Spoiler Alert:
Processing of Sentence Final Particles in Dutch

Maxime A. Tulling

s1111205

Leiden University

Faculty of Humanities

Research Master in Linguistics

MA thesis

Supervisors: Dr. L. Pablos Robles

Dr. S. Gryllia

Second reader: Prof. dr. N.O. Schiller

19 July 2016

ABSTRACT

Sentence final particles (SFPs) play an important role in the every-day spoken communication of various languages. For example, the addition of a Dutch intentional particle *hè* or *hoor* to a bare declarative utterance such as *het is lekker weer* 'the weather is great' can make the difference between the sentence being interpreted as an agreement-seeking question, or a correction. Still, we know very little about the psycho- and neurolinguistic properties of the processing and production of final particles. The purpose of this thesis is to generate more research on the psycholinguistics side and to deepen the theoretical knowledge we have by gathering experimental data.

There are theoretical reasons to assume that intentional SFPs play an important role from the beginning of sentence formulation. The SFP-head selects for the entire proposition as its complement, so it is possible that speakers plan the particle ahead before they start producing the rest of the sentence. This hypothesis also makes sense from a psycholinguistic perspective, as it is presumed that the intention of the speaker is already determined before he/she starts uttering a sentence. In this thesis, I focus on sentences in isolation, and investigate the production, planning and perception of Dutch pragmatic particles (i.e. SFPs) that convey the speaker's intention. The question I pursued to answer is whether Dutch sentence final particles are planned in advance, or whether they are inserted at the final moment. To investigate the potential planning of intentional SFPs I conducted three experiments.

In a production experiment (Experiment 1) I investigated whether the speaker already starts encoding the intention of the message with prosodic cues preceding the intentional particle. Such cues would indicate that the speaker is already building up the illocutionary force of the sentence before the particle. Results indicate that there are such cues, and that even though they are sometimes quite small, they are used quite consistently across participants. In Experiment 2, the gating-technique is used in a perception experiment to investigate whether these prosodic cues preceding the particle could possibly help the listener anticipate for the intention or attitude expressed by an utterance. The results of this experiment indicate that participants were not that good at anticipating the end of sentences containing the final particles *hè* and *hoor* in the given task. Experiment 3 directly addresses the question whether the speaker plans a final particle ahead or whether they integrate the particle at a later stage of production. This question is about how incremental and how far ahead a sentence is planned in production. In this experiment, I examined the production process of the intentional SFPs *hè* and *hoor* in Dutch with a variant on the picture-word-interference task to investigate whether the particles are planned in advance, or not. I created an experiment that manipulates the prime preceding colored pictures, which are associated in a training task with a particular final element. The target condition of the experiment contains sentences with the final elements *hè* and *hoor*, which are intentional final particles. The effect of the distractor prime on the target condition was compared to a control condition. A congruent prime was assumed to facilitate sentence production at speech onset, only if the speaker is already planning the particle congruent with the prime. In the control conditions, in which it is assumed that speakers are not yet planning ahead for the final element of the sentence at speech onset, facilitation was assumed not to take place at speech onset. The results obtained for this experiment were not significant, due to high standard deviations. In future research it would be interesting to see whether the paradigm of this experiment could be adjusted, to gain more reliable results that can answer the main question we pursued.

ACKNOWLEDGMENTS

I would like to express my deep gratitude to my supervisors Leticia Pablos Robles and Stella Gryllia for their help, useful comments,and remarks during the entire process of this master thesis. Especially their encouragements were invaluable to the research and writing of this thesis. Furthermore I would like to thank Prof. Schiller for agreeing to be a second-reader. Also, I want to thank Shotta Momma, Robert Slevc and Colin Phillips from the University of Maryland for discussing their research with me, and providing useful suggestions for the third experiment of this thesis. Special thanks to Jos Pacilly for his technical support, and to all the people that voluntarily devoted their time to participate in my experiments. Also I would like to thank Sofja Volkova, Isabella Jordanoska and Ruby Sleeman for their comments and emotional support. Last but not least, I would like to thank Gouming Martens for proofreading my thesis and his continuous support throughout the entire process.

CONTENTS

1. INTRODUCTION

In everyday conversation, contextual information plays a major role in helping speech participants communicate. In order to communicate successfully, speech participants should not only be able to understand the message their interlocutor is uttering, but they should also be able to interpret it in that specific context. Take for example the sentence displayed in (1):

(1). It's a dog.

Without any contextual information, every speaker of English is able to understand that the speaker conveys a message claiming that there is an animate being present, which should be recognized as a dog. However, without any additional information, we could think of several reasons why this sentence is being uttered. Maybe it is just the answer to a question ('What is that?'), but it could also be a reminder for example ('Why is it barking?' -> it is a dog, they behave that way) or an accusation (if a shopkeeper hears barking in his store, while dogs are not allowed to be inside). Searle (1969, 1985) argued for a distinction between these different types of meaning, and distinguished the propositional *content* of a sentence from the illocutionary *force*. Though the propositional content of an utterance is often linguistically encoded in such a way that we can infer the content from its surface form, contextual cues play a huge role in deciding what speakers actually intent with their utterance. This is illustrated in (2), where a context is provided for the example in (1):

(2). *A and B see an animal chasing a cat into the tree.*

A: Look! That rabbit chasing the cat is huge!

B: You think that's a rabbit? **It's a dog**.

From the context provided in (2), it becomes clear that the illocutionary force of the statement 'it's a dog' is that of a 'correction'. The previous example illustrates that context provides an important clue as to what the intention of an utterance is. However, speakers are not completely at loss when they are provided with a spoken sentence like (1) in isolation. In English, prosodic cues are used to convey para- and non-linguistic information[1], such as the speech act of an utterance. For example, there is a huge difference in saying (1) as an answer to the question 'What is that?', with respect to the way (1) would be pronounced in (2), where 'dog' probably would be emphasized.

Östman (1991) suggested a typological distinction between languages that use either prosodic means, like intonation, or verbal means, like pragmatic particles, as their primary means for implicitly expressing modality, attitudes, politeness and other pragmatic aspects. The link between prosody and pragmatic particles has been noticed many times before (Fujisaki & Hirose, 1993; Kirsner et al., 1996; Kwok, 1984; Wakefield, 2010; Zhang, 2014). For instance, in a language like English there are mainly prosodic cues that indicate para-linguistic information. Contrary to English, Cantonese relies heavily on the use of pragmatic particles (sentence final particles, SFPs) to express this information. However, there is a third group of languages that have the option of using either intonational means or pragmatic particles to express para-linguistic information, where one such language is Dutch. An example of a Dutch sentence with an SFP is displayed in (3).

[1] There is some confusion in the literature about the definition of para-linguistic and non-linguistic information (Schötz, 2002). In this paper I will refer to para-linguistic information when referring to prosodic information that conveys information about the intention and attitude of an utterance. The prosodic effects of emotion and physical attributes of the speaker (e.g. age or health), on the other hand, are defined as non-linguistic.

(3). Dat is een hond **hoor**.
that is a dog SFP
'That's a dog.' (Interpretation: You are wrong thinking it could be something else.)

The sentence in (3) shows that the SFP *hoor* adds an additional pragmatic interpretation to the sentence. Loosely paraphrased, this could be interpreted as 'you are wrong', and this sentence could easily be used in a context like (2). The SFP *hoor*, always appears in sentence final position. Since intentional final particles, such as *hoor*, are closely connected to the illocutionary force of the sentence, there are reasons to hypothesize that they are planned in advance, even though they appear at the end of the sentence.

From a theoretical point of view this hypothesis gains support from syntactic analyses which analyze the SFP as a head located in the CP-domain (Law, 2002; Munaro & Poletto, 2003; Speas & Tenny, 2003; Sybesma & Li, 2007), which will be discussed in more detail in Section 1.1.2. Since this sentence final particle head selects for the entire proposition as its complement, it would be reasonable to think that speakers plan the particle ahead before they start producing the rest of the sentence.

This hypothesis also makes sense from a psycholinguistic perspective, as it is presumed that the intention of the speaker is already determined before he/she starts uttering a sentence (Bock & Levelt, 2002; Gambi & Pickering, 2016). Final particles express this intention (illocutionary force), and therefore one would assume that they are planned in advance. From a production perspective, the question is whether this intention is immediately mapped into a linguistic form (e.g. the pragmatic particle is mentally activated before the lexical items of the rest of the sentence are) or whether it is only implemented at the end of the sentence, where the particle is positioned. There is an ongoing debate in the sentence production literature about the extent to which sentences are built-up incrementally, and the amount of preparation speakers engage in. There are questions about the size of chunks used for planning (e.g. Allum & Wheeldon, 2007; Brown-Schmidt & Konopka, 2008; Garrett, 1980) and the nature of message planning (Brown-Schmidt & Konopka, 2015). The literature on sentence production does not agree on whether conceptualization takes place incrementally (i.e. in a linear fashion, which implies that speakers assemble mental constructs in the same order as the word-order), or whether messages are planned holistically for the entire communicative intention, and it is the conceptual framework which guides subsequent lexical encoding.

The fact that the intentional particle appears sentence-finally is also interesting from a perception point of view. As listeners, in contrast to speakers, are dependent on the linear order of the sentence they perceive, the fact that the elements expressing the illocutionary force of a sentence appear at the end of the sentence, raises questions about the consequences this might have for sentence comprehension. Is it the case that listeners can only interpret the illocutionary force of a sentence containing these particles at the very end of the sentence, or are there any strategies listeners can use to anticipate the speaker's intention? Are there strategies based on prosody alone? We know that intention can be expressed by both intonation and particles in Dutch. It is also the case, that prosody interacts with the interpretation of the final particles. One possibility, which this thesis wants to investigate, is whether the speaker already starts encoding the intention of the message with prosodic cues preceding the intentional particle, which would mean the speaker is already building up the illocutionary force of the sentence before the particle.

In this thesis I focus on sentences in isolation, and investigate the strategies used by Dutch speakers and listeners to express and detect para-linguistic information within an utterance. I address questions about the production, planning and perception of Dutch pragmatic particles (i.e. SFPs) that convey the speaker's intention. Both pragmatic particles and prosody are understudied when it comes to their psycholinguistic processes. This is not that surprising, as both phenomena are often vague and difficult to grasp, especially when they occur together. However, studying the production and processing of pragmatic particles might provide us with a linguistic anchor to investigate the interaction between particles and prosody conveying para-linguistic information. This thesis aims to address this issue by starting out with a very basic question, concerning the relationship between the pragmatic particles and the rest of the sentence. The question I pursue to answer in the current thesis is whether Dutch sentence final particles are planned in advance, or whether they are inserted at the final moment. This question hereby contributes to the body of work that investigates the incrementality of speech production (e.g. Brown-Schmidt & Konopka, 2015; Ferreira & Swets, 2002; Momma et al., 2016, Smith & Wheeldon, 1999 among many others).

The research conducted in this thesis thus considers three different questions about the processing of intentional sentence final particles in Dutch: 1) Do speakers plan SFPs ahead or are these particles only implemented at the end? 2) Are there acoustic cues in the speech signal predicting the upcoming SFP? 3) Can listeners anticipate the upcoming SFP from certain acoustic cues, or can these particles only be interpreted at the end? These questions are addressed with three separate experiments: a variation on the picture-word-interference naming task targeting the naming of particle sentences, a production task examining the prosodic properties of sentences containing final particles, and a perceptual auditory gating experiment.

In Section 1.1, I discuss the semantic and structural properties of sentence final particles, and their relation to intonation, providing an overview of Dutch SFPs and their intonational properties. In Section 1.2, I discuss some of the existing literature on sentence production and processing in relation to this thesis, and discuss some of the issues surrounding the literature about sentence planning and anticipation. Finally, in Section 1.3, I discuss the design of the experiments performed in this thesis, in order to investigate the main question about the planning of Dutch sentence final particles.

1.1 SENTENCE FINAL PARTICLES

Originally, the term *sentence final particle* was used to describe a specific type of discourse markers in Sinitic languages (e.g. Li & Thompson, 1981), in which these particles occur clause finally. Sometimes this type of particles are also referred to as *(final) utterance particles* or *postclausal discourse markers* (Gupta, 2006). The labels that are used are descriptive, indicating that the particles discussed here occur sentence (or utterance) final, and have a discourse function. Sentence final particles play an important role in the every-day spoken communication of various languages. The meaning of a sentence final particle is often difficult to describe, even though native speakers have a strong sense for the appropriateness of a particle in a certain context. In the linguistic literature there is a lot to find about sentence final particles, but there is very little research focusing on the psycho- and neurolinguistic properties of the processing and production of final particles. The purpose of this thesis is to generate more research on the psycho/neurolinguistic side and to deepen the theoretical knowledge we have by gathering experimental data. A specific topic within the SFP-literature, the final position of SFPs, is for example both interesting from a theoretical perspective, as it is interesting from a production/processing perspective. Interestingly, the final particle does not only appear at the end in a head final language such as Japanese (Tsuchihashi, 1983), but also in a strict head initial language such as Mandarin Chinese (Law, 2002). This fact raises several questions. First of all, the question why sentence final particles 'break' the linearity principles of head-inital languages (Haegeman, 2014; Hsieh & Sybesma, 2011), but also a more trivial question: why do they appear at the end? As these final particles add an important interpretation, or intention to the proposition, thereby modifying the interpretation of the entire proposition, it is a bit puzzling that this piece of information could be saved for the end.

Before I discuss the experimental approach of this paper, I first (1.1.1) discuss in more detail the meaning and properties of the sentence final particles *hè* and *hoor*. In Subsection 1.1.2 I discuss the structural positioning of sentence final particles and provide a clear definition for the type of sentence final particles discussed throughout this thesis. In Subsection 1.1.3, I review the issue of sentence finality and the structural representation of sentence final particles in various languages. Finally, in Subsection 1.1.4 I discuss the existing literature that examines the relationship between sentence final particles and prosody.

1.1.1 Dutch Sentence Final Particles

Even though Dutch is a language that uses intonation, it is also well known for its extensive use of particles to express emotions and nuances (van der Wouden, 2006). In Dutch, particles can appear in various positions of the sentence (van der Wouden, 2006), though this study focuses only on those particles that occur in sentence final position.[2,3] The semantic and prosodic properties of Dutch sentence final particles have received a considerate amount of attention from the literature (e.g. Kirsner, 2003; Kirsner & Deen, 1990; Kirsner & van Heuven, 1996; Kirsner et al., 1994; van der Wouden, 2006; van der Wouden & Foolen, 2011; Mazeland & Plug, 2010). Nevertheless, the structural properties of Dutch SFPs have not been researched in that much detail. The research conducted by Kirsner (2003), Kirsner & Deen (1990) and Kirsner et al. (1994) mainly focused on the final particle *hoor,* though it also mentioned the particles *hè, joh* and *zeg.* Van der Wouden & Foolen (2011) analyzed the syntactic properties of SFPs in the right periphery, considering, unlike Kirsner and van Heuven (1996), a much larger set of particles that appear in the final position of the sentence. They did not propose a specific syntactic position for the SFPs in Dutch due to it being difficult to generalize one unique position for the high number of particles they examined. With regard to the aim of this thesis, I will limit myself to the Dutch final particles *hè* and *hoor*.

Hè and *hoor*

The particles *hè* and *hoor* establish a certain epistemic knowledge of the speaker. Kirsner and van Heuven (1996) claim that *hè* (4a) and *hoor* (4b) establish a relationship between the listener and speaker.

(4).	a.	Jij	komt	morgen	ook,	hoor.	
		you	come	tomorrow	too,	SFP	
		'You be sure to come tomorrow!'					(K & van H, 1996, 1a)
	b.	Jij	komt	morgen	ook,	hè?	
		you	come	tomorrow	too,	SFP	
		'You're coming tomorrow too, aren't you?'					(K & van H, 1996, 1b)

[2] Note that a particle like Dutch *toch* (comparable to a tag 'is/isn't it') can appear both sentence finally, and elsewhere in the sentence. Particles like *hè* and *hoor* however, can only appear sentence finally. In this thesis I only call particles 'sentence final' if they are from the latter type.

[3] Note that there are elements that can follow final particles, like *hè* and *hoor*, such as vocatives (I).

(I)	Jan	komt	naar	huis	hè,	Max/ *(Max, hè)?
	Jan	comes	to	home	SFP	Max
	'Jan will come home right, Max?'					

Vocatives however, do not alter the interpretation of the sentence in any way, instead they are only used to address the listener. Therefore, it seems reasonable to assume that they are detached from the sentence.

Though Kirsner and van Heuven (1996) rightfully argue that *hè* and *hoor* have conflicting semantics, I do not completely agree with their analysis with regard to the precise semantic content of sentences containing these particles. Kirsner and van Heuven (1996) claim that *hoor* and *hè* form a pair of semantic opposition, much like the English discourse markers *now* and *then* or *I mean* and *y'know*. Whereas *hè* would ask the hearer for some sort of confirmation, *hoor* would indicate that nothing of the kind is needed or wanted. However, this semantic opposition is not as clear as they state. For one, Kirsner and van Heuven (1996) compare the *hè* particle to English tag-questions, and argued it to be confirmation seeking. However, more than being a confirmation-seeking particle it is an establishment of the high degree of confidence the speaker has in its own utterance. If someone disagrees with the *hè* statement, the speaker will be surprised, as a negative answer is highly unexpected. It might therefore be better to consider it an 'agreement-seeking' particle, as has also been observed by Englert (2010). In addition, when drawing the parallel between *hè* and the English tag-questions, Kirsner and van Heuven (1996) ignore sentences in which the *hè* particle cannot be directly translated with a tag-question. (This is for example the case in 'reminding' *hè*-sentences, which will be discussed in more detail below in the subsection about multiple usages of *hè* and *hoor.*)

Also, though it is true that *hoor* indicates a high level of confidence on the side of the speaker, and does not ask for any confirmation (Kirsner & van Heuven, 1996), contra Kirsner and van Heuven (1996), the particle *hoor* does not always indicate that no confirmation is wanted or needed. In fact, in some cases a confirming response would be very natural, e.g. in 'emphasizing' *hoor*-sentences (also discussed in the following subsection).

According to Tulling (2015) the SFPs *hè* and *hoor* can both be considered to operate at an epistemic level, displaying the speaker's thoughts about probability and predictability of information. The crucial difference between the two particles is, however, that the *hè* particle is speaker-oriented while the *hoor* particle is hearer-oriented. Basically *hè* conveys 'I am right, don't you agree?' while *hoor* just says 'you are making the wrong assumption'. This contrast becomes more clear in the following examples (5) and (6):

(5). Hij houdt niet van taart **hè**?
he loves not of cake **SFP**
'He doesn't like cake, right?' (Tulling, 2015, 1)

(6). Hij houdt niet van taart **hoor**!
he loves not of cake **SFP**
'He doesn't like cake!' (Tulling, 2015, 2)

The sentence in (5) could be uttered by the host who is serving cake to her guests. She could, for example, say this to the mother of a little boy, when she remembers that the boy does not like cake. In (5) she is not really asking the mother whether the child likes cake or not, since she is quite confident herself that she remembers it correctly. With *hè* she indicates this confidence, providing an estimation about the probability of her own utterance. She believes that the chances are very high that her proposition 'The boy does not like cake.' represents the truth. A negative answer from the mother, denying the proposition that the boy does not like cake, is highly unexpected in this scenario. The sentence in (6), on the other hand, can only be uttered as a reaction to either a previous utterance or an observation, it cannot be uttered out of the blue. In a similar birthday setting, the mother of the boy could say this to the host when she observes the woman offering her son some cake. The SFP *hoor* in (6) also signals an estimation of the speaker, but this time it is the estimation that the listener is making some wrong assumptions. The host probably thinks that the boy would like some cake, but in fact he does not like cake at all.

It is important to note however, that the particles can also be interpreted differently from the scenarios that I have described above. This is discussed in more detail for the particles *hè* and *hoor* in the following section.

***Multiple Usages of* hè *and* hoor**

Besides marking the epistemic content of the speaker, the particles *hè* and *hoor* are used to express certain speech acts. In the section discussion above, I have provided examples of the most canonical examples of *hè*, as an agreement-seeker, and *hoor* as a correction. The core semantics of the particle is one of epistemic assessment (Tulling, 2015). The particle *hè* can be paraphrased as 'what I am saying is part of the common ground, don't you agree?' while the particle *hoor* can be paraphrased as 'what I am saying is not part of the common ground, since you are making the wrong assumptions'. These core meanings can have different interpretations depending on the context and prosody of the utterance. In Table 1 I displayed the 'micro-variation' within the functions of *hè* and *hoor*. For extra information and examples of these sub-functions of *hè* I refer to Appendix A.

Table 1. *Functions of Dutch Sentence Final Particles*

Particle	Main Function	Sub-functions
hè	epistemic speaker-oriented →marks that something is or should be part of the common ground	agreement-seeking confirmation reminding urging down-playing
hoor	epistemic hearer-oriented → marks that something is not part of the common ground, but the truth	correction warning reassurance emphasis command

1.1.2 The Structural Position of Sentence Final Particles

As discussed in the previous section, the Dutch particles *hè* and *hoor* are intentional particles, reflecting the speaker's epistemic estimations towards the proposition and maintaining/initiating a speaker-hearer relationship. The vast majority of literature on discourse particles agrees that final discourse particles that have such a function (i.e. expressing attitude, intention or speech acts) are located somewhere within the (split) CP-domain, occupying a head position[4]. The Split-CP hypothesis, in which the CP projection is expanded and separated into separate layers, was originally proposed by Rizzi (1997). The initial argumentation behind this split is based on

[4] This head-status can be checked with a syntactic test, proposed by Munaro and Poletto (2003), in which a parenthetical is placed between the supposed head (the particle) and its complement (the preceding utterance). Munaro & Poletto (2003) argue that it is well-known that parentheticals cannot intervene between a head and its specifier (IIa), while they can intervene between two maximal projections (IIb). The sentence in (II) shows that a parenthetical cannot intervene in Dutch:

(II).	a.	*Het	heeft	geregend,	denk	ik,	hè/hoor.
		it	has	rained	think	I	SFP
	b.	Het	heeft,	denk ik,		geregend,	hè/hoor.
		it	has	think I		rained	SFP
		'It has rained, I believe.'					

Though note (Law, 1991) and (Chan, 2013) who argue that SFPs are not heads.

different selectional properties of the CP. The complementizer is selected by the main verb, since different verbs select for different clause types (e.g. *believe* selects declarative clauses while *wonder* selects interrogative clauses). However, the complementizer selects a certain subordinate clause itself, being sensitive to finiteness (some C's select for finite clauses, while others select nonfinite clauses). Rizzi (1997) argued that these two different functions should be expressed as properties of two different heads in the CP system. In order to be selected as a complement of a main verb, a CP has to be specified for its clause type, its 'Force', while the selection requirement of the finiteness of the complementizer clause is specified in the lower projection, the Finite Phrase (FinP). Ever since then, syntacticians have found reasons and arguments to split up the CP-layer in multiple layers, not always consistent with each other. It is far beyond the scope of this thesis to go into detail about the different proposals that have been made about the different layers and ordering of these layers within CP.

What is important to realize about the split-CP domain (also called the left periphery), is that it has the structure like that of an onion. An example of an abstract representation of the split-CP hierarchy is displayed in Figure 1. The propositional content of a sentence is expressed in IP. The first layers surrounding this IP domain are functional layers related to the sentential domain (denoted with '1' in the picture), connecting the sentence to the world (place and time). Surrounding those layers, is a propositional-discourse domain ('2'), in which processes related to information structuring take place. Then, going up a level, we reach the speech act domain ('3'), where clause-typing is assumed to take place (Cheng, 1991, Munaro & Poletto, 2003). Rizzi (1997) has argued that ForceP, type-clausing the sentence, is the highest node. However, it has been argued that a distinction should be made between sentential force, which types the clause (e.g. question, imperative, declarative) and the illocutionary force projection, indicating that the clause is performed as speech act of a certain type (Chierchia & McConnel-Ginet, 1990; Zanuttini & Portner, 2003). Such distinction is very intuitive, as it is a well-known fact that for example questions can have the illocutionary force of a request (e.g. 'Can you pass the salt?'). Some would call the level above ForceP the illocutionary force domain (Chierchia & McConnel-Ginet, 1990; Güneş, 2015). Other researchers would again split this 'illocutionary force domain' ('4') into 'speaker-oriented' levels and 'hearer-oriented' levels, in which speaker-oriented particles follow hearer oriented-particles (Hill, 2013) [5].

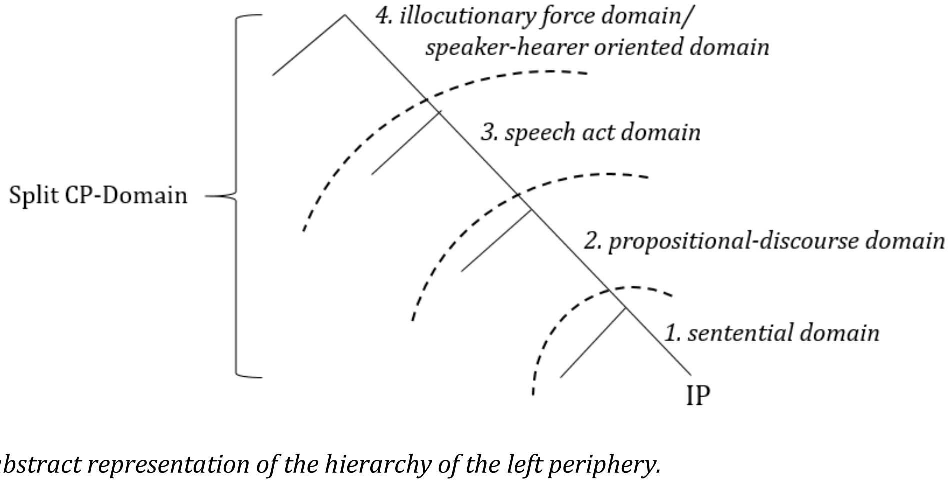

Figure 1. *An abstract representation of the hierarchy of the left periphery.*

[5] The specific details of the representation of this speaker-hearer relationship differ between researchers. Haegeman (2014) proposed that there are two speech act projections present, in which the higher projection is more directly related to the performative aspect of the speech act, initiating a hearer-speaker relation, while the lower projection modulates the (established) relationship between the speaker and hearer. Sybesma and Li (2007) propose that there are two epistemic levels present above the ForceP, also conveying information about the speaker's assumptions and opinions about the proposition and listener.

