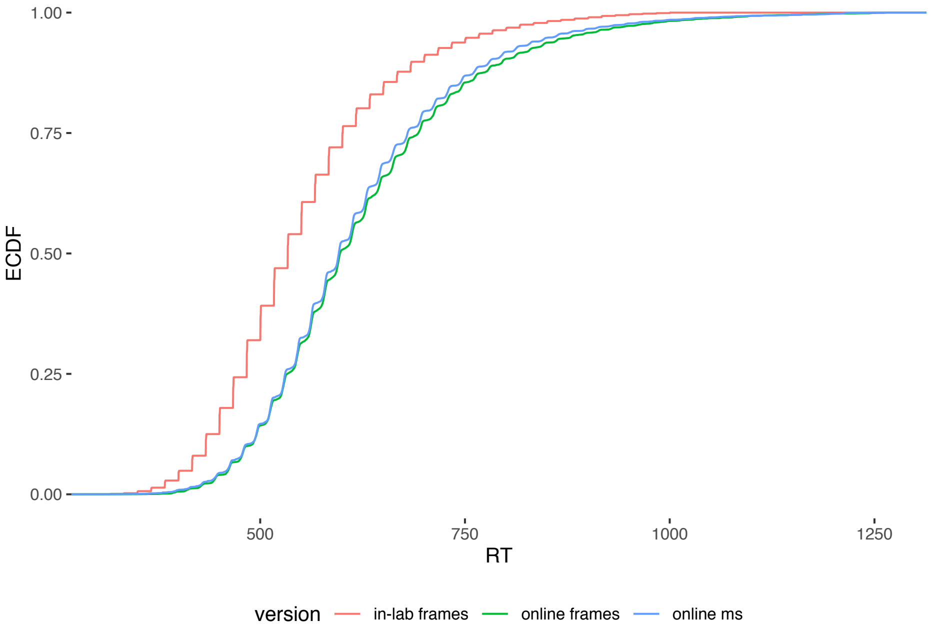

Figure 2.7.: Empirical Cumulative Density Function of the RT distribution across the three versions of experiment 1.

2.3.5. Results

Figure 2.8 and Table 2.7 report the descriptive statistics of the three versions of experiment 1. As also seen in sec. 2.2, the *in-lab frames version* shows significant priming effects for the identity, transparent, and opaque conditions and non-significant priming effects for the syllabic and non-syllabic conditions. This pattern of results is our baseline. We now compare the results of the two on-line versions as well. The *online ms version* seems to show significant priming effects for the identity and the transparent conditions and no priming effects for the opaque, syllabic, and non-syllabic conditions. The *online frames version* shows significant priming effects for all conditions. In general, we notice that priming magnitudes and effect sizes tend to be smaller in the two on-line version in all conditions, except for the transparent condition.

Both ANOVA and BF analyses were performed on the three versions of experiment 1 in a similar way. In each version, we constructed a series of LMER and BF models for each subdesign (that is, identity, transparent, opaque, syllabic, and non-syllabic; Baayen, 2008; Barr et al., 2013). In each model, the dependent variable was *raw RT* in the *in-lab frames version* and *log RT* in the two versions; the fixed factor was RELATEDNESS (2 levels: related vs. unrelated) and the random factors (intercepts only) were SUBJECT and ITEM in all three versions. Estimates of each parameter of the models were determined using the restricted maximum likelihood (REML) criterion using the `lmerTest` R-package (Kuznetsova et al., 2017). P-values were estimated using the Satterthwaite approximation of degrees of freedom. We also estimated Bayes Factors for each subdesign of each versions using the R package `BayesFactor`

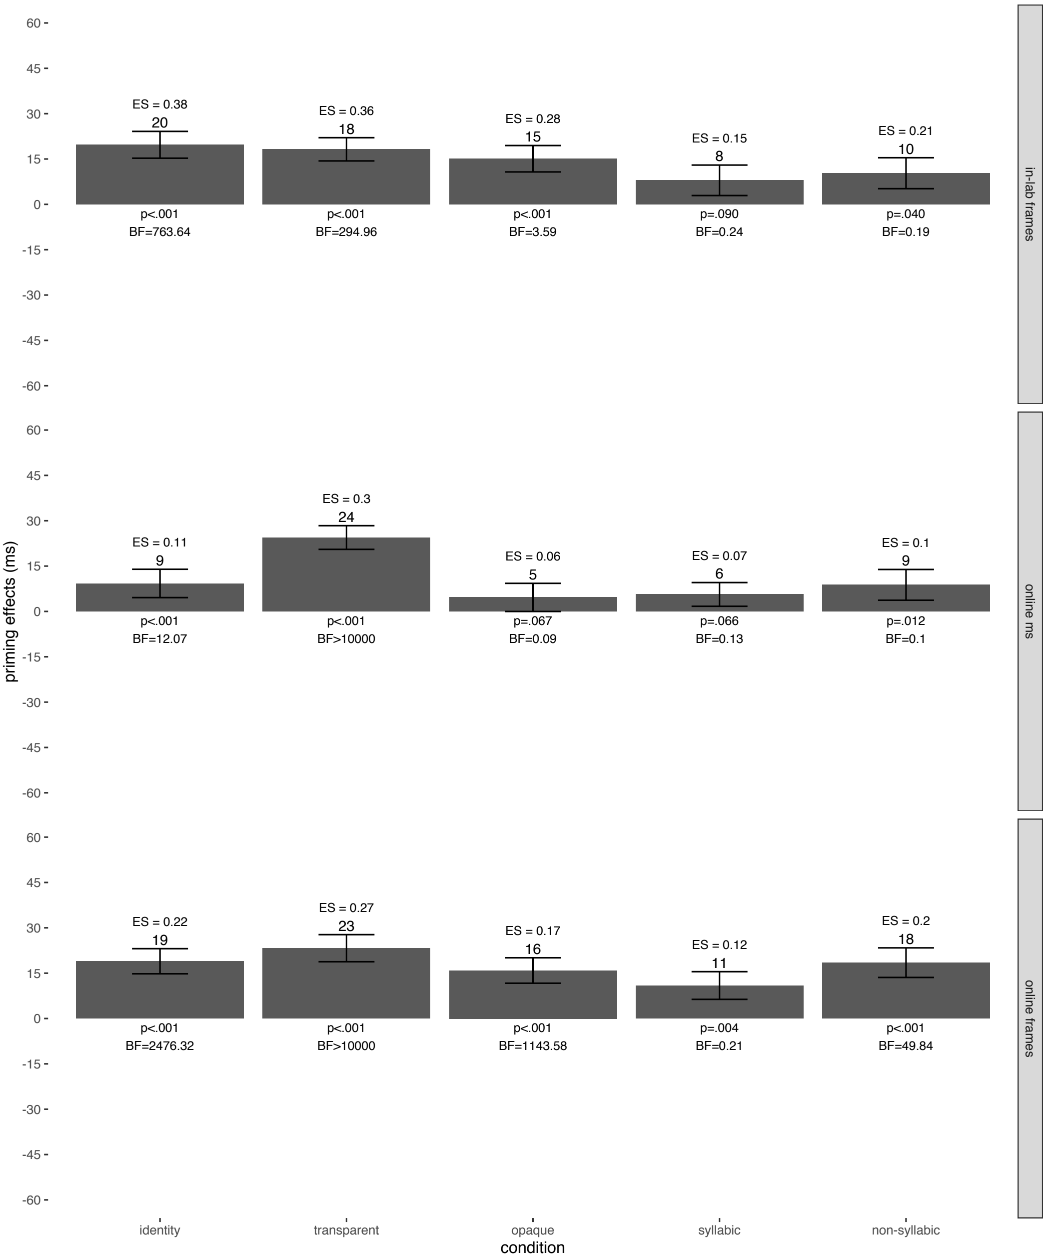

Figure 2.8.: Experiment 1 – syllabicity/root priming. Summary of the priming effects across the three versions. The numbers over the bars are the priming magnitudes and Cohen's d (effect sizes, ES); the numbers below the bar are the p- and $BF_{1,0}$-values of the LMER and BF analyses, respectively.

(Morey, 2018); for the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.5. The details of the LMER and the BF analyses across the three versions are reported in Table 2.8 below. As also seen in sec. 2.2, in the *in-lab frames version* significant priming effects were found in the identity, transparent, and opaque conditions. The priming effects of the syllabic condition were not significant in both analyses. The

(a) *in-lab frames version*

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
identity	550	530	20	0.38
transparent	558	540	18	0.36
opaque	571	558	15	0.28
syllabic	567	561	8	0.15
non-syllabic	556	546	10	0.21

(b) *online ms version*

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
identity	610	601	9	0.11
transparent	624	660	24	0.30
opaque	628	623	5	0.06
syllabic	631	625	6	0.07
non-syllabic	626	617	9	0.10

(c) *online frames version*

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
identity	619	600	19	0.22
transparent	633	609	23	0.27
opaque	652	636	16	0.17
syllabic	647	636	11	0.12
non-syllabic	639	621	18	0.20

Table 2.7.: Experiment 1 – syllabicity/root priming. Mean RTs, priming magnitudes, and Cohen's *d* (effect sizes, ES) across the three versions.

priming effects of the non-syllabic condition were instead significant in the LMER analysis, but were found to substantially support the null hypothesis in the BF analyses. In the *online ms version*, both analyses were consistent in indicating significant priming effects in the identity and transparent conditions, and non-significant priming effects in the opaque and syllabic conditions. The priming conditions in the non-syllabic condition were significant in the LMER analysis, but were found to substantially support the null hypothesis in the BF analysis. In the *online frames version*, both analyses were consistent in indicating significant priming effects in the identity, transparent, opaque, and non-syllabic conditions only. The priming effects in the syllabic condition were significant in the LMER analysis, but were found to substantially support the null in the BF analysis.

(a) *in-lab frames version*

CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$ interpretation
identity	28.94	<.001	763.63	extreme for H_1
transparent	24.02	<.001	294.95	extreme for H_1
opaque	11.98	<.001	3.59	substantial for H_1
syllabic	2.83	.09	0.24	substantial for H_0
non-syllabic	4.37	.04	0.19	substantial for H_0

(b) *online ms version*

CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$interpretation
identity	19.74	<.001	12.07	substantial for H_1
transparent	60.65	<.001	>10000	extreme for H_1
opaque	3.34	.067	0.09	extreme for H_0
syllabic	3.39	.066	0.13	substantial for H_0
non-syllabic	6.28	.01	0.10	substantial for H_0

(c) *online frames version*

CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$interpretation
identity	42.36	<.001	2476.32	extreme for H_1
transparent	50.58	<.001	>10000	extreme for H_1
opaque	19.14	<.001	1143.58	extreme for H_1
syllabic	8.23	.004	0.21	substantial for H_0
non-syllabic	27.75	<.001	49.84	strong for H_1

Table 2.8.: Experiment 1 – syllabicity/root priming. Summary of the statistical results across the three experiments.

Pairwise comparisons of priming magnitudes were also performed between conditions in all combinations in each version of the experiment 1. Table 2.9 reports both Dunn-corrected p-values and uncorrected $BF_{1,0}$-values for each combination in each version. In the *in-lab frames version* and in *online frames version*, none of the comparisons were found to be significant. In the *online ms version*, all comparisons involving the transparent conditions were found to strongly support the alternative hypothesis.

(a) *in-lab frames version.*

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	$p = 1$	$BF_{1,0}$=0.18
transparent	*syllabic*	$p = 1$	$BF_{1,0}$=0.2
transparent	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.35
transparent	*identity*	$p = 1$	$BF_{1,0}$=0.59
opaque	*syllabic*	$p = 1$	$BF_{1,0}$=0.30
opaque	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.43
opaque	*identity*	$p = 1$	$BF_{1,0}$=0.2
syllabic	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.18
syllabic	*identity*	$p = .7$	$BF_{1,0}$=0.71
non-syllabic	*identity*	$p = 1$	$BF_{1,0}$=0.43

(b) *online ms version.*

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	$p = .2$	$BF_{1,0}$=21.41
transparent	*syllabic*	$p = .03$	$BF_{1,0}$=27.09
transparent	*non-syllabic*	$p = .09$	$BF_{1,0}$=3.69
transparent	*identity*	$p = .15$	$BF_{1,0}$=2.60
opaque	*syllabic*	$p = 1$	$BF_{1,0}$=0.16
opaque	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.17
opaque	*identity*	$p = 1$	$BF_{1,0}$=0.19
syllabic	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.16
syllabic	*identity*	$p = 1$	$BF_{1,0}$=0.18
non-syllabic	*identity*	$p = 1$	$BF_{1,0}$=0.15

(c) *online frames version.*

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	$p = 1$	$BF_{1,0}$=0.30
transparent	*syllabic*	$p = .03$	$BF_{1,0}$=0.20
transparent	*non-syllabic*	$p = .5$	$BF_{1,0}$=0.88
transparent	*identity*	$p = 1$	$BF_{1,0}$=0.19
opaque	*syllabic*	$p = 1$	$BF_{1,0}$=0.21
opaque	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.17
opaque	*identity*	$p = 1$	$BF_{1,0}$=0.18
syllabic	*non-syllabic*	$p = 1$	$BF_{1,0}$=0.29
syllabic	*identity*	$p = 1$	$BF_{1,0}$=0.35
non-syllabic	*identity*	$p = 1$	$BF_{1,0}$=0.16

Table 2.9.: Experiment 1 – syllabicity/root priming. Pairwise comparisons of the priming effects across conditions in each version.

Finally, we compare the priming effects within each subdesign across the three versions to look for differences in effect sizes. We fit a one-way ANOVA and a BF models that had *logRT* as the dependent variable and VERSION (3 levels: in-lab frames, online ms, online frames) as the independent variable. Estimates of each parameter of the models were determined using the restricted maximum likelihood (REML) criterion using the `lmerTest` R-package (Kuznetsova et al., 2017). P-values were estimated using the Satterthwaite approximation of degrees of freedom. We also estimated Bayes Factors for each subdesign of each version using the R package `BayesFactor` (Morey, 2018); for the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.5. The details of the ANOVA and the BF analyses across the three versions are reported in Table 2.10. Both analyses were consistent in suggesting that there was no effect of version in the distributions of priming effects elicited in each design ($p > .15$, $BF_{1,0} < 0.33$).

CONDITION	$F(df_1, df_2)$	p	$BF_{1,0}$	$BF_{1,0}$ interpretation
identity	1.76 (2,263)	.17	0.20	substantial for H_0
transparent	0.57 (2,266)	.57	0.06	strong for H_0
opaque	2.07 (2,252)	.13	0.27	substantial for H_0
syllabic	0.37 (2,253)	.69	0.05	strong for H_0
non-syllabic	1.10 (2,249)	.33	0.11	strong for H_0

Table 2.10.: Experiment 1 – syllabicity/root priming. Summary of the statistical results of the priming effects as an effect of VERSION.

2.4. General discussion

The experiment reported in this chapter had two goals. The first and foremost goal was to test the hypothesis that the brothel and the slegrack/flexire effects are due to a phono-orthographic procedure of syllabification occurring during early decomposition. We entertained this hypothesis as a way to explain the two effects without running afoul of the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). The second goal was to validate the use of online data collection performed on the new online platform *PsychoJS* for priming experiments. This was a necessary preliminary step before using the online method for the experiments reported in following chapters. We discuss each goal in detail in the next two subsections.

2.4.1. Decomposition in contact with the lexicon

As discussed in sec. 2.1, the brothel non-effect and the slegrack/flexire effect pose theoretical and computational problems for all models of morphological decomposition. The brothel non-effect by itself challenges any connectionist-oriented model of decomposition because it suggests that a simple activation mechanism of morpho-orthographic units (namely, in the parlance of Rastle and Davis, 2008, islands of regularity) is not enough to explain the conditions under which decomposition seems to occur. Decomposition seems to need a more abstract level of computation referring to morphemes and the information associated to them; this goes against the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). Additionalle, the slegrack/flexire effect seems to further challenge the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3), since it suggests that non-words (*slegrack, flexire*) are treated differently from real words (*brothel*) from early stages of word processing. If decomposition can distinguish between a word and a non-word, it can arguably make contact with the lexicon and access the information therein. The experiments reported here was designed to overcome this theoretical impasse by entertaining an alternative, connectionism-friendly mechanism that could explain both effects without challenging MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). We hypothesized that an online procedure of phono-orthographic syllabification may occur within the decomposition procedure, along the line of Taft (2003). The influence of phono-orthographic syllabification in early morphological decomposition may indeed explain both the brothel non-effect and the slegrack/flexire effect seamlessly. On one hand, the slegrack/flexire effect would be the direct result of syllabification: non-words like *slegrack, flexire* would be loosely syllabified as $ s l e g . r a c k $, $ f l e x . i r e $, thus triggering priming on the correspondent targets (*RACK, FLEX*). On the other hand, the brothel non-effect would be a confound resulting from the averaging across the RTs to prime words of a different syllabic nature, but presented in the same condition. Primes containing a syllabic, non-morphemic ending are broken down into smaller syllabic chunks: e.g., $ b r o t h . e l $, which would trigger priming onto *BROTH*; primes containing a non-syllabic, non-morphemic ending are not further broken down: e.g., $ a g a i n s t $, which would not trigger priming onto *AGAIN*.

The results of the two experiments reported in this chapter suggest that syllabification does not drive in visual decomposition, thus further validating the brothel non-effect. Our results showed that priming arose in morphologically transparent pairs (*aiming-AIM*), in morphologically opaque pairs (*belly-BELL*), and in identity pairs (*fuss-FUSS*); priming instead did not arise in the orthographically-related syllabic (*ban.jo-BAN*) and non-syllabic (*starch-STAR*) pairs. As already discussed in sec. 1.4.2, these results cannot be fully explained in any of the models considered in this dissertation. Here we briefly summarize again whether and how each of the four models explains the brothel non-effect.

A. BIN MODEL

In the bin model, the masked environment of the priming design allows the visual stimulus to go through the entry-opening stage only. At this stage, all entries that are orthographically similar to the stimulus are opened. Under the assumption that the entry-opening mechanism triggers priming effects onto the target stimulus, the syllabic and non-syllabic conditions are expected to prime as much as the transparent and the opaque conditions. The prediction is not met however, since both the syllabic and non-syllabic conditions do not trigger significant priming effects. The bin model therefore can explain the priming effects in the identity condition (*fuss-FUSS*), but fails to explain the differential effect between morphologically related pairs and morphological unrelated pairs, which has been consistently reported in the priming literature and confirmed in the current experiment.

B. RACE MODEL

In the race model, the mechanisms underlying the parsing route are not explicit enough to explain these results. In particular, after the segmentation stage blindly decomposes the visual stimulus, it sends the segmented strings to the licensing stage, where any matched lexical entry is activated. Therefore, our orthographic conditions (namely, the syllabic and non-syllabic condition) should trigger priming effects, along with our morphological (transparent and opaque) conditions.

C. FULL-DECOMPOSITION MODEL

In the full-decomposition model, when a bimorphemic word (*aiming*) or a morphologically opaque word (*belly*) is presented, they are decomposed in the decomposition stage depending on morpho-orthographic islands of regularity (Rastle and Davis, 2008; *aiming*→aim-ing, *belly*→bell-y). The segmented strings are then sent to the lookup stage, where they match with and activate the corresponding existing entries. By the time the related target is presented (*AIM, BELL*), the relative entry is already activated, thus leading to priming. When monomorphemic prime words that do not contain a potential root (e.g., *fuss*) are presented, they do not decompose in the decomposition stage because one of segmented strings is low-probability island of regularity (e.g., \$ s s \$). The whole string is therefore sent as-is to the lookup stage, in which the relative lexical entry {*fuss*} activates and leads to priming onto the related target (*FUSS*). When monomorphemic primes that contain a potential root (e.g., *banjo, starch*) are presented, they do not decompose in the decomposition stage because one of segmented strings is a low-probability cluster (e.g., \$ j o \$, \$ c h \$). When the contained root target (*BAN, STARCH*) is presented, priming does not arise.

D. MORPHO-ORTHOGRAPHIC MODEL

In the morpho-orthographic model, a decomposability constraint is implemented within the morpho-orthographic level (although Crepaldi et al., 2010 propose no computational implementation of it), and allows the morpho-orthographic segmentation to only break down words that are entirely decomposable into morphemes. Therefore, morphologically transparent and opaque prime words are broken down in the morpho-orthographic level and sent to the orthographic lexicon; at this level, they activate the related nodes ({*aim*}, {*bell*}), which in turn trigger priming onto the related target (*AIM, BELL*). Conversely, monomorphemic words (*fuss, banjo, starch*) are not further broken down in the morpho-orthographic level because they are not fully decomposable, in compliance with the decomposibility constraint. They are therefore sent as-is to the orthographic lexicon, in which the whole-form nodes are activated. If the target is identical to the corresponding prime word (as in the identity condition), priming arises because the relative lexical entry was indeed activated; if the target is the contained root (as in the syllabic and non-syllabic conditions), priming does not arises because the relative lexical node ({*ban*}, {*star*}) was not previously activated.

It is worth underlining once again that the decomposition procedure argued in the full-decomposition model (C) and the morpho-orthographic model (D) hinges on a descriptive, rather than *mechanistic*, explanation for the brothel non-effect. Both models claim that the non-morphemic string \$ e l \$ inhibits activation of the morphemic string \$ b r o t h \$ (i.e., Crepaldi et al., 2010's "decomposability constraint"). No actual mechanism has indeed been proposed yet, as it is computationally difficult to implement in the connectionist infrastructure these models are built on. Crepaldi et al. (2010) point out that implementation difficulties arise especially when orthographic identification and integration are assumed to occur *serially* from left to right (e.g., Rumelhart and McClelland, 1982). When a word like *banjo* is presented, letters are decoded from left to right and sequentially sent to the word processing system. Therefore, it would be hard for, say, the string jo to prevent the previous string ban from activating the corresponding lexical entry {*banjo*}. Here, we additionally point out that similar difficulties arise if assuming a *parallel* model of orthographic processing (as growing evidence seems to support; for a review, Grainger, 2018).[7] Even assuming the full orthographic string is decoded in parallel before being sent to the word processor, the substring $broth$ is still expected to activate the corresponding morpho-orthographic node and the substring el has no way to inhibit it. Their different morphological statuses – in other words, a "decomposability constraint" of some sort – seem to be somehow necessary.

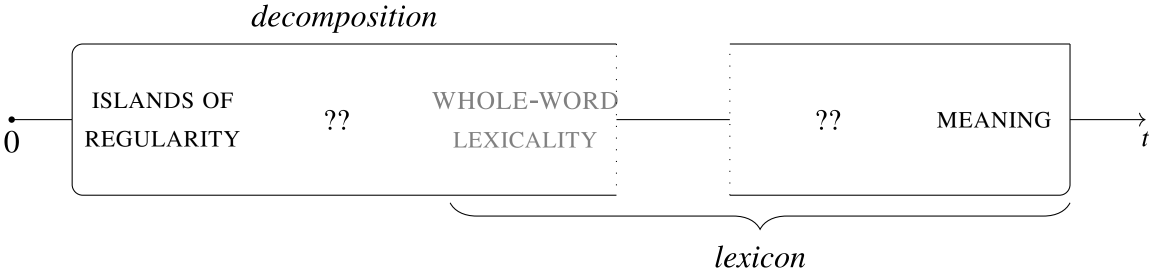

Figure 2.9.: Experiment 1 – syllabicity/root priming. Time-course of the unfolding of word information in word processing (updated version after ch. 2). The property of whole-word lexicality is grayed-out because although the experiment was not designed to test it directly, it does indirectly support it.

In rejecting the hypothesis that syllabification drives early decomposition, these results further validate the brothel non-effect and, indirectly, the slegrack/flexire effect. As a consequence, the possible explanation for the two effects that we are left with involves a procedure of decomposition that has access to the lexicality status of the whole-string stimulus (as shown in Fig. 2.9), in contrast with the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). When a mono-morphemic, root-containing word is presented (e.g., *brothel*), decomposition does not happen (brothel ↛ broth-el) in compliance with the decomposability constraint discussed above (whatever its computational implementation is); therefore, no priming arises onto the target root (*BROTH*). When a root-containing non-word is presented (e.g., *slegrack, flexire*), decomposition seamlessly happens (slegrack → sleg-rack, flexire → flex-ire), so that priming can arise onto the target root (*RACK, FLEX*). These results further suggest that some lexical properties may be available early on in visual word processing, and ultimately validates the present investigation. In the light of these results, the following chapters will indeed explore the sensitivity of decomposition to additional linguistic and lexical properties.

---

[7] As a side note, we want to point out that this parallel view of orthographic processing, which neatly ties in with visual and auditory modalities, may follow different computational and neuronal pathways, as the literature on neighborhood density/phonotactic probability seem to suggest (Vitevitch and Luce, 1998, 1999).

2.4.2. The in-lab and the on-line experimental environments

The second goal of this chapter was to report a validation study regarding the use of the online platform *PsychoJS* for priming experiments. Running experiments online allows for faster data collection from a larger sample size. However, no priming experiments run online had yet been tested, so a thorough comparison between the two environments was necessary to validate the new method.

Our analyses compared the distribution of three different versions of experiment 1: *in-lab frames version*, *online frames version*, and *online ms version*. The three versions of experiment 1 showed both similarities and differences in both the distributions and the statistical results taken into account here. First, the online experiments presented the prime duration issue, which forced us to remove, on average, about 20% of each dataset. The fluctuations were likely due to two factors. First, parallel execution of multiple programs substantially slows down CPU-specific routines, such as frame flipping. This could be prevented in the in-lab environment, but could not in the online environment. Second, unlike PsychoPy, PsychoJS first calls the browser requesting the flip, which then calls the operating system; that additional step might have added further delays in the stimulus presentation. Pairwise comparisons further revealed that the two online versions were also different from one another. A two-sample t-test on the two online versions confirmed that ($t(63314)=23.14$, $p<.0001$; 95% CIs: [1.77, 2.10]; $BF_{1,0}>10000$). This result is not surprising if we take into account that the analysis included a large number of datapoints. Type-I errors (i.e., false positive results) are more likely with such a large sample size. As a consequence, we suggest that the significant difference among the two versions may not be actually meaningful, as also suggested by the similar distributions (see Figure 2.3).

Additionally, we reported an increase of sample variance in the two online versions, as compared to the in-lab version. The high sample variance in the online versions might be connected to the different tools for RT recording. In the in-lab version, subjects used the same fully mechanic keyboard, which is able to record RTs quite reliably; in the online version, subjects used their own keyboards, which are usually semi-mechanic and may be unreliable for RT recording.

Finally, the statistical analyses on the two online versions of experiment 1 revealed some commonalities as well as some differences with the in-lab version. In all versions, both the identity and the transparent conditions primed, and the syllabic condition did not prime. The *online ms version* showed no priming effects for the opaque and non-syllabic conditions, while the *online frames version* showed priming effects for the opaque condition and the non-syllabic condition (see Table 2.11). Effect sizes in the online versions were generally lower than the effect sizes of the in-lab version (see Table 2.12). All of the above was also the reason why the sample size of the online version was nearly twice as much the sample size of the in-lab version.

	IDENTITY	TRANSPARENT	OPAQUE	SYLLABIC	NON-SYLLABIC
in-lab version	✓	✓	✓	✗	✗
ms online version	✓	✓	✗	✗	✗
frame online version	✓	✓	✓	✗	✓

Table 2.11.: Summary of the priming effects elicited in the three versions of experiment 1. *Legend*: ✓, significant priming effects; ✗, no priming effects.

	IDENTITY	TRANSPARENT	OPAQUE	SYLLABIC	NON-SYLLABIC
in-lab version	0.38	0.36	0.28	0.15	0.21
ms online version	0.11	0.30	0.06	0.07	0.10
frame online version	0.22	0.27	0.17	0.12	0.20

Table 2.12.: Summary of the effect sizes in the three versions of experiment 1.

The differences mentioned above are not surprising when all the issues described in sec. 2.3.4 are taken into account altogether. All of these issues – prime duration fluctuations, the unreliable RT recorder, and high sampled variance – contributed to an increase in noise in the data. This also explains the effect sizes being substantially lower in the online versions than in the in-lab version. On one hand, the increased noise in the data seems inevitable in the current state of development of the program and needs to be taken into account when interpreting the results obtained in the online environment. Even though the differential results obtained for the smallest effects (i.e., in the opaque and the non-syllabic conditions) could be traced back to the noise in the online data, it may still suggest that the online environment suffers from inherent limitations that make it less reliable than the in-lab environment. On the other hand, the consistency of the results for the largest effects (i.e., in the identity and transparent conditions) obtained across the three versions is reassuring because it suggests that these effects are reliable and replicable even under the adverse circumstances that online experiments are affected by. As our goal is to elicit large effect sizes, we therefore deemed *PsychoJS* reliable enough to be used for the questions investigated in the following chapters. In the next experiments, the prime duration was set in milliseconds for two reasons. First, it ensured compatibility of the design across monitors with potentially different refresh rates. Second, the significant results obtained in the *online ms version* (in the identity and the transparent conditions only) were stricter than the significant results in the *online frames version* (in the identity, transparent, opaque, and non-syllabic conditions), and therefore more similar to the results obtained in the in-lab version.

Chapter 3.

Morphological decomposition and *dominance*

3.1. Introduction

In Chapter 2, we tested the hypothesis that decomposition is driven by an on-line procedure of syllabification, as a way out of the theoretical impasse revealed by the brothel and slegrack/flexire effects. Our results, however, rejected this hypothesis and suggested that decomposition may, indeed, have access to at least some properties that are commonly considered to be stored in the lexicon, in contrast with the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). In the upcoming three chapters, we build on this main finding and explore the types of linguistic information that decomposition has access to. In this chapter, we test the sensitivity of decomposition to frequency-based asymmetries within *morpho-syntactic alternations*. In morpho-syntactic alternations, alternants realize opposing morpho-syntactic features. For example, in English, the number features [SG] ∼ [PL] are realized in the nominal domain as the alternation *ø∼s*: e.g., *window-ø∼window-s* .[8]

(7) a. world- ø
$\sqrt{\text{WORLD}}$ SG
b. world- s
$\sqrt{\text{WORLD}}$ PL

(8) a. window- ø
$\sqrt{\text{WINDOW}}$ SG
b. window- s
$\sqrt{\text{WINDOW}}$ PL

Lexical decision studies have shown that recognition of plural forms as in (7b)-(8b) depends on the surface frequencies of the singular and the plural forms of the same noun (*dominance*; Baayen et al., 2007, 2002, 1997). Plural forms such as (8b), for which the surface frequency of the plural form *windows* is higher than the surface frequency of the singular form *window*, are visually recognized faster than plural forms such as (7b), for which the surface frequency of the plural form *worlds* is lower than the surface frequency of the singular form *world.* This asymmetrical effect has been used as an argument in favor of the dual-route race model proposed by

---

[8] The formalization of number as different features is merely for clarity reasons. Here we do not commit to a specific position in the long-standing theoretical debate regarding the use of binary features (e.g., [±SG]/ [±PL]) and unary/privative features (e.g., [SG]/[PL]), as it is not relevant for the current purposes.

Schreuder and Baayen (1995). Recall that in this model words are visually recognized through a process in which two independent routes race in parallel against one another: the *parsing route*, in which words are decomposed according to language-specific phono-ortographic and morpho-orthographic rules before being recognized; and the *storage route*, in which words are recognized as whole units without getting decomposed (see sec. 1.4.2 for further details). In the storage route, the whole stimulus is searched through the lexicon, in which the lexical entries are organized by frequency. In the parsing route, the visual stimulus is first automatically decomposed into morphemes (decomposition step), which are then checked with their specific subcategorization properties (e.g., syntactic categories and affixal selectional restriction; licensing step). Even though all stimuli go though both routes, the response time at which a word is recognized as such reflects the winning route. Since the storage route looks through the lexicon by frequency, high-frequency words are recognized faster than low-frequency words because they are recognized as whole units without being decomposed. The same mechanism can be applied for singular- and plural-dominant plural forms. When a singular-dominant plural form (e.g., *worlds*) is presented, the parsing route wins the race because the low-frequency plural form is decomposed via the parsing route faster than it is looked up via the storage route. When a plural-dominant plural form (e.g., *windows*) is presented, the storage route wins the race because the high-frequency plural form is looked up via the storage route faster than it is decomposed via the parsing route.

In this chapter, we take on the issue regarding the potential impact of number dominance onto decomposition. To this end, in the experiment reported below, both singular-dominant (e.g., *worlds*) and plural-dominant (e.g., *windows*) plural words were presented as primes. All primes were followed by singular-dominant plural target words (e.g., *HEAVENS, GODS*), so to ensure that all targets could equally benefit from potential priming effects arising from prime presentation. If dominance impinges on decomposition, priming is expected to arise only when singular-dominant plural targets are preceded by singular-dominant plural primes; no priming is expected to arise instead when singular-dominant plural targets are preceded by plural-dominant plural primes. Among the four models of decomposition considered in this dissertation (sec. 1.4.2), only the bin model predicts such an effect. Recall that, in this model, entries are opened in a frequency-based orderly fashion. Accordingly, singular-dominant plural form entries are opened later than the corresponding singular form entries, whereas plural-dominant plural form entries are opened earlier than the corresponding singular form entries. The difference in the entry-opening timing would then bring about the dominance effect. On the other hand, the remaining models taken into consideration here predict that no dominance effect arises in masked priming. In the race model, whole-word frequency (and therefore dominance) is expected to affect the lexical search occurring in the storage route, and ultimately, impinge on recognition (and therefore RT); however, it is not expected to impinge on the decomposition procedures occurring in the parsing route, which are argued to occur automatically. As such, in a masked priming environment, prime words with different dominance (e.g., *windows, worlds*) are therefore equally expected to decompose and therefore trigger priming onto the corresponding related target (*GODS, HEAVENS*). Finally, in both the full-decomposition and the morpho-orthographic models, decomposition is argued to exclusively rely on orthographic statistical regularities and whole-word frequency is expected to affect word processing at later stages. In contrast with the previous predictions, our results suggest that dominance indeed modulates the priming response onto English regular plural forms. Singular-dominant plural prime words triggered priming onto singular-dominant plural target words, while plural-dominant plural prime words did not. These results ultimately seem to challenge the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3), by which decomposition should occur before lexical properties (such as whole-word frequency) are accessed. In sec 3.2.5, we further discuss the theoretical implications of these findings and some potential modifications to the models to try to account for these effects; finally, we consider some follow-up questions that would help further understand the impact of dominance on decomposition.

3.2. Experiment 2 – dominance/affix priming

3.2.1. Materials

One-hundred and sixty pairs were selected from the Worldlex corpus (WL: Gimenes and New, 2016) and the English Lexicon Project corpus (ELP: Balota et al., 2007), and organized in eight conditions of 20 pairs each. The conditions were divided in three different groups. The conditions of group 1 consisted of words that were all unequivocally [z]-ending plurals forms. All targets were singular-dominant to ensure that all targets decomposed and could benefit from the potential priming facilitation. Primes were either singular-dominant or plural-dominant, in order to look at potential effects elicited by manipulating their dominance.

1. Group 1
    a) In the *sgdom-sgdom condition*, pairs consisted of singular-dominant primes and singular-dominant targets (e.g., *worlds-HEAVENS*).
    b) In the *pldom-sgdom condition*, pairs consisted of plural-dominant primes and singular-dominant targets (e.g., *windows-GODS*).

We calculated *dominance* δ as the difference between the log frequency of the plural form $log f_{\text{PL}}$ and the log frequency of the corresponding singular form $log f_{\text{SG}}$. A given plural form was considered *'plural-dominant'* if the difference was greater than 0; e.g., *windows*. A given plural form was instead considered *'singular-dominant'* if the difference was lower than 0; e.g., *worlds*. Figure 3.1 shows the mean dominance ratios of the primes and the targets used in the two conditions of group 1. In the plot, only the mean dominance ratio of the primes of the pldom-sgdom conditions is above 0 (which indicates plural-dominance), whereas the mean dominance ratio of the prime of the sgdom-sgdom condition is below 0 (which indicates singular-dominance).

While group 1 elicited suffix priming while manipulating dominance of the prime word, group 2 tested the general effect size of suffix priming effects. The *er*-condition tested morphological suffix priming; the rhyme condition tested orthographic priming onto the right-side portion of the letter string.

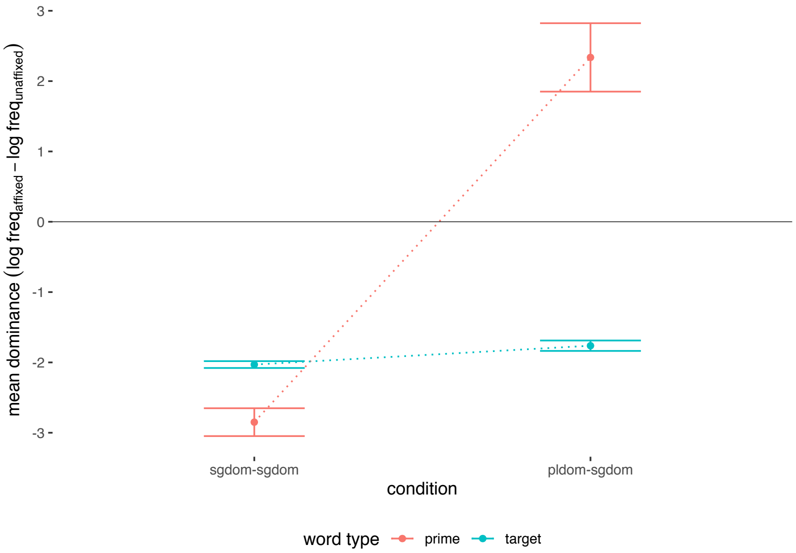

Figure 3.1.: Experiment 2 - dominance/affix priming. Plot of the mean dominance ratios in the primes and targets of the sgdom-sgdom and pldom-sgdom conditions.

2. Group 2

   c) In the er-*condition*, pairs shared the -er suffix (e.g., *scorer-WORKER*).

   d) In the *rhyme condition*, the prime and the target words rhyme with one another (e.g., *taper-VAPOR*).

Finally, group 3 tested the typical root priming pattern as shown in Rastle et al. (2004). This was done as a “sanity check”, to make sure the online environment in which the experiment was run did not affect the priming response in any way. The conditions of group 3 consisted of the word pairs used in experiment 1 (see sec. 2.2.1).

3. Group 3

   e) In the *transparent condition*, the pairs carried a semantically and morphologically transparent relationship (e.g., *boneless-BONE*).

   f) In the *opaque condition*, the pairs carried an apparent morphological relationship, but not a semantic one (e.g., *belly-BELL*).

   g) In the *syllabic condition*, the pairs consisted of prime words (e.g., *banjo*) that were made of the corresponding target word (*BAN*) plus an additional syllabic pseudo-suffix (-*jo*).

   h) In the the *non-syllabic condition*, the pairs consisted of prime words (*starch*) that were made of the corresponding target word (*STAR*) and an additional consonantal, non-syllabic pseudo-suffix (-*ch*).

Calculations of the lexical properties of the words were performed within each group and not across groups. Primes and targets in group 1 were matched as closely as possible on frequency and length. Conditions in group 2 were matched in frequency only because of inherent orthographic differences across the words used. The tables below show the mean values of the relevant lexical properties in group 1 ans 2 (for the mean values for the conditions of group 3, see the relative sections).

PROPERTY		CONDITIONS		STATISTICAL ANALYSIS
		sgdom-sgdom	*pldom-sgdom*	
PRIMES	WL frequency (log10)	0.805	0.715	$F(1,38)=.18, p=.66$ $BF_{0,1}=3$
	HAL frequency (log10)	6.74	6.88	$F(1,38)=.07, p=.78$ $BF_{0,1}=3.14$
	length	7.35	7.4	$F(1,38)=.01, p=.89$ $BF_{0,1}=3.21$
TARGETS	WL frequency (log10)	1.50	1.18	$F(1,38)=2.94, p=.1$ $BF_{0,1}=1.02$
	HAL frequency (log10)	7.34	7.74	$F(1,38)=.88, p=.35$ $BF_{0,1}=2.27$
	length	6.8	6.65	$F(1,38)=.12, p=.72$ $BF_{0,1}=3.07$

Table 3.1.: Experiment 2 – dominance/affix priming (group 1). Summary of the lexical properties of the stimuli used.

