Bilingual sentence production and code-switching: Neural network simulations

Chara Tsoukala

The educational component of the doctoral training was provided by the International Max Planck Research School (IMPRS) for Language Sciences. The graduate school is a joint initiative between the Max Planck Institute for Psycholinguistics and two partner institutes at Radboud University – the Centre for Language Studies, and the Donders Institute for Brain, Cognition and Behaviour. The IMPRS curriculum, which is funded by the Max Planck Society for the Advancement of Science, ensures that each member receives interdisciplinary training in the language sciences and develops a well-rounded skill set in preparation for fulfilling careers in academia and beyond.

More information can be found at www.mpi.nl/imprs

ISBN: 978-94-92910-27-1

Cover design: Amalia Tsichla

Printed by: Ipskamp printing

Layout: This thesis was typeset with LaTeX 2$_\varepsilon$. It uses the *Clean Thesis* style developed by Ricardo Langner.

This thesis was funded by the Netherlands Organisation for Scientific Research (NWO) Gravitation Grant 024.001.006 to the Language in Interaction Consortium.

Bilingual sentence production and code-switching: Neural network simulations

Proefschrift

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. dr. J.H.J.M. van Krieken,
volgens besluit van het college van decanen
in het openbaar te verdedigen op

woensdag 21 april 2021
om 14.30 uur precies

door

Chara Tsoukala

geboren op 18 april 1986
te Athene (Griekenland)

Promotor
Prof. dr. A.P.J. van den Bosch

Copromotoren
Dr. S.L. Frank
Dr. M. Broersma

Manuscriptcommissie
Prof. dr. J.J.M. Schoonen (voorzitter)
Prof. dr. A.M. Backus (Tilburg University)
Prof. dr. F. Chang (Kobe City University of Foreign Studies, JP)
Prof. dr. J.G. van Hell (Pennsylvania State University, VS)
Dr. M.C. Parafita Couto (Universiteit Leiden)

Bilingual sentence production and code-switching: Neural network simulations

Doctoral Thesis

to obtain the degree of doctor
from Radboud University Nijmegen
on the authority of the Rector Magnificus prof. dr. J.H.J.M. van Krieken,
according to the decision of the Council of Deans to be defended in public on

Wednesday, April 21, 2021
at 14.30 hours

by

Chara Tsoukala

born on April 18, 1986
in Athens (Greece)

Supervisor
Prof. dr. A.P.J. van den Bosch

Co-supervisors
Dr. S.L. Frank
Dr. M. Broersma

Doctoral Thesis Committee
Prof. dr. J.J.M. Schoonen (chair)
Prof. dr. A.M. Backus (Tilburg University)
Prof. dr. F. Chang (Kobe City University of Foreign Studies, JP)
Prof. dr. J.G. van Hell (Pennsylvania State University, U.S.A.)
Dr. M.C. Parafita Couto (Leiden University)

Contents

1 Introduction

I first became aware of the notion of code-switching when I was living in Portugal and working in an English-speaking environment. I spoke English, Greek (my native language) and a little bit of Portuguese on a daily basis, but almost every time I tried to switch from Portuguese to Greek, I would end up uttering English words instead. This simple incident sparked an ongoing curiosity as to how the code-switching process works. It made me realize that code-switching is done by many people who speak more than one language, and not only by balanced bilingual speakers who grew up in bilingual communities; it is an important aspect of human language. Despite that, we still do not fully understand the underlying processes of this phenomenon.

Even though it was the *mistakes* in language selection that drew my initial attention, such as producing English rather than Greek words, I soon became interested in the patterns that bilingual, and multilingual, speakers produce when speaking their languages, especially when they mix them in a single sentence. The overarching goal of this dissertation is to come closer to understanding how code-switching works; I aim to do so by developing a computational cognitive model of code-switching to investigate certain production patterns. Using this novel method to study code-switching, I aim to bridge corpus studies, (psycho)linguistic theories and experimental research, while contributing to further progress in cognitive theories on the (bilingual) sentence production system and, hopefully in the future, Natural Language Processing (NLP) applications.

1.1 Code-switching

Code-switching occurs when a multilingual speaker alternates (switches) between their languages in a single conversation context. Muysken (2000) suggested the following three code-switching categories, which I will be referring to throughout this dissertation; all examples are in Spanish–English[1]:

1. Insertional switching
   Insertions of single words or fixed expressions:
   - lexical (e.g., noun): e.g., "*He leído los* terms" (I have read the terms)
   - fixed expressions/ interjections/ idiomatic expressions: e.g., "*Oh my god*, estamos sin palabras" (we are speechless)
2. Alternational switching
   Switch from one language to the other for multi-word sequences, either between or within sentences:
   - Inter-sentential switching: e.g., "*¿Por qué no está aquí?* What is going on?" (Why is (s)he not here?)
   - Intra-sentential switching: e.g., "I had a hard time finding *mis llaves esta mañana*." (my keys this morning)
3. Congruent lexicalization or dense code-switching
   If the languages share syntactic structures and cognate words (i.e., words between languages that have the same root and have similar meaning, for instance 'intelligent' in English and '*inteligente*' in Spanish), it is possible to use the shared syntax and insert lexical items from either language: e.g., "*Bueno*, in other words, *el* flight *que sale de* Chicago around three o'clock" (Fine, ... the flight which leaves from ...) (Pfaff, 1979)

Code-switching emerges in places where two or more languages come in contact, e.g., Spanish–English among the Spanish-speaking communities in the United States (Deuchar et al., 2014; Poplack, 1980), Arabic–French in Morocco (Bentahila, 1983), and Hindi–English in India (Malhotra, 1980). As a language contact phenomenon, code-switching initially caught the attention of sociolinguists and theoretical linguists (e.g., Bullock and Toribio, 2009a; Lipski, 1978; MacSwan, 2014; Muysken, 2000; Poplack, 1980), but in the past decades it has also been studied by researchers from other domains such as psycholinguistics and Natural Language Processing (NLP) (e.g., Dussias, 2003; Fernandez, Litcofsky, and van Hell, 2019; Guzzardo Tamargo, Valdés Kroff, and Dussias, 2016; Isurin, Winford, and De Bot, 2009; Litcofsky and van Hell, 2017; Toribio, 2001). These studies revealed that code-switches do not

[1] I will use this term to denote both "the mix of Spanish constituents in English" and "the mix of English constituents in Spanish".

occur randomly but follow systematic patterns, which bilingualism researchers had already hypothesized before. For instance, as early as the 1970s, Timm (1975) argued that Spanish-English bilingual speakers do not switch between the subject or object pronoun and the verb (as in "*she* swims" or "ella *nada*").

1.2 Factors affecting code-switching

In a loose manner, code-switching follows grammatical rules (or patterns), and it is additionally affected by the following interconnected factors: i) individual preferences and psychological factors, e.g., one's linguistic identity, ii) social aspects, such as the community's view on code-switching and on the languages involved, or the community's code-switching practices, and iii) contextual factors, such as the domain of the conversation, the languages of the interlocutor and the formality of the conversation.

More specifically, individual and psychological factors include: a) the bilingual speaker's attitudes towards code-switching and towards the individual languages, especially because one of the languages may be a minority language in the society, as is often the case with the native language of heritage speakers (Dewaele and Wei, 2014b); b) the speaker's proficiency in each of the languages; and c) the language they prefer to use for a specific domain (e.g., a Spanish–English speaker in the U.S. may use English at school and Spanish for food; Fishman, Cooper, and Newman, 1971). The social aspect includes society's general attitude towards code-switching (Dewaele and Wei, 2014a), as well as community-based practices. The latter have such a strong influence on code-switching that it has been observed that two communities that code-switch using the same language pair may have opposite code-switching preferences (e.g., Balam, Parafita Couto, and Stadthagen-González, 2020; Blokzijl, Deuchar, and Parafita Couto, 2017).

In this dissertation, I focus mainly on the first factor[2], namely, code-switching patterns caused by the interaction of the two languages with the cognitive system. This has the advantage that one can disentangle internal factors from the effects of exposure to code-switching (which touches upon the second factor as well), and investigate whether certain switching patterns can be explained by the distribution of the languages involved, or whether they are strictly a community-based practice that can only be explained by exposure to those particular patterns of code-switches. Furthermore, I focus on Spanish–English code-switching and cross-linguistic phenom-

[2]The model described in the following chapters also has a conversational setting, which is part of the third factor, but in the simulations the conversational setting is set to either "monolingual" or "bilingual" (i.e., allowed to utilize both languages); I did not otherwise aim to simulate different conversational contexts.

ena because of the large bilingual population that code-switches between these two languages (for instance, in several locations in the U.S. as mentioned above), and because it is a well-documented language pair with freely available corpora, such as the Bangor Miami corpus (Deuchar et al., 2014).

Grammatical constraints on code-switching: Linguistic theories

Different theories have attempted to explain the grammar of code-switching. The most prominent ones are the following, each paired with a Spanish–English example that would be considered *invalid* by the corresponding theory:

- Constraint-based models
  - Models proposed by Poplack (1980) and Sankoff and Mainville (1986): The Equivalence Constraint (EC) (or equivalence of structure constraint) states that code-switching tends to appear at points in which the two languages have identical syntax, so that the resulting code-switch does not violate the syntax of either language (e.g., an invalid combination would be "*linea* blue" or "*azul* line"). The free-morpheme constraint (FMC) states that code-switching cannot occur between a lexical stem and its morpheme (e.g., "*está* sleep*ando*" for "is sleeping").
  - The Functional Head Constraint (FHC) (Belazi, Rubin, and Toribio, 1994) claims that no code-switch can occur between a functional head and its complement (e.g., "*Los niños han* eaten their vegetables", meaning "the kids have eaten ...").
- The Matrix Language Frame (MLF) model (Joshi, 1982; Myers-Scotton, 1993) suggests that a code-switched sentence consists of a dominant Matrix Language (ML) that provides abstract grammatical frames in which the Embedded Language (EL) is inserted in the form of single content words or larger constituents. (e.g., "I would also join *los* slow ones", i.e. "the slow ones" (Callahan, 2002), is disallowed because the Matrix Language is English and only the Spanish function word (determiner) *los* is inserted in the sentence.)

All proposed grammatical models have been challenged with counter-examples from several language pairs. For instance, MacSwan (2005) reviewed counter-arguments against each theory and presented a counter-proposal that the principles governing code-switching are the same as those governing monolingual processing, i.e., there are no constraints other than those of the mixed grammars; this is known as the "constraint-free approach".

In the computational model that I will present in the next chapters, I have not imposed any constraints on code-switched production. This decision was made to step away from assumptions, and to observe instead whether such constraints emerge naturally from the combination of the two languages.

1.3 Methods to study code-switching patterns

Munarriz-Ibarrola, Parafita Couto, and Vanden Wyngaerd (2018) present a collection of methodological suggestions for studying intra-sentential switching. Most of the suggested methods are on code-switching comprehension, using, for instance, EEG and/or eye-tracking while reading code-switched sentences. Sentence production, on the other hand, is more challenging to explore experimentally compared to sentence comprehension. In this dissertation, I will focus on syntactic patterns in sentence production using computational cognitive modeling.

Apart from the computational method I will propose, there are a few empirical methods to study code-switching in production. These include: i) corpus analysis, i.e., analyses of spontaneous dialogues collected in bilingual communities (e.g., Balam, Parafita Couto, and Stadthagen-González, 2020; Deuchar et al., 2014; Poplack, 1980), ii) shadowing, a task in which participants repeat heard stimuli as quickly as possible (e.g. Bultena, Dijkstra, and van Hell, 2015; Lipski, 2019), and iii) confederate priming, where one of the participants is in fact a confederate with a script, who provides primes for the actual participant (e.g. Kootstra, van Hell, and Dijkstra, 2010; Kootstra, van Hell, and Dijkstra, 2012).

Additionally, there are several computational models of code-switching, but very few are cognitive models. In the last decade there has been an increased interest in NLP applications that can detect (e.g., Solorio and Liu, 2008; Solorio, Blair, et al., 2014) or translate (e.g., Johnson et al., 2017; Menacer et al., 2019) code-switched text, recognize code-switched speech (e.g., Adel, Vu, and Schultz, 2013; Shan et al., 2019; Yılmaz, van den Heuvel, and van Leeuwen, 2016) or even detect emotions in code-switched text (Wang et al., 2016). However, these models are not cognitive models, in the sense that they do not aim to explain the code-switching process or to inform or test linguistic theories, but to develop useful applications for multilingual text and speech processing. There are a few language models, though, that have been based on linguistic theories of grammatical constraints on code-switching (see Section

1.2). Li and Fung (2014) incorporated the FHC into their language model to improve code-switched detection in automatic speech recognition, while Pratapa et al. (2018) used the EC to synthetically generate code-switched data in order to improve a code-switched language model[3]. The latter work also showed that randomly generating code-switched data by combining constituents of monolingual sentences decreases the performance of the language model compared to a baseline model that uses only monolingual or real (non-generated) code-switched data; this is to be expected given that code-switching does not occur randomly, but follows patterns. Last, Bhat, Choudhury, and Bali (2016) simulated Hindi–English inter-sentential switching based on the EC and MLF models using grammar-based models. They observed that both models are highly constrained and do not allow certain patterns that are commonly observed in Hindi–English code-switching.

As mentioned above, rather than incorporating hard constraints based on a linguistic theory, I will follow a different modeling approach. The computational cognitive model that I will present does not allow for incorporating linguistic constraints: any constraint has to be learned from the input or to be implicit in the cognitive assumptions.

1.4 Models of code-switching and bilingualism

Computational modeling has helped advance the field of cognitive science in general, and language processing in particular, by i) implementing theoretical models, thereby specifying all aspects of a verbal model and identifying hidden assumptions, and ii) testing hypotheses that are difficult to investigate experimentally. Some of the most prominent computational models in the field of bilingualism are the Bilingual Interactive Activation (BIA; Dijkstra and van Heuven, 1998) and BIA+ (Dijkstra and van Heuven, 2002) models of word recognition, the Multilink model (Dijkstra, Wahl, et al., 2019) of word translation, the Devlex-II model (Li, Zhao, and MacWhinney, 2007) of lexical development that simulates children's word acquisition, the French (1998) model of bilingual processing (next-word prediction), and the Self-Organising Model of Bilingual Processing (SOMBIP; Li and Farkaš, 2002) (see Frank (2021) and Thomas and van Heuven (2005) for overviews of influential computational models of bilingual processing).

However, regarding code-switching, computational cognitive modeling as a means to investigate this phenomenon is still in its infancy. Filippi, Karaminis, and

[3]Performance measured with a standard evaluation metric in language modeling: perplexity

Thomas (2014) developed two connectionist models to simulate an Italian–English word naming task, and more specifically the switch cost, which is larger when switching to the native language (L1) compared to the second language (L2) (Meuter and Allport, 1999). Filippi, Karaminis, and Thomas initially extended the Cohen, Dunbar, and McClelland (1990) model of the Stroop task to allow it to learn mappings between orthography, lexical-semantics, and phonology for a number of Italian (L1) and English (L2) words. This model learned to name words in both languages, and it was faster when producing L1 words. However, it did not produce a larger L2-to-L1 switch cost. The authors proceeded to create a second computational model, based on the Gilbert and Shallice (2002) model of task switching that assumes interference from a carryover of the previous task set into the switch trials. Filippi, Karaminis, and Thomas combined this model with a cross-linguistic morphological acquisition model by Karaminis and Thomas (2010); the combined model replicated the intended asymmetry in switch costs, and showed that the asymmetry depended on the degree of language imbalance, the word class (i.e., cognates and homographs vs. unique words in each language), and the interaction between these factors. Additionally, this model produced a novel prediction that word naming should be slower for words with different orthographic cues compared to words with orthographic cues common to both languages, but that the former induce smaller switch costs because language-specific cues reduce competition effects.

Goldrick, Putnam, and Schwarz (2016) employed the Gradient Symbolic Computation framework (Smolensky, Goldrick, and Mathis, 2014), a grammar-based model using symbolic representations, and used it with probabilistic grammars of two languages to simulate code-switched processing. Using the model, the authors showed that the combination of the grammatical principles of the languages involved, as well as the graded co-activation of the mental representation of the languages, can simulate aspects of English–Tamil intra-sentential switching.