Essentially, what is relevant for this thesis is not how we specifically denote the domain located above the type-clausing domain. Important is that above the ForceP, which determines the surface force of a proposition, there is a domain that does not necessarily alter the form of the sentence or the grammatical structures within this sentence. Instead it alters the intention, attitude and feeling of a sentence. This is the domain in which speakers not only express their attitudes and evaluations about the propositional content of the sentence in relation to the world and the discourse, but also their attitude and thoughts about the speaker, and the speaker-hearer relationship. Based on previous literature positioning sentence final particle within the highest levels of the CP-domain, it is safe to assume that the Dutch particles in this thesis are also heads located within that domain.

1.1.3 Sentence Final Particles and Head Finality

As we have established in Section 1.1.2, Dutch final particles can be analyzed as syntactic heads located within the CP-domain. However, the final position of this particle raises some questions, as usually Dutch C-heads (like the complementizer) are head initial, while the sentence final particle would have to be head-final, as displayed in (7):

(7)	Marie	weet	$[_{CP1}[_{CP2}$ dat	Max	rookt]	hè]?
	Marie	knows	that	Max	smokes	SFP
	'Marie knows Max smokes, right?'					

This is not a problem to be found in Dutch linguistics only, but is also a widely discussed problem within Chinese linguistics. For Sinitic languages, various authors have argued that the final position of the sentence final particle head is actually the consequence of IP/CP-raising, in which the entire proposition moves to the Spec of the projection where the particle head is located[6] (e.g. Huang, 2007; Hsieh & Sybesma, 2011; Simpson & Wu, 2002). Similarly, this has also been suggested for other Indo-European languages, such as the Italic Veneto dialects (Munaro & Poletto, 2002; 2003), and for West-Flemish particles (Haegeman, 2014). In Figure 2, both options are illustrated (the head final view on the left, and the head first view on the right).

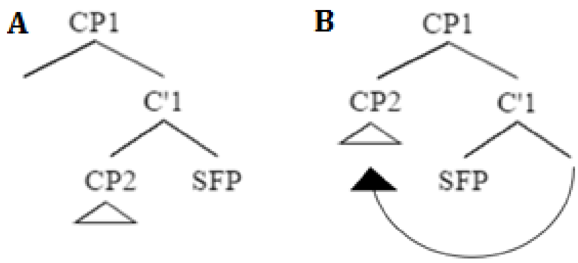

Figure 2. *The possible syntactic structures of SFP-sentences in a head-initial language.* A: *Head-final CP layer.* B: *Head-initial CP layer with CP-movement to Spec CP.*

For the aim of this thesis no assumptions are made about the 'right' underlying representation of sentences containing the final particles *hè* and *hoor* (whether they are head-initial or head-final). Both representations are compatible with the a particle planning view, assuming that sentences are produced top-down (as argued for by Phillips, 2003). However, if we find that particles are not planned in advance, then the head-initial view becomes less likely to be the correct

[6] The motivation for this type of movement is often provided in terms of antisymmetry (Kayne, 1994). Kayne (1994) argues for a universal underlying order of Specifier-Head-Complement. In situations where this is not the case, it is argued that the complement has overtly moved to the Spec position. The reason why the CP/IP would move towards the spec-position is due to symmetry breaking. According to Kayne (1994) only elements in an asymmetric c-command relationship can be linearized.

representation of final particle sentences, as planning would be expected if the particle underlyingly precedes the proposition.

1.1.4 The Prosodic Properties of Sentence Final Particles

The Dutch sentence final particles form one intonational unit with the utterance they are attached to (Haegeman, 2014). Kirsner et al. (1996) show that the particles *hoor* and *hè* are usually associated with a high boundary tone, and argue that this is due to the semantics of *hoor* and *hè,* which are compatible with the presumed meaning of such a tone. They argue that H% indicates an 'appeal' from speaker to hearer (in contrast to L%, indicating 'no appeal'), which can be interpreted as a request for the hearer's continued attention or reply.

As discussed in Section 1.1.1, the precise interpretation of Dutch sentences with a final particle is both dependent on the context, as well as on the prosody of the sentence. This is also the case for other languages. While the relationship between overall sentence prosody and the interpretation of SFPs has not been discussed in the literature for Dutch particles, there is literature on this topic for French and Japanese (Aubergé et al., 1997; Beyssade & Marandin, 2006; Fujisaki & Hirose, 1993; Iwata & Kobayashi, 2012). Beyssade and Marandin (2006) and Fujisaki and Hirose (1993) discuss the role of intonation in transmitting the speaker's attitude and intention in French and Japanese respectively. Both studies comment on the prosodic variety available in the language of research, and show that certain attitude and intentions are associated with specific intonation contours. Fujisaki and Hirose (1993) mention that para-linguistic modification is often found to be expressed towards the end of an utterance in Japanese, whereas Beyssade and Marandin (2006) focused on the final contours of the sentence in French. In the case of Japanese, differences between different attitudes and intentions are mainly found in the f0-contours and segmental durations (Fujisaki & Hirose, 1993). A study by Iwata & Kobayashi (2012) investigated the expression of speaker's intentions through combinations of sentence final particles and intonation in conversational speech. They found that the intonation on SFPs plays a significant role in expressing intention in Japanese sentences. Although Iwata & Kobayashi (2012) note that besides the sentence-final intonation, the intonation of the whole sentence also varies depending on the speaker's intention or attitude. However, it is not clear how important the role of the intonational cues that precede the final particle and which express attitude and intention is for the identification of intention and attitude of an utterance. While it has been reported for French and Dutch that participants are able to identify the speaker's attitude/intention before the end of the sentence (Aubergé et al., 1997; van Heuven & Haan 2000; 2002), there is also evidence that the prosodic information of the final particle 'overrides' the preceding prosodic cues if the two are in conflict (van Heuven & Haan, 2002; and as mentioned[7] in Iwata & Kobayashi, 2012).

1.1.5 Summary

To conclude this section about the functions of Dutch sentence final particles and the theoretical assumptions underlying SFPs in general, there are some key points to take into consideration. First of all, I discussed the semantics and functions of two different Dutch final particles: *hè, hoor*. Based on previous literature in a variety of languages, we can assume that Dutch final particles are syntactic heads located in (the highest levels of) the CP-domain, selecting the proposition as its complement. This is supported by distributional facts, and means that the particle has a tight connection with the sentence it is attached to. The head-finality of this particle is a puzzling feature in several languages that are head-initial. The possibility has been discussed that the particle is

[7] Iwata & Kobayashi (2012) report about a Japanese study executed by Sugito (2001) that found that when the speech segments of the sentence final particles were cut off and swapped, listeners perceived the intention exposed by the segment of the final particle even though the overall f0-contour of the utterances differed from each other.

actually head-initial, and that the CP moves into the specifier position of the projection that hosts the particle. In this thesis, I take no stance in this discussion. The precise interpretation of a sentence with a particle is highly dependent on the context it appears in and on the prosodic properties of the sentence. Particularly, the final-intonation on the particle is important for the interpretation of its attitude and intention, though it has been reported that there are also cues expressing attitude and intention preceding it for languages such as French and Japanese. Based on the facts described in this section, it seems that Dutch particles could be very much related to the intention of the sentence. Regardless of whether the particle is underlying head-initial or not, it is a fact that its realization is at the end of the sentence. This could have consequences for both the production and planning of the particle, as well as for its interpretation. In the following section I review some of the relevant literature on planning in sentence production and processing to examine what the consequences of the theoretical implications discussed in the current section could be. If the sentence final intentional particle is a syntactic head that selects for the proposition as its complement, do we find effects of this selection in speech production and/or processing?

1.2 Planning and Anticipation of Intentional Final Particles in Dutch

Some of the questions addressed in this thesis are concerned with the planning of intention. How strong is the connection between the linguistic intention and a Dutch final particle like *hoor*? Does the speaker plan ahead for the usage of *hoor* to express this intention, or is its usage decided upon at the end of the utterance? Take for example a situation where somebody walks into an antique store, looking around. At some point, a beautiful chair catches his/her eye and he/she starts looking for a price tag. Suddenly a voice behind him/her says (8):

(8).	Die	is	niet	te	koop,	hoor!
	that	is	not	to	buy	SFP
	'That one is not for sale!'					

The utterance in (8) could be uttered by a shop owner. Possibly, the shop owner saw that the customer was looking for a price tag, and assumed that this person would be interested in buying the chair. However, as this particular chair is not for sale, the shop owner utters (8) in order to correct the assumption that the customer has (or to warn him not to get his hopes up). It becomes immediately clear that the utterance in (8) is not just an assertion, but that the speaker means something by what he says, and wants to convey this message to the listener. The sentence in (8) is thus formulated with the *intention* of a 'correction/warning', added to an assertion of the fact that 'That chair is not for sale.'. In (8) the linguistic intention seems to be mainly expressed with the final particle *hoor* (discussed in Section 1.1.1), but intuitively the shop owner first comes up with the intention of his message (the why), before he formulates the exact components of this utterance (the what).

Unlike the speaker, the listener is completely dependent on the sequence of the speech signal. If the intention of the utterance is mainly expressed with the final particle, does this mean the listener can only interpret this intention at the end of the utterance? Or can he already infer the intention from the sentence from other cues, e.g. the prosody of the utterance? If prosodic cues are available, they could help the listener to anticipate upon the continuation of the sentence.

I provide an overview on planning in sentence production in Section 1.2.1. Next, some literature about sentence parsing and anticipation from the listener is discussed in Section 1.2.2.

1.2.1 Planning in Sentence Production

Not much is known about how speakers plan the messages that are translated into utterances during language production (Brown-Schmidt & Tanenhaus, 2006). Brown-Schmidt & Konopka

(2015) describe that there are two influential ideas about message planning that form the basis of modern views on sentence conceptualization: A sequential view, in which speaking begins as soon as any message element is available (Paul, 1880), or a holistic approach, in which speakers first generate a holistic plan for the entire communicative intention, before the linguistic encoding begins (Wundt, 1900). However, for practical reasons, it is difficult to tear these two ideas apart, as it is challenging to distinguish message planning from sentence planning. Even in studies that demonstrate incremental production planning (e.g. Griffin, 2001; Ferreira & Swets, 2002; Smith & Wheeldon, 1999), it is possible that a complete message may have been set before speech onset first, and only linguistic encoding proceeded incrementally after that (Brown-Schmidt & Konopka, 2015). For this reason, studies that try to investigate the interaction between message planning and linguistic encoding, have mainly focused on sentence repairing, situations in which speakers prepare messages that change over time and where new message-level information must be incorporated into the sentence (Brown-Schmidt & Konopka, 2008; 2015; Brown-Schmidt & Tanenhaus, 2006). These studies show that sentences are planned in a continuous incremental fashion, allowing message updates to be fluidly integrated into linguistic encoding (Brown-Schmidt & Konopka, 2015). However, the fact that sentences can be repaired without consequences does not exclude the possibility of a prior holistically planned message that allows for smooth reparations or adjustments.

Most current models of language production propose that speaking is an incremental process, and that speakers plan the lexical content of what they want to say in small chunks, rather than in whole sentences (see Wheeldon (2013) for an overview). However, there is still a wide debate about how big the planning units at different stages of production actually are. A strong version of incrementality states that the sentence plan is developed in the same order as words of an utterance are produced and that words are often planned just before they are uttered. Other views have proposed either that the whole clause is a unit of planning (supported by word exchange errors, such as *this spring has a seat in it*; Garrett, 1980) or that the scope of planning is phrasal (e.g. Allum & Wheeldon, 2007; Smith & Wheeldon, 1999), although for the phrasal-planning view, it is unclear whether this concerns the syntactic, or the phonological phrase (Wheeldon & Lahiri, 2002).

Since verbs play a very important role in sentences, several models of speech production have adopted the view that the verb's syntactic representation guides structural processes (i.e. the lemma of the verb is selected before the relevant structural processes are performed, Bock & Levelt, 2002; Ferreira, 2000; Kempen & Hoenkamp, 1987). Specifically, Ferreira (2000) argued that the selection of the verb takes place before the phonological encoding of the first phrase of a sentence is finalized. However, Schriefers et al. (1998) argue against this view. Making use of an extended version of the picture-word interference paradigm, Schriefers et al. hypothesized that if verb selection in German occurs before the speech onset of an utterance, then the semantic interference effect caused by a semantically related distractor should occur both in verb initial (VS) and verb final (SV) utterances. However, Schriefers et al. showed the semantic interference effect was only observed in verb-initial utterances, suggesting that verb selection is not necessarily performed in advance. These results are both compatible with a word-by-word and a phrase-by-phrase (planning) based incrementality, but not by an account that argues that the verb (and its thematic relations) guide structural processes.

Recent research by Momma et al. (2016) on Japanese argues that subjects and objects of a verb might behave asymmetrically. Theoretically, it is often assumed that the object has a tighter connection with the verb than the subject (as the object is analyzed to be a complement of the verb and the subject is argued to be an external argument of the verb, e.g. Marantz, 1984; Kratzer, 1996). Furthermore, the study by Schriefers et al. (1998) only considered the planning of the verb in relation to the subject, and Momma et al. (2016) hypothesized that verb selection might only be required before object articulation, but not necessarily before subject articulation. In order to test this hypothesis, Momma et al. (2016) adapted the extended picture-word interference paradigm by Schriefers et al. (1998) for German and tested Japanese OV and SV sentences. As Japanese is a head-final language, the verb is in final position, occurring both after the subject and object. Similar to Schriefers et al. (1998), Momma et al. (2016) argued that semantic interference before speech onset (which results in slower onset latencies) would only occur if the verb was already planned before speaking. Interestingly, Momma et al. (2016) found that this semantic interference effect took place in the OV condition (before the onset of an object) but not in the SV condition (before the onset of a subject) in Japanese. Thus it seems that Japanese speakers have not yet planned the verb they will use when producing the subject of the sentence, but they have committed to this verb when they produce the object before this verb. These findings are in conflict with the strong incrementality approach and suggest some structural guidance of lexical encoding. Other research by Momma et al. (2015) also tested this effect for passive versus active sentences in English (in which the subject of the passive sentence is assumed to be an internal argument of the verb, while the subject of the transitive sentence is an external argument). They reported similar results to the previous Japanese study: only the verb of a passive sentence showed semantic interference effects at the subject's speech onset.

Last but not least, several studies have shown that the incrementality of speech planning is dependent on various factors, such as the nature and time limit of a task (Ferreira & Swets, 2002), language-specific linguistic features (e.g. different phrasal word orders; Brown-Schmidt & Konopka, 2008) and the availability of cognitive resources (Wagner et al., 2010).

1.2.2 Anticipation in Sentence Processing

There is more literature on sentence processing and anticipation than there is on language production and planning. Unlike sentence production, where one goes from conceptualization to output, sentence perception starts out with an auditory or orthographic input, which then needs to be processed until conceptualization of the input. This type of input is used to determine how humans are able to understand and produce sentences in conversation with a striking speed and accuracy, and many linguists ascribe this ability to prediction (e.g. Chang et al, 2006; Pickering & Garrod, 2013). Accurate expectations about the world would allow humans to make sense of incoming stimuli and respond in efficient ways, allowing them to anticipate what others will say before they say it and to prepare their own response, reducing the computational burden of quickly interpreting their utterance (Kutas et al, 2011; Pickering & Garrod, 2013; Rabagliati et al., 2016). Prediction (or anticipation) could take place at various levels: 1) at the semantic level (e.g., which word is expected to fill a certain slot based on meaning; Zwitserlood, 1989), 2) at the syntactic level (e.g. 'garden-path' sentences[8]; Bever, 1970) or 3) at the phonological level (such as the anticipation of a specific phonological form of an indefinite article in English; DeLong et al., 2005[9]). However, while there are various accounts that support prediction of the upcoming speech input (e.g. DeLong et al., 2005; Foucart et al., 2015; Kamide, 2008; Kamide et al, 2003; Wicha, et al., 2004; van Berkum et al., 2005), there is still some discussion about what such anticipation actually means for theories of speech processing. Roughly, the literature on prediction in speech processing can be divided into a side that considers prediction to be a local

[8] Garden-path sentences are sentences like *the horse raced past the barn fell*, in which *raced* is initially analyzed as a main verb, and not a participle (Bever, 1970).

[9] DeLong et al. (2005) showed that when participants heard a sentence like *the day was breezy so the boy went outside to fly… a kite/an airplane … in the park* the participants presented with *an airplane* in this context elicited a N400 response, already at the indefinite article. This indicates that participants did not only anticipate on a specific concept (e.g. it is more likely to fly a kite than an airplane), but they even anticipated on the phonological form of the article connected to this concept.

event and syntactic processing to occur incrementally (e.g. Altmann & Steedman, 1998; Kutas et al., 2011; Pickering, 1994; Tanenhaus et al., 1995), and a side that assumes that the sentence processor forms explicit structural or thematic expectations about the upcoming material (Aoshima et al., 2004; Ferreira & Clifton, 1986; Ford, 1983; Frazier & Fodor, 1978; Miyamoto & Takahashi, 2002; Stowe, 1986).

Most literature on sentence prediction focuses on grammatical relationships and linguistic sequences. However, it has been suggested, that it would also be beneficial for a listener to predict the speaker's intent (DeLong et al., 2014). Due to the difficulty of creating a good experimental paradigm that really taps into prediction and anticipation, it is not surprising that there is little research focusing on the prediction of the speaker's intent, as it is the case that the intention of the speaker is namely often expressed by means of prosody or pragmatic particles (Section 1.1.4). Some literature that comes close to investigating the anticipative processing of intentional information is that focusing on the prosodic properties of sentences that express communicative intent or attitude by the means of prosody (Aubergé et al., 1997; Sugito, 2001; van Heuven & Haan, 2002). As mentioned in Section 1.1.4, Aubergé et al. (1997), Sugito (2001) and van Heuven and Haan (2002) all showed that listeners are able to use prosodic cues in the sentence to predict the end of the sentence it (differentiating different intentions and/or attitudes). However, such studies use an experimental paradigm that encourages participants to engage in prediction, and therefore cannot tell us anything about the natural predictive processing of intentional and prosodic information. Hopefully, experimental paradigms tapping into that specific question will be developed in the future.

1.3 THIS STUDY

As discussed in the previous section, speech production models assume that every utterance starts with a certain intention and the formulation of a conceptual message (e.g., Bock & Levelt, 2002; Gambi & Pickering, 2016). However many important cues about the intention and attitude with which a sentence is uttered appear at the end, namely, in the form of a discourse particle or a boundary tone. This raises interesting questions about the planning involved when expressing intention. Segmental elements expressing intention and attitude like final discourse particles provide a linguistic anchor to investigate the processes involved in the implementation of a speaker's intent into an utterance. As discussed in Section 1.1, there are theoretical reasons to assume that sentence final discourse particles play an important role from the beginning of sentence formulation, selecting the entire proposition as their complements. If it indeed is the case that intentional discourse markers are hosted in a syntactic head within the CP-domain, thus selecting the rest of the utterance, it is expected that particles are planned in advance before the encoding of the whole utterance takes place. (Even though they are overtly realized at the end of the utterance.) In the current thesis, this hypothesis is put to the test, aiming to answer whether the Dutch sentence final intentional particles *hè* and *hoor* are planned in advance, i.e. is the speaker's intent directly mapped onto a linguistic item, or whether these particles are integrated at the end of a sentence.

To investigate the potential planning of intentional final discourse particles I conducted three experiments. In this thesis, I examine the prosodic and "message-encoding" properties of the final particles *hè* and *hoor* in Dutch against other Dutch elements which also appear in sentence final position, but which serve a different non-intentional function (such as the deictic marker *zo*, the addressee marker *man* and the focus particle *wel*) and are thus expected to behave differently. While the intentional final particles *hoor* and *hè* add the speaker's intention and thought to the entire utterance, this is not the case for the other final elements discussed in this thesis. The hypothesis underlying the three experiments of this thesis is that the non-intentional elements should behave differently from the intentional particles *hè* and *hoor* with regard to planning, if it is true that intentional particles select the entire utterance as their complement. This question is addressed in Experiment 3 (Chapter 4) with a priming-picture-naming task, inspired on Momma et al. (2015). Furthermore, two additional experiments were conducted in order to examine the prosodic properties of sentences with different final elements (Experiment 1), and to examine via a gating task whether any different prosodic cues can be used by the speaker before the final particle to anticipate the final element of the sentence (Experiment 2). Overall, the results of these experiments should shed some light on the role intentional Dutch final particles play in the sentence, and about the strategies concerning the perception and production of these particles in relation to planning and prediction strategies.

2. EXPERIMENT 1: THE PRODUCTION OF DUTCH FINAL PARTICLES

2.1 INTRODUCTION

This first experiment looks at the prosodic properties of sentences containing sentence final particles. Since SFPs like Dutch *hè* and *hoor* directly express the intention with which an utterance is spoken, the prosodic features of the sentence preceding the final particles might already contain features of that attitude or intention. Though such prosodic cues cannot prove that speakers have already planned ahead on the use of a final particle, it would show that the speaker-related force intended by the sentence (as discussed in Section 1.1.2, Figure 1) is already available before the final particle appears. In this experiment, the intentional SFPs *hè* and *hoor* are compared against two elements appearing in final position that do not express intentional content: the deictic particle *zo* 'like this' and the addressee marker *man* 'man'[10].

The meaning of *hè* and *hoor* has been discussed in detail in Section 1.1. Examples of sentences containing *hè* and *hoor* are displayed in (9) and (10), repeated from (5) and (6) respectively. Sentences containing the final elements *man* and *zo* accompanied by a possible context they could appear in are provided in (11) and (12) respectively.

(9). *A wants to hand out cake to a boy, but then remembers the boy doesn't like cake. He/she asks the mother:*

A. Hij houdt niet van taart **hè**?
he loves not of cake SFP
'He doesn't like cake, right?' (=(5))

(10). *Somebody tries to give A's son cake. However, A's son doesn't like cake at all. He/she says:*

A. Hij houdt niet van taart **hoor**!
he loves not of cake SFP
'He doesn't like cake!' (=(6))

(11). *Two friends are outside in a park. One of them spots a bird.*

A. Wat is dat voor-'n vogel?
what is that for-a bird
'What kind of bird is that?'

B. Dat is een Vlaamse gaai, **man**.
that is a Flemish jay man
'That's a Flemish jay, man.'

(12) *Two friends are outside in a park. They just discovered a Flemish jay sitting in a tree.*

A. Oh, hoe klinkt die vogel, dan?
oh how sounds that bird then
'Oh, how does that bird sound?'

B. De Vlaamse gaai klinkt **zo**: *maakt vogelgeluid*
the Flemish jay sounds like.this:
'The Flemish jay sounds like this: makes bird sounds"

[10] The elements *zo* and *man* can appear in, but are not restricted to, the final position of the sentence.

The final elements in (9)-(12) provide different communicative functions. They can be distinguished from each other at the following levels: the speaker/hearer relationship, intentionality, and response requirements. The final particle *man* (11), for example, is similar to the particles *hè* (9) and *hoor* (10) in the way it addresses a speaker-hearer relationship, while this is not the case for the deictic element *zo* (12), which is purely descriptive. However, while the particles *hoor* and *hè* both add a specific communicative intent from the speaker to the hearer, the addressee marker *man* merely asks for the hearer's attention, and marks the utterance as being relevant to him/her. Furthermore, the particles *hè* and *hoor* differ from each other in the expectations they communicate about a possible response: the particle *hè* clearly indicates that a response is wanted, while *hoor* does not necessarily require a response from the hearer. Table 2 summaries these different communicative properties of the elements in (9)-(12).

Table 2. *The communicative properties of the Dutch elements* hè*,* hoor*,* man *and* zo.

Type	Function	S/H Relation	Intentional	Response	Periphrasis
X-*hè*	agreement-seeking	+	+	+	You should agree with me that X is true.
X-*hoor*	correction	+	+	-	You should know that, despite what you might be thinking, X is true.
X-*man*	adressee marker	+	-	-	You should pay attention to X.
X-*zo*	deictic element	-	-	-	X happens like this

**X stands for a declarative transitive sentence, e.g.* de ridder wast een draak *'the knight is washing a dragon'.*

The aim of this experiment is to identify the prosodic properties of sentences that contain *hè*, *hoor*, *man*, and *zo* respectively. Do speakers use prosodic cues to mark the upcoming of a sentence final intentional particle like *hè* and *hoor*, and are there cues to distinguish these two intentional SFPs from each other? Furthermore, is the prosodic marking of sentences containing an intentional SFP different from those sentences containing other final elements, such as *man* and *zo*? In order to address these questions, sample sets of sentences containing these four different types of final elements are recorded and their prosody is analyzed.

2.2 METHODS

2.2.1 Participants

Twelve native speakers of Dutch (3 male, 9 female), participated in this experiment on voluntary basis. Their age range was 19-26 years old ($\bar{x}$= 22.25; SD = 1.82). All of them were students affiliated with Leiden University.

2.2.2 Stimuli

Sixteen base sentences were constructed which were similar in prosodic and syntactic structure. The sentences were controlled for variance, in order to avoid any possible confounds of sentence length, syllable structure or syntactic structure. All sentences were transitive sentences, consisting of an animate definite human subject and an action/perception verb that takes an indefinite direct object. All constituents had the same syllable structure (CV 'CV.CVC CVCC VC CCVC), and consisted of as many sonorants as possible (to make the detection of f0-values easier). An example of a base sentence is given in (13).