One hundred and six unrelated prime words were selected for each target word; these words were orthographically, morphologically and semantically unrelated to targets and were matched as closely as possible on frequency (Group 1. WL: $t(39)=0.31$, $p=.75$, $BF_{0,1}=5.88$; HAL: $t(39)<.0001$, $p=1$, $BF_{0,1}=5.86$. Group 2. WL: $t(39)=-0.58$, $p=.56$, $BF_{0,1}=8.11$; HAL: $t(39)=-1.03$, $p=31$, $BF_{0,1}=7.73$) and length (Group 1. $t(39)=-1.22$, $p=22$; $BF_{0,1}=2.92$. Group 2 was not controlled for length). By-condition statistical results are reported below in Table 3.3). The stimuli used in the experiment can be found in Appendix II.

Finally, a set of 160 pseudo-word targets were chosen, which matched to word targets of groups 1 and 3 in length ($F(3, 216)=1.35$, $p=.25$; $BF_{0,1}=4.47$); they were preceded by unrelated suffixed word primes that were not used in the experiment.

Target words from each condition were counterbalanced and divided at random into two versions of equal number of pairs. In each version, half of the target words were preceded by

PROPERTY		CONDITIONS		STATISTICAL ANALYSIS
		er	*rhyme*	
PRIMES	WL frequency (log10)	0.58	0.79	$F(1,38)=2.10$, $p=.15$ $BF_{0,1}=1.41$
	HAL frequency (log10)	6.50	6.46	$F(1,38)=0.06$, $p=.79$ $BF_{0,1}=3.15$
TARGETS	WL frequency (log10)	1.36	1.44	$F(1,38)=.31$, $p=.57$ $BF_{0,1}=2.85$
	WL frequency (log10)	8.41	8.43	$F(1,38)=.01$, $p=.82$ $BF_{0,1}=3.22$

Table 3.2.: Experiment 2 – dominance/affix priming (group 2). Summary of the lexical properties of the stimuli used.

PROPERTY	CONDITIONS			
	group 1		group 2	
	sgdom-sgdom	*pldom-sgdom*	*er*	*rhyme*
WL frequency (log10)	$t(19)=-0.03$, $p=.97$ $BF_{0,1}=4.30$	$t(19)=0.44$, $p=.66$ $BF_{0,1}=3.94$	$t(19)=-1.31$, $p=.20$ $BF_{0,1}=2.04$	$t(19)=0.49$, $p=.63$ $BF_{0,1}=3.86$
HAL frequency (log10)	$t(19)=0$, $p=1$ $BF_{0,1}=4.30$	$t(10)<0.0001$, $p=1$ $BF_{0,1}=4.30$	$t(21)=-1.04$, $p=.31$ $BF_{0,1}=2.65$	$t(21)=-0.47$, $p=.64$ $BF_{0,1}=3.73$
length	$t(19)=-0.84$, $p=.41$ $BF_{0,1}=3.14$	$t(19)=-0.91$, $p=.37$ $BF_{0,1}=2.97$	–	–

Table 3.3.: Experiment 2 – dominance/affix priming. Within-condition statistical results of the lexical properties of related and unrelated primes across conditions.

a related prime and half by an unrelated prime. Participants received only one version of the word list, so that they saw each target word exactly once.

3.2.2. Participants & procedure

One hundred and thirty-one participants (53 females, 78 males; mean age: 33.34, s.d: 9.38) were recruited through Amazon Mechanical Turk and received monetary compensation for their participation. They were all native speakers of American English.

Stimulus presentation and data recording were performed on-line through PsychoJS (Peirce et al., 2019), the Javascript equivalent of PsychoPy, on the Pavlovia platform (`www.pavlovia.org`). As in the lab environments, the subjects recruited online were asked to read the capitalized letter strings on the display and decide as quickly and accurately as possible whether or not each string is a word. To do that, they used their own monitor and keyboard. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case for 33 ms. The target word immediately followed in uppercase, and remained on the screen until a response was made. The order of pairs was chosen randomly across participants.

Participants were given 10 practice pairs before the actual experiment began. A total of 320 pairs were presented to each participant. During the experiment, participants were also given the possibility to take 7 brief breaks. To detect bots, subjects were also asked to answer three open-ended questions immediately after a break. To help subjects refocus on the main task post-break, we also made sure that first five trials presented after each break were pseudo-word trials.

3.2.3. Predictions

Two main scenarios are under consideration here (Figure 3.2). Similar priming arising in both the sgdom-sgdom and the pldom-sgdom conditions (Figure 3.2, scenario 1) suggests that decomposition occurs independently of number dominance effects, as predicted by the race model, the full-decomposition model, the morpho-orthographic model, and in compliance with the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). Priming effects arising in the sgdom-sgdom condition only (Figure 3.2, scenario 2) suggests instead that decomposition is affected by dominance, as predicted by Forster (1999)'s bin model of lexical access.

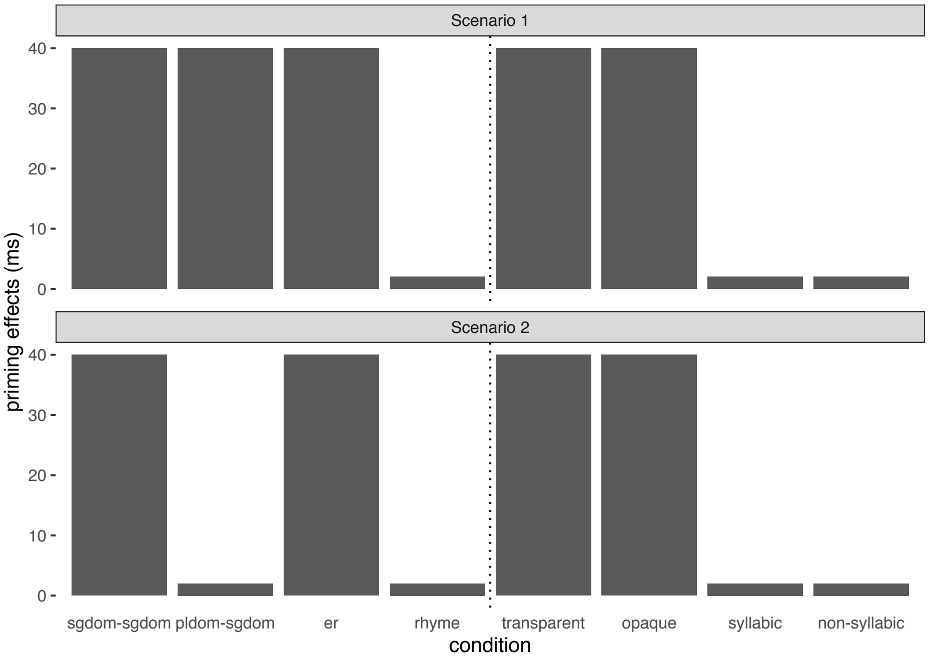

Figure 3.2.: Predicted results for experiment 2 – dominance/affix priming.

The purpose of the remaining conditions was just to ensure cross-comparison of the effect sizes elicited. The rhyme, syllabic, and non-syllabic conditions are expected not to prime, as they all involve orthographically related pairs. The *er*-, transparent, and opaque conditions are instead expected to prime as they all involve morphologically related pairs, in line with the results reported in the literature. Similar to scenario 3 above, if any of these expected results were not met, it may suggest that some methodological issue has occurred during data collection.

3.2.4. Results

In compliance with the pipeline adopted for online experiments, a preliminary trimming was performed to remove trials in which the duration of the prime was longer or shorter than 34±8 ms (see 2.3 for further details). In this experiment, 7,781 trials (18.56% of the dataset) were removed because of this. The usual pipeline was followed. First, we calculated by-subject error rates for words and pseudo-words separately. Since the means of the two distributions did not vary significantly ($t(260)=0.24$, $p=.80$; $BF_{0,1}=7.16$), we calculated by-subject overall error rates (that is, including words and pseudo-words) and removed all subjects whose error scores were higher than 20%. Words were also removed if their error rate was above 20% and 30%, respectively. Incorrect responses and fillers (word and pseudo-word) were excluded from analysis. Then, RTs were log-transformed to guarantee near-Gaussian distribution as suggested by Baayen (2008) and reduce the high between-subject variability. Individual log RTs were excluded from the analysis if they were more than 2.5 standard deviations away from the by-subject and overall log mean RT. Outlier rejection resulted in excluding a total of 542 points (3.64% of the dataset). A total of 14,339 datapoints were included in the analysis.

GROUP	CONDITION	MEAN RT		PRIMING	ES
		unrelated	*related*		
1	sgdom-sgdom	664	647	17	0.19
	pldom-sgdom	629	636	-7	-0.09
2	er	640	637	3	0.03
	rhyme	635	624	11	0.14
3	transparent	601	578	23	0.31
	opaque	608	596	12	0.16
	orthographic	592	594	-2	-0.03

Table 3.4.: Experiment 2 – dominance/affix priming. Summary table of the priming effects.

The means plotted in Figure 3.3 and Table 3.4 seem to suggest that, in group 1, the sgdom-sgdom condition elicited facilitation priming but the pldom-sgdom condition elicited inhibition priming; in group 2, the rhyme condition elicited priming but the *er*-condition did not; finally,

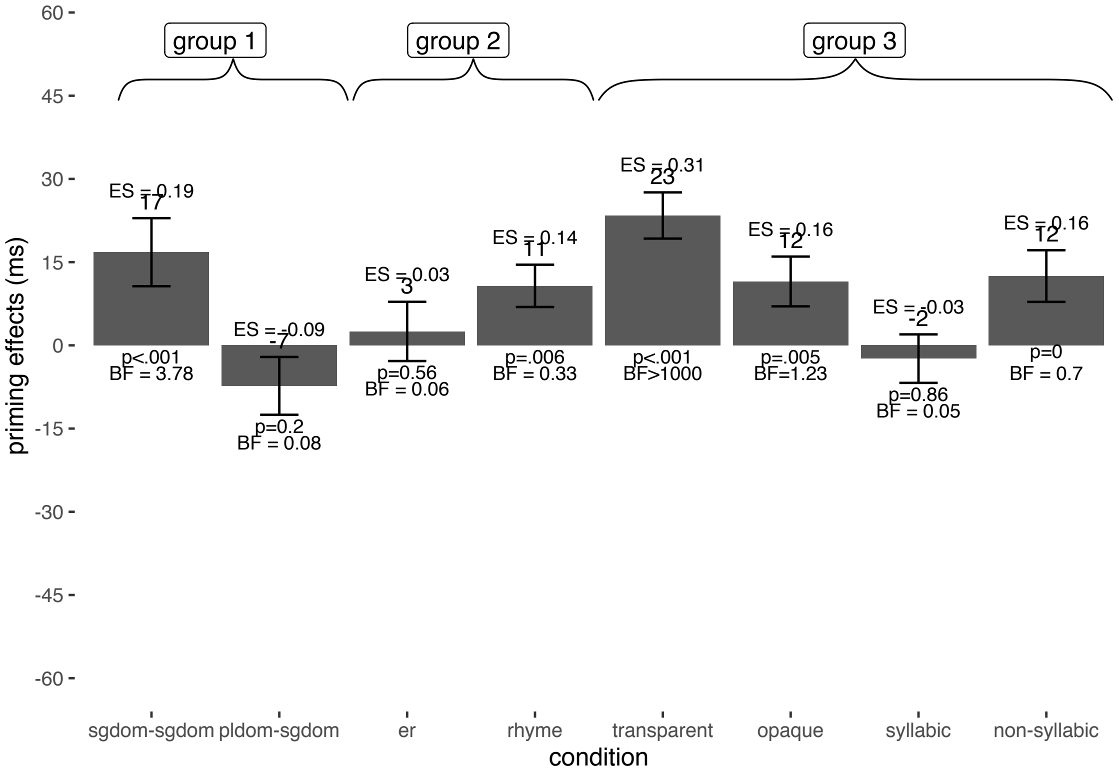

Figure 3.3.: Priming effects for experiment 2 – dominance/affix priming. Summary of the priming effects. The numbers over the bars are the priming magnitudes and Cohen's *d* (effect sizes, ES); the numbers below the bar are the *p*- and $BF_{1,0}$-values of the LMER and BF analyses, respectively.

in group 3, the transparent and opaque conditions elicited priming, but the orthographic condition did not.

We constructed a series of linear mixed-effect regression (LMER) models for each subdesign (that is, identity, transparent, opaque, syllabic, and non-syllabic; Baayen, 2008; Barr et al., 2013). Each model had *log RT* (in ms) as the dependent variable, RELATEDNESS (2 levels: related vs. unrelated) as the fixed factor, and SUBJECT and ITEM as random factors (intercept only). P-values were estimated using the Satterthwaite approximation of degrees of freedom (using the `lmerTest` R-package). For each subdesign, we also estimated Bayes Factors using the R package `BayesFactor` (Morey, 2018). For the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.5. The details of the LMER and the BF analyses are reported in Table 3.5 below.

GROUP	CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$ interpretation
1	sgdom-sgdom	12.09	<.001	3.78	substantial for H_1
	pldom-sgdom	1.63	.2	0.08	strong for H_0
2	er	0.34	.56	0.06	strong for H_0
	rhyme	10.72	.006	0.32	strong for H_0
3	transparent	36.96	<.001	10371.53	extreme for H_1
	opaque	10.04	.005	1.23	anecdotal for H_1
	orthographic	0.03	.86	0.05	strong for H_0

Table 3.5.: Experiment 2 – dominance/affix priming. Summary of the statistical results.

Group 1

The priming effects for the sgdom-sgdom condition were highly significant in the LMER analysis ($p < .0001$). The BF analyses confirmed that the data strongly supported a model with an effect of relatedness (that is, the alternative hypothesis; $BF_{1,0} > 10$). On the other hand, the priming effects for the pldom-sgdom condition were not significant in the LMER analysis ($p = .28$). The BF analyses revealed that the data strongly supported a model without an effect of relatedness (that is, the null hypothesis; $\text{BF}_{1,0} < 0.1$).

Group 2

The priming effects for the *er*-condition were not significant in the LMER analysis ($p = .51$). The BF analyses revealed that the data strongly supported the null hypothesis ($BF_{1,0} < 0.1$). The priming effects for the rhyme condition were found significant in the LMER analysis ($p = .003$). However, the BF analyses revealed that the data anectodally supported the null hypothesis ($1 > BF_{1,0} > 0.33$).

Group 3

The priming effects for the transparent- condition were highly significant in the LMER analysis ($p < .0001$). The BF analyses confirmed that the data extremely supported the alternative hypothesis ($BF_{1,0} > 100$). The priming effects for the opaque condition were found significant in the LMER analysis ($p = .002$). However, the BF analyses revealed that the data anectodally supported the alternative hypothesis ($3 > BF_{1,0} > 1$). The priming effects for the syllabic condition were not significant in the LMER analysis ($p = .88$). The BF analyses confirmed that the data strongly supported the null hypothesis ($BF_{1,0} < 0.1$). Finally, the priming effects for the non-syllabic condition were found significant in the LMER analysis ($p = .01$). However, the BF analyses revealed that the data substantially supported the null hypothesis ($BF_{1,0} < 0.33$).

Finally, pairwise comparisons of priming magnitudes were performed across conditions within each group. Table 3.6 below reports both Dunn-corrected *p*-values and uncorrected $BF_{1,0}$-values for all combinations within each group. Both analyses showed that the priming effects for the sgdom-sgdom condition and the pldom-sgdom condition, and the priming effects for the transparent condition and the syllabic condition, were significantly different from each other (ps= .003; $BF_{1,0}$s> 3). All other combinations supported the null hypothesis in both analyses (ps> .07; $BF_{1,0}$s< 1). We also compared the priming effects across groups. Most combinations were found non-significant (ps> .2, $BF_{1,0}$s< 1), and are not reported here. Here we only mention that the priming effects to the sgdom-sgdom condition and the transparent condition were not significantly different from one another ($p = 1$, $BF_{1,0} = 0.17$). The priming effects to the transparent condition and the pldom-sgdom condition were instead significantly different ($p = .02$, $BF_{1,0} = 328$).

(**a**) Group 1.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
sgdom-sgdom	*pldom-sgdom*	.003	8.75

(**b**) Group 2.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
rhyme	*er*	.08	0.64

(**c**) Group 3.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	.43	0.69
transparent	*syllabic*	.003	30.74
transparent	*non-syllabic*	.31	0.73
opaque	*syllabic*	.57	0.65
opaque	*non-syllabic*	1	0.15
syllabic	*non-syllabic*	.75	0.45

Table 3.6.: Experiment 2 – dominance/affix priming. Pairwise comparisons of priming effects in the three groups of conditions.

3.2.5. Discussion

In this chapter, we investigated whether dominance affects decomposition procedures. We defined *dominance* as the ratio between the frequency of the plural form (e.g., *windows*) and the frequency of the corresponding singular form (*window*; among others, Baayen et al., 1997, 2002). To this end, we elicited priming in response to singular- and plural-dominant plural prime words when both are followed by a singular-dominant plural target word. We found priming in the sgdom-sgdom condition, in which both primes and targets were singular-dominant plural forms (that is, the frequency of the stem form was higher than the frequency of the corresponding [z]-affixed form; e.g., *worlds-HEAVENS*). We instead found no priming in the pldom-sgdom condition, in which the primes were plural-dominant (that is, the frequency of the stem form was lower than the frequency of the corresponding [z]-affixed form; e.g., *windows-GODS*). Furthermore, we elicited priming in two additional groups of conditions, as a way to monitor the general effect size of suffix priming (group 2) and issues coming from the online data collection procedure (group 3). In these groups, the LMER and BF analyses were in conflict with one another; in interpreting these contradictory results, we chose to prefer the BF analysis over the LMER analysis, since it provides a clearer (and more straightforward) measure of the evidence for the alternative hypothesis over the null hypothesis without any additional inferential process (which is instead necessary in null-hypothesis testing approaches such as LMERs). In group 2, both the *er*-condition and the rhyme condition did not elicit priming. In group 3, the transparent condition (*boneless-BONE*) elicited significant priming effects, whereas the opaque (*belly-BELL*) and orthographic (*banjo-BAN*) conditions did not. In the next three subsections, we discuss in detail the theoretical implications of the results obtained in each group.

Group 1

In group 1, our results suggest that decomposition is affected by dominance. This is a novel result that has never been reported before in the masked priming literature. As already discussed above, only Forster (1999)'s bin model is able to account for these results. As we did in the previous chapter, we here provide a detailed description of where each model succeeds or fails in explaining the results reported above. For each failing model, modifications that would be needed to explain the effects are also proposed.

A. BIN MODEL

In the bin model, the entry-opening mechanism goes through orthographically similar candidates in a frequency-based orderly fashion. Therefore, when a singular-dominant plural prime (*worlds*) is presented, the decomposed entries {*world*}-{*s*} are opened earlier than the whole-form, less frequent entry {*worlds*}. When the singular-dominant plural target (*HEAVENS*) is presented, the decomposed entry {*s*} was already opened, which therefore speeds recognition. When a plural-dominant plural prime (*windows*) is presented, the decomposed entries {*window*}-{*s*} are opened later than the whole-form, more frequent entry {*windows*}. When the singular-dominant plural target (*GODS*) is presented, the decomposed entry {*s*} is already closed because it is suppressed by the more highly ranked whole-form unit; as a consequence, no priming arises.

While explaining the dominance-driven asymmetries reported above, the bin model is unable to account for the fact that priming arises in morphologically related pairs (e.g., sgdom-sgdom condition: *worlds-HEAVENS*; transparent condition: *boneless-BONE*) but does not arise in phonologically and orthographically related pairs (e.g., rhyme condition: *taper-VAPOR*; syllabic condition: *banjo-BAN*; non-syllabic condition: *starch-STAR*). Since in the bin model entries are opened depending on their orthographic similarity with the stimulus, orthographically related pairs are mistakenly expected to trigger priming effects, akin to morphologically related pairs.

B. RACE MODEL

In the race model, words are parsed through both the parsing and the storage route in parallel. Therefore, the prime word should decompose in the parsing route, while being searched through the lexicon in the storage route. Due to its short duration, the prime word is expected not to be fully recognized as a whole unit in the storage route; however, it is expected to decompose and therefore trigger priming onto the relative target, regardless of its dominance ratio. As such, the race model mistakenly predicts both conditions to trigger similar priming effects.

To explain the effects reported above, the race model needs to argue that the storage route activates the lexical entry of whole-form plural forms at early stages of processing, provided that their frequency is high enough. When a singular-dominant plural prime (*worlds*) is presented, it decomposes in the parsing route ({*world*}-{*s*}) and is searched through in the storage route. When the singular-dominant plural target (*HEAVENS*) is presented, it is primed because the prime word decomposes in the parsing route. When a plural-dominant plural prime (*windows*) is presented, it decomposes in the parsing route ({*window*}-{*s*}), though, at the same time, is recognized as a whole-word entry in the storage route ({*windows*}). When the singular-dominant plural target (*GODS*) is presented, the activated whole-word entry {*windows*} suppresses priming onto target recognition. Under this perspective, the dominance effects reported above result from activation of the whole-word lexical entry occurring within the first tens of milliseconds post stimulus onset. This implicitly suggests an incredibly fast process of lexical access, which has never been shown before and therefore unlikely to be true.

C. FULL-DECOMPOSITION MODEL

Similar to the race model, the full-decomposition model expects all plural forms to decompose at early stages of processing regardless of their dominance ratio. In the decomposition stage, the full-form string is expected to decompose on the basis of orthographic statistical regularities (islands of regularity; Rastle and Davis, 2008). In the lookup stage, the putative decomposed constituents are recognized; whole-word frequency (and therefore, dominance) is not expected to affect word processing until the constituents are put back together (recombination stage).

There seems to be no modification to the full-decomposition model able to capture the effects above without contradicting the main tenets of the model. For example, one way to explain these effects would be that in the lookup stage, high-frequency multimorphemic words are encoded as independent full-form nodes, while low-frequency multimorphemic words are not. However, this solution would refute the foundational argument of the model, whereby all words decompose and only morphemes (and not fully derived/inflected forms) are recognized in the lookup stage.

D. MORPHO-ORTHOGRAPHIC MODEL

Similarly to the race and the full-decomposition models, the morpho-orthographic model also expects dominance to not affect decomposition. At the morpho-orthographic segmentation, decomposition of the full-form string relies on Rastle and Davis (2008)'s islands of regularity. The decomposed constituents then activate the nodes in the orthographic lexicon. At this level, nodes only contains orthographic information associated to bare roots, and inflected and derived forms. Therefore, word frequency is not predicted to impinge on decomposition.

To explain the effects above, the morpho-orthographic model needs to assume that the dominance effects reported above are driven by the different resting activation levels of the nodes in the orthographic lexicon (Figure 3.4). In computational modeling, the *resting activation level* of a node is its energy level before activating, and usually correlates with the frequency of node activation: high-frequency nodes have a higher resting activation level and therefore activate faster than low-frequency nodes. Nodes compete with one another, so that the node that activates first inhibits the other competing nodes (*lateral inhibition*). Recall that the orthographic lexicon includes the orthographic information associated bare roots, and fully inflected and derived full forms (see sec. 1.4.2); in our case, the orthographic lexicon contains both singular and plural forms. When a singular-dominant plural form (*worlds*) is presented as a prime, it decomposes in the morpho-orthographic segmentation level (*worlds*→world-s), and its constituents activate the matching nodes in the orthographic lexicon. The orthographic node {*world*} activates faster than the competing orthographic whole-unit node {*worlds*} because the frequency of the former is higher than the frequency of the latter; thus, {*world*} inhibits {*worlds*}. Activation of the orthographic node {*world*} and the consequential inhibition of the whole-unit orthographic node {*worlds*} indirectly substantiate the segmentation pattern $ w o r l d - s $. When the singular-dominant target is presented (*HEAVENS*), the segment node $ s $ has already activated in the morpho-orthographic segmentation level, thus speeding target recognition and eliciting priming. When a plural-dominant plural form (*windows*) is presented as a prime, it decomposes in the morpho-orthographic segmentation level (*windows*→window-s), and its constituents activate the matching nodes in the orthographic lexicon. Here, the whole-unit node {*windows*}, however, activates faster than the root node {*window*} because the frequency of the former is higher than the frequency of the latter. When the singular-dominant target is presented (*GODS*), it goes through decomposition in the morpho-orthographic segmentation level (*gods*→god-s), but the segment node $ s $ has not activated yet; therefore, target recognition cannot benefit from any facilitation and priming does not arise. The mechanism argued for here is functionally the same as the one briefly sketched for the full-decomposition model above with a crucial difference. In the full-decomposition model, the mechanism occurs in the lookup stage and goes against the obligatory decomposition tenet of the model that the lookup stage contains morphemes and not full forms. In the morpho-orthographic model, the mechanism occurs in the orthographic lexicon instead, where only orthographic information (and, within the proposed modification above, frequency) is included; in this model, more abstract properties are stored in the higher lemma level. The division of labor between two separate levels of computation of the morpho-orthographic model (the orthographic lexicon and the lemma level) allows us to explain the dominance effects while not necessarily challenging the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3).

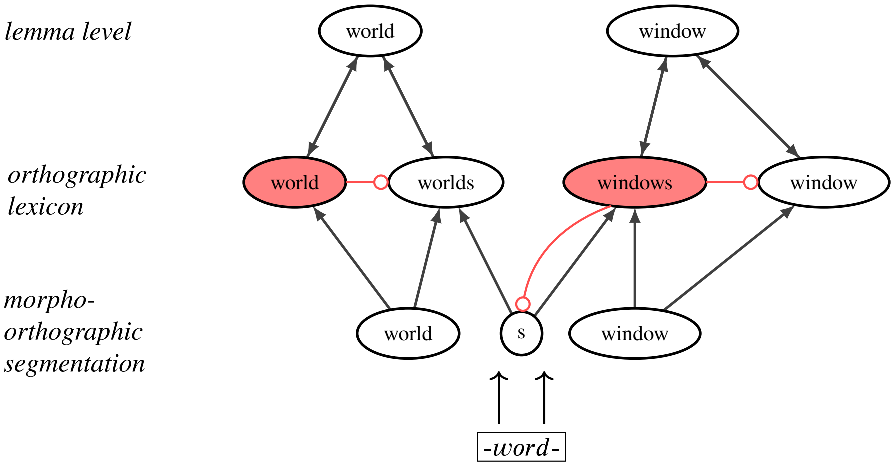

Figure 3.4.: Experiment 2 – dominance/affix priming. Graphic representation of the morpho-orthographic model (adapted from (Crepaldi et al., 2010) while explaining the results for experiment 2. Red nodes have a high resting activation and therefore takes less to activate (thus inhibiting competing nodes). The red circle-ending lines signal lateral inhibition (whereas arrows signal activation).

Out of the four models and modifications thereof, only the morpho-orthographic model seems able to successfully account for the results of experiment 2 without necessarily challenging the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). This is allowed exclusively thanks to the fact that the lexical and linguistic properties are *distributed* over two separate levels of computation – the orthographic lexicon and the lemma level (see sec. 1.4.2 for further details). On the other hand, we want to underline that similarly to the other models, the morpho-orthographic model completely glosses over the implementational problems that have arisen from the brothel and slegrack/flexire effects (see Chapter 2). We are therefore at a crossroads. The first possibility is that we maintain the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3) at the cost of leaving experiment 1 unaccounted for. The second possibility is that we argue against the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3), in favor of a model in which the whole-form unit (e.g., {*windows*}) and the decomposed-form unit ({*window*} - {*s*}) compete for activation with one another *within* the decomposition stage (rather than *across*) routes. In this dissertation, we explore this second possibility, thus allowing for decomposition to access whole-word frequency, in addition to whole-word lexicality as claimed in Chapter 2 (see Fig. 3.5). For the moment, we hold off on the actual mechanism accounting for these results. We will come back to this issue and provide the full details in Chapter 6.

In next chapters, we will turn to investigate the sensitivity of decomposition to other linguistic properties (i.e., phonological conditioning in alternations: ch. 4; syntactic restrictions to affixation: ch. 5). Besides, it is important to mention that the novelty of the dominance effects reported in these chapters opens a series of follow-up questions exploring the role of dominance asymmetries in morphological decomposition more in depth – we can mention at least two. The first question asks whether dominance impinges on root decomposition as much as it does on affix decomposition. Under the assumption that morphemes are decomposed and recognized regardless of their morphological status (i.e., regardless of being roots or affixes,

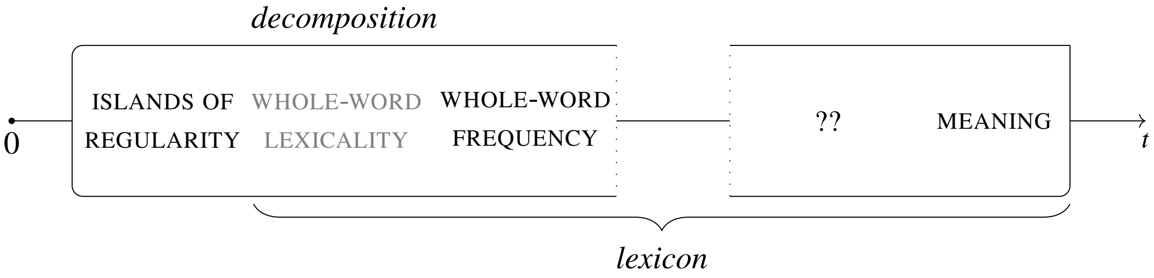

Figure 3.5.: Experiment 2 – dominance/affix priming. Time-course of the unfolding of word information in word processing (updated version after ch. 3). Whole-word lexicality is greyed-out because it was not directly addressed in this dissertation.

or free or bound morphemes), dominance is expected to equally impinge on decomposition of roots and affixes. For example, a plural-dominant plural prime (*windows*) should not prime the corresponding root (*WINDOW*) for the same reason that it does not prime a word with the same suffix (*HEAVENS*); conversely, a singular-dominant plural prime (*worlds*) should prime the corresponding root (*WORLD*) for the same reason that it primes a word with the same suffix (*GODS*). The second question asks whether dominance effects impinge on decomposition of other affixes as they do on decomposition of the plural suffix *-s*. We have begun to explore this question while dealing with the unexpected results obtained in group 2 (see below).

Group 2

In group 2, neither the *er*-condition nor the rhyme condition showed significant priming effects. On one hand, the lack of priming for the rhyme condition is expected. This condition is similar to the traditional orthographic condition (tested in Rastle et al., 2004; see also the orthographic condition in group 3 of this experiment: e.g., *banjo-BAN*); both consist of orthographically related pairs, with the phono-orthographic overlap being located in the rightward and leftward portions of the letter string in the rhyme and orthographic conditions, respectively. It is therefore unsurprising that priming did not occur in either condition. On the other hand, the lack of priming for the *er*-condition is quite surprising. This condition is similar to the traditional morphological transparent condition (tested in Rastle et al., 2004; see also the transparent condition in group 3 of this experiment: e.g., *boneless-BONE*); as they both consist of morphologically related pairs, they were expected to trigger similar priming effects. By way of explaining these results, we explored the possibility that dominance might have inhibited decomposition of the affix *-er*, inasmuch as it inhibited decomposition of the affix *-s*. This hypothesis extends the notion of number dominance (Baayen et al., 1997) to all other affixes and defines *affixal dominance* as the ratio between the surface frequency of the affixed form (e.g., *stalker*) and the surface frequency of the correspondent form after being stripped off of the relevant affix (*stalk*). Affixal dominance was not controlled for while constructing the materials, so suppression of *er*-priming might have been triggered by *er*-suffixed prime words (e.g., *runner*) having a lower frequency than the corresponding affixless target forms (e.g., *RUN*).

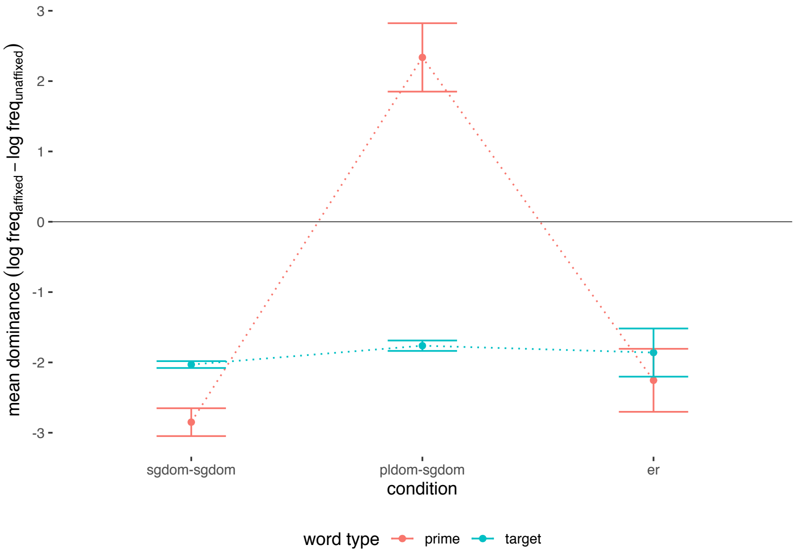

Figure 3.6.: Experiment 2 – dominance/affix priming. Plot of the mean dominance ratios in the words of the *er*-condition, as compared to the words of the sgdom-sgdom and pldom-sgdom conditions.

The plot in Figure 3.6 above shows that the dominance ratio of both primes and targets used in the *er*-condition was below 0, which indicates that the surface frequency of the affixless form (i.e., *run*) was actually higher than the surface frequency of the affixed form (*runner*). The distribution of the dominance ratio of the prime words varied across the sgdom-sgdom, pldom-sgdom, and *er*-conditions ($F(2,57)=50.72$, $p<.0001$; $BF_{1,0}>10000$). Post-hoc comparisons revealed that the dominance ratios of the primes were not different in the sgdom-sgdom and *er*-conditions (Dunn-corrected $p=.89$, uncorrected $BF_{0,1}=3$). These results suggest that the dominance effects that contributed to the asymmetrical priming effects in the sgdom-sgdom and pldom-sgdom conditions above could not have also contributed to the absence of priming in the *er*-condition. The lack of priming in the *er*-condition therefore raises an interesting follow-up question about whether all suffixes prime. The experiments reported in following chapters seem to confirm this trend (see in particular sec. 4.2.5).

Group 3

In group 3, we elicited the priming response to the conditions used in Rastle et al. (2004): i.e., the morphologically transparent condition (*boneless-BONE*), opaque condition (*belly-BELL*), and the orthographic condition (*banjo-BAN*). This was done primarily to test the reliability of the online platform *PsychoJS* (Peirce et al., 2019). The results were in line with the results obtained in the online re-run of experiment 1 (sec. 2.3). On one hand, the transparent conditions elicited priming effects and the orthographic condition did not, as reported in the literature (Rastle et al., 2004). On the other hand, the opaque condition did not elicit priming. This was indeed expected, given the technical issues encountered in online data collection: prime fluctuations, higher variance, and lower RT reliability (sec. 2.3.4). Replicability of the same results across multiple runs further validates online data collection through *PsychoJS* for priming experiments for large effect sizes (i.e., for morphologically transparent priming, but not for morphologically opaque priming; see sec. 2.4.2).

Chapter 4.

Morphological decomposition and *morpho-phonological alternations*

4.1. Introduction

The experiments reported in the two previous chapters suggest that decomposition has access to at least three linguistic properties: the orthographic realization of morphemes, lexicality, and dominance. This seems to contradict the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3), which expects decomposition to be exclusively sensitive to morpho-orthographic islands of regularity. It instead supports a model in which some properties are available at early stages of processing, thus potentially impinging on decomposition. In this chapter, we continue our investigation on the degree of linguistic sophistication of decomposition and turn to test whether decomposition is sensitive to phonologically-conditioned *morpho-phonological alternations*.

A morpho-phonological alternation involves two or more realizations of the same abstract morpheme alternating on the basis of specific morphological and/or phonological conditions. Several studies have looked at decomposition of *irregular morpho-phonological alternations* such as the ablauted forms of the English past tense (e.g., *fall*~*fell*; *think*~*thought*; *rise*~*rose*; Halle and Chomsky, 1968). These alternating forms are treated either as independent lexical entries (among others, Pinker, 1991) or as resulting from sub-regular morpho-phonological operations (among others, Embick, 2015). Crepaldi et al. (2010) report that irregular past tense forms prime the correspondent base form (e.g., *fell*→*fall*) similarly to regular past tense forms (e.g., *walked*→*walk*); the same results have also been replicated in a recent neuromagnetic study in Fruchter et al. (2013). These results seem to suggest that inflected forms decompose at early stages of processing, even when the morphological boundaries are not orthographically clear (e.g., *fe̲ll* cannot be linearly decomposed, as the realization of the past tense is the change of the root-internal vowel). In this chapter, we deal with *regular morpho-phonological alternations*, in which the alternants realize the same value of a given morpho-syntactic feature and can be accounted for by assuming language-wide, regular phonological operations triggered by the surrounding phonological context. In particular, we tested decomposition of the two alternants of the derivational prefix *in-*. The realization of this prefix alternates between [im] and [ɪn] as in (9), depending on whether the following root-initial phonological segment is a labial consonant or not.

(9) a. [ɪn]-*t*olerable; [ɪn]-*e*legant
b. [ɪm]-*p*ossible

These alternations are generally dealt with as resulting from general phonological processes (*rules* and/or *constraints*) operating on a single underlying form. In the example above, it is generally assumed that the underlying form of the prefix *in-* is [ɪn], which becomes [ɪm] via an assimilation process triggered in front of labial sounds (in our case, [p]).

Experiment 3 was designed to test whether decomposition is sensitive to the phonologically-based asymmetrical relationship between these two alternants. The experiment consisted of four main conditions. In the first two conditions, both prime and target words were prefixed with the same orthographic alternants of the prefix *in-*. In the *in*-condition, the prime and the target words were all prefixed with the alternant *in* (e.g., *inelegant-INTREPID*); in the *im*-condition, the prime and the target words were all prefixed with the alternant *im* (*immature-IMPURE*). The main purpose of thesese two conditions was to elicit priming in pairs of words sharing the same orthographic form of the prefix *in*; this was necessary in light of the results from experiment 2, which suggest that some affixes (e.g., *-er*) may not decompose in the first place (sec. 3.2.5).[9] In the remaining two conditions, the prime and the target words were prefixed with either of the two orthographic alternations of the prefix *in-*. In the *in-im* condition, the prime and the target words were *in*-prefixed and *im*-prefixed, respectively (e.g., *inelegant-IMPURE*); finally, in the *im-in* condition, the prime and the target words were *im*-prefixed and *in*-prefixed, respectively (e.g., *immature-INTREPID*). These two conditions were designed to explore two logically separate questions. The first question asked whether decomposition is able to detect the morpho-phonological *relationship* between alternating realizations of the same morpheme. If this were the case, we would expect priming effects to arise in at least one of the last two conditions just described. The second question asked whether decomposition is able to detect the morpho-phonological *asymmetry* between phonologically-conditioned alternants, whereby one alternant (i.e., *im*) phonologically derives from the other (i.e., *in*). If this were the case, we would expect priming effects to be modulated by manipulating the alternating form being presented as a prime.

For the experiment reported in this chapter, the four models considered in this dissertation make the following predictions (see sec. 4.2.3 for further details). The bin model predicts priming to arise across all conditions, as they all share a a similar number of letters. As we have already seen in the previous chapters, the priming response results from a low-level, orthographic-based comparison between the presented stimulus and the form of the lexical entries; thus this model does not provide a mechanism that specifically teases apart facilitation effects due to orthographic relatedness and facilitation effects due to morphological relatedness. The remaining three models – namely, the race, full-decomposition, and morpho-orthographic

---

[9] The two conditions were not controlled for affixal dominance as defined in sec. 3.2.5 because of the limited number of words available in the English lexicon. However, we controlled for potential dominance-driven effects in the post-hoc analysis described in sec. 4.2.5.

models – consistently predict priming to arise in the *in* and *im* conditions, but not in the *in-im* and *im-in* conditions. In these models, decomposition is indeed argued to rely on the morpho-orthographic information associated to each alternant and should therefore be unable to access the phonologically-driven relatedness among alternating forms of the same morpheme. If priming were to arise in all of the four conditions above, it would argue for a model in which decomposition may be affected by the phonological relationship among alternants of the same morpheme as described above. Our results showed that only the *im*-condition triggered significant priming effects, while all other conditions (i.e., the *in*, *in-im*, and the *im-in* conditions) did not. These results seem to suggest that the form *im* decomposes, but the form *in* does not. In sec. 4.2.5, we discuss these surprising results in exploring the factors that might have been at play and the implications for theories of decomposition.