So far I have mentioned different types of models, such as grammar-based and connectionist models. Connectionist models, which view the mind as a number of (neural) networks of interconnected units, are by far the most popular model type in cognitive sciences. They can be either localist (e.g., Dijkstra and van Heuven, 1998; Dijkstra and van Heuven, 2002; Dijkstra, Wahl, et al., 2019) or distributed (e.g., Filippi, Karaminis, and Thomas, 2014; French, 1998; Li, Zhao, and MacWhinney, 2007), which are also known as Artificial Neural Networks (ANNs). In the localist approach, connectivity is set by hand and the model is not trained, i.e., no learning occurs. Distributed models, on the other hand, can learn a task simply by being exposed to training examples and by creating their own representations. For instance, a Recurrent Neural Network (RNN), which is the most commonly used connectionist model in sentence processing, can learn about syntax simply by being exposed to sentences and by predicting the next word in the sentence. This means that the number of assumptions is reduced compared to the localist models, because the model does not work with given syntactic categories such as noun, verb, etc., but it learns instead (during training) that certain words occur in similar contexts. For an explanation of the architecture and an overview of the benefits and usage examples of RNNs for language acquisition and processing, see Frank, Monaghan, and Tsoukala (2019).

In this dissertation I will be working with this architecture. To be able to investigate code-switching (the specific research questions are given in Section 1.6), an important first step is to develop a computational cognitive model of bilingual and code-switched sentence production by extending an existing RNN-based model of sentence production. After reviewing implemented models of sentence production, I chose to work on the Dual-path architecture (Chang, 2002) because, as shown below, the Dual-path model has been employed to explain a wide range of phenomena in several languages.

1.5 Neural network models of sentence production[4]

Surprisingly, there has been little work done on computational models of bilingual sentence production, and on sentence production models in general compared to comprehension models.

A few production models were constructed by running a sentence comprehension model backwards (Calvillo, Brouwer, and Crocker, 2016; Hinaut et al., 2015). However, the most successful and empirically validated sentence production models were specifically designed to simulate production.

The first neural network model of sentence production, the so-called structural priming model (Dell, Chang, and Griffin, 1999; Chang, Dell, Bock, and Griffin, 2000), was developed to simulate syntactic priming: the tendency of speakers to repeat the structure of recently spoken or heard sentences (Bock, 1986). The model assumes a

[4]Part of this Section has been adapted from Section 4.2 of Frank, Stefan L., Padraic Monaghan, and Chara Tsoukala (2019). "Neural network models of language acquisition and processing". In: *Human Language: from Genes and Brains to Behavior*. Ed. by Peter Hagoort. The MIT Press, pp. 277–281.

close link between sentence comprehension and production; comprehension of what has been said or heard so far influences the production of a sentence. The model encodes the intended meaning (or "message") by units that represent role-concept pairs (e.g. "agent-CHILD" or "patient-MAN"), which are given as input to the model until the whole message has been produced. The output-layer units represent words, where the most active unit is regarded as the produced word and is fed back into the network, which thereby receives information about what has been produced so far. This model was able to successfully account for several structural priming phenomena; for instance, if "boys chase dogs" was used as prime (i.e., active rather than passive voice), the model would produce the message "agent-GIRL; action-FEED; patient-CAT" as "girl feeds cat" instead of "cat is fed by girl" (Dell, Chang, and Griffin, 1999). However, it failed to show priming between transitive locatives ("boys chase dogs near car") and prepositional datives ("boys give dog to girl") which is empirically shown by Bock and Loebell (1990). Another limitation of the model was that, because each concept-role pair is represented in a single unit, the agent-concept MAN is different from the patient-concept MAN. Consequently, the model is unable to generalize its ability to produce "the man is chasing a dog" to the ability to produce "the dog is chasing a man". This violates the property of systematicity, which Fodor and Pylyshyn (1988) argued is a fundamental feature of human cognition that neural networks do not possess.

Chang (2002) proposed and compared two neural network models of sentence production, Prod-SRN and Dual-path. Prod-SRN was a simple extension of the structural priming model. It was tested on a more advanced morphology and it was closer to a Simple Recurrent Network (SRN; Elman 1990), which is the most basic type of RNN, but the model still lacked systematicity. Dual-path, which is still the most influential neural network model of sentence production and the one I will base this dissertation on (see Chapter 2 for a detailed description of the architecture), was the first to overcome the limitation of generalization. It does so by creating temporary bindings between a layer for roles and a layer for concepts, so that there is only one unit for MAN, irrespective of its semantic role. These bindings, along with the event semantics (information about tense and aspect, e.g., PRESENT SIMPLE), form the model's semantic path. The model has a second path (hence its name), the syntactic path, which is an SRN that allows the model to learn syntactic categories. This way, the model was not only able to generalize words to new positions, but also to generalize a noun as a verb; this is something that speakers usually do with proper nouns, e.g., "Skype" becomes "skyping". Chang (2002) compared Prod-SRN to the Dual-path model, and the latter was able to generalize 82% of the time whereas Prod-SRN only reached 6%. The models were also tested on unseen adjective-noun pairs and identity construction (e.g., "a cat is a cat"); Dual-path outperformed Prod-SRN in all tests. The model also expanded on Gordon and Dell's (2003) simple model of aphasic production, offering a natural explanation of two different types of aphasia, agrammatism and anomia (see also Dell and Chang, 2014).

Chang, Dell, and Bock (2006) applied the Dual-path model to a wider range of structural priming phenomena. The model displayed similar priming whether the prime had been produced or only comprehended. It was also able to account for long-term priming, as the extent of structural priming was not dependent on the number of fillers between the sentences. Furthermore, Chang, Dell, and Bock (2006) showed that the strong but short-lived tendency to repeat previously heard or said words (the so-called lexical boost) is due to a different mechanism from structural priming. This prediction was confirmed experimentally two years later (Hartsuiker et al., 2008).

To test whether the model could handle a language that is typologically very different from English, Chang (2009) tested Dual-path on Japanese. Despite the different word orders between these two languages, the model was able to exhibit similar levels of grammaticality (93% for English and 95% for Japanese). Furthermore, the model was able to explain differences in production preferences between speakers of these two languages. For instance, in English long phrases are usually placed after short ones (e.g., "The woman sent a book to the man that she met while traveling" is preferred over "The woman sent the man that she met while traveling a book"). This phenomenon is called heavy NP shift (Ross, 1967), and English exhibits a short-before-long bias whereas Japanese does the opposite. The model was able to account for this cross-linguistic difference. Chang (2009) showed that the phenomenon is caused by a difference in the relative importance of the meaning in the positions ("choice points") where the word orders differed. For English, the choice point was right after the verb, whereas in Japanese the choice point is at the beginning of the sentence as verbs tend to occur at the end of the sentence.

The Dual-path model was further able to account for cross-linguistic differences in lexical/conceptual accessibility between English and Japanese (Chang, 2009). English speakers tend to prefer using animate elements early in the sentence (MacDonald, 2013), which can lead to the usage of less common structures such as passives (e.g., "the man was almost hit by a car"). At the same time, they do not have animacy preference in conjunctions: speakers of English find "the man and the car" as acceptable as “the car and the man”. Therefore, it was hypothesized that animacy can influence the functional level but not the positional level from Garrett’s (1988) theory of sentence production. However, if this were the case, animacy would not affect word order in Japanese as this language uses case-markers to indicate roles and repositioning words does not affect the meaning (e.g., in the passive sentence above, “man” would be marked as the subject that receives the action “hit”, regardless of word order). Nevertheless, it has been shown that animacy affects word position in Japanese (Branigan, Pickering, and Tanaka, 2008). Using Dual-path, Chang (2009) noticed that these preferences were related to the frequency of the input; by giving it sentences where animate words were used early in the sentence, the model learned stronger connections between concepts and words for animates than for inanimates.

More recently, the Dual-path model was employed to explain aspects of sentence comprehension (Fitz and Chang, 2019) and language acquisition (Fitz and Chang, 2017; Janciauskas and Chang, 2018; Twomey, Chang, and Ambridge, 2014). Specifically, Fitz and Chang (2019) simulated nine studies on event-related potentials (ERPs), namely the N400 and P600, and suggested that these signals arise as side effects of an error-based learning mechanism that is part of linguistic adaptation and language acquisition. Twomey, Chang, and Ambridge (2014) provided an account as to why children erroneously overgeneralize certain structures even though these errors are not available in their input. Using the Dual-path model, the authors explained that the overgeneralizations are caused by structural biases in language. Fitz and Chang (2017) modeled the acquisition of auxiliary inversion and argued that these structures arise because of how the message is structured. Last, by providing the model first with Korean and then with English sentences, Janciauskas and Chang (2018) used the model to explain the effect of input and age of acquisition in second language learning, and more specifically in the L2 English acquisition by L1 Korean speakers.

Because the validity of the Dual-path model and its ability to explain a wide range of phenomena are well established, I chose to extend this model to handle more than one language simultaneously and to code-switch.

1.6 Research questions and dissertation overview

In **Chapter 2**, which is loosely based on Tsoukala, Frank, van den Bosch, et al. (2019), I will introduce the Bilingual Dual-path model: a Python implementation based on the Dual-path architecture. The goal of this chapter is to investigate whether code-switching can be (at least partially) attributed to internal factors and explained by the distributions of the two languages involved, or whether it is strictly a community-based practice that can only be explained by exposure to code-switches. More specifically, I hypothesize that if the Bilingual Dual-path receives training input in two languages (Spanish and English), without any code-switched input, it will nevertheless be able to produce code-switched sentences by combining patterns from the two languages it has been trained on. The ultimate goal is to simulate balanced Spanish–English bilinguals and their ability to code-switch.

Building upon the previous chapter, in **Chapter 3** (Tsoukala, Broersma, et al., 2021) I will extend the simulations to model late speakers of English who have Spanish as a dominant native language, and late speakers of Spanish who have English as a dominant native language, as well as balanced speakers. The goal of this chapter is to explore how code-switching patterns differ between balanced and non-balanced bilinguals, and to compare the patterns produced by the simulations to two corpora of spontaneous bilingual speech. The hypothesis is that during the early stages of L2 acquisition the Spanish-dominant and English-dominant (non-balanced) models will code-switch more from their non-proficient L2 into their dominant native language rather than the other way around. At the later stages of acquisition, however, when the non-balanced models are also somewhat proficient in their L2, the hypothesis is that these models will be able to produce code-switches from the L1 into their L2, even though the frequency and the patterns of code-switches are expected to be different between the balanced and non-balanced simulation conditions. Additionally, to validate the model output, I will run an exploratory analysis comparing the simulated code-switches to bilingual speech corpora.

In **Chapter 4** (Tsoukala, Frank, van den Bosch, et al., 2020) I will employ the model to investigate a code-switching asymmetry, known as the auxiliary phrase asymmetry; this chapter showcases the advantage of computational modeling and of the ability to manipulate the miniature languages. Additionally, I will be able to show in practice that with this method we can investigate whether a specific pattern is caused by internal factors or due to community practices, as mentioned above. In this chapter, the hypothesis is that the model will produce the asymmetry, even without exposing it to any code-switched input, showing that the asymmetry can derive from the properties of the Spanish and English languages. Using the model, I will test whether the auxiliary phrase asymmetry exists because of the lack of semantic weight of one of the Spanish auxiliary verbs. To test this, I will run a simulation similar to the first one, but this time artificially adding semantic weight to the auxiliary verb in question. The hypothesis is that by doing so, the asymmetry will disappear. Third, I will run a final simulation controlling for the relatively low occurrence frequency of the aforementioned Spanish auxiliary.

In **Chapter 5** I will introduce cognates to the model and test their effect on code-switching, as well as the effect that the percentage of cognates in a language pair has on code-switching. The hypothesis is that sentences that contain a cognate word are more likely to be code-switched than sentences that do not contain a cognate, thus exhibiting the so-called cognate triggering effect. Additionally, I will explore whether there is a different triggering effect in balanced versus non-balanced models, and whether the cognate percentage of a language pair affects the overall probability of code-switching.

In **Chapter 6** (Tsoukala, Frank, and Broersma, 2017) I will employ the model to simulate a language contact phenomenon other than code-switching. In this chapter, I will focus on a cross-linguistic transfer error that has been observed among native Spanish speakers of second-language English.

Last, **Chapter 7** will conclude this dissertation with a summary of the findings, overall conclusions, future steps and final remarks.

2

Simulating Spanish-English code-switching

Multilingual speakers are able to switch from one language to the other ("code-switch") between or within sentences. Because the underlying cognitive mechanisms are not well understood, in this study we use computational cognitive modeling to shed light on the process of code-switching. We employed the Bilingual Dual-path model, a Recurrent Neural Network of bilingual sentence production (Tsoukala, Frank, and Broersma, 2017) and simulated sentence production in simultaneous Spanish-English bilinguals. Our goal was to investigate whether the model would code-switch without being exposed to code-switched training input. The model indeed produced code-switches even without any exposure to such input and the patterns of code-switches are in line with earlier linguistic work (Poplack, 1980). To our knowledge, this is the first computational cognitive model that aims to simulate code-switched sentence production.

This chapter is loosely based on: Tsoukala, Chara, Stefan L. Frank, Antal van den Bosch, Jorge Valdés Kroff, and Mirjam Broersma (2019). "Simulating Spanish-English code-switching: El modelo está generating code-switches". In: *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*. Minneapolis, Minnesota: Association for Computational Linguistics, pp. 20–29

2.1 Introduction

People who speak several languages are able to switch from one language to the other, between or within sentences, a process called code-switching. Code-switching has been studied for decades by theoretical linguists and sociolinguists (e.g., Poplack, 1980; Muysken, 2000) and more recently by psycholinguists (e.g., Bullock and Toribio, 2009c). In the past few years it has started being studied with a computational methodology, and it has garnered attention among the natural language processing (NLP) research community. Several NLP applications have emerged, e.g., to detect code-switches (Solorio and Liu, 2008; Guzmán et al., 2017), or to automatically recognize code-switched speech (Yılmaz, van den Heuvel, and van Leeuwen, 2016; Gonen and Goldberg, 2019). Moreover, there are a small number of cognitive computational models relevant to code-switching: Filippi, Karaminis, and Thomas (2014) developed a model of code-switched word production and Janciauskas and Chang (2018), while simulating age of acquisition effects on native Korean speakers of English, reported that the models that had been exposed to English later produced code-switches, i.e., occasionally used Korean words in their predominantly English production.

The underlying mechanisms of code-switching, however, are still not well understood. Therefore, we suggest using computational cognitive modeling to simulate code-switching behavior in multilinguals with the goal of gaining more insight into the process of code-switching. In this work, we have employed a model of bilingual sentence production (Tsoukala, Frank, and Broersma, 2017) and tested whether it can produce spontaneous code-switches without being trained on code-switched sentences. We wanted to test whether code-switching can be (partially) attributed to internal factors and explained by the distributions of the two languages involved, or whether it is strictly a community-based practice that can only be explained by exposure to code-switches. To test the former, we hypothesized that a model that receives training input in two languages but no code-switched sentences, will nevertheless be able to produce code-switched sentences by combining patterns from the two languages it has been trained on.

To our knowledge, this is the first computational cognitive model that aims to simulate code-switched sentence production.

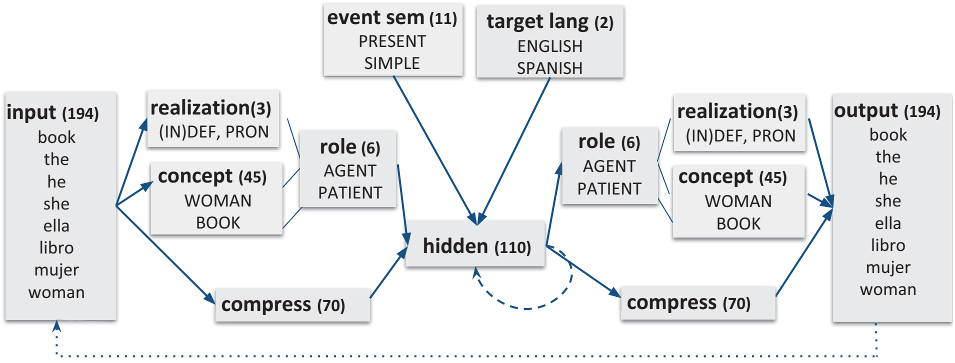

Figure 2.1.: Bilingual Dual-path, the model used in these simulations, is a next-word prediction model that converts messages into sentences. It is based on an SRN (the lower path, via the 'compress' layers) that is augmented with a semantic stream (upper path) that contains information about concepts, thematic roles, event semantics, and the target language.

2.2 Model

To simulate code-switched sentence production, we first needed to simulate bilingual production. For that purpose, we employed the Bilingual Dual-path model (Tsoukala, Frank, and Broersma, 2017) and trained it to simulate simultaneous Spanish-English bilinguals, i.e., speakers who acquired both Spanish and English from infancy.