(13).	De	ridder	wast	een	draak.
	the	knight	washes	a	dragon
	'The knight is washing a dragon.'				

Each base sentence was combined with *hè*, *hoor*, *man* and *zo*. Thus obtaining a total of 64 target stimuli. Additionally, 48 pseudo-fillers were added. Pseudo-fillers were constructed by turning the base sentences into declaratives, question-declaratives and exclamatives. An example set is provided in Table 3. A list of all base sentences can be found in Appendix B.

Table 3. *Example set of sentences in the seven different conditions.*

Condition		**Example Sentence**
With SFP:	*hè* 'agreement-seeking'	De ridder wast een draak, **hè**?
	hoor 'correction'	De ridder wast een draak, **hoor**!
Addressee marker:	*man* 'man'	De ridder wast een draak, **man**!
Deictic marker:	*zo* 'in this way'	De ridder wast een draak **zo**.
Without particle:	Declarative	De ridder wast een draak.
(pseudo-fillers)	Question-declarative	De ridder wast een draak?
	Exclamative	De ridder wast een draak!

2.2.3 Procedure

Recordings were made in a sound proof booth in the Phonetics lab at Leiden University. Participants were instructed to utter the sentences that were displayed on a computer screen, using ExperimentMFC Praat (Boersma & Weenink, 2016). The stimuli were presented into seven separate blocks of 16 sentences each (i.e., 4 blocks with the target conditions, 3 blocks with filler sentences). Within blocks, sentences were presented in a pseudo-random order. Sentences of the same type were presented within the same block, because this way it was thought to be easier for the speaker to stay within one mode (e.g. question mode) of speaking. The order in which the different blocks were displayed was randomized across participants. The experimenter controlled when the next target sentence would appear. Participants were asked to repeat hesitations or 'wrong' intonation contours (e.g. when people said the *zo* of 'time', and not the *zo* of 'manner') a second time. Utterances were directly recorded on the PC (44100 Hz, 16 bits).

2.2.4 Analysis

Acoustic analysis

Two Dutch native speakers (including myself) inspected all 640 utterances for naturalness. After inspection we excluded 72 sentences (11,2%) as they contained hesitations, mistakes or sounded unnatural. The remaining utterances were manually segmented setting the syllable boundaries in Praat (Boersma & Weenink, 2016) in which the sentences were segmented per word and syllable. The onset of the utterance was set at the burst of release of the plosive [d] (ignoring the medial phase (silence) of the plosive). This is illustrated in Figure 3. Raw f0 curves were stylized using straight lines: at every juncture a pitch point was set. This resulted in 11 pitch points per utterance, as shown in Figure 4. Pitch point 1-4 correspond to the subject (the first pitch accent), pitch point 5 and 6 correspond to the verb, pitch point 7-9 correspond to the object (second pitch accent) and pitch point 10 and 11 correspond to the final element of the sentence. Duration and f0 data were extracted using a Praat script.

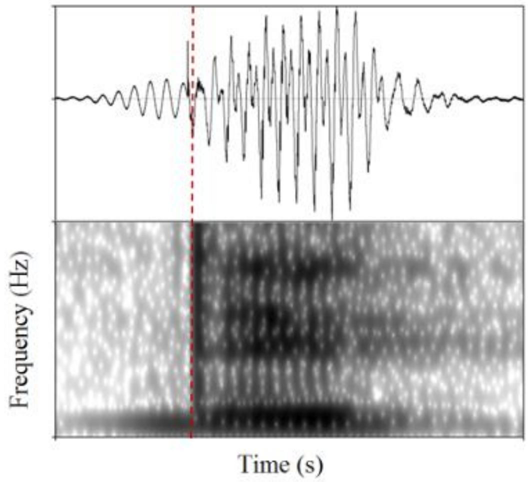

Figure 3. *Red dashed line indicates starting point measuring onset of the utterance. In oscillogram (top) and spectrogram (bottom) of* de *'the'.*

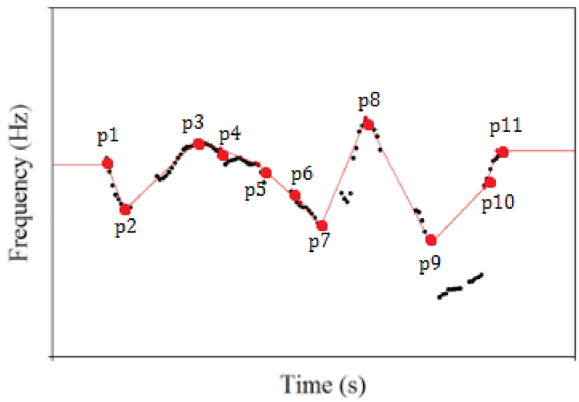

Figure 4. *Raw f0-contour (black) with stylized f0-contour superimposed (red). The red dots indicate the pitch points (p1-p11).*

Statistical analysis

The obtained data were further analyzed with the statistics software R (R Core Team, 2013) and *lme4* (Bates et al., 2015) to perform a linear mixed effects analysis of the relationship between pitch and duration with sentence type (*hè*, *hoor*, *man* or *zo*). This analysis was chosen because the linear mixed effects analysis takes into account random effects that are due to individual differences, and spoken sentences from a varying pool of individuals inherently deals with a lot of individual variation. Sentence type was included into the model as a fixed effect, and the random effects had intercepts for subjects (10 participants) and items (the 16 different sentences). The *t*-values of the fixed effects were reported to be significant at a .05 level if they exceeded 2.0. The data was plotted in figures with help of SPSS (IBM Corp, 2012).

2.3 Results

In order to gain a general overview of the data, both duration and f0 data was plotted and then analyzed statistically. The two measures are discussed independently below.

2.3.1 Duration

The average duration of the different sentence constituents (subject, verb, object and final element, denoted here as 'particle') for all participants and sentences is plotted per condition in Figure 5. The corresponding mean duration values and standard deviations for each constituent are displayed in Table 4. As shown in Figure 5, we can observe the following results: 1) the duration of *zo* differs from the other three constituents for all the constituents, 2) the verbs in the *man*-condition are shorter, compared to the verbs in the other three conditions, and finally, 3) *hè*-sentences have shorter object and particle durations compared to *man*- and *hoor*-sentences.

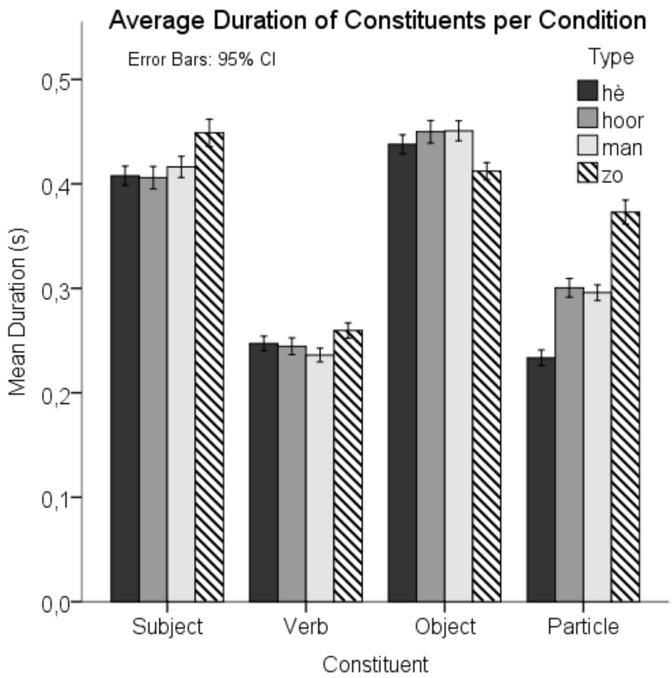

Figure 5. *The average duration (s) of the four different constituents (subject, verb, object and particle) of the 16 different sentences for all participants (N=10) , grouped per condition (hè, hoor, man and zo).*

Table 4. *Mean (SD) durations for each constituent in milliseconds, per sentence type ending with* hè*,* hoor*,* man *or* zo*.*

	Constituent			
Condition	Subject	Verb	Object	Particle
hè	408 (56)	247 (43)	439 (57)	234 (46)
hoor	406 (63)	245 (48)	450 (64)	301 (53)
man	416 (64)	236 (42)	451 (61)	296 (48)
zo	449 (74)	260 (41)	412 (47)	373 (65)

The linear mixed effect analysis reported estimates (in s) and standard errors of each constituent (subject, verb, object and particle) from the duration data separated by condition (*hè, hoor, man, zo*). This data is summarized in Table 5 below. In rows are the referents against the conditions (in columns), e.g. for the subject, the *hoor* duration-estimate is 2 ms longer than the *hè*-condition. A minus sign means that the difference is negative (shorter than the referent). The corresponding *t*-values can be found in Appendix C.1. As can be observed from the data displayed in Table 5, all constituents of the *zo*-condition differ significantly from the other three conditions in duration. Compared to the other conditions, the object constituent is significantly shorter in *zo*-sentences, while the other three constituents are significantly longer. As for the other three conditions, there is no significant difference in subject duration, though there are some marginal differences later in the sentence. For example, the *man*-condition has significantly shorter verbs, compared to the other three conditions, and the *hè*-sentences have shorter object and particle durations compared to *man*- and *hoor*-sentences.

Table 5. *Reported estimates (in s) and standard errors of the duration data of the constituent (subject, verb, object and particle) per condition (*hè, hoor, man, zo*)*.

Subject	**hè**	**hoor**	**man**	**zo**
hè	X	.002 (.014)	.009 (.005)	.038 (.006)*
hoor		X	.007 (.005)	.036 (.006)*
man			X	.029 (.005)*
Verb	**hè**	**hoor**	**man**	**zo**
hè	X	-.003 (.004)	-.011 (.003)*	.009 (.004)*
hoor		X	-.009 (.004)*	.011 (.004)*
man			X	.020 (.004)*
Object	**hè**	**hoor**	**man**	**zo**
hè	X	.012 (.005)*	.010 (.004)*	-.034 (.005)*
hoor		X	-.002 (.004)	-.046 (.005)*
man			X	-.043 (.005)*
Particle	**hè**	**hoor**	**man**	**zo**
hè	X	.068 (.005)*	.060 (.005)*	.134 (.005)*
hoor		X	-.008 (.005)	-.066 (.005)*
man			X	.074 (.005)*
standard error = in parenthesis				

* = significant at .05

2.3.2 F0

Figure 6 illustrates the overall stylized pitch contour per condition for all participants and sentences, based on the averaged f0-values of the 11 measure points. The corresponding mean f0-values and standard deviations for each pitch point are displayed in Table 6. As can be observed from raw observation of the data, the pitch contour of the *zo*-condition differs the most from the other three conditions, especially at the end. The *zo*-conditions ends with a fall, while the other three conditions end in a rise. This rise is the most prominent for the *hè*-condition, followed by the *hoor*- and *man*-condition. Overall the other three conditions are closer together, especially in the part preceding the final element, though some differences can be observed. However, there is quite some variation in the prominence of certain distinctions between conditions across participants (see Appendix D for an overview of averaged pitch contours per participant).

Table 6. *Mean (SD) f0-*values *for each measure point (11) in Hertz, per sentence type (ending with hè, hoor, man or zo).*

		Condition			
	Pitch Point	*hè*	*hoor*	*man*	*zo*
Subject	**1**	190 (45)	191 (52)	184 (47)	173 (45)
	2	174 (41)	175 (45)	168 (41)	162 (42)
	3	263 (65)	265 (74)	262 (70)	247 (72)
	4	208 (55)	209 (61)	200 (55)	184 (51)
Verb	**5**	211 (54)	217 (64)	207 (54)	187 (50)
	6	201 (49)	205 (58)	201 (51)	191 (49)
Object	**7**	171 (42)	172 (48)	163 (43)	169 (44)
	8	248 (64)	250 (67)	241 (67)	178 (44)
	9	164 (45)	157 (42)	151 (40)	196 (66)
Particle	**10**	165 (43)	159 (40)	148 (36)	194 (71)
	11	272 (71)	252 (59)	219 (58)	147 (45)

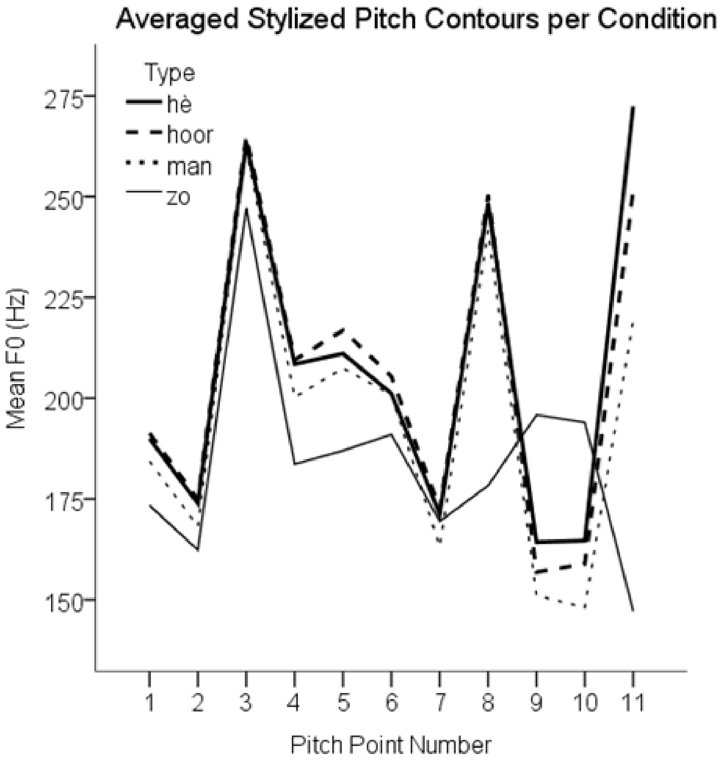

Figure 6. *The averaged stylized pitch contours of the 16 different sentences, grouped per condition (hè, hoor, man and zo), for all participants (N=10). The pitch contours are stylized and f0 and time data are displayed for 11 pitch points, located at the major transition points of the f0-contour. Roughly, pitch point 1-4 correspond to the subject (the first pitch accent), pitch point 5 and 6 corresponds to the verb, pitch point 7-9 correspond to the object (second pitch accent) and pitch point 10 and 11 correspond to the final element of the sentence.*

The estimates and standard errors obtained from the mixed linear regression are displayed in Table 7. The estimate provides information about the modeled difference (in Hertz) between two conditions at one measuring point. In rows are the referents against the conditions (in columns), e.g. for point 1 (p1) the *hoor* f0-estimate is 1.70 Hz higher than the *hè*-condition. A minus sign means that the difference is negative (lower than the referent). Significance is indicated with an asterisk at the .05 level, based on the height of the *t*-values (these can be found in Appendix C.2).

As can be observed from Table 7, the differences between the *zo*-condition and the other three conditions is significant at almost all measure points, with the exception of pitch point 2 (for the difference with the *man*-condition), pitch point 6 (for the difference with the *hè*-condition) and pitch point 7 (for the *hoor*-condition). The *zo*-condition also differs from the other three condition in the preceding part of the sentence, with a lower overall frequency and smaller peaks and rises on the subject, verb and object. This effect is particularly striking on the object (pitch point 7-9), as there is no pitch-peak related to the object. The other three conditions are more closely related, though some differences can still be observed (more subtle than the differences with the *zo*-condition). The man-condition seems to have a lower overall frequency, though this difference is the smallest on the subject nucleus (pitch point 3), as *man* does not differ significantly from *hè* or *hoor* here. At all other points *man* is significantly lower than *hoor*. The *man* and *hè* pattern alike on the pitch nuclei (pitch point 3 and 8) and on the verb (pitch point 5 and 6). The conditions *hè* and *hoor* are the most similar in their overall pitch contour preceding the final element. The only clear differences in pitch that can be observed between these two conditions (*hè* and *hoor*) are on the verb (pitch point 5 and 6) and in the fall before the particle (pitch point 9).

Table 7. *Reported estimates (in Hz) and standard errors (in parenthesis) mixed linear regression of the f0-values per pitch measure point (p1-11) and condition (*hè, hoor, man, zo*).*

p1	**hè**	**hoor**	**man**	**zo**
hè	X	1.70 (1.50)	-3.16 (1.53)*	-7.44 (1.58)*
hoor		X	-4.89 (1.51)*	-9.13 (1.62)*
man			X	-4.28 (1.55)*
p2	**hè**	**hoor**	**man**	**zo**
hè	X	0.95 (1.25)	-4.14 (1.21)*	-5.18 (1.29)*
hoor		X	-5.09 (1.24)*	-6.13 (1.32)*
man			X	-1.03 (1.27)
p3	**hè**	**hoor**	**man**	**zo**
hè	X	3.70 (2.90)	0.37 (2.78)	-7.39 (2.98)*
hoor		X	-3.33 (2.85)	-11.10 (3.05)*
man			X	7.76 (2.93)*
p4	**hè**	**hoor**	**man**	**zo**
hè	X	2.96 (2.30)	-6.17 (2.22)*	-16.01 (2.37)*
hoor		X	-9.13 (2.27)*	-18.93 (2.43)*
man			X	-9.85 (2.33)*
p5	**hè**	**hoor**	**man**	**zo**
hè	X	8.47 (2.35)*	-1.71 (2.27)	-15.57 (2.43)*
hoor		X	-10.18 (2.32)*	-24.04 (2.48)*
man			X	-13.86 (2.39)*
p6	**hè**	**hoor**	**man**	**zo**
hè	X	6.61 (2.00)*	1.41 (1.93)	-2.70 (2.07)
hoor		X	-5.21 (1.98)*	-9.31 (2.12)*
man			X	-4.11 (2.04)*
p7	**hè**	**hoor**	**man**	**zo**
hè	X	2.46 (1.95)	-6.14 (1.88)*	5.06 (2.01)*
hoor		X	-8.60 (1.93)*	2.60 (2.06)
man			X	11.20 (1.98)*
p8	**hè**	**hoor**	**man**	**zo**
hè	X	4.67 (3.58)	-5.05 (3.45)	-63.72 (3.70)*
hoor		X	-9.72 (3.54)*	-68.38 (3.78)*
man			X	-58.67 (3.63)*
p9	**hè**	**hoor**	**man**	**zo**
hè	X	-6.90 (3.27)*	-11.67 (3.15)*	37.17 (3.36)*
hoor		X	-4.77 (3.23)	44.07 (3.44)*
man			X	48.84 (3.31)*
p10	**hè**	**hoor**	**man**	**zo**
hè	X	-5.55 (3.30)	-15.12 (3.18)*	34.10 (3.40)*
hoor		X	-9.57 (3.26)*	39.65 (3.48)*
man			X	49.22 (3.34)*
p11	**hè**	**hoor**	**man**	**zo**
hè	X	-18.11 (4.00)*	-51.76 (3.86)*	-121.62 (4.12)*
hoor		X	-33.65 (3.95)*	-103.5 (4.22)*
man			X	-69.86 (4.06)*
standard error = in parenthesis				

* = significant at .05

2.4 DISCUSSION

In this production experiment I investigated whether there are early prosodic cues (i.e., before the final element) that differentiate utterances with various final elements from each. Results indicate that there are such cues, and that even though they are sometimes quite small, they are used quite consistently across participants. Interestingly, the amount of cues in which the conditions differ from each other seem to be related to the semantic/pragmatic properties that were set out in Table 2 of Section 2.1 for each of the four final elements. The *zo*-condition is the most deviant from all other conditions, in almost every measuring domain (i.e., in the duration of all constituents and the f0-values of all pitch measure points). The *man*-condition differs at several points from the *hè* and *hoor*-condition, with a lower overall frequency (except in the peak-values) and a shorter verb duration. Sometimes *man*-sentences pattern with the *hè*-sentences (e.g. with the f0-values on the verb), whereas in other cases, they patters more with *hoor* (e.g. in object and verb duration). The measurements of the *man*-condition are often in-between the *zo*-condition and the *hè*/*hoor*-condition, in both pitch and duration. The *hè* and *hoor*-conditions differ from each other only a little in the part preceding the final particle. The clearest differences can be found on the verb, as f0-values are higher for *hoor* than for *hè*, and the object duration of *hè* is shorter than the object duration of *hoor*-sentences.

It can thus be concluded that there are prosodic differences in the sentence material preceding the four different final elements *hè*, *hoor, man* and *zo*. However, many of the consistent differences between the sentences that address the speaker-hearer relationship (*hè*, *hoor* and *man*) are not that evident. The question is whether listeners are able to detect these subtle differences, or whether they are too small to be noticeable. I addressed this question in the following perception experiment.

3. EXPERIMENT 2: THE PERCEPTION OF DUTCH FINAL PARTICLES

3.1 INTRODUCTION

In the previous chapter, the prosodic properties of sentences containing four different final particles were discussed. The results of this experiment showed that sentences ending with the deictic element *zo* differed significantly from the other three conditions (sentences ending with the intentional SFPs *hè* and *hoor* and sentences ending with *man*) at almost every measuring point in both duration and pitch. The three final elements that address the speaker-hearer relationship (*hè*, *hoor* and *man*) pattern more alike in both pitch and duration, though some subtle significant differences can be reported. It is the case that sentences ending with *man* tend to be lower in pitch than sentences ending with *hè* and *hoor*, and they also have a shorter verb duration. The intentional SFPs *hè* and *hoor* differ the least from each other in the sentence parts preceding the final element. The object of *hè*-sentences has shorter durations, and lower f0-values in the falls. Verbs also have lower f0-values in *hè* than in *hoor*. As we observed that there are several consistent differences between the sentences preceding different particles, the current experiment focuses on whether Dutch listeners are able to detect these cues, and predict what final particle will follow.

As discussed in Chapter 1, different SFPs express different expectations from the speaker about the continuation of the conversation. For example, Kirsner and van Heuven (1996) noted that the particle *hè* explicitly requests acknowledgment from the hearer, while the particle *hoor* seems to indicate that nothing of the kind is needed or wanted. However, if the speaker's intention only becomes explicit at the end of the sentence with the particle, the listener has not much time to anticipate its response. It is well-known that prosody in Dutch can also signal the speaker's intention and attitude, and that prosodic cues preceding the particle could possibly help the listener anticipate for the intention or attitude expressed by an utterance. In this perception experiment we put this hypothesis to the test by exposing participants to a gating task.

The gating-technique (Grosjean, 1980) has been used in previous research to trace the ability of listeners to anticipate upcoming events. A spoken language stimulus is presented into segments of increasing duration, starting at the beginning of the stimulus. The acoustic stimuli used for this type of experiment can vary, and the gating paradigm has been used for units as small as sounds or as big as sentences (Grosjean, 1996). In this experiment, the gating paradigm is used to investigate whether listeners can anticipate upon the speaker's intention or attitude (expressed by the final particle) before the occurrence of this particle. The gating technique has been used before to research the anticipation of speaker attitude. Van Heuven & Haan (2000; 2002) for example, used a gating task to show that Dutch listeners could differentiate between a statement and a declarative question before listening to the boundary tone. Also Aubergé et al. (1997) used this paradigm to test whether French speakers could perceive attitudes before the end of sentences. In both cases, speakers were able to decide upon the attitude of the speaker before the end of the sentence, when they were forced to make a decision out of a set of different attitudes presented to them after listening to a gated signal.

In the following experiment, I examined whether Dutch listeners can anticipate the upcoming of an SFP (i.e., *hè* or *hoor*), an addressee marker (*man*), or a deictic element (*zo*). In other words: I investigated whether the prosodic differences before the particle observed in Experiment 1 are used by listeners to anticipate the continuation of a sentence containing one of the four final particles.

3.2 METHODS

3.2.1 Participants

Twenty-four Dutch native speakers (12 male, 12 female), participated in this experiment on voluntary basis. Their age range was 18-29 years old ($\bar{x}$ = 22.0; SD = 2.5). All of them were students affiliated with Leiden University. None of the participants had participated in Experiment 1 and none of them reported any hearing disorders. Participants were randomly assigned to one of the four pseudo-randomized versions of the experiment.

3.2.2 Stimuli

For constructing the audio stimuli, I inspected once more the recordings obtained in the production experiment and selected a speaker who was the most clear in articulation and who lacked any distinctive dialectal accent. This speaker, who was a 21 year-old female from the area of Harlem (North-Holland), was invited to come back to the Phonetics Lab for an additional recording session. Therefore, 16 new base sentences (Appendix B) combined with *hè*, *hoor*, *man* and *zo* were recorded. All stimuli were standardized to a mean intensity of 60 dB. An example of a set is given in (14).

(14). a. De ridder wast een draak, **hè**? HÈ-CONDITION
the knight washes a dragon SFP
'The knight is washing a dragon, right?.'

a. De ridder wast een draak, **hoor**! HOOR-CONDITION
the knight washes a dragon SFP
'The knight is washing a dragon!' (you are wrong thinking differently)

a. De ridder wast een draak, **man**! MAN-CONDITION
the knight washes a dragon SFP
'The knight is washing a dragon, man!.'

a. De ridder wast een draak **zo**. ZO-CONDITION
the knight washes a dragon SFP
'The knight is washing a dragon like this?.'

All audio stimuli were segmented into three gates: I) the first gate containing the subject of the sentence (e.g. *de ridder* 'the knight') with an average duration of about 400 ms, II) the second gate containing both the subject and verb of the sentence (e.g. *de ridder wast* 'the knight is washing') with an average duration of about 630 ms, and finally, III) the third gate containing the subject + verb + object of the sentence (e.g. *de ridder wast een draak* 'the knight is washing a dragon') with an average duration of about 1400 ms. The third stimuli type was cut of right before the onset of the final sentence element, and it was made sure that there were no anticipating cues for the following segment (initial sound of the final element). The precise duration of the stimuli depended on both sentence type and condition. The average duration and f0-values for each condition for the 16 sentences are displayed in Figure 7 and 8 respectively. When the stimuli duration and pitch-values of the current recordings are compared the average duration and pitch-values obtained from Experiment 1 (see Figure 5-6, Section 2.3) a similar overall pattern is observed. The difference between *zo* and the other three sentence conditions is the most obvious, in both pitch and duration measurements. Moreover, the differences between *hè* and *hoor* that we observed in Experiment 1 are even more prominently present in the recordings of the current stimuli (with longer subject and object durations for *hè* and with a lower pitch for *hè* than *hoor*). In the stimuli recordings *man*-sentences pattern a bit more like *hoor* than they did in the recordings of Experiment 1. This is because the *man*-recordings from the stimuli are pronounced with more exclamation than the average *man*-recordings obtained from Experiment 1.