4.2. Experiment 3 – *in-im*/affix priming

4.2.1. Materials

Seventy-two pairs were selected from the English Lexicon Project corpus (ELP: Balota et al., 2007, which refers to the HAL corpus: Lund and Burgess, 1996) and distributed over six conditions.

1. Group 1.
   a) In the [in]-*condition*, the prime and the target words share the alternant [in] (e.g., *inelegant-INTREPID*).
   b) In the [im]-*condition*, the prime and the target words share the alternant [im] (e.g., *immature-IMPURE*).
   c) In the [in]-[im] *condition*, the prime word takes the alternant [in] and the target word takes the alternant [im] (*inelegant-IMPURE*).
   d) In the [im]-[in] *condition*, the prime word takes the alternant [im] and the target word takes the alternant [in] (*immature-INTREPID*).
   e) The dis-*condition*, primes and targets share the *dis*- prefix (e.g., *disembark-DISABLE*). This condition had a least two purposes. First, together with the the *in* and *im* conditions, the *dis*-condition served the general purpose to control for the magnitude of the priming effects triggered by the repeated presentation of the same morphological alternant within the same condition. Second, it was designed to elicit priming effects for a prefix that was different from *in*- and does not show any relevant phono-orthographic alternation; this condition, finally, could also contribute to the side investigation on decomposition of single affixes we are undertaking in parallel with our main goals (see sec. 3.2.5).

f) In the *cohort condition*, primes and targets share the same beginning (e.g., *banter-BANJO*). This condition serves as baseline for effects potentially arising in response to mere phonetic overlap, as a way to distinguish them from effects arising in response to morphological relatedness.

We were forced to have a fewer number of items per condition than previous experiments because of the limited number of *in-/im-*prefixed words. For this reason, in the conditions (a–d) above, we could not control for a number of possible confounding factors as much as we wished. First, the root of some words were bound (e.g., *clement* in the word *inclement*). Second, the words of the same pair were always morphologically related, but often were semantically vaguely related with one another (e.g., *immediate* and *mediate* are only vaguely related from the semantic point of view). These factors were overlooked because previous results suggest neither to matter at all (at least in the visual modality; Rastle et al., 2004; Stockall and Marantz, 2006; Solomyak and Marantz, 2010). The shortage of relevant items to choose from made the word selection difficult also because some of the words used are not well known; the concern would be that subjects may not know the word at all and take more time to make the decision. To prevent this from happening, we set up the pairs in such a way that the lowest frequency words were presented as primes rather than as targets. For the very same reason, the words could also not be controlled for affixal dominance (as defined in sec. 3.2.5); potential dominance-driven effects are further investigated in sec. 4.2.5 as post-hoc analysis.

All words (primes and targets, separately) were matched as much as possible in frequency, morphological length, syllabic length, and orthographic length. The mean values across the conditions and the relative statistical results are shown in Table 4.1 below.

In addition to the six experimental conditions above, we also prepared a second group of conditions. Similarly to group 3 in experiment 2, group 2 in this experiment ensured that the online environment in which the experiment was run did not affect the priming response in any way.

2. Group 2.

g) In the *transparent condition*, pairs carried a semantically and morphologically transparent relationship (e.g., *silky-SILK*).

h) In the *opaque condition*, pairs carried an apparent morphological relationship, but not a semantic one (e.g., *ponder-POND*).

i) In the [-M+S+O]-*condition*, pairs were semantically (+S) and orthographically related (+O), but morphologically unrelated (-M; e.g. *screech-SCREAM*; see sec. 1.3). These pairs were directly taken from Rastle et al. (2000).

The two groups of conditions were not matched for frequency, morphological length, or syllabic length with one another. This was because the two groups had inherently different properties that made the match impossible. Conditions in group 2 were controlled for frequency only (see Table 4.2).

PROPERTY		CONDITIONS						STATISTICAL ANALYSES
		in	*im*	*in-im*	*im-in*	*dis*	*cohort*	
PRIMES	HAL frequency (log10)	4.49	5.10	4.79	4.20	4.6	4.69	$F(5,66)=.59, p=.7$ $BF_{0,1}=9.57$
	morphological length	2.92	2.92	3	2.67	2.42	2.58	$F(5,66)=1.87, p=.11$ $BF_{0,1}=1.58$
	syllabic length	3.75	3.92	3.83	3.83	3.17	3.42	$F(5,66)=1.97, p=.1$ $BF_{0,1}=1.38$
	orthographic length	10	10	9.75	9.75	9.5	9.75	$F(5,66)=0.17, p=.97$ $BF_{0,1}=17.04$
TARGETS	HAL frequency (log10)	6.48	6.88	6.07	7.17	6.94	6.61	$F(5,66)=.82, p=.53$ $BF_{0,1}=6.87$
	morphological length	2.67	2.58	2.67	2.67	2.25	2.33	$F(5,66)=1.16, p=.33$ $BF_{0,1}=4.27$
	syllabic length	3.33	3.33	3.42	3.92	3.42	3.17	$F(5,66)=1.52, p=.19$ $BF_{0,1}=2.58$
	orthographic length	9	8.83	9.42	9.75	9.92	9.67	$F(5,66)=1.21, p=.31$ $BF_{0,1}=4.02$

Table 4.1.: Experiment 3 – *in-im*/affix priming (group 1). Summary of the lexical properties of the stimuli used.

log10 frequency (log10)	CONDITIONS			STATISTICAL ANALYSES
	transparent	*opaque*	*[-M+S+O]*	
PRIMES	6.18	6.32	6.56	$F(2,33)=.22, p=.8$ $BF_{0,1}=4.54$
TARGETS	7.54	8.4	7.52	$F(2,33)=2.22, p=.12$ $BF_{0,1}=1.23$

Table 4.2.: Experiment 3 – *in-im*/affix priming (group 2). Summary of the lexical properties of the stimuli used.

One hundred and eight unrelated prime words were selected, one for each target word. These words were phonetically, morphologically, and semantically unrelated to the targets. They were matched as closely as possible to the correspondent related prime on frequency ($t(107)=-.02$, $p=.97$; $BF_{0,1}=9.37$), morphological length ($t(107)=-.62$, $p=.53$; $BF_{0,1}=7.76$),

(a) Group 1.

PROPERTY	CONDITIONS		
	in	*im*	*in-im*
HAL freq. (log10)	$t(11)=0.77, p=.45$ $BF_{0,1}=2.69$	$t(11)=-1.17, p=.26$ $BF_{0,1}=1.97$	$t(11)=-0.24, p=.81$ $BF_{0,1}=3.39$
morph. length	$t(11)=0, p=1$ $BF_{0,1}=3.48$	$t(11)=-0.36, p=.72$ $BF_{0,1}=3.29$	$t(11)=0, p=1$ $BF_{0,1}=3.48$
syll. length	$t(11)=0, p=1$ $BF_{0,1}=3.48$	$t(11)=-0.32, p=.75$ $BF_{0,1}=3.33$	$t(11)=0, p=1$ $BF_{0,1}=3.48$

PROPERTY	CONDITIONS		
	im-in	*dis*	*cohort*
HAL freq. (log10)	$t(11)=0.07, p=.94$ $BF_{0,1}=3.47$	$t(11)=0.01, p=.99$ $BF_{0,1}=3.48$	$t(11)=-0.01, p=.98$ $BF_{0,1}=3.48$
morph. length	$t(11)=0, p=1$ $BF_{0,1}=3.48$	$t(11)=0, p=1$ $BF_{0,1}=3.48$	$t(11)=0, p=1$ $BF_{0,1}=3.48$
syll. length	$t(11)=0, p=1$ $BF_{0,1}=3.48$	$t(11)=0, p=1$ $BF_{0,1}=3.48$	$t(11)=0, p=1$ $BF_{0,1}=3.48$

(b) Group 2.

PROPERTY	CONDITIONS		
	transparent	*opaque*	*[-M+S+O]*
HAL freq. (log10)	$t(11)=0.52, p=.61$ $BF_{0,1}=3.10$	$t(11)=-0.08, p=.93$ $BF_{0,1}=3.47$	$t(11)=0.001, p=.99$ $BF_{0,1}=3.48$

Table 4.3.: Experiment 3 – *in-im*/affix priming. Within-condition statistical results of the lexical properties of related and unrelated primes across the two groups of conditions.

and syllabic length ($t(107)=-.18$, $p=.85$; $BF_{0,1}=9.22$). By-condition statistical results are reported below in Table 4.3. The stimuli used in the experiment are available in Appendix III.

An additional set of one hundred and eight pseudo-word targets were chosen; they were preceded by unrelated suffixed word primes (not used in the experimental conditions).

Target words from each condition were counterbalanced and divided at random into two versions of equal numbers of pairs. In each version, half of the target words were preceded by the unrelated prime and half by the related prime. Participants received only one version of the experiment, so that they saw each target word exactly once. Each participant saw a total of 216 word pairs.

4.2.2. Participants & procedure

One hundred and fifty-seven participants (76 females, 81 males; mean age: 35.16, s.d: 9.97) were recruited through Amazon Mechanical Turk and received monetary compensation for their participation. They were all native speakers of American English.

Stimulus presentation and data recording were performed on-line through PsychoJS (Peirce et al., 2019), the Javascript equivalent of PsychoPy, on the Pavlovia platform (`www.pavlovia.org`). They were asked to read the capitalized letter strings on the display and decide as quickly and accurately as possible whether or not each string is a word. To do that, they used their own monitor and keyboard. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case for 33 ms; the target word immediately followed in uppercase, and remained on the screen until a response was made. The order of pairs was chosen randomly across participants.

Participants were given 10 practice pairs before the actual experiment began. A total of 216 pairs were presented to each participant. During the experiment, participants were also given the possibility to take 6 brief breaks. To detect bots, subjects were also asked to answer three open-ended questions immediately after a break. To help subjects refocus on the main task post-break, we also made sure that first five trials presented after each break were pseudo-word trials.

4.2.3. Predictions

The four models considered in this dissertation makes different predictions for the conditions tested in this experiment. Table 4.4 summarizes the predictions for each model.

The bin model subsumes decomposition under the orthographically-based comparison mechanism within the entry-opening stage (see sec. 1.4.2). As a consequence, any word prime is expected to trigger priming onto an orthographically similar target, regardless of whether they are morphologically related (as in the *in*-, *im*, and *dis*-, transparent and the opaque conditions) or not (as in the cohort and $[-M+S+O]$-conditions). Therefore for the bin model predicts all conditions tested to show priming effects.

MODELS	GROUP 1						GROUP 2		
	in	*im*	*in-im*	*im-in*	*dis*	cohort	transparent	opaque	$[-M+S+O]$
bin	✓	✓	✓	✓	✓	✓	✓	✓	✓
race	✓	✓	✗	✗	✓	✗	✓	✓	✗
full-decomposition	✓	✓	✗	✗	✓	✗	✓	✓	✗
morpho-orthographic	✓	✓	✗	✗	✓	✗	✓	✓	✗

Table 4.4.: Experiment 3 – *in-im*/affix priming. Predicted results across models. (Legend. ✓: priming effects; ✗: no priming effects.)

The race and the full-decomposition models assume a procedure of morphological decomposition based on morpho-orthographic islands of regularity (Rastle and Davis, 2008) and in compliance with the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3)(see sec. 1.1). As a consequence, priming is predicted to arise in the following conditions: *in*, *im*, and *dis* (group 1); transparent and opaque (group 2). In these conditions, the pairs involve words that are morphologically related, sharing either the same root (transparent, opaque) or the same affix (*in*, *im*, *dis*). No priming is instead predicted to arise in the remaining conditions (Figure 4.2, scenario 1). The cohort condition and the $[-M + S + O]$-condition involve pairs that are orthographically, but not morphologically related. As morphological decomposition is assumed to rely on the morpho-orthographic information, it should not occur and priming should not arise in these conditions either. The *in-im* and *in-im* conditions involve words that share the same abstract morpheme (i.e., the affix *in-*), but differ in their orthographic form. Since decomposition is argued to not possibly detect any deeper information of a morpheme other than its form (i.e., the abstract phonological relationship between alternants of the same morpheme), decomposition should not occur in these conditions and priming should not arise.

Finally, the morpho-orthographic model hinges on the same assumptions as the race model (B) and the full-decomposition model (C) above. As such, the morpho-orthographic model makes similar predictions (Figure 4.2, scenario 1). The *in*, *im*, *dis*, transparent, and opaque conditions are expected to prime because they all involve pairs sharing either the same root (transparent, opaque) or the same affix (*in*, *im*, *dis*). Instead, no priming is predicted to arise in the remaining conditions. The cohort condition and the $[-M+S+O]$-condition involve pairs that are orthographically, but not morphologically related. As morphological decomposition is assumed to rely on the morpho-orthographic information, it should not occur, and priming should not arise in these conditions either. The morpho-orthographic model provides a more mechanistically sophisticated motivation as to why the *in-im* and *in-im* conditions should not prime. In the morpho-orthographic model, derived words have independent nodes in the orthographic lexicon. Therefore, the two alternants *in*~*im* are both expected to decompose, but eventually activate the whole-word nodes in the orthographic lexicon; activation of independent whole-word nodes in the orthographic lexicon prevents priming from arising (see Figure 4.1).

If priming effects were to be found in both/either the *in-im* condition and/or the *im-in* condition, they would argue against all of the models of decomposition considered here, in favor of a

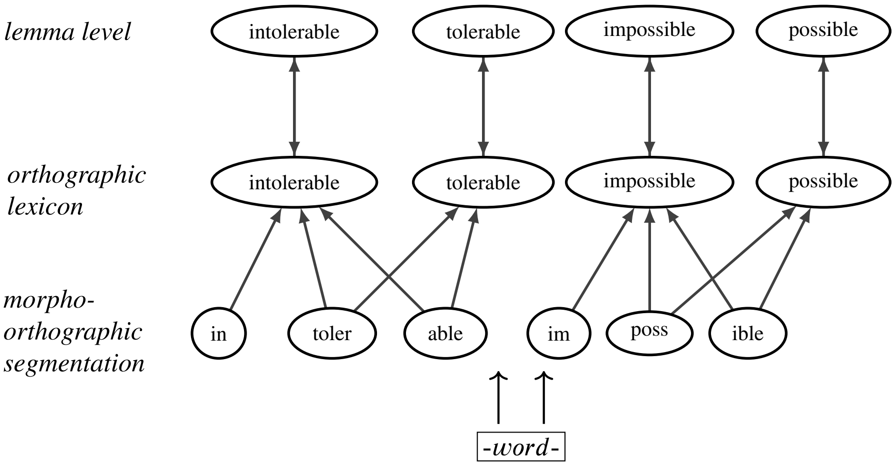

Figure 4.1.: Experiment 3 – *in-im*/affix priming. Graphic representation of the morpho-orthographic model (adapted from (Crepaldi et al., 2010) while predicting the results of the conditions experiment 3.

model in which decomposition is sensitive to the morpho-*phonological* relationship between alternants of the same morpheme, and not just to morpho-*orthographic* forms. If priming effects were the same in both conditions (Figure 4.2; scenario 2), it would suggest that phonologically-conditioned morphological alternations are recognized as different realizations of the same morpheme, with no specific asymmetrical relationship between them. If priming effects were to be found in only one of the two conditions, we may suggest that morpho-phonological alternations of the same morpheme are instead asymmetrically organized. In an asymmetrical organization of lexical entries, alternants that are underspecified (i.e, lacking some phonological and/or morphological information) are in a subset relation with specified alternants; in our example, the segment [n] of the alternant *in* is usually considered to be underspecified for place (since it is assumed to be the elsewhere condition), whereas the segment [m] of the alternant *im* is usually considered to be specified for place as [+labial]. In this sense, either alternant being presented as a prime may or may not facilitate subsequent recognition of the other alternant being presented as the related target, depending on their respective featural specifications. Priming effects in the *in-im* condition (but not in the *im-in* condition) suggest that recognition of the specified alternant *im* is facilitated by prior presentation of the underspecified alternant *in* (scenario 3). Priming effects in the *im-in* condition (but not in the *in-im* condition) instead suggest that recognition of the underspecified alternant *in* is facilitated by prior presentation of the specified alternant *im* (scenario 4).

4.2.4. Results

In compliance with the pipeline adopted for online experiments, a preliminary trimming was performed to remove trials in which the duration of the prime was longer or shorter than 34±8 ms (see 2.3 for further details). In this experiment, 7,781 trials (18.56% of the dataset) were removed because of this. The usual pipeline was followed. First, we calculated by-subject error rates for words and pseudo-words separately. Since the means of the two distributions

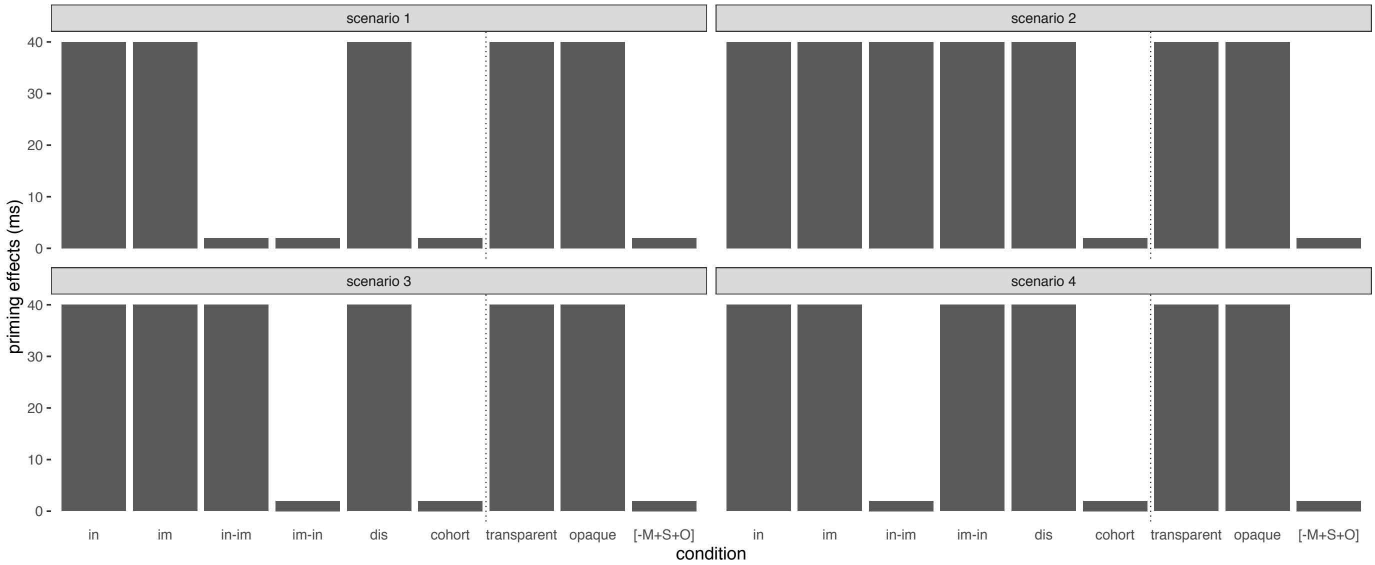

Figure 4.2.: Experiment 3 – *in-im*/affix priming. Predicted results

did not vary significantly ($t(221.2)=0.47, p=.63$; $BF_{0,1}=7.10$), we calculated by-subject overall error rates (that is, including words and pseudo-words) and removed all subjects whose error scores were higher than 20%. Words were also removed if their error rate was above 20% and 30%, respectively. Incorrect responses and fillers (word and pseudo-word) were excluded from analysis. Then, RTs were log-transformed to guarantee near-Gaussian distribution as suggested by (Baayen, 2008) and reduce the high between-subject variability. Individual log RTs were excluded from the analysis if they were more than 2.5 standard deviations away from the by-subject and overall log mean RT. Outlier rejection resulted in excluding a total of 319 points (3.81% of the dataset). A total of 8,053 datapoints were included in the analysis.

The means plotted in Figure 4.3 and reported in Table 4.5 suggest that in group 1, only the priming megnitude of the *im*-condition trend positive. The priming magnitude of the *in*-condition seem to be around 0, and the priming magnitudes of the *dis*-, *in-im* and *im-in* conditions trend instead negative. In group 2, the priming effects of all conditions trend positive, with the effects of the transparent condition being larger than the effects of the opaque condition and $[-M+S+O]$-condition.

We then constructed a series of linear mixed-effect regression (LMER) models for each sub-design (that is, identity, transparent, opaque, syllabic, and non-syllabic; Baayen, 2008; Barr et al., 2013). Each model had *log RT* as the dependent variable, RELATEDNESS (2 levels: related vs. unrelated) as the fixed factor, and SUBJECT and ITEM as random factors (intercept only). P-values were estimated using the Satterthwaite approximation of degrees of freedom (using the `lmerTest` R-package). For each subdesign, we also estimated Bayes Factors using the R package `BayesFactor` (Morey, 2018). For the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.5. The details of the LMER and the BF analyses are reported in Table 4.6 below.

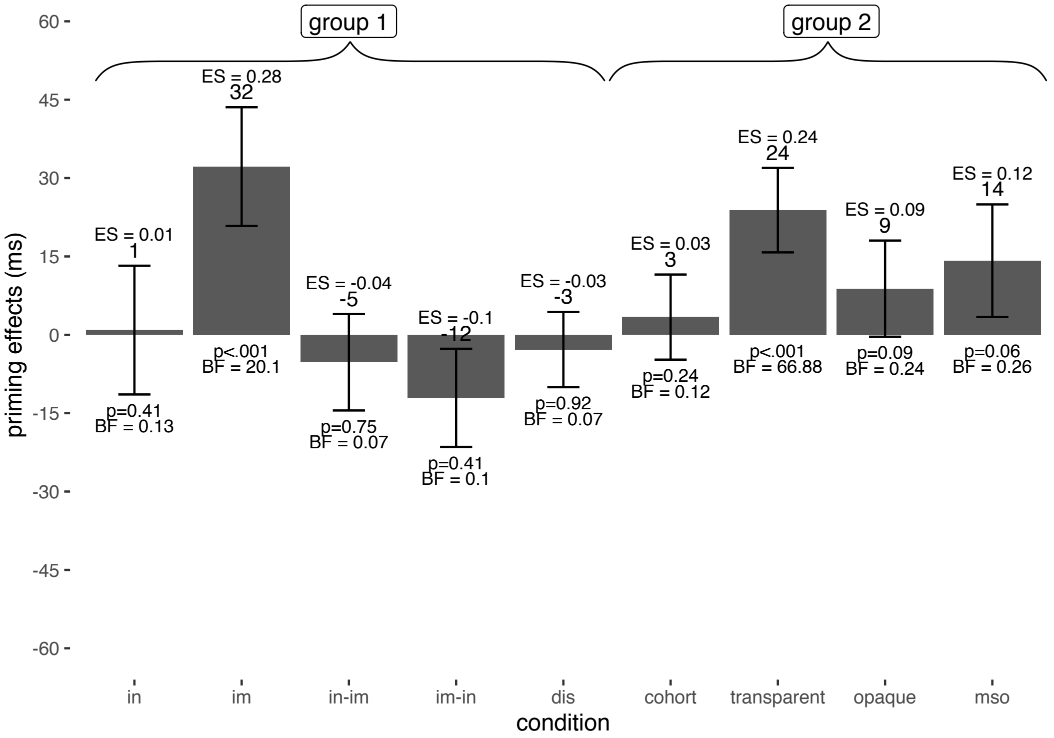

Figure 4.3.: Experiment 3 – *in-im*/affix priming. Summary of the priming effects. The numbers over the bars are the priming magnitudes and Cohen's *d* (effect sizes, ES); the numbers below the bar are the *p*- and $BF_{1,0}$-values of the LMER and BF analyses, respectively.

Group 1

In the LMER analysis, the *in-*, *in-im*, *im-in*, *dis-*, and cohort conditions were found to be non-significant ($p \geq .4$). In the BF analysis, the *in-im*, *im-in*, and *dis*-conditions were also found to strongly support the null hypothesis ($BF_{1,0} \leq 0.1$); the *in-* and cohort conditions were found to substantially support the null hypothesis ($BF_{1,0} < 0.33$). Finally, the *im*-condition was found to obtain significant priming effects in both analyses (p=.003; $BF_{1,0}$=20.10).

Group 2

Only the transparent condition triggered significant priming effects in both analyses (p<.0001; $BF_{1,0}$=66.88). Both the opaque and $[-M+S+O]$ conditions did not; in particular, they were found to substantially support the null hypothesis in the BF analysis ($BF_{1,0} < 0.33$).

Pairwise comparisons

Pairwise comparisons of priming magnitudes were also performed across conditions. Table 4.7 below reports both Dunn-corrected *p*-values and uncorrected $BF_{1,0}$-values for each combina-

(**a**) Group 1.

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
in	710	709	1	0.01
im	690	658	32	0.28
in-im	703	708	-5	-0.04
im-in	681	693	-12	-0.1
dis	678	681	-3	-0.03
cohort	693	690	3	0.03

(**b**) Group 2.

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
transparent	680	656	24	0.24
opaque	661	652	9	0.09
$[-M+S+O]$	676	662	13	0.12

Table 4.5.: Experiment 3 – *in-im*/affix priming. Mean RT, priming magnitudes and effects sizes (ES).

CONDITION		F	p	$BF_{1,0}$	$BF_{1,0}$interpretation
GROUP 1	*in*	0.67	.41	0.13	substantial for H_0
	im	13.85	.0002	20.10	strong for H_1
	in-im	0.10	.75	0.07	strong for H_0
	im-in	0.68	.41	0.10	substantial for H_0
	dis	0.01	.92	0.07	strong for H_0
	cohort	1.39	.23	0.12	substantial for H_0
GROUP 2	transparent	21.53	<.0001	66.88	strong for H_1
	opaque	2.95	.09	0.24	substantial for H_0
	$[-M+S+O]$	3.52	.06	0.26	substantial for H_0

Table 4.6.: Experiment 3 – *in-im*/affix priming. Summary of the statistical results.

tion. The two analyses were consistent with one another in suggesting that most combinations supported the null hypothesis (ps>.15; $BF_{1,0}$s<1). The only exception was the combination *im* vs. *im-in*, which instead seems to support the alternative hypothesis ($p<.05$; $BF_{1,0}>3$).

(**a**) Group 1.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
in	*im*	.17	0.18
in	*in-im*	1	0.43
in	*im-in*	1	0.42
in	*dis*	1	0.34
in	*cohort*	1	0.16
im	*in-im*	1	0.43
im	*im-in*	.04	3.49
im	*dis*	.32	0.76
im	*cohort*	1	0.23
in-im	*im-in*	1	0.17
in-im	*dis*	1	0.15
in-im	*cohort*	1	0.19
im-in	*dis*	1	0.21
im-im	*cohort*	1	0.32
dis	*cohort*	1	0.25

(**b**) Group 2.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
transparent	*opaque*	1	0.16
transparent	$[-M+S+O]$	1	0.20
opaque	$[-M+S+O]$	1	0.15

Table 4.7.: Experiment 3 – *in-im*/affix priming. Pairwise comparisons of priming effects in the two groups of conditions.

4.2.5. Discussion

Experiment 3 looked at the decomposition of the morpho-phonological orthographic alternants of the English prefix *in-*, *in*∼*im*. To this end, a total of nine conditions (split in two separate groups) were tested. Group 1 consisted of five morphologically related conditions and one orthographically-related condition. In the morphologically related conditions, prime and target words either shared the same form of the affix (*in-*, *im-*, and *dis*-conditions) or combined the two phonologically-conditioned alternants of the affix *in-* (*in-im* and *im-in* conditions). In the orthographic-related condition (cohort condition), the prime and the target words shared the same beginning and it was used as baseline for orthographic effects. Group 2 consisted of Rastle et al. (2000, 2004)'s traditional conditions (transparent, opaque, and $[-M + S + O]$) and ensured that the online environment did not affect the response. The results showed that only the *im*-condition (group 1) and the transparent condition (group 2) primed, but remaining conditions (group 1: the *in*, *in-im*, *im-in*, *dis-*, and cohort conditions; group 2: opaque and $[-M + S + O]$ conditions) did not. These results challenge all the models considered here (see also sec. 4.2.3). In the bin model (sec. 1.4.2A), decomposition occurs as part of a orthographically-based comparison mechanism within the entry-opening stage. As a consequence, all of the conditions tested should prime because all word pairs therein are orthographically related – including the cohort, *in-*, *dis-*, *in-im*, and *im-in* conditions. In the race, full-decomposition, and morpho-orthographic models (sec. 1.4.2B-D), morphological decomposition operates on morpho-orthographic islands of regularity (Rastle and Davis, 2008) and in compliance with the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). As a consequence, priming is predicted to arise at least in the conditions where primes and targets involve the same affix, namely *in-*, *im-*, and *dis*-conditions.

This surprising pattern of results prevents us from properly interpreting the effects for the *in-im* and *im-in* conditions as sketched in sec. 4.2.3 at this moment, and validates the need for a replication of experiment 3 in the near future. It is worth highlighting that the differential response for the *in-* and *im*-condition is particularly surprising because the forms *in-* and *im-* are similar to one another at a number of levels of analysis. At the orthographic level, they share everything but a single grapheme: $\$\,\mathrm{n}\,\$$ and $\$\,\mathrm{m}\,\$$, which in turn share several orthographic features too. At the morpho-phonological level, they are context-dependent realizations of the same morpheme (see sec.4.1). At the semantic level, both *in* and *im* negate the stem they attach to (e.g., *intolerable*, which refers to something that is *not* tolerable; *impossible*, which refers to something that is *not* possible). For this reason, the remainder of the section discusses the null priming effects of the three basic morphological conditions of the experiment – namely, *in-, im-*, and *dis*-conditions. We tackle this in the following two subsections. The first subsection explores potential confounding factors that might have affected the priming response in the three conditions. As none of these factors were found to have impinged on the priming response, the second subsection discusses the theoretical consequences of this experiment for the models considered here and, more generally, for the segmentation procedure underlying morphological decomposition.

Potential confounding factors

We identified two confounding factors that might have affected the results reported above. First, the results could have been due to the fact that decomposition procedures (i) might not have been performed at all or (ii) might have been performed by a subset of our subject sample. As for (i), the pattern of results of group 2 suggests that at least some decomposition procedure was indeed performed during the experiment. In group 2, the transparent condition primed and the $[-M + S + O]$ condition did not, in line with the results widely reported in the visual morphological priming literature (for a review, see sec. 1.3; Rastle et al., 2000). (The opaque condition not priming was also expected, and was likely due to the technical issues encountered in online data collection, which prevented smaller effect sizes from arising – see in particular sec. 2.4.2.) We therefore looked into (ii), i.e. in the possibility that the sample recruited was not uniform in performing decomposition procedures. More concretely, we explored the hypothesis that the subjects recruited for the experiment might have belonged to two distinct populations - one that decomposes and one that does not. If so, averaging together the RT response of the two populations might have brought about the null results reported above.

To this end, we checked the modality of the priming distribution across conditions. The modality of a distribution is determined by the number of peaks (or modes), namely the value (or values) that occurs more frequently in the data. Unimodal distributions have only one mode and suggest that the data were from a single population; multimodal distributions have more than one mode and suggest that the were drawn from more than one population. As Figure 4.4 shows, the density distribution of priming effects is unimodal in all conditions but the *im*-condition, which shows instead a bump in the right tail. Similar to a non-unimodal distribution, fat-tailed distributions may suggest that two different populations were recruited for the experiment. We therefore performed two different test-statistics to assess whether the significant priming effects in the *im*-condition were due to distribution modality and/or a fat tail. To check whether the priming distribution was non-unimodal across conditions, we performed Hartigans' *dip test*, which measures non-unimodality in a sample by calculating the maximal difference D between the empirical distribution function and the unimodal distribution function (Hartigan and Hartigan, 1985); the closer D is to 1 the more likely it is that the distribution is multimodal. Therefore, D is expected to approximate to 1 for all non-unimodal distributions. We fed the by-subject priming estimates into the dip test analysis by using the `diptest` R-package. The p-values were calculated from a Monte Carlo simulation of a uniform distribution with 10,000 replicates, and measured the probability that the distribution of the data was unimodal (the null hypothesis). The priming distribution of all conditions were found to be unimodal (*in*: $D = 0.031, p = .93$; *im*: $D = 0.027, p = .97$; *in-im*: $D = 0.016, p = 1$; *im-in*: $D = 0.025, p = .97$; *dis*: $D = 0.038, p = .42$; *cohort*: $D = 0.029, p = .86$). This suggests that all subjects recruited for the experiment indeed belonged to the same population. Furthermore, we explored the extent to which the significant effects in the *im*-condition were due to the positive fat-tailed priming distribution. To test this, we cut off the tails of the RT distribution of the *im*-condition by excluding data points that were 1.5 standard deviations away from the overall log mean RT (75 data points, 10.6% of the dataset). If the significant priming effects reported above were due to a fat-tailed distribution, significance in the statistical results should disappear after the tails are trimmed off. As Figure 4.5 shows, cutting the tails removed the fat right tail shown in Figure 4.4. While the priming magnitude estimate (the dotted vertical line in Figure 4.4 was lower than the priming magnitude originally estimated (see Figure 4.3 above), the LMER and BF analyses (performed as usual, see sec. 4.2.4 for further details) still showed an effect of relatedness: $F(1, 543.82) = 8.19, p = .004$; $BF_{1,0}$=6.55. This suggests that the significant priming effects reported above for the *im*-condition were due to a true shift in the priming distribution.

Second, the results could have been affected by affixal dominance. We have already explored this possibility in sec. 3.2.5 as a way to explain the null results for the *er*-condition. The hypothesis being tested here is that, similar to plural forms (see ch. 3), affixed words (e.g., *impossible*) may not decompose at all if their frequency is higher than the frequency of the corresponding stem form (e.g., *possible*). As dominance was not controlled for while preparing the materials for the experiments, the *in-* and *dis-* conditions not priming might have been due *in-* and *dis*-prefixed prime words having a higher frequency than the corresponding stem

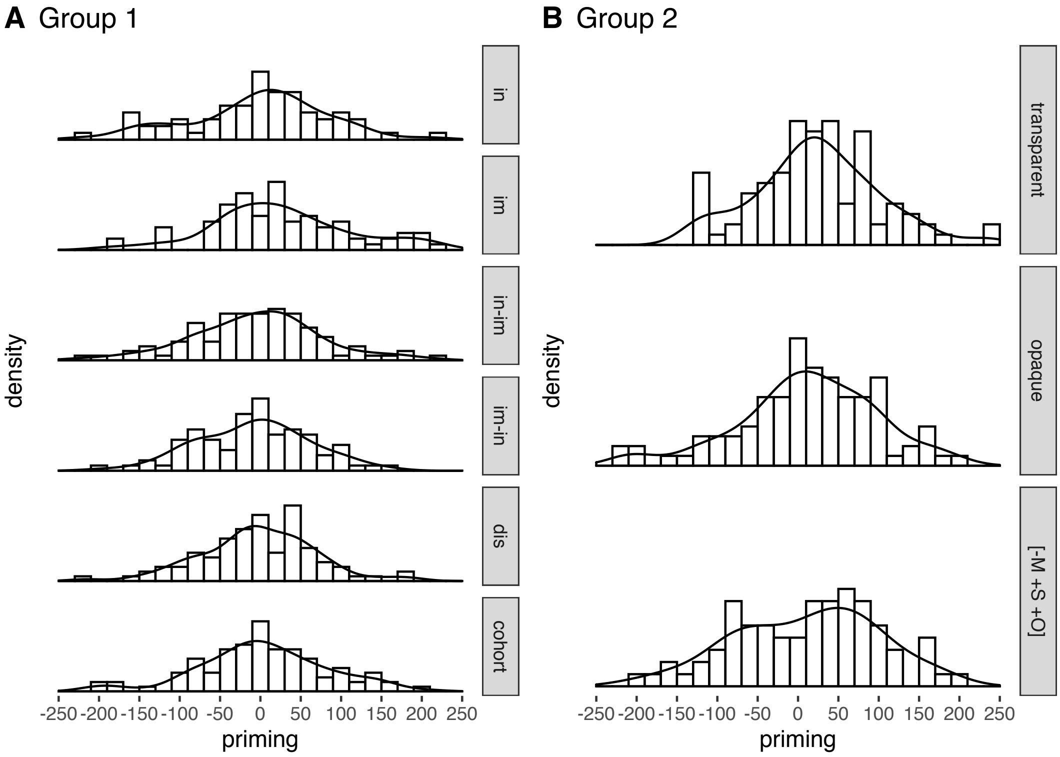

Figure 4.4.: Experiment 3 - *in-im*/visual. Density distribution and priming effects across conditions.

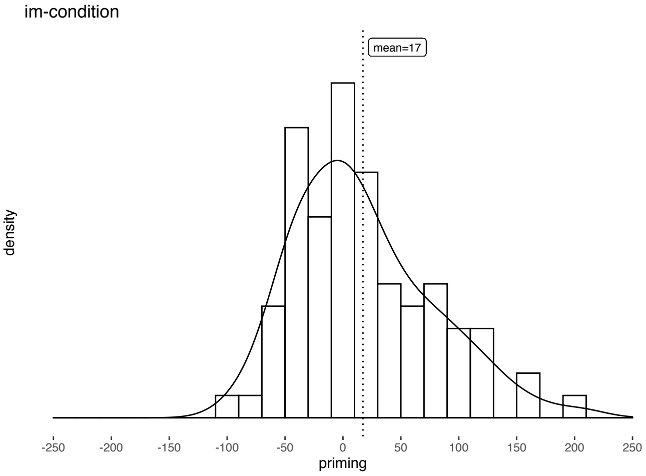

Figure 4.5.: Experiment 3 – *in-im*/affix priming. Density distribution and priming in the *im*-condition after cutting the tails at 1.5 SD. The dotted line represents the mean priming magnitude.

form; this could have in turn inhibited decomposition and, ultimately, priming. We calculated the affixal dominance ratio of the primes and targets used in the *in-*, *im-*, *in-im*, *im-in*, and *dis*-conditions used in this experiment, similarly to what we have done for the *er*-condition in the previous chapter (sec. 3.2.5). Recall that affixal dominance δ_a is calculated by subtracting the log frequency of the stem form (e.g., *possible*) from the log frequency of the affixed form (e.g., *impossible*): $\delta_a = logfreq_{affixed} - logfreq_{stem}$. By this formula, $\delta_a > 0$, if the frequency of the affixed form is higher than the frequency of the unaffixed form (AFFIXED-DOMINANT); and $\delta_a < 0$, if the frequency of the affixed form is lower than the frequency of the unaffixed form (STEM-DOMINANT). If dominance impinged on the priming response of the conditions considered, we would expect two things: (i) the words used in the *im*-condition should be stem-dominant, so to trigger priming; and (ii) the words in all other conditions should be affixed-dominant, so to inhibit priming. We calculated δ_a for all prime and target words used; bound roots (e.g., *famy* from *infamy*) were assigned a default frequency of 1. Figure 4.6 below suggests a pattern that is opposite to the predicted one. On average, (a) target words were stem-dominant ($\delta_a < 0$) in all conditions, while (b) prime words were affixed-dominant in the *im*-condition ($\delta_a > 0$), and stem-dominant ($\delta_a < 0$) in all other conditions. This pattern would predict that all conditions primed and the *im*-condition did not, which is exactly the opposite of what we report above. To assess whether the dominance ratios statistically varied across conditions, we ran a one-way ANOVA and the corresponding BF test, in which *dominance ratio* δ_a was the dependent variable and CONDITION (5 levels: in, im, in-im, im-in, dis) was the independent variable. Our analyses revealed that the data did not significantly vary across conditions (primes: $F(4,55)=1.86$, $p=.13$; targets: $F(4,55)=0.78$, $p=.54$) and supported the null hypothesis (anectodally, for primes: $BF_{1,0}=0.61$; substantially, for targets: $BF_{1,0}=0.17$). None of the post-hoc pairwise comparisons were found significant (Dunn-corrected ps>.12) or supporting the alternative hypothesis ($BF_{1,0}$s≤1). These results suggest that the dominance ratio of the words tested was statistically similar across conditions. Therefore, if dominance had affected priming, we would have mistakenly expected all conditions to either prime or not prime. As such, dominance could not have contributed to the asymmetrical response we reported across conditions.