The Bilingual Dual-path model is a modified version of Dual-path (Chang, 2002). We chose to work with, and extend, the Dual-path model because it is one of the most successful and empirically validated cognitive models of sentence production. It has been used to explain a wide range of phenomena in various languages; for an overview see Section 1.5.

2.2.1 Bilingual Dual-path model

The Bilingual Dual-path model (Figure 2.1) is a Recurrent Neural Network (RNN) based on the Simple Recurrent Network (SRN; Elman, 1990) architecture. It learns to convert a message into a sentence by predicting the sentence word by word. Dual-path (Chang, 2002) got its name because of its two pathways that influence the production of each word: i) the meaning, or semantic, system that learns to map words onto concepts (and their realization, see below and Section "Message"), thematic roles, event semantics and, in the bilingual version, the intended language ("target language"), and ii) the sequencing, or syntactic, system that is an SRN that learns to abstract syntactic patterns. Both paths influence the next word prediction (the "output" layer).

To express a new message (see Section 2.2.2 for examples of messages), the following items are fixed and influence the production of the first word: the to-be-expressed semantic roles have fixed connections with their concepts and realizations, and the relevant "event semantics" and "target language" units are activated. Additionally, the hidden layer's context units are reset to a default value (0.5 in our simulations).

The output word is determined as the word with the highest activation in the output layer. Once an output word has been produced, it is fed back as input (to the "input" layer). During the training phase, the target word is given as input instead of the (potentially different) output word.

The sequencing system is a regular SRN that has one recurrent hidden layer (of 110 units in our simulations) and two 70-unit "compress" layers that are placed between the input word and the hidden layer, and between the hidden layer and the output word.

The meaning system learns to map the input word onto a concept and, whenever relevant, the realization of that concept (PRON for pronoun, INDEF for an indefinite article and DEF for definite articles; see Section "Message" for concrete examples).

A difference between this architecture and other RNNs is that whenever a new message needs to be expressed, the network receives fixed connections between concepts and roles; this allows their separation (instead of having a single unit for, e.g., 'AGENT-WOMAN') and, in turn, enables the model to generalize and to produce words in novel roles. For instance, if the concept 'WOMAN' has only been seen as an AGENT in the training set, it can still be correctly expressed in novel roles (PATIENT, RECIPIENT) during the test phase (Chang, 2002).

All layers use the *tanh* activation function, except the output and predicted role layers that use *softmax*[1]. The model is built in Python and can be found at `https://github.com/xtsoukala/CMCL19`[2].

[1] The original Dual-path model has an additional "comprehended role copy" layer that helps keep track of the roles that have been processed. In our implementation it did not improve the performance of the model; we therefore decided not to use it for reasons of simplicity and symmetry between "comprehension" (the left part in Figure 2.1) and "production" (the right part).

[2] The latest version can be found at `https://github.com/xtsoukala/dual_path`

2.2.2 Input languages

In order to simulate Spanish-English bilingual sentence production, we generated input with relevant properties of the two languages. The sentences (and their messages, see Section "Message") are generated before the training starts, and they are based on the allowed structures (Section "Structures"). For each part of speech (POS) a randomly selected lexical item (from that POS and target language) is sampled from the lexicon (Section "Bilingual lexicon"). The advantage of using artificial (miniature) languages is that we can manipulate the frequency and grammar of the input and isolate (and thereby study) the phenomenon of interest.

Message

The model is trained using generated sentences (as described above) paired with their message that consists of semantics and their realization, event-semantics, and target language, which will be explained in turn below.

In these simulations, the semantics contains information regarding 45 unique concepts and 6 thematic roles: AGENT, AGENT-MODIFIER, PATIENT, ACTION-LINKING, RECIPIENT, and ATTRIBUTE.

ACTION-LINKING is a combined thematic role that can be used for all verb types: action (e.g., 'throws'), linking ('is') and possession ('has'). ATTRIBUTE is an attribute expressed with a linking verb concept ('BE'). AGENT and RECIPIENT can be expressed only with animate nouns.

A concept (e.g., WOMAN for the English word 'woman' or Spanish word 'mujer') is assigned to each thematic role (during the sentence generation process) along with a realization attribute (PRON for pronoun, DEF for definite article, and INDEF for indefinite article) according to the meaning that needs to be expressed. For instance, in the sentence "the woman runs" the message would include "AGENT=WOMAN, DEF", whereas "a woman" would be encoded as "AGENT=WOMAN, INDEF", and "she" as "AGENT=WOMAN, PRON".

Furthermore, the message contains event semantic information (denoted as EVENT-SEM), which gives information regarding the tense (PRESENT, PAST) and aspect (SIMPLE, PERFECT or PROGRESSIVE). The EVENT-SEM layer also contains information regarding the roles needed for that particular message; the model needs to keep track of the roles expressed and make sure that if, e.g., the role of RECIPIENT is activated then the recipient has also been expressed.

Additionally, the message contains information about the target language so that the model knows whether it is learning to produce an English or Spanish sentence.

Structures

The allowed structures for both languages are the following in our simulations:

- SV: Subject - Verb, e.g., "a happy dog runs"; "un perro feliz corre"
- SVO: Subject - Verb - Object, e.g., "the boy is carrying a book"; "el niño está llevando un libro"
- SVDOPP: Subject - Verb - Direct Object - Prepositional Phrase, e.g., "she shows a book to the girl"; "ella muestra un libro a la niña"
- SVIODO: Subject - Verb - Indirect Object - Direct Object: e.g., "she shows the girl a book" (Structure occurs only in English)
- SVPPDO: Subject - Verb - Prepositional Phrase - Direct Object: e.g., "ella muestra a la niña un libro" ("she shows to a girl a book"). Structure only used in Spanish.

The roles can be expressed using either a Noun Phrase (NP) with definite (DEF) or indefinite (INDEF) article (e.g., 'the woman', 'a woman'). Additionally, AGENT can be expressed with a pronoun (PRON, e.g. 'she'). NPs optionally contain a modifier (an adjective, e.g., 'a tall woman').

The verbs are either intransitive (e.g., 'sleep'), transitive ('carry'), double transitive ('show'), linking ('is') or possession verb ('has'). The tense is present or past and the aspect is simple, progressive, or perfect. Only the simple past was used whereas the present tense is used with all three aspects:

- simple present: "the man cooks"; "el hombre cocina"
- present progressive: "the man is cooking"; "el hombre está cocinando"
- present perfect: "the man has cooked"; "el hombre ha cocinado"

Bilingual lexicon

The lexicon consists of 194 words (Table 2.1): 88 English words, 105 Spanish words, and the shared period ('.') that marks the end of the sentence. The Spanish lexicon is bigger because this language is gendered: for instance, 'tired' is either 'cansado', if it modifies a masculine noun, or 'cansada' for a feminine noun. Syntactic category information (such as 'adjective', 'participle') is not given explicitly; the model learns it through the syntactic path during training.

Table 2.1.: POS in bilingual lexicon (Spanish in italics)

POS	n	Examples
Verbs	64	
auxiliary	4	is, has, *está, ha*
intransitive	32	walked, swims, *nada*
transitive	12	carries, push, *lleva*
double transitive	12	gives, throws, *da*
possession	4	has, had, *tiene, tenía*
linking [1]	4	is, was, *está, estaba*
Participles [2]	57	eating, eaten, *comido*
Nouns	46	
animate	10	uncle, aunt, *tío, tía*
inanimate	36	pen, book, *libro*
Adjectives	22	busy, *ocupado*
Determiners	6	a, the, *un, una, el, la*
Prepositions	2	to, *a*
Pronouns	4	he, she, *él, ella*

[1] Three of these overlap with the auxiliary verbs.
[2] Nine of these have the same form as a verb; e.g., 'walked' is either a perfect participle or a verb.

Input examples

To illustrate the input, here is an example of the message (excluding the target language):

AGENT=WOMAN, INDEF

AGENT-MOD=TALL

ACTION-LINKING=GIVE

PATIENT=BOOK, INDEF

RECIPIENT=GIRL, DEF

EVENT-SEM=SIMPLE, PRESENT, AGENT, AGENT-MOD, PATIENT, RECIPIENT

This message would be expressed linguistically in the following manner in English and Spanish:

- a tall woman gives the girl a book .
- una mujer alta da a la niña un libro . (word-by-word translation: "a woman tall gives to the girl a book")

If the aspect was PROGRESSIVE instead of SIMPLE, on the other hand, the corresponding sentences would be “a tall woman is giving the girl a book”; “una mujer alta está dando a la niña un libro”.

The linking verb messages were encoded in the following manner:

AGENT=WOMAN, DEF

ACTION-LINKING=BE

ATTRIBUTE=TIRED

EVENT-SEM=SIMPLE, PRESENT, AGENT, ATTRIBUTE

and expressed as “the woman is tired”; “la mujer está cansada”.

2.2.3 Training

The model was trained on a total of 3040 randomly generated sentence-message pairs in English and Spanish (training set; 50% [1520 pairs] per language). Recall that no code-switched sentences were given as input.

We ran 60 simulations using different input and different random initial weights per simulation, as the input and the weights are the only non-deterministic parts of the model. The models were trained for 30 epochs, where 1 epoch corresponds to a full iteration of the training set (3040 sentences). At the beginning of each epoch, the training set was shuffled.

The “realization–role” and “role–realization” connection weights were set to 10, and the “concept–role” and “role–concept” to 30. The initial learning rate was 0.10 and linearly decreased over 10 epochs until it reached 0.02; the momentum was set to 0.9. None of the hyper-parameters was optimized for the task, and they do not play a crucial role in the results. We selected the values from Tsoukala, Frank, and Broersma (2017) and increased the “concept–role” connections because this resulted in slightly better performance (the current experiments use more concepts).

2.2.4 Evaluation and performance threshold

The correctness of a sentence is determined by whether the correct (and complete) semantic meaning has been expressed in a grammatical sentence but not necessarily in the target syntactic structure. For instance, if the target sentence is “a sad grandfather is showing the book to a girl” and the produced sentence is “a sad grandfather is showing the *pen* to a girl” it is counted as incorrect, whereas if the produced sentence is “a sad grandfather is showing *a girl the book*” it is counted as correct even though it was expressed with a different syntactic structure than the target one. If it is expressed with a different aspect (e.g., perfect instead of progressive) or realization (e.g., pronoun instead of an NP with an indefinite article) it is also marked as incorrect. If the sentence contains code-switches it is marked as correct as long as it expresses the correct meaning, is expressed in one of the allowed structures (Section “Structures”) and the POS sequence of each phrase (NP, Verb Phase [VP], Prepositional Phrase [PP]) is valid in either language.

For all the experiments, we excluded from the analysis simulations that did not learn to produce at least 75% of the messages correctly according to the criteria above.[3]

2.3 Experiment: Code-switching

In this study, we investigate whether the Bilingual Dual-path would produce code-switched sentences if trained on Spanish and English (but not code-switched) sentences. We investigate the occurrence of different patterns of code-switching that have been observed in the language use of human bilingual speakers.

2.3.1 Background

As mentioned in Section 1.1, Muysken (2000) suggested the following three types of code-switches: insertional, alternational, and congruent lexicalization. In congruent lexicalization, the two languages “share a grammatical structure which can be filled lexically with elements from either language” (Muysken, 2000, p. 6), whereas alternation is a true switch from one language to the other. The difference, however, between congruent lexicalization and the occurrence of consecutive alternations has not been formally operationalized. For that reason, in the simulations reported in this thesis if there are multiple switches (e.g., Spanish-English-Spanish or vice-versa) we treat it as alternational switching and do not differentiate it from congruent lexicalization.

Code-switching by syntactic category

In a seminal study, Poplack (1980) observed the Puerto-Rican community in the US. She found that balanced bilinguals produced mostly complex code-switches, such as intra-sentential ones, and fewer insertions. Switches at the NP were more frequent

[3] In later studies (Chapters 3-6) we decided against this strict criterion and did not exclude any simulations.

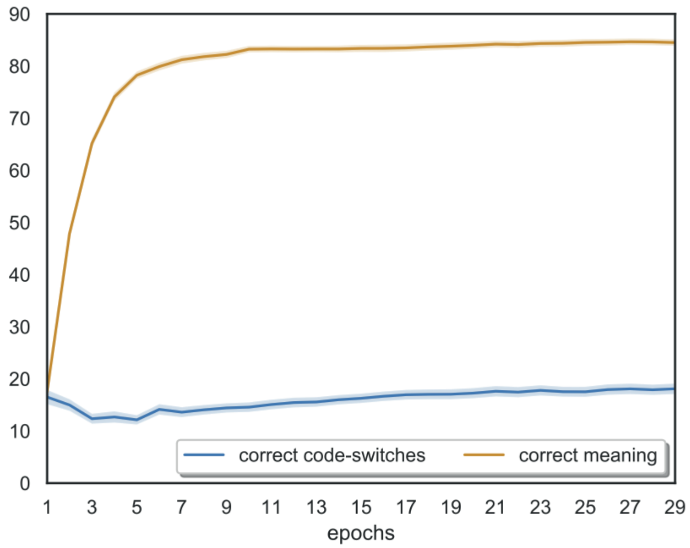

Figure 2.2.: Percentage of correctly produced sentences and of code-switches among those sentences. The shaded area shows the Standard Error of the Mean (SEM)[a] computed over 56 simulations.

[a]Means and standard errors are less appropriate for percentage data. In Chapters 3, 4, and 5 I plot bootstrapped Confidence Intervals instead, which is more appropriate for this type of data. Chapters 2 and 6 were published before the other chapters and I decided to leave the originally published plots which show the SEM.

than switches at the VP and PP, and noun insertions were the most frequent lexical insertion whereas determiner insertions occurred rarely.

2.3.2 Method

To simulate code-switching, we trained the model as described in Section 2.2.3 and tested it on 760 unseen sentences (test set) that were randomly generated in the same manner as the training set.

During the test ("production") phase we manipulated the model's language control by activating a target language only at the beginning, before the production of the first word, so as to indicate the conversational setting (intended language). After the first word had been produced, we activated both target language nodes, thus allowing the model to produce the sentence in either language or to code-switch.

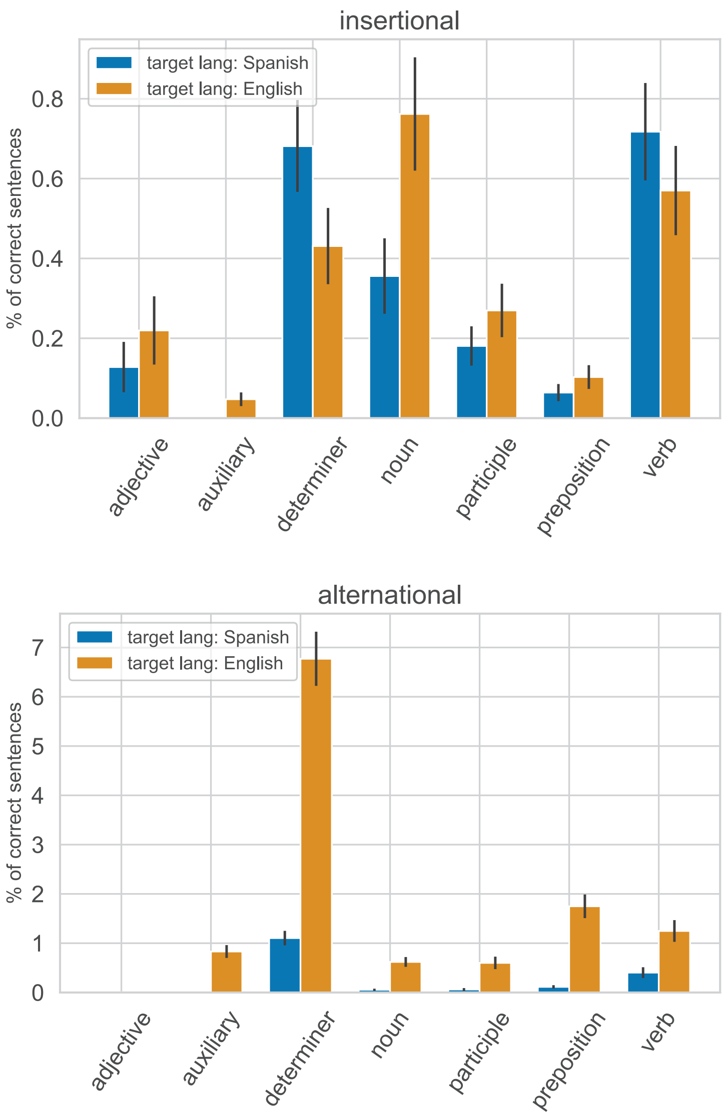

Figure 2.3.: Types of insertional switching (upper row) and alternational switching (lower row). For alternational switches, the POS indicates the first point of switch. All values designate the percentage of correctly produced sentences. The error bars show the SEM computed over 56 simulations.