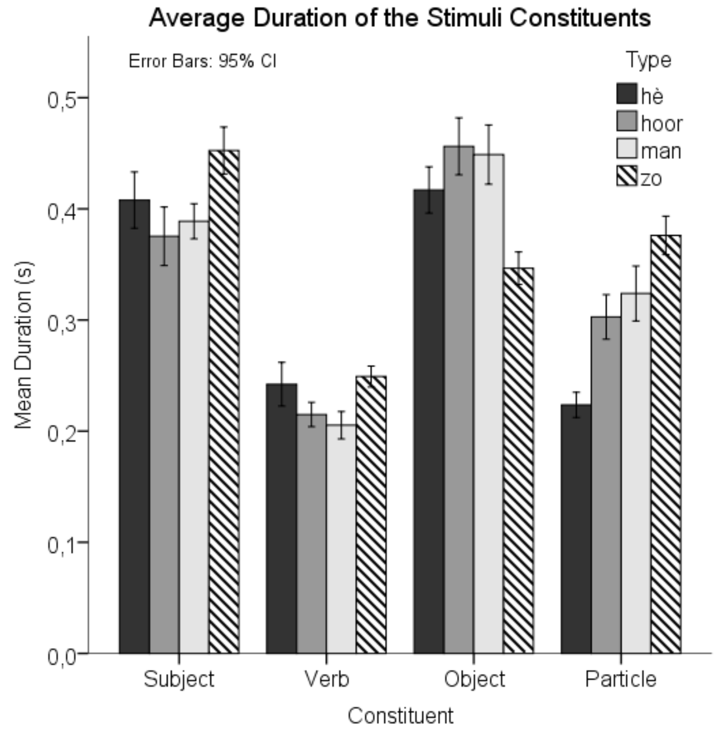

Figure 7. *The average duration (s) of the four different constituents (subject, verb, object and particle) of the 64 different stimuli sentences, grouped per condition (16 sentences per condition).*

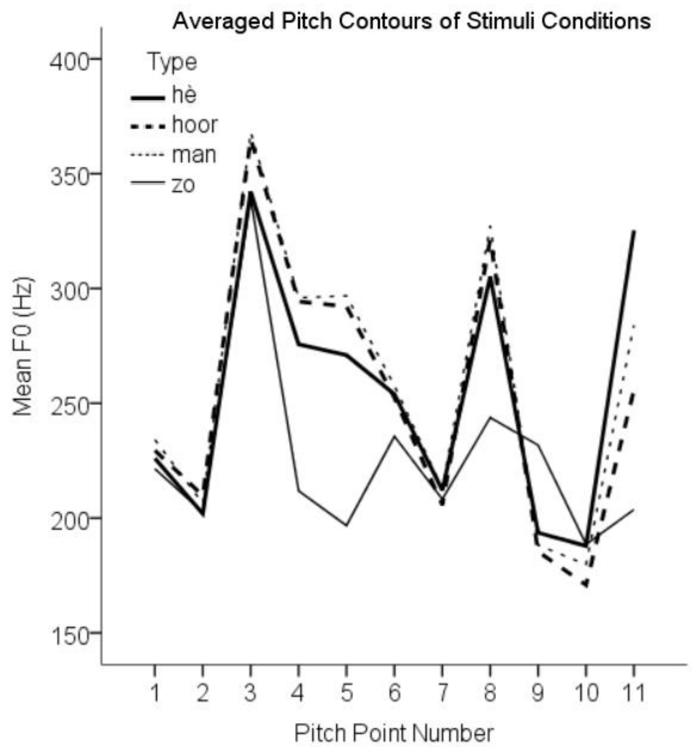

Figure 8. *The averaged stylized pitch contours of the complete 64 different stimuli sentences, grouped per condition (16 sentences per condition). The pitch contours are stylized and f0 is displayed for 11 pitch points, located at the major transition points of the f0-contour. Roughly, pitch point 1-4 correspond to the subject (the first pitch accent), pitch point 5 corresponds to the verb, pitch point 6-8 correspond to the object (second pitch accent) and pitch point 9-11 correspond to the final element of the sentence.*

3.2.3 Procedure

Before the experiment started, participants were presented with a document containing four different stories serving as the context for either *hè*, *hoor*, *man* or *zo*. This was to ensure that each participant would interpret the final element in a uniform way (e.g. that they interpreted *zo* as a deictic element reporting 'manner' and not as the temporal adverb 'later', which is a possible interpretation for *zo* in Dutch). The different contexts and their translations are provided in the Appendix E.

The gating experiment was conducted using ExperimentMFC, Praat (Boersma & Weenink, 2016). Participants were seated comfortably in front of a computer, wearing headphones, in a quiet area, and performed the experiment by clicking with the computer mouse on the appropriate answer after listening to each of the gates. The experiment consisted of four blocks: one practice block (consisting of six trials) to gain familiarity with the task at hand, and three experimental blocks (consisting of 64 trials, each 16 different sentences x 4 conditions). The first block contained all sentences and conditions from stimuli type I (only subject), the second block contained all stimuli from type II (subject + verb) and the third block consisted of all stimuli from type III (subject + verb + object; the entire sentence up to the onset of the final element). All participants were presented with the same ordering of experimental blocks, in order to rule out any facilitating effects of first being presented with a longer sequence (e.g. the stimuli used in Block 3) before being encountered with a shorter sequence (e.g. the stimuli used in Block 1). This is in order to avoid the possibility that participants use prosodic information about the longer sequence to recognize a shorter sequence. Within blocks, the trials were pseudo-randomized, meaning that there were four different versions of the experiment with different randomization lists of the trials.

With each trial, participants were first presented with an audio file playing the sound stimulus. After listening to this stimulus, a screen appeared in which the upper part showed the written continuation of the sentence in text followed by the four possible sentence endings *hè*, *hoor*, *man* and *zo*. Participants were asked to listen to the sound stimulus, read the continuation, and click on the sentence element they thought the sentence would end with. There was no time limit for making their decision, though participants were urged to choose intuitively, and not to overthink their answer. After making this decision, participants were asked to indicate how confident they were about their response on a five point scale (1 = very uncertain, 5 = very certain). The precise length of the sound stimulus and continuation varied per block (see Table 8), and the position of the four possible answers (i.e., *hè*, *hoor, man* and zo) varied per trial. The different positions were counterbalanced such that every position (upper, lower-upper, lowest and upper-lowest) contained the same amount of instances of each condition, and the same amount of right matches. The next trial started as soon as the participant clicked on the *OK* button. In between blocks, there was room for a short break. Completing the experiment took approximately 20-30 minutes.

Table 8. *Trial-sequence per block, with the example sentence* 'De ridder wast een draak.'
'The knight is washing a dragon'

	Sound	Continuation	Decision
Block 1	🔊 De ridder	.. wast een draak	- hè - hoor - man - zo
Block 2	🔊 De ridder wast	… een draak	
Block 3	🔊 De ridder wast een draak	…..	

3.2.4 Analysis

In total, 4608 responses (3 blocks * 64 trials *24 participants) were obtained from this experiment. Accuracy data (cases in which the participant's continuation choice matched with the speaker's continuation), the actual responses, and certainty indications were collected for further analysis. The data of the actual responses and correct responses per condition was plotted in contingency tables (one table per block), and a Pearson chi-square test was conducted to test whether all responses were equally distributed. For those cases in which the chi-square value reported a significant effect, a post hoc analysis was conducted to investigate which responses were responsible for this effect. The post hoc analysis was based on the Standardized Residual Method (Beasley & Schumacker, 1995), in which residuals (adjusted *z*-scores) are calculated for each actual count (versus the expected count), and then translated to precise estimates of *p*-values. Since the residual is calculated for each cell in the contingency table, a Bonferroni correction was applied to the *p*-values depending on how many cells there were within a table. Also, I looked at the confusion matrix, i.e. what listeners confused with what when they made an inaccurate choice, and ran a multinomial logistic regression to investigate the influence of trial condition (*hè*-target, *hoor*-target, *man*-target or *zo*-target) on the nominal dependent variable 'actual response' (*hè*, *hoor*, *man* or *zo*) for all three blocks. This way it could be investigated whether the response proportions differed significantly across conditions. The relationship between accuracy and certainty was investigated using a Spearman's Rho correlation test.

3.3 RESULTS

3.3.1 Response Distribution

In order to gain an overview of the most important results of this experiment, Figure 9-11 show the percentage of actual responses on stimuli of the four different conditions, separated per block.

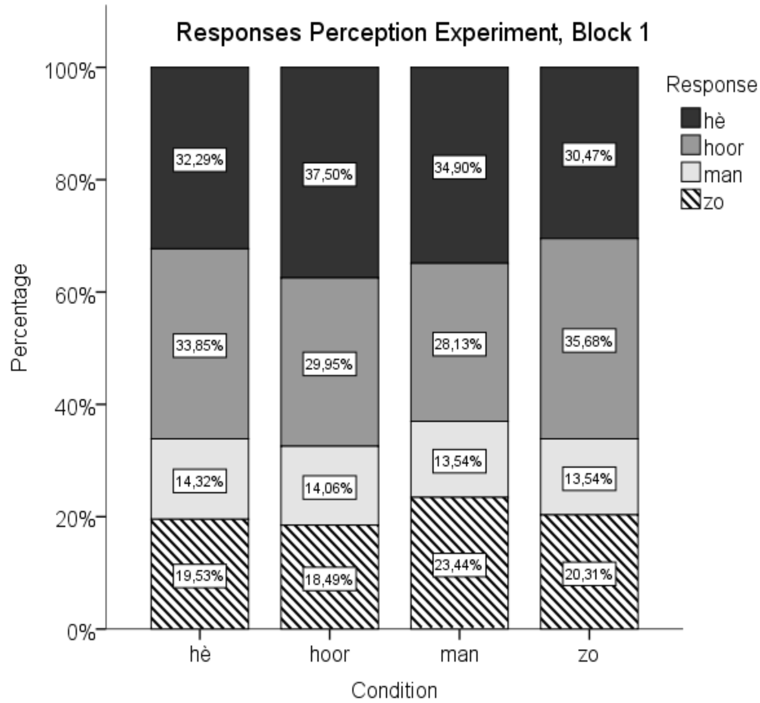

Figure 9. *Responses (in percentages) of Block 1, distributed per condition (*hè, hoor, man *and* zo*). The audio stimuli only consisted of the subject of the sentence.*

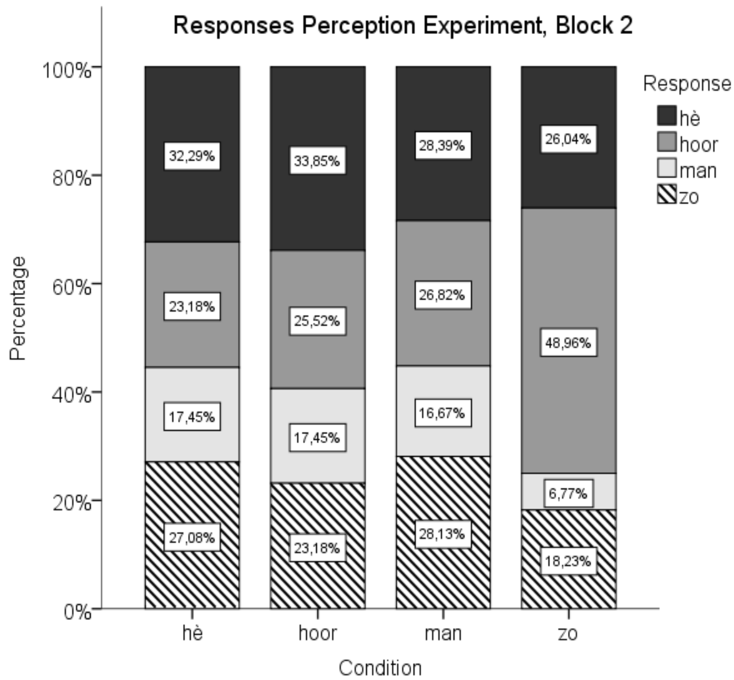

Figure 10. *Responses (in percentages) of Block 2, distributed per condition (hè, hoor, man and zo). The audio stimuli consisted of the subject + verb of the sentence.*

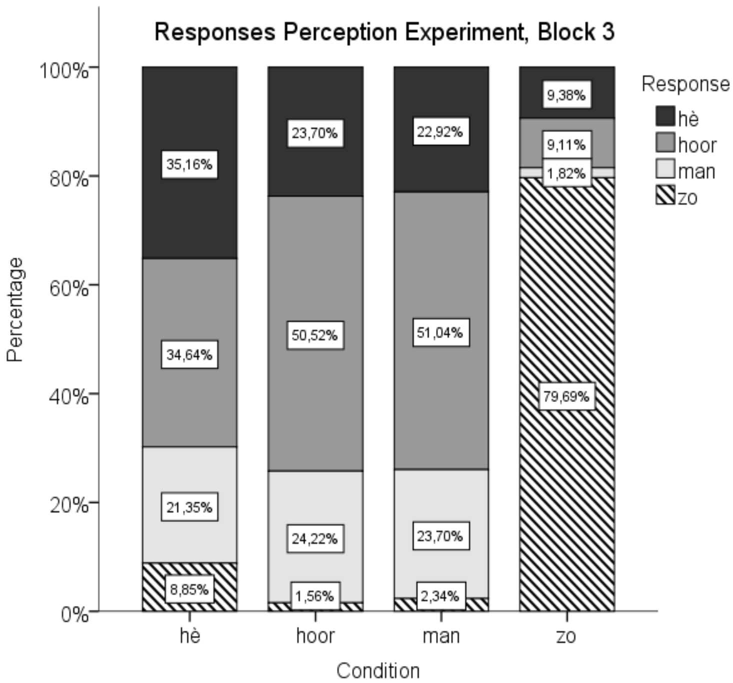

Figure 11*: Responses (in percentages) of Block 3, distributed per condition (*hè, hoor, man *and* zo*). The audio stimuli consisted of the subject + verb + object of the sentence.*

The figures above indicate that participants could not anticipate on the correct sentence finalization for all conditions. In Block 1, the distribution of responses (*hè*, *hoor*, *man* and *zo*) did not differ significantly per condition, nor was one response preferred significantly above another, X^2 (9, N = 1536) = 10.25, $p > .05$. However, a multinomial logistic regression shows that the proportion of *hoor* responses in comparison to *hè* responses is significantly less in the conditions '*hoor*' and '*man*', compared to the proportion of these response in the *zo*-condition (respectively:

β1 = .682 (95% CI .482 – 0.966), Wald = 4.651, p < 0.031; and β1 = 0.688 (95% CI 0.483 – 0.980), Wald = 4.284, p < 0.038). In other words, the response *hoor* is chosen significantly more often in the *zo*-condition ($p < .05$), than in the *hoor*- or *man*-conditions.

This effect is also observed for Block 2. The preferred responses for the different conditions obtained in Block 2 were significantly not equally distributed, X^2 (9, N = 1536) = 89.21, $p <.05$. Further post hoc analysis was conducted, making use of the standardized residuals. As the standardized residual was calculated for each cell separately, a Bonferroni correction was applied to the *p*-value, resulting in $p = .05/48$ (.0001). For Block 2, the post hoc analysis showed that the significant effect of the chi-square test could be attributed to the proportion of *hoor* responses in the *zo*-condition. This proportion was significantly higher than the expected value, with a residual of 8.7, corresponding to $p < .0001$. Further analysis of the data with a multinomial logistic regression shows that the proportion of *hoor* responses is significantly larger than the proportion of *hè* responses in the *zo*-condition compared to the other three conditions: *hè*-condition, β1 = .382 (95% CI .265 – .550), Wald = 26.783, p < 0.000; *hoor*-condition, β1 = .401 (95% CI .281 – .573), Wald = 25.142, p < 0.000; and *man*-condition, β1 = .503 (95% CI .350 – .722), Wald = 13.835, p < 0.000. This increase in *hoor*-responses in the *zo*-condition goes hand in hand with a significant decrease in *man*-responses (compared to *hè*-responses) for the *zo*-condition compared to the *hè*- and *man*-condition, respectively, β1 = 2.078 (95% CI 1.231 – 3.509), Wald = 7.489, p < .006 and β1 = 2.258 (95% CI 1.175 – 3.343), Wald = 9.058, p < .003. For the *hoor*-condition, this effect is only marginally significant, β1 = 1.982 (95% CI 1.329 – 3.839), Wald = 6.587, p < .010 (p > .05). Thus, the shared results of Block 1 and Block 2 thus indicate that the response distribution of *hè*, *hoor* and *man* are equally distributed. However, participants were significantly more likely to respond *hoor* for stimuli belonging to the condition '*zo*' compared to the other three conditions.

As for Block 3, there were multiple significant deviations from equally distributed proportion values, X^2 (9, N = 1536) = 967.7, $p < .05$. The most obvious effect (which can also be observed in Figure 11) is the highly significant preference for the response *zo* in the condition *zo*, scoring well below the *p*-value of .0001, with a residual value of 30.4. This is also reflected in the results from the multinomial logistic regression. There are significantly more *zo*-responses than *hè*-responses in the condition '*zo*' compared to the other three conditions: *hè*-condition, β1 = .030 (95% CI .018 – .049), Wald = 182.469, p < .000; *hoor*-condition, β1 = .008 (95% CI .003 – .019), Wald = 113.135, p < .000; and *man*-condition, β1 = .012 (95% CI .006 – .026), Wald = 127.266, p < .000. Correspondingly, the other possible responses for the *zo*-condition, *hè, hoor* and *zo*, scored significantly below the expected values, with residuals of -7.2, -12.8 and -9.4 respectively, $p < .0001$. Based on these results, it is clear that participants really recognized the *zo* particle to belong to the base sentence of the *zo*-condition, as the responses for *zo* to the other 3 conditions, *hè* (residual = -7.7), *hoor* (residual = -11.6) and *man* (residual = -11.1), dropped drastically, scoring significantly below the expected proportion values of these conditions, $p < .0001$.

There are three more significant deviations from a random distribution to report. First of all, in both the *hè*- and *hoor*-condition, the correct response (i.e., *hè* in the hè-condition, and *hoor* in the hoor-condition) was chosen significantly more than one would expect from a random distribution (with a standardized residual of 6.7, $p < .0001$). However, the amount of correct responses come by far not near the effect we can observe from the response *zo* in the *zo*-condition, and it is also not the case that we observe a drop in the proportions of *hè* and *hoor* in non-corresponding conditions (as we did for *zo*), indicating that the participant group could not uniformly recognize the base sentence of the *hoor* and *hè*-conditions as such. Also, it is the case that the multinomial logistic regression analysis shows that there are significantly more *hoor* (than *hè*) responses in the conditions '*hoor*' and '*man*' compared to the condition '*zo*', respectively β1 = 2.193 (95% CI 1.294 – 3.717), Wald = 8.504, p < .004 and β1 = 2.291 (95% CI 1.350 – 3.888), Wald = 9.437, p < .002. This is not the case for the proportion of *hè* responses, which is not significantly larger in the *hè*-condition compared to the *zo*-condition, β1 = 1.013 (95% CI .600 – 1.710), Wald = .014, p < .960 (p > .05). This means that even though the proportion of *hè*-responses in Block 3 is larger than the proportions of *hè* in the previous two blocks, the likelihood of participants choosing *hè* as the right answer for the *hè*-condition does not deviate from the likelihood of participants responding *hè* to the *zo*-condition. However, for the proportion of *hoor*-responses, there is an increase of the response *hoor* in the conditions '*hoor*' and '*man*' compared to the previous two blocks. The likelihood of a participant choosing *hoor* in the *hoor*- and *man*-condition is significantly bigger than the likelihood of responding *hoor* in the *zo*-condition.

It is thus the case that *hoor* is not only responded more to sentences of the *hoor*-condition, but also to sentences actually belonging to the *man*-condition. In fact, the response *hoor* (50.52%) is preferred over the response *man* in this condition (24,22%). It is the case that for the *man-condition*, the response *hoor* is chosen more than one would expect from a random distribution, with a standardized residual of 6.9, p <.0001, while this is absolutely not the case for the actual correct response *man*, standardized residual = 3.5, p > .05.

3.3.1 Accuracy

The accuracy data is visualized in Figure 12, which mainly shows effects reflecting the results that were discussed above. However, it should be noted, that while we found no significant effect of the Chi-square test in Block 1 for the actual responses, there is a significant effect for the mean accuracy data of Block 1, X^2 (3, N = 1536) = 45.60, p < .05, as well as for the other two blocks, Block 2, X^2 (3, N = 1536) = 31.24, p < .05 and Block 3, (3, N = 1536) = 267.9, p < .05. Again, we conducted a post-hoc analysis to find which cells of the contingency table were responsible for these significant effects. Standardized residuals were again calculated, but this time a Bonferroni correction of .05/24 was applied (as there were less cells in this table) resulting into a corrected significant p-value of .003.

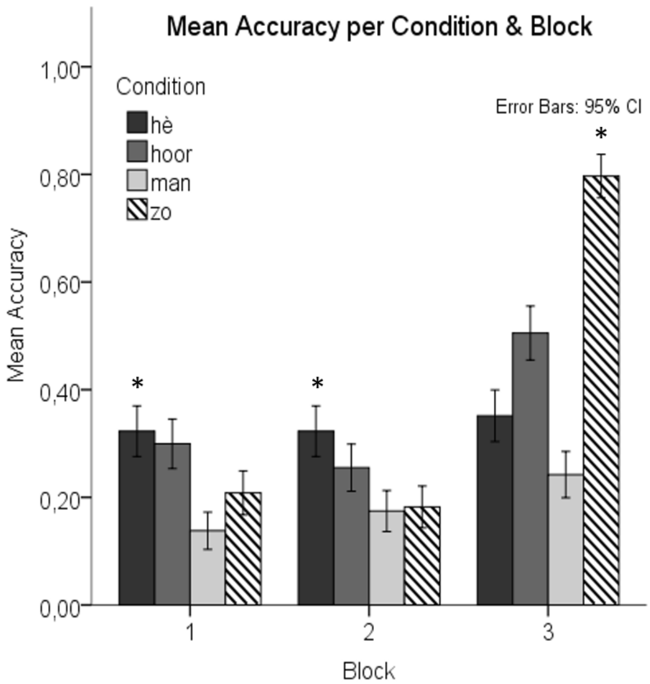

Figure 12. *Mean accuracy (in percentage) displayed per condition (*hè, hoor, man *and* zo*) and block (1-3). Stars indicate a significant above-chance performance (based on a post-hoc standardized residual analysis of chi square contingency-table test results, taking into account a Bonferroni correction of .05/24).*

The post-hoc analysis reveals that the Chi-square effect of Block 1 can be attributed to the mean accuracy proportions of *hè*, which were higher than expected (residual = 4.3, p<.003) and the accuracy proportions of *man*, which were significantly lower than expected (residual = -5.5, $p < .003$). This indicates that, while participants performed at chance-level with the conditions *zo* and *hoor*, they scored slightly above chance-level for the condition *hè*, and below chance level for the condition *man*. Also in Block 2, the mean accuracy of the *hè*-condition is responsible for the significant effect reported for the Chi-square test of Block 2, residual = 4.8, $p < .003$. This is in stark contrast with the results obtained in the final block. Here, it again becomes clear that participants are good at anticipating the proper continuation of the *zo*-condition, residual = 14.6, $p < .003$. Participants perform at chance-level for the *hoor*-condition (with a standardized residual of 1.4, $p > .05$) and even below chance-level for the conditions *hè* and *zo*, respectively, with a standardized residual of 5.5, $p < .003$, and 10.5, $p < .003$.

3.3.1 Certainty

In Figure 13, the certainty judgments are plotted against the response accuracy, separated per block. A Spearman's Rho correlation test shows that there is a significant strong negative correlation between accuracy and certainty in Block 1, $r_S = -.067$, $p < .05$; no correlation between accuracy and certainty in Block 2, $r_S = -.007$, $p < .05$; and a very strong positive correlation for certainty and accuracy in Block 3, $r_S = -.228$, $p < .05$. So only in Block 3 the accurate responses correspond to the answers participants were most confident about.