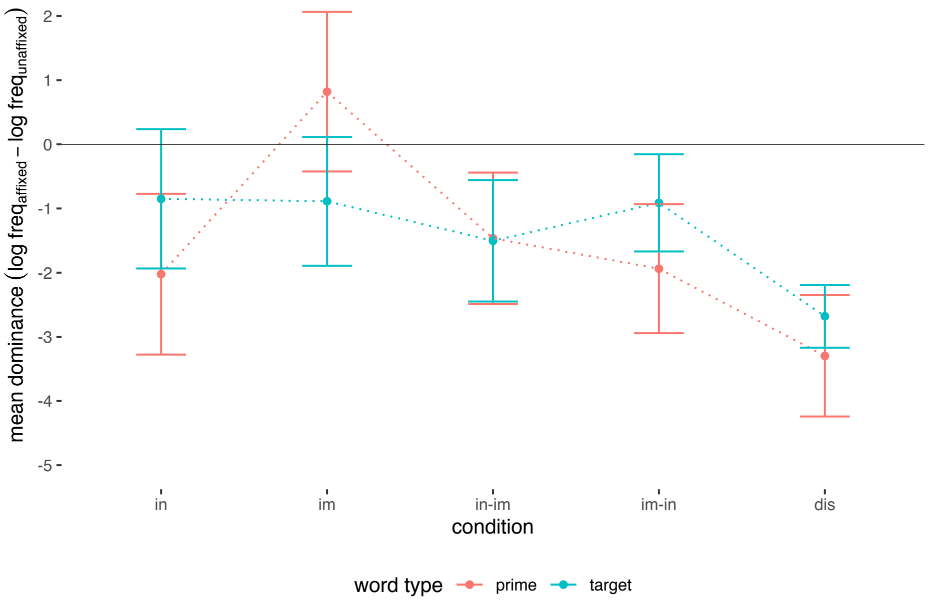

Figure 4.6.: Experiment 3 - *in-im*/visual. Mean affixal dominance ratio across the morphological conditions of group 1.

In the lack of additional factors that might have reasonably affected the priming response, the analyses above further validate the results reported above, at least until future replications (and potential follow-up studies thereby) prove otherwise. As such, we now turn to discuss the main theoretical implications these results have on theories of decomposition.

Towards a reassessment of morpheme segmentation

Taken at the face value, the results of the *in-, im-*, and *dis*-conditions seem to suggest that the affix *im-* decomposes, but the affixes *in-* and *dis-* do not. This represents a rather unexpected pattern, which none of the models of decomposition considered in this dissertation is factually able to account for. As we have done in previous chapters, in what follows, we work these results through each model. As we will see, while the effects reported above could not be accounted for in the bin model, they might have been explained in either of the remaining three models as resulting from the morpho-orthographic segmentation mechanism that decomposition is assumed to rely on (Rastle and Davis, 2008). We show, however, that the morpho-orthographic patterns predicted by this model are not borne out, which ultimately calls for a reassessment of morpheme segmentation and, more generally, decomposition.

A. BIN MODEL

In the bin model, decomposition is implemented as part of the orthographic comparison mechanism in the entry-opening stage. Therefore, under the assumption that priming effects may be linked to early decomposition procedures, our results suggest that the orthographic comparison needs to be modified in two ways: (i) *im*-prefixed word entries are opened, so that priming could be triggered; and (ii) *in-* and *dis*-prefixed words entries are not opened, so that priming could be inhibited. However, the model does not seem able to provide a mechanism to implement this.

At the verification stage, all opened entries are further analyzed in more detail until the entry matching with the stimulus is selected. As such, the verification stage could not have impinged on the priming effects, as it is, by definition, blocked by the priming masking (Forster, 1998).

B. RACE MODEL

In the race model, decomposition-via-priming effects are generally assumed to arise at the segmentation stage in the parsing route. However, as already discussed at length in previous chapters (see sec. 1.4.2 B, in particular), the race model does not provide a full-fledged description of the segmentation mechanism, though it may be assumed to rely on orthographic statistical regularities (similar to the full-decomposition and the morpho-orthographic model; see right below). Therefore, our results could be explained in the race model if, at early stages, the string $\$ \mathrm{i\,m} \$$ is identified as a morpho-orthographic unit, but the strings $\$ \mathrm{i\,n} \$$ and $\$ \mathrm{d\,i\,s} \$$ are not. We explore this possibility in detail further below.

The remaining two stages (namely, the licensing stage and the composition stage) seem unlikely to have impinged on the priming response. At the licensing stage, the morpho-syntactic information associated with the respective morphemes is properly accessed. This stage could not have impinged on priming because all prime words presented in this experiment are morphosyntactically well-formed; thus, the lexical entries of the respective morphemes should therefore all be properly accessed. At the composition stage, the morphemes are put back together and properly interpreted. This stage could not have impinged on priming because it is generally assumed to be blocked by prime masking.

C. FULL-DECOMPOSITION MODEL

In the full-decomposition model, decomposition-via-priming effects are assumed to occur at the decomposition stage, where the visual stimulus is broken down according to orthographic statistical regularities, as discussed in sec. 1.4.2 C ("islands of regularity"; Rastle and Davis, 2008). Therefore, similar to the race model above, the full-decomposition model could explain the results if, at the decomposition stage, the string $ i m $, but not the strings $ i n $ and $ d i s $, is parsed as a morpho-orthographic unit. We explore this possibility in detail further below.

The remaining two stages (namely, the lookup stage and the recombination stage) seem unlikely to have affected the priming response. At the lookup stage, access to morpheme lexical entries occurs and is expected to be affected by morpheme frequency. While morpheme frequency has been shown to affect RTs in traditional (i.e. non-priming) lexical decision tasks, it is unlikely to impinge on the priming response, as prior presentation of a related prime should outweigh any frequency-driven advantage in target recognition. At the recombination stage, the morphemes are recombined in order to be properly interpreted. This stage could not have impinged on priming because it is generally assumed to be blocked by prime masking.

D. MORPHO-ORTHOGRAPHIC MODEL

In the morpho-orthographic model, decomposition-via-priming effects are assumed to occur at the morpho-orthographic segmentation level where the visual stimulus is broken down on the basis of Rastle and Davis (2008)'s morpho-orthographic islands of regularities (sec. 1.4.2 D). As such, the morpho-orthographic model aligns with the the full-decomposition and the race models above in explaining the results as indicative of the fact that the string $ i m $, but not the strings $ i n $ and $ d i s $, is identified as a morpho-orthographic unit. We explore this possibility in detail further below.

The remaining two computation levels (namely, the orthographic lexicon and the lemma level) seem unlikely to have affected the priming response. The orthographic lexicon stores the orthographic form of roots, inflected, and derived complex forms and, potentially, their orthographic frequency. In particular, as we said for the lookup stage in the full-decomposition model above, orthographic frequency could unlikely affect the priming response (unless it is implemented as lateral inhibition affecting activation of competing forms; see sec. 3.2.5 for an implementation in this direction). The lemma level stores all lexical entries (and the associated information), and is assumed to be blocked by prime masking.

Out of the four models just discussed, the race model (B), the full-decomposition model (C), and the morpho-orthographic model (D) all assume a segmentation procedure based on morpho-orthographic sub-lexical units (i.e., Rastle and Davis, 2008's "islands of regularities"). In the remainder of this section, we therefore take the chance to assess the reliability of the morpho-orthographic segmentation procedure assumed by the models (B-D) above, as the potential source of the unexpected priming results. It is important to keep in mind that the *seg-*

mentation procedure is *logically separate* from the *decomposition* procedure: the former is the process whereby morpho-orthographic units are acquired and retained in the long-term memory; the latter is the real-time process of identifying morpho-orthographic units within the visual stimulus. It is important to keep in mind that the investigation explore above assumes that decomposition relies on the morpho-orthographic units that are *exclusively* identified through the segmentation algorithm. Under this hypothesis, if a morphemic letter string is not identified as a morpho-orthographic unit through the segmentation algorithm, it is not expected to prime. The morpho-orthographic units are argued to be learned during reading development in a similar way as word segmentation during language acquisition. For this reason, the strategies that are thought to be involved in the formation of morpho-orthographic units are drawn from the computational and psycholinguistic literature on spoken word segmentation (among others, Saffran et al., 1996a,b; Elman, 1990; Davis, 1999; Seidenberg, 1987; Jusczyk, 1997; Davis et al., 2002; for a review, Davis, 2003). The algorithm consists of evaluating the likelihood of sound/letter sequences in a given stimulus forming a unit at a given level (syllables, morphemes, words, utterance) on the basis of their transition probability (henceforth, TP) of co-occurrence, namely the probability of a given string given the preceding one ($p(b|a)$, where a, b are two strings that form the sequence ab). Usually, the sequences being analyzed by the algorithm are composed of two or three elements (respectively called bigrams and trigrams, if the elements are letter strings).[10] In these terms, in a given complex words such as *undo*, the TP of the bigram \$ u n \$ forming the morpho-orthographic unit *un-*, the TP of the bigram \$ n d \$ straddling a morpheme boundary, and the TP of the bigram \$ d o \$ following the morpheme boundary are expected to exhibit a trough pattern, in which the TP of the cross-boundary bigram \$ n d \$ is lower than the TPs of both the adjacent bigram strings \$ u n \$ and \$ d o \$ (Seidenberg, 1987; Rastle et al., 2004): $\text{TP}_{un} > \text{TP}_{nd} < \text{TP}_{do}$. A formalization of the trough pattern is given in (10).

(10) A given 4-gram string follows a TROUGH PATTERN if and only if the TPs of the first bigram (formed by the first and the second letter of the string, and formalized as 1:2) and the third bigram (formed by the third and the fourth letter of the string, and formalized as 3:4) are greater than the TP of the second bigram (formed by the second and third letter of the string, and formalized as 2:3):

$$N_{\vee} = \text{TP}_{1:2} > \text{TP}_{2:3} < \text{TP}_{3:4}$$

Notice that the trough patterning just described operates on the basis of the *relative* TP trend across subsequent bigrams, in the sense that it makes no reference to a given predefined, absolute threshold. Therefore, if the bigram sequence \$ u n \$ - \$ n d \$ - \$ d o \$ follows a trough pattern, a boundary is placed between \$ n \$ and \$ d \$, regardless of the actual TPs of each bigram, thus triggering the morpho-orthographic segmentation \$ $\boxed{\text{u n}}$ $\boxed{\text{d o}}$ \$, where each box represents a morpho-orthographic unit.

---

[10] For the sake of clarity of exposition, we chose not to calculate trigram TPs here. This reason for this comes from the fact that the current morpheme segmentation literature actually gives no specific detail about the criteria whereby the system decides the length of the gram string to look in the search of morpho-orthographic units.

The segmentation algorithm above predicts that "those affixes that *consistently surface* in words with reliable low-level segmentation cues (e.g., a robust bigram trough separating the affix from its stem) would be more readily segmented by readers and hence produce more reliable masked priming effects than would those affixes that do not surface in the context of such segmentation cues" (Rastle and Davis, 2008, p. 955). We purposefully emphasized the term "*consistently surface*", because it seems to suggest that the system forms morpho-orthographic units by generalizing over the total number of trough patterns in which the same bigram occurs (e.g., *undo, unveil, unchain, unworn, ...*). The literature, on the other hand, does not provide enough details about how the system quantifies the "robustness" of a given trough patterning. Therefore, here we interpret it in terms of relative frequency; that is, a morpho-orthographic unit of an affix is formed as long as the number of the relative bigram sequences that follow a trough pattern ($N_{\vee}$) is greater than the number of bigram sequences that do not ($N_{\not\vee}$)

(11) ROBUSTNESS OF THE TROUGH PATTERN
A trough pattern is *robust* if $N_{\vee} > N_{\not\vee}$.

Assuming (11) is true as part of the segmentation algorithm above, we have the following possibilities: either (i) a boundary may be mistakenly placed between letters forming an actual morphemic n-gram; or (ii) a morpho-orthographic unit may be mistakenly detected for a non-morphemic n-gram. Under the assumption that decomposition relies on morpho-orthographic units that are exclusively detected through the segmentation algorithm, scenarios (i)–(ii) make specific predictions regarding the priming response.[11] The affix *im-* priming suggests that, when a *im*-prefixed prime word is presented (*immature*), the string $\$\ \mathrm{i\ m}\ \$$ is marked as a morpho-orthographic unit ($\$\ \boxed{\mathrm{i\ m}}\ \$$), so that when the related target word is presented (*IMPURE*), the overlap in the word-initial segmented morpho-orthographic unit triggers priming. As such, we expect that the a robust trough pattern *consistently* arises across several bigram sequences involving the string $\$\ \mathrm{i\ m}\ \$$, with the trough peak occurring at the morpheme boundary between the last letter of the prefix *im-* and the stem-initial letter (12a). On the other hand, the affixes *in-, dis-* not priming suggests that, when a *in-* or *dis*-prefixed prime word is presented (*inelegant, disembark*), the string $\$\ \mathrm{i\ n}\ \$$ or $\$\ \mathrm{d\ i\ s}\ \$$ is not identified as separate morpho-orthographic units (e.g., $\$\ \boxed{\mathrm{i\ n}}\ \$$, $\$\ \boxed{\mathrm{d\ i\ s}}\ \$$); therefore, when the related target word is presented (*INTREPID, DISABLE*), there is no morpho-orthographic overlap with the prime word and priming does not arise. As such, we expect that the a robust trough pattern does not *consistently* arise across several bigram-to-bigram sequences involving the strings $\$\ \mathrm{i\ n}\ \$$ (12b) and $\$\ \mathrm{d\ i\ s}\ \$$ (12c).

(12) a. $N_{im,\vee} > N_{im,\not\vee}$

---

[11] Rastle and Davis (2008) argue that the segmentation procedure may additionally rely on "form-meaning regularities", which are assumed to be fully developed even in beginner readers as part of their adult-like speaking proficiency (at least in typical populations). It is worth pointing out that, under this assumption, the null results reported in this chapter are even more surprising, as they would be expected to further strengthen the morpho-orthographic representations of affixes, and therefore boost correct segmentation, which instead seems to be challenged by the priming results reported in this chapter (see also the discussion further below).

b. $N_{in,\vee} \not> N_{in,\not\vee}$
c. $N_{dis,\vee} \not> N_{dis,\not\vee}$

To test the predictions (12) above, we calculated the TPs of the bigrams involving the affixes *in, im, dis*, and all the possible bigram continuations. We additionally calculated the TPs of the bigrams and trigrams involving the affixes tested in the transparent condition (group 2), as a way to compare TP trends across the main morphological conditions of experiment 3. As the transparent condition primed similarly to the *im*-condition, at least some of the affixes tested in the transparent condition (i.e., *-ly, -ed, -ing, -ance, -ful*) are expected to follow a robust trough pattern as in (11).[12] The TP analysis is based on a database constructed on the Google Books Ngram Viewer dataset.[13] From the database, we first selected all 4-grams beginning with the two letters of the affixes that are at the edge of the morpheme boundary. So for the prefixes *-in, -im, -dis* (group 1), the 4-grams extrapolated were `$in--$`, `$im--$`, and `$(d)is--$` (where the dash '-' is a placeholder for any letter): e.g., `$inat$`, `$inco$`, `$intr$`, etc. For the six suffixes of the transparent condition (group 2), the 4-grams extrapolated were: `$--ly$`, `$--ed$`, `$--ic(e)$`, `$--in(g)$`, `$--an(ce)$`, and `$--fu(l)$`; e.g., `$ally$`, `$tily$`, `$ngly$`. The purpose here was two-fold. On one hand, we extracted the right (for prefixes) and left (for suffixes) edge of tri- and four-gram affixal strings in order to focus on the TP-patterns occurring across morpheme boundaries. On the other hand, we selected 4-gram strings from the ngram database as an approximate solution to have an acceptable set of all possible gram sequences involving each affixal portion, while abiding by the English phono-orthographic constraints. From each 4-gram, we then extracted the bigram sequence and assigned a letter position coordinate to each bigram. For example, the four-gram string `$inat$` was broken down into the bigrams `$in$` - `$na$` - `$at$`, which were assigned the letter positions 1:2 (first letter : second letter), 2:3 (second letter : third letter), and 3:4 (third letter : fourth letter), respectively. We then calculated TP of each bigram `$ab$` as shown in (13).[14]

(13) $$p(b|a) = \frac{p(ab)}{p(a)} = \frac{freq-of-ab}{freq-of-a}$$

Figures 4.7-4.8 show the bigram TPs of the sequences involving the prefixes and suffixes tested in the experiment, respectively. In each figure, panel A graphically represents the bigram

---

[12] The suffix *-y* was not included in the following calculations because it is a monogram and could therefore not show a clear trough pattern as shown above in (10).

[13] The N-gram database can be freely downloaded at http://norvig.com/mayzner.html.

[14] Our TP calculations follow Saffran et al. (1996b), though we ought to mention an alternative, slightly different, way to calculate TPs, which consists of dividing the frequency of a given string `$ab$` by the summed frequency of all bigrams beginning with `$a$`:

(i) $$p(b|a) = \frac{p(ab)}{p(a)} = \frac{freq-of-ab}{freq-of-bigrams-beginning-with-a}$$

These two ways reflect the theoretical question about whether the algorithm looks only at the frequency of bigrams as in (i); or, if it looks at the frequency of all grams involving the first letter of the string as in (13). As the TPs calculated as in (i) did not differ from the TPs calculated as in (13), they will be reported here.

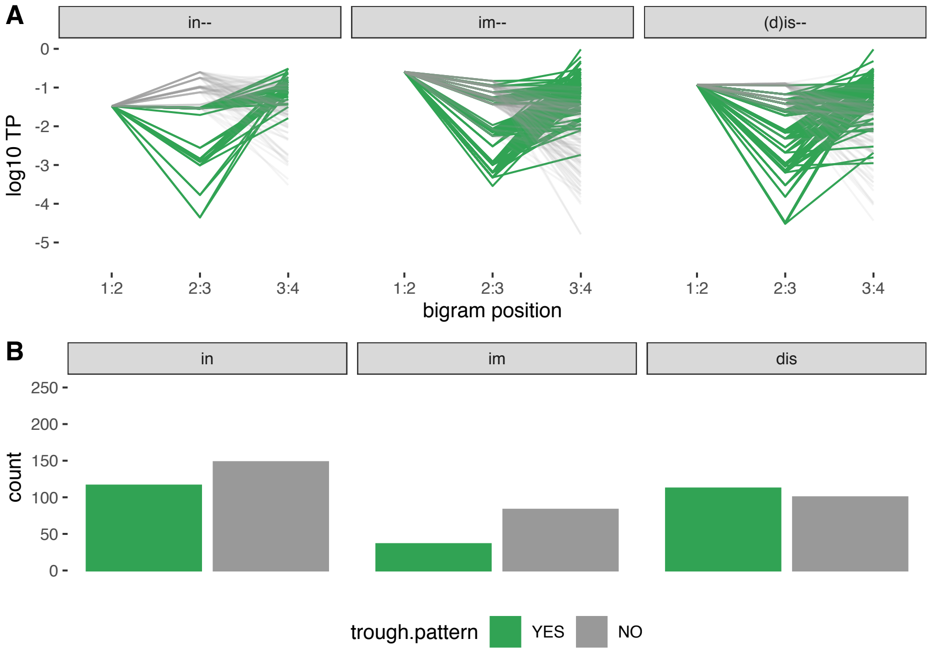

Figure 4.7.: Experiment 3 - *in-im*/affix priming. TP-patterns of the bigram sequences involving the affixes tested in group 1. In panel A, the number coding on the x-axis refers to the position of the letters occurring in a given string. The green sequences follow the trough pattern as sketched in (12a); the grayed-out sequences do not.

TP patterns for each affix; panel B reports the number of bigram sequences that did and did not follow the trough pattern as defined in (10). Our analyses show that the predictions (12a)-(12c) above are not borne out (Figure 4.7). For the affixes *in, im-*, the number of bigram sequences following the trough pattern is smaller than the number of bigram sequences not following the trough pattern, which, under the decision criterion (11), seems to suggest that the strings $ in $ and $ im $ are not identified as morpho-orthographic units. These results predict no priming to arise in both the *in-* and *im*-conditions, in contrast with our results. For the affix *dis-*, the number of bigram sequences following the trough pattern is, though minimally, greater than the number of bigram sequences not following the trough pattern, in contrast with the relative predictions. This suggests that $ dis $ might have been detected as a separate morpho-orthographic unit ($ $\boxed{\mathrm{d\,i\,s}}$ $), which wrongly predicts priming to arise in the *dis*-condition, again in contrast with our results. Finally, as for the suffixes tested in the transparent condition (Figure 4.8), only (11) is satisfied only for two of the six affixes analyzed – i.e., *-ful* and *-ly*. These results seem to suggest that most of the suffixes tested in the transparent condition are not detected as morpho-orthographic units, which therefore predicts priming to be inhibited for most of the prime-target pairs tested in the transparent condition. This prediction does not match the results: priming effects were indeed elicited by the *ful*-suffixed primes (*resentful-RESENT, sorrowful-SORROW*) and two of the three *ing*-suffixed primes (*departing-DEPART, reckoning-RECKON*); however, no priming effects were elicited by the *ly*-suffixed prime (*fluently-FLUENT*; Figure 4.9). The segmentation procedure assumed in the literature seems therefore unable to account for the priming effects on suffixed words (which are widely

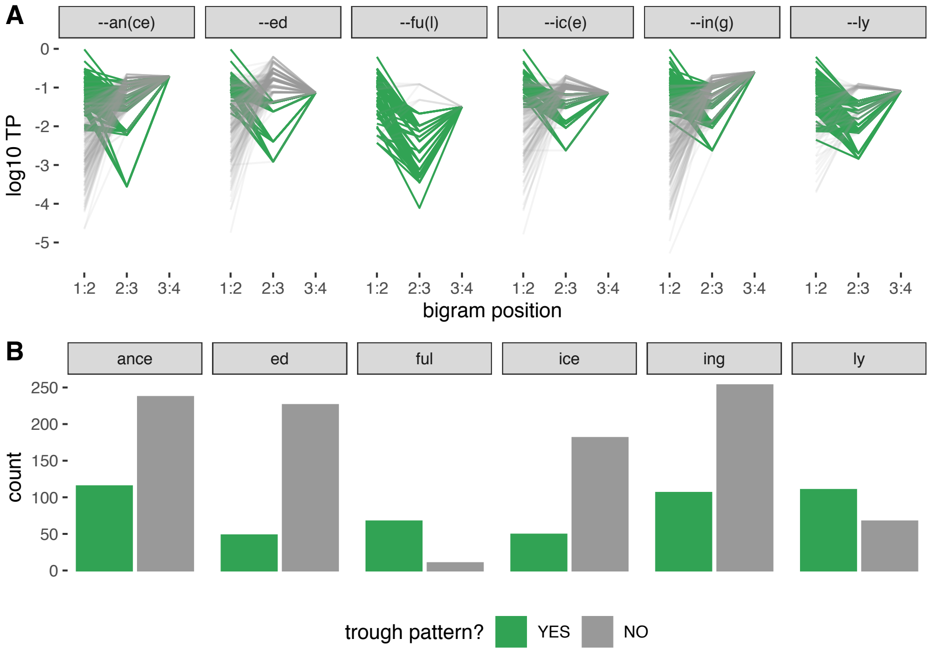

Figure 4.8.: Experiment 3 - *in-im*/affix priming. TP-patterns of the bigram sequences involving the affixes tested in the transparent condition (group 2). In panel A, the number coding on the x-axis refers to the position of the letters occurring in a given string. The green sequences follow the trough pattern as sketched in (12a); the grayed-out sequences do not.

confirmed in the literature), as well as the priming effects on the prefixed words tested in the experiment.

Tacking stock

The goal of this chapter was to investigate, via masked priming, the extent to which early visual decomposition detects the phonologically conditioned morphological alternation *in~im*. The analysis of the priming data reported in this chapter revealed priming effects for one alternant (*im*), but not for the other (*in*). This pattern of results is quite surprising, as all current models of decomposition predict all extant morphemes of a given language to equally decompose (see sec. 4.2.3). As such, we could not attend to the question that experiment 3 was originally designed for; rather, we devoted the current section to discussing potential explanations for such surprising results. After excluding potential confounding factors, we closely explored the possibility that the results were due to the morpho-orthographic segmentation algorithm, which is assumed to rely on orthographic statistical regularities across adjacent letters (i.e., transitional probability: TP). According to this algorithm, the reading system makes morpho-orthographic units out of letters frequently occurring together ("islands of regularity"; Rastle and Davis, 2008) and places morpheme boundaries between letters not frequently occurring together. In the last portion of the section, we therefore tested this algorithm and calculated the TP of all the bigram sequences involving the affixes tested in the conditions of experiment 3:

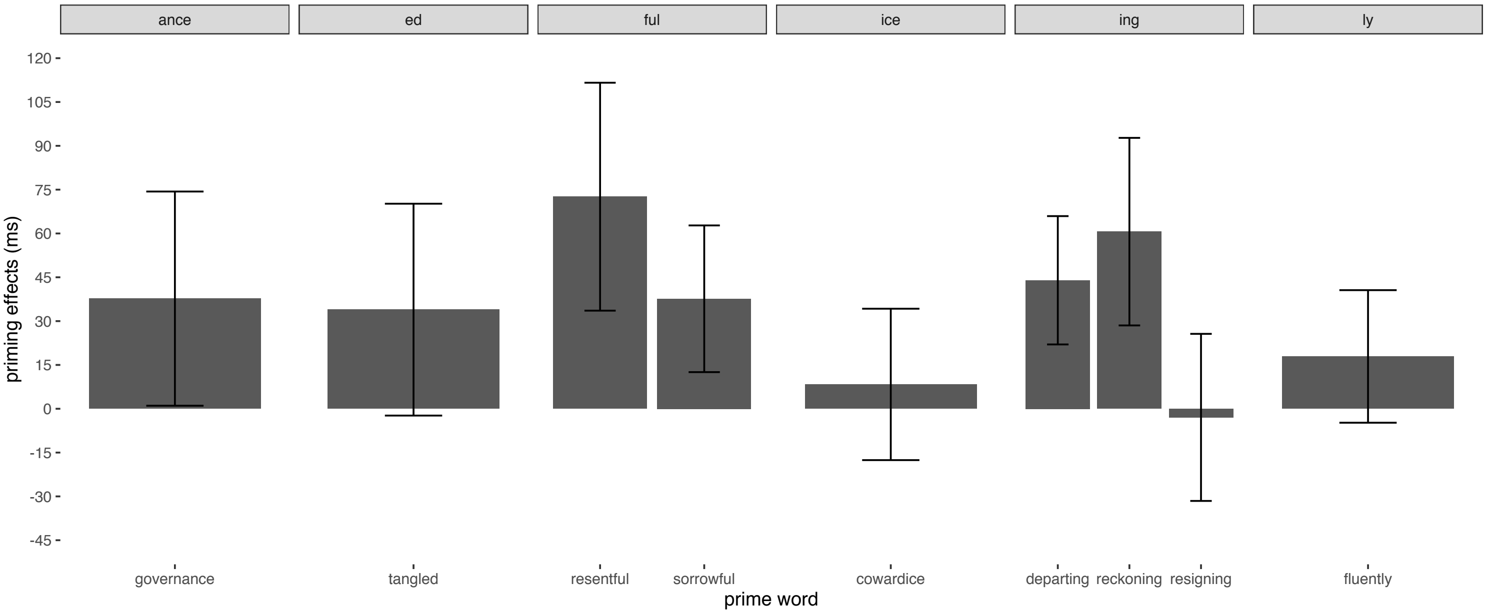

Figure 4.9.: Experiment 3 - *in-im*/affix priming. Priming effects triggered by the prime words tested in the experiment.

in-, im-, dis- (group 1) and *-ance, -ed, -ful, -ice, -ing, -ly* (group 2). Our calculations turned out to run against the predictions made by the morpho-orthographic mechanism of segmentation. Notice that the segmentation algorithm proposed by Rastle and Davis (2008) claims that the low-level TP-based segmentation analysis may be aided by 'form-meaning regularities' in the formation of morpho-orthographically reliable islands of regularities (see fn. 11). For this reason, it would be expected to "see an influence of higher-level factors such as the proportion of semantically transparent forms, affix consistency, and productivity on the effectiveness of morpho-orthographic segmentation for specific affixes" (Rastle and Davis, 2008, p. 959). This expectation also seems to be disconfirmed by our results. In general, we cannot help but notice that the segmentation issue has consistently been taken for granted in the field, with the result that, as far as we know of, a thorough investigation on the morpheme segmentation mechanism is still missing.

From this perspective, the results reported in this chapter (and more generally in this dissertation) call for a reassessment of the morpheme segmentation procedure that is assumed to be at play during early visual decomposition. Such a procedure will need to define the algorithm that, while describing how morpho-orthographic units are acquired by the reading system, is able to account for both the priming effects reported in the literature (that is, the root priming results reported for morphologically transparent words) and those reported in this chapter (if at all replicated). Admittedly, we did not expect the morpho-orthographic segmentation procedure to be as inadequate as described above. For this reason, we are unable at the moment to provide an alternative proposal in this sense. More data will have to be collected on individual morphemes to help generate testable segmentation algorithms.

It is important to underline that, our results still have crucial implications on decomposition, regardless of what the new segmentation algorithm may look like. If the results were true and replicable, they would indeed suggest that decomposition of the prefixes *in-, im-, dis-* do not rely on their morpho-orthographic status; if that were the case, both affixes *in-* and *dis-* would have been found to decompose similarly to *im-*. Rather, they would suggest that decomposition may occur for *some* morphemes only. One way to encode this is to allow decomposition to have access to what we call a "*morpheme repository*." Such a repository is somewhat reminiscent of the orthographic lexicon implemented in Crepaldi et al. (2010)'s morpho-orthographic model, but different in its nature, given that it includes free roots (*cat*), bound roots (*vir-* as in *viral*), and affixes (*-s, -al, in-*), as well as (a subset of) the associated properties (whereas the orthographic lexicon includes the orthographic form of bare roots and fully inflected and derived words). The morpheme repository contains the actual orthographic form of a subset of the morphemes of the language (along the lines of Dikker et al., 2009) and may therefore license their decomposition. So far, the nature of the morphemes included in the repository as suggested by our data seem to escape any abstract, linguistically-relevant categorization; further testing ib individual morphemes would therefore also help identify the actual mechanism underlying these results. The next chapter starts to accommodate the need for more data, while asking a completely different question about the linguistic sophistication of decomposition. The experiments reported in the chapter indeed inform on the priming response to a few other derivational affixes (i.e., *-able, -ity, -ment, -ness*), which may be beneficial for the segmentation and the decomposition issues raised here.

Chapter 5.

Morphological decomposition and *syntactic affixal restrictions*

5.1. Introduction

Chapters 2 through 4 reported a series of experiments that seem to support the distributed view of the mental lexicon and run against the modular view maintained by the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). In this chapter, we continue our investigation of the degree of linguistic sophistication of decomposition and finally test whether decomposition is affected by violations to *syntactic restrictions* associated with affixes. The term "syntactic restrictions" refer to the selectional restrictions by which affixes may only attach to stems of a given syntactic category. For example, the suffix *-ful* may only attach to nominal stems (e.g., *blissful*, but **fastful*); conversely, the suffix *-ness* can only attach to adjectival stems (e.g., *calmness*, but **blissness*). The investigation below investigates the decomposition of syntactically violating forms affixed with one of the following four suffixes: *-able, -ity, -ment* and *-ness*. These suffixes were chosen because they all do not attach to nominal stems: *-able* and *-ment*[15] attach to verbal stems; the affixes *-ity* and *-ness* attach to adjectival stems (see Table 5.1).

AFFIX	ATTACHES TO...	YIELDS...	*example*
able	V	A	detect]$_V$ able]$_A$
ity	A	N	pure]$_A$ ity]$_N$
ment	V (or A)[15]	N	settle]$_V$ ment]$_N$
ness	A	N	weak]$_A$ ness]$_N$

Table 5.1.: Syntactic restrictions of the four suffixes tested. *Legend.* A: adjective; N: noun; V: verb.

---

[15] The affix *-ment* very rarely attaches to adjectives too: e.g., *merry-ment, odd-ment, content-ment*.

While doing so, we acknowledge two major properties of English affixation that require preliminary attention.[16] The first property concerns a special type of affixation, namely the affixation involving the use of the phonologically-null morpheme, the "*zero-morpheme*" [ø]. The ø-morpheme is an abstract device used in theoretical morphology to account for the fact that stems may change their syntactic category without an overtly realized affix (*zero-derivation*). For example, in English, the word *escape* may be used as a verb or a noun without a phonological realization signaling the category change (14).[17]

(14) $\sqrt{\text{ESCAPE}}$:
  a. verb: to *escape-∅*
  b. noun: an *escape-∅*

In current theories of morphology, zero-derivation has virtually no constraint, thus potentially leading to overgeneration.[18] For example, the illicit suffixed forms such as **blissable* could be considered grammatical, if an adjective-yielding zero-morpheme is posited between the root *bliss* and the suffix *-able*:

---

[16] An additional property that is involved in English affixation is morphological blocking, i.e. the phenomenon whereby specific affixed forms are preferred over others on a purely lexical basis. For example, the suffix *-ness* could legally attach to adjectival roots like *wide, strong*, and *long*, but the corresponding *th*-suffixed forms are generally preferred – namely, *width, strength, length* (among others, Aronoff, 1976; Plag, 1999; Embick and Marantz, 2008). Additionally, native speakers seem to accept blocked *ness*-affixed words like *wideness, strongness*, and *longness* to underline the quality of being wide, strong, and long rather than the actual measurement. We will not deal with this issue, as these differences involve accessing the semantics of these complex forms, which is not expected to affect early decomposition (see sec. 1.3).

[17] The operation of zero-derivation changing nouns to adjective and adjectives to nouns is more controversial and rare. For example, adjectives can be turned to nouns: e.g., *the wealthy, the moderate, the poor*. Grammars label these cases as *partial conversion* (where conversion is just an alternative way to refer to zero-derivation), because these zero-derived forms cannot be fully inflected (Quirk et al., 1985, p. 1562). De-nominal adjectives are even more rare (*ib.*):

(i) $\sqrt{\text{BRICK}}$:
  a. a *brick-∅* garage
  b. The garage is *brick-∅*.

(ii) $\sqrt{\text{COTTON}}$:
  a. a *cotton-∅* dress
  b. This dress is *cotton-∅*.

While stating that zero-derivation involving nouns and adjectives is, in principle, possible, these examples may be considered special cases analyzable as involving elliptical constructions.

[18] Supporters of Kiparsky (1982)'s Lexical Morphology and Phonology framework (LMP) have tried to constrain zero-derivation within the level ordering hypothesis (Siegel, 1974). According to this theory, affixes may be grouped in two classes: Level-I affixation occurs first, triggers stress-assignment, and may involve bound roots (e.g., *-able, dis-, -ic, in-, -ion, -ity*); Level-II affixation follows, is stress-neutral, and only involves bare roots (e.g., *-less, -ly, -ment, -ness*). Allen (1979) argues, with independent empirical generalizations, that zero-derivation must follow Level-I affixes; thus, zero-derivation in forms such as (15) should not be allowed and be therefore judged ungrammatical. (As we will see in the following pages, this prediction is only in part confirmed by our data, which shows that *ity*-forms are ungrammatical, but *able*-forms are not.) Following the criticisms raised against the LMP framework (among others, Fabb, 1988), such a theoretical implementation of zero-derivation fell through too. As far as we know, the phenomenon seems now relatively unconstrained in all current theories of morphology.

(15) *blissable*: bliss$]_{N}$ ø$]_{V}$ able $]_{A}$

Application of zero-morphemes as in (15) may be a serious confound for the current purposes, as zero-derivation may license *any* kind of affixation, thus making the question of this chapter moot at best. We address this first concern in Section 5.2, where we explore the extent to which zero-derivation occurs in syntactically violating forms and makes them grammatical. To this end, in experiment 4, we asked subjects to rate syntactically violating forms such as **blissable, *blissity, *blissment, *blissness*. Results show that all illicit forms were judged phonologically grammatical, but *able-, ment-,* and *ness*-suffixed forms were rated higher than *ity*-suffixed forms. We propose two possible interpretations for these differential effects. The first interpretation suggests that the low ratings for the *ity*-suffixed forms is indicative of their ungrammaticality. The second interpretation suggests that the low ratings for the *ity*-suffixed forms were due to an uncontrolled-for violation to a morpho-phonactic constraint that forbids *-ity* from attaching to consonant-final roots. Under this hypothesis, the *ity*-suffixed forms are to be interpreted as grammatical on a par with the other forms and may be not apt for investigating the question about the syntactic sensitivity of decomposition.

The second property we look into concerns selective decomposition, i.e. the fact that decomposition seems to occur only for a subset of the English affixes (see in particular secc. 3.2.5 and 4.2.5). As such, selective decomposition may compromise interpretability of the present investigation. If, for example, syntactically licit *ity*-suffixed words (e.g., *purity*) did not prime and decompose, the lack of priming for syntactically illicit *ity*-suffixed pseudo-words (e.g., *blissity*) could not be interpreted as answering the question about the sensitivity of visual decomposition to syntactic restrictions. We address this concern in Section 5.3. There, we report experiment 5, which elicited the root priming response to *able-, ity-, ment-, ness*-suffixed *licit* words in separate conditions (e.g., *detectable-DETECT*) as a way to address this issue and provide more data that could ultimately contribute to the identification of the mechanism underlying selective decomposition. Results show that all words tested primed, thus suggesting that all four of the affixes involved decompose.

Building on the results of experiments 4-5, we finally turn to addressing the question about the syntactic sensitivity of decomposition. To this end, we report experiment 6, which elicited the root priming response to syntactically *illicit* forms suffixed with one of the four affixes explored in this chapter (e.g., *blissable-BLISS, blissity-BLISS, blissment-BLISS, blissness-BLISS*). We found that the *ness*- and *ity*-conditions trigger significant (and large) priming effects; the *able*- and *ment*-conditions instead elicited trending (and smaller) priming effects. We discuss two possible interpretations of these results, depending on the hypothesis being adopted for the ratings of *ity*-suffixed forms (see above). In particular, under the hypothesis that *ity*-ratings are indicative of ungrammaticality of those forms, the priming effects elicited in experiment 6 seem to suggest that decomposition is not sensitive to syntactic affixal restrictions. Section 5.5 concludes the chapter and discusses the major theoretical consequences for decomposition and visual processing.

5.2. Experiment 4 – illicit/rating

Experiment 4 explores the extent to which zero-derivation may grammaticalize syntactically violating forms such as **blissable*. The term *zero-derivation* refers to the abstract operation whereby word derivation occurs by attaching the ø-morpheme to the stem. As zero-derivation may freely occur, it may potentially license syntactically violating forms, thus ultimately making our question untestable. We directly addressed this concern by exploring the extent to which syntactically illicit suffixed forms can be accepted by native English speakers as potential new word formations via zero-derivation.

In this experiment, we asked subjects to rate, on a typical Likert scale ranging from 1 to 7, the acceptability of syntactically violating forms such as **blissable, *blissity, *blissment, *blissness*, organized in separate conditions (each presenting a form with one of the four affixes: *-able, -ity, -ment, -ness*). Our starting assumption is that acceptability is, as defined in (Sprouse, in press), "the conscious report of the perception of an error signal that arises automatically during the processing" of a linguistic stimulus (word form, sentence). In this sense, acceptability may inform us on the grammaticality of any kind of linguistic stimulus we want to test, as long as the other factors possibly impinging on it are controlled for. Furthermore, we assume that, although numerical acceptability rating/judgment tasks (such as Likert scaling, but also magnitude estimation) yield continuous values, it does not necessarily mean that they tap into a gradient model of the grammar. Rather, gradient values of acceptability may result from other interacting factors that yield gradience effects (e.g., working memory load, frequency, semantic plausibility, prototypicality; see Schütze and Sprouse, 2014; Schütze, 2016). Gradient acceptability may therefore be compatible with a categorical model of grammar, in which a given linguistic expression (let it be a word form or a whole sentence) is either grammatical or ungrammatical.