We excluded from the analysis four models that did not pass the 75% performance threshold (as explained in Section 2.2.4). The reported results are from the remaining 56 simulations.

2.3.3 Results

As hypothesized, the model produced code-switches even though it had not been exposed to code-switched input. The model code-switched in 18.09% of the correctly produced sentences (at the last epoch, see Figure 2.2).

Typology of code-switching in the model's output

Figure 2.3 shows the insertions per POS and the alternational code-switched types (per POS at which the first language switch occurred) that were produced by the model at the end of the training (30th epoch). The model produced alternational switches more frequently than insertional switches (13.57% vs 4.52%).

Examples of code-switched sentences

Insertional code-switches of different syntactic categories are illustrated below:

- Noun insertion:
  Target: un anfitrión feliz ha pateado un bolígrafo . (English: a happy host has kicked a pen)
  Output: un anfitrión feliz ha pateado un *pen* .

- Verb insertion:
  Target: un camarero llevó la llave . (English: a waiter carried the key)
  Output: un camarero *carried* la llave .

- Determiner insertion:
  Target: he is showing the book to the father .
  Output: he is showing *el* book to the father .

- Adjective insertion:
  Target: a man is sad .
  Output: a man is *triste* .

Examples of alternational switches are provided below:

- Alternation at the determiner (Noun Phrase):
  Target: the uncle has shown a father the toy .
  Output: the uncle has shown *un padre* the toy .

- Alternation at the noun:
  Target: the short boy shows a brother a book .
  Output: the short boy shows a *libro a un hermano* .

- Alternation at the preposition (Prepositional Phrase):
  Target: the tall waiter has given a brother a book .
  Output: the tall waiter has given *a un hermano un libro* .

- Alternation at the auxiliary verb (Auxiliary Phrase):
  Target: the short waiter is showing a dog a toy .
  Output: the short waiter *está mostrando a un perro un juguete* .

Note that in the third example (Prepositional Phrase) the model inserted a preposition when switching, thus adhering to Spanish grammar: The double object does not exist with the double noun phrase form in Spanish. This cross-linguistic difference is even more relevant in the fourth example (Auxiliary Phrase switch) because the verb is in Spanish and the sentence would have been entirely ungrammatical if the model had not inserted a preposition ("**a** un perro").

2.3.4 Discussion

The model produced spontaneous code-switches through the manipulation of the target language, without being exposed to code-switched input. This supports the hypothesis that code-switches can occur due to internal and distributional factors, and not only because of exposure to code-switching.

Simulating a balanced bilingual speaker, the model produced mostly alternational switches as opposed to insertional ones. This is in line with Poplack's (1980) observation. Furthermore, alternations at the NP (alternational switch at the determiner) were more likely than alternations at the VP (alternational switch at the verb) or PP (alternational switch at the preposition), which is also in line with the patterns observed by Poplack. However, the model also produced code-switching patterns that are not attested in humans. For instance, the model inserted determiners (1.11% of the correctly produced sentences), especially English determiners in an otherwise Spanish sentence (0.68% of correctly produced sentences). We hypothesize that the model has this preference because determiners in English are not gendered. This means both that the model does not need to select a gendered article and that it prefers to use the English determiner which has twice the frequency of the Spanish ones (as, e.g., 'the' is the translation of both 'el' and 'la' that are the Spanish definite determiners for masculine and feminine nouns respectively).[4]

In bilingual environments where both languages are used, bilingual speakers start with an intended language that is defined by the conversational environment, but they are capable of communicating using either of their languages, or by code-switching (Grosjean, 2001). The top-down language control manipulation in the model (i.e., activating both target languages) is analogous to manipulating the conversational setting in which a speaker is interacting. Spontaneous code-switches occur when there is no target language preference. We only activate a target language right before the production of the first word so as to set the conversational environment.

2.4 Conclusion

We have presented a novel method to test hypotheses in code-switched sentence production. This computational cognitive model can easily be modified to simulate code-switched production of a different language pair. Additionally, the generated input allows for manipulations that help test other hypotheses about code-switching, for instance the idea that cognates can trigger code-switched speech (Clyne, 1980), which is investigated in Chapter 5.

[4]In Chapter 3 that reports on both balanced Spanish–English bilinguals and non-balanced bilinguals with L1 Spanish, this pattern is no longer exhibited by the model.

3 Simulating code-switching in balanced and non-balanced bilinguals

Code-switching is the alternation from one language to the other during bilingual speech. We present a novel method of researching this phenomenon using computational cognitive modeling. We trained a neural network of bilingual sentence production to simulate early balanced Spanish–English bilinguals, late speakers of English who have Spanish as a dominant native language, and late speakers of Spanish who have English as a dominant native language. The model produced code-switches even though it was not exposed to code-switched input. The simulations predicted how code-switching patterns differ between early balanced and late non-balanced bilinguals; the balanced bilingual simulation code-switches considerably more frequently, which is in line with what has been observed in human speech production. Additionally, we compared the patterns produced by the simulations to two corpora of spontaneous bilingual speech and identified noticeable commonalities and differences. To our knowledge, this is the first computational cognitive model simulating the code-switched production of non-balanced bilinguals and comparing the simulated production of balanced and non-balanced bilinguals to that of human bilinguals.

This chapter is based on: Tsoukala, Chara, Mirjam Broersma, Antal van den Bosch, and Stefan L. Frank (2021). “Simulating code-switching using a neural network model of bilingual sentence production”. In: *Computational Brain & Behavior* 4, pp. 87–100

3.1 Introduction

Bilingual speakers are able to switch from one language to the other, between or within sentences, when conversing with other bilinguals who speak the same languages. This process is called code-switching and it is common among communities where two languages come in contact. For instance, Spanish–English code-switching occurs frequently in the United States among Puerto Rican-Americans (Poplack, 1980) and Mexican-Americans (Pfaff, 1979), French–Arabic code-switching is common in Morocco (Bentahila, 1983) and Algeria (Heath, 1984), and Hindi–English in India (Malhotra, 1980).

It is incorrect to think of code-switching as a speech error; bilinguals only code-switch when conversing with others who speak the same languages. Grosjean (1997) suggested that bilinguals utilize their languages differently depending on whom they talk to: when they converse with someone with whom they only share one language, they are in a monolingual language mode. When, on the other hand, they are in a setting in which everybody speaks the same languages, they are in a bilingual mode which allows them to code-switch. The amount of code-switching differs per speaker, depending on their personality as well as on the environment and the context of the conversation (Dewaele and Wei, 2014b).

Another misconception is that bilingual speakers mostly code-switch to fill in lexical gaps; this is not the case for proficient speakers (Romaine, 1986). However, in the early stages of language acquisition speakers code-switch more from their less proficient language into their dominant one, rather than vice-versa, because they lack the linguistic structures and lexicon needed to communicate; this has been observed both in child bilingual acquisition (Genesee, Nicoladis, and Paradis, 1995; Petersen, 1988) and second language (henceforth L2) acquisition (Sert, 2005).

Our goal is to simulate code-switching using a computational cognitive model, with the ultimate aim to further our understanding of the underlying cognitive processes. In Chapter 2 (Tsoukala, Frank, van den Bosch, et al., 2019) we showed that a neural network model designed and validated for monolingual sentence production can generate realistic code-switches when extended to the bilingual case, by training it with syntactic properties and lexical items from two languages and by equipping it with a language control node that allows for the production of either language. Interestingly, the code-switches occurred in the simulations even without exposure to code-switched sentences. In the current study, the aim is to expand on the previous study in the following three ways: First, we will test the robustness of the model's code-switching behavior; we will do this by replicating the study while randomly varying free parameters. Second, we will simulate balanced and non-balanced bilinguals, and shed light on the code-switching patterns of each simulated group. During the early stages of L2 learning we hypothesize that the non-balanced bilingual models will code-switch more from their non-proficient L2 into their dominant native language rather than the other way around, possibly because of gaps in their knowledge of the L2. At the later stages of acquisition, however, when the non-balanced models are more proficient in their L2, we hypothesize that these models will also be able to produce code-switches into their L2, even though we expect to find differences in the frequency and distributions of code-switches between the balanced and non-balanced simulation conditions. Third, we will investigate to what extent the simulated code-switches correspond to what is observed in bilingual speech corpora; this is an exploratory analysis with the goal to validate the model patterns.

Code-switching in balanced and non-balanced bilinguals

Code-switching has been studied mainly in early bilinguals, specifically in i) early balanced bilinguals, i.e., people who have acquired both languages from birth or in early childhood, and ii) heritage speakers whose home language is a minority language (e.g., Spanish in the U.S.) and whose dominant language is usually the one spoken in the community (e.g., English) (see, e.g., Poplack, 1978; Bullock and Toribio, 2009b). Speakers who are exposed to an L2 at a later age (e.g., during adulthood) can also code-switch, although the frequency and patterns are known to be different in early balanced bilinguals. Globalization and greater mobility have caused an increase in the numbers of late non-balanced speakers and there is no social pressure to refrain from code-switching (Matras, 2013).

Most studies comparing balanced and non-balanced bilinguals have focused on comprehension rather than production, specifically on reading comprehension (e.g. using eye-tracking) or grammaticality judgement of code-switched sentences. For instance, Lederberg and Morales (1985) asked different groups of Spanish–English bilinguals to rate the grammaticality of code-switches, and correct them if needed; they compared bilingual children, early balanced bilingual adults, and late non-balanced bilingual speakers who had Spanish as a native language (hereinafter referred to as L1) and moved to the US as adults where they acquired their L2 English. They found that the late non-balanced bilinguals showed differences in the (grammaticality judgment) acceptance rates compared to the balanced bilinguals; however, the code-switching patterns that the groups followed were similar, which led to the conclusion that the rules governing code-switching are not based on extensive exposure to code-switching, but rather on "knowledge of the grammars of the two code-switched languages in combination with some general linguistic knowledge" (Lederberg and Morales, 1985, p.134).

Guzzardo Tamargo and Dussias (2013) studied the reading processing of Spanish-to-English auxiliary phrase switches by balanced and non-balanced bilinguals and found no fundamental differences in the processing patterns between the two groups either, even though the non-balanced bilingual group was slower.

Unlike comprehension studies discussed above, production studies do report differences in the code-switched patterns of early balanced and late non-balanced bilinguals. Poplack (1980) analyzed the code-switching production patterns of Spanish–English bilinguals with varying degrees of proficiency who live in the Puerto-Rican community in New York. She observed that balanced bilinguals produced more complex code-switches (e.g., in the middle of the sentence) whereas speakers who were less proficient in their L2 were more likely to switch only for idiomatic expressions, tags (e.g., "you know"), and fillers (e.g., "I mean"). Similarly, Lantto (2012) analyzed the speech patterns of 22 Basque-English bilinguals (10 among them were early balanced bilinguals) and observed clear differences between the early balanced and late non-balanced bilingual groups, with the balanced group producing a wider variety of switch patterns. Psycholinguistic studies (e.g., Gollan and Ferreira, 2009, using a picture naming task) have also observed that balanced bilinguals code-switch more frequently.

In the current study we simulate balanced and non-balanced bilinguals using a sentence production model. We investigate whether the simulations yield differences between the balanced and non-balanced bilingual groups that are similar to those observed in the aforementioned linguistic studies on the production of code-switches in human speech. We thus assess whether the non-balanced bilingual models show a large likelihood to code-switch in the early stages of L2 acquisition from L2 to L1 and whether in later stages of bilingual production the likelihood to code-switch is higher for the balanced than for the non-balanced models.

The switch directionality of a code-switch (i.e., whether a switch is from the L1 to the L2 or vice-versa), can be determined either in an incremental (i.e., linear, as in, e.g., Broersma and De Bot (2006)) or a hierarchical manner (using, e.g., one of the most influential grammatical models, the Matrix Language-Frame (MLF) model (Myers-Scotton, 1993)). In the following simulations we have taken a linear approach

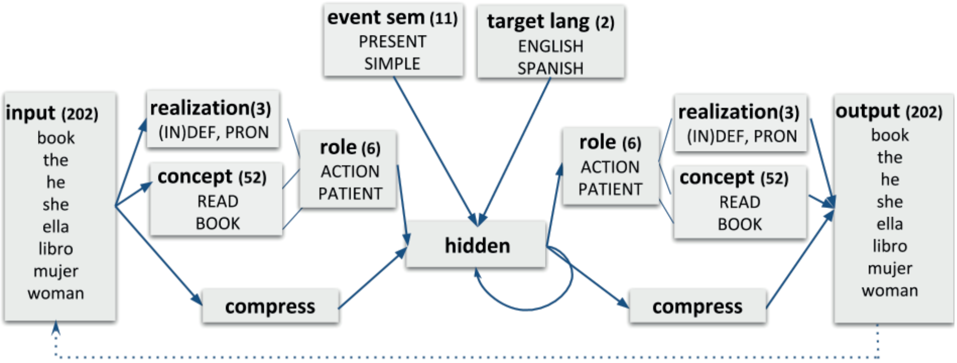

Figure 3.1.: The Bilingual Dual-path model receives messages (see Section 3.2.2 for examples of messages) and expresses them in sentences, word-by-word. The model is based on the Simple Recurrent Network architecture (the syntactic path, via the 'compress' layers), which is augmented with a semantic path that contains information about concepts and their realization, and thematic roles. Additionally, the model receives information on the event semantics and the target language (conversational setting). The numbers in the parentheses indicate the size of each layer (e.g., 52 concept units); the sizes of the hidden and compress layers vary with each training repetition (see Section 3.2.6). The solid arrows denote trainable connection weights, whereas the lines between roles, realization, and concepts correspond to connections that are given as part of the message-to-be-expressed (e.g., PATIENT is connected to BOOK in a particular message). The dotted arrow indicates that the produced word is given back as input, influencing the production of the next word.

because we do not want to make assumptions about the way structural relations affect processing.

3.2 Bilingual Dual-path model

The computational cognitive model we employed for this task is the Bilingual Dual-path model (Tsoukala, Frank, and Broersma, 2017) which is an extension of the Dual-path model (Chang, 2002) of monolingual sentence production. The Dual-path model has successfully modeled a wide range of phenomena over the past years: e.g., structural priming and syntax acquisition in English (Chang, Dell, and Bock, 2006; Fitz and Chang, 2017) and in German (Chang, Baumann, et al., 2015), cross-linguistic differences in word order preference between English and Japanese (Chang, 2009), and input and age of acquisition effects in L2 learning (Janciauskas and Chang, 2018).

We chose to base our model on the Dual-path architecture not only because of its success in modeling sentence production, but also because it is a learning model (a

Recurrent Neural Network, RNN), which therefore allows us to investigate whether code-switched production can emerge from exposure to non-code-switched sentences. Note, however, that a next-word generator, i.e. a simple language model based on an RNN alone or any statistical language model, is very unlikely to produce code-switches without being exposed to code-switched sentences, as the transitional probability between two words in different languages would be zero. The Dual-path, on the other hand, contains a semantic stream and a language control on top of the syntactic stream (the RNN); therefore, it might in theory, and does in practice as our work shows, learn to code-switch even without exposure to code-switched input.

3.2.1 Model architecture

The Bilingual Dual-path model (Figure 3.1) learns to express a given message word-by-word (see Section 3.2.2 for examples of messages). The model assumes that there are two paths influencing sentence production: i) a syntactic path (the lower path in Figure 3.1, via the 'compress' layers), which is a Simple Recurrent Network (Elman, 1990), and ii) a semantic path (the upper path in Figure 3.1), which contains information about the thematic roles (e.g., AGENT, RECIPIENT), the concepts they are connected to, and their realization. The syntactic path learns the syntactic patterns of each language, whereas the semantic path learns to map concepts onto words. Additionally, the model receives information about i) the event semantics that define when an event takes place (e.g. PAST, PROGRESSIVE), and ii) the target language, through the corresponding node, which acts as the only language control of the model. Specifically, the target language node simulates the conversational setting in which a speaker is interacting (i.e., one target language in a monolingual setting, both languages in a bilingual setting). All layers use the tanh activation function, except for role and output that use softmax.

The two streams, along with the event semantics and target language, work together to produce grammatically correct sentences that express a specific message.

3.2.2 Messages

A message is represented by i) a target language, ii) event-semantic information, iii) pairs of thematic roles and concepts, and iv) pairs of thematic roles and realizations (pronoun, definite, indefinite) whenever applicable in the case of noun phrases.

The target languages are ENGLISH and SPANISH. The event semantics contain information regarding the aspect (SIMPLE, PROGRESSIVE, PERFECT) and tense (PRESENT, PAST), as well as the thematic roles that are used in each message.