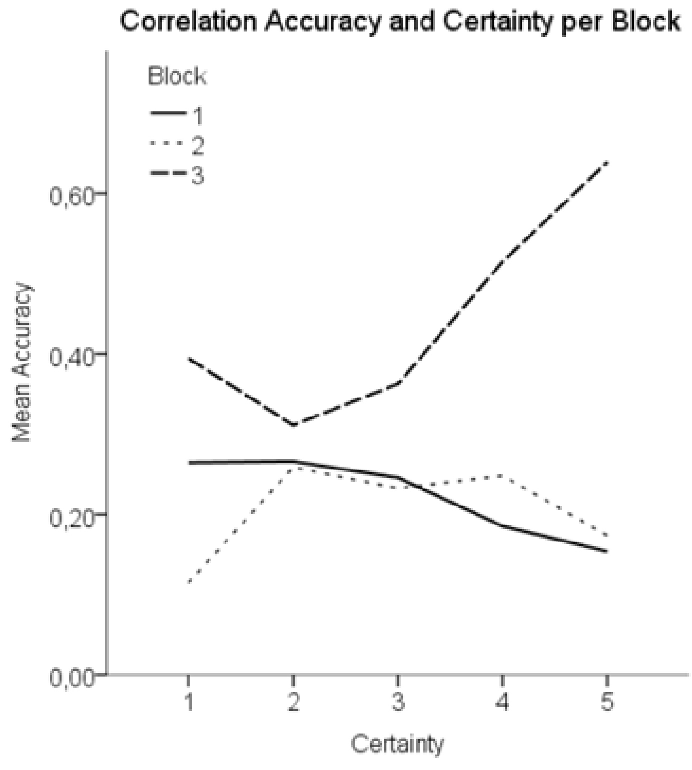

Figure 13. *Correlation of mean accuracy with certainty judgments provided on a Likert-scale of 1-5 (1= very uncertain, 5=very certain), displayed per experimental block (1-3).*

In Figure 14, the certainty judgments for Block 3 are plotted against the response accuracy, separated per response type. When looking at the correlation between accuracy and certainty in Block 3 for each response type separately, it can be observed that the correlation is the strongest for the response '*zo*'. Also, the responses of *hè* stand out, as the accuracy is actually the lowest for the responses that scored highest on the Likert-scale (the responses reflecting a high level of certainty from the participants on choosing the correct answer).

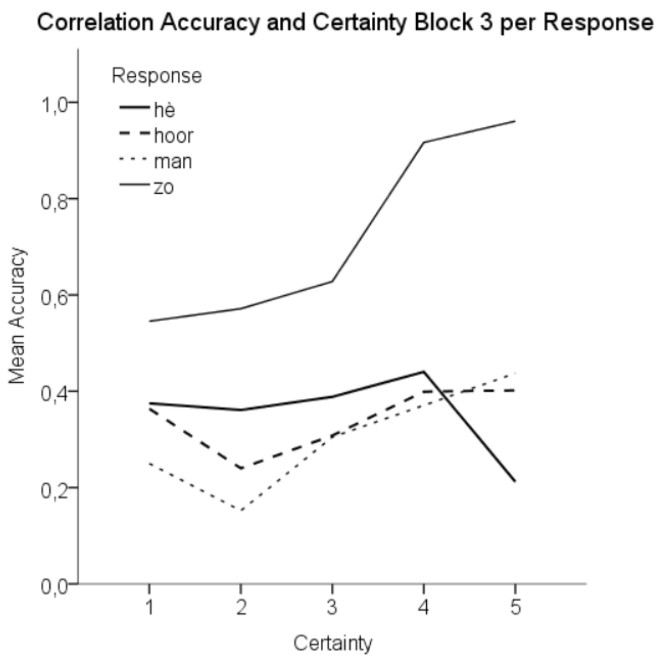

Figure 14. *Correlation of mean accuracy with certainty judgments provided on a Likert-scale of 1-5 (1= very uncertain, 5=very certain) from experimental Block 3, displayed per response type.*

3.4 DISCUSSION

In this experiment I investigated whether Dutch listeners make use of early prosodic cues to anticipate the upcoming of *hè*, *hoor*, *man* or *zo*. Based on the prosodic and semantic properties of the four types of sentences, it was hypothesized that listeners would be able to anticipate upon the right sentence-ending, as there are already consistent prosodic cues available in the information preceding the final element (see Section 2.3).

The results of the experiment indicate that participants were not that good at anticipating the end of the sentence across conditions. In Blocks 1 and 2, participants scored close to chance-level, only scoring above chance-level for the condition *hè*. Note, though, that the small above-chance effect of the correct *hè*-responses does not go hand in hand with a decrease of that choice in the other conditions, suggesting that participants did not uniformly associate this continuation with sentences from the *hè*-condition. Also the probability that *hoor* was responded for the *zo*-condition was higher than the likelihood that *hoor* was responded for other conditions.

However, these above-chance effects are not that big, and possibly due to the fact that *hè* and *hoor* are more frequently[11] used in sentence final position. As indicated by the results, *hè* and *hoor* are overall chosen more often than the other two conditions, and could therefore be used as a 'default' option when there is no clear indication for other elements. Also, the certainty results indicate that there is no significant correlation between the first two blocks and certainty indications, suggesting that participants could not reliably recognize conditions at this point, that is, after hearing the subject and verb of the sentence.

In Block 3, there is a clear-cut recognition of sentences from the *zo*-condition, where participants not only strongly preferred the *zo*-continuation above other possible continuations, but the amount of *zo*-responses in other conditions clearly dropped. This indicates that participants really could tell that the *zo*-continuation belonged to one particular type of sentence.

[11] This claim is based on the intuition of my participants. As far as I know there are no frequency ratios for man and zo in final position. However, for hoor and hè I refer to Schelfhout et al. (2005).

When considering the confusion matrix of the responses, it seems that participants are able to anticipate upon *hoor* to some extent, when provided with all the utterance material up to the particle (Block 3). Participants were more likely to correctly choose *hoor* as the continuation of *hoor*-sentences, than they were likely to choose *hoor* in other conditions (such as the condition *zo*). The certainty data shows that there is a correlation in Block 3 for the conditions *hè*, *hoor* and *zo*, in which accuracy correlated with certainty. This means that, even though participants did not uniformly correctly anticipate upon the end of sentences from the *hoor*- or *man*-condition, confident responses correlated with being the correct response. This means that, at least for some items, some people could use the prosodic information before the final element to identify the correct continuation. This is not the case for the *hè*-condition, where participants scored at chance-level in Block 3. Participants thus failed to anticipate upon the *hè*-particle, and confused it with other final elements.

The task at hand was not undoable, as can be observed from the reliable recognition of *zo*-sentences, so the difficulty the participants had could be due to the fact that the prosodic cues distinguishing *hè*-sentences from *hoor* and *man*-sentences are not strong enough for listeners to pick up. This is interesting, because the tasks forces the participants to listen carefully for any hint that can help them to predict the correct answer, and I did identify prosodic cues that differentiate the different type of sentences in the previous experiment. Possibly, these cues can be associated to a specific condition if we emphasize them by manipulating them artificially, which would be an idea for future research.

For the *man* and *hoor* cases, the response *hoor* was significantly preferred above the actual correct response *man* in the *man*-conditions, but the confusion was not reversed (so the response of *man* did not increase in the *hoor*-condition). This indicates that participants confused *man*-type of sentences for *hoor*-type of sentences, but not the other way around. This is possibly due to a preference based on the attitude of a sentence, instead of the intention[12]. After carefully re-listening the stimuli used in this experiment, several *man*-sentences where uttered with an undertone of annoyance implying something like 'you should know'[13]. This attitude is not only audible in some of the *man*-sentences, but also in many of the *hoor*-sentences, as the 'you should know' attitude is very common (but by no means necessary) for corrective sentences with *hoor*. If participants identify *hoor*-sentences based on the 'you should know' undertone of the 'annoyance/impatience' attitude, it makes sense that they mistake *man*-sentences with this undertone for sentences with a *hoor* continuation. This would mean that, though participants are able to pick up on certain attitudes before the end of the sentence, they could not differentiate between different intentions at this point. This attitude could also play a role in the inability to distinguish the *hè*-condition from the *man*- and *hoor*-condition. As discussed in Chapter 1.1, the particles *hè* and *hoor* can have multiple functions (see Appendix A). It could be that participants imagined a different type of *hè/hoor* continuation than what was targeted specifically for this experiment (e.g. imagined a 'reminding' *hè* instead of the agreement-seeking *hè*). Previous research has shown that the specifics of the final particle can overrule prosodic cues before that (e.g., Heuven & Haan, 2002; Iwata & Kobayashi, 2012). So imagining a continuation that does not match with the base sentence might overrule the prosodic cues in the base sentence that mismatch with those of the imagined final element. In subsequent research, it might be useful to present the continuations auditorily, to avoid such possible effects.

[12] The intention of a sentence indicates the goal of the speech utterance (e.g. a correction), the attitude with which this intention is uttered adds pragmatic meaning to the sentence, e.g., a gentle or a hostile attitude.

[13] A concrete example could be example at a bakery, where someone is waiting to order, when another person jumps the queue. The person who was there waiting before him could utter something like: *ik was hier eerst hoor!* 'I was here first!', a bit angrily or annoyed.

4. EXPERIMENT 3: THE PLANNING OF DUTCH FINAL PARTICLES

4.1 INTRODUCTION

In the previous chapters, I discussed experiments that examined the prosodic cues available in the linguistic signal before a particle occurs (Chapter 2: Experiment 1) and whether Dutch listeners are able to pick up on these cues and anticipate with which final element a sentence ends (Chapter 3: Experiment 2). I found that the prosodic cues that could be used to distinguish out-of-the-blue sentences with different final elements (i.e. *hè*, *hoor*, *man* and *zo*) were not used to anticipate upon a final element even when the experimental paradigm biased for anticipation (with a gating task). However, even though the prosodic differences between the different sentence types were not used by the listener to anticipate upon the following final element, these differences were present in the speech output. Could this mean that speakers already anticipate upon the final element they are going to use, and do these prosodic differences found in Experiment 1 reflect some form of planning?

The question remains whether the speaker plans a final particle ahead or whether they integrate the particle at a later stage of production. This question is about how incremental and how far ahead a sentence is planned in production, as discussed in Section 1.2.1 of Chapter 1.2. While there is clear evidence that we do not need to plan whole sentences before we start speaking (Brown-Schmidt & Konopka, 2008; 2015; Brown-Schmidt & Tanenhaus, 2006; Griffin, 2001; Levelt & Meyer, 2000; Schriefers et al., 1998; Zhao & Yang, 2016), several accounts argue against a strong incremental approach in language production, which assumes that the sentence plan is developed in the same order as the words of an utterance are produced. Momma et al. (2015; 2016) showed that while English and Japanese speakers are not yet committed to the verb when they start uttering external arguments of the verb (like the subject of a transitive clause), they were committed to the verb when they started their utterance with an internal argument of the verb (such as the object of a transitive clause or the subject of a passive sentence). They conducted a Picture-word-interference (PWI) paradigm and showed that semantic interference of the verb, only occurred when participants started their utterance with an internal argument of the verb, and not with an external argument. This suggests that verb look-ahead selectively occurs before the articulation of internal arguments of the verb, reflecting the close linguistic relation between the verb and its internal argument. Thereby the results by Momma et al. support what has often been assumed in theoretical research (e.g. in generative syntax the object of the verb is assumed to be the complement of the verb, while the subject of the verb is projected in the specifier position of *v*P). Thus, Momma et al. (2015) suggest that the scope of planning is influenced by linguistic dependency, which would imply that the parser is respecting the rules of the grammar and that the correspondence between grammar and parser is very tight (as suggested by Phillips, 2003).

As discussed extensively in Chapter 1.1.2 SFPs such as *hè* and *hoor* are analyzed as syntactic heads that take the entire sentence as their complement. In this sense, the relationship between a verb and its object is not that different from the relationship between a particle head and its sentential complement, the only difference being that the particle head-complement relationship resembles dependencies at a higher level of the sentence. However, while the relationship between a verb and its internal argument is a strong thematic relationship, the relationship between a sentence final particle and its sentential complement is a pragmatic one and involves no linguistic dependencies. If both relationship types function alike (hypothesis 1), some planning could be involved in sentences with SFPs where the particle occurs linearly later than its complement (as observed for the verb head and its object complement by Momma et al., 2016). Whereas, if both types do not work alike, then it could be the case that SFPs do not require planning, and that the head-complement relationship at the pragmatic level functions differently (hypothesis 2). It could thus be imaginable that the selection relationship between a final particle and its preceding sentence resemble that of a verb with its internal argument, and that some planning is involved when the particle occurs later than its complement. However it could also be the case that such planning is not necessary at the pragmatic level of the sentence. The latter could have implications for syntactic accounts that incorporate pragmatic and discourse information into the syntactic structure, mirroring lower level grammatical relations (e.g. Haegeman, 2014; Sybesma & Li, 2007; Munaro & Poletto, 2003), as such an approach incorporating both pragmatic and grammatical information into the syntactic structure in a uniform way might not be justified by empirical data.

In this experiment, I examined the production process of the intentional SFPs *hè* and *hoor* in Dutch to investigate whether the particles are planned in advance, or not. The current study was inspired by research by Momma et al. (2016). However, since my research topic varies in multiple aspects from the original study, several adjustments needed to be made to the research paradigm in order to fit this study. One of the adjustments is the type of target utterance, where Momma et al. (2016) used short utterances containing either a subject and a verb (SV) or an object and verb (OV), this study targets complete transitive sentences, including a final element. This adjustment is not problematic for the picture-word interference design, and full transitive English sentences have been used successfully by Momma et al. (2015). I required my participants to name event pictures to ensure that the utterances of the participants are similar in phonetic and structural complexity. Also, I added one more dimension to the pictures (color). The color dimension was implemented because the semantic PWI paradigm proved to be impossible for the SFP-containing sentences I wanted to elicit, due to it being difficult to find appropriate semantic competitors for SFPs. As a replacement, a paradigm is proposed that makes use of associative facilitation. The current paradigm exploited the color-particle association that participants were trained on during a training phase of the experiment by using related and unrelated color terms as distractors. Colors have been used in previous production studies to symbolize a specific language mode (e.g. in bilingual studies: Christoffels et al., 2007) or a constituent type (Allum & Wheeldon, 2007). In the current experiment, participants were trained to associate the color of a picture with a specific experimental condition (e.g. with an SFP like *hè* or *hoor*).

If the connection between a specific color and an experimental condition is strong, a color term prime matching the color of the picture is expected to influence the speech onset latency of the sentence production if the particle (either SFP or something else) associated with that color is processed prior to the speech onset. This is only hypothesized for SFPs such as *hè* and *hoor*, but not for the other final elements that operate at a structurally lower level, included in the experiment for comparison with SFPs (i.e., the deictic marker *zo* and the focus particle *wel*[14]). These final elements that are not considered SFPs, modify the predicate of the utterance (e.g. *wast een draak* 'washes a dragon'), but not the entire sentence. Therefore a difference in speech onset latency is expected between these two particle types. These expectations are illustrated in Figure 15 below.

[14] The observant reader might notice that the experimental condition *man* from Experiment 1 and 2 is replaced here for the focus particle *wel*. There are two reasons for this: First of all, from the results of Experiment 2, it seems that there is some overlap between the conditions *hoor* and *man* (as speakers confused stimuli ending in *man* with stimuli ending in *hoor*). To avoid any possible confounds from the two in this experiment, the condition *man* was excluded. Also, since this experiment addresses production processes, the structural position of the control condition should be taken into consideration. While in the previous experiments final elements were selected based on their semantic meaning, in this experiment control conditions are chosen based on structural position. While the addressee marker *man* adjoins high in the structure, the focus particle *wel* adjoins the VP in the type of sentences used for this experiment.

In Figure 15, the hypothesized process of naming the target picture in four different experimental conditions (two SFP targets with *hè* and *hoor*; and two final element controls with *zo* and *wel*) is schematized for two different types of distractor: a related (matching in color) prime or an unrelated (mismatching in color) prime. When presented with a color-term prime, the concept of that color is activated.

In the current paradigm, we assume that, after an extensive learning period in which participants associate colors to the four different experimental conditions, the concept of the color also activates the concept of the related condition. Then, when the picture (the knight is washing a dragon) is presented later, the participant recognizes the picture, and activates concepts of different elements within the picture (e.g. *draak* 'dragon', *wassen* 'to wash', *ridder* 'knight' the color of the picture). If the condition of the prime is related to the color of the picture, we assume that recalling the concept of the experimental condition is facilitated (produced faster) compared to recalling the concept of a color from a picture following an unrelated prime. This facilitation is expected for both target and control conditions. The crucial manipulation of this experiment is therefore the timing of this facilitation. We expect that if SFPs such as *hè* and *hoor* are planned in advance (hypothesis 1) then this facilitation should have an effect at speech onset in the target SFPs conditions, resulting in faster speech onset latencies. On the other hand, if *hè* and *hoor* are not planned in advance but incorporated later in the sentence (hypothesis 2), then facilitation should have an effect later, and not at speech onset. For the elements in the control condition, we assume that they will not show any facilitation effect at speech onset in any condition.

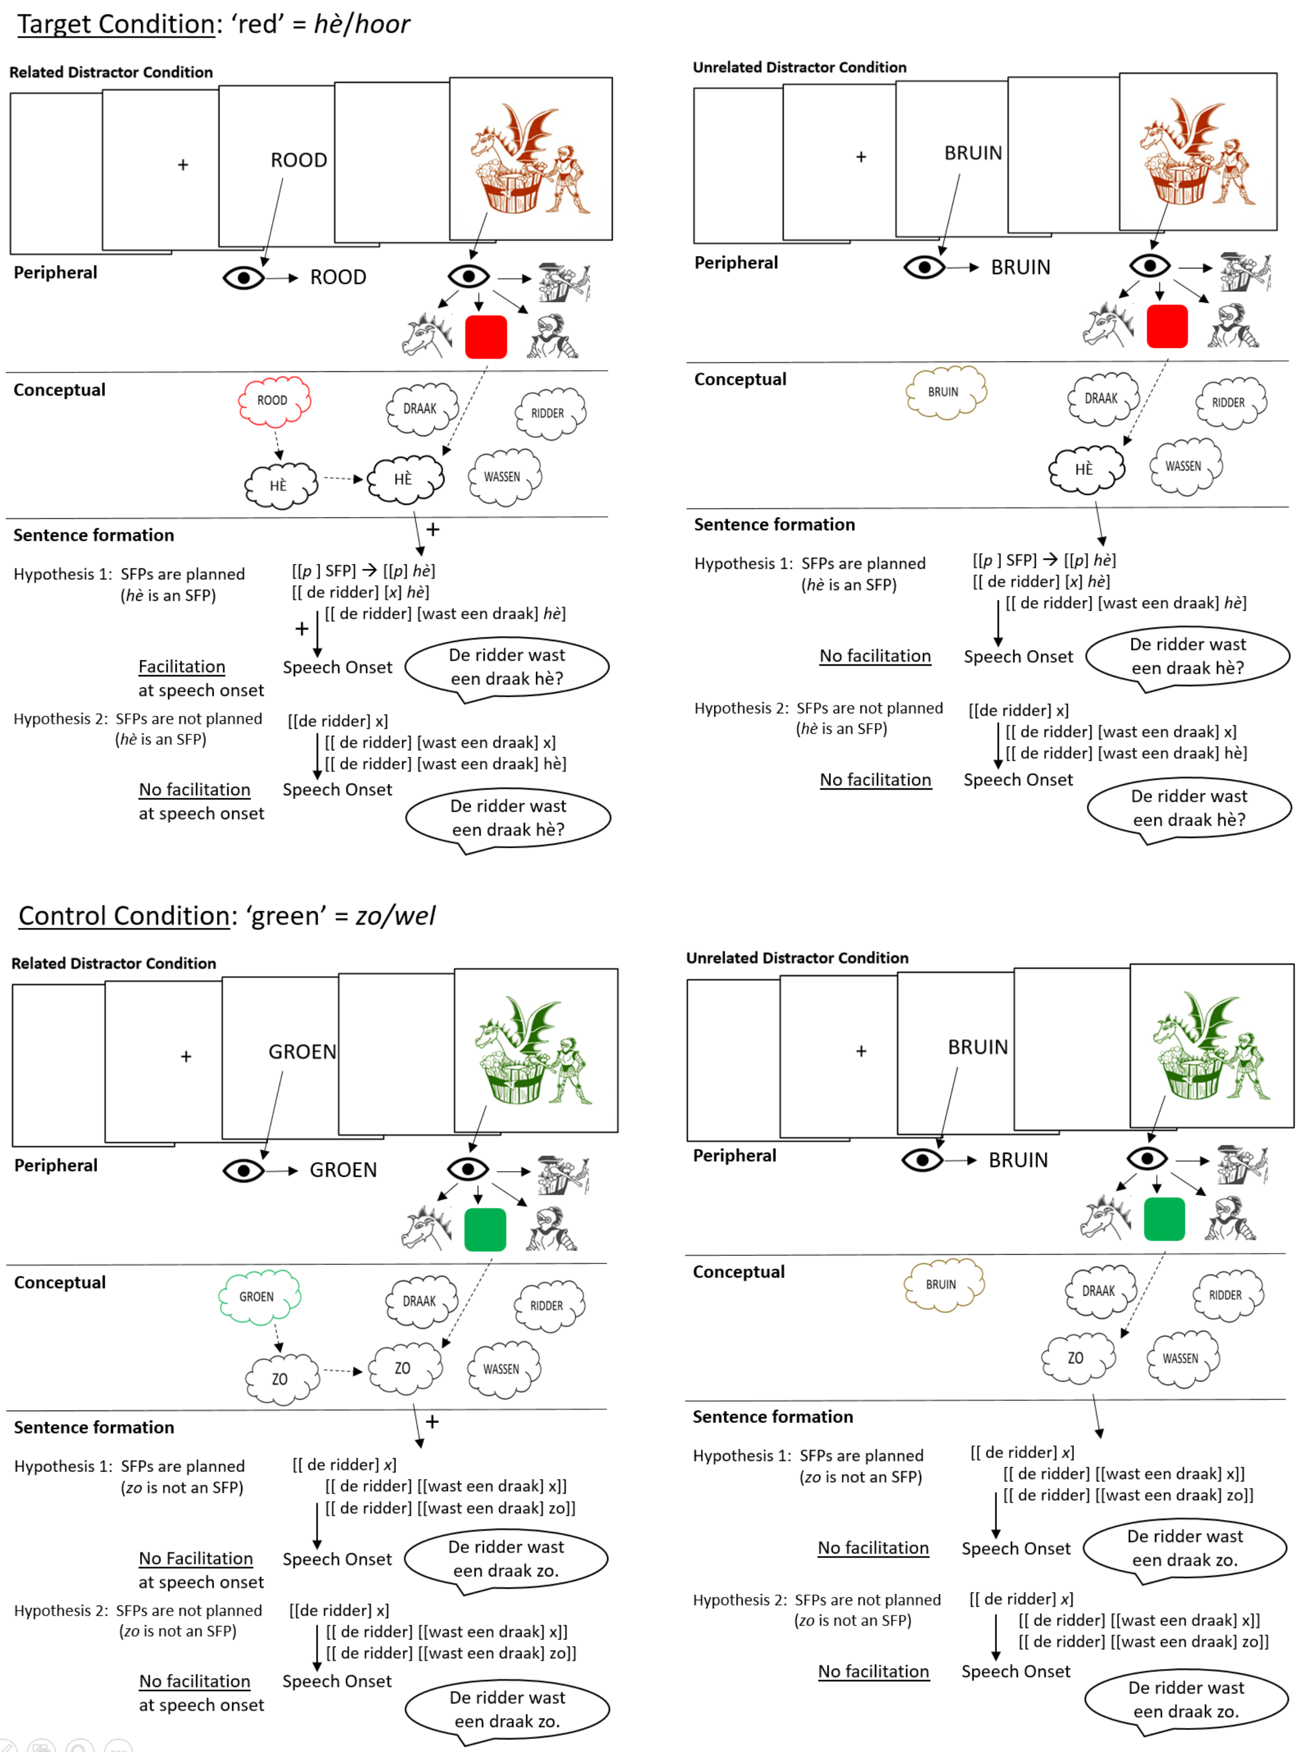

Figure 15. *Expectations about the effects of associative facilitation priming in the current experimental paradigm with respect to target and control conditions.*

4.2 METHODS

4.2.1 Participants

Twenty-four Dutch native speakers (9 male, 15 female), participated on voluntary basis in this experiment. Their age range was 17-32 years old ($\bar{x}$ = 23.08; SD = 3.97). Most of them were students affiliated with Leiden University. Nine participants had already participated in Experiment 1. This is not a problem for the results of Experiment 3, as any prior experience with the target sentences vanishes through the extensive training session every participant goes through in the current experiment.

4.2.2 Stimuli

Sounds: The recordings from Experiment 1 and 2 were used in the training phase of this experiment, where the sound of the sentence denoted by the picture was played to participants. Recordings of the same female speaker, as those used for Experiment 2 were used. The recordings for *hè*, *hoor*, and *zo* from Experiment 2 were included. Additionally 16 stimuli sentences were recorded the focus particle *wel* (e.g. *De ridder wast een draak wel* 'The knight IS washing a dragon.', see Appendix B). The recordings of the sentences without final element from Experiment 1 were used for the filler pictures. All audio stimuli were standardized to a mean intensity of 60 dB.

Pictures: Of each of the 16 sentences (Appendix B) a picture that displayed the action described in the sentence was self-designed, with the drawing software ArtRage Studio Pro (Version 3.4.5). Each of the 16 pictures was colored into five distinctive colors: red, yellow, green, blue and grey and, rather than adding a colored border as in previous studies using color (e.g., Christoffels et al., 2007), the entire picture was colored to ensure that participants viewed the color and the picture as a whole. This way, I tried to avoid strategies in which a participant could first start naming the picture without looking at the border by considering the color of the picture only at the end. The object of the sentence was always clearly distinguishable and detached from the central character or subject. A set of all pictures is included in Appendix F. All pictures were sized to 600x600 pixels.

4.2.3 Procedure

The experiment took place in the sound proof booth of the phonetics lab at Leiden University. Participants wore headphones. A microphone was positioned at an approximately 30 cm distance. The experiment was run using the presentation and recording software E-prime (Version 2.0).

The experiment was separated in two halves (separated by a big break of 5-10 minutes). In the first half, participants learned to name a set of pictures with one target item, one control item and one filler item. Then they performed a picture-naming task of the learned pictures. In the second half, participants followed the same procedure, for a different set of colored pictures with the other target and control condition (i.e. if they learned *hè* in the first half, then they learned *hoor* in the second half). An example stimulus set is provided in Figure 16. In order to avoid any influence that the color of the picture might have on the learning process of this color association, which color association was related with each experimental condition was counterbalanced across participants. The counterbalancing of color[15], half and condition resulted in 8 different versions of the experiment. The counterbalance scheme of the eight different versions can be observed in Appendix G.