As a way to facilitate the interpretation of (inherently continuous) acceptability ratings in terms of categorical grammaticality, we also recorded acceptability ratings of phonotactically and morphologically illicit filler items, which were chosen depending on their rating as reported in the literature (Hayes and White, 2013; Daland et al., 2011; Albright and Hayes, 2003). The actual ratings of these fillers are used to construct a reference acceptability scale ranging from 1 to 7 of the Likert scale. This scale is then used to interpret the acceptability ratings of the target items in terms of grammaticality. To do this, we set two distinct thresholds: the threshold based on the phonological fillers measures phonological/phonotactic goodness, while the threshold based the morphological fillers measures morphological/lexical goodness. Both thresholds are set at the mean rating of the filler items with the expected midpoint rating of the scale. In particular, the phonological threshold will be set at the midpoint between the mean ratings of the fillers with an expected rating of 3 and 5; the morphological threshold will be set at the mean rating of the fillers with an expected rating of 4. Target items will be considered phonologically and/or morpho-phonologically grammatical only if the relative ratings are above the respective threshold. Such a "crude" linking hypothesis is generally inadequate for syntactic acceptability judgments, since deviations from it are well known (e.g., ungrammatical sentences being rated high, and grammatical sentences being rated low; for a review, see Schütze and Sprouse, 2014). However, we still assume the hypothesis above in this experiment, as the items here are not likely to show such an exceptional behavior. Our results show that syntactically illicit items in the *able-, ment-* and *ness*-conditions received ratings that were above the morphological threshold, whereas syntactically illicit items in the *ity*-condition received ratings that were below it. These results suggest that only the *ity*-suffixed illicit forms may be tested for the purpose of this chapter. We however point out that an alternative interpretation of the results is possible. The low acceptability ratings in the *ity*-condition might have been due to the fact that most of the noun roots used in the experiments violate the morpho-phonotactic constraint whereby *-ity* attaches to Latinate, monosyllabic, and vowel-final stems. If this alternative conclusion is entertained, *ity*-suffixed forms tested are to be interpreted grammatical akin to the remaining suffixed forms, which suggests that zero-derivation may not be excluded as impinging on the acceptability of syntactically illicit forms.

5.2.1. Materials

One hundred and five roots were selected from the English Lexicon Project corpus (ELP: Balota et al., 2007). The roots were either disyllabic or monosyllabic, and did not end with substrings that could potentially be considered suffixed (e.g., *butter*). We selected roots that are predominantly used as nouns. Each root was then matched with the four suffixes: *-able*, *-ment*, *-ity*, *-ness*, for a total of 420 suffixed forms. We were particularly careful in selecting the roots, so that the resulting suffixed forms did not raise any orthographic or phonological issue across boundaries (e.g., *theory*: **theoriable*; *pasta*: *pastaable*; *kite*: *kiteity*). All roots selected generally ended with a consonant, so that morphological boundaries were easily identifiable.

From this set of bare noun roots, twenty-four were then selected for the rating experiment depending on their HAL log10 frequency (taken from the ELP database: Balota et al., 2007), and arranged in three frequency bins. This was done to prevent root frequency effects from potentially impinging on the rating response. The first bin had roots with low frequency (<7; $mean = 6.53, s.d. = 0.16$); the second bin had roots with mid frequency (<9; $mean = 8.18, s.d. = 0.06$); the third bin had roots with high frequency (>9; $mean = 10.11, s.d. = 0.85$). We ensured that the mean frequencies of the three bins were maximized while controlling for the mean frequency difference between neighboring lists ($mean\,freq_{bin2} - mean\,freq_{bin1} = 1.65; mean\,freq_{bin3} - mean\,freq_{bin2} = 1.72$) and orthographic length across the three bins ($mean = 4.16, s.d. = 0.63; F(2,21) = 0.7, p = .51$; $BF_{0,1}$=2.7). The roots, their frequency and length, and the resulting illicit suffixed forms can be found in Appendix IV. The resulting 96 suffixed forms were then arranged in 6 different lists in a Latin-Square fashion. Each list consisted of 16 target items, of which 5 were taken from each frequency bin and the remaining item was randomly taken from one of frequency bins.

As a way to establish a reference rating scale to base the interpretation of the target ratings on, thirty-two fillers were then added to each list. Half of them were taken from phonotactic rating studies (e.g., *nlezzig*; Hayes and White, 2013; Daland et al., 2011) and the other half from morphological rating studies (e.g., *gezzed*; Albright and Hayes, 2003). The phonological fillers were non-word forms which instantiated violations to ortho-phonotactic constraints of English. The morphological fillers were phonotactically licit past-tense forms, in which a non-existing

FILLER TYPE	EXPECTED RATING	EXAMPLE	REFERENCE
phonotactic	1	*tighw*	Hayes and White (2013)
phonotactic	2	*zreppid*	Daland et al. (2011)
phonotactic	3	*pwudge*	Hayes and White (2013)
morphological		*snold*	Albright and Hayes (2003)
morphological	4	*gleeded*	Albright and Hayes (2003)
phonotactic	5	*sneck*	Hayes and White (2013)
morphological		*stinned*	Albright and Hayes (2003)
phonotactic	6	*kilp*	Hayes and White (2013)
morphological		*stired*	Albright and Hayes (2003)
phonotactic	7	*trisk*	Hayes and White (2013)
morphological		*wissed*	Albright and Hayes (2003)

Table 5.2.: Experiment 4 - illicit/rating. Examples of the filler items used in the expected, listed by illicitness type (phonotactic/morphological) and expected rating as reported in the referenced study.

root was attached to either the regular (e.g., *stinned* < *stin*) or irregular (ablauting: e.g., *snold* < *snell*) past tense morpheme; therefore, the morphological fillers instantiated violations to the phonological context in which the morphological rules of past tense morpheme can apply. The mean ratings from the previous studies were taken as expected ratings, and used to evenly distribute fillers across the 7 points of the Likert scale. Five fillers were selected for the expected ratings of 1, 3, 5, and 7; four fillers were selected for the expected ratings of 2, 4, and 6. We also ensured that the two types of fillers (i.e., phonotactic and morphological items) were as evenly balanced across the expected rating points as possible. This had two main side effects, which will have important consequences in the calculations of the phonological and morphological thresholds (see sec. 5.2.3 below). First, none of the fillers with an expected rating of 4 were phonological. Second, the fillers with the lowest expected ratings (1-3) were predominantly phonotactically illicit (there was only one morphological filler item with an expected rating of 3); this is because all of the morphological items used obeyed English phonotactic constraints. Table 5.2 reports one example for each type (i.e., phonological/morphological) and expected Likert rating. The complete list of fillers can be found in Appendix IV.

5.2.2. Participants & procedure

Sixty English native speakers were recruited on Amazon Mechanical Turk and participated in the experiment. Participants were asked to rate how likely they thought a given form was to be a possible new word in English. To perform the task, they used a 7-point Likert scale, in which 1 was labelled as ‘least possible’ and 7 is ‘most possible’. Each participant was randomly assigned only one list. The experiment took 5-8 minutes on average. Participants received monetary compensation after completing the experiment.

The rating experiment was created on PCIbex Farm (Zehr and Schwarz, 2018). In a typical experimental session, subjects would first go through 9 practice items. These items were taken from previous rating experiments and were not used in the actual experiment. One item was selected for each expected rating from 1 to 7 and one more for both rating 1 and rating 7. The order was fixed for all subjects, regardless of the list they were assigned. Right after the practice, subjects were presented with the 48 items of the assigned list. The presentation order was pseudo-randomized, so that (a) suffixed items were evenly interspersed between fillers and (b) fillers of the same expected rating were not presented one after the other. As a way to screen potential bots, subjects were asked to answer three open-ended questions interspersed in the experiment.

5.2.3. Predictions

We will use the ratings of the (phonological and morphological) filler items as the reference scale against which we interpret the ratings of the syntactically violating forms across the four target conditions. The filler ratings will be split in two different scales. The ortho-phonotactic fillers form the phonological scale as a measure of phonotactic goodness; the morphological fillers form the morphological scale as a measure of morpho-phonological goodness. In each scale, we will then set a threshold at the mean rating of the filler items whose expected rating is equal to the midpoint value in the 7-point Likert scale. The phonological scale does not have fillers with an expected rating of 4 (i.e., the median of 7), so the phonological threshold will be set at the midpoint between the mean ratings of the fillers with an expected rating of 3 and 5. The morphological scale does not have fillers with an expected rating lower than 3 due to the inherent nature of the items (i.e., they are also phonotactically licit); however, it does have fillers with an expected rating of 4, which will therefore be used to set the morphological threshold. The target ratings will be compared to both thresholds and interpreted accordingly.

Phonotactic and morphological grammaticality of the four target conditions will be assessed on the basis of both thresholds. We expect target items to be phonotactically grammatical and therefore be rated above the phonological threshold, as they all follow English phonotactic constraints. Since all target items were created to be ortho-phonotactically licit and able to attach to all of the suffixes tested, the position of the ratings with respect to the morphological threshold will be interpreted as indicating morpho-syntactic grammaticality. Therefore, ratings below the morphological threshold will suggest morpho-syntactic ungrammaticality, thus implying no application of zero-derivation. Under this circumstance, those forms may be apt for an investigation on the sensitivity of decomposition to syntactic restrictions. Conversely, ratings above the morphological threshold will suggest morpho-syntactic grammaticality presumably due to application of zero-derivation. Under this circumstance, those forms may not be apt for the question being asked in this chapter.

5.2.4. Results

Before running any analysis, one subject was removed because all forms were rated with the same Likert point. The remaining datapoints were analyzed without any outlier rejection procedure. First, we performed by-subject z-transformation of the raw (that is, Likert-point) ratings of the filler items. Figure 5.1 plots the distribution of the filler items in both the raw and z-transformed scales, and shows that both scales converge on similar results. The plots also report the phonological and morphological thresholds in both scales, as described in the section above. The convergence across the two scales suggests that the potential risk of scale bias in the raw scale is fairly negligible. For this reason, we chose to primarily use raw ratings in the analyses below, as their interpretation is more transparent to the reader and closer to our predictions and logic (see sec. 5.2.3). The z-score ratings are nonetheless reported for completeness.

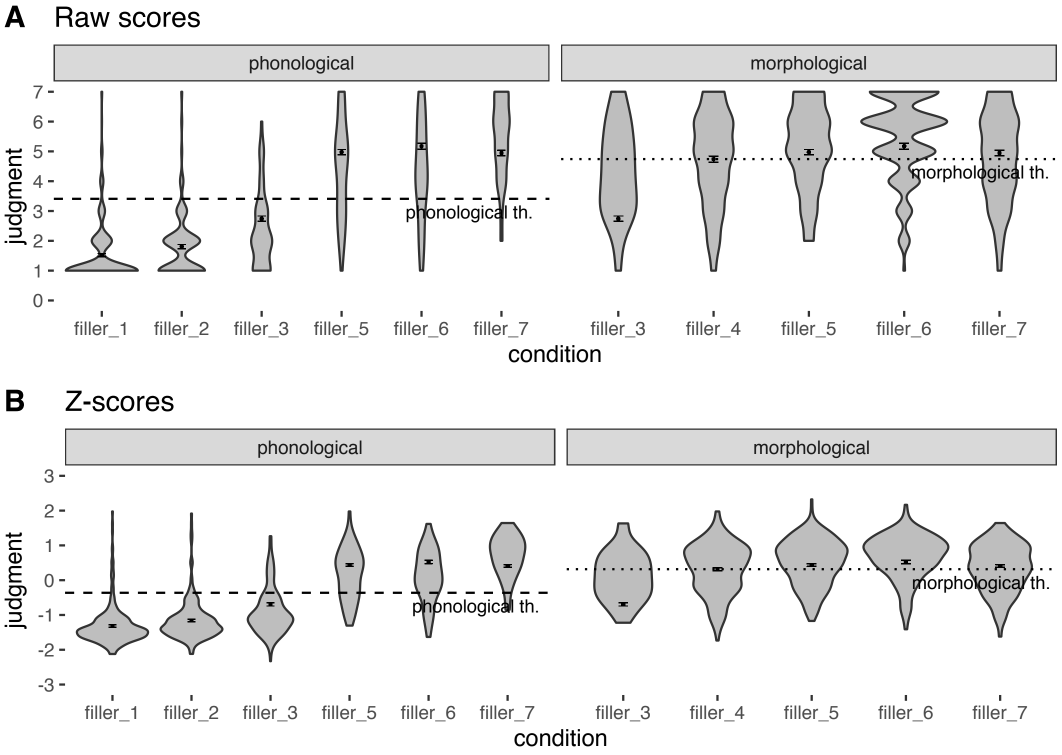

Figure 5.1.: Experiment 4 – illicit/rating. Distribution and means of the raw and the z-transformed ratings of the filler items. The labels on the x-axis refer the corresponding expected rating: e.g., $filler\_1$ groups filler items with an expected raw Likert rating of 1.

Figure 5.2 shows the distribution and the mean ratings of the four conditions with respect to the phonological and morphological thresholds. On one hand, the mean ratings of all target items seem to be well above the phonological threshold, as expected. On the other hand, the mean ratings of the four conditions seem to have a different behavior with respect to the morphological threshold. The mean ratings of the *able-, ment-* and *ness*-items were above it, but the *ity*-items were rated below it. The plot also seems to suggest that the ratings of the

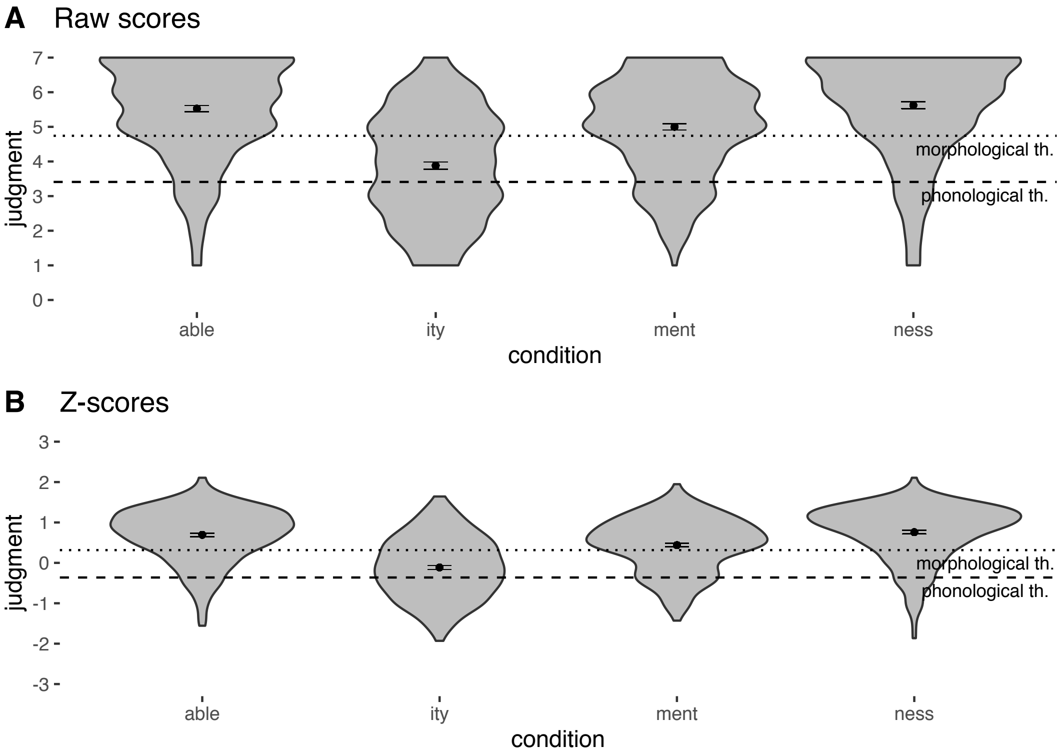

Figure 5.2.: Experiment 4 – illicit/rating. Distribution and means of the raw and the z-transformed ratings of the target items.

ment-condition were lower than the *able*- and *ness*-conditions, which suggests that some extra-grammatical factor is at play. In the analysis below, we will only focus on the above/below threshold differences; potentially significant differences across target conditions will not be further looked into, as not strictly relevant to the current purposes.

We ran two different statistical calculations on the dataset to assess the statistical difference between the target ratings of each condition and the filler ratings of phonological and morphological thresholds. To do this, we built two separate Linear Mixed-Effect Regression models (using the R-package `lme4`), whose *p*-values were estimated using the Satterthwaite approximation of degrees of freedom (using the `lmerTest` R-package). We compared the *raw scores* of the four target conditions to the phonological threshold (*phon*) and the morphological threshold (*morph*). In each model, we also added ITEM and SUBJECTS as random effects (intercept only). Similarly, we performed a series of one-sample Bayes Factor tests by using the R-package `BayesFactor` (Morey, 2018); for the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.5. All conditions were statistically different from both the *phon* and *morph* thresholds (ps<.05, $BF_{1,0}$>10000). The details of the LMER and BF analyses are reported in Table 5.3.

	CONDITION	β	t	df	p	$BF_{1,0}$	$BF_{1,0}$ interpretation
PHON VS.	able	2.10	13.59	37.89	<.0001	>10000	strong for H_1
	ity	0.47	3.81	36.24	.003	>10000	strong for H_1
	ment	1.61	10.62	39.90	<.0001	>10000	strong for H_1
	ness	2.20	11.48	41.51	<.0001	>10000	strong for H_1
MORPH VS.	able	0.76	4.97	37.08	<.0001	>10000	strong for H_1
	ity	-0.86	-5.8	36.24	<.0001	>10000	strong for H_1
	ment	0.28	1.82	39.9	.08	>10000	strong for H_1
	ness	0.87	4.55	41.51	<.0001	>10000	strong for H_1

Table 5.3.: Experiment 4 – illicit/rating. Summary of the statistical results.

5.2.5. Discussion

The main goal of the rating experiment was to address the extent to which zero-derivation possibly impinges on the decomposition of syntactically violating forms. English derivation may indeed occur via *zero-derivation*, in which the ø-morpheme can attach to any stem and turn it to any given category; for example, the word *escape* may be freely used as verb, noun, or adjective, without the need for a phonologically-overt affix signaling the category change. As virtually unconstrained, zero-derivation may in principle license illicit affixation, thus potentially compromising the feasibility of the question regarding the syntactic sensitivity of decomposition. If zeroes may be freely inserted if need be, there may be no way to test the sensitivity of decomposition to syntactic affixal restrictions without zero-derivation possibly interacting. We explored the extent of this risk by testing the acceptability of syntactically illicit items such as **blissable, *blissity, *blissment, *blissness* in four separate conditions (each including illicit forms sharing the same affix). The logic behind the experiment was that if zero-derivation applies, syntactically violating forms are grammatical and should, therefore, be rated high; if zero-derivation does not apply, they are ungrammatical and should be rated low. To this end, we used the ratings of the filler items as reference to interpret the ratings of the target items. The fillers were split in two different scales, in each of which a goodness threshold was identified in the mean of the ratings of the fillers whose expected rating corresponded to the median value of the corresponding scale (i.e., the ratings of the fillers with an expected rating of 3 and 5 in the phonological scale; the ratings of the fillers with an expected rating of 4 in the morphological scale). The phonological fillers formed the reference scale of phonotactic goodness and were used to assess ortho-phonological grammaticality of the target items. All target items were found to be rated significantly higher than the phonological threshold, thus suggesting that they were all ortho-phonologically grammatical. This was indeed expected, as all target items abide by English phonotactic constraints. The morphological fillers formed the reference scale of morpho-phonological goodness, which was used to assess the morpho-syntactic grammaticality of the target item. All target items were indeed prepared so not to raise any orthographic, phonological, and/or morphological issue (see sec. 5.2.1). Therefore, any variation in the ratings from the morphological threshold may be only interpreted as due to violations to syntactic affixal restrictions. The ratings in the *able-*, *ment-* and *ness*-conditions were found to be above the morphological threshold; the ratings in the *ity*-condition were instead found to be below it.

These results seem to suggest that (i) acceptability of new word formations may indeed be affected by zero-derivation procedures and (ii) zero-derivation-driven grammaticality is contingent on the affix being involved. At the current stage of the research, however, we have no plausible explanation for the complete pattern of results reported above. On one hand, no lexical properties that are usually associated to affixes are able to account for the pattern of results reported here. First, neither affix frequency nor family size seem to correlate with the acceptability of the forms tested here (see Table 5.4). The table also reports Baayen and Lieber (1991)'s affix productivity index $\wp^*$, which is defined as the ratio between the number of hapax legomena with a given affix (namely, the number of types occurring only once in the sample) and the total number of tokens of all words with that affix. At first glance, affix productivity is a seemingly good potential source of the effects reported above. We however ought to point out that productivity as usually defined in the literature does not seem able to positively correlate with acceptability (i.e., in the sense that high productivity correlates with high acceptability). On one hand, productivity is generally seen as the probability of a given affix to be involved in the coinage of new formations that are expected to be seamlessly acceptable and *interpretable* by other hearer-speakers (see for example, Schultink, 1961). On the other hand, acceptability does not always correlate with *interpretability* (for a review, Schütze, 2016): some sentences that are not semantically interpretable (e.g., *More people have been to Russia than me*) are generally found to receive high ratings; whereas some sentences that are semantically interpretable but syntactically hard to parse (e.g., *The nurse the patient the dog bit cured yelled*) are instead found to receive relatively low ratings. Therefore, we believe that the linking hypothesis between productivity and acceptability, if it exists, is hard to define; regardless, it should not include the notion of interpretability. Nevertheless, even if productivity may be encoded in the mechanisms underlying acceptability, the $\wp^*$ indices reported in Table 5.4 would still be unable to explain the effects. For these reasons, we may interpret the results as suggesting that, out of the four syntactically violating form types tested, only the *ity*-form type may be used to directly explore the sensitivity of decomposition to syntactic affixal restrictions. Accordingly, we predict that, if decomposition is affected by violations to syntactic affixal restrictions, priming (i) should arise for the grammatical, but syntactically violating forms affixed with *-able, -ment-, and -ness*, similarly to the corresponding syntactically licit forms; but (ii) should not arise for the ungrammatical and syntactically violating forms affixed with *-ity*.

AFFIX	HAL frequency (log10)	Family size	$\wp^*$
able	6.08	872	0.006
ity	6.21	580	0.004
ment	6.15	288	0.002
ness	5.25	1243	0.016

Table 5.4.: Experiment 4 – illicit/rating. Lexical properties of the four affixes tested (Sánchez-Gutiérrez et al., 2018).

On the other hand, here we point out that a specific non-syntactic selectional restriction may explain at least the relative low rating of the *ity*-condition. It is indeed well known that the affix *-ity* selects for Latinate stems — that is, stems with a Greek-/Latin-like morphological pattern that have come into English through borrowing from French (Aronoff and Fuhrhop, 2002; see also, Marchand, 1969; Aronoff, 1976). In this sense, the ungrammaticality of the *ity*-suffixed forms might have been triggered by the fact that the noun free stems used in the experiment were primarily non-Latinate. There is, however, some debate regarding the way speakers may detect the etymological origin of morphemes without referring to the associated lexical information. While syllabic length does not seem able to unequivocally characterize *ity*-attaching Latinate stems (Anshen et al., 1986), the Reverse English Dictionary (Muthmann, 2010) actually suggests that *-ity* may only attach to vowel-final, but not consonant-final, monosyllabic stems: out of the 600 *ity*-suffixed words, only 2 were indeed found to have consonant-final, monosyllabic stems (i.e., *oddity, nullity*). Since 19/24 noun stems used in the experiment were consonant-final, we cannot exclude that the low acceptability mean rating of the *ity*-condition might have been due to violations of this morpho-phonotactic generalization rather than violations to syntactic restrictions. Further testing is required to clarify the role and the nature of such a morpho-phonotactic constraint in the acceptability of these forms, especially since the *ity*-suffixed items were all judged phonologically grammatical. At any rate, if we entertain this possibility, we cannot exclude that the *ity*-ratings might have been caused by the aforementioned (and uncontrolled-for) morpho-phonotactic constraint, which has a crucial impact on the interpretation of the results reported above. When being factored in, it suggests that all target items – including the *ity*-suffixed items – may be possible grammatical words, which therefore may be not apt for our investigation on the syntactic sensitivity of decomposition.

In the remainder of the chapter, we will entertain both interpretations of the rating results above for the interpretation of the priming effects onto syntactically illicit forms (sec. 5.4). We turn to address the second property of English affixation namely, selective decomposition.

5.3. Experiment 5 – licit/root priming

The results of experiments 2 & 3 in chapters 3 & 4 unexpectedly revealed that visual affix priming did not occur systematically for all affixes. In light of these results, experiment 5 tests which of the four affixes tested in this chapter decomposes and primes. To this end, we elicited root priming in bimorphemic real (i.e., syntactically licit) words suffixed with the four suffixes tested in this chapter (namely, *-able, -ity, -ment, -ness*) in separate conditions. In addition to the four conditions above, an additional condition was tested, in which *ful*-suffixed primes were presented before the corresponding target roots (*ful*-condition: e.g., *successful-SUCCESS*). This condition was an additional morphologically-transparent, single-affix condition, which will be used as the control condition for licit morphological priming in experiment 6 reported below.

All four models of decomposition here considered expect all words to prime and therefore decompose, regardless of the affix used. In these models, decomposition relies on letter clustering on the basis of low-level orthographic comparison (the bin model) or morpho-orthographic statistical regularities (the race, full-decomposition, and morpho-orthographic models). Our results indeed show significant priming effects for all conditions, though the effects for the *able*- and *ness*-conditions were larger than the effects for the *ity*- and *ment*-conditions.

5.3.1. Materials

One-hundred and forty-four word pairs were selected from English Lexicon Project (ELP: Balota et al., 2007) database, and organized in the following six conditions of 24 pairs each.

1. Pairs in the *identity condition* included monomorphemic words that were presented both as prime and target (e.g., *thumb-THUMB*).
2. Pairs in the *ful condition* involved *ful*-suffixed words as primes and the corresponding roots as targets (e.g., *successful-SUCCESS*).
3. Pairs in the *able condition* involved *able*-suffixed words as primes and the corresponding roots as targets (e.g., *detectable-DETECT*).
4. Pairs in the *ity condition* involved *ity*-suffixed words as primes and the corresponding roots as targets (e.g., *purity-PURE*).
5. Pairs in the *ment condition* involved *ment*-suffixed words as primes and the corresponding roots as targets (e.g., *settlement-SETTLE*).
6. Pairs in the *ness condition* involved *ness*-suffixed words as primes and the corresponding roots as targets (e.g., *weakness-WEAK*).

The identity condition was added as the baseline priming condition; the *ful*-condition was added as an additional morphologically-transparent condition, which will work as baseline of licit morphological priming effects in experiment 6. We ensured that both primes and targets maintained the same syntactic category within the same conditions. Moreover, all of the primes were bi-morphemic words. Targets were matched as closely as possible in frequency and length. Similarly, primes were matched as closely as possible in frequency and orthographic length, though the inherent orthographic differences between suffixes could not be circumvented. The primes of the identity condition were excluded from the analyses of the primes. The table below shows the mean values of the relevant lexical properties of the items used in each condition.

One hundred and forty-four unrelated prime words were selected for each target word; these words were orthographically, morphologically, and semantically unrelated to targets and were matched as closely as possible on frequency ($t(143)$=.005, p=.99; $BF_{0,1}$=10.77) and length ($t(143)$=.32, p=.74; $BF_{0,1}$=10.24). By-condition statistical results (Table 5.6) revealed significant effects of frequency in the identity condition and marginal effects of length in the *ment*-condition; we will control for this during analysis of the results by adding prime frequency and prime length as covariates in the model. All unrelated primes were suffixed (bimorphemic) words. The stimuli used in the experiment can be found in Appendix V.

PROPERTY		CONDITIONS						STATISTICAL ANALYSES
		identity	able	ful	ity	ment	ness	
PRIMES	HAL frequency (log10)	–	6.37	6.55	7.00	7.48	6.66	$F(4,115)=2.00, p=.1$ $\mathrm{BF}_{0,1}=2.05$
	length	8.25	9.12	8.17	7.83	9.67	8.88	$F(4,115)=12.38, p=<.0001$ $\mathrm{BF}_{0,1}<0.0001$
TARGETS	HAL frequency (log10)	8.74	9.18	9.11	8.45	8.17	9.01	$F(5,138)=1.50, p=.19$ $BF_{0,1}=4.80$
	length	5.12	5.42	5.17	5.12	5.38	5.67	$F(5,138)=1.96, p=.09$ $BF_{0,1}=2.2$

Table 5.5.: Experiment 5 – licit/root priming. Summary of the lexical properties of the stimuli used.

Property	Conditions		
	identity	*ful*	*able*
HAL frequency (log10)	$t(23)=-3.99, p=.0005$ $BF_{0,1}=0.01$	$t(23)=-0.009, p=.99$ $BF_{0,1}=4.66$	$t(23)=-0.003, p=.99$ $BF_{0,1}=4.66$
length	$t(23)=0.77, p=.44$ $BF_{0,1}=3.54$	$t(23)=-1.86, p=.07$ $BF_{0,1}=3.39$	$t(23)=-0.84, p=.41$ $BF_{0,1}=1.08$

Property	Conditions		
	ity	*ment*	*ness*
HAL frequecy (log10)	$t(23)=0.001, p=.99$ $BF_{0,1}=4.66$	$t(23)=0.16, p=.87$ $BF_{0,1}=4.60$	$t(23)=-0.003, p=.99$ $BF_{0,1}=4.66$
length	$t(23)=-0.26, p=.79$ $BF_{0,1}=4.51$	$t(23)=1.9, p=.07$ $BF_{0,1}=1$	$t(23)=1.39, p=.17$ $BF_{0,1}=2$

Table 5.6.: Experiment 5 – licit/root priming. Within-condition statistical results of the lexical properties of related and unrelated primes across the five conditions.

Finally, a set of 144 pseudo-word targets were chosen, in order to match the orthographic length of the word targets as much as possible ($F(6,281)=2.08, p=.06$; $BF_{0,1}=2.74$); they were preceded by unrelated suffixed word primes that were not used in the experiment.

Target words from each condition were counterbalanced and divided at random into two versions of equal number of pairs. In each version, half of the target words were preceded by a related prime word and half by an unrelated prime word. Participants received only one version of the word list, so that they saw each target word exactly once.

5.3.2. Participants & procedure

One hundred and thirty-eight participants (48 females, 91 males; mean age: 36.30, s.d: 10.49) were recruited through Amazon Mechanical Turk and received monetary compensation for their participation. They were all native speakers of American English.

Stimulus presentation and data recording were performed on-line through PsychoJS (Peirce et al., 2019), the Javascript equivalent of PsychoPy, on the Pavlovia platform (`www.pavlovia.org`). The subjects recruited online were asked to read the capitalized letter strings on the display and decide as quickly and accurately as possible whether or not each string is a word. To do that, they used their own monitor and keyboard. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case for 33 ms;[19] the target word immediately followed in uppercase, and remained on the screen until a response was made. The order of pairs was chosen randomly across participants.

Participants were given 10 practice pairs before the actual experiment began. A total of 288 pairs were presented to each participant. During the experiment, participants were also given the possibility to take 6 brief breaks. To detect bots, subjects were also asked to answer three open-ended questions immediately after a break. To help subjects refocus on the main task post-break, we also ensured that first five trials presented after each break were pseudo-word trials.

5.3.3. Predictions

First, we expect to find ceiling priming effects in the identity condition. A lack of priming effects in the identity condition may potentially suggest a methodological problem with the design, which would therefore impede any further interpretation of the results. Should priming be found in the identity condition, the priming effects of the remaining conditions might then be interpreted.

Regarding the remaining conditions, we expect them to all trigger priming, as indeed predicted by the models of decomposition considered here (Figure 5.3, scenario 1). If priming is found in none of the single-affix conditions tested (Figure 5.3, scenario 2), we would not be able to put forward any theoretical interpretation in the terms being sketched below. This might also suggest that priming is suppressed when the same affix is presented in the same condition. Such an effect might be due to the suppression effect that is triggered in pairs of primes and targets that are phonologically/orthographically overlapping in their onsets – much like orthographic and phonological suppression effects that have been widely reported in the literature (see sec. 1.3).

---

[19] Notice that in the online version of the script durations of the mask and the prime were purposefully coded in milliseconds rather than in number of frames/flips in order to ensure compatibility across different monitors with possibly different refresh rates.

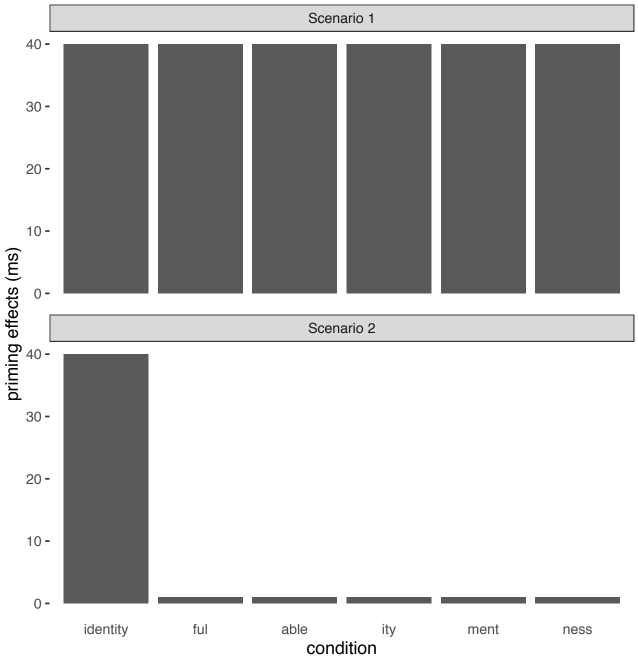

Figure 5.3.: Experiment 5 - licit/root priming. Predicted results.

5.3.4. Results

Response times (RTs) were measured from target onset and cleaned of outliers according to the following criteria. First, we calculated by-subject error rates for words and pseudo-words separately. Given that the means of the two distributions did not vary significantly ($t(221.34)=1.75$, $p=.08$; $BF_{0,1}=1.74$), we calculated by-subject overall error rates (that is, including words and pseudo-words) and removed all subjects whose error scores were higher than 20%. Second, items were excluded from the analysis if their overall error rate was higher than 30%. Incorrect responses and fillers (word and pseudo-word) were excluded from analysis. Finally, RTs were first log-transformed to guarantee near-Gaussian distribution as suggested by (Baayen, 2008); then, individual log RTs were excluded from the analysis if they were more than 2.5 standard deviations away from the by-subject and overall log mean RT. Outlier rejection resulted in excluding a total of 342 datapoints (3.47% of the dataset). A total of 9,514 datapoints were included in the analysis.

The plot (Figure 5.4) seems to suggest that all conditions triggered priming effects. We then constructed a series of linear mixed-effect regression (LMER) models (Baayen, 2008; Barr et al., 2013) for each sub-design. Each model had *log RT* as the dependent variable, RELATEDNESS (2 levels: related vs. unrelated) as the fixed factor, and SUBJECT and ITEM as random factors (intercept only); for some of the designs, we also added the lexical properties that could not be controlled for in the materials tested (see sec. 5.3.1). P-values were estimated using the Satterthwaite approximation of degrees of freedom (using the `lmerTest` R-package). For each subdesign, we also estimated Bayes Factors using the R package `BayesFactor` (Morey, 2018).

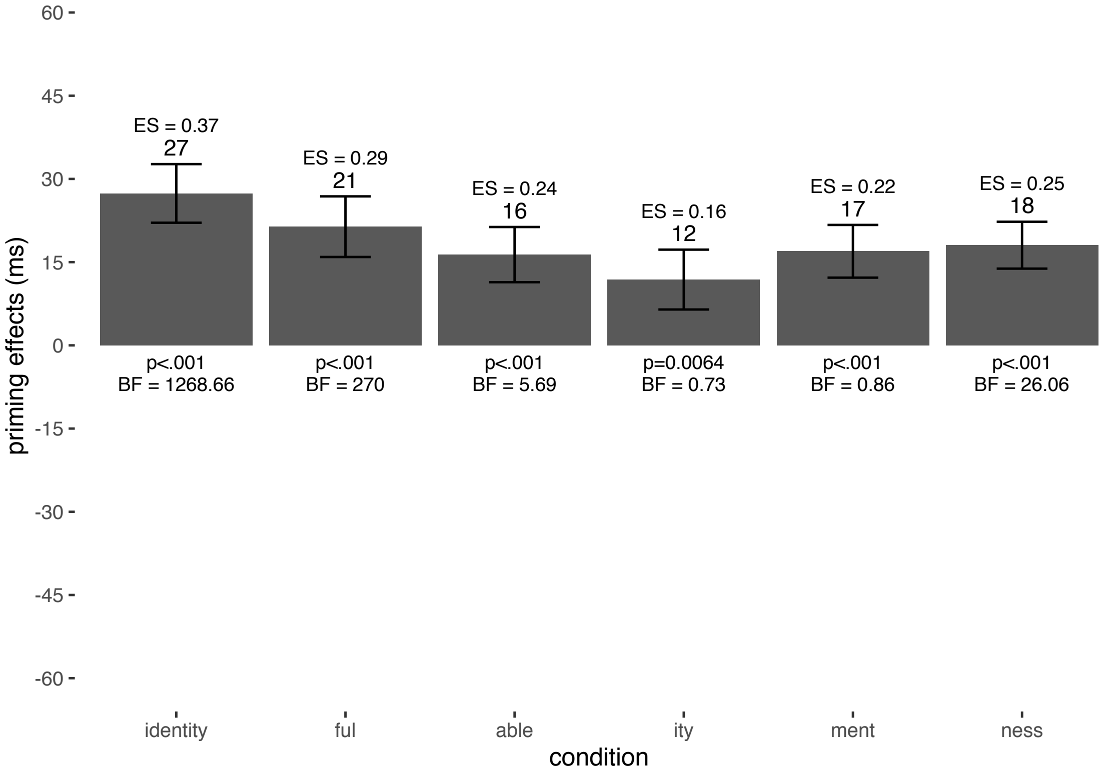

Figure 5.4.: Experiment 5 - licit/root priming. Summary of the priming effects. The numbers over the bars are the priming magnitudes and Cohen's *d* (effect sizes, ES); the numbers below the bar are the *p*- and $BF_{1,0}$-values of the LMER and BF analyses, respectively.

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
identity	623	596	27	0.37
ful	607	586	21	0.29
able	616	600	16	0.24
ity	625	613	12	0.16
ment	632	615	17	0.22
ness	593	611	18	0.25

Table 5.7.: Experiment 5 - licit/root priming. Mean RTs, priming effects, and Cohen's *d* (ES, effect sizes).

For the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.2. The details of the LMER and the BF analyses are reported in Table 5.8 below. Both analyses were consistent in suggesting that the priming effects were significant in following conditions: identity, ful, able, and ness. While the LMER analysis suggested that the priming effects for the *ity*- and *ment*-conditions were significant, the BF analysis instead revealed that the data for the two conditions anectodally supported the null hypothesis, while approximating to neutrality ($0.33 < BF_{1,0}$s≈ 1).

CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$interpretation	COVARIATES ADDED
identity	36.24	<.0001	1268.66	extreme for H_1	prime frequency
ful	28.53	<.0001	270	extreme for H_1	–
able	15.25	.0001	5.59	substantial for H_1	–
ity	7.46	.006	0.73	anectodal for H_0	–
ment	13.14	.0003	0.86	anectodal for H_0	prime length
ness	17.08	<.0001	26.06	strong for H_1	–

Table 5.8.: Experiment 5 – licit/root priming. Summary of the statistical results.