The following simulations make use of 52 unique concepts and six thematic roles: AGENT, PATIENT, AGENT-MODIFIER, RECIPIENT, ACTION- LINKING, and ATTRIBUTE. The roles AGENT and RECIPIENT are only paired with animate nouns (e.g., 'son', 'cat'). ACTION-LINKING is a combined thematic role that can be used for all main verb types: action (e.g., 'shows'), linking ('is') and possession ('has'). ATTRIBUTE is an attribute expressed only with a linking verb.

Additionally, AGENT, PATIENT and RECIPIENT are connected to their realization: pronoun (e.g., 'he' for the concept BOY), and definite or indefinite article for concepts that are expressed as a noun phrase (e.g., 'the boy' or 'a boy'). These roles contain optionally a modifier (an adjective, e.g., 'a happy dog'). Note that in English the adjective comes before the noun ("the intelligent woman") whereas in Spanish the modifier comes after the noun ("la mujer inteligente"). This knowledge is learned by the model through the training examples and not through explicit syntactic labels.

3.2.3 (Message-to-)Sentence production

For a message to be expressed, the following nodes need to be activated in the model: the event-semantics (e.g., PRESENT, PAST) and the target language (ENGLISH, SPANISH) that specifies the intended output language. Furthermore, the semantic roles (e.g., ACTION, PATIENT) are connected to their respective concepts (e.g., READ, BOOK) and realizations (e.g., INDEF for an indefinite article). For example, if the message is:

AGENT = WOMAN, DEF

ACTION = GIVE

PATIENT = BOOK, INDEF

RECIPIENT = FATHER, INDEF

TARGET-LANG = ENGLISH

EVENT-SEM = PRESENT, PROGRESSIVE,
AGENT, PATIENT, RECIPIENT

the model would learn to express it in English as "the woman is giving a book to a father", and if the target language was Spanish as "la mujer está dando un libro a un padre". Following Chang, Dell, and Bock (2006), to express the recipient before the patient ("the woman is giving a father a book" or "la mujer está dando a un padre un libro") the PATIENT receives less activation through the event semantics, thus prioritizing the RECIPIENT. In the current example, the event semantics would be: PRESENT, PROGRESSIVE, AGENT, PATIENT:0.5, RECIPIENT.

When the model receives a message, it produces it word-by-word. The produced word is the output word with the highest activation. Each produced word is then given as input in the next time step, and it influences the production of the next word. The period ('.') works as an end-of-sentence marker and the model stops producing words when it outputs the period or if it has exceeded the length of the target sentence, plus 2 extra words. We allow extra words because the model might produce a different structure than the target one; for instance, the message of the sentence 'the boy is giving the girl a key' (double object) could also be expressed as 'the boy is giving the key to the girl' (prepositional dative).

3.2.4 Miniature Languages

In order to simulate Spanish–English sentence production, we generated training sentences that are derived from a small subset of the syntactic properties (Section "Tense and aspect") and the lexica (Section "Bilingual lexicon") of the two languages. Note that the model does not contain a phonological level because we are only focusing on the interaction between semantics and syntax, and not on restrictions imposed by phonology.

Tense and aspect

The allowed tenses used in the structures are past and present, and the aspects: simple, progressive, and perfect. The past tense is only used in simple aspect sentences (e.g., "the girl jumped"), whereas the present tense applies to all three aspects. The allowed structures for the two languages and all tenses and aspects can be found in Table 3.1.

Bilingual lexicon

The bilingual lexicon (Table 3.2) is an extension of the lexicon used in the Tsoukala, Frank, van den Bosch, et al. (2019) study (Chapter 2). It contains 202 words: 92 English words, 109 Spanish words, and the shared end-of-sentence marker ('.'). Recall that the Spanish lexicon is larger because Spanish is a gendered language. For instance, nouns and adjectives are usually expressed differently depending on whether they

Table 3.1.: Allowed structures for English and Spanish. Abbreviations: S: Subject, V: Verb, O: Object, D: Direct object, I: Indirect object.

Aspect & Tense	English example	Spanish example
Simple Past		
SV	the brother sneezed	el hermano estornudó
SVO	a gentleman pushed the chair	un señor empujó la silla
Simple Present		
SV	the girl swims	la niña nada
SVO	the grandmother kicks the toy	la abuela patea el juguete
SVO (linking)	the hostess is happy	la anfitriona está feliz
SVO (possession)	a dog has the ball	un perro tiene la pelota
SVDI	the man throws a book to the aunt	el hombre tira un libro a la tía
SVID	the man throws the aunt a book	el hombre tira a la tía un libro
Perfect Present		
SV	he has walked	él ha caminado
SVO	a woman has pushed the pen	una mujer ha empujado el bolígrafo
SVDI	a girl has thrown a key to the lady	una niña ha tirado una llave a la señora
SVID	a girl has thrown the lady a key	una niña ha tirado a la señora una llave
Progressive Perfect		
SV	she is jumping	ella está saltando
SVO	he is kicking a chair	él está pateando una silla
SVDI	the father is giving the toy to a girl	el padre está dando el juguete a una niña
SVID	the father is giving a girl the toy	el padre está dando a una niña el juguete

Table 3.2.: Syntactic categories in the bilingual lexicon (Spanish in italics)

Syntactic category	n	Examples
Verbs	66	
auxiliary	4	is, has, *está, ha*
intransitive	16	walked, swims, *nada*
transitive	12	carries, pushed, *lleva*
double	12	throws, gives, *tira*
possession	4	has, had, *tiene, tenía*
linking [1]	2	is, *está*
Participles [2]	56	
progressive	28	walking, *caminando*
perfect	28	eaten, *comido*
Nouns	52	
animate	40	uncle, aunt, *tío, tía*
inanimate	12	pen, book, *libro*
Prepositions	2	to, *a*[3]
(Predicate) adjectives	26	busy, *ocupado*
Determiners	6	a, the, *un, una, el, la*
Pronouns	4	he, she, *él, ella*

[1] Both linking verbs overlap with the auxiliary verbs.
[2] Nine of these have the same form as a verb; e.g., 'walked' is either a perfect participle or a verb.
[3] The Spanish preposition '*a*' and the English indefinite article 'a' are differentiated, in that they have separate nodes in the lexicon.

modify a masculine noun or a feminine one. We also included four common-gendered Spanish adjectives such as 'feliz' ('happy') that do not change depending on the noun it modifies.

The verbs are either intransitive (e.g., 'swims'), transitive ('carries'), double ('throws'), linking ('is', 'está'), or possession verb ('has', 'tiene'). The two linking verbs ('is', 'está') and the English possession verb ('has') were also used as auxiliary verbs for the progressive and perfect forms, respectively, as was the Spanish perfect-form auxiliary verb 'haber'. Following the allowed structures, each verb had four forms: simple present, simple past, present participle and past participle.

Note that syntactic category information (such as 'noun', 'participle') is not given explicitly; the model learns through training (via the syntactic path) that words that occur in similar context tend to be of the same syntactic category.

3.2.5 Message-sentence pair examples

We hereby illustrate how a message corresponds to (and is expressed with) a sentence. For instance, the following message:

> AGENT = WAITER, DEF
>
> AGENT-MOD = TALL
>
> ACTION-LINKING = SNEEZE
>
> EVENT-SEM = SIMPLE, PAST, AGENT, AGENT-MOD

corresponds to the following sentences in English and Spanish:

- the tall waiter sneezed .
- el camarero alto estornudó . (literally: “the waiter tall sneezed .”)

Changing the tense of the message to PRESENT instead of PAST would correspond to the sentences “the tall waiter sneezes” and “el camarero alto estornuda”, whereas further changing the aspect to PROGRESSIVE instead of SIMPLE would correspond to “the tall waiter is sneezing” and “el camarero alto está estornudando”.

Messages that contain direct and indirect objects can be expressed with the thematic roles of PATIENT and RECIPIENT respectively. For instance, the following message:

> AGENT = FATHER, PRON
>
> ACTION-LINKING = THROW
>
> PATIENT = BALL, DEF
>
> RECIPIENT = DOG, INDEF
>
> EVENT-SEM = SIMPLE, PRESENT, AGENT,
> PATIENT, RECIPIENT

is expressed as “he throws the ball to a dog” or “él tira la pelota a un perro”.

Finally, messages that contain linking verbs are encoded using an attribute:

> AGENT = MAN, DEF
>
> AGENT-MOD = KIND
>
> ACTION-LINKING = BE

ATTRIBUTE = TIRED

EVENT-SEM = SIMPLE, PRESENT, AGENT, AGENT-MOD, ATTRIBUTE

which is expressed as "the kind man is tired" or "el hombre amable está cansado".

3.2.6 Model training

The model learns through supervised training. A message is given as input and the network tries to generate a sentence word-by-word; after a word has been produced, it is compared to the target word and the weights are adjusted according to the backpropagation algorithm. All networks were trained for 40 epochs using 2,000 message-sentence pairs.

The backpropagation parameters were the same across all simulations: the momentum was set to 0.9 and the initial learning rate was 0.10, which linearly decreased for 10 epochs until it reached 0.02, at which point it was held constant from epoch 11 onward. This applies to both the Balanced and the Non-balanced bilingual models (Sections 3.3.1 and 3.3.2, respectively). Note that the Non-balanced models are exposed to their L2 around the 15th epoch; therefore, they start learning the L2 with a decreased learning rate (0.02).

To increase the generalizability of the reported results and to test the robustness of the results reported in Tsoukala, Frank, van den Bosch, et al. (2019), we trained 40 networks per simulation, randomizing all free parameters (as seen below), excluding the backpropagation parameters (i.e., the momentum and learning rate) and the training set size. The parameters were randomized per training repetition (i.e., for each of the 40 networks), but the same parameter values were kept across the three different simulations: e.g., the first training repetition of the balanced bilingual simulation had the same initialized weights as the first training repetition of the non-balanced bilingual model(s).

First, the message-sentence pairs were randomly generated for each simulation before the training started. The sentences were constrained by a set of allowed structures (Section "Tense and aspect") and for each syntactic category a randomly selected word was sampled from the bilingual lexicon (Section "Bilingual lexicon"). Note that the target sentences were never code-switched.

Second, when sampling the bilingual training set for the balanced and non-balanced bilingual simulations, we varied the percentage of English and Spanish. The percentage of English was sampled from a normal distribution with a mean of

50% (standard deviation: 8) and the rest was Spanish. Third, weights of trainable connections were initialized using Xavier initialization (Glorot and Bengio, 2010). Last, the weights of the connections between thematic roles and concepts ('concept'–'role' and 'predicted role'–'predicted concept' in Figure 3.1), which are not trained, were integer values sampled between 10 and 20 (the exact value was randomized once per training repetition and was the same for all these connections).

The hidden layer size was also sampled per training repetition (between 70 and 90 units) and the compress layer size was set to the closest integer to 77% of the hidden layer size.

3.2.7 Code-switching

As mentioned above, the target sentences did not contain any code-switches. To allow the model to code-switch, we manipulated the model's language control (target language node) when testing, which simulates the conversational setting, or language mode (Grosjean, 1997). Only one of the target languages was activated before the production of the first word, and the network was thereby biased towards producing the first word in that language, but once the first word was produced, both languages were activated. This allowed the model to continue in the same language or to code-switch.

With regard to the code-switching types, in the current simulations we look at two types of code-switches, which Muysken (2000) calls (lexical) insertions and (intra-sentential) alternations[1]:

- Insertional switching (insertion of single words)
  e.g.: "He gave the *libro* to my niece." (He gave the book to my niece.)

- Alternational switching (intra-sentential switching)
  e.g.: "*María prefiere hacer el viaje* by train instead." (Maria prefers to make the trip by train instead.)

[1]Muysken also identified other types of code switches (i.e., congruent lexicalization) and sub-categories of the insertions and alternations (e.g., insertions of fixed expressions, idioms and tags, and alternations between sentences called "inter-sentential switching"), but these fall beyond the scope of the model because the model produces single sentences without context and without the usage of fixed expressions and tags.

3.2.8 Correctness of sentence production

A produced sentence is considered grammatically correct if it consists of an allowed sequence of syntactic categories, i.e., if the sequence exists in the training set. The criterion for correct meaning is that the sentence is grammatical and that all thematic roles are expressed correctly, even if they are code-switched, but with no omitted or extra attributes (e.g., not "dog" instead of "big dog" or vice versa). In some cases, the meaning can be correct even if the produced sentence is different than the target. For instance, if a double object sentence ("the woman gives the cat a ball") is expressed as a prepositional dative ("the woman gives a ball to the cat"), the meaning is counted as correct because the message is expressed correctly.

3.3 Method: Simulations and Corpus Analysis

We address the three goals of this study by running three sets of simulations. First, having expanded the lexicon and having varied almost all free parameters compared to the Tsoukala, Frank, van den Bosch, et al. (2019) study (see 3.2.6), we run 40 training repetitions, with different parameters each, to investigate i) whether the model again produces code-switched sentences, and ii) the sensitivity of this behavior to the random parameter settings and initial weights.

Second, we simulate balanced and non-balanced Spanish–English bilinguals and compared their production patterns with respect to code-switching. Specifically, we measure i) how often a sentence is code-switched in total and per switch direction[2] (Spanish-to-English vs English-to-Spanish), ii) what kind of code-switches (alternational, insertional) are produced and at which syntactic point, and iii) how the patterns vary with the amount of training and exposure to the two languages. Note that each epoch corresponds to the amount (time) of learning, not the amount of training examples per language; the non-balanced bilingual models are initially exposed only to their L1, whereas the balanced bilingual model directly receives bilingual message-sentence pairs, thus receiving approximately half the exposure per epoch to each individual language. Third, we test the validity of the simulated patterns by comparing them to human data, i.e., code-switched utterances in bilingual speech corpora.

To address the first two goals, we run one early balanced bilingual model and two late non-balanced bilingual models with different L1 (English, Spanish).

[2]As mentioned in Section 3.1, our approach is linear. We start from the first word of a code-switched output sentence; if the word is in Spanish, we mark the switch direction as "Spanish-to-English", whereas if the first word is English we count it as "English-to-Spanish"

For the third goal, we analyze the Bangor Miami corpus (Deuchar et al., 2014) to obtain code-switched patterns of Spanish–English bilingual speech.

3.3.1 Balanced bilingual model (Balanced model)

The Balanced model was simultaneously exposed to both languages (roughly 50% per language as described in Section 3.2.6), therefore simulating balanced bilinguals. The Balanced model was trained for 40 epochs using 2,000 message-sentence pairs and tested on 500 messages. The training and test sets were unique per training repetition (40 training repetitions in total) and the distribution of Spanish and English in the test messages was the same as in the training messages.

3.3.2 Non-balanced bilingual models (L1 English and L1 Spanish models)

The non-balanced bilingual models were first exposed only to their L1 for roughly 15 epochs. Specifically, the L1 English model was trained with English-only sentences (2,000 message-sentence pairs) for about 15 epochs (the exact number of epochs was randomly sampled between 13 and 17), whereas the L1 Spanish model was initially trained on Spanish-only sentences (2,000 message-sentence pairs). For the remaining epochs (making a total of 40) the networks were exposed to the same 2,000 message-sentence pairs as the Balanced model and tested on the same 500 messages. Once again, there were 40 training repetitions per model and the message-sentence pairs and test messages were different for each run.

3.3.3 Corpus analysis

To compare the simulated patterns to human data, we analyzed the transcriptions of the Bangor Miami corpus (Deuchar et al., 2014)[3] that consists of 56 spontaneous and informal conversations between two-to-five speakers, living in Miami, Florida. Out of the 84 speakers, 60 were equally fluent in English and Spanish[4]. Each word in the conversation file has been automatically tagged with a language code (English or Spanish) and a syntactic category (e.g., noun). We selected the sentences that contained more than one language code, resulting in 2,796 code-switched sentences, which is

[3] http://bangortalk.org.uk/speakers.php?c=miami

[4] Fluency was measured by self-reported “Spanish ability” and “English ability”. The questionnaire results can be found on the corpus website.