[15] The color-combinations were restricted, as the picture color within a half either consisted of: grey, blue and yellow; or grey, red and green (and not for example, grey, green and yellow). Possible color combinations were chosen on basis of visual discriminability as the difference between green and red, and yellow and blue are more salient than the color combinations blue and green or yellow and red.

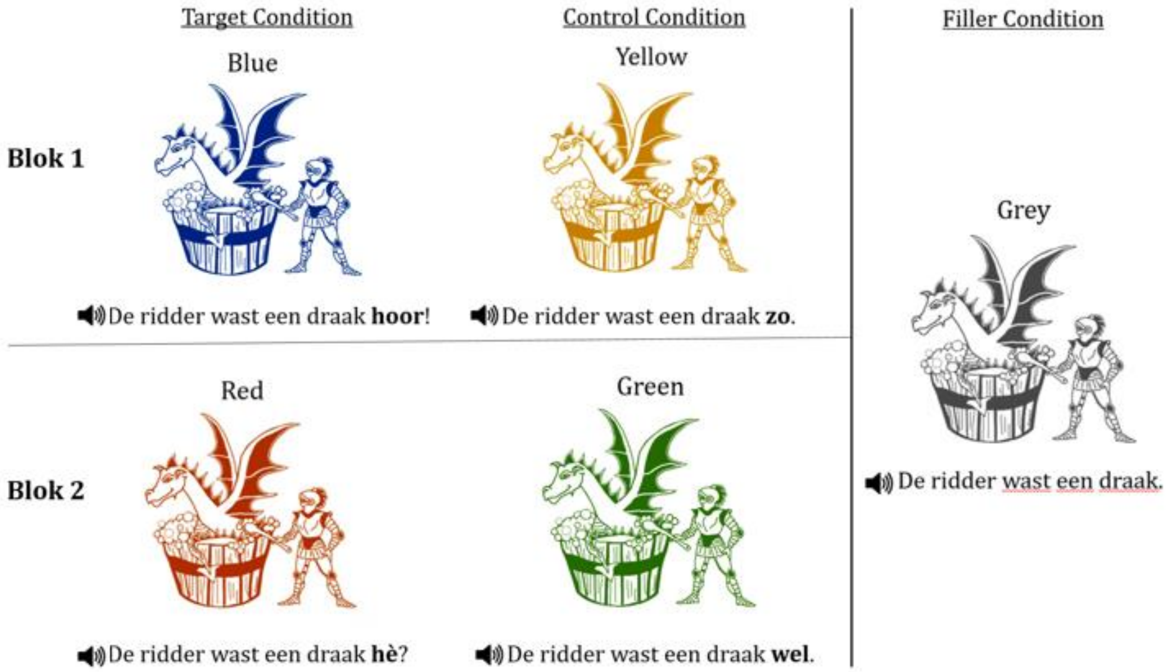

Figure 16. *Possible stimulus set of the sentence:* De ridder wast een draak*. 'The knight is washing the dragon'. In this version, target pictures are blue (*hoor*) and red (*hè*), and control pictures yellow (*zo*) and green (*wel*). The fillers are always in grey.*

Each half of the experiment consisted of 4 different components. First a training phase, participants were presented with three different tasks where they were encouraged to learn the right utterance of the pictures, and make the right color-condition associations. Secondly, in a learning task, participants were presented with 48 pictures (16 target, 16 control and 16 fillers) in random order. The color of the picture differed per condition, and participants learned to associate the three experimental conditions (target/filler/control) with three different colors. Together with the appearance of the picture, participants heard simultaneously a recording in which a female voice described the event in the picture, which depending on the picture's condition, was uttered with or without a final element. Participants were instructed to listen carefully, and to learn to name the picture as they heard it on the recording with respect to its content words and structure[16]. In this learning phase, the participant pressed the space bar to continue and to see the next picture and there was no time limit. During the second part of the training phase, participants were trained on naming the pictures (48 items) they had just learned. The procedure was the same as in the preceding phase, only now participants did not hear the utterance, but had to name it themselves. They could name the pictures in their own pace. If participants had forgotten the right utterance for a picture, they could press 'l' to listen again to the recording. Finally, to conclude the training phase, and to check whether participants had learned the right color-condition association, we conducted a pre-test. During the pre-test, participants were presented with the learned pictures, and a recording of the picture utterance (24 trials). The recording was either congruent (12 trials) or incongruent (12 trials) with the picture. Participants indicated whether the picture and sound were a correct match ('V') or

[16] Participants were explicitly instructed that they did not need to remember the proper 'prosody' of a sentence, and that they could use different articles (definite/indefinite) if that came more naturally to them. For example, some participants preferred to use definite articles for the objects, even though an indefinite article was used in the audio files. For this experiment, as long as it is monosyllabic, the precise content of the determiner is not important.

whether it mismatched ('X'). There was no time limit to this task either, and feedback was provided after each trial. The training phase lasted approximately 15 minutes in total.

In the last and fourth component of the first half of the experiment, participants were asked to name the pictures they had learned, and their productions were recorded. In this naming-task, participants were first provided with a practice block to get used to the task, before the real experiment started. The trials consisted of one of the learned pictures in one of the different conditions (target, control and filler), presented together with a distractor: a related distractor (which for the experimental conditions was the color term for the picture color of the condition), or an unrelated distractor (a monosyllabic color term not associated to the experiment). To provide a concrete example, for blue pictures associated with the target *hoor*, for example, the related distractor was BLAUW 'blue' and the unrelated distractors were PAARS 'purple' or BRUIN 'brown'. For yellow pictures associated with the control *zo,* for example, the related distractor was GEEL 'yellow' and the unrelated distractors also PAARS 'purple' or BRUIN 'brown'. Fillers (grey pictures) were preceded by all possible distractors[17]. In total there were 120 experimental items: 30 targets (15 pictures X 2 distractor conditions), 30 controls (id.) and 60 fillers. The experimental procedure of the naming task was divided into two balanced blocks, to ensure that the same picture (with different distractors) would not appear within the same block. Trials were selected from these blocks at random. For the practice trials, a picture[18] that did not appear in the real experiment was presented six times (2 target +1 filler X 2 distractor conditions).

The second half of the experiment was built up the same as described for the first half of the experiment above (first a training phase consisting of: a learning task, a practice picture-naming task, and a pre-test, then the picture-naming task). This time however, participants learned the other half of the experimental conditions, e.g. in the first half of the experiment they learned the target *hoor* and the control *zo*, so in the second half of the experiment they learned the target *hè* and the control *wel*.

The structure of a trial is displayed in Figure 17. First a fixation point of 750 ms appeared on the middle of the screen. Next, the distractor (font size: 48, color: black) appeared for 200 ms, followed by an inter-stimulus-interval of 200 ms[19]. Then the picture which participants had to name appeared. Recordings were made with the voice-key function of the SRR box of E-prime. After 3500 ms the picture disappeared, and the next trial automatically started after an inter-trial-interval of 1800 ms.

---

[17] This was done to ensure that participants could not use the congruent prime in the target and control conditions to predict the color of the upcoming picture. With the distractor distribution used in this experiment, the distractors GEEL 'yellow', ROOD 'red', BLAUW 'blue' and GROEN 'green' appear as often before a picture with a matching color, as they appear before a picture with a mismatching color.

[18] This picture was *de hacker maakt een plan* 'the hacker makes a plan'.

[19] These priming intervals are based on Bloem et al., (2004) and chosen because a -400 ms SOA seems long enough to allow for access to semantic/conceptual information, and a 200 ms display allows for enough time to perceive the prime, although it seems it might not be long enough to use the prime strategically.

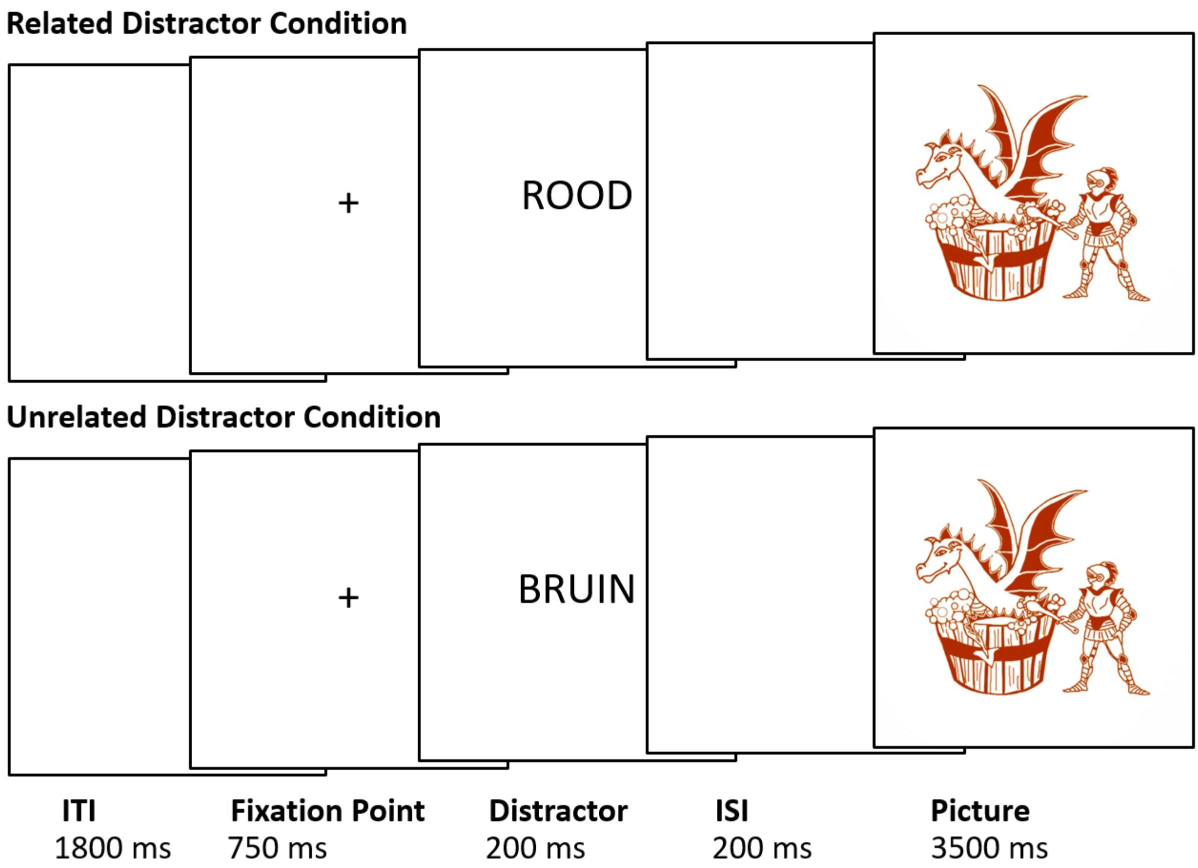

Figure 17. *Example trials of a red image depicting: De ridder wast een draak. 'The knight is washing the dragon' (associated to a target or control item), in a related distractor condition (ROOD 'red') or in the unrelated distractor condition (BRUIN 'brown').*

4.2.4 Analysis

Reaction time and accuracy data was obtained from the pre-test to check whether participants had learned the right associations. We did not exclude participants based on this information since the lowest accuracy score was 87.5%, and we considered that to be high. The recordings from the target and control conditions of the picture-naming task were inspected and those that contained considerable hesitations, mistakes or disfluency were excluded from further analysis (16.2 %). For the analysis, following the filtering of the data applied by of Momma et al. (2016) in their study, we excluded reaction times that deviated more than 2 standard deviations from the mean, which resulted in an additional rejection rate of 3.8% (leaving 2306 items for further analysis). The speech onset latency for each recording was manually determined using Praat (Boersma & Weenink, 2016). The onset of the utterance was set at the burst of release of the plosive [d] (ignoring the medial phase (silence) of the plosive). When determining the onset latency of a condition, I was blind to the distractor conditions of the recordings. The mean onset latencies were visually inspected for normality and analyzed per experimental condition, distractor condition and type with SPSS (IBM Corp, 2012). Since the original onset latency data was not normally distributed, the data was log-transformed for statistical analyses. This data was further analyzed with the statistics software R (R Core Team, 2013) and *lme4* (Bates, et al., 2015) to perform a linear mixed effects analysis of the relationship between the onset-latency and sentence type, experimental condition and distractor condition. The *t*-values of the fixed effects were reported to be significant at a ,05 level if they exceeded 2.0. Also the dispersion from the mean was analyzed with an independent t-test, to find whether there was a difference between the related and unrelated distractor condition.

4.3 RESULTS

Since the raw onset latency data was negatively skewed (as displayed in Figure 18a) we conducted a log-transformation, which resulted in a more normal distribution (Figure 18b).

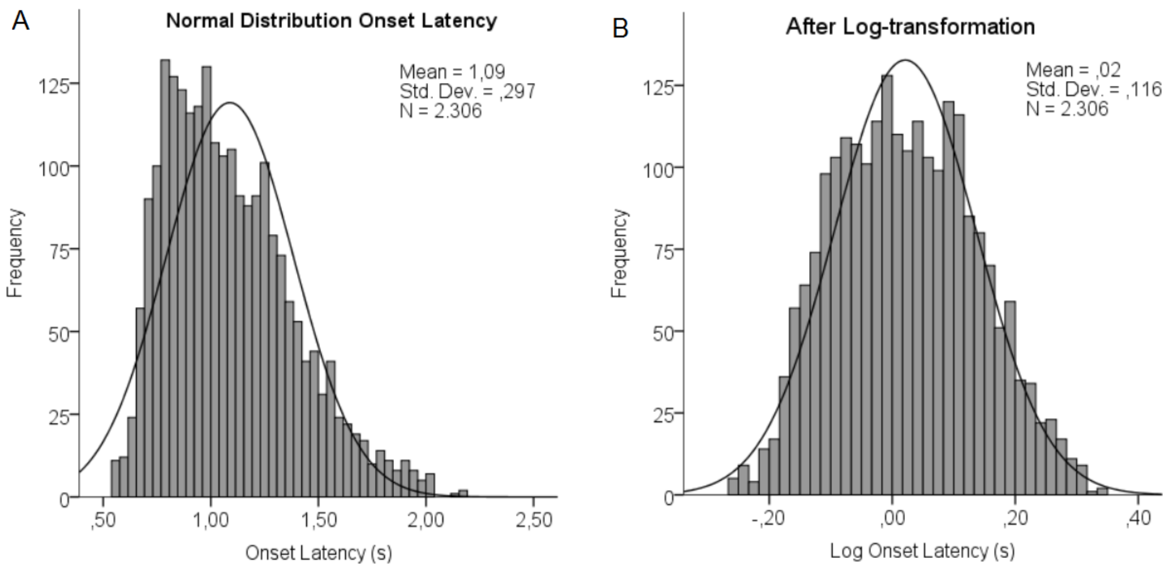

Figure 18. A: *The distribution of the onset latencies obtained from Experiment 3.* B: *The distribution of onset latencies after log-transformation.*

The essential manipulation of the experiment is the distractor condition, which was either related to the final element or unrelated. The prime was hypothesized to facilitate the onset latency at the target condition (sentences with *hè* and *hoor*), but not in the control condition (sentences with *zo* and *wel*). Data showing the effect of the related and unrelated distractor is displayed in Figure 19.

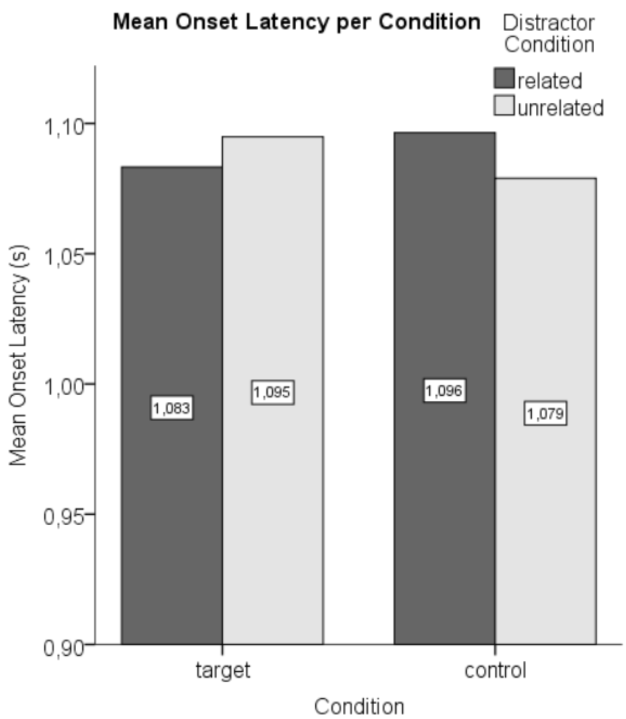

Figure 19. *Mean onset latency (s) per condition (x axis: target or control) grouped in target and control and distractor condition (black and grey: related or unrelated prime).*

First, there is a marginal effect of distractor condition (related or unrelated) on both the target and control condition as can be observed in Figure 19. In the target condition, the related distractor has a facilitating effect ($\bar{x}$=1.083, SD =.288) compared to the unrelated distractor ($\bar{x}$=1,095, SD=.309), while the effect is reversed for the control condition, in which the unrelated distractor evokes shorter onset latencies ($\bar{x}$=1.079, SD=.287) than the related distractor ($\bar{x}$=1.096, SD=.304). A mixed linear model with the interaction between experimental condition (target/control) and distractor condition (related/unrelated) was constructed with participant number and sentence item number as random slopes. The model shows no significant effect of distractor condition (b=.004, SE=.006, t=0.61) experimental condition (b=.003, SE=.006, t=0.56), or the interaction thereof (b=.008, SE=.008, t=-1.02) on the log onset latency (p>.05). The standard deviations of these averaged distributions are on the high side. We find standard deviations of +-300ms, while, for example, the standard deviations reported for sentence production in Momma et al. (2015) range between 151-228 ms. This indicates that there is a lot of variance within the data, which can be mostly attributed to variation between and within participants. Within-subject variation can be attributed to multiple experimental factors, such as the color of the picture or the experiment's half (one target and control condition appeared in the first half of the experiment, while the other occurred in the second half). While these factors are balanced out by the counter-balancing scheme in the averaged data, it could have an effect on within-subject variation, since some of the experimental items are excluded from analysis. To see whether experimental half has also an influence on the grouped data, the data was plotted per half (see Figure 20).

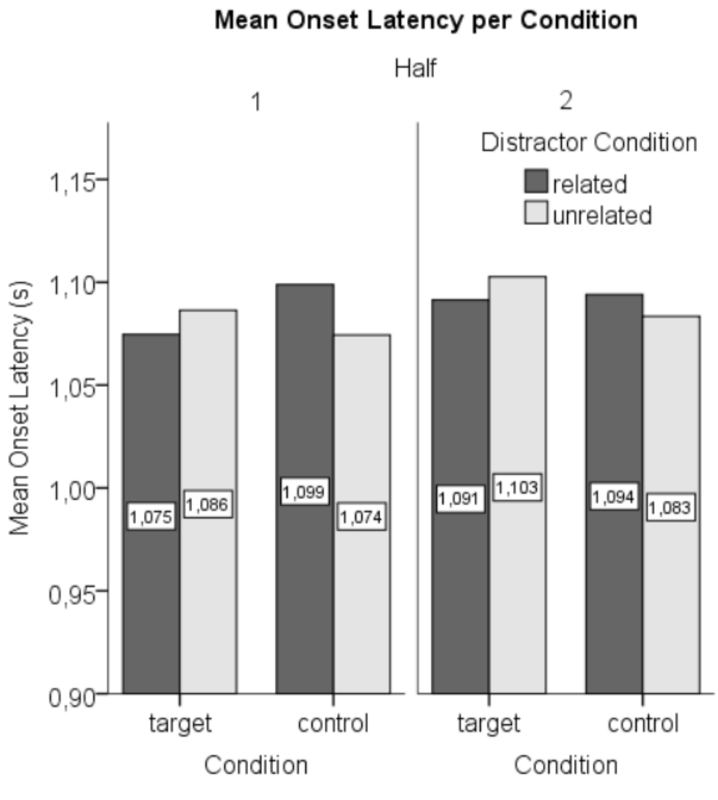

Figure 20. *Mean onset latency (s) per condition (target or control) and distractor condition (related or unrelated prime) separated into first and second half of the experiment.*

The experimental half has no significant effect on the distribution of the experimental condition relative to the distractor condition, as the data in Figure 20 shows. This is supported by linear mixed model that takes experimental half into account, in addition to the interaction between distractor condition and experimental condition (b=.000, SE=.003, t=-0.04). If we look within the condition to *hè*, *hoor*, *man* and *wel* respectively, there is no significant interaction with either distractor condition (b=.006, SE=.010, t=0.66) or experimental half (b=.005, SE=.011, t=0.52).

We checked with multiple *t*-tests whether the standard deviations in the related ($\bar{x}$=.166, SD=.150) and unrelated ($\bar{x}$=.167, SD=.146) conditions were significantly different in the target condition (*hè* and *hoor*), and whether the related ($\bar{x}$=.159, SD=.134) and unrelated ($\bar{x}$=.161, SD=.141) conditions were significantly different in the control condition (*man* and *wel*). This was for neither condition the case, target: $t(449) = -.095$, p=.925; control: $t(489) = -.205$, p=.838, meaning that the standard deviations for the related and unrelated distractors did not differ significantly from each other in both target and control conditions.

4.4 DISCUSSION

There is not that much research on sentence planning, let alone on the planning of pragmatic intentional final particles. The research conducted in this experiment is thus quite pioneering, and several adjustments needed to be made from existent paradigms (such as that in Momma et al., 2016) to create an experiment that addresses the question of this thesis about the planning of SFPs in Dutch. We created an experiment that manipulates the prime preceding colored pictures, which are associated in a training task with a particular final element. The target condition of the experiment contains sentences with the final elements *hè* and *hoor*, which are intentional final particles. The effect of the distractor prime on the target condition was compared to the control condition. As reported in the results section (Section 4.3) there is a mean difference between the related and unrelated distractor condition. In the target condition, the related distractor seems to facilitate sentence production (-12 ms compared to the unrelated distractor mean), while the opposite holds in the control condition, in which the related distractor is associated with longer onset latencies (+ 17 ms compared to the unrelated distractor mean). This effect is expected for the target condition (as displayed in Figure 15), as a congruent prime is assumed to facilitate sentence production at speech onset, only if the speaker is already planning the particle congruent with the prime. In the control conditions, in which it is assumed that speakers are not yet planning ahead for the final element of the sentence at speech onset, facilitation cannot take place. In fact, it might even be the case that related prime in the control condition slows the participant down, as a concept is activated that will not yet be used at that moment.

However, the mean effects we observed for the data of this experiment are not significant. This is not necessarily due to the fact that the difference between the conditions is too small, but can be accounted for by the fact that the standard deviations from the mean are very high. This is the case both within and across participants. Though the differences within participants can be explained through the experimental paradigm (in which color and order of appearance of the conditions is counterbalanced with eight different versions, see Appendix G), the differences across participants are particularly big. In other words: participants vary in the effect the related and unrelated primes have on their onset latencies. This can mean two things: first, this variance can mean that the distractor primes have an effect (if the distractor primes would have had no effect on the onset latencies, we would expect participants to show more or less the same distribution for the related and unrelated condition). Second of all, this can mean that something is causing this dispersion. Though it is not clear what precisely underlies the between-participant differences, I speculate it might have something to do with individual learning effects or the micro-variation of the sentence final particles. The most crucial manipulation of this experiment is the color-condition association, in which participants are taught to associate a particular color with a particular final element of either the target or control condition. Though the pre-test showed that all participants reliably learned these associations, participants differed from each other in how well and quickly they could use this association to name the pictures during the experimental task.

So while all participants have acquired the intended color-condition associations, the strength of this association could vary between participants. Since supposed related primes in this experiment use this association (they target a color concept, which should be associated to the learned particle) differences in the strength of this association could have consequences for the effect the prime has on the onset latency. As for the micro-variation of the particles *hè* and *hoor* (discussed in Section 1.1.1, and further illustrated in Appendix A), it is possible that participants interpreted or produced particles different from the intended interpretations (*hè* as an agreement-seeking particle, and *hoor* as a correction) giving rise to variation among speakers. Though we tried to avoid this effect by having the participants hearing the target sentences, rather than reading them, during the training phase, this possibility cannot be excluded. Other factors, such as individual task-dependent strategies, could also have contributed to individual variation. An example of a task-dependent strategy is how participants name the pictures, i.e. whether they first consciously decide what the color-condition association of the picture is, or whether they start naming the picture immediately before they consider the color of the picture. This strategy could be influenced by a natural property, i.e. whether you are quick in responding or slow. Additionally, we used a within-participant design, in which all participants were presented with all experimental conditions. This means that all items were repeated many times, which might have influenced the semantic facilitation effect greatly (all conditions were highly activated already). The filler in this experiment depicted the same event, but was uttered without a final particle. No unrelated fillers (completely different pictures, depicting a different sort of sentences, e.g. intransitives) were used in this experiment, as this would have made the experiment too long (the experiment lasted an hour). Thus, it is very likely that the continuous repetition of the same 15 pictures could have contributed to the highly activated state of all experimental items.

Though we obtained no significant results from this experiment, we think it is worth to investigate the different properties and set-up of the research paradigm proposed in this chapter further in the future. Especially as a similar paradigm has proven to be effective in other sentence planning research (Momma et al., 2015; 2016). First of all, a longer or more spread out training phase might help to gain a more homogenous color-condition association across participants. In addition, an in-between-design, in which not all participants are exposed to all conditions, might help to make the experiment more feasible and would allow for the inclusion of unrelated filler items and less repetitions. Furthermore, an adjusted paradigm in which participants are provided with some contextual cues for the target sentences during the training phase might help to ensure one homogenous interpretation across participants, avoiding the problem of micro-variation of *hè* and *hoor*. One could also consider altering the manner of priming used for this experiment, i.e. by displaying an actual color as prime (e.g. RED in red letters) or by using a different way of priming/distracting. In this experiment, I chose not to use a repetition prime, like a congruent repetition of color or lexical item, to make sure that the outcome of the experiment could be attributed to a conceptual (planning) reason. When using repetitive primes, you cannot be certain whether the effect you observe is due to an effect at the conceptual level or an effect at a lower level of processing (such as the phonological level, or visual-recognition level). However, it might be the case that the learned-associative priming we aimed for in this paradigm is simply too far-fetched, and an alternative method needs to be sought to detect planning mechanisms in sentence production.