Pairwise comparisons of priming magnitudes were also performed across conditions. Table 5.9 below reports both Dunn-corrected p-values and uncorrected $BF_{1,0}$-values for each combination. None of the comparisons were significant.

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
identity	*able*	1	0.48
identity	*ful*	1	0.22
identity	*ity*	.46	1.19
identity	*ment*	1	0.45
identity	*ness*	1	0.40
able	*ful*	1	0.21
able	*ity*	1	0.20
able	*ment*	1	0.18
able	*ness*	1	0.17
ful	*ity*	1	0.25
ful	*ment*	1	0.21
ful	*ness*	1	0.19
ity	*ment*	1	0.22
ity	*ness*	1	0.25
ment	*ness*	1	0.18

Table 5.9.: Experiment 5 – licit/root priming. Pairwise comparisons of the priming effects across conditions.

5.3.5. Discussion

The primary goal of the experiment was to assess whether or not any of the affixed word forms that were tested decompose. While all current models of decomposition expect decomposition to occur regardless of the affix used, we take this issue on in light of the results reported in the previous chapters. If decomposition were to occur selectively for a subset of the affixes of the language, it may be a confounding factor in the investigation undertaken in this chapter. If any of the affixes being tested here (i.e., *-able, -ity, -ment, -ness*) did not decompose in syntactically licit forms in the first place, we would be interpret the potential effects elicited for the corresponding illicit forms accordingly. To this end, experiment 5 elicited root priming in five different morphologically-transparent conditions, each testing one of the four suffixes tested in this experiment (i.e., *-able, -ity, -ment*, and *-ness*), plus the suffix *-ful*, which was used as the control condition for licit morphological priming in experiment 6. Our statistical analyses reported a mixed pattern. While the LMER analysis showed all conditions having significant priming effects (*p*s<.05), the BF analysis revealed a mixed pattern: the *identity, ful-, able-,* and *ness*-conditions strongly supported the alternative hypothesis ($BF_{1,0}$s> 3), while the *ity-* and *ment*-conditions anectodally supported the null hypothesis (with $BF_{1,0}$approximating 1, i.e. neutrality between the two hypotheses). This seems to indicate that the *ity-* and *ment*-conditions did not not provide enough data to detect their small effect sizes.

Accordingly, both the LMER and BF statistical analyses clearly suggest that the affixes *-ful, -able, -ment,* and *-ness* prime and decompose. However, the neutral result in the BF analysis in the *ity-* and *ment*-conditions suggests smaller effects for these conditions. Nonetheless, the BF results were paired with the highly significant *p*-values calculated in the LMER analysis, which leads us to cautiously claim that decomposition does occur for the affixes *-ity* and *-ment* as well. It is indeed possible that the neutral $BF_{1,0}$'s were due to the noise brought about by the online data collection methodology (see sec. 2.3). We therefore conclude that these results confirm that all affixes tested decompose. While being aware of potential difference in the priming response to specific affixes due to (possibly uncontrollable) external factors, we finally turn to explore the potential impact of syntactic violations onto decomposition.

5.4. Experiment 6 – illicit/root priming

In the previous two sections, we explored two properties of English affixation that might influence the sensitivity of decomposition to syntactic affixal restrictions. Experiment 4 explored the extent to which the operation of zero-derivation makes syntactically violating forms (e.g., *blissable*) grammatical. Experiment 5 addressed the issue that some affixes do not decompose in grammatical words (e.g., *detectable*) at all. First, we found that *able-, ment-, and ness*-suffixed illicit forms were rated higher than *ity*-suffixed illicit forms. One possible conclusion for these results suggests that the the *able-, ment-,* and *ness*-suffixed illicit forms are actually grammatical to native speakers, but *ity*-suffixed forms are not. An alternative conclusion is that the low acceptability ratings for the *ity*-suffixed items were due to the fact that most *ity*-suffixed items violate the affix-specific morpho-phonotactic constraint whereby *ity-* only attaches to vowel-final roots. Second, we found that all affixes tested decompose, although *able*- and *ness*-suffixed words were found to trigger larger priming effects than *ity*- and *ment*-suffixed words. In experiment 6, we elicited root priming of the same target roots we used in experiment 5, but preceded by syntactically illicit affixed pseudo-words as arranged in four experimental conditions, each involving one of the four affixes explored in this chapter (i.e., *-able, -ity, -ment, -ness*). We also tested two additional conditions: an identity condition, in which the same bare root word was presented as both prime and target and a morphologically-transparent condition, in which syntactically licit *ful*-suffixed prime words were paired with the corresponding root targets (*ful*-condition; e.g., *graceful-GRACE*). These two conditions are used as baseline conditions to compare the priming effects of the experiment conditions.

The predictions and the results of this experiment build on the results of experiments 4 & 5 just summarized. In particular, the results of experiment 4 open to two possible interpretations for the results of this experiment. If we consider the low *ity*-ratings as indicative of ungrammaticality, we may use the *ity*-suffixed illicit forms (and no other forms) to explore the question of the syntactic sensitivity of decomposition. If decomposition is affected by syntactic violations, *ity*-priming is expected to be inhibited. If we instead consider the low *ity*-ratings as resulting from violations to the uncontrolled-for morpho-phonotactic selectional constraint mentioned above, we may then have to claim that all syntactically illicit affixed forms may be rated as plausible words by native speakers, thus making it impossible for us to rule out zero-derivation as a potential confound in our investigation. Our results show significant priming effects for the *ity*- and *ness*-conditions, but no priming for the *able*- and *ment*-conditions. Though not straightforwardly interpretable, these findings seem to suggest that decomposition occurs for forms that are non-words possibly due to affix-specific violations of syntactic restrictions.

5.4.1. Materials

One hundred and two roots were chosen from the set of roots used in the rating experiment (experiment 4). All roots were disyllabic or monosyllabic nouns; for further information, see sec. 5.2.1. The roots were used as target words that were preceded by one of the following six types of primes (with 17 pairs each): (1) identical primes (*identity condition*; e.g., *skit-SKIT*); (2) *able*-suffixed primes (able-*condition*; e.g., *skittable-SKIT*); (3) *ity*-suffixed primes (ity-*condition*; e.g., *skitity-SKIT*); (4) *ment*-suffixed primes (ment-*condition*; e.g., *skitment-SKIT*); (5) *ness*-suffixed primes (ness-*condition*; e.g., *skitness-SKIT*); and (6) unrelated primes (*unrelated condition*; e.g., *trainee-SKIT*). The root targets of these conditions were divided into 6 lists and arranged in a Latin-Square design so that in each list one-sixth of the roots appeared in exactly one condition. The Latin Square design ensured that the lexical properties of the targets were naturally controlled over the six lists (see sec. 1.2); for this reason, no statistical analysis of the lexical properties is provided for the conditions above.

A seventh condition was added outside of the Latin Square arrangement described above. In this condition, root targets were preceded by the corresponding *ful*-affixed word primes (e.g., *graceful-GRACE*; (licit) ful-*condition*). The purpose of this condition was to be the baseline for syntactically licit priming effects. The targets of the *ful*-condition were matched with the targets of the other conditions in length ($F(6,112)=0.46$, $p=.83$; $BF_{0,1}=69.09$). However, due to the overall low number of *ful*-suffixed words in the HAL corpus, it was not possible to match the frequency of the targets of the *ful*-condition with the target of the other conditions ($F(7,162)=2.61$, $p=.01$; $BF_{0,1}=0.25$). For the same reason, it was not possible to control for syntactic category across conditions. The thirty-four related primes of the *ful*-condition were then matched with an unrelated, bimorphemic prime in frequency ($t(33)=-0.01$, $p=.98$; $BF_{0,1}=5.44$) and length ($t(33)=-0.13$, $p=.89$; $BF_{0,1}=5.39$). All primes (related and unrelated) of the *ful*-subdesign were also matched with the unrelated condition above in frequency ($F(2,82)=0.01$, $p=.98$; $BF_{0,1}=9.02$) and length ($F(2,82)=0.10$, $p=.89$; $BF_{0,1}=8.44$). The target words in the *ful*-condition were then counterbalanced and divided at random into two versions of equal numbers of pairs. In each of these versions, half of the target words were preceded by a related prime word and half by an unrelated prime word. Either version of the *ful*-condition was added to each of the six Latin-Square lists created with the six conditions described above; each resulting wordlist had 136 word pairs. All of the stimuli used in the experiment can be found in Appendix VI.

Finally, a set of 136 pseudo-word targets were chosen, in order to match the orthographic length of the word targets as much as possible ($F(8,297)=0.43$, $p=.90$; $\mathrm{BF}_{0,1}=89.56$); they were preceded by unrelated suffixed word primes that were not used in the experiment. Each participant was randomly assigned a wordlist, so that they saw each target word exactly once (for a total of 272 words).

5.4.2. Participants & procedure

One hundred and forty-one participants (59 females, 82 males; mean age: 36, s.d: 9.71) were recruited through Amazon Mechanical Turk and received monetary compensation for their participation. They were all native speakers of American English.

Stimulus presentation and data recording were performed on-line through PsychoJS (Peirce et al., 2019), the Javascript equivalent of PsychoPy, on the Pavlovia platform (`www.pavlovia.org`). The subjects recruited online were asked to read the capitalized letter strings on the display and decide as quickly and accurately as possible whether or not each string is a word. To do that, they used their own monitor and keyboard. Each prime was preceded by a 500ms-long forward mask (#######) and presented in lower case for 33 ms. The target word immediately followed in uppercase, and remained on the screen until a response was made. The order of pairs was chosen randomly across participants.

Participants were given 10 practice pairs before the actual experiment began. A total of 288 pairs were presented to each participant. During the experiment, participants were also given the opportunity to take 7 brief breaks. To detect bots, subjects were also asked to answer three open-ended questions immediately after a break. To help subjects refocus on the main task post-break, we also made sure that first five trials presented after each break were pseudo-word trials.

5.4.3. Predictions

First, we expect to find ceiling priming effects in the identity condition. A lack of priming effects in the identity condition suggests a methodological problem with the design, which would therefore prevent any further interpretation of the results. Priming effects in the *ful-* (licit) condition are also expected, as also reported in experiment 5. The predictions for the remaining conditions build on the results of the previous experiments reported in this chapter. In particular, given the results of experiment 4, two potential interpretations of the results are possible. First, if the low *ity*-ratings are considered indicative of ungrammaticality, the *ity*-condition can inform us on the potential sensitivity of decomposition to syntactic affixal restrictions (Fig. 5.5A). If *ity*-priming arises, it may suggest that decomposition is not affected by syntactic affixal restrictions. This is also what all models of decomposition considered here predict, as they assume decomposition to occur before accessing any lexical property (including syntactic affixal restrictions; see sec. 1.4.2). If *ity*-priming does not arise, it may instead suggest that decomposition is affected by violations to the syntactic affixal restrictions of the affix *-ity*. The priming response elicited in the remaining conditions can instead inform us on the decomposition of non-lexicalized, grammatical forms (Fig. 5.5B). If we find priming effects in all of the conditions, it would suggest that decomposition occurs regardless of lexicality (i.e., whether or not a given form is a real word of the language), as also predicted by models of decomposition. If we find no priming effects in the conditions, it would suggest that decomposition is indeed affected by stimulus lexicality.

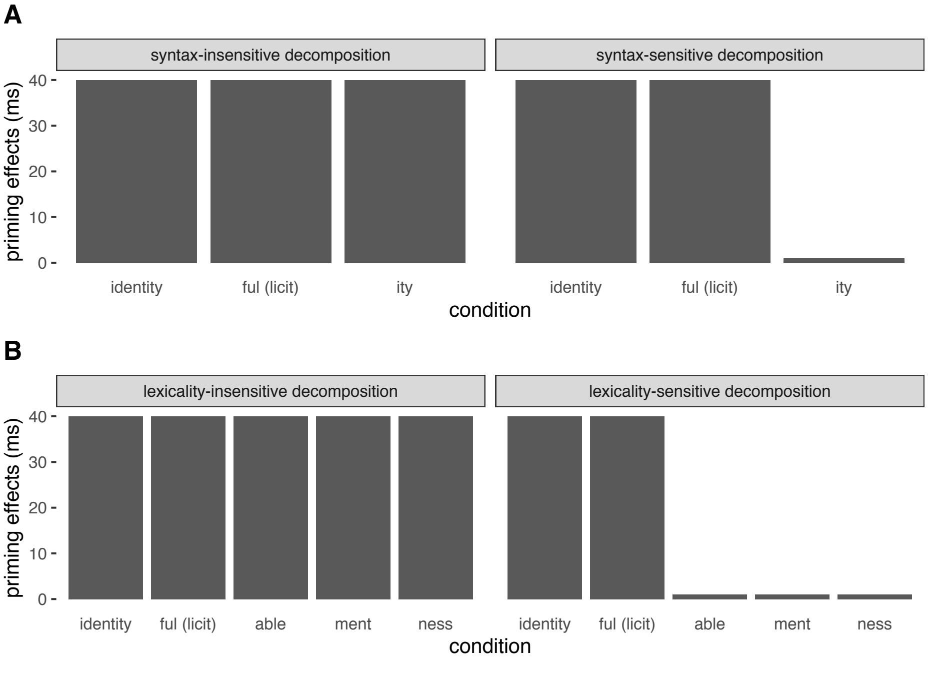

Figure 5.5.: Experiment 6 – illicit/root priming. Predicted results.

Second, if the low *ity*-ratings are considered to result from violations to the *ity*-specific morpho-phonotactic constraint whereby *ity-* only attaches to vowel-final roots, the potential effect of zero-derivation in decomposition cannot be ruled out in any of remaining conditions tested. As a consequence, all conditions tested are predicted to prime and decompose, as a result of either/both the syntactic insensitivity of decomposition or/and zero-derivation. If, however, priming does not arise in any of the conditions tested, it may suggest that zero-derivation does not play a role in decomposition, which may, therefore, argued to be sensitive to syntactic affixal restrictions.

5.4.4. Results

Response times (RTs) were measured from target onset and cleaned of outliers. First, we calculated by-subject error rates for words and pseudo-words separately. Since the means of the two distributions did not vary ($t(251.02)=0.74, p=.45$; $BF_{0,1}=5.88$), we calculated by-subject overall error rates (that is, including words and pseudo-words) and removed all subjects whose error scores were higher than 20%. Second, items were excluded from the analysis if their overall error rate was higher than 30%. Incorrect responses and fillers (word and pseudo-word) were excluded from analysis. Finally, RTs were first log-transformed to guarantee near-Gaussian distribution as suggested by (Baayen, 2008); then, individual log RTs were excluded from the analysis if they were more than 2.5 standard deviations away from the by-subject and overall log mean RT. Outlier rejection resulted in excluding a total of 531 datapoints (4.47% of the dataset). A total of 11,336 datapoints were included in the analysis.

(a) Latin-Square arrangement.

CONDITION	MEAN RT	PRIMING	ES
identity	607	28	0.7
able	625	10	0.26
ity	619	16	0.45
ment	625	9	0.26
ness	619	16	0.46
unrelated	635	–	–

(b) Self-contained subdesign arrangement.

CONDITION	MEAN RT		PRIMING	ES
	unrelated	*related*		
ful	622	598	24	0.64

Table 5.10.: Experiment 6 – illicit/root priming. Mean RTs, priming effects and Cohen's *d* (ES, effect sizes).

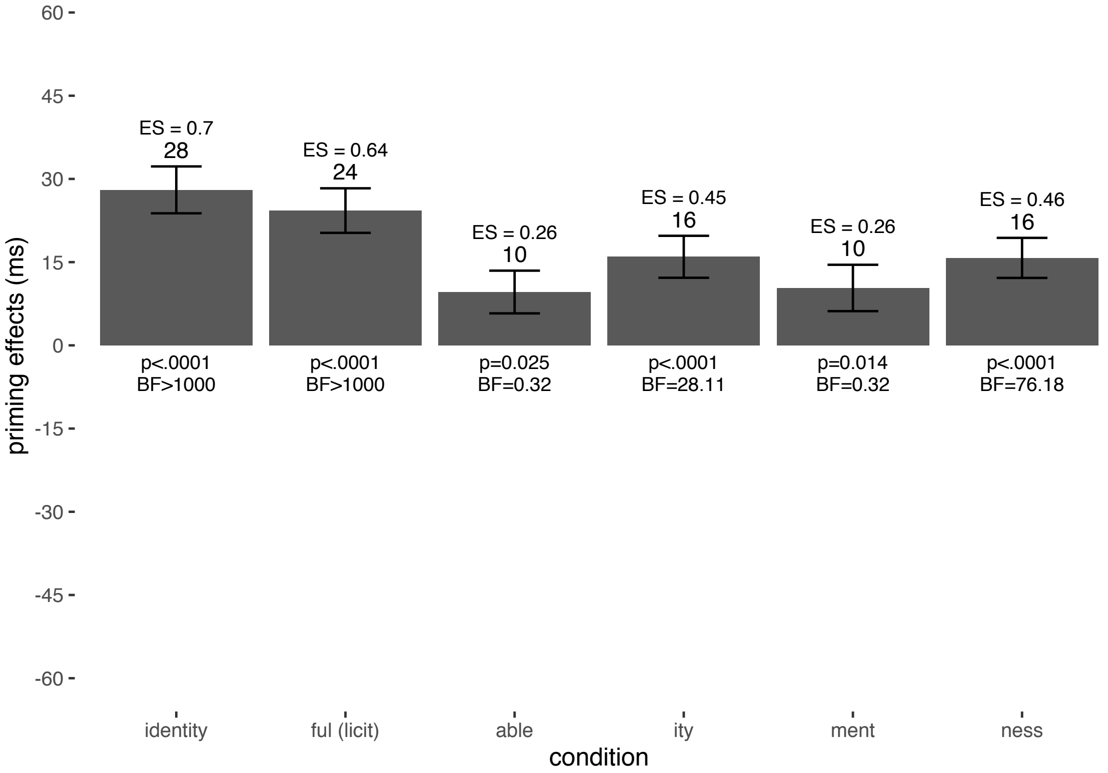

Figure 5.6.: Experiment 6 – illicit/root priming. Summary of the priming effects. The numbers over the bars are the priming magnitudes and Cohen's *d* (effect sizes, ES); the numbers below the bar are the *p*- and $BF_{1,0}$-values of the LMER and BF analyses, respectively.

Priming effects were calculated depending on the design adopted for each condition (see sec. 5.4.1). The identity, *able-, ity-, ment-* and *ness*-conditions were constructed in a Latin-Square arrangement; therefore, the priming response to each of these conditions was calculated by subtracting the average RT to each condition from the unrelated condition. The *ful*-condition was instead constructed within a self-contained arrangement; thus, the *ful*-priming response was calculated by subtracting the average RT to the related pairs from the average RT to the unrelated pairs within the same condition (see Table 5.10). As Figure 5.6 shows, all conditions seem to trigger priming (though the magnitude and effect sizes seem to vary across conditions).

We therefore constructed a total of 6 separate linear mixed-effect regression (LMER) models, one for each condition (Baayen, 2008; Barr et al., 2013). The models for the identity, *able-, ity-, ment-*, and *ness*-conditions had *log RT* as the dependent variable, CONDITION (2 levels: related and unrelated) as the fixed factor, and SUBJECT and ITEM as random effects (intercepts only). The model for the *ful*-condition had *log RT* as the dependent variable, RELATEDNESS (2 levels: related vs. unrelated) as the fixed factor, and SUBJECT and ITEM as random factors (intercept only). P-values were estimated using the Satterthwaite approximation of degrees of freedom (using the `lmerTest` R-package). For each subdesign, we also estimated Bayes Factors using the R package `BayesFactor` (Morey, 2018). For the interpretation of BFs, we refer to the interpretive table suggested by Jeffreys (1961) and reported in sec. 1.2. The details of the LMER and the BF analyses are reported in Table 5.11 below. In the LMER analysis, the priming effects of all conditions were found to be significant. However, the BF analysis revealed that only the identity, *ful-* (licit), *ity-*, and *ness*-conditions supported the alternative hypothesis. The *able-* and *ment*-conditions instead substantially support the null hypothesis (able: $BF_{1,0} < 0.33$).

CONDITION	F	p	$BF_{1,0}$	$BF_{1,0}$interpretation
identity	59.96	<.0001	>10000	extreme for H_1
ful (licit)	67.58	<.0001	>10000	extreme for H_1
able	5.03	.025	0.32	substantial for H_0
ity	21.68	<.0001	28.11	strong for H_1
ment	6.11	.01	0.32	substantial for H_0
ness	20.44	<.0001	76.18	strong for H_1

Table 5.11.: Experiment 6 – illicit/root priming. Summary of the statistical results.

Pairwise comparisons of priming magnitudes were also performed across conditions. Table 5.12 below reports both Dunn-corrected p-values and uncorrected $BF_{1,0}$-values for each combination. The Dunn-corrected t-analysis found significance in the following comparisons: identity vs. able and identity vs. ment. The uncorrected BF-analysis suggested that the following comparisons supported the alternative hypothesis ($BF_{1,0}$s>3): identity vs. able, identity vs. ment, and able vs. ful. The following comparisons anecdotally supported the alternative hypothesis (1<$BF_{1,0}$s<3): identity vs. ity, identity vs. ness, ful vs. ment, and ful vs. ness. The following comparisons anecdotally supported the null hypothesis (1>$BF_{1,0}$s>0.33): ful vs. ity and ful vs. ness. All remaining comparisons substantially supported the null hypothesis ($BF_{1,0}$s<0.33).

5.4.5. Discussion

Experiment 6 was designed to explore the main question of the chapter: the syntactic sensitivity of decomposition. To this end, we elicited the priming response to a total of six conditions. The identity and *ful*-conditions were the baseline conditions for identity priming and morphological priming, respectively. The remaining four conditions consisted of syntactically illicit words in which an existent bare noun root was affixed with one of the four affixes tested in this chapter (i.e., *-able, -ity, -ment, -ness*). The statistical results were different depending on the analysis. In the LMER analysis, all conditions were found to significantly reject the null hypothesis. In the BF analysis, the identity, *ful-* (licit), *ity-* and *ness*-conditions strongly supported the alternative hypothesis, while the *able-* and *ment*-conditions substantially supported the null hypothesis. As per our strict experimental methodology (see sec. 1.5), we deal with the contradictory results of the two statistical analyses performed on the *able-* and *ment*-conditions

CONDITION 1	CONDITION 2	Dunn-corrected p	uncorrected $BF_{1,0}$
identity	*able*	.016	18.43
identity	*ful*	1	0.19
identity	*ity*	.47	1.31
identity	*ment*	.025	9.32
identity	*ness*	.43	1.53
able	*ful*	.13	3.96
able	*ity*	1	0.31
able	*ment*	1	0.16
able	*ness*	1	0.30
ful	*ity*	1	0.46
ful	*ment*	.19	2.33
ful	*ness*	1	0.51
ity	*ment*	1	0.25
ity	*ness*	1	0.16
ment	*ness*	1	0.25

Table 5.12.: Experiment 6 – illicit/root priming. Pairwise comparisons of the priming effects across conditions.

by deferring to the BF analysis, which suggests that the data substantially supported the null hypothesis (H_0, namely the absence of priming). As a consequence, we lean towards claiming that priming did not arise in the *able-* and *ment*-conditions.

The interpretation of these results strictly hinges on the results of the previous two experiments. While experiment 5 confirmed that all of four affixes decompose (see sec. 5.3.5), experiment 4 reported ambiguous results. The main issue regards the intepretation of the low mean rating of the *ity*-condition, for which we entertain two possible hypotheses. The first hypothesis suggests that *ity*-suffixed illicit forms (e.g., **skittity*) are ungrammatical, and, therefore, may index the potential sensitivity of decomposition to syntactic affixal restrictions. As such, the significant priming effects in the *ity*-condition suggests that decomposition occurs regardless of the violations to the syntactic constraints associated with the suffix *-ity*. This is also expected in all models of decomposition considered in this dissertation. When an *ity*-suffixed illicit form is presented as a prime (e.g., *skittity*), it is decomposed so that the root (*skit*) is separated from the suffix (*ity*); when the corresponding root is presented as a target (*SKIT*), it is recognized faster and priming arises. The alternative hypothesis suggests that the *ity*-suffixed illicit forms were rated lower than the remaining forms because of an uncontrolled-for morpho-phonotactic selectional restriction on *ity*-affixation (see sec. 5.2.5). Therefore, under this hypothesis, we cannot completely rule out the potential impact of zero-derivation in the acceptability and, for our current purposes, in the decomposition of the forms tested.

Finally, we point out that the lack of priming effects to the *able-* and *ment-*conditions remains unclear, regardless of the interpretive hypothesis being adopted for the *ity-*ratings. Under both hypotheses, both conditions are expected to prime, similarly to the syntactically licit counterparts tested in experiment 5. The statistical results of the two conditions are also contradictory: the LMER analysis indicates the presence of significant priming effects ($p < .05$), whereas the BF analysis indicates that the data substantially supported the null hypothesis ($BF_{1,0} < .033$). If we strictly stuck to our experimental methodology (sec. 1.5), we would have to rely on the BF analysis and therefore interpret the data for the *able-* and *ment-*conditions as eliciting no priming. However, we point out that the contradictory results in two statistical tests might have been brought about by the high amount of noise in the data – a known side effect of the online data collection methodology (see sec. 2.3).[20] We therefore chose to temporarily suspend judgment on the matter. We expect to carry out in-lab replications of this experiment (as well as of all of the other experiments reported in this dissertation) in the near future, which will help clarify the source of the contradictory statistical results reported here.

5.5. General discussion

This chapter explored the extent to which early visual decomposition is sensitive to syntactic affixal restrictions. It is well-known that across languages affixes do not attach to roots freely; rather, they adhere to specific restrictions regarding the syntactic category of the attaching root. These restrictions are usually considered to be part of the lexical properties of the affixes; for example, *-able* and *-ment* attach to verbal stems, *-ity* and *-ness* attach to adjectival stems, but none of them attach to nominal nouns. All of the four models of decomposition considered in this dissertation expect decomposition to not be sensitive to such restrictions (see sec. 1.1). The experiments reported in this chapter were designed to test this prediction by eliciting root priming from syntactically illicit non-words such as **blissable, *blissment, *blissity*, and **blissness*.

Before directly looking at decomposition of these syntactically violating forms, we first identified and addressed two properties of English affixation that might potentially hinder an investigation on the question above. The first property concerned the fact that English derivation may involve the phonologically-null ø-morpheme with no apparent grammatical restriction, which could potentially make syntactically violating forms (e.g., *blissable*) grammatical, and therefore legitimately decomposable new formations, akin to syntactically licit words (e.g., *detectable*). We tested this in the rating experiment reported in sec. 5.2 (experiment 4), in which we asked subjects to rate illicit forms as possible new word formations. Results showed that all forms were phonologically grammatical, though *able-, ment-,* and *ness-*suffixed illicit forms were rated higher than *ity-*suffixed illicit forms. We interpreted these differential effects as indicating two possible conclusions. On one hand, the low ratings for the *ity-*suffixed items may suggest that *able-, ment- and ness-*suffixed illicit forms were grammatical, but *ity-*

---

[20] Relatedly, we note that the inconsistent $BF_{1,0}$'s for the *ity-*conditions of experiment 5 ($BF_{1,0} \approx 1$) and experiment 6 ($BF_{1,0} > 10$) may be indicative of an accidental increase of noise in the former experiment due to the online data collection methodology (see sec. 2.3).

suffixed forms are not. Under this conclusion, only the *ity*-suffixed illicit forms could be used to directly test the potential impact of syntactic restriction on decomposition. Alternatively, the low ratings for the *ity*-suffixed items might have been due violations to an *ity*-specific morpho-phonotactic constraint that was not controlled for during material preparation. Under this circumstance, all forms must be considered grammatical, thus suggesting that the potential impact of zero-derivation on the decomposition of these forms could not be fully ruled out.

The second property concerned the argument advanced in previous chapters that decomposition may selectively occur for only a subset of the extant morphemes of the language (see in particular secc. 3.2.5 and 4.2.5). As a consequence, we needed to find out which of the four affixes tested in this chapter, if any, do not decompose when licitly attached to a root word (e.g., *detectable*). We tested this in the masked priming experiment reported in sec. 5.3 (experiment 5). Our results showed priming for all of the four affix conditions tested, although the priming effects of the *able*- and *ness*-conditions were statistically larger than the priming effects of the *ity*- and *ment*-conditions.

Section 5.4 reported experiment 6, which tested whether or not decomposition is affected by violations to syntactic affixal restrictions. Our results showed that priming arose in response to the *ity*- and *ness*-suffixed illicit forms, but did not arise in response to *able*- and ment-suffixed illicit forms. Under the hypothesis that *ity*-suffixed forms are ungrammatical due to violations of syntactical selectional restrictions, the priming effects found for the *ity*-condition suggest that decomposition is not affected by syntactic affixal restrictions. This claim is represented in Figure 5.7, where syntactic restrictions are accessed at post-decomposition stages of processing. While compatible with a distributed view of lexical access, this clear-cut distinction might be due to the fact that both syntactic processing (i.e., the process whereby syntactic categories and affixal restrictions are retrieved and checked) and semantic processing (i.e., the process whereby whole-word meaning is accessed) require at least two mechanisms to occur: (i) recombination of decomposed morphemes (along the lines of Taft, 2004), and, eventually, (ii) access to compatible syntactic and semantic representations. Given that these are fairly complex, these representations indeed take more time to access fully, which, therefore, may not occur until after decomposition.

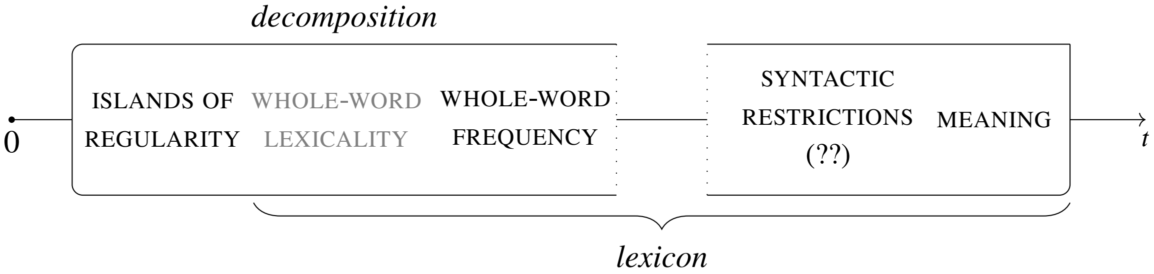

Figure 5.7.: Experiment 6 – illicit/root. Time-course of the unfolding of word information in word processing (updated version after ch. 5). Whole-word lexicality is greyed-out because it was not directly addressed in this dissertation. The question mark symbol '??' refers to provisional claims that need further testing.

Under the hypothesis that *ity*-suffixed forms are, instead, grammatical, the complete priming pattern reported for experiment 6 remains unexplained and calls for further testing. In particular, the differential effects found for the *ity*- and *ness*-conditions (priming) on one hand, and the *able*- and *ment*-conditions (not priming) on the other hand, makes us unable to make a specific claim regarding the question being asked in the chapter. Further research will need to clarify whether the pattern of results reported here is replicated in the in-lab environment or was simply due to the high amount of noise generated by the online data collection environment (see sec. 2.3).

Chapter 6.

Concluding remarks

6.1. Introduction

Building on the wealth of evidence for an early procedure of morphological decomposition (see sec. 1.3), this dissertation has explored the extent to which such a procedure may be sensitive to any linguistic and/or lexical properties, in addition to (Rastle and Davis, 2008)'s morpho-orthographic *islands of regularities*, as maintained by the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3).

(16) EMPIRICAL QUESTION ON DECOMPOSITION
In addition to the morpho-orthographic form of morphemes, what type of linguistic properties, if any, affects decomposition?

In asking (16), this dissertation has also indirectly engaged with the long-standing debate on the mechanism of *lexical access*, and, in particular, the way lexical properties are accessed during visual word recognition (17). While the traditional view of lexical access assumes that all lexical and linguistic properties of morphemes are stored in one play, we have pointed out that current models of decomposition adopt a distributed view, in which different properties may become available at different points in the word recognition process, yet assuming nonetheless the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3) (see sec. 1.4.2).

(17) THEORETICAL QUESTION ON DECOMPOSITION
What does (16) tell us about the unfolding of lexical/linguistic properties in word processing?

To this end, we probed a representative subset of linguistic and lexical properties of morphemes, spanning from phonology, to morpho-phonology, to syntax. In this chapter, we finally take stock of the investigation reported in previous chapters. In Section 6.2, we briefly summarize the results reported in the previous chapters. In Section 6.3, we discuss a new set of questions about decomposition raised by the present investigation, which lays out a research program we expect to undertake in the near future.

6.2. Beyond *islands of regularities*?

Each of the previous chapters has tested the sensitivity of decomposition to one of the four properties chosen to address the questions (16)-(17): phonological properties (syllabification; experiments 1), lexical properties (dominance; experiment 2), morpho-phonological properties (phonologically-conditioned alternations; experiment 3), and syntactic properties (syntactic affixal restrictions; experiments 4-6). In this section, we briefly synthesize the results reported in each of previous chapters (see also Table 6.1) and propose a model of decomposition accordingly.

In chapter 2, we explored the potential sensitivity of decomposition to phono-orthographic SYLLABIFICATION (experiment 1). We entertained this hypothesis as a way to explain the apparent contradiction between the brothel non-effect (no priming for *brothel-BROTH*: Rastle et al., 2004) and the slegrack/flexire effect (priming for *slegrack-RACK*: Fiorentino et al., 2015; and for *flexire-FLEX*: Morris et al., 2011). Crucially, early visual ortho-phonological syllabification would have explained the contradiction, while still maintaining the MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS (3). However, our results suggested that syllabification does not occur during decomposition.

In chapter 3, we explored the decomposition's sensitivity to asymmetries in the number DOMINANCE ratio across English regular plural forms. Lexical decision studies have shown plural forms are recognized faster if if the frequency of the stem form is lower than the frequency of the plural form (plural-dominant forms: e.g., *windows*; Baayen et al., 1997). This finding seems to support Schreuder and Baayen (1995)'s dual-route race model; plural-dominant plural forms are recognized faster than stem-dominant plural forms because the high frequency of the whole-word forms guarantees that the storage route wins the race without having to go through the steps of the parsing route. Crucially, the race model, as well as all of the other models considered here, assumes that all words obligatorily go through the decomposition stage – at least at early stages. Therefore, in all models, dominance effects are not expected to affect decomposition; therefore, masked priming effects should occur in all plural forms, regardless of their dominance ratio. However, this prediction was contradicted by the results (experiment 2), which instead showed dominance-driven priming effects. This finding seems to suggest that whole-word frequency (via dominance asymmetries among forms of the same paradigm) may be accessed at early stages.

In Chapter 4, we asked whether decomposition is sensitive to the PHONOLOGICAL CONDITIONING in morphological alternations such as *in*∼*im*. The results reported (experiment 3) could not be interpreted accordingly, since one of the morphological single-affix control conditions (the *in*-condition: *inelegant-INTREPID*) unexpectedly did not show priming. As a side investigation, we explored the possibility that these results might have been due to the segmentation algorithm, i.e. the process that is responsible for the acquisition and formation of morpho-orthographic units. The currently assumed segmentation algorithm involves a evaluation procedure of the probability of a given bigram to be grouped together on the basis of it transition probability (TP) with respect to the adjacent bigrams; if a bigram has a lower TP than the flanking bigrams, a trough pattern arises, and a morpheme boundary is placed at the

PRIMING TYPE	CONDITION	EXAMPLE	PRIME LEXICALITY	PRIMING?	EXP. #, REFERENCE
ROOT PRIMING	identity	*fuss-FUSS*	W	Y	exp. 1, ch. 2
	transparent	*alarming-ALARM*	W	Y	
	opaque	*belly-BELL*	W	Y	
	syllabic	*canvas-CAN*	W	N	
	non-syllabic	*starch-STAR*	W	N	
	ful	*successful-SUCCESS*	W	Y	exp. 5, sec. 5.3
	able	*detectable-DETECT*	W	Y	
	ity	*purity-PURE*	W	Y	
	ment	*settlement-SETTLE*	W	Y	
	ness	*weakness-WEAK*	W	Y	
	able	*skittable-SKIT*	NW	N	exp. 6, sec. 5.4
	ity	*skittity-SKIT*	NW	Y	
	ment	*skitment-SKIT*	NW	N	
	ness	*skittness-SKIT*	NW	Y	
SUFFIX PRIMING	sgdom-sgdom	*worlds-HEAVENS*	W	Y	exp. 2, ch. 3
	pldom-sgdom	*windows-GODS*	W	N	
	er	*scorer-WORKER*	W	N	
	rhyme	*taper-VAPOR*	W	N	
PREFIX PRIMING	*in*	*inelegant-INTREPID*	W	N	exp. 3, ch. 4
	im	*immature-IMPURE*	W	Y	
	in-im	*inelegant-IMPURE*	W	N	
	im-in	*immature-INTREPID*	W	N	
	dis	*disembark-DISABLE*	W	N	
	cohort	*banter-BANJO*	W	N	

Table 6.1.: Summary of results of visual masked priming experiments reported in this dissertation.

low-TP peak (the "trough pattern" Seidenberg, 1987). If a trough pattern "consistently surfaces" across many letter sequences involving the same bigram, a morpho-orthographic unit is formed (also called "islands of regularity"; Rastle and Davis, 2008, p. 955). Our analyses revealed the inadequacy of this algorithm, as only few of the affixes under consideration were indeed found to show consistent trough-pattern behavior (see secc. 4.2.5 and 5.4.5) and call for a reassessment of the segmentation process. Nonetheless, we claimed that, if replicable, our results suggest that visual morphological decomposition may be *selective*, in the sense that it only occurs for a subset of the morphemes of the language, which we tentatively claimed to be listed in a "*morpheme repository*" (along the lines of Dikker et al., 2009). On one hand, if true, selective decomposition raises a number of questions regarding the systematicity of the decomposition procedures. At this time, we are not able to provide any specific argument on this matter – namely, how and why some morphemes decompose and others do not. On the other hand, we ought to point out that these unexpected results might have been brought about the online data collection environment (see sec. 2.3). More data – in both in-lab and online environments – will need to be collected to validate the issue and further understand its underpinning source (see sec. 6.3).

Finally, Chapter 5 asked whether decomposition is sensitive to SYNTACTIC AFFIXAL RESTRICTIONS, whereby affixes may only attach to specific stem categories (e.g. *-ity* attaches to adjectives: pure]$_A$ ity]$_N$, but may not attach to nouns: *bliss]$_N$ ity]$_N$). Unfortunately, our investigation revealed that the impact of zero-derivation on syntactically illicit forms could not be properly factored out, which prevented us from making any strong claim about the potential sensitivity of decomposition to syntactic affixal restrictions.

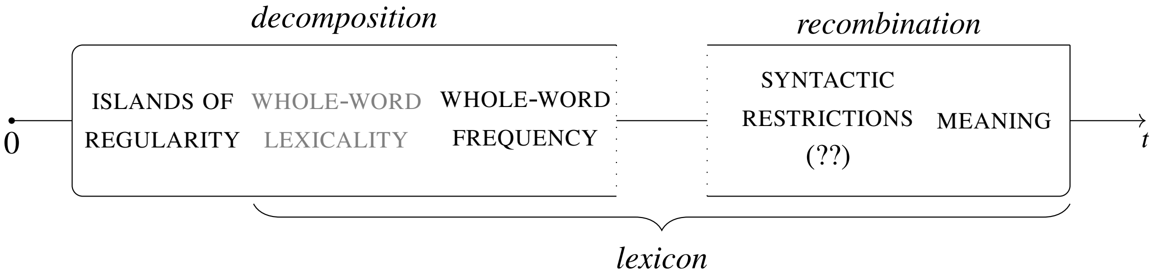

Figure 6.1.: Updated time-course of the unfolding of word information in word processing (final version). Whole-word lexicality is greyed-out because it was not directly addressed in this dissertation. The question mark symbol '??' refers to provisional claims that need further testing.