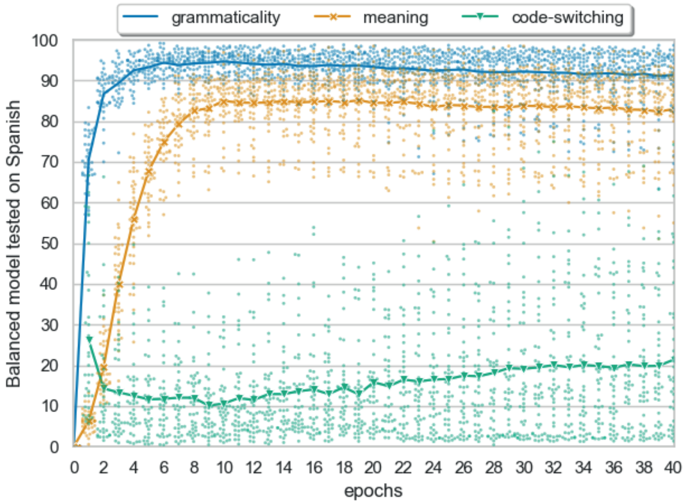

a: Balanced model tested on Spanish

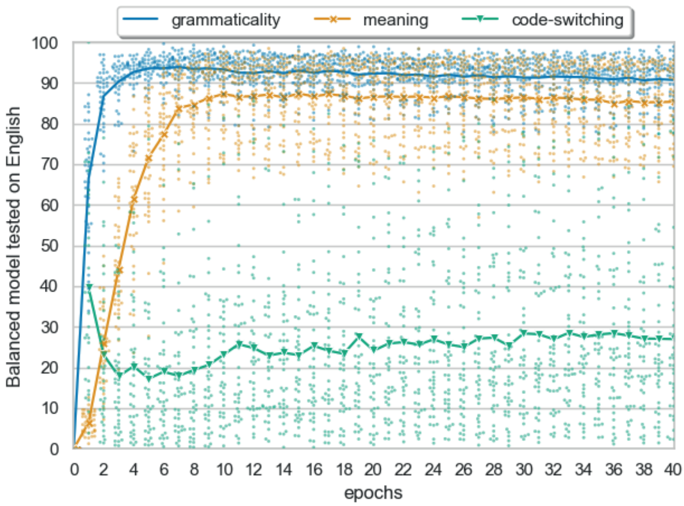

b: Balanced model tested on English

Figure 3.2.: Mean grammaticality, correctness of meaning, and code-switch percentage of the balanced bilingual model, tested on Spanish (a) and English (b). The dots are jittered and represent each individual training repetition.

6.2% of the corpus (45,289 sentences in total). We then divided the code-switches into alternations, in case the sentence continued in the code-switched language, and insertions, if the code-switches were single words that were inserted (once or several times) in the sentence. Meanwhile, we corrected the syntactic categories of erroneous or missing tags[5].

From the 2,796 code-switches observed in the Miami corpus, we included only the 1,369 that occur at syntactic categories that are relevant to our model; for instance, we excluded interjection insertions because interjections are out of the scope of the model.

As an additional corpus, we compared the model's patterns to the code-switches observed in Poplack (1980). Note that Poplack's corpus is not publicly available; therefore, we could not re-analyze the data. Likewise, out of the 1,835 code-switching instances observed in Poplack's study we have only included the syntactic categories that are relevant for our study.

Note that in both corpora most switches are so-called extra-sentential, which are not grammatically or semantically related to any other part of the sentence (e.g., tag insertions, such as "you know" and "right?") and are therefore not included in the model.

3.4 Results

3.4.1 Model performance

Balanced model

Figure 3.2 shows the performance (i.e., percentage of sentences with correct grammar and with correct grammar and meaning) of the balanced bilingual model on its two native languages: Spanish (3.2a) and English (3.2b). Both languages are learned equally well: the mean percentage of sentences that are produced with correct meaning at the last epoch (hereinafter: correct meaning) is 83% for Spanish and 85.4% for English.

[5]The scripts used, as well as the resulting sentences, can be found at `https://osf.io/vd3wa/`.

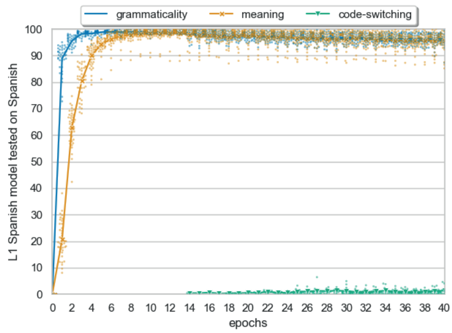

a: L1 Spanish tested on Spanish

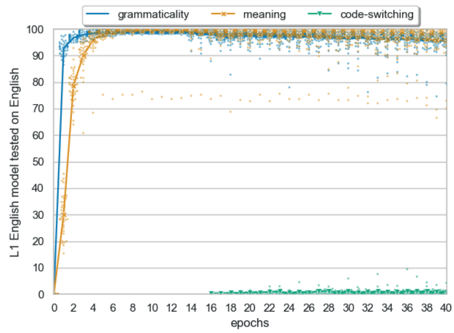

b: L1 English tested on English

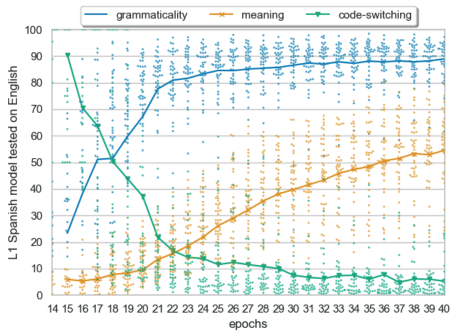

c: L1 Spanish tested on L2 English

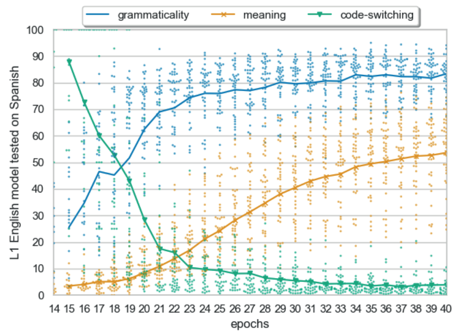

d: L1 English tested on L2 Spanish

Figure 3.3.: Mean grammaticality, correctness of meaning and code-switch percentage of the non-balanced bilingual models, tested on their L1 (upper row) and L2 languages (lower row) over 40 training repetitions (L1 Spanish model: left column, L1 English model: right column). The dots are jittered and represent each individual training repetition.

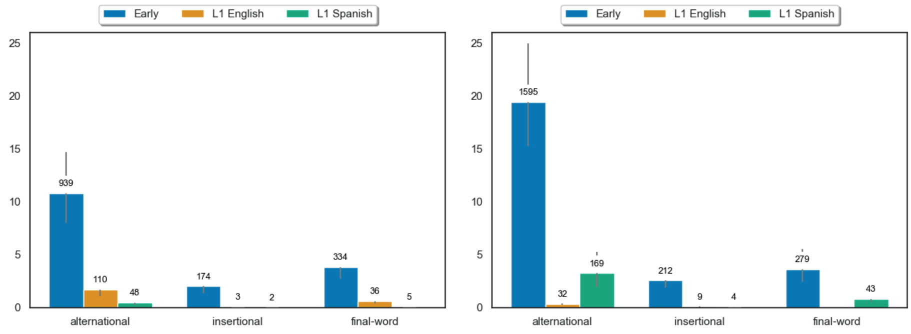

a: Spanish → English **b:** English → Spanish

Figure 3.4.: Model comparison of code-switching types. The percentage is against all correctly produced sentences of that target language (left: Spanish; right: English) and of that model, and the numbers denote the absolute number of sentences with that switch type. The error bars show the 10,000-sample bootstrapped 95% Confidence Interval.

Table 3.3.: Percentage of switch types per model over all correctly produced sentences after 25 epochs of exposure to both languages (25th epoch for the Balanced model and 40th for the L1 English and L1 Spanish models). The numbers in the brackets show the 10,000-sample bootstrapped 95% Confidence Interval.

Switch type	Balanced	L1 English	L1 Spanish
Alternation	15.0% [12.7, 17.9]	1.7% [1.0, 3.1]	1.1% [0.8, 1.6]
Insertion	2.3% [1.8, 2.9]	0.1% [0.1, 0.2]	0.0% [0.0, 0.1]
Final-word	3.7% [2.9, 5.0]	0.3% [0.2, 0.5]	0.2% [0.1, 0.3]
Total	21.0% [17.9, 24.9]	2.1% [1.4, 3.6]	1.3% [1.0, 1.8]

L1 English and L1 Spanish models

Figure 3.3 (upper row) shows the performance of the L1 Spanish (3.3a) and L1 English (3.3b) models on their native language. Note that around the 15th epoch the L2 is introduced which slightly affects the production of the L1.

The lower row of Figure 3.3 indicates the performance of the non-balanced bilingual models on their L2, starting from the epoch in which the L2 was introduced. Note that because the exact starting epoch varies per training repetition, only after the 18th epoch are all 40 training repetitions introduced to the L2; until then, the plot shows the mean only of the training repetitions that have already been exposed to the L2. Figure 3.3c shows the performance of L2 English in the L1 Spanish model and Figure 3.3d shows the L2 Spanish performance of the L1 English model.

3.4.2 Research goal 1: Code-switching in the models

In the final epoch, the Balanced model (Figure 3.2) produces 21.4% Spanish-to-English and 27.0% English-to-Spanish code-switches, out of all correctly produced sentences. Examples of code-switched sentences include[6]:

- a boy pushed *la silla* ("the chair")
- a happy cat *tiene una pelota* ("has a ball")
- the uncle *está triste* ("is sad")
- a dog *corrió* ("ran")

The non-balanced models' code-switching patterns develop over time: at the early stages of L2 learning they produce very few L2 sentences correctly, most of which contain code-switches into the L1; for instance, on the 14th epoch the L1 English model produces 3.5% of Spanish sentences correctly, out of which 87.9% contain a code-switch into English. Respectively, the L1 Spanish model produces 5.9% of English sentences correctly, with 90.5% of these containing a switch into Spanish. Over time, though, the models become more proficient in their L2 and stop reverting to their L1: at the end of the training the L1 Spanish model reaches 55.1% in meaning accuracy of English sentences and produces 5.3% switches into Spanish, whereas the L1 English model reaches 55.9% accuracy on Spanish and switches back into English 3.9% of the time. Code-switches from the L1 into the L2 are more steady throughout acquisition: the L1 English model code-switches 0.9% from English into Spanish, and the L1 Spanish model code-switches 1.3% of the time from Spanish into English.

3.4.3 Research goal 2: Balanced versus non-balanced model comparisons

The second goal of this study is to investigate the differences in code-switch types produced by the balanced and non-balanced models at the late stages of acquisition, when both models have been exposed to the bilingual input for 25 epochs (i.e., the 25th epoch for the balanced model vs. the 40th epoch for the non-balanced models). Table 3.3 presents the percentage of the total code-switch types (alternational, insertional, and final-word, in case the switch is at the end of the sentence and it is therefore unclear whether it is an insertion or an alternation) for the three models (balanced

[6] The full list of output sentences per model can be viewed online at `https://osf.io/vd3wa/`

Spanish–English bilingual, non-balanced bilingual with L1 English, non-balanced bilingual with L1 Spanish). The balanced bilingual model code-switches much more frequently than the L1 Spanish and L1 English models.

Figure 3.4 compares the three models with respect to the switch type and switch direction. The percentages shown here are against all correctly produced sentences *of that target language*, not of all sentences as in Table 3.3.

Additional information on the exact code-switching patterns per switch type (alternational, insertional, final-word), language direction (English-to-Spanish and Spanish-to-English) and syntactic category in which the switch occurs can be found at `https://osf.io/vd3wa/` under `results/supplementary_plots`.

3.4.4 Research goal 3: Model versus corpus comparison

To test the validity of the produced patterns, we compared the simulated code-switched patterns to the ones observed in the Miami corpus, as well as in the patterns observed in Poplack's (1980) study. The results can be found in Table 3.4.

Both corpora contain code-switches from all participants, both balanced bilinguals and Spanish-dominant. To compare the model results to the corpora that contain both balanced and Spanish-dominant speakers, Table 3.4 reports switches from the Balanced and L1 Spanish models combined. The simulations produce a high percentage of noun phrase alternations, which is also the case in Poplack's study and the Miami corpus. Furthermore, both the corpora and the model display a substantial (but small) amount of verb alternations. The other phrase alternations are fewer in the models than in Poplack's data, which is probably due to the fact that with the current artificial languages we have only simulated prepositional phrases whereas the phrase alternation in the corpus include other types of phrases as well (i.e., adjective, adverb and infinitive phrases). Both the simulations and human bilinguals, especially in Poplack's study, disprefer preposition insertions.

There are also clear differences between simulated and empirical code-switches. Unlike human bilinguals, the model seems to favor determiner insertion, and more specifically Spanish determiners. The most striking difference between corpora and model results, however, is found in noun insertions: both corpora showcase that noun insertion is a major switching category among human bilinguals. The model, however, produces less than 1% of noun insertions, whereas the two corpus data report over 30% of noun insertions.

Table 3.4.: Number of code-switches by syntactic category and language in i) the balanced bilingual and L1 Spanish models ('Sim.') ii) the Miami corpus ('Miami'), and iii) Table 2 of Poplack (1980) ('Pop.'), adapted to include only syntactic categories the model produces. In alternational switching, the code-switching starts at the syntactic category presented in the leftmost column and continues in the non-target language. For instance, an adjective alternation within a noun phrase means that the adjective of the noun phrase was the first code-switched item of that sentence.

Syntactic category of CS	# English to Spanish			# Spanish to English			# Total CS			% of total CS		
	Sim.	Miami	Pop.	Sim.	Miami	Pop.	Sim.	Miami	Pop.	Sim.	Miami	Pop.
insertions:												
determiner	102	7	3	49	3	0	151	10	3	4.8	0.7	0.5
noun	21	120	34	6	367	141	27	487	175	0.9	35.6	31.4
auxiliary	32	0	0	51	0	0	83	0	0	2.6	0.0	0.0
verb	44	25	6	58	9	13	102	34	19	3.2	2.5	3.4
adjective	0	24	3	2	58	12	2	82	15	0.1	6.0	2.7
preposition	15	9	2	2	10	0	17	19	2	0.5	1.4	0.4
alternations:												
noun phrase			106			103			209			37.5
determiner	954	77		517	51		1471	128		47.0	9.4	
noun	38	25		8	178		46	203		1.5	14.8	
adjective	0	30		9	87		9	117		0.3	8.5	
verb phrase			27			13			40			7.2
verb	297	104		183	57		480	161		15.3	11.8	
auxiliary	304			241			545			17.4		
participle	71			28			99			3.2		
prepositional phrase	100	68	55	1	60	39	101	128	94	3.2	9.3	16.9
Total	1978	489	236	1155	880	321	3133	1369	557	100.0	100.0	100.0

3.5 Discussion

3.5.1 Code-switching in the models

The first goal of this paper was to verify the robustness of the code-switch model presented in Tsoukala, Frank, van den Bosch, et al. (2019). Having varied almost all free parameters in the current simulations, and using an expanded lexicon, we tested again whether the bilingual model is able to produce code-switches that are attested in bilingual speech, even without having been exposed to code-switched input. The models indeed produced code-switches, thus confirming that code-switching can partially be explained by the distribution of the two languages involved (in combination with the cognitive architecture of the model, in our simulations). This is in line with Lederberg and Morales (1985), who claimed that (extensive) exposure to code-switching is not needed for a bilingual speaker to code-switch.

Importantly, the model is able to code-switch by merely having (and manipulating) a language control ("target language") node that sets the conversational setting and allows the model to produce in either language. No other cognitive control was required for the model to code-switch.

It is interesting to observe the huge variance in the amount of code-switching between training repetitions; at the last epoch of the balanced bilingual model tested on English (Figure 3.2b), the percentage of code-switches produced by the 40 models ranges from 2.3% to 80.8%. Large individual variance is something that has also been observed among human bilinguals (Dewaele and Wei, 2014b).

As mentioned in Section 3.2, an RNN alone trained on non-code-switched data is unlikely to produce code-switched sentences. As a case in point, we trained the SRN-only part of the model (i.e., the syntactic stream alone) using the same input and settings as described in the Methods section. It is difficult to directly compare the Dual-path to an SRN-only version because the former expresses a specific message; for an approximate comparison, we gave the SRN-only model the first word of the target message and let it produce any sentence. The SRN-only model learned to produce grammatical sentences but it did not produce any code-switched sentences.

3.5.2 Balanced versus non-balanced model comparisons

The second aim of this study was two-fold: First, to investigate the development of code-switches over time in the non-balanced bilingual models. Second, to compare the production patterns of balanced and non-balanced bilinguals and per switch direction.