While the current experimental paradigm was not sufficient to answer the questions addressed in this chapter concerning the planning of intentional final particles in Dutch, I believe variants of this paradigm should be considered before rejecting the method for answering this question.

5. General discussion

Speech production models assume that every utterance starts with a certain intention and the formulation of a conceptual message (e.g. Bock & Levelt, 2002; Garrett, 1980; Gambi & Pickering, 2016). However, it has also been noted, that many important cues about the intention and attitude of a sentence occur sentence-finally, e.g. sentence final intonation (boundary tones) and sentence final pragmatic particles. This raises interesting questions about the planning and expression of intentional linguistic elements. This thesis focused on intentional final particles in Dutch, and aimed to shed light on this process formulating a very basic question: Are sentence final intentional particles planned in advance (i.e. is the speaker's intent directly mapped onto a linguistic item) or are these particles only integrated at the end of a sentence? The validity of this question is supported by theoretical work on the syntactic properties of sentence final intentional particles. It is often assumed that the final particle selects for the entire utterance as its complement, and thus plays an important role from the beginning of sentence formulation. It has even been argued, that though the final particle appears at the end of an utterance, it actually appears head-initially in the underlying syntactic structure (Haegeman, 2014; Hsieh & Sybesma, 2011; Munaro & Poletto, 2003). If such theoretical notion touches upon the truth, it is expected that even though they are overtly realized at the end of the utterance, particles are planned in advance before the encoding of the whole utterance takes place. In this thesis, this question was put to the test.

Three experiments were conducted to investigate questions about the production, perception and planning of the intentional final particles *hè* and *hoor* in Dutch. In the first experiment, the production of sentences with different final elements, *hè*, *hoor*, *man* and *zo* were recorded and compared. The final elements can all appear in the final position of the sentence, but vary in their meaning and function. The sentence final particles *hè* and *hoor* both express intentional meaning, expressing the illocutionary force of the sentence. They differ from one another in the type of intention they convey, and the response that is expected from the hearer (agreement-seeking *hè* requires a response, while *hoor* does not). To see whether these final particles differ from each other in the sentence elements preceding them, their duration and f0-values where obtained and compared against each other, and against the non-intentional elements *man* and *zo*. The results obtained from the first experiment show that there are differences in both duration and f0-values before the final elements. In fact, the distribution of the different sentence types matches the relatedness in meaning and function: sentences with the particles *hè* and *hoor* are the most alike, both not differing too much from the prosodic structure of *man-sentences*. The sentences with the final particle *zo* differ significantly from the other three conditions, at almost every measure. The results of this experiment indicate that there already is some expression of intention/pragmatic function of the sentence before the final element appears. This indicates that the speaker is already concerning him/herself with the illocutionary message of the sentence before the end of the utterance. The question remains whether the speaker already anticipates upon the final element of the sentence, and that these prosodic differences could indicate a form of sentence planning, or that the speaker just expresses the intention of the sentence without having planned for a specific final element. This question is addressed in the third experiment of this thesis.

The second experiment of this thesis investigated whether the prosodic cues that were identified in the first production experiment are also perceivable by listeners, and can be used to anticipate upon the end of the sentence. With a gating task, participants were presented with partial sentence fragments of increasing size, and had to decide whether the fragment they heard belonged to a sentence ending in *hè*, *hoor*, *man* or *zo*. The results of this experiment show that, while there are cues that could help listeners to distinguish different sentence types in the material preceding the object of the sentence, none of these cues are reliably used to identify the sentence correctly. In fact, participants only started to reliably identify the correct sentence of the *zo*-condition in the third experimental block, i.e. when they heard the entire utterance preceding the final particle. Not even in the third block were they able to reliably distinguish sentences from the conditions *hè*, *hoor* and *man*. Possibly the cues distinguishing these conditions are too subtle to consciously perceive the distinction. Another possibility could be that the experimental paradigm allowed participants to deviate from the intended final particle (e.g. the participants imagined a reminding *hè* final particle after a *hoor* sentence, while the experiment actually only targets the agreement-seeking *hè*). Both explanations suggest an important role of the final particle in the identification of sentence type, a finding that is congruent with earlier research (van Heuven & Haan, 2002; Iwata & Kobayashi, 2012).

Knowing now a little bit more about the prosodic properties of the sentence final particles, the third experiment directly addressed the question of the planning of *hè* and *hoor*. Using a picture-interference-naming task, the sentence final particles *hè* and *hoor* were compared to the deictic marker *zo* and the focus particle *wel*. While the final particles *hè* and *hoor* both alter the complete sentence at a pragmatic level, the elements *zo* and *wel* are modifying the predicate of the sentence at an information structural level. In a training task, the different final elements were associated with different colored pictures, intending to substantiate a conceptual connection between a particular color and final element. Later, in the naming experiment, color words were presented as primes before each target picture. The expectation was that a color term that matches the color of the following picture would facilitate naming of the picture before speech onset, only if speakers are already concerned with the final particle they are going to use at the end before they start speaking. This was expected to be the case for the target conditions (for the intentional final particles *hè* and *hoor*) but not for the other two conditions, if participants had correctly formed a connection between the particles and their associated colors. If the final particles *hè* and *hoor* are not planned in advance, no effects of the different priming conditions were expected to occur before speech onset. However, due to a large variation between participants, not showing a consistent effect on speech onset latencies for the different prime conditions (related and unrelated color words), our research question cannot be answered reliably. Possibly, the variation in the prime's effect on the speech onset latency of different participants is due to the set-up of the experiment in which a trained artificial connection is manipulated by the prime conditions. Since the nature of the associative prime is based on an association that is learned throughout the experiment, there is no control for the strength of this connection. Conceivably, individual learning factors cause different association strengths across participants, which could have consequences for the effect of the prime used in this experiment. Other methodological factors could also have contributed to the non-significant effects of the experiment. In future research it would be interesting to see whether the paradigm of this experiment could be adjusted, to gain more reliable results that can answer the main question we pursued.

So, although this thesis cannot make any claims at the moment about the question of intentional particle planning in Dutch, it explored a new area of sentence production research, namely the interface between prosody and pragmatic sentence level processing, by implementing a novel sentence production paradigm. I hope to have raised interesting questions about the mechanisms underlying the formulation of intention, concerning its planning and expression through SFPs, and hope that future research will further explore this intriguing area of research.

REFERENCES

Allum, P. H., & Wheeldon, L. R. (2007). Planning scope in spoken sentence production: The role of grammatical units. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(4), 791-810.

Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, *30*(3), 191-238.

Aoshima, S., Phillips, C., & Weinberg, A. (2004). Processing filler-gap dependencies in a head-final language. *Journal of Memory and Language*, *51*(1), 23-54.

Aubergé, V., Grepillat, T., & Rilliard, A. (1997). Can we perceive attitudes before the end of sentences? The gating paradigm for prosodic contours. In: *Proceedings of EuroSpeech 1997,* Rhodes, Greece, 871-877.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software, 67*(1), 1-48.

Beasley, T. M., & Schumacker, R. E. (1995). Multiple regression approach to analyzing contingency tables: Post hoc and planned comparison procedures. *The Journal of Experimental Education*, *64*(1), 79-93.

Bever, T. G. (1970). The cognitive basis for linguistic structures. In: Hayes, R. (Ed.), *Cognition and Language Development* (pp. 279-362). New York, NY: Wiley & Sons, Inc.

Beyssade, C., & Marandin, J-M. (2006). French Intonation and Attitude Attribution. In: Denis P. (Ed.), *Proceedings of the 2004 Texas Linguistics Society Conference* (pp. 1-12). Somerville, MA: Cascadilla Proceedings Project.

Bloem, I., van den Boogaard, S., & La Heij, W. (2004). Semantic facilitation and semantic interference in language production: Further evidence for the conceptual selection model of lexical access. *Journal of Memory and Language*, *51*(2), 307-323.

Bock, K., & Levelt, W. (2002). Language production. *Psycholinguistics: Critical Concepts in Psychology*, *5*, 405-452.

Boersma, P., & Weenink, D. (2016). Praat: doing phonetics by computer [Computer program]. Version 6.0.18, retrieved 23 May 2016 from http://www.praat.org/.

Brown-Schmidt, S., & Konopka, A. E. (2008). Little houses and casas pequeñas: Message formulation and syntactic form in unscripted speech with speakers of English and Spanish. *Cognition*, *109*(2), 274-280.

Brown-Schmidt, S., & Konopka, A. E. (2015). Processes of incremental message planning during conversation. *Psychonomic Bulletin & Review*, *22*(3), 833-843.

Brown-Schmidt, S., & Tanenhaus, M. K. (2006). Watching the eyes when talking about size: An investigation of message formulation and utterance planning. *Journal of Memory and Language*, *54*(4), 592-609.

Chan, B. H-S. (2013). Sentence-Final Particles, Complementizers, Antisymmetry, and the Final-over-Final Constraint. In: Biberauer, T., Sheehan, M. (Eds.), *Theoretical Approaches to Disharmonic Word Order* (pp. 445-468). Oxford, England: Oxford University Press.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological review*, *113*(2), 234-272.

Chierchia, G. & McConnell-Ginet, S. (1990). *Meaning and grammar: an introduction to semantics*. Cambridge, MA: MIT Press.

Christoffels, I. K., Firk, C., & Schiller, N. O. (2007). Bilingual language control: An event-related brain potential study. *Brain research, 1147*, 192-208.

DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature neuroscience*, *8*(8), 1117-1121.

DeLong, K. A., Troyer, M., & Kutas, M. (2014). Pre-processing in sentence comprehension: sensitivity to likely upcoming meaning and structure. *Language and Linguistics Compass*, *8*(12), 631-645.

Englert, C. (2010). Questions and responses in Dutch conversations. *Journal of Pragmatics*, *42*(10), 2666-2684.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of memory and language*, *25*(3), 348-368.

Ferreira, V. S., & Dell, G. S. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology, 40*(4), 296-340.

Ferreira, F., & Swets, B. (2002). How incremental is language production? Evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language*, *46*(1), 57-84.

Ford, M. (1983). A method for obtaining measures of local parsing complexity throughout sentences. *Journal of Verbal Learning and Verbal Behavior*, *22*(2), 203-218.

Foucart, A., Ruiz-Tada, E., & Costa, A. (2015). How do you know I was about to say "book"? Anticipation processes affect speech processing and lexical recognition. *Language, Cognition and Neuroscience*, *30*(6), 768-780.

R Core Team (2013). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from http://www.R-project.org/.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, *6*(4), 291-325.

Fujisaki, H., & Hirose, K. (1993). Analysis and perception of intonation expressing para-linguistc information in spoken Japanese. *Lund Linguistics Working Papers*, *41*, 254-257.

Gambi, C., & Pickering, M. J. (2016). Predicting and imagining language. *Language, Cognition and Neuroscience*, *31*(1), 60-72.

Garrett, M. F. (1980). Levels of processing in sentence production. In B. Butterworth (Ed.), Language production (Vol. 1, pp. 177-220). London, England: Academic Press.

Griffin, Z. M. (2001). Gaze durations during speech reflect word selection and phonological encoding. *Cognition, 82*, B1-B14.

Grosjean, F. (1980). Spoken word recognition processes and the gating paradigm. *Perception & Psychophysics*, *28*(4), 267-283.

Grosjean, F. (1996). Gating. *Language and Cognitive Processes*, *11*(6), 597-604.

Güneş, G. (2015). *Deriving prosodic structures*. Utrecht, the Netherlands: LOT.

Gupta, A. F. (2006). Epistemic modalities and the discourse particles of Singapore. In: K, Fisher, (Ed.), *Approaches to discourse particles* (pp. 244-263). Amsterdam, the Netherlands: Elsevier.

Haegeman, L. (2014). West Flemish verb-based discourse markers and the articulation of the speech act layer. *Studia Linguistica*, *68*(1), 116-139.

Hill, V. (2013). *Vocatives: How syntax meets with pragmatics*. Leiden, the Netherlands: Brill.

Hsieh, F. F., & Sybesma, R. (2011). On the Linearization of Chinese Sentence-final Particles. *Korea Journal of Chinese Language and Literature, 1(1), 53-90*.

Huang, X. Y. K. (2007). Initialness of sentence-final particles in Mandarin Chinese. In: *Proceedings of the 21st Pacific Asia Conference on Language, Information and Computation* (pp. 182-191).

IBM Corp. (2012). IBM SPSS Statistics for Windows, Version 21.0. Armonk, NY: IBM Corp.

Iwata, K., & Kobayashi, T. (2012). Expressing speaker's intentions through sentence-final intonations for Japanese conversational speech synthesis. In: *13th Annual Conference of the International Speech Communication Association 2012*, *INTERSPEECH 2012* (Vol. 1, pp. 442-446). Baixas, France: International Speech Communications Association.

Kamide, Y. (2008). Anticipatory processes in sentence processing. *Language and Linguistics Compass*, *2*(4), 647-670.

Kamide, Y., Altmann, G. T., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, *49*(1), 133-156.

Kayne, R. (1994). *The antisymmetry of syntax*. Cambridge, MA: MIT Press.

Kempen, G., & Hoenkamp, E. (1987). An incremental procedural grammar for sentence formulation. *Cognitive science*, *11*(2), 201-258.

Kirsner, R. S. (2003). On the Interaction of the Dutch Pragmatic Particles hoor and hè with the Imperative and Infinitivus Pro Imperativo. *LOT Occasional Series*, *1*, 59-96.

Kirsner, R. S., & Deen, J. Y. (1990). Het mes snijdt aan twee kanten. On the semantics and pragmatics of the Dutch sentence-final particle hoor. *The Low Countries: Multidisciplinary Studies*, 1-11.

Kirsner, R. S., & van Heuven, V. J. (1996). Boundary tones and the semantics of the Dutch final particles hè, hoor, zeg and joh. *Linguistics in the Netherlands*, *13*(1), 133-146.

Kirsner, R. S., van Heuven, V. J., & van Bezooijen, R. (1994). Interaction of particle and prosody in the interpretation of factual Dutch sentences. *Linguistics in the Netherlands*, *11*(1), 107-118.

Kratzer, A. (1996). Severing the external argument from its verb. In: Rooryck, J., Zaring, L. (Eds.), *Phrase structure and the lexicon* (Vol. 33, pp. 109-137). Springer Netherlands.

Kutas, M., DeLong, K. A., & Smith, N. J. (2011). A look around at what lies ahead: Prediction and predictability in language processing. In: Bar, M. (Ed.), *Predictions in the brain: Using our past to generate a future* (pp. 190-207). Oxford, England: Oxford University Press.

Kwok, H. (1984). *Sentence Particles in Cantonese.* Hong Kong: Centre of Asian Studies, The University of Hong Kong.

Law, A. (2002). Cantonese sentence-final particles and the CP domain. *UCL Working Papers in Linguistics*, *14*, 375-398.

Law, S-P. (1991). *The syntax and phonology of Cantonese sentence-final particles* (Unpublished doctoral dissertation). Boston University, Boston, MA.

Levelt, W. J., & Meyer, A. S. (2000). Word for word: Multiple lexical access in speech production. *European Journal of Cognitive Psychology*, *12*(4), 433-452.

Li, C., N., & Thompson, S., A. (1981). *Mandarin Chinese: a Functional Reference Grammar.* Berkeley, CA: University of California Press.

Marantz, A. (1984). *On the nature of grammatical relations.* (Linguistic Inquiry Monographs 10). Cambridge, MA: MIT Press.

Mazeland, H., & Plug, L. (2010). Doing confirmation with ja/nee hoor Sequential and prosodic characteristics*. Prosody in Interaction, 23,* 161-188.

Miyamoto E., & Takahashi S. (2002). The processing of wh-phrases and interrogative complementizers in Japanese. In: Akatsuka N., Strauss S., (Eds.), *Japanese/Korean Linguistics* (Vol. 10, pp. 62-75). Stanford, CA :CSLI Publications.

Momma, S., Slevc, L. R., & Phillips, C. (2015). The timing of verb planning in active and passive sentence production. In *Poster presented at the 28th annual CUNY Conference on Human Sentence Processing, Los Angeles, CA, March 19–21*.

Momma, S., Slevc, L. R., & Phillips, C. (2016). The timing of verb selection in Japanese sentence production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(5), 813-824.

Munaro, N., & Poletto, C. (2002). Ways of clausal typing. *Rivista di Grammatica Generativa, 27*, 87-105.

Munaro, N., & Poletto, C. (2003). Sentential particles and clausal typing in the Veneto dialects. *Working Papers in Linguistics*, *13*, 127-154.

Östman, J-O. (1991). On the language-internal interaction of prosody and pragmatic particles. In: J. Blommaert & J. Verschueren (Eds.), *The Pragmatics of International and Intercultural Communication: Selected papers from the International Pragmatics Conference, Antwerp, August 1987. Volume 3: The Pragmatics of International and Intercultural Communication* (Vol. 6, No.3, pp. 203-222). Amsterdam, the Netherlands: John Benjamins Publishing.

Paul, H. (1880). Prinzipien der Sprachgeschichte [Principles of the history of language]. Leipzig, Germany: Niemeyer.

Phillips, C. (2003). Linear order and constituency. *Linguistic inquiry*, *34*(1), 37-90.

Pickering, M. J. (1994). Processing local and unbounded dependencies: A unified account. *Journal of Psycholinguistic Research*, *23*(4), 323-352.

Pickering, M. J., & Garrod, S. (2013). How tightly are production and comprehension interwoven?. *Frontiers in psychology*, 4: 238.

Rabagliati, H., Gambi, C., & Pickering, M. J. (2016). Learning to predict or predicting to learn? *Language, Cognition and Neuroscience, 31*(1), 94-105.

Rizzi, L. (1997). The fine structure of the left periphery. In: Haegeman, L. (Ed.), *Elements of grammar: A handbook of generative syntax* (pp. 281–327). Dordrecht, the Netherlands: Kluwer.

Schelfhout, C. R. M., Coppen, P. A. J. M., Oostdijk, N. H. J., & Slik, F. V. D. (2005). Interjections in Dutch: A corpus-based approach. *Neerlandistiek.NL,* 05.05.

Schötz, S. (2002). Phonetic Variation in Speaker Recognition & Text-to-Speech Synthesis. In *Speech Technology, GSLT.* Lund, Sweden: Lund University.

Schriefers, H., Teruel, E., & Meinshausen, R. M. (1998). Producing simple sentences: Results from picture–word interference experiments. *Journal of Memory and Language*, *39*(4), 609-632.

Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language (Vol. 626)*. Cambridge, England: Cambridge University Press.

Searle, J. R. (1985). *Expression and meaning: Studies in the theory of speech acts.* Cambridge, England: Cambridge University Press.

Simpson, A., & Wu, Z. (2002). IP-raising, tone sandhi and the creation of S-final particles: evidence for cyclic spell-out. *Journal of East Asian Linguistics*, *11*(1), 67-99.

Smith, M., & Wheeldon, L. (1999). High level processing scope in spoken sentence production. *Cognition*, *73*(3), 205-246.

Speas, P., & Tenny, C. (2003). Configurational properties of point of view roles. *Asymmetry in grammar*, *1*, 315-345.

Stowe, L. A. (1986). Parsing WH-constructions: Evidence for on-line gap location. *Language and Cognitive Processes*, *1*(3), 227-245.

Sugito, M. (2001). Shuujoshi 'ne' no imi, kinoo to intoneeshon [Meanings, functions and intonation of sentence-final particles 'ne'], In: Spoken Language Working Group (Ed.), *Speech and Grammar III* (pp. 3-16). Tokyo, Japan: Kurosio Publishers [in Japanese].

Sybesma, R.P.E., & Li, B. (2007). The dissection and structural mapping of Cantonese sentence final particles. *Lingua, an International Review of General Linguistics, 117*, 1739-1783.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268* (5217), 1632-1634.

Tsuchihashi, M. (1983). The speech act continuum: An investigation of Japanese sentence final particles. *Journal of Pragmatics*, *7*(4), 361-387.

Tulling, M. A. (2015). Divide and Conquer: Split CP Hypothesis in Rhetorical Questions, the Case of SFPs in RQs. *Student Undergraduate Research E-journal!*, *1*(1).

Van Berkum, J. J., Brown, C. M., Zwitserlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating upcoming words in discourse: evidence from ERPs and reading times. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(3), 443-467.

van Heuven, V. J., & Haan, J. (2000). When and how do we hear whether a Dutch speech utterance is a statement or a declarative question. *Linguistics in the Netherlands*, *17*(1), 93-104.

van Heuven, V.J., & Haan, J. (2002). Temporal Distribution of Interrogativity Markers in Dutch: A perceptual study. In: Gussenhoven, C. & Warner, N. (Eds.), *Papers in Laboratory Phonology 7* (pp. 119-143). Dordrecht, the Netherlands: Kluwer.

van der Wouden, T. (2006). Partikels: woordjes die het Nederlands markeren. Van der Sijs, N., Stroop, J. & Weerman, F. (ed.) *Wat iedereen van het NEDERLANDS moet weten en waarom*. Amsterdam, the Netherlands: Uitgeverij Bert Bakker, 119-129.

van der Wouden, T., & Foolen, A. (2011). Pragmatische partikels in de rechterperiferie. *Nederlandse Taalkunde*, *16*(3), 307-322.

Wakefield, J. C. (2010). *The English equivalents of Cantonese sentence-final particles: A contrastive analysis* (Unpublished doctoral dissertation). The Hong Kong Polytechnic University, Hong Kong, China.

Wakefield, J. C. (2012). A floating tone discourse morpheme: The English equivalent of Cantonese lo1. *Lingua*, *122*(14), 1739-1762.

Wicha, N. Y., Moreno, E. M., & Kutas, M. (2004). Anticipating words and their gender: An event-related brain potential study of semantic integration, gender expectancy, and gender agreement in Spanish sentence reading. *Journal of Cognitive Neuroscience*, *16*(7), 1272-1288.

Wheeldon, L. (2013). Producing spoken sentences: the scope of incremental planning. In: Perrier, P. & Lang, P. (Eds.), *Cognitive and Physical Models of Speech Production, Speech Perception, and Production-Perception Integration* (pp. 93-114). Berlin, Germany: Springer-Verlag GmbH.

Wheeldon, L. R., & Lahiri, A. (2002). The minimal unit of phonological encoding: prosodic or lexical word. *Cognition*, *85*(2), B31-B41.

Wundt, W. (1900). *Völkerpsychologie: Die Sprache* [Folk psycholoy: Language] (Vols. 1-2). Leipzig, Germany: Kröner-Engelmann.

Zanuttini, R., & Portner, P. (2003). Exclamative clauses: At the syntax-semantics interface. *Language* 79(1), 39-81.

Zhang, L. (2014). Segmentless sentence-final particles in Cantonese: An experimental study. *Studies in Chinese Linguistics*, *35*(2), 47-60.

Zhao L-M., & Yang Y-F. (2016) Lexical Planning in Sentence Production Is Highly Incremental: Evidence from ERPs. PLoS ONE 11(1): e0146359.

Zwitserlood, P. (1989). The locus of the effects of sentential-semantic context in spoken-word processing. *Cognition*, *32*(1), 25-64.

APPENDIX

A. Micro-variation of *hè* and *hoor*

I. Variation within hè

A. Agreement-seeking

The particle *hè* can be used to ask for confirmation or agreement from the hearer. While it is supposed to be a confirmation-seeking device, there is a clear preference for a convergent answer, agreeing with the propositional content provided by the speaker.

(1).	Het	is	lekker	weer,	**hè**?
	it	is	nice	weather	SFP
	'The weather is great, right?'				

B. Confirmation

Almost contrary to what has been previously stated, *hè* can also be used to display agreement with the proposition of the other speaker and to create common ground.

Conversation between two friends:

(2).	A:	Zij	heeft	hele	mooie	schoenen!
		she	has	really	beautiful	shoes
		'Her shoes are really pretty!'				
	B:	Echt,	**hè**?			
		for.real	SFP			
		'Yeah, right?'				
	A:	Ja!				
		'Yes!'				

C. Reminding

The 'reminding' function of *hè* allows the speaker to manipulate the 'common ground' property of *hè*. In other words: the speaker reminds the listener of *p*, by presenting *p* as if it is part of the common ground.

(3).	A:	Ik	kan	mijn	sleutels	nergens	vinden!	
		I	can	my	keys	nowhere	find.INF	
		'I can't find my car keys anywhere!'						
	B:	Je	hebt	ze	aan	Marie	gegeven	**hè**?
		you	have	them	to	Marie	given	SFP
		'You gave them to Mary, remember?'						

D. Urging

Related to the 'reminding' function, a proposition with *hè* can also urge someone to do something. When you run out the door with a light t-shirt, in the middle of the winter, and your mother shouts something like (4) for example, she wants you to put on a coat.

(4).	Het	is	koud	**hè**!
	it	is	cold	SFP
	'It is cold!'			

E. Downplaying

Downplaying is almost like a reminder. The difference is that with downplaying the speaker implies that *p* is the case, and not something else. This is illustrated in (5), in which the speaker makes a comment after he sees his friend packing a huge suitcase for a weekend trip.

(5).	Je	gaat	maar	twee	dagen	weg,	**hè**?
	you	go	only	two	days	away	SFP

'You are only gone for two days, not all week!'