As a whole, our results suggest that there is a strong divide between the properties that are accessible during the decomposition stage and properties that are, instead, accessible after (e.g., at Taft, 1994's recombination stage; see Figure 6.1). On one hand, this investigation supports the hypothesis that decomposition is essentially morpho-orthographic, as maintained by current models of decomposition. On the other hand, decomposition seems to be additionally affected by whole-word properties such as FREQUENCY and (possibly) LEXICALITY. This is instead unexpected in any of the models of lexical access being considered in this dissertation because both properties are generally assumed to be accessed at the end point of word processing. Here we propose a model that is able to account for these facts, while still maintaining the morpho-orthographic nature of decomposition. The model proposed here assumes a model of word processing consisting of three major stages (along the lines of Taft, 1994's full-decomposition model): (i) the orthographic stage, in which the physical stimulus is decoded into a parsable orthographic representation; (ii) the decomposition stage, in which the orthographic representation is break down into smaller units; and (iii) the recombination stage, in which the units are put back together. Here, we specifically focus on the decomposition stage, which we claim consists of two separate steps. In the first step (*breakdown step*), the incoming stimulus is broken down into all possible decomposition patterns on the basis of morpho-orthographic segmentation. This operation is quite fast and only requires that one of the subset segments of the visual stimulus be identified as a morpho-orthographic unit. Whole strings may also be among the potential decomposition patterns, as long as they match with a morpho-orthographic unit. For example, for a mono-morphemic words such as *canvas* (or *starch*, i.e. the words used as prime words in the syllabic and non-syllabic orthographically-related conditions of experiment 1), the following decomposition patterns are generated: `$canvas$` and `$can-vas$`, which are triggered by the morpho-orthographic units therein (i.e., {*canvas*}, {*can*}, respectively). In the case of non-words such as *slegrack* (as well as *flexire*), the only possible decomposition pattern is `$sleg-rack$`, since no other pattern contains a segment that matches with a morpho-orthographic unit (e.g., the whole-string form may not be a pattern because it is not a morpho-orthographic unit).[21] In the second step (evaluation step), the generated decomposition patterns are evaluated with respect to the the number of the segments therein that actually match with a morpho-orthographic unit. Evaluation is inherently binary, in the sense that each pattern may be assigned a score of either 0 or -1, depending on whether or not all the segments within the pattern match with a morpho-orthographic unit. The candidate with the highest score is chosen as the optimal one (see Fig. 6.2a). In *canvas*, while the candidate `$can-vas$` receives a score of -1 because of the morpho-orthographically incompatible segment `$vas$`, the candidate `$canvas$` receives a score of 0 and is selected as the optimal candidate; it then activates, therefore inhibiting activation of the target root unit {*can*}. In *slegrack*, the only available decomposition pattern `$sleg-rack$` is automatically selected, even though it has a score of -1; it therefore activates, and elicits priming onto the target root *RACK*). This binary evaluation system is able to provide a mechanistic gap that we identified in previous models of decomposition in which the brothel non-effect was descriptively explained as resulting from a "decomposability constraint" (Crepaldi et al., 2010; see also secc. 1.4.2 and 2.4.1). Even though our results do not clearly suggest this to be the case since a potential impact of zero-derivation could not completely be factored out (see ch. 5), in sec. 5.4.5 we suggested that syntactic affixal restrictions are only available at the recombination stage and therefore cannot hinge on decomposition procedures (Fig. 6.1). If this hypothesis were to be further supported in future studies, it would be seamlessly explained in the present model (Fig. 6.2a). Syntactically illicit non-words such as *blissity* are broken down to form only one candidate: `$bliss-ity$`, as it is the only pattern that has at least one segment that matches with a morpho-orthographic unit. Next, in the evaluation step, it is automatically selected as the optimal candidate (as there is no other competing candidates to compare it to). Both units therein activate (i.e., {*bliss*}, {*ity*}); thus, priming onto the target arises (*BLISS*).

If more than one candidate is assigned a score of 0 (i.e., at least two candidates have the same number of segments that match with a morpho-orthographic unit), they compete with one another via *lateral inhibition*. By this mechanism, units with the highest resting activation level activate faster and inhibit the competing units. It is generally assumed that the resting activation level of a morpho-orthographic unit correlates with its frequency; high-frequency units have a high resting activation level, low-frequency units have a low resting activation level (see sec. 1.2.1). On one hand, in multi-morphemic prime words in which the root unit has a higher resting activation level than the whole-string unit, the former unit activates at the expense of the latter unit. This is what usually happens for the words tested in the typical morphologically related condition tested in root priming studies (for a review, see 1.3). To

---

[21] Depending on the actual directionality of the scanning system of decomposition, one may argue that the alternative pattern `$s-leg-rack$` is also generated. This, however, does not change the result, since {*rack*} is still identified as a morpho-orthographic unit, and therefore is still expected to trigger priming onto the related target *RACK*.

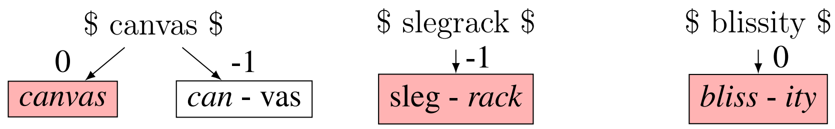

(a) Root-containing mono-morphemic word and non-word visual stimuli.

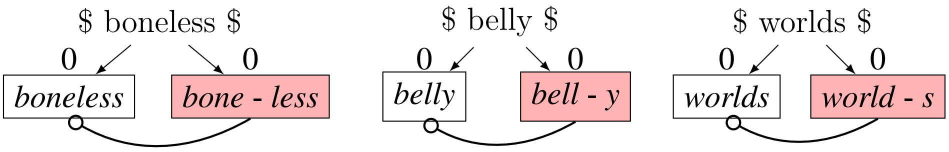

(b) Stem-dominant bi-morphemic (transparent and opaque) word visual stimuli.

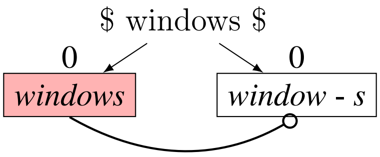

(c) Affixed-dominant bi-morphemic word visual stimuli.

Figure 6.2.: A competition-based model of decomposition. Breakdown and evaluation across the different types of word stimuli tested in this dissertation. Morpho-orthographically compatible segments are *italicized*. Activated patterns are colored in red; the circle-ending arrow signals inhibition.

exemplify, let's take the prime words used for the morphologically transparent and opaque conditions of experiments 1 and 5 (e.g., *boneless, belly*); and the prime words used for the sgdom-sgdom condition of experiment 2 (e.g., *worlds*).[22] For each of these words, the breakdown step first generates morpho-orthographically compatible candidates (see above). Then, the evaluation step selects both the whole-string candidate (e.g., $ boneless $, $ belly $, $ worlds $) and the decomposed candidate ($ bone-less $, $ bell-y $, $ world-s $), since they have the same number of segments matching with morpho-orthography. For each word, the two candidates compete for activation. As Figure 6.3 shows, all prime words have a dominance ratio smaller than 0 in all conditions, which indicates stem-dominance (see 3.2.1 for the dominance formula used in this dissertation). This is also confirmed by a series of one-sample t-tests: *transparent condition*: $\beta = -1.87, t(31) = -8.74, p < .0001$; *opaque condition*: $\beta = -0.84, t(31) = -2.31, p = .008$; *sgdom-sgdom condition*: $\beta = -2.85, t(19) = -14.38, p < .0001$; *able-condition*: $\beta = -2.53, t(23) = -15.54, p < .001$; *ity-condition*: $\beta = -2.81, t(23) = -7.29, p < .0001$; *ment-condition*: $\beta = -0.88, t(23) = -2.81, p = .01$; *ness-condition*: $\beta = -2.53, t(23) = -15.54, p < .001$. The negative β-estimates confirm

---

[22] We do not mention the the prime words used for the conditions of experiment 3 here, since the experiment was generally affected by the bigger issue regarding the potential selectivity of decomposition (see section below for further details).

that the mean dominance ratio of the prime words used in the three conditions is significantly smaller than 0. In such cases, as having a higher resting activation level, the root unit (i.e., {*bone*}, {*bell*}, {*world*}) activates faster than the competing whole-string unit (i.e., {*boneless*}, {*belly*}, {*worlds*}), which is therefore inhibited (see Fig. 6.2b). When the related target root is presented (*BONE, BELL, WORLD*), priming arises.

On the other hand, in multi-morphemic prime words in which the root unit has a lower resting activation level than the whole-string unit, the latter unit activates at the expense of the former unit. This is what happens, for example, for the prime words used in the pldom-sgdom condition in experiment 2 (e.g., *windows*). As Figure 6.3 shows, the mean dominance ratio of this condition is significantly greater than 0, which indicates affixed-dominance. The one-sample *t*-test performed for this condition seems to confirm this: $\beta = 2.33, t(19) = 4.8, p = .0001$, where the β-estimate for this condition is instead positive. In such cases, the higher-frequency whole-string candidate \$windows\$ activates faster than the decomposed candidate \$win-dow-s\$, which is therefore inhibited (see Fig. 6.2c). When the related, singular-dominant target root is presented (*GODS*), priming does not arise.

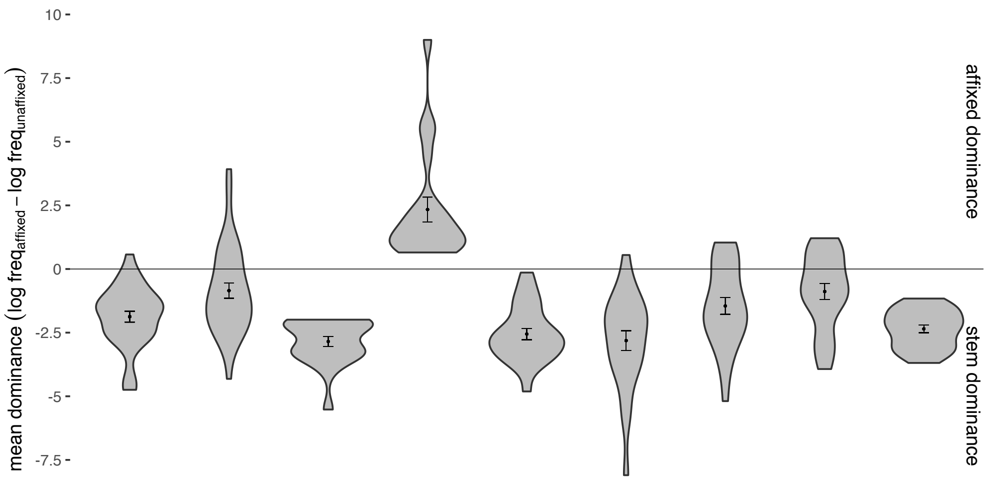

Figure 6.3.: Dominance ratios of the prime words used in the morphological conditions across all the experiments reported in this dissertation.

In the upcoming subsection, we set out a research program aiming at testing and possibly further extending the model just sketched.

6.3. *Decomposition matters*: a research program

While addressing the question in (16), our investigation ultimately lays the groundwork for future work on morphological decomposition. In this final section, we briefly outline a set of programmatic questions we anticipate to be addressed in the near future. The program we envision below is concerned with both theoretical and methodological avenues of research, each described in the two following subsections. In section 6.3.1, we discuss potential ways to further explore the mechanisms underlying the segmentation algorithm and the lexically-informed decomposition procedure we have argued for in this dissertation. In subsection 6.3.2, we sketch an agenda for exploring the mechanics of both behavioral (e.g., masked priming in both visual and auditory modalities) and electrophysiological (e.g., correlational source and space analyses on high-density EEG) experimental methodologies for investigations of morphological decomposition.

6.3.1. Theoretical issues

At the end of the previous section, we proposed a parallel model of decomposition, in which multiple decomposition patterns of the same visual stimulus are generated and evaluated in parallel. Here, we identify five main research questions aiming to further testing the model.

1. Selective decomposition: The first question interrogates the properties of form-based, selective decomposition. Our results show that priming effects may only arise for specific affixes: i.e., *-s* (experiment 2), *im-* (experiment 3), and *-able, -ity, -ful, -ment*, and -ness (experiment 5), but not for others – i.e., *-er* (experiment 2), *in-, dis-* (experiment 3). As such, these results seem to suggest a *selective* procedure of visual decomposition, which may or may not occur depending on the affix being involved. On one hand, such a claim, if true, would challenge all theories of early decomposition, in which all morphemes are predicted to undergo decomposition systematically. In previous chapters, we then tentatively proposed that decomposition may be selectively licensed by a "*morpheme repository*," which includes only the orthographic form of the decomposable morphemes (along the lines of Dikker et al., 2009). At this stage of research, we ought to say that there seems to be no theoretically-relevant, testable rationale behind the distinction between decomposable morphemes (i.e., included in the repository) and non-decomposable morphemes (i.e., excluded from the repository). We are therefore led to believe that the asymmetrical priming effects across affixes might have been due to the high amount of noise in the data, as brought about by the online data collection environment (see sec. 2.3). We anticipate that the experiments reported in this dissertation will be run in an in-lab environment, which guarantees more control over common sources of noise (i.e., prime duration fluctuations, poor RT reliability, high variance).

2. Dominance/whole-word frequency: The second question regards the impact of whole-word frequency onto decomposition. In the previous section, we proposed a model of decomposition in which activation of morpho-orthographic units at early stages of processing may be affected by their resting activation level, which is generally assumed to correlate with their frequency. We proposed this (i) as a way to explain the dominance-driven asymmetrical priming effects reported in Chapter 3 (experiment 2), while (ii) still maintaining the generally-accepted claim about early decomposition. As outlined above, the model predicts the priming response to be systematically affected by dominance asymmetries between the affixed and the stem forms, regardless of the affix being used and of the priming type being elicited (i.e., root/affix priming). One way to test this prediction may be to conduct an experiment eliciting the masked priming response to stem- and affixed-dominant words sharing the same affix, but arranged in separate conditions (much like we did for plural forms in experiment 2).

3. Phonological conditioning: The third question concerns the question of whether decomposition is affected by phonologically-conditioned morphological alternations. We meant to start exploring this question with experiment 3, though methodological confounding issues (i.e., selective decomposition, see question 1 above) prevented us from doing that at this time. In the future, we plan to resume addressing this question. First, we anticipate that experiment 3 will be rerun in the in-lab environment, which reduces the amount of noise in the data (see sec. 2.3) and, therefore, may help interpretation of the data in the terms sketched in sec. 4.2.3. Second, we also expect to run more experiments looking at the decomposition mechanisms at play in the processing of similar phenomena. Unfortunately, the English language has no other similar phonologically-conditioned alternation (i.e., in which the alternating realizations are orthographically marked), and, therefore, could only be used to explore the decomposition mechanisms in the auditory modality, once the methodological issues of the masked priming design in the auditory modality are resolved (see sec. 6.3.2 below).

Alternations such as [ɪn]~[ɪm] represent the simplex case of phonological conditioning in morphology, as they involve exactly *one* phonological operation; in our example, the underlying representation /in/ becomes [ɪm] by a single operation of place assimilation. As discussed in detail in Chapter 4, investigating the decomposition of such alternations may allow us to understand whether the phonological asymmetry between allomorphs of the same morpheme is detected at all during decomposition. If we found it is the case, priming-elicited decomposition may also be probed as a way to test competing phonological accounts for more complex phonologically-conditioned alternations, in which *more than one* phonological operation interact with one another. Derivational theories assume that phonological processes are ordered serially so that the surface form results from multiple, intermediate stages in which the underlying form undergoes rule in a fixed order (SPE: Halle and Chomsky, 1968). Representational theories assume, instead, that phonological processes all occur in parallel and compete with one another for output selection (*Optimality Theory*: Prince and Smolensky, 2008). These two theories of phonology therefore make different predictions with respect to the number of grammatical operations involved in the computation, which have been argued to affect processing time and therefore response times (*derivational theory of complexity*, DTC; Miller and Chomsky, 1963; Phillips, 1996). On one hand, derivational theories predict that the parser analyzes the stimulus by applying the rules "in reverse" to go back to the underlying representation (which is assumed to be the long-term representation). As such, decomposition procedures are predicted to take more time depending on the the number of rules to undo to get to the corresponding long-term representation. On the other hand, representational theories predict that multiple phonological processes are all evaluated in parallel. Therefore, decomposition procedures are predicted to be unaffected by the number of processes involved.

5. Syntactic restrictions: Finally, the fifth question concerns the impact of syntactic restrictions on decomposition. In the section above we tentatively claimed that decomposition is not affected by violations to syntactic restrictions. To reiterate, this claim could only be tentative because zero-derivation could not be ruled out as a potential interacting factor in the investigation reported in Chapter 5(see in particular sec. 5.4.5). Therefore, more data will have to be collected in the future to test the claim made here, while ensuring that the potential impact of zero-derivation is controlled for. In particular, we envision two avenues of research to this end. First, the impact of syntactic restrictions on decomposition will need to be investigated for all remaining affixes, upon prior investigation of the potential interpretive risks of zero-derivation and selective decomposition. For example, we can anticipate that the affix *-ful* may be an optimal candidate for a future follow-up study; since we have already confirmed that it decomposes (experiment 5), one would just need to make sure that these forms (e.g., **fastful*) are not judged as grammatical due to zero-derivation before actually eliciting decomposition via masked priming (e.g.,*fastful-FAST*). Other potential affixes to test could be, for example, *-less* (*fastless-FAST*), *-ize* (*fastize-FAST*), and *al* (*fastal-FAST*). Second, the same investigation will also have to be conducted in affix priming designs. In previous months, we indeed conducted two suffix priming experiments along the lines of the two root priming experiments reported above (e.g., licit: *detectable-CURABLE*; illicit: *skittable-CURABLE*). However, we chose not to report the results of these experiment here, as evident methodological issues (i.e., lack of priming in the identity control condition) prevented any sort of theoretical interpretation of the results retrieved.[23]

6.3.2. Methodological issues

We anticipate that three different methodological projects will be carried out in the near future. All three projects have the main goal of making up a comprehensive set of experimental (behavioral and electrophysiological) tools for linguists interested in morphological processing.

Prime processing time in visual masked priming

In line with the literature on visual morphological decomposition, we have made great use of the visual masked priming design to explore the question in (16). Regardless, we should acknowledge that the whole enterprise undertaken in this dissertation hinges on the assumption that the masked priming design ensures that the processing of prime stimulus stops as soon as it is visually replaced by the target stimulus on the screen ('no-overlapping theory of priming'). This assumption is crucial for us to interpret the priming data as indicative of the availability of specific properties during decomposition of the prime stimulus. An alternative plausible theory of masked priming assumes that the prime stimulus keeps being analyzed in parallel

---

[23] Relatedly, a preliminary methodological study will also have to be conducted to define the mechanism behind the difference between root and suffix masked priming. While the root priming has consistently been reported to be more reliable than the affix priming (sec. 1.3; see also Amenta and Crepaldi, 2012), the difference is essentially left unexplained in all models of decomposition and requires a much more thorough investigation to be fully understood.

with the target and priming effects may therefore be triggered post target onset ('overlapping theory of priming' Forster et al., 2003). From this perspective, the masked priming response cannot be used to explore the mechanisms of early morphological decomposition and, more generally, on early stages of visual word processing. One way to assess the contribution of the pre-target-onset prime processing onto the priming response is to implement a backward mask in the masked priming design. The idea is not new (Forster et al., 2003; Morris et al., 2007; Grainger et al., 2003; Holcomb and Grainger, 2006) and involves the use of a mask (i.e., #####, usually lasting around 20-30 ms) being sandwiched between the prime and the target stimuli (so the presentation of a typical prime-TARGET pair flows as follows: ##### (500 ms) → *fuss* (30-40 ms) → ##### (20-30 ms) → *FUSS*). The backward mask is meant to prevent processing of the prime stimulus from overlapping with processing of the target stimulus and therefore allow us to quantify the amount of processing of the prime occurring within the duration of prime presentation.[24] The investigation will need to run a series of masked priming experiments across conditions with different relatedness dimensions (e.g., identity pairs: *fuss-FUSS*; morphologically transparent and opaque related pairs: *boneless-BONE, belly-BELL*; orthographically related pairs; *canvas-CAN, starch-STAR*; dominance-asymmetric pairs; *worlds-HEAVENS, windows-GODS*), both with a constant backward mask (backward-mask condition) and without it (no-backward-mask condition). The no-overlapping theory of priming predicts the priming response to be the same for all conditions, regardless of the presence of the backward mask; for example, *boneless-BONE* should therefore prime in both the no-backward-mask and the backward-mask conditions. The overlapping theory of priming, instead, predicts the priming response to be a function of the presence of the backward mask; for example, *boneless-BONE* should prime in the no-backward-mask condition, but should not prime in the backward-mask condition.

Auditory masked priming

The initial plan for this dissertation envisaged an investigation on the decomposition's sensitivity to linguistic properties in both the visual and the auditory modalities as a way to investigate the potentially different, modality-driven mechanisms of decomposition. To this end, we ran three auditory experiments in total. One experiment tested the sensitivity of auditory decomposition to syllabification (much like the visual counterpart described in Chapter 2). The other two experiments tested the sensitivity of auditory decomposition to two different sets of phonologically-conditioned alternations: [z]∼[ɪz] (i.e., two of the three realizations of the English regular plural morpheme *-s*) and [ɪn]∼[ɪm] (akin to the visual counterpart described in Chapter 4).

We used the auditory masked priming technique pioneered in Kouider and Dupoux (2005). This technique is, in principle, the auditory counterpart of the visual masked priming technique. In the auditory masked priming design, the typical prime-target sound stream had the following characteristics (see Figure 6.4). All primes are first intensity-attenuated and then time-compressed, so that their original duration was reduced to usually between 35% and 40%

---

[24] One potential complication is that the presence of the backward mask has been occasionally reported to have the side effect of increasing priming visibility and therefore weaker priming effects (Forster et al., 2003, p. 5). This fact suggests that particular attention will be needed when choosing the duration of the backward mask.

of their original duration. Randomly selected, attenuated, and time-compressed primes are also time-reversed to create auditory masks. Each time-compressed prime is preceded by one mask and followed by four other different masks; the target word, which remained at a normal speech rate and non-attenuated, was over-imposed on the sequence of masks and prime so that the onset of the target coincided with the offset of the prime. Participants usually report hearing a clear word (i.e., the target word sound) surrounded by unintelligible babble. The results suggested a number of methodological issues, which led us to ultimately decide to not report them in this dissertation. The issues seem to be connected to the interplay between the time-compression rate and the phonological length of the prime stimuli, which have a substantial impact on the subliminal intelligibility of the prime – and, ultimately, to priming elicitation. A full-length study will be conducted in the near future that explores the interdependence of compression duration, the linguistic/acoustic factors just mentioned, and priming effects, with the ultimate goal of defining the minimum compressed duration at which priming can be obtained.

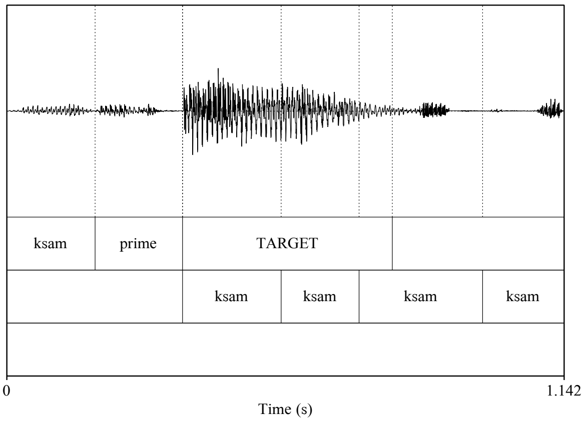

Figure 6.4.: Schematic description of a typical trial in Kouider and Dupoux (2005)'s auditory masked priming paradigm. The attenuated and time-compressed prime is immediately preceded and followed by attenuated, time-compressed, and time-reversed masks (as indicated by the label *ksam*, i.e. the word *mask* in reverse). The target is played at normal speech rate and is not attenuated; it immediately follows the prime (i.e., its onset coincides with the offset of the prime) and is superimposed on the rightward mask stream.

Electrophysiological methodologies

In the last decade, a series of magnetoencephalographic (MEG) studies have suggested that the brain response to morphological properties originates from a region of the inferior temporal cortex known as 'Visual Word Fom Area' (VWFA; Cohen et al., 2000, 2002) and peaks at about 150 ms post stimulus onset (the M170 component). This area of the brain has been shown to be sensitive to the morpho-orthographic statistical regularities at play during the early stages of the processing of non-affixed (e.g., heteronyms: *wind* Solomyak and Marantz, 2009), affixed (e.g., bi-morphemic words involving free and bound stems: *taxable, tolerable*; Solomyak and Marantz, 2010), and pseudo-affixed (Lewis et al., 2011, e.g., *brother*) words, but not to semantic properties, which seem to be indexed by a response in the superior temporal and Silvyan Fissure region peaking at about 350 ms (the M350 component; Pylkkänen and Marantz, 2003). These studies have developed a novel approach to MEG data analysis, which is based on the correlational data analysis approach pioneered by Hauk et al. (2006) for ERP sensor-space data analysis. While traditional data analysis pipelines group together stimuli sharing some properties in different conditions and compare them to one another, this new approach correlates a given variable (e.g., orthographic TP) of *each stimulus* with the brain response arising at a given time window. This kind of technique is particularly beneficial for investigations on morphological phenomena, which often consists of a small number of stimuli, which would be impossible to control for all the necessary variables. Unfortunately, not many morphologists that do not have access to highly expensive MEG machines, as there are, indeed, only a few available worldwide.

A still ongoing project we have recently taken up explores the possibility of applying the MEG correlational technique just mentioned to high-density EEG data as a way to provide a reliable and a cost-effective methodology for morphologists interested in conducting electrophysiological research. So far, we have collected high-density EEG data for an identical replication of Lewis et al. (2011). In this study, the M170 response to both transparent (e.g., *driver*) and opaque (e.g., *brother*) words is shown to correlate with the morpho-orthographic TP between pseudo-stem (*broth*) and pseudo-affix (*er*), which supports the morpho-orthographic nature of early visual decomposition. We are now in the process of defining the EEG-adapted pipeline and comparing the two sets of results to one another. The hope is to obtain an agreed-on pipeline for correlational data analysis, which may guarantee comparability across types of signal (i.e., MEG vs. EEG) and boost electrophysiological research in morphology.

References

Albright, Adam and Bruce Hayes (2003). “Rules vs. analogy in English past tenses: A computational/experimental study”. In: *Cognition* 90.2, pp. 119–161.

Allen, Margaret Reece (1979). “Morphological investigations”. PhD thesis. Storrs, CT: University of Connecticut.

Amenta, Simona and Davide Crepaldi (2012). “Morphological Processing as We Know It: An Analytical Review of Morphological Effects in Visual Word Identification”. In: *Frontiers in Psychology* 3, p. 232.

Andoni Duñabeitia, Jon, Manuel Perea, and Manuel Carreiras (2008). “Does darkness lead to happiness? Masked suffix priming effects”. In: *Language and Cognitive Processes* 23.7-8, pp. 1002–1020.

Anshen, Frank, Mark Aronoff, Roy J. Byrd, and Judith Klavans (1986). “The Role of Etymology and Word Length in English Word Formation”. Ms., Stony Brook.

Aronoff, Mark (1976). *Word formation in generative grammar*. Linguistic Inquiry Monographs. Cambridge, MA: MIT Press.

Aronoff, Mark and Nanna Fuhrhop (2002). “Restricting suffix combinations in German and English: closing suffixes and the monosuffix constraint”. In: *Natural Language & Linguistic Theory* 20, pp. 451490.

Baayen, R. Harald (2008). *Analyzing Linguistic Data: A Practical Introduction to Statistics using R*. Cambridge, UK: Cambridge University Press.

Baayen, R. Harald and Rochelle Lieber (1991). “Productivity and English derivation: a corpus-based study”. In: *Linguistics* 29, pp. 801–843.

Baayen, R. Harald, Ton Dijkstra, and Robert Schreuder (1997). “Singulars and plurals in Dutch: Evidence for a parallel dual-route model”. In: *Journal of Memory and Language* 37.1, pp. 94–117.

Baayen, R. Harald, Robert Schreuder, Nivja De Jong, and Andrea Krott (2002). “Dutch inflection: the rules that prove the exception”. In: *Storage and computation in the language faculty*, pp. 61–92.

Baayen, R. Harald, W. Levelt, Robert Schreuder, and Mirjam Ernestus (2007). “Paradigmatic structure in speech production”. In: *Proceedings from the annual meeting of the Chicago linguistic society*. Vol. 43. 1. Chicago Linguistic Society, pp. 1–29.

Balota, David A., Melvin J. Yap, Keith A. Hutchison, Michael J. Cortese, Brett Kessler, Bjorn Loftis, James H. Neely, Douglas L. Nelson, Greg B. Simpson, and Rebecca Treiman (2007). “The English Lexicon Project”. In: *Behavior Research Methods* 39.3, pp. 445–459.

Barr, Dale J., Roger Levy, Christoph Scheepers, and Harry J. Tily (2013). “Random effects structure for confirmatory hypothesis testing: Keep it maximal”. In: *Journal of memory and language* 68.3, pp. 255–278.

Bayes, Thomas (1763). “LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, FRS communicated by Mr. Price, in a letter to John Canton, AMFR S”. In: *Philosophical transactions of the Royal Society of London* 53, pp. 370–418.

Chateau, Dan, Erin V Knudsen, and Debra Jared (2002). “Masked priming of prefixes and the influence of spelling–meaning consistency”. In: *Brain and Language* 81.1-3, pp. 587–600.

Cohen, Laurent, Stanislas Dehaene, Lionel Naccache, Stéphane Lehéricy, Ghislaine Dehaene-Lambertz, Marie-Anne Hénaff, and François Michel (2000). “The visual word form area: Spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients”. In: *Brain* 123, pp. 291–307.

Cohen, Laurent, Stéphane Lehéricy, Florence Chochon, Cathy Lemer, Sophie Rivaud, and Stanislas Dehaene (2002). “Language-specific tuning of visual cortex? Functional properties of the Visual Word Form Area”. In: *Brain* 125.5, pp. 1054–1069.

Coltheart, Max, Kathleen Rastle, Conrad Perry, Robyn Langdon, and Johannes Ziegler (2001). “DRC: a dual route cascaded model of visual word recognition and reading aloud”. In: *Psychological review* 108.1, pp. 204–256.

Crepaldi, Davide, Kathleen Rastle, Max Coltheart, and Lyndsey Nickels (2010). “‘Fell’ primes ‘fall’, but does ‘bell’ prime ‘ball’? Masked priming with irregularly-inected primes”. In: *Journal of Memory and language* 63, pp. 83–99.

Crepaldi, Davide, Lara Hemsworth, Colin J. Davis, and Kathleen Rastle (2016). “Masked suffix priming and morpheme positional constraints”. In: *The Quarterly Journal of Experimental Psychology* 69.1, pp. 113–128.

Daland, Robert, Bruce Hayes, James White, Marc Garellek, Andrea Davis, and Ingrid Norrmann (2011). “Explaining sonority projection effects”. In: *Phonology* 28.2, pp. 197–234.

Davis, Colin J. (1999). “The self-organising lexical acquisition and recognition (SOLAR) model of visual word recognition”. PhD thesis. University of New South Wales.

Davis, M. H., W. D. Marslen-Wilson, and M. G. Gaskell (2002). “Masked priming of prefixes and the influence of spelling–meaning consistency”. In: *Journal of Experimental Psychology: Human Perception and Performance* 28, pp. 218–244.

Davis, Matt H. (2003). “Connectionist modelling of lexical segmentation and vocabulary acquisition”. In: *Connectionist models of development: Developmental processes in real and articial neural networks*. Ed. by Quilan P. Hove, UK: Psychology Press.

Dehaene, S., L. Cohen, M. Sigman, and F. Vinckier (2005). “The neural code for written words: A proposal”. In: *Trends in Cognitive Science* 9.7, pp. 335–341.

Dehaene, Stanislas and Laurent Cohen (2007). “Cultural recycling of cortical maps”. In: *Neuron* 56, pp. 384–398.

Dikker, Suzanne, Hugh Rabagliati, and Liina Pylkkänen (2009). “Sensitivity to syntax in visual cortex”. In: *Cognition* 110.3, pp. 293–321.

Dominguez, Alberto, Maira Alija, Javier Rodriguez-Ferreiro, and Fernando Cuetos (2010). “The contribution of prefixes to morphological processing of Spanish words”. In: *European Journal of Cognitive Psychology* 22.4, pp. 569–595.

Elman, J. L. (1990). “Finding structure in time”. In: *Cognitive Science* 14, pp. 179–211.

Embick, David (2015). *The morpheme: A theoretical introduction*. Vol. 31. Boston, MA/Berlin: Walter de Gruyter.

Embick, David and Alec Marantz (2008). “Architecture and blocking”. In: *Linguistic inquiry* 39.1, pp. 1–53.

Fabb, Niegel (1988). “English Suffixation Is Constrained Only by Selectional Restrictions”. In: *Natural Language & Linguistic Theory* 6.4, pp. 527–539.

Fiorentino, Robert and Ella Fund-Reznicek (2009). “Masked morphological priming of compound constituents”. In: *The Mental Lexicon* 4.2, pp. 159–193.

Fiorentino, Robert and David Poeppel (2007). “Compound words and structure in the lexicon”. In: *Language and Cognitive processes* 22.7, pp. 953–1000.

Fiorentino, Robert, Stephen Politzer-Ahles, Natalie S Pak, María Teresa Martínez-García, and Caitlin Coughlin (2015). “Dissociating morphological and form priming with novel complex word primes: Evidence from masked priming, overt priming, and event-related potentials”. In: *The mental lexicon* 10.3, pp. 413–434.

Forster, Kenneth I. (1981). “Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing”. In: *Quarterly journal of Experimental Psychology* 33, pp. 465–495.

— (1998). “The pros and cons of masked priming”. In: *Journal of psycholinguistic research* 27.2, pp. 203–233.

— (1999). “Microgenesis of priming effects in lexical access”. In: *Brain and Language* 68, pp. 5–15.

Forster, Kenneth I. and Chris Davis (1984). “Repetition priming and frequency attenuation in lexical access.” In: *Journal of experimental psychology: Learning, Memory, and Cognition* 10.4, p. 680.

Forster, Kenneth I., Kathleen Mohan, and Jo Hector (2003). “The mechanics of masked priming”. In: *Masked priming: The state of the art*. Ed. by Sachiko Kinoshita and Stephen J. Lupker. New York, NY/Hove, UK: Psychology Press, pp. 3–37.

Fruchter, Joseph and Alec Marantz (2015). “Decomposition, lookup, and recombination: MEG evidence for the Full Decomposition model of complex visual word recognition”. In: *Brain and language* 143, pp. 81–96.

Fruchter, Joseph, Linnaea Stockall, and Alec Marantz (2013). “MEG masked priming evidence for form-based decomposition of irregular verbs”. In: *Frontiers in Human Neuroscience* 7, p. 798.

Gimenes, Manuel and Boris New (2016). “Worldlex: Twitter and blog word frequencies for 66 languages”. In: *Behavior Research Methods* 48.3, pp. 963–972.

Grainger, Jonathan (2018). “Orthographic processing: A ‘mid-level’ vision of reading: The 44th Sir Frederic Bartlett Lecture”. In: *Quarterly Journal of Experimental Psychology* 71.2, pp. 335–359.

Grainger, Jonathan and Arthur M. Jacobs (1996). “Orthographic processing in visual word recognition: a multiple read-out model”. In: *Psychological review* 103.3, pp. 518565.

Grainger, Jonathan, Kevin Diependaele, Elsa Spinelli, Ludovic Ferrand, and Fernand Farioli (2003). “Masked repetition and phonological priming within and across modalities.” In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29.6, p. 1256.

Halle, Morris and Noam Chomsky (1968). *The Sound Pattern of English*. Harper & Row.

Hartigan, J. A. and P. M. Hartigan (1985). “The Dip Test of Unimodality”. In: *The Annals of Statistics* 13.1, pp. 70–84.

Hauk, O., Davis. M. H., M. Ford, F. Pulvermüller, and W. D. Marslen-Wilson (2006). “The time course of visual word recognition as revealed by linear regression analysis of ERP data”. In: *Neuroimage* 30.4, pp. 1382–1400.

Hayes, Bruce and James White (2013). “Phonological naturalness and phonotactic learning”. In: *Linguistic Inquiry* 44.1, pp. 45–75.

Holcomb, Phillip J and Jonathan Grainger (2006). “On the time course of visual word recognition: An event-related potential investigation using masked repetition priming”. In: *Journal of cognitive neuroscience* 18.10, pp. 1631–1643.

Jeffreys, Harold (1961). *Theory of probability*. Oxford, UK: Oxford University Press.

Jusczyk, P. W. (1997). *The discovery of spoken language*. Cambridge, MA: MIT Press.

Kiparsky, Paul (1982). “From cyclic phonology to lexical phonology”. In: ed. by Harry van der Hulst and N. Smith. Vol. 1. Foris: Dordrecht, pp. 131–175.

Kouider, Sid and Emmanuel Dupoux (2005). “Subliminal speech priming”. In: *Psychological Science* 16.8, pp. 617–625.

Kriesel, David (2007). *A Brief Introduction to Neural Networks*. URL: `http://www.dkriesel.com/en/science/neural_networks`.

Kruschke, John K (2011). *Doing Bayesian data analysis: A tutorial with R and BUGS. Burlington, MA*.

Kuznetsova, Alexandra, Per B. Brockhoff, and Rune H. B. Christensen (2017). “`lmerTest` Package: Tests in Linear Mixed Effects Models”. In: *Journal of Statistical Software* 82.13, pp. 1–26.

Landauer, Thomas K. and Susan T. Dumais (1997). “A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge.” In: *Psychological review* 104.2, pp. 211–240.

Lehtonen, Minna, Victor A Vorobyev, Kenneth Hugdahl, Terhi Tuokkola, and Matti Laine (2006). “Neural correlates of morphological decomposition in a morphologically rich language: An fMRI study”. In: *Brain and Language* 98.2, pp. 182–193.

Lehtonen, Minna, Toni Cunillera, Antoni Rodríguez-Fornells, Annika Hultén, Jyrki Tuomainen, and Matti Laine (2007). “Recognition of morphologically complex words in Finnish: Evidence from event-related potentials”. In: *Brain research* 1148, pp. 123–137.

Lewis, Gwyneth, Olla Solomyak, and Alec Marantz (2011). “The neural basis of obligatory decomposition of suffixed words”. In: *Brain and language* 118.3, pp. 118–127.

Lund, Kevin and Curt Burgess (1996). “Producing high-dimensional semantic spaces from lexical co-occurrence”. In: *Behavior Research Methods, Instruments, & Computers* 28.2, pp. 203–208.

Marchand, Hans (1969). *The Categories and Types of Present-Day English Word-Formation*. München: C. H. Becksche Verlagsbuchhandlung.

Marslen-Wilson, William D. and Alan Welsh (1978). “Processing interactions and lexical access during word recognition in continuous speech”. In: *Cognitive psychology* 10.1, pp. 29–63.

McClelland, James L. and Jeffrey L. Elman (1986). “The TRACE model of speech perception”. In: *Cognitive psychology* 18.1, pp. 1–86.

McClelland, James L. and David E. Rumelhart (1981). “An interactive activation model of context effects in letter perception: Part I. An account of basic findings”. In: *Psychological review* 88.5, pp. 375–407.

Miller, George A. and Noam Chomsky (1963). “Finitary models of language users”. In: *Handbook of Mathematical Psychology*. Ed. by R. Duncan Luce, Robert R. Bush, and Eugene Galanter. Vol. 2. New York, NY: John Wiley & Sons, pp. 269–321.

Morey, Richard D. (2018). *Using the `BayesFactor` package, version 0.9.2+*. URL: `https://richarddmorey.github.io/BayesFactor/`.

Morris, Joanna, Tiffany Frank, Jonathan Grainger, and Phillip J. Holcomb (2007). "Semantic transparency and masked morphological priming: An ERP investigation". In: *Psychophysiology* 44, pp. 506521.

Morris, Joanna, James H Porter, Jonathan Grainger, and Phillip J Holcomb (2011). "Effects of lexical status and morphological complexity in masked priming: An ERP study". In: *Language and cognitive processes* 26.4-6, pp. 558–599.