On the one hand, at the early stages of L2 acquisition, the non-balanced models have not been exposed enough to their L2 and they strongly prefer to switch back into their L1 (i.e., over 87% of the time). This preference is in line with what has been observed in bilingual language acquisition by children (e.g., Petersen, 1988). When comparing, on the other hand, the balanced and non-balanced models after an equal amount of exposure to both languages (25 epochs), the patterns change: the balanced bilingual model code-switches considerably more frequently than the non-balanced bilingual models, which is in line with what has been observed in humans (e.g., Poplack, 1980; Gollan and Ferreira, 2009).

Note that the non-balanced bilingual models perform better in their L1 compared to the balanced bilingual model: 95.5% accuracy in meaning in the last epoch for the L1 Spanish model and 95.9% for the L1 English model on their L1 (Figure 3.3), as opposed to 85.4% for English and 83% for Spanish accuracy in the Balanced model (Figure 3.2). The reason behind this discrepancy is that the non-balanced bilingual models receive double the input in their L1 (for the first 15 epochs) compared to the balanced bilingual model that has two native languages. As mentioned above, an epoch corresponds to the amount of learning time, not the input received.

In the current simulations we have assumed that the L1 is the dominant language. However, a large proportion of bilingual speakers in communities that code-switch are heritage speakers who, as mentioned in Section 3.1, are more exposed to (and fluent in) their L2, the majority language of the country they live in, rather than the L1 that is mostly spoken at home. Heritage speakers could also be simulated in the model, by first exposing the model to the L1 only (similar to the non-balanced models) and then introducing bilingual input in which the L2 is much more frequent than the L1, reflecting heritage speakers' exposure.

3.5.3 Model versus corpus comparison

The third goal of this paper was explorative, aiming to validate the model by investigating to what extent the simulated patterns correspond to bilingual speakers' behavior. We cannot expect a perfect match between the model's code-switching patterns and the corpus data because the simulations use an artificial micro version of English and Spanish. Nevertheless, some patterns are similar to what human bilinguals produced in the two corpora (see Table 3.4 for percentages of patterns of all three studies). There are also noticeable differences between the modeled and human code-switched patterns, with the most striking one being the high amount of noun insertions in the two corpora compared to the simulations. A possible explanation for this discrepancy is that human bilinguals tend to prefer a specific language depending on the domain, for instance, Spanish for food, English for school- and work-related terms (Fishman, Cooper, and Newman, 1971). Additionally, bilinguals align with their collocutors and repeat syntactic structures and utterances (Fricke and Kootstra, 2016). In the Miami corpus, for instance, when we analyzed the noun insertions per dialogue (chat transcription), we found 117 repetitions out of the 487 insertions. The model, on the other hand, simulates individual sentences and has no context of what has been produced before, nor a notion of domain-specific terms; the only context given is the language control, which specifies whether the setting is monolingual or bilingual. Another minor reason behind the small number of noun insertions in the model simulations is that we have excluded from our analysis final-word switches, as we are unable to determine whether they are insertions or alternations. In the corpus analysis, on the other hand, we counted 296 final-word noun switches as noun insertions.

3.6 Conclusion

We have shown that the Tsoukala, Frank, van den Bosch, et al. (2019) results are robust: the Bilingual Dual-Path model can produce code-switching patterns without exposure to such code-switched patterns, by only having a language control that allows the model to produce in either language. Furthermore, we simulated three groups of bilinguals and showed the differences between the early balanced and late non-balanced simulated bilingual populations, as well as the development of code-switches over acquisition for the non-balanced bilingual models. Third, we explored how the patterns of the simulated groups compare to code-switching patterns extracted by two corpora that contain spontaneous utterances from Spanish–English bilingual populations.

Having established that the model reliably produces code-switched sentences, we argue that it can be employed to explain the role of syntax and semantics in specific code-switching phenomena. As a case in point, in the following chapter (Tsoukala, Frank, van den Bosch, et al., 2020) we employed the model to shed light on a well-known effect of verb aspect on Spanish-to-English code-switch probability. The current study's results show that the model can also account for differences in the code-switching patterns between balanced and non-balanced bilinguals.

4

Modeling the auxiliary phrase asymmetry in code-switched Spanish–English

Spanish–English bilinguals rarely code-switch in the perfect structure between the Spanish auxiliary *haber* ("to have") and the participle (e.g., "*Ella ha* voted"; "She has voted"). However, they are somewhat likely to switch in the progressive structure between the Spanish auxiliary *estar* ("to be") and the participle ("*Ella está* voting"; "She is voting"). This phenomenon is known as the "auxiliary phrase asymmetry". One hypothesis as to why this occurs is that *estar* has more semantic weight as it also functions as an independent verb, whereas *haber* is almost exclusively used as an auxiliary verb. To test this hypothesis, we employed a connectionist model that produces spontaneous code-switches. Through simulation experiments, we showed that i) the asymmetry emerges in the model and that ii) the asymmetry disappears when using *haber* also as a main verb, which adds semantic weight. Therefore, the lack of semantic weight of *haber* may indeed cause the asymmetry.

This chapter is based on: Tsoukala, Chara, Stefan L. Frank, Antal van den Bosch, Jorge Valdés Kroff, and Mirjam Broersma (2020). "Modeling the auxiliary phrase asymmetry in code-switched Spanish–English". In: *Bilingualism: Language and Cognition*. DOI: 10.1017/S1366728920000449

4.1 Introduction

Multilingual speakers are able to switch from one language to the other ("code-switch") between or within utterances. For instance, a Spanish–English speaker might produce a sentence such as "*Los niños están* playing in the front yard" ("The kids are playing in the front yard"). As mentioned in Section 1.1, studies have revealed that code-switches do not occur randomly but follow systematic patterns.

The aim of the current work is twofold. First, we present a novel method of researching code-switched sentence production using computational cognitive modeling. To that end, we employ the Bilingual Dual-path model (Tsoukala, Frank, and Broersma, 2017) that can produce sentences in two languages, including code-switched ones. Second, using this method, we shed light on a production phenomenon that has been observed in the Spanish–English-speaking community in the US and is known as the "auxiliary phrase asymmetry" (Dussias, 2003; Guzzardo Tamargo, Valdés Kroff, and Dussias, 2016; Lipski, 1978; Pfaff, 1979; Poplack, 1980). This asymmetry is observed in the frequency of Spanish-to-English code-switches at the participle in progressive and perfect structures. On the one hand, Spanish–English bilinguals rarely produce a code-switch between the Spanish auxiliary *haber* ("to have") and the participle. On the other hand, they are likely, albeit only moderately so, to code-switch in the progressive structure between the Spanish auxiliary verb *estar* ("to be") and the participle. Thus whereas Sentence 1 is attested, Sentence 2 is very infrequent and dispreferred:

1. *Las personas están* voting (The people are voting)

2. *Las personas han* voted (The people have voted)

Furthermore, a switch at the auxiliary (i.e., the first word that is switched is the auxiliary) is approximately equally likely for both structures: "*Las personas* are voting", "*Las personas* have voted".

The auxiliary phrase asymmetry has been found both in speech production and in reading. With respect to speech production, several quantitative analyses of Spanish–English corpora have reported code-switches at the participle in progressive sentences but none in perfect sentences (Lipski, 1978; Muysken, 2000; Pfaff, 1979; Poplack, 1980). Guzzardo Tamargo, Valdés Kroff, and Dussias (2016) performed a systematic analysis on two corpora: the Miami corpus (Deuchar et al., 2014) that contains spontaneous conversations, and a corpus extracted by Guzzardo Tamargo, Valdés Kroff, and Dussias (2016) from entries of an online column in a Gibraltarian newspaper

("Gibraltar corpus"), which contains fictional code-switched written dialogue. As Table 4.1 indicates, although switches are infrequent in either structure, the asymmetry is reported in both corpora.

Table 4.1.: Absolute and relative frequencies in the Miami and Gibraltar corpora. Values reported in Guzzardo Tamargo, Valdés Kroff, and Dussias (2016), Table 1

	Oral corpus (Miami)				Written corpus (Gibraltar)			
	Progressive		Perfect		Progressive		Perfect	
All code-switches	93	100%	28	100%	106	100%	150	100%
Switch at auxiliary	7	7.53%	3	10.71%	8	7.55%	14	9.33%
Switch at participle	7	7.53%	0	0.00%	8	7.55%	1	0.667%

As Dussias (2003) pointed out, because a switch within the auxiliary verb phrase is not a common phenomenon to begin with, it is difficult to obtain sufficient evidence from corpus data to establish whether the reported asymmetry is real. To resolve that, she performed an eye-tracking-while-reading study; she asked English–Spanish bilinguals to read Spanish-to-English code-switched sentences that i) were either in the progressive or perfect structure, and that ii) contained a switch either at the auxiliary (e.g., "*La madre sabe que los chicos* are going to the park"; English translation: "The mother knows that the children are going to the park") or at the participle ("*La madre sabe que los chicos están* going to the park"). The analysis showed that for the perfect-form structures, participants took significantly longer to read switches at the participle compared to switches at the auxiliary, whereas for the progressive-form structures they did not show a significant preference for a switch position. Dussias concluded that auxiliary verb phrase switches are processed differently depending on the structure. Guzzardo Tamargo, Valdés Kroff, and Dussias (2016) also ran an eye-tracking-while-reading study in which they confirmed the results reported in the previous study (Dussias, 2003). Furthermore, Giancaspro (2015) found evidence for the auxiliary asymmetry from an acceptability judgment study. Specifically, he asked English–Spanish bilinguals to rate the grammaticality of code-switches; the participants rejected perfect participle switches and accepted progressive participle ones, thus supporting the asymmetry.

Two (non-mutually-exclusive) explanations have been proposed for this phenomenon; the "grammaticalization account" and the "exposure-based account" (Guzzardo Tamargo, Valdés Kroff, and Dussias, 2016). According to the grammaticalization account, the source of this asymmetry is the difference in semantic weight between the auxiliary verbs. Namely, that *estar* has more semantic weight and is syntactically more independent as it also functions as a linking verb (e.g., "el enfermero está cansado"; "the nurse is tired"), whereas *haber* is highly grammaticalized because it is almost exclusively used as an auxiliary. The verb of possession in Spanish is *tener* ("el enfermero tiene un libro"; "the nurse has a book"), while *haber* is only used as an auxiliary verb or in archaic formulations. The exposure-based account is an alternative hypothesis, which was suggested, but not attested, by Guzzardo Tamargo, Valdés Kroff, and Dussias; it states that the asymmetry emerges from community-supported practice, that is, bilingual speakers learn the asymmetry from exposure to this pattern in the community. In this study we focus on the grammaticalization account to determine whether grammaticalization is a plausible reason why the asymmetry emerged. In human bilinguals, exposure also plays a role, as experience with the language influences production patterns (e.g., MacDonald, 2013).

The grammaticalization account is difficult to test experimentally with psycholinguistic methods, especially in production. Common experimental paradigms for production studies, such as shadowing (where participants repeat stimuli as quickly as possible, e.g. Lipski, 2019) or confederate priming (in which one of the participants is in fact a confederate with a script, who provides primes for the participant, e.g. Kootstra, van Hell, and Dijkstra, 2010), could perhaps confirm the presence of the auxiliary asymmetry. However, to test the grammaticalization account in production, which states the asymmetry is caused by the lack of semantic weight of the Spanish auxiliary verb *haber*, we need to know whether the asymmetry would persist if *haber* did have additional syntactic and semantic functions and was used more frequently, as in the case of the (main and auxiliary) English verb "to have". It is difficult to envision a traditional technique that can test explicitly the role of semantic function of the Spanish auxiliary verb. One could potentially employ artificial language learning that mimics the acquisition and production of code-switched Spanish and English progressive- and perfect-forms. However, this would require a complex setting which would be very challenging for the participants as it would entail advanced learning of the two artificial languages.

Computational cognitive modeling, on the other hand, allows us to make modifications to the vocabularies and the language structures of the modeled languages while keeping everything else the same, thus enabling us to focus on the phenomenon of interest. In this study, we will showcase how we can use computational modeling to add and remove semantic weight from the Spanish auxiliary verb, thus investigating whether the asymmetry could be derived from the properties of Spanish and English. For that reason, we have employed the Bilingual Dual-path model (Tsoukala, Frank, and Broersma, 2017), a connectionist model of bilingual sentence production that produces code-switches.

The model is not exposed to any code-switched sentences; it is only trained on English sentences and Spanish sentences. Therefore, if this particular asymmetry emerges in the model it will be due to the distributional patterns of the two languages (as claimed by the grammaticalization account), in interaction with the properties of the model, and not because of exposure to the asymmetry (i.e., exposure-based account). Such a result would show that the distributional patterns are, in principle, sufficient to lead to the asymmetry.

To investigate whether the asymmetry can emerge in the model and to test the grammaticalization account, we run three sets of simulations. First, we simulate the production of participle switches for the progressive and perfect structures; we hypothesize that this simulation will produce more participle switches in the progressive structure than the perfect one, thus exhibiting the asymmetry. Second, we hypothesize that artificially adding semantic weight to the Spanish auxiliary verb will make the asymmetry disappear because it is caused by the lack of semantic weight of *haber*. We explicitly test this in the second set of simulations by using the Spanish main verb *tener* ("to have") also as an auxiliary verb. Third, we aim to test whether the asymmetry exists due to the relatively low occurrence frequency of *haber*. Namely, because *haber* only functions as an auxiliary verb for the perfect-form sentences, it has lower frequency than the three other auxiliary verbs (be, have, *estar*) that are also used as independent verbs. In the third simulation we correct for this confound.

4.2 Method

4.2.1 Obtaining corpus frequencies

As the occurrence frequency of the two structures of interest could influence the asymmetry, we made sure that we used realistic relative percentages of the progressive and perfect structures for each of the two languages in the Bilingual Dual-path model. To achieve that, we ran a corpus analysis on a Spanish–English corpus. We analyzed the transcriptions of the Bangor Miami corpus (Deuchar et al., 2014)[1] that consists of 30 hours of spontaneous and informal conversations between two or more speakers (84 speakers in total), living in Miami, Florida.

[1] http://bangortalk.org.uk/speakers.php?c=miami

For each of the 56 conversations, we separated the sentences into English only, Spanish only, and code-switched. The corpus is predominantly English. There are 27,835 fully English sentences (61.5% of the whole corpus), 14,631 fully Spanish sentences (32.3% of the corpus), and 2,823 code-switched sentences (6.2% of the corpus).

First, we extracted i) the progressive-form sentences that contain the verb "to be" in the third singular person ("is" for English sentences and "está" for Spanish) followed by a present-tense participle, or by an adverb and a present-tense participle, and ii) the perfect-form sentences that contain the third singular person of the verb "to have" ("has" for English, "ha" for Spanish) followed by an optional adverb and a past-tense participle. From the English sentence candidates that were selected for the progressive form, we excluded the ones that were used to indicate the future form ("is gonna" and "is going to" followed by a verb, e.g., "it is going to rain"). Then, we inspected each extracted sentence manually and further excluded the ones that had been mislabeled (e.g., "disgusting" in "is disgusting" was marked as a participle, not a (participial) adjective, and is therefore not a progressive-form sentence). The results are shown in Table 4.2: for English, the progressive form is about twice as frequent as the perfect form, whereas for Spanish the two forms are more balanced. Furthermore, the simple form is considerably more frequent than the progressive and perfect ones for both languages.

There are enough sentences of each type in the corpus to obtain a somewhat reliable estimate of their frequencies. Therefore, these percentages will be used to generate the corresponding structures of the languages that the model learns (see Section 4.2.3).

4.2.2 Bilingual Dual-path model

Model architecture

The Bilingual Dual-path model[2] (Figure 4.1; Tsoukala, Frank, and Broersma, 2017) is an extension of the Dual-path model (Chang, 2002) of monolingual sentence production[3]. The model is called Dual-path because of its two pathways that influence sentence production: the syntactic path that learns to abstract the syntactic patterns

[2]The model and the full training and test sets and simulation results can be found at `https://osf.io/ba5ru/`.

[3]Independently from our work (Tsoukala, Frank, and Broersma, 2017), the original Dual-path implementation was used by Janciauskas and Chang (2018) to simulate bilingual processing, and more specifically the age of acquisition effects in native Korean speakers of English.