<u>II. Variation within hoor</u>

A. Correction

The most canonical application of *hoor* is the one used in corrective assertions, such as (6). Here *hoor* functions as an epistemic particle, as it expresses the speaker's assessment of the truth value of the expressed proposition. The speaker of (11) not only believes that the proposition *Chinees is een toontaal* 'Chinese is a tone language' is true, but it also implies that the speaker believes that the listener assumes something else (namely, that Chinese is not a tone language). This assumption of the speaker can be based on something the hearer has previously said, or by some non-verbal or contextual behavior that gave the speaker the idea that the hearer does not know that Chinese is a tone language (e.g. he is searching for a book about Chinese in front of a bookcase labeled 'non-tone languages'). This correction can range from gentile to hostile, depending on the context of the utterance and its prosody.

(6).	Chinees	is	een	toontaal,	**hoor**.
	Chinese	is	a	tone.language	SFP

'Chinese is a tone language. ' (Interpretation: you are wrong thinking differently)

B. Warning

It is very intuitive that something that can be used as a correction, is also used to correct behavior, or to prevent certain things from happening. Take for example the utterance in (7), which is uttered by a dog owner, when a child tries to pet their dog.

(7).	Hij	kan	bijten,	**hoor**!
	he	can	bite	SFP

'He can bite!'

Again, just like in the 'correction' example, the attitude and tone with which this utterance is spoken influences the integration of the sentence. It could, for instance, be a gentle warning from an old lady, but it could also be interpreted as a threat from a grumpy landlord. In this context the sentence with *hoor* could be paraphrased as 'don't think he won't bite, cause he will'.

C. Reassurance

A bit paradoxically (considering that *hoor* can be used as a threat), the 'you are wrong' interpretation of *hoor* can also be used to reassure someone. Similar to the 'warning' function, in this type of sentences *hoor* indicates something like 'don't think *p*, because actually *q* is the case'. This is exemplified in (8), which could be uttered by a doctor to its patient.

(8).	Het	komt	helemaal	weer	goed,	**hoor**!
	it	come	completely	again	good	SFP

'It will be completely fine again!' (Interpretation: don't think anything bad (don't worry), because actually it will be just fine).

D. Emphasis
Hoor can also be used to put emphasis on an utterance. When someone says *leuk hoor!* 'nice!' after a great performance, *hoor* does not necessarily correct someone who thinks it was actually not a good performance. In fact, *hoor* is used to put emphasis on the exclamation, pushing away any objections anyone might have.

E. Command
Hoor can also occur with imperatives, but only in combination with the softener *maar*, as in (9):

(9).	Doe	de	deur	maar	dicht,	**hoor**!
	do	the	door	soft	shut	hoor
	'You can close the door!'					

The particle *hoor* in (9) makes the command sound either urging, implying, don't leave the door open!, or it could sound very friendly, e.g. if somebody looks at the door doubting whether he should close it or not, then someone could respond with (9) to them in a friendly manner. When *hoor* is uttered in a friendly imperative, it is not really a command, but more of a way to grant permission[20].

F. Confirmation/Rejection
While *hoor* in the examples above has a clear function of expressing disagreement with the speaker's assessment of the hearer's believes and thoughts, it can also be used in agreements and disagreeing responses. Very commonly used occurrences of *hoor* are: *Ja, hoor.* 'Yes.' and *Nee, hoor* 'No.'. Again, in these sentences, the precise interpretation of the particle depends on the context it occurs in, though there is not always a 'you are wrong' meaning linked to the particle in these cases. This becomes especially clear in the following example (10), which you could easily overhear in a supermarket:

(10). *A customer pays for his groceries at the cash register.*

Cashier:	Wilt	u		het	bonnetje?
	want	you		the	receipt
	'Would you like the receipt?				
Customer:	Nee,	**hoor**.	Bedankt.		
	no	SFP	thanks		
	'No, thanks.'				

It could be the case that in confirmations/rejection scenarios, the particle *hoor* is lexicalized as some sort of turn-taking discourse marker (Mazeland & Plug, 2010), and may be different from the *hoor* in SFPs. Especially since it seems that *hoor* sometimes forms one phonological word with the yes/no particle, and that the prosodic properties of *hoor* in *ja/nee hoor* (discussed elaborately by Mazeland & Plug, 2010) differ from the prosodic properties of *hoor* in complete utterances (in which *hoor* bears a F0 rising pitch).

[20] One could for example tell a child who is looking at the cookie jar: *Neem maar een koekje, hoor!* meaning 'Go ahead and take a cookie!'. Possibly, this pragmatic usage (i.e. granting permission) of *hoor* is due to the fact that *hoor* is rejecting the possible presupposition of the listener. In this case the speaker assumes that the child presupposes that he/she cannot have a cookie, which the speaker refutes with *hoor*.

B. Base Sentences used for recording, production Experiment 1

ID	Sentence	Translation
1.	*De bakker vult een plaat.*	'The baker is filling up a plate.'
2.	*De duiker vindt een schat.*	'The diver finds a treasure.'
3.	*De hacker maakt een plan.*	'The hacker is thinking of a plan.'
4.	*De jager raakt een zwijn.*	'The hunter shoots a wild pig.'
5.	*De jongen koopt een bloem.*	'The boy is buying a flower.'
6.	*De koning leest een brief.*	'The king is reading a letter.'
7.	*De lasser maakt een blok.*	'The welder is making a block.'
8.	*De leraar ruikt een drol.*	'The teacher smells a turd.'
9.	*De monnik weegt een kraal.*	'The monk is weighing a beat.'
10.	*De ridder wast een draak.*	'The knight is washing a dragon.'
11.	*De ruiter hoort een slang.*	'The rider hears a snake.'
12.	*De trainer gooit een fluit.*	'The trainer is throwing a whistle.'
13.	*De viking zoekt een schip.*	'The viking is looking for a ship.'
14.	*De visser vangt een schoen.*	'The fisherman is catching a shoe.'
15.	*De zanger wint een prijs.*	'The singer is winning a prize.'
16.	*De ziener helpt een vrouw.*	'The prophet is helping a woman.'

C. Corresponding *t*-values from the mixed linear model of Experiment 1.

1. *Reported* t-*values mixed linear regression of the f0*-values *per pitch measure point (p1-11) and condition (*hè, hoor, man, zo*).*

p1	**hè**	**hoor**	**man**	**zo**
hè	X	1.11	-2.14*	-4.71*
hoor		X	-3.21*	-5.65*
man			X	-2.75*
p2	**hè**	**hoor**	**man**	**zo**
hè	X	0.76	-3.43*	-4.01*
hoor		X	-4.11*	-4.63*
man			X	-0.81
p3	**hè**	**hoor**	**man**	**zo**
hè	X	1.28	0.13	-2.48*
hoor		X	-1.17	-3.64*
man			X	2.65*
p4	**hè**	**hoor**	**man**	**zo**
hè	X	1.29	-2.78*	-6.75*
hoor		X	-4.02*	-7.82*
man			X	-4.22*
p5	**hè**	**hoor**	**man**	**zo**
hè	X	3.60*	-0.76	-6.42*
hoor		X	-4.38*	-9.69*
man			X	-5.81*
p6	**hè**	**hoor**	**man**	**zo**
hè	X	3.30*	0.73	-1.30
hoor		X	-2.62*	-4.40*
man			X	-2.01*
p7	**hè**	**hoor**	**man**	**zo**
hè	X	1.26	-3.26*	2.51*
hoor		X	-4.46*	1.26
man			X	5.66*
p8	**hè**	**hoor**	**man**	**zo**
hè	X	1.30	-1.47	-17.26*
hoor		X	-2.75*	-18.11*
man			X	-16.17*
p9	**hè**	**hoor**	**man**	**zo**
hè	X	-2.11*	-3.71*	11.05*
hoor		X	-1.48	12.8*
man			X	14.76*
p10	**hè**	**hoor**	**man**	**zo**
hè	X	-1.68	-4.76*	10.03*
hoor		X	-2.94*	11.40*
man			X	14.72*
p11	**hè**	**hoor**	**man**	**zo**
hè	X	-4.53*	-13.42*	-29.49*
hoor		X	-8.51*	-24.53*
man			X	-17.22*

* = significant at .05

2. *Reported* t-*values mixed linear regression of the constituent (subject, verb, object and particle) duration data and condition (*hè, hoor, man, zo*).*

Subject	**hè**	**hoor**	**man**	**zo**
hè	X	0.29	1.68	6.83*
hoor		X	1.35	6.40*
man			X	5.35*
Verb	**hè**	**hoor**	**man**	**zo**
hè	X	-0.70	-3.23*	2.34*
hoor		X	-2.44*	2.95*
man			X	5.45*
Object	**hè**	**hoor**	**man**	**zo**
hè	X	2.65*	2.26*	-7.15*
hoor		X	0,48	-9.50*
man			X	-9.41*
Particle	**hè**	**hoor**	**man**	**zo**
hè	X	13.35*	12.20*	25.55*
hoor		X	-1.63	-12.29*
man			X	14.37*

* = significant at .05

D. Individual averaged stylized pitch contours, production experiment 1. (N=10)

Legend: hè hoor man zo

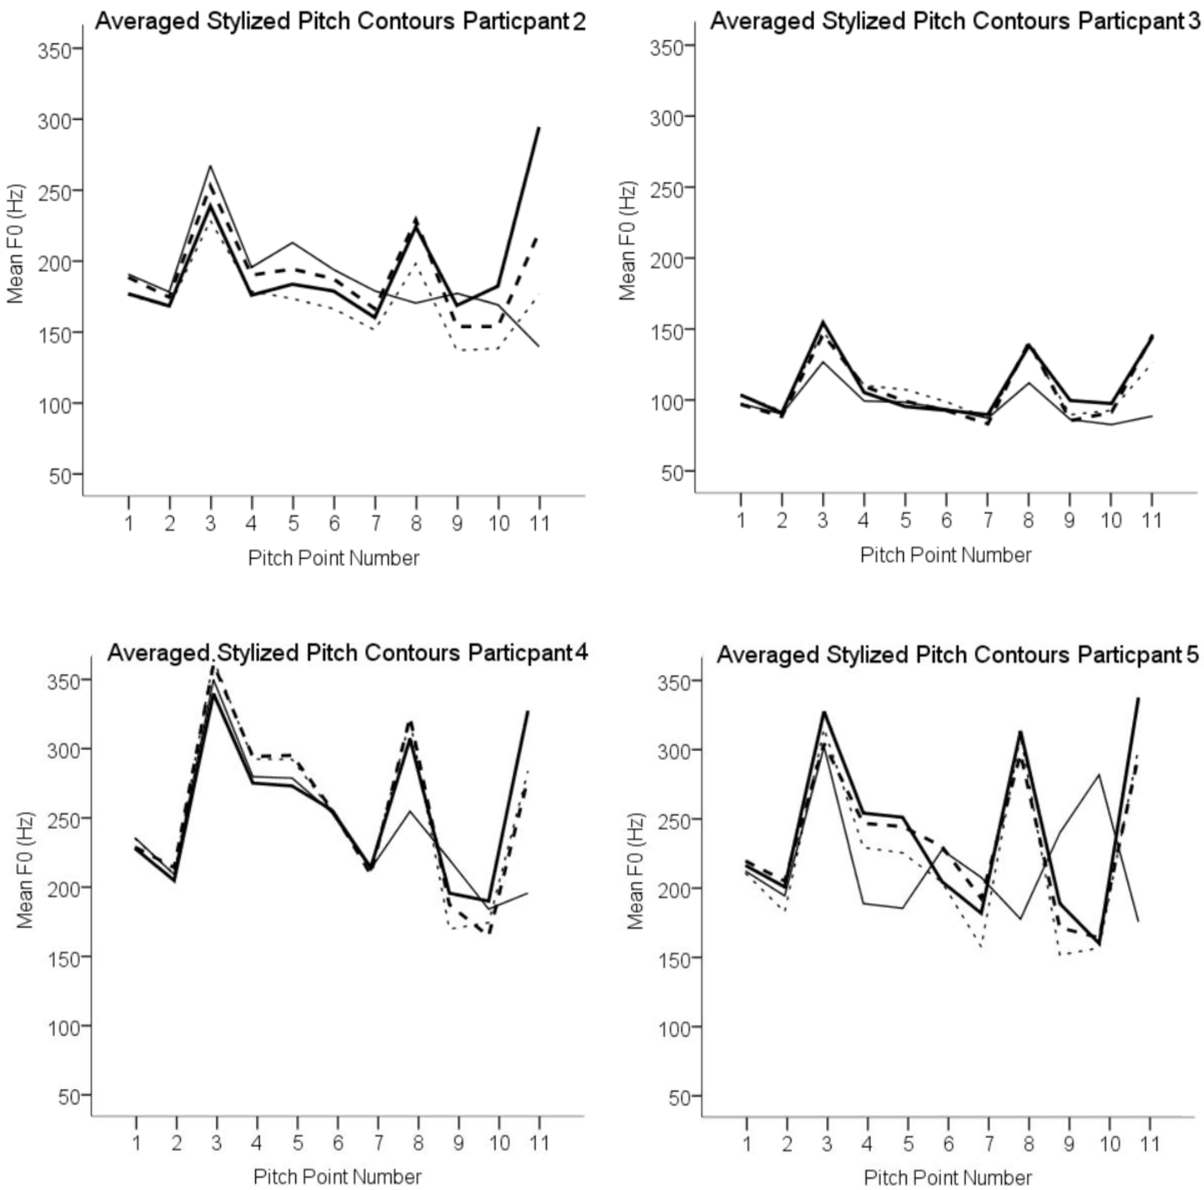

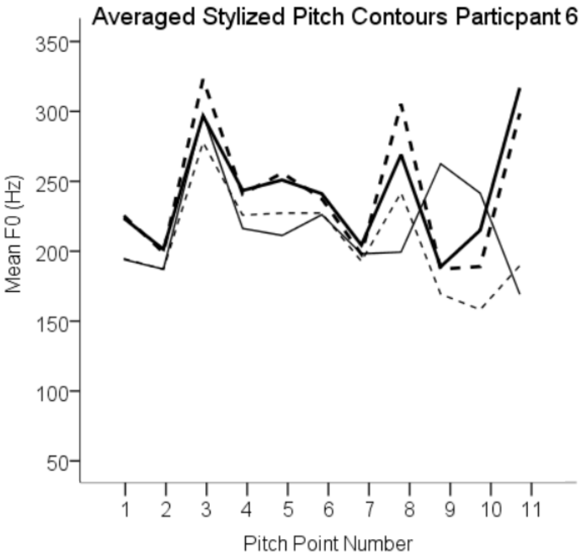
Averaged Stylized Pitch Contours Particpant 6
Mean F0 (Hz)
350
300
250
200
150
100
50
1 2 3 4 5 6 7 8 9 10 11
Pitch Point Number

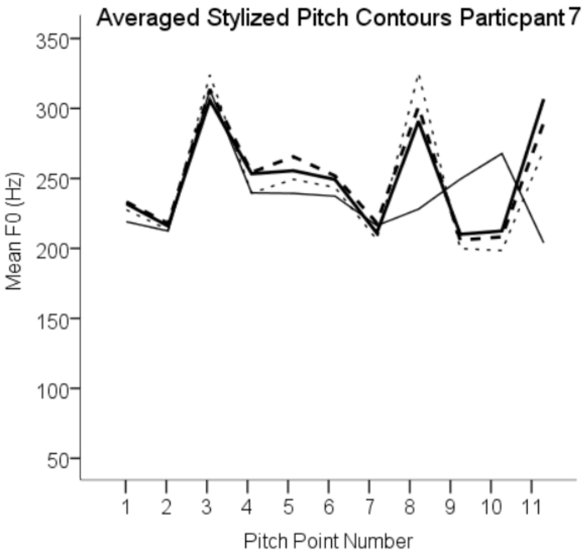
Averaged Stylized Pitch Contours Particpant 7
Mean F0 (Hz)
350
300
250
200
150
100
50
1 2 3 4 5 6 7 8 9 10 11
Pitch Point Number

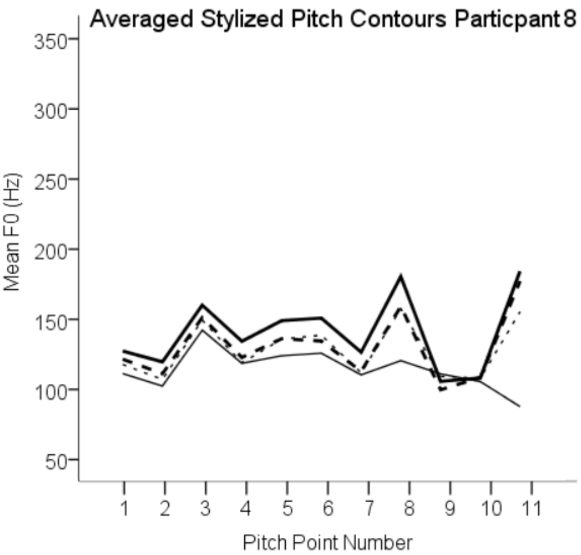
Averaged Stylized Pitch Contours Particpant 8
Mean F0 (Hz)
350
300
250
200
150
100
50
1 2 3 4 5 6 7 8 9 10 11
Pitch Point Number

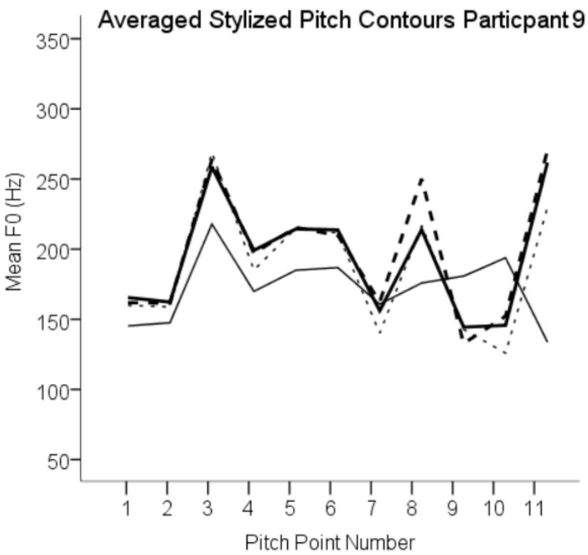
Averaged Stylized Pitch Contours Particpant 9
Mean F0 (Hz)
350
300
250
200
150
100
50
1 2 3 4 5 6 7 8 9 10 11
Pitch Point Number

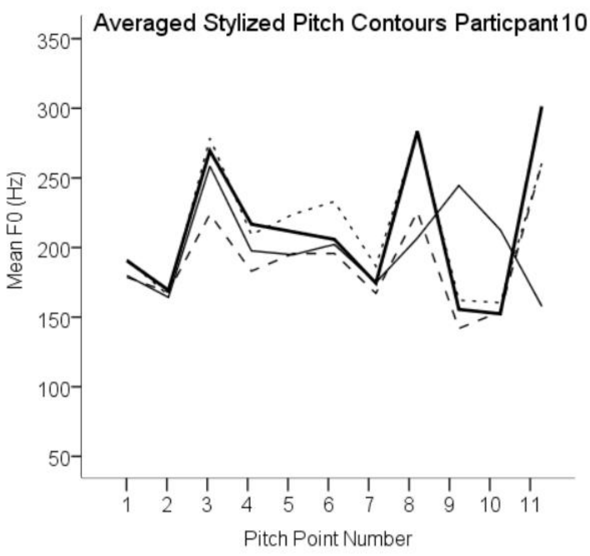
Averaged Stylized Pitch Contours Particpant 10
Mean F0 (Hz)
350
300
250
200
150
100
50
1 2 3 4 5 6 7 8 9 10 11
Pitch Point Number

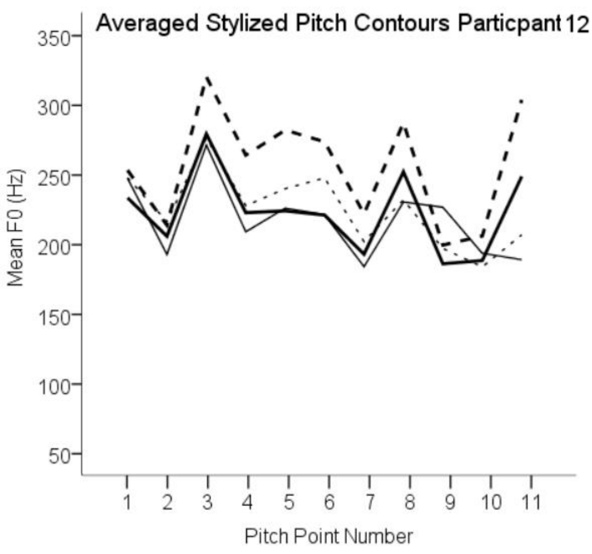
Averaged Stylized Pitch Contours Particpant 12
Mean F0 (Hz)
350
300
250
200
150
100
50
1 2 3 4 5 6 7 8 9 10 11
Pitch Point Number

E. Example contexts for the four different conditions used in Experiment 2

Original (Dutch): Context for ***hoor***:

Janneke kijkt televisie met haar broer. Haar broer kijkt heel intensief naar de televisie, en Janneke krijgt het idee dat haar broer denkt dat hij naar een film aan het kijken is. Dat is echter niet het geval en Janneke zegt:

"Dit is gewoon reclame hoor!"

Translation:

Janneke is watching television with her brother. Her brother is watching the television screen so intensively, so that Janneke starts thinking that her brother believes he is watching a movie. However, this is not the case and Janneke says:

"This is just a commercial *hoor*!"

Original (Dutch): Context for ***man***:

Janneke zit in de woonkamer. Haar broer komt binnen, net terug van zijn werk. Haar broer vertelt het volgende:

"Ik heb een nieuwe collega ontmoet vandaag. Dat is een leuk meisje man!"

Translation:

Janneke is sitting in the living room. Her brother, who just got back from work, walks in. Janneke's brother tells her the following:
"I have met a new colleague today. That is nice girl *man*!"

Original (Dutch): Context for ***hè*:**

Janneke gaat hardlopen met haar broer. Ze rennen langs een boerderij. Opeens hoort Janneke een geluid, dat ze meent te herkennen. Ze zegt tegen haar broer:

"Dat is een koe hè?"

Translation:

Janneke goes out running with her brother. They are running past a farm. Suddenly Janneke hears a sound that she thinks she recognizes. She tells her brother:

"That is a cow *hè*?"

Original (Dutch): Context for ***zo***:

Janneke weet niet hoe ze een hond trimt. Ze gaat naar vriendin Shirley die een professioneel hondentrimster is, en vraagt aan haar hoe het moet. Shirley laat zien hoe je een hond trimt, en zegt daarbij:

"Een hond trim je zo."

Translation:

Janneke does not know how to groom a dog. She goes to her friend Shirley, who is a professional dog groomer, and asks her how to do it. Shirley shows her how to groom a dog, and says:

"You groom a dog like this."

F. Stimulus pictures, Experiment 3. Pictures are displayed in black here, in the actual experiment pictures appeared in the following colors: red, blue, green, yellow and grey. Original pictures are 600 x 600 pixels.

Picture 1. De bakker vult een plaat.

Picture 2. De duiker vindt een schat.

Picture 3. De hacker maakt een plan.

Picture 4. De jager raakt een zwijn.

Picture 5. De jongen koopt een bloem

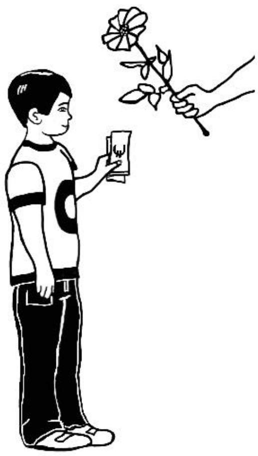

Picture 6. De koning leest een brief.

Picture 7. De lasser maakt een blok.

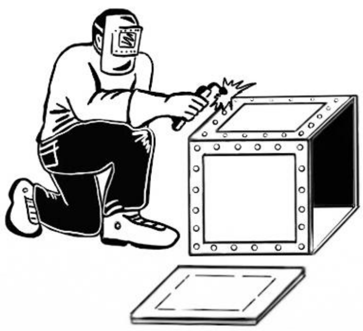

Picture 8. De leraar ruikt een drol.

Picture 9. De monnik weegt een kraal.

Picture 10. De ridder wast een draak.

Picture 11. De ruiter hoort een slang.

Picture 12. De trainer gooit een fluit.

Picture 13. De viking zoekt een schip.

Picture 14. De visser vangt een schoen.

Picture 15. De zanger wint een prijs.

Picture 16. De ziener helpt een vrouw.

G. Counterbalance scheme, Experiment 3

Version	***Condition***	***Block 1***		***Block 2***	
		Type	**Color**	**Type**	**Color**
1.	Target	*hè*	red	*hoor*	blue
	Distractor	*zo*	green	*wel*	yellow
	Filler	decl	grey	decl	grey
2.	Target	*hè*	yellow	*hoor*	green
	Distractor	*zo*	blue	*wel*	red
	Filler	decl	grey	decl	grey
3.	Target	*hoor*	red	*hè*	blue
	Distractor	*wel*	green	*zo*	yellow
	Filler	decl	grey	decl	grey
4.	Target	*hoor*	yellow	*hè*	green
	Distractor	*wel*	blue	*zo*	red
	Filler	decl	grey	decl	grey
5.	Target	*hoor*	blue	*hè*	red
	Distractor	*zo*	yellow	*wel*	green
	Filler	decl	grey	decl	grey
6.	Target	*hoor*	green	*hè*	yellow
	Distractor	*zo*	red	*wel*	blue
	Filler	decl	grey	decl	grey
7.	Target	*hè*	blue	*hoor*	red
	Distractor	*wel*	yellow	*zo*	green
	Filler	decl	grey	decl	grey
8.	Target	*hè*	green	*hoor*	yellow
	Distractor	*wel*	red	*zo*	blue
	Filler	decl	grey	decl	grey