Muthmann, Gustav (2010). *Reverse English Dictionary*. Berlin/Boston: Mouton De Gruyter. DOI: `/10.1515/9783110806847`.

Neely, J. H. (1991). "Semantic priming effects in visual word recognition: A selective review of current findings and theories". In: *Basic processes in reading: Visual word recognition*. Ed. by D. Besner and G. Humphreys. Hillsdale, NJ: Erlbaum, pp. 264–336.

Peirce, Jonathan, Jeremy R. Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv (2019). "PsychoPy2: Experiments in behavior made easy". In: *Behavior Research Methods* 51.1, pp. 195–203.

Peirce, Jonathan W. (2009). "Generating stimuli for neuroscience using PsychoPy". In: *Frontiers in neuroinformatics* 2, p. 10.

Phillips, Colin (1996). "Order and structure". PhD thesis. Massachusetts Institute of Technology.

Pinker, Steven (1991). "Rules of language". In: *Science* 253.5019, pp. 530–535.

Plag, Ingo (1999). *Morphological productivity: Structural constraints in English derivation*. Vol. 28. Topics in English Linguistics. Berlin/New York, NY: Mouton de Gruyter.

Prince, Alan and Paul Smolensky (2008). *Optimality Theory: Constraint interaction in generative grammar*. John Wiley & Sons. DOI: `10.1002/9780470759400`. URL: `https://onlinelibrary.wiley.com/doi/book/10.1002/9780470759400`.

Pulvermüller, F., R. Assahollahi, and T. Elbert (2001). "Neuromagnetic evidence for early semantic access in word recognition". In: *European Journal of Neuroscience* 13.1, pp. 201–205.

Pylkkänen, Liina and Alec Marantz (2003). "Tracking the time course of word recognition with MEG". In: *Trends in Cognitive Science* 7, pp. 187–189.

Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik (1985). *A comprehensive grammar of contemporary English*.

Rastle, Kathleen and Matthew H. Davis (2008). "Morphological decomposition based on the analysis of orthography". In: *Language and Cognitive Processes* 23.7-8, pp. 942–971.

Rastle, Kathleen, Matt H. Davis, William D. Marslen-Wilson, and Lorraine K. Tyler (2000). "Morphological and semantic effects in visual word recognition: A time-course study". In: *Language and cognitive processes* 15.4-5, pp. 507–537.

Rastle, Kathleen, Matthew H. Davis, and Boris New (2004). "The broth in my brothers brothel: Morpho-orthographic segmentation in visual word recognition". In: *Psychonomic Bulletin & Review* 11.6, pp. 1090–1098.

Rumelhart, David E. and James L. McClelland (1982). "An interactive activation model of context effects in letter perception: II. The contextual enhancement effect and some tests and extensions of the model." In: *Psychological review* 89.1, p. 60.

Saffran, Jenny R., Richard N. Aslin, and Elissa L. Newport (1996a). “Statistical learning by 8-month-old infants”. In: *Science* 274.5294, pp. 1926–1928.

Saffran, Jenny R., Elissa L. Newport, and Richard N. Aslin (1996b). “Word segmentation: the role of distributional cues”. In: *Journal of memory and language* 35, pp. 606–621.

Sánchez-Gutiérrez, Claudia H., Hugo Mailhot, Hél‘ene S. Deacon, and Maximiliano A. Wilson (2018). “MorphoLex: A derivational morphological database for 70,000 English words”. In: *Behavior Research Methods* 50, pp. 1568–1589. DOI: `10.3758/s13428-017-0981-8`.

Schreuder, Robert and R. Harald Baayen (1995). “Modeling morphological processing”. In: *Morphological aspects of language processing*. Ed. by Laurie Beth Feldman. Hillsdale, NJ: Lawrence Erlbaum, pp. 131–154.

Schultink, Henk (1961). “Produktiviteit als morfologisch fenomeen”. In: *Forum der letteren* 2, pp. 110–125.

Schütze, Carson T. (2016). *The empirical base of linguistics: Grammaticality judgments and linguistic methodology*. Vol. 3. Classics in Linguistics. Language Science Press.

Schütze, Carson T and Jon Sprouse (2014). “Judgment data”. In: *Research methods in linguistics*. Ed. by Robert J. Podesva and Devyani Sharma. Cambridge University Press Cambridge, pp. 27–50.

Seidenberg, Mark S. (1987). “Sublexical structures in visual word recognition: Access units or orthographic redundancy?” In: *Attention and performance 12: The psychology of reading*. Ed. by M. Coltheart. Hillsdale, NJ: Erlbaum.

Seidenberg, Mark S. and James L. McClelland (1989). “A distributed, developmental model of word recognition and naming.” In: *Psychological review* 96.4, p. 523.

Sereno, S. C., K. Rayner, and M. I. Posner (1998). “Establishing a time-line of word recognition: Evidence from eye movements and event-related potentials”. In: *NeuroReport* 9.10, pp. 2195–2200.

Siegel, Dorothy Carla (1974). “Topics in English morphology.” PhD thesis. Massachusetts Institute of Technology.

Solomyak, Olla and Alec Marantz (2009). “Lexical access in early stages of visual word processing: A single-trial correlational MEG study of heteronym recognition”. In: *Brain and Language* 118, pp. 191–196.

— (2010). “Evidence for early morphological decomposition in visual word recognition”. In: *Journal of Cognitive Neuroscience* 22.9, pp. 2042–2057.

Sprouse, Jon (in press). “A user’s view of the validity of acceptability judgments as evidence for syntactic theories”. In: *Linguistic Intuitions*. Ed. by Samuel Schindler, Anna Drozdzowicz, and Karen Brøcker. Oxford University Press.

Stockall, Linnaea and Alec Marantz (2006). “A single route, full decomposition model of morphological complexity: MEG evidence”. In: *The mental lexicon* 1.1, pp. 85–123.

Stump, Gregory T. (2001). *Inflectional morphology: A theory of paradigm structure*. Vol. 93. Cambridge University Press.

Taft, Marcus (1979). “Recognition of affixed words and the word frequency effect”. In: *Memory & Cognition* 7.4, pp. 263–272.

— (1994). “Interactive-activation as a framework for understanding morphological processing”. In: *Language and cognitive processes* 9.3, pp. 271–294.

— (2003). “Morphological representation as a correlation between form and meaning”. In: *Reading complex words*. Ed. by E. Assink and D. Sandra. Amsterdam: Kluwer, pp. 113137.

Taft, Marcus (2004). “Morphological decomposition and the reverse base frequency effect”. In: *The Quarterly Journal of Experimental Psychology Section A* 57.4, pp. 745–765.

Taft, Marcus and Kenneth I. Forster (1975). “Lexical storage and retrieval of prefixed words”. In: *Journal of verbal learning and verbal behavior* 14.6, pp. 638–647.

Taft, Marcus and Gail Hambly (1986). “Exploring the cohort model of spoken word recognition”. In: *Cognition* 22.3, pp. 259–282.

Taft, Marcus and Minh Nguyen-Hoan (2010). “A sticky stick? The locus of morphological representation in the lexicon”. In: *Language and Cognitive Processes* 25.2, pp. 277–296.

Vinckier, F., S. Dehaene, A. Jobert, J. P. Dubus, M. Sigman, and L. Cohen (2007). “Hierarchical coding of letter strings in the ventral stream: Dissecting the inner organization of the visual word-form system”. In: *Neuron* 55.1, pp. 143–156.

Vitevitch, Michael S. and Paul A. Luce (1998). “When words compete: Levels of processing in perception of spoken words”. In: *Psychological science* 9.4, pp. 325–329.

— (1999). “Probabilistic phonotactics and neighborhood activation in spoken word recognition”. In: *Journal of memory and language* 40.3, pp. 374–408.

Vitevitch, Michael S., Paul A. Luce, David B. Pisoni, and Edward T. Auer (1999). “Phonotactics, neighborhood activation, and lexical access for spoken words”. In: *Brain and language* 68.1-2, pp. 306–311.

Wagenmakers, Eric-Jan (2007). “A practical solution to the pervasive problems ofp values”. In: *Psychonomic bulletin & review* 14.5, pp. 779–804.

Zehr, Jeremy and Florian Schwarz (2018). *PennController for Internet Based Experiments (IBEX)*. DOI: `10.17605/OSF.IO/MD832`.

Zweig, Eytan and Liina Pylkkänen (2008). “A visual M170 effect of morphological complexity”. In: *Language and Cognitive Processes* 24, pp. 412–439.

Appendices – Wordlists

I. Experiment 1 - syllabicity/root priming

IDENTITY CONDITION			TRANSPARENT CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
ace	gig	ACE	acidic	touchy	ACID
angel	mount	ANGEL	aiming	mildly	AIM
ant	cue	ANT	boneless	courting	BONE
bay	oil	BAY	clueless	avoiding	CLUE
cow	ink	COW	crying	shower	CRY
dust	wave	DUST	dearly	coping	DEAR
eat	fan	EAT	dipping	erupted	DIP
elm	beg	ELM	dreadful	labeling	DREAD
eye	fit	EYE	drying	sturdy	DRY
fail	bird	FAIL	egoism	weeded	EGO
fund	tank	FUND	fatter	patted	FAT
fuss	deed	FUSS	flawless	sweating	FLAW
hat	cod	HAT	foggy	widen	FOG
hip	bog	HIP	foolish	clarify	FOOL
jam	ill	JAM	harmless	rotating	HARM
joke	pure	JOKE	honorable	strangely	HONOR
knot	plea	KNOT	joyful	cooker	JOY
lap	wax	LAP	lawful	binder	LAW
lock	suit	LOCK	muddy	aided	MUD
pot	hay	POT	oddly	rises	ODD
pub	guy	PUB	padding	lodging	PAD
push	core	PUSH	plugging	uniquely	PLUG
query	skate	QUERY	ranking	amended	RANK
rat	gap	RAT	sadness	freshly	SAD

rug	kin	RUG	sharper	poorest	SHARP
shirt	prize	SHIRT	skiing	jumper	SKI
spin	rent	SPIN	spying	hunted	SPY
spot	rose	SPOT	tagging	hurtful	TAG
tip	row	TIP	madness	sticker	MAD
toy	sit	TOY	trimming	promoter	TRIM
tub	eel	TUB	weaken	nicest	WEAK
wig	rum	WIG	wisely	loosen	WISE

OPAQUE CONDITION			SYLLABIC CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
archer	ranges	ARCH	apartheid	closeness	APART
belly	hated	BELL	banjo	silky	BAN
brandy	softer	BRAND	boycott	leaning	BOY
butter	handed	BUTT	bravo	faded	BRA
choppy	subtly	CHOP	bungee	wearer	BUN
corny	taker	CORN	canvas	rented	CAN
department	management	DEPART	captain	highest	CAP
earnest	crooked	EARN	carcass	halting	CAR
figment	vividly	FIG	cashew	sicker	CASH
flicker	cheaply	FLICK	counterfeit	structuring	COUNTER
former	moving	FORM	dogma	poses	DOG
gullible	creeping	GULL	enterprise	accounting	ENTER
gutter	raided	GUT	extradite	absorbent	EXTRA
hammer	masses	HAM	ginseng	induces	GIN
inner	races	INN	goblin	poster	GOB
irony	saver	IRON	gumbo	wiper	GUM
kitten	tasted	KIT	haggis	bagger	HAG
ladder	greedy	LAD	hummus	exited	HUM
luster	truism	LUST	lettuce	blender	LET
matter	saying	MAT	maintain	employer	MAIN
muggy	gazer	MUG	mantis	setter	MAN
optic	finer	OPT	marquis	sensual	MAR
petty	loser	PET	panda	bulky	PAN

pigment	quieter	PIG	parka	roomy	PAR
ponder	bosses	POND	plaintiff	columnist	PLAIN
rustic	eluded	RUST	plantain	availing	PLAN
tender	openly	TEND	pundit	redden	PUN
topic	older	TOP	sincere	tightly	SIN
tractable	gradation	TRACT	summit	dating	SUM
treaty	cosmic	TREAT	tabloid	sleeper	TAB
wander	sucker	WAND	tandem	baking	TAN
witness	visited	WIT	tennis	wasted	TEN

NON-SYLLABIC CONDITION			WORD-WORD FILLERS		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
bark	suns	BAR	–	leaps	SIR
beet	sobs	BEE	–	phony	HERB
blurb	urges	BLUR	–	sands	RAY
charge	places	CHAR	–	frogs	WORM
dense	wiped	DEN	–	marker	FEE
drawl	casks	DRAW	–	noses	FOND
farce	monks	FAR	–	floods	POD
fleece	swears	FLEE	–	fences	TIDE
furl	eked	FUR	–	echoes	PIE
gasp	hubs	GAS	–	cooled	LEAN
growl	soils	GROW	–	skirts	HUG
hemp	tits	HEM	–	fluffy	NAIL
overt	moons	OVER	–	ounces	SHY
pawn	jaws	PAW	–	paints	DULL
peer	buys	PEE	–	danced	GEM
pinch	joins	PIN	–	lacked	CLAY
prompt	stored	PROM	–	clicks	CAB
ramp	dots	RAM	–	hopping	FLESH
runt	mops	RUN	–	shaving	STEAL
scarce	merged	SCAR	–	carrots	KID
seam	nuns	SEA	–	broadly	ASH
sight	holds	SIGH	–	tissues	PAY

starch	tacked	STAR	–	wrongly	OAK
stunt	gangs	STUN	–	bananas	WEB
tart	iced	TAR	–	hopeful	BOW
tease	rainy	TEA	–	hardness	DIG
twinge	silken	TWIN	–	indicted	POUR
warmth	smiles	WAR	–	boxer	EAR
weep	figs	WEE	–	chatting	CHESS
whirl	moths	WHIR	–	booking	TIN
yawn	lays	YAW	–	stunned	JET
zoom	sins	ZOO	–	drowned	RAW

II. Experiment 2 - dominance/affix priming

SGSG-CONDITION			PLSG-CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
sarcasms	canonize	GENDERS	windows	painted	GODS
panoramas	flautist	CHAPTERS	mittens	arching	FESTIVALS
careers	bundled	WAYS	gallows	neuronal	ASSASSINS
manners	eternity	FELLOWS	dioramas	excitable	PERIODS
morals	vaguely	EMPERORS	halogens	royalist	VIOLINS
carpenters	constipated	TAVERNS	sandals	modular	RIVERS
goods	mailed	TOWNS	qualms	crusty	MANSIONS
valleys	sobbing	DILEMMAS	artisans	classify	KINDS
worlds	assumed	HEAVENS	trousers	optimist	OCEANS
councils	inscribed	CEILINGS	stairs	dental	BISHOPS
colonels	possessor	CAMELS	scissors	follower	BEERS
kitchens	vengeful	ROADS	utensils	murmured	WAGONS
acorns	rougher	LIZARDS	binoculars	persuasive	CANALS
cancers	booming	PERSONS	lentils	coinage	DRAGONS
fountains	tormented	DIAMONDS	pajamas	rearing	CRYSTALS
humans	managed	CREEDS	nostrils	frighten	SOULS
domains	collision	DOGMAS	caverns	polished	GUITARS
husbands	coastal	MASONS	ancestors	prevention	DAUGHTERS
operas	snowy	PASSIONS	mammals	brushed	LUNGS
umbrellas	adornment	PROTEINS	thighs	removal	COUSINS

er-CONDITION			RHYME CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
scorer	bulging	STALKER	ballad	muffin	SALAD
forger	butchered	DANCER	beetle	pupil	NEEDLE
shipper	outing	BINDER	casket	noxious	BASKET
jogger	excused	FREEZER	cuddle	midget	SUBTLE
booker	hostess	HUNTER	petal	convey	MEDAL
grabber	brainless	PREACHER	deceit	cripple	RECEIPT
whiner	sealing	TRAINER	fennel	mutton	KENNEL
thriller	chopping	CALLER	ferry	violin	CHERRY

reaper	weaning	MARKER	hurdle	apron	TURTLE
cooker	grassy	DRUMMER	lactic	garnish	TACTIC
rancher	toxic	BUILDER	lotion	fissure	NOTION
golfer	steely	BREEDER	marrow	gradient	NARROW
gunner	likeness	RUNNER	zealous	devour	JEALOUS
stroller	panicked	SENDER	mumble	sundry	HUMBLE
drinker	fondness	STICKER	rattle	violet	CATTLE
learner	stressing	SCANNER	ravage	troupe	SAVAGE
hatter	tickled	FOUNDER	saloon	purport	BALLOON
solver	warping	JUMPER	taper	mimic	VAPOR
wrangler	sewage	LOVER	tumor	satin	RUMOR
sniper	bedding	WINNER	wallow	torrent	HOLLOW

TRANSPARENT CONDITION			OPAQUE CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
aiming	mildly	AIM	archer	ranges	ARCH
clueless	avoiding	CLUE	belly	hated	BELL
dearly	coping	DEAR	brandy	softer	BRAND
dipping	erupted	DIP	butter	handed	BUTT
drying	sturdy	DRY	choppy	subtly	CHOP
fatter	patted	FAT	corny	taker	CORN
flawless	sweating	FLAW	earnest	crooked	EARN
foggy	widen	FOG	former	moving	FORM
foolish	clarify	FOOL	hammer	masses	HAM
joyful	cashed	JOY	inner	races	INN
madness	treasury	MAD	irony	saver	IRON
muddy	aided	MUD	kitten	tasted	KIT
oddly	rises	ODD	luster	truism	LUST
sadness	freshly	SAD	petty	loser	PET
sharper	poorest	SHARP	pigment	quieter	PIG
spying	hunted	SPY	ponder	bosses	POND
tagging	hurtful	TAG	tender	openly	TEND
trimming	promoter	TRIM	topic	older	TOP
weaken	nicest	WEAK	treaty	cosmic	TREAT

wisely	loosen	WISE	wander	sucker	WAND

SYLLABIC CONDITION			NON-SYLLABIC CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
banjo	silky	BAN	bark	suns	BAR
boycott	leaning	BOY	beet	sobs	BEE
bravo	faded	BRA	blurb	casts	BLUR
bungee	wearer	BUN	drawl	casks	DRAW
canvas	rented	CAN	farce	monks	FAR
captain	highest	CAP	fleece	swears	FLEE
carcass	halting	CAR	furl	eked	FUR
cashew	sicker	CASH	gasp	hubs	GAS
dogma	poses	DOG	dense	soils	DEN
gumbo	wiper	GUM	overt	moons	OVER
lettuce	blender	LET	pawn	rescued	PAW
maintain	employer	MAIN	peer	buys	PEE
mantis	setter	MAN	pinch	joins	PIN
panda	bulky	PAN	prompt	stored	PROM
pundit	redden	PUN	ramp	dots	RAM
sincere	tightly	SIN	runt	mops	RUN
summit	dating	SUM	seam	nuns	SEA
tabloid	sleeper	TAB	starch	tacked	STAR
tandem	baking	TAN	tart	iced	TAR
tennis	wasted	TEN	tease	grinds	TEA

III. Experiment 3 - *in*-*im*/visual

in-CONDITION			*im*-CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
incautious	submersion	INSINCERE	imperishable	oversimplify	IMPETUOUS
inconstant	freeloader	INFERTILE	impassable	superceded	IMPROPER
indiscreet	detraction	INVARIANT	impotent	synonyms	IMMORAL
invariable	desecrated	INHUMANE	immodest	castaway	IMPOSSIBLE
inclement	uselessly	INELIGIBLE	impermanent	admonishing	IMMEDIATE
indignity	embodying	INANIMATE	impeccable	antithesis	IMPLACABLE
incontinent	premonition	INFREQUENT	immaculate	antimatter	IMMATERIAL
inedible	emanated	INCOHERENT	immemorial	asymmetric	IMPURE
inattention	exportation	INTREPID	impregnable	expropriate	IMPROBABLE
inoffensive	multiracial	INEPT	impecunious	additionally	IMPIOUS
inscrutable	malevolence	INDIFFERENT	impunity	emissary	IMPRACTICAL
inexorable	eradicated	INDEFINITE	impropriety	telegraphic	IMPATIENT

in-im CONDITION			*im-in* CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
indisposed	realigning	IMPERMEABLE	impair	upshot	INTOLERANCE
insensible	commissary	IMPENETRABLE	imbalance	evergreen	INEFFICIENT
insolence	walkabout	IMPRECISE	imperturbable	reinstitution	INDIRECT
indistinct	nonchalant	IMPRUDENT	imponderable	demilitarize	INACTIVE
inexact	alloted	IMPERVIOUS	impudent	ferryman	INACCURATE
infamy	awaken	IMMATURE	impolitic	sublimity	INCAPABLE
incoherence	expositions	IMPOLITE	implausible	obnoxiously	INSECURE
inoperative	denominated	IMMODERATE	impertinent	visitations	INCONSIDERATE
incalculable	teleological	IMPERSONAL	immovable	reelected	INFLEXIBLE
indecision	antecedent	IMPARTIAL	immeasurable	cohabitation	INADEQUACY
inequity	monotony	IMMORTAL	immobile	undersea	INFALLIBLE
incorruptible	premeditation	IMPERFECT	immutable	anarchism	INDELIBLE

dis-CONDITION			COHORT CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
dislocate	resurgent	DISADVANTAGE	scorekeeper	transpiring	SCORPIONS

displease	blockhead	DISCONTENT	distributive	attractively	DISTRACTION
disentangle	overhauling	DISALLOW	immersion	weakening	IMMIGRANT
disembark	redoubled	DISAPPOINT	matriarchy	veneration	MATERIALIZE
disillusion	equidistant	DISRESPECT	personified	derivations	PERSISTENT
discordant	headlining	DISHONEST	appraisals	forerunner	APPALLING
disabuse	reverent	DISBELIEF	corrugated	intramural	CORNERSTONE
disloyal	bookworm	DISCOMFORT	venerated	antipathy	VENTILATION
disorient	convexity	DISABILITY	meditative	ornamental	MEDICATION
discontinuous	transposition	DISASSEMBLE	tendering	vainglory	TENTACLES
disengage	interlude	DISINFECT	boarded	blowout	BOREDOM
dishonor	eviction	DISINTEREST	abstinent	underpass	ABSURDITY

TRANSPARENT CONDITION			OPAQUE CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
fluently	lectured	FLUENT	corny	widen	CORN
tangled	cyclist	TANGLE	former	moving	FORM
resentful	parentage	RESENT	gutter	raided	GUT
silky	greasy	SILK	luster	truism	LUST
governance	suggestive	GOVERN	pigment	quieter	PIG
sketchy	weaning	SKETCH	figment	vividly	FIG
crunchy	weeping	CRUNCH	muggy	gazer	MUG
departing	regretted	DEPART	tractable	gradation	TRACT
reckoning	luxurious	RECKON	skewer	berate	SKEW
cowardice	headlight	COWARD	seedy	fuses	SEED
resigning	stealer	RESIGN	crafty	welded	CRAFT
sorrowful	scrambler	SORROW	ponder	bosses	POND

$[-M + S + O]$ CONDITION		
related prime	*unrelated prime*	*TARGET*
flood	punch	FLOAT
flutter	draught	FLURRY
freeze	roof	FROST
ghost	clay	GHOUL
gleam	nymph	GLINT
plunge	kettle	PLUMMET
scald	drawl	SCORCH
scrape	thrash	SCRATCH
screech	wrinkle	SCREAM
shelve	lentil	SHELF
shrivel	pontiff	SHRINK
slither	startle	SLINK

IV. Experiment 4 - illicit/rating

Roots and target items

	ROOT	ABLE	ITY	MENT	NESS	FREQUENCY	LENGTH
LOW FREQUENCY BIN $mean = 6.53, s.d. = 0.16$	loft	loftable	loftity	loftment	loftness	6.62	4
	moth	mothable	mothity	mothment	mothness	6.24	4
	abbot	abbotable	abbotity	abbotment	abbotness	6.42	5
	wasp	waspable	waspity	waspment	waspness	6.67	4
	glee	gleeable	gleeity	gleement	gleeness	6.45	4
	dusk	duskable	duskity	duskment	duskness	6.68	4
	gait	gaitable	gaitity	gaitment	gaitness	6.48	4
	brow	browable	browity	browment	browness	6.69	4
MID FREQUENCY BIN $mean = 8.18, s.d. = 0.06$	trunk	trunkable	trunkity	trunkment	trunkness	8.09	5
	dough	doughable	doughity	doughment	doughness	8.13	5
	tray	trayable	trayity	trayment	trayness	8.13	4
	poet	poetable	poetity	poetment	poetness	8.15	4
	moss	mossable	mossity	mossment	mossness	8.22	4
	flaw	flawable	flawity	flawment	flawness	8.24	4
	crypt	cryptable	cryptity	cryptment	cryptness	8.26	5
	silk	silkable	silkity	silkment	silkness	8.26	5
HIGH FREQUENCY BIN $mean = 10.11, s.d. = 0.85$	folk	folkable	folkity	folkment	folkness	9.39	4
	lung	lungable	lungity	lungment	lungness	9.51	4
	desk	deskable	deskity	deskment	deskness	9.52	4
	debt	debtable	debtity	debtment	debtness	9.56	4
	skill	skillable	skillity	skillment	skillness	9.82	5
	soul	soulable	soulity	soulment	soulness	10.38	4
	luck	luckable	luckity	luckment	luckness	11.01	4
	law	lawable	lawity	lawment	lawness	11.74	3

Filler items

FILLER	ITEM TYPE	RAW LIKERT	REFERENCE	EST. LIKERT
tighw	phonotactic	1	Hayes and White (2013)	1
sadekp	phonotactic	1.2	Hayes and White (2013)	
kmeppid	phonotactic	1.3	Daland et al. (2011)	
lmottiff	phonotactic	1.3	Daland et al. (2011)	
nlezzig	phonotactic	1.35	Daland et al. (2011)	
vnet	phonotactic	1.88	Hayes and White (2013)	2
zreppid	phonotactic	2.05	Daland et al. (2011)	
zhmat	phonotactic	2.05	Hayes and White (2013)	
zrossip	phonotactic	2.1	Daland et al. (2011)	
pwudge	phonotactic	2.89	Hayes and White (2013)	3
sweegiff	phonotactic	3	Daland et al. (2011)	
shleebid	phonotactic	3.1	Daland et al. (2011)	
snold	morphological	2.83	Albright and Hayes (2003)	
krezzig	phonotactic	3.3	Daland et al. (2011)	
spuck	morphological	3.96	Albright and Hayes (2003)	4
dapt	morphological	4	Albright and Hayes (2003)	
gleeded	morphological	4.22	Albright and Hayes (2003)	
shurnt	morphological	4.22	Albright and Hayes (2003)	
glitted	morphological	5	Albright and Hayes (2003)	5
splung	morphological	5.45	Albright and Hayes (2003)	
bized	morphological	5.3	Albright and Hayes (2003)	
stinned	morphological	5.3	Albright and Hayes (2003)	
sneck	phonotactic	5.13	Hayes and White (2013)	
tunked	morphological	5.67	Albright and Hayes (2003)	6
kilp	phonotactic	5.9	Hayes and White (2013)	
stired	morphological	6	Albright and Hayes (2003)	
panked	morphological	6.3	Albright and Hayes (2003)	
trisk	phonotactic	6.95	Hayes and White (2013)	7
wissed	morphological	6.57	Albright and Hayes (2003)	
murned	morphological	6.57	Albright and Hayes (2003)	
shurned	morphological	6.57	Albright and Hayes (2003)	
gezzed	morphological	6.61	Albright and Hayes (2003)	

V. Experiment 5 - licit/root priming

IDENTITY CONDITION			*ful*-CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
seal	bind	SEAL	doubtful	intriguing	DOUBT
doll	lazy	DOLL	successful	protection	SUCCESS
vague	mason	VAGUE	truthful	citizenry	TRUTH
oppose	ghoul	OPPOSE	fearful	remedies	FEAR
orbit	verse	ORBIT	hateful	padding	HATE
shore	assign	SHORE	faithful	routinely	FAITH
grasp	cloud	GRASP	scornful	bloodless	SCORN
cheat	genius	CHEAT	sorrowful	pyramidal	SORROW
canon	tear	CANON	shameful	rhythmic	SHAME
habit	buck	HABIT	skillful	lecturing	SKILL
patrol	burst	PATROL	graceful	thrower	GRACE
thumb	punk	THUMB	cheerful	shuffling	CHEER
grain	lamp	GRAIN	lustful	busiest	LUST
meal	raid	MEAL	spiteful	sprouted	SPITE
abroad	blonde	ABROAD	youthful	narrower	YOUTH
arrest	nurse	ARREST	blissful	shredder	BLISS
maze	paste	MAZE	purposeful	scandalous	PURPOSE
prism	bang	PRISM	wasteful	avoidance	WASTE
rabbit	carbon	RABBIT	tasteful	secretive	TASTE
climb	tennis	CLIMB	mournful	trifling	MOURN
yield	invite	YIELD	dreadful	scratched	DREAD
thesis	barrel	THESIS	fruitful	sniffing	FRUIT
panic	reed	PANIC	frightful	inhibitory	FRIGHT
garage	clause	GARAGE	stressful	diffusion	STRESS

able-CONDITION			*ity*-CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
laughable	attendant	LAUGH	sanity	dangers	SANE
advisable	parenthood	ADVISE	complexity	emerging	COMPLEX
honorable	referral	HONOR	serenity	wrinkled	SERENE
amendable	wallowed	AMEND	absurdity	spherical	ABSURD

washable	theorize	WASH	lucidity	inquirer	LUCID
quotable	humorist	QUOTE	agility	nicest	AGILE
commendable	vacationing	COMMEND	virginity	exploiting	VIRGIN
variable	greatest	VARY	maturity	broader	MATURE
detectable	timeless	DETECT	purity	nominal	PURE
curable	frontage	CURE	fluidity	tabulate	FLUID
accountable	dimensional	ACCOUNT	humidity	sentiments	HUMID
debatable	courageous	DEBATE	stupidity	forbidden	STUPID
dependable	currencies	DEPEND	obesity	flawless	OBESE
traceable	silencer	TRACE	divinity	poisoned	DIVINE
suitable	resistance	SUIT	fragility	fugitives	FRAGILE
attainable	affliction	ATTAIN	captivity	abdominal	CAPTIVE
observable	possessing	OBSERVE	density	sooner	DENSE
favorable	adequately	FAVOR	chastity	parodies	CHASTE
admirable	descendant	ADMIRE	rigidity	slippage	RIGID
pardonable	spiritless	PARDON	oddity	airless	ODD
definable	lengthwise	DEFINE	solidity	populism	SOLID
manageable	imposition	MANAGE	validity	symbolic	VALID
cashable	improviser	CASH	rarity	louder	RARE
adaptable	ascension	ADAPT	scarcity	thematic	SCARCE

ment-CONDITION			*ness*-CONDITION		
related prime	*unrelated prime*	*TARGET*	*related prime*	*unrelated prime*	*TARGET*
appeasement	abortionist	APPEASE	weakness	container	WEAK
shipment	planetary	SHIP	numbness	saintly	NUMB
judgement	surprising	JUDGE	correctness	disposal	CORRECT
allotment	registries	ALLOT	thickness	swallowed	THICK
commandment	funniest	COMMAND	baldness	childless	BALD
puzzlement	relishing	PUZZLE	awareness	distinction	AWARE
movement	followed	MOVE	blindness	strictest	BLIND
assessment	friendship	ASSESS	harshness	sparsely	HARSH
amazement	ruthless	AMAZE	fitness	quicker	FIT
abridgment	prosodic	ABRIDGE	brightness	traveler	BRIGHT
placement	specialist	PLACE	calmness	vicarious	CALM

arrangement	employed	ARRANGE	rudeness	mobilize	RUDE
adornment	umbrellas	ADORN	soreness	molesting	SORE
advancement	attraction	ADVANCE	toughness	visionary	TOUGH
agreement	columns	AGREE	gentleness	voiceless	GENTLE
harrassment	seduction	HARASS	fairness	yearly	FAIR
annulment	murmuring	ANNUL	aloofness	flutist	ALOOF
amusement	ignition	AMUSE	crispness	victimize	CRISP
attachment	explosive	ATTACH	stubbornness	accentuate	STUBBORN
achievement	questioned	ACHIEVE	madness	spectral	MAD
adjustment	veterans	ADJUST	darkness	election	DARK
settlement	structural	SETTLE	politeness	snapping	POLITE
atonement	trombonist	ATONE	awkwardness	subjection	AWKWARD
alignment	compliance	ALIGN	boldness	fiendish	BOLD

VI. Experiment 6 - illicit/root priming

[RP: Maybe it would be enough to put the roots and the unrelated primes in two separate tables? It will certainly make the list shorter.]

IDENTITY	ABLE	ITY	MENT	NESS	UNRELATED	*TARGET*
moth	mothable	mothity	mothment	mothness	trainee	MOTH
wool	woolable	woolity	woolment	woolness	patronize	WOOL
skull	skullable	skullity	skullment	skullness	unionist	SKULL
cartel	cartelable	cartelity	cartelment	cartelness	divinely	CARTEL
frog	frogable	frogity	frogment	frogness	brighten	FROG
throat	throatable	throatity	throatment	throatness	shaky	THROAT
brow	browable	browity	browment	browness	lessen	BROW
tray	trayable	trayity	trayment	trayness	falsify	TRAY
creek	creekable	creekity	creekment	creekness	sharpen	CREEK
ox	oxable	oxity	oxment	oxness	frontal	OX
virtue	virtueable	virtueity	virtuement	virtueness	ruthless	VIRTUE
folk	folkable	folkity	folkment	folkness	straighten	FOLK
carrot	carrotable	carrotity	carrotment	carrotness	heaviest	CARROT
soup	soupable	soupity	soupment	soupness	childless	SOUP
debt	debtable	debtity	debtment	debtness	selfless	DEBT
porch	porchable	porchity	porchment	porchness	overtly	PORCH
yeast	yeastable	yeastity	yeastment	yeastness	lengthen	YEAST
sparrow	sparrowable	sparrowity	sparrowment	sparrowness	trainee	SPARROW
reef	reefable	reefity	reefment	reefness	patronize	REEF
mall	mallable	mallity	mallment	mallness	unionist	MALL
spinach	spinachable	spinachity	spinachment	spinachness	divinely	SPINACH
moss	mossable	mossity	mossment	mossness	brighten	MOSS
relief	reliefable	reliefity	reliefment	reliefness	shaky	RELIEF
yacht	yachtable	yachtity	yachtment	yachtness	lessen	YACHT
grief	griefable	griefity	griefment	griefness	falsify	GRIEF
cult	cultable	cultity	cultment	cultness	sharpen	CULT
calf	calfable	calfity	calfment	calfness	frontal	CALF
gulf	gulfable	gulfity	gulfment	gulfness	ruthless	GULF
luck	luckable	luckity	luckment	luckness	straighten	LUCK

flint	flintable	flintity	flintment	flintness	heaviest	FLINT
ghoul	ghoulable	ghoulity	ghoulment	ghoulness	childless	GHOUL
almond	almondable	almondity	almondment	almondness	selfless	ALMOND
salad	saladable	saladity	saladment	saladness	overtly	SALAD
dirt	dirtable	dirtity	dirtment	dirtness	lengthen	DIRT
skit	skitable	skitity	skitment	skitness	trainee	SKIT
poet	poetable	poetity	poetment	poetness	patronize	POET
bread	breadable	breadity	breadment	breadness	unionist	BREAD
squid	squidable	squidity	squidment	squidness	divinely	SQUID
wrist	wristable	wristity	wristment	wristness	brighten	WRIST
ear	earable	earity	earment	earness	shaky	EAR
hen	henable	henity	henment	henness	lessen	HEN
rhythm	rhythmable	rhythmity	rhythmment	rhythmness	falsify	RHYTHM
symbol	symbolable	symbolity	symbolment	symbolness	sharpen	SYMBOL
idol	idolable	idolity	idolment	idolness	frontal	IDOL
elf	elfable	elfity	elfment	elfness	ruthless	ELF
cohort	cohortable	cohortity	cohortment	cohortness	straighten	COHORT
bliss	blissable	blissity	blissment	blissness	heaviest	BLISS
sheep	sheepable	sheepity	sheepment	sheepness	childless	SHEEP
wasp	waspable	waspity	waspment	waspness	selfless	WASP
trunk	trunkable	trunkity	trunkment	trunkness	overtly	TRUNK
bull	bullable	bullity	bullment	bullness	lengthen	BULL
tabloid	tabloidable	tabloidity	tabloidment	tabloidness	trainee	TABLOID
herb	herbable	herbity	herbment	herbness	patronize	HERB
chess	chessable	chessity	chessment	chessness	unionist	CHESS
chef	chefable	chefity	chefment	chefness	divinely	CHEF
lizard	lizardable	lizardity	lizardment	lizardness	brighten	LIZARD
desk	deskable	deskity	deskment	deskness	shaky	DESK
conch	conchable	conchity	conchment	conchness	lessen	CONCH
oak	oakable	oakity	oakment	oakness	falsify	OAK
gait	gaitable	gaitity	gaitment	gaitness	sharpen	GAIT
violin	violinable	violinity	violinment	violinness	frontal	VIOLIN
fraud	fraudable	fraudity	fraudment	fraudness	ruthless	FRAUD
loft	loftable	loftity	loftment	loftness	straighten	LOFT

bulb	bulbable	bulbity	bulbment	bulbness	heaviest	BULB
shirt	shirtable	shirtity	shirtment	shirtness	childless	SHIRT
trough	troughable	troughity	troughment	troughness	selfless	TROUGH
crypt	cryptable	cryptity	cryptment	cryptness	overtly	CRYPT
fruit	fruitable	fruitity	fruitment	fruitness	lengthen	FRUIT
pact	pactable	pactity	pactment	pactness	trainee	PACT
affair	affairable	affairity	affairment	affairness	patronize	AFFAIR
lung	lungable	lungity	lungment	lungness	unionist	LUNG
mule	muleable	muleity	mulement	muleness	divinely	MULE
lamp	lampable	lampity	lampment	lampness	brighten	LAMP
maggot	maggotable	maggotity	maggotment	maggotness	shaky	MAGGOT
syrup	syrupable	syrupity	syrupment	syrupness	lessen	SYRUP
arrow	arrowable	arrowity	arrowment	arrowness	falsify	ARROW
dusk	duskable	duskity	duskment	duskness	sharpen	DUSK
dough	doughable	doughity	doughment	doughness	frontal	DOUGH
devil	devilable	devility	devilment	devilness	ruthless	DEVIL
crab	crabable	crabity	crabment	crabness	straighten	CRAB
flaw	flawable	flawity	flawment	flawness	heaviest	FLAW
snail	snailable	snaility	snailment	snailness	childless	SNAIL
nomad	nomadable	nomadity	nomadment	nomadness	selfless	NOMAD
turtle	turtleable	turtleity	turtlement	turtleness	overtly	TURTLE
hat	hatable	hatity	hatment	hatness	lengthen	HAT
culprit	culpritable	culpritity	culpritment	culpritness	trainee	CULPRIT
grain	grainable	grainity	grainment	grainness	patronize	GRAIN
abbot	abbotable	abbotity	abbotment	abbotness	unionist	ABBOT
farce	farceable	farceity	farcement	farceness	divinely	FARCE
mood	moodable	moodity	moodment	moodness	brighten	MOOD
quartz	quartzable	quartzity	quartzment	quartzness	shaky	QUARTZ
fossil	fossilable	fossility	fossilment	fossilness	lessen	FOSSIL
fleet	fleetable	fleetity	fleetment	fleetness	falsify	FLEET
lint	lintable	lintity	lintment	lintness	sharpen	LINT
silk	silkable	silkity	silkment	silkness	frontal	SILK
pearl	pearlable	pearlity	pearlment	pearlness	ruthless	PEARL
elk	elkable	elkity	elkment	elkness	straighten	ELK

spouse	spouseable	spouseity	spousement	spouseness	heaviest	SPOUSE
mud	mudable	mudity	mudment	mudness	childless	MUD
pint	pintable	pintity	pintment	pintness	selfless	PINT
tomb	tombable	tombity	tombment	tombness	overtly	TOMB
skill	skillable	skillity	skillment	skillness	lengthen	SKILL