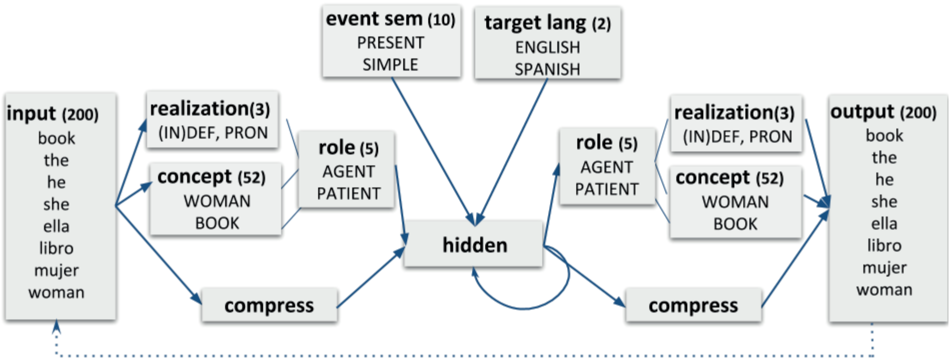

Figure 4.1.: The Bilingual Dual-path model generates sentences word-by-word that express a given message. It is based on a Simple Recurrent Network architecture (the syntactic stream, via the 'compress' layers) that is augmented with a semantic stream (upper path) that contains information about concepts and their realization, thematic roles, event semantics, and the target language. The numbers in the parentheses indicate the size of each layer (e.g., 52 concept units); the sizes of the hidden and compress layers vary with each model run (see Section 4.2.4). The solid arrows denote connections with weights that change during training, whereas the lines between roles, realization, and concepts correspond to connections that are given as part of a message-to-be-expressed (e.g., the AGENT is connected to WOMAN in a particular message). The dotted arrow indicates that once a word is produced, it is given back as input thus contributing to the production of the next word.

Table 4.2.: Absolute frequencies per language in the Miami corpus

Language	Structure type	n
English	Auxiliary Structures	248
	Perfect Present	81
	Progressive Present	167
	Simple Structures	2,451
	Simple Past	722
	Simple Present	1,729
Spanish	Auxiliary Structures	187
	Perfect Present	83
	Progressive Present	104
	Simple Structures	2,959
	Simple Past	629
	Simple Present	2,330

of a language, and the semantic path that receives event semantic information and learns to map concepts onto words. It is a computational cognitive model based on the Simple Recurrent Neural Network architecture (Elman, 1990), and it learns to produce sentences given a message to be expressed (see Section 4.2.3 for an explanation of messages and for an example of how a message is given and is then expressed as a sentence).

We chose to work on and extend the monolingual Dual-path model because the architecture has been employed to explain a wide range of phenomena, for example: structural priming in English (Chang, Dell, and Bock, 2006) and German (Chang, Baumann, et al., 2015), and cross-linguistic differences between English and Japanese (Chang, 2009). Importantly, previous studies using the Dual-path model have focused on semantic weight effects; Chang (2002) simulated different types of aphasia by testing the effect of syntactic-path or semantic-path lesions on production of words with heavy vs no semantic weight (e.g., content words versus function words, heavy verbs vs light verbs).

Sentence production

As mentioned above, the model generates sentences that express a given message. To express a message, the following items are provided to the model and influence production: the to-be-expressed semantic roles (e.g., 'AGENT', 'ACTION') are connected to their concepts (e.g., 'DOG', 'SWIM') and realizations (e.g., 'INDEF' for an indefinite article). The relevant "event semantics" (EVENT-SEM, e.g., 'PRESENT', 'PROGRESSIVE') and "target language" ('ENGLISH', 'SPANISH') units are activated. For instance, the model learns to express the message "AGENT=DOG, DEF; ACTION=SWIM; EVENT-SEM=PRESENT, PROGRESSIVE" in English as "the dog is swimming." and in Spanish as "el perro está nadando."[4]. The language units are included as a means to exert language control: a single language is activated in monolingual contexts and both languages are activated in bilingual contexts.

When the model is given a message, it produces a sentence one word at a time; the produced word ("output") is considered the one with the highest output activation. Each output word is subsequently provided as input in the next time step, and it contributes to the next word production.

Correctness of produced sentences

A produced sentence is regarded as correct if it is grammatical and conveys the target meaning. In some cases, a sentence could be grammatical but incorrect. For instance, if the target sentence is "the mother is pushing a toy" it is grammatically correct to produce the sentence "the tired mother is pushing a toy", even if the meaning is counted as incorrect because of the extra information that was expressed. The same applies to incomplete semantics ("the mother" instead of "the tired mother"); the sentence is counted as incorrect but grammatical.

Code-switching

To allow the model to produce in either language or to code-switch, when testing the model we manipulated the model's language control. Specifically, because we were interested in the Spanish-to-English switch direction, we activated the Spanish language at the beginning, before the production of the first word, so as to indicate the conversational setting[5] (intended language) and to bias towards the production of Spanish utterances. Immediately after the first word was produced, we activated both target language nodes, thus allowing the model to continue in the same language or to code-switch.

4.2.3 Miniature Languages

The sentences that the model learns to produce are derived from miniature versions of natural languages. As we are studying the auxiliary phrase asymmetry which is observed in Spanish–English bilinguals, we focus on Spanish and English sentence production. Hence, we generated sentences based on the relevant properties of the two

[4]All sentences end with a period, even if this is not shown explicitly in the examples.

[5]Note that the sentences are produced independently from one another; the language of the (last word of the) previously produced sentence does not influence the subsequent sentence production.

languages, constrained by the corpus analysis (Section 4.2.1). The advantage of using artificial (miniature) languages is that we can manipulate their structural frequencies, and even grammar, which in turn allows us to isolate and study the phenomenon of interest. For instance, in the case of the auxiliary phrase asymmetry, we can change the frequency and semantic weight of the Spanish auxiliary verb *haber* and see whether the asymmetry persists (see Section 4.3 for an explanation of this process).

Bilingual Lexicon

The lexicon consists of 200 lexical items (Table 4.3): 91 English words, 108 Spanish words, and the shared period ('.') which indicates the end of the sentence. The Spanish lexicon is larger because Spanish is a gendered language. For instance, the adjective 'tall' is either 'alto', if it modifies a masculine noun, or 'alta' for a feminine noun. We also used four common-gendered Spanish adjectives such as 'inteligente' ('intelligent') that do not change depending on the noun it modifies. Note that syntactic category information (such as 'adjective', 'verb') is not given explicitly; the model learns during training that words of the same syntactic category occur in similar contexts.

Table 4.3.: Parts of speech (POS) in bilingual lexicon (Spanish in italics)

POS	n	Examples
Verbs	66	
auxiliary	4	is, has, *está, ha*
intransitive	32	walked, swims, *nada*
transitive	24	carries, pushed, *lleva*
possession	4	has, had, *tiene, tenía*
linking [1]	2	is, *está*
Participles [2]	56	
progressive	28	eating, *comiendo*
perfect	28	eaten, *comido*
Nouns	52	
animate	40	uncle, aunt, *tío, tía*
inanimate	12	pen, book, *libro*
Adjectives	26	busy, *ocupado*
Determiners	6	a, the, *un, una, el, la*
Pronouns	4	he, she, *él, ella*

[1] Both overlap with the auxiliary verbs.
[2] Nine of these have the same form as a verb; e.g., 'walked' is either a perfect participle or a verb.

Structures

For the two languages we used the present and past tense, and three aspects: simple, progressive, and perfect. The past tense is only used in simple aspect sentences (e.g., "the man swam") whereas the present tense applies to all three aspects. The allowed structures for the two languages and all tenses and aspects are Subject - Verb (SV) and Subject - Verb - Object (SVO) (Table 4.4).

Table 4.4.: Allowed structures for English and Spanish

Structure	English example	Spanish example
Present perfect		
SV	she has swum	ella ha nadado
SVO	a man has thrown the key	un hombre ha tirado la llave
Present progressive		
SV	a happy dog is running	un perro feliz está corriendo
SVO	the boy is carrying a book	el niño está llevando un libro
Simple past		
SV	the girl ran	la niña corrió
SVO	he threw a book	él tiró un libro
Simple present		
SV	the grandmother sneezes	la abuela estornuda
SVO	the tall uncle kicks the toy	el tio alto patea el juguete
SVO (linking)	the aunt is focused	la tia está atenta
SVO (possession)	the cat has a ball	la gata tiene una pelota

The grammatical roles can be expressed using either a Noun Phrase (NP) with definite (DEF) or indefinite (INDEF) article (e.g., 'the woman', 'a woman'). Additionally, the subject can be expressed with a pronoun (PRON, e.g., 'she'). NPs optionally contain a modifier (an adjective, e.g., 'a tall woman'). Note that in English the adjective comes before the noun ("the intelligent woman") whereas in Spanish the modifier comes after the noun ("la mujer inteligente"). As mentioned above, the model learns all this through the training examples and not through explicit syntactic labels.

The verbs are either intransitive (e.g., 'swims'), transitive ('carries'), linking ('is', 'está') or possession verb ('has', 'tiene'). The two linking verbs ('is', 'está'[6]) and the English possession verb ('has') were also used as auxiliary verbs for the progressive and perfect forms respectively. As mentioned before, the Spanish perfect-form auxiliary verb is *haber* ('ha' in the 3rd person singular form), which does not function as a main

[6]The Spanish language has two linking verbs (*estar* and *ser*) that are commonly translated as 'to be' in English. For purposes of simplification, in the simulations reported here we have employed only attributes that are expressed with the former linking verb, *estar* (*está* in the 3rd person singular).

verb. Following the allowed structures, each verb had four forms: simple present, simple past, present participle and past participle.

Messages

The goal of the model is to express a specified message using a grammatical sentence, such as the ones described above. A message is represented by (a) a target language, (b) event-semantic information, (c) pairs of thematic roles and concepts, and (d) pairs of thematic roles and realizations (pronoun, definite, indefinite) whenever applicable in the case of noun phrases.

The target languages are ENGLISH and SPANISH. The event semantics contain information regarding the aspect (SIMPLE, PROGRESSIVE, PERFECT) and tense (PRESENT, PAST), as well as the thematic roles that are used in each message.

The following simulations make use of 52 unique concepts and five thematic roles: AGENT, AGENT-MODIFIER, PATIENT, ACTION-LINKING, and ATTRIBUTE. The AGENT is only paired with animate nouns. ACTION-LINKING is a combined thematic role that can be used for all main verb types: action (e.g., 'shows'), linking ('is') and possession ('has'). ATTRIBUTE is an attribute expressed only with a linking verb.

Additionally, AGENT and PATIENT are not only connected to concepts but also to their realization: pronoun (e.g., 'he' for the concept MAN), and definite or indefinite article for concepts that are expressed as a noun phrase (e.g., 'the man' or 'a man' respectively).

Message-sentence pair examples

To incorporate and illustrate all the information given above (Lexicon, Structures, and Messages), here is an example of how a message is expressed as a sentence:

AGENT=WOMAN, INDEF

AGENT-MOD=TALL

ACTION-LINKING=CARRY

PATIENT=BOOK, INDEF

EVENT-SEM=PRESENT, PERFECT, AGENT, AGENT-MOD, PATIENT

The corresponding sentences in English and Spanish are:

- a tall woman has carried a book
- una mujer alta ha llevado un libro (word-by-word translation: "a woman tall has carried a book")

If the aspect was PROGRESSIVE instead of PERFECT, the corresponding sentences would be “a tall woman is carrying a book”; “una mujer alta está llevando un libro”. Similarly, for the SIMPLE aspect, the corresponding sentences are “a tall woman carries a book”; “una mujer alta lleva un libro”.

Linking verb messages are encoded using an attribute:

AGENT=GIRL, DEF

ACTION-LINKING=BE

ATTRIBUTE=TIRED

EVENT-SEM=SIMPLE, PRESENT, AGENT, ATTRIBUTE

and expressed as “the girl is tired” or “la niña está cansada”, depending on the target language.

A message with a possession verb is encoded similar to a message with a transitive verb:

AGENT=GRANDFATHER, DEF

AGENT-MOD=SHORT

ACTION-LINKING=HAS

PATIENT=CHAIR, INDEF

EVENT-SEM=SIMPLE, PRESENT, AGENT, AGENT-MOD, PATIENT

The preceding message would be expressed as “the short grandfather has a chair” and “el abuelo bajo tiene una silla”.

Note that auxiliary verbs do not have an explicit concept and are not assigned to a thematic role (e.g., for “has carried” the ACTION-LINKING is CARRY). When “have” (or “tiene”) is used as a possessive verb, as in the example above, the ACTION-LINKING is HAS which indicates that the verb has a semantic property. The Spanish auxiliary verb *haber* does not function as a main verb and, therefore, has no semantic weight in the model: it is never connected to ACTION-LINKING. This is in contrast to the verbs “have”, “is”, “está”, which have semantic weight (through a connection to ACTION-LINKING) when used as main verbs.

4.2.4 Model training

Connectionist models have trainable connection weights that are adapted during the learning process. During training, the network sees examples (targets) of messages and sentences (Section 4.2.3). Before training, the network produces random words (e.g., "cat cat cat"). After each word has been produced, the model receives feedback as to whether the word was correct or not, and the connections change their weights depending on the mismatch between the produced and the target word. Gradually, through exposure to the examples, the model learns to successfully express a message using the corresponding sentences.

A set of training examples always contained 2,000 message-sentence pairs. All three simulations (see Section 4.3) were trained for 40 epochs, where each epoch corresponds to one pass through the training set. The connection weights were updated using the backpropagation algorithm after each output word.[7]

A number of random factors influence the message-to-sentence production and the overall performance of the model. In order to minimize the risk of choosing parameters that are either too specific (i.e., resulting in an effect that does not generalize) or improper (i.e., causing failure or low performance in the overall sentence production), we decided to train several networks per simulation. Specifically, for each simulation we trained 60 different networks (model runs) for 40 epochs while randomizing all free parameters per network, as explained below, except the training set size and backpropagation parameters.

The target message-sentence pairs (see Section 4.2.3 for examples of messages) are randomly generated before the training starts, and the sentences are constrained by the set of allowed structures (Section 4.2.3). For each part of speech (POS) a randomly selected lexical item (from that POS and target language) is sampled from the bilingual lexicon (Section 4.2.3).

As mentioned in the section on corpus analysis (Section 4.2.1), the simple tense occurs considerably more frequently than the progressive and perfect-form constructions. Since in the current simulations we are mainly interested in the latter two forms, to ensure the model encounters these forms sufficiently we increased their percentage in the training set by downsampling the simple form. At the same time, we made sure to keep the relative frequencies of the progressive and perfect forms intact. The percentages used to produce structures in the model are shown in Table 4.5.

Additionally, the percentage of English and Spanish varied slightly: the goal was to simulate balanced bilinguals, but as it is very unlikely that a human bilingual receives truly balanced input, we sampled the percentage of English using a normal

[7]Backpropagation (Rumelhart, Hinton, and Williams, 1986) is a learning algorithm typically used in neural networks. In our simulations, the momentum was set to 0.9 and the initial learning rate was 0.10 and linearly decreased after each training sample over 10 epochs until it reached 0.02.

Table 4.5.: Structure frequencies in the model training sentences

Language	Structure type	Percentage
English	Auxiliary Structures	62%
	Perfect Present	20%
	Progressive Present	42%
	Simple Structures	38%
	Simple Past	6%
	Simple Present	32%
Spanish	Auxiliary Structures	62%
	Perfect Present	27%
	Progressive Present	35%
	Simple Structures	38%
	Simple Past	6%
	Simple Present	32%

distribution with a mean of 50% and a standard deviation of 8, the rest being Spanish. Importantly, the target sentences were never code-switched.

Furthermore, the network's connection weights were randomly initialized from a normal distribution centered at zero, and the (non-trainable) weights of the connections between the thematic roles and the concepts ('concept'–'role' and 'predicted role'–'predicted concept', see Figure 4.1) were integer values that were sampled for each simulation between the values of 10 and 20, whereas the unused roles and concepts are not connected. These connections are not trained. The size of the hidden layer was also random, between 90 and 110 units, and the size of the compress layer was set to roughly 77% of the size of the hidden layer.

4.3 Simulations

To test the grammaticalization account of the auxiliary phrase asymmetry, we ran three sets of simulations, consisting of 60 model runs each.

4.3.1 Simulation 1: "haber model"

In the first set of simulations, we tested whether the auxiliary phrase asymmetry can emerge only from the distributional patterns of the two languages, which would indicate that exposure to the asymmetry is not necessary to explain the phenomenon. To test that, we trained the model ("haber model") on 2,000 sentence-message pairs using the generated examples described in Section 4.2.3 that contain progressive-, perfect- and simple-tense sentences. We then tested it on 700 novel messages that had Spanish as a target language. Only Spanish messages are included because we are interested in Spanish-to-English code-switches; therefore, we activated the Spanish language unit until the first word was produced, after which both languages were activated, allowing the model to continue in the same language or to code-switch. Of these 700 messages, 350 had a PROGRESSIVE aspect (e.g., “the boy is kicking a ball”) and 350 were the PERFECT-form equivalent of those sentences (“the boy has kicked a ball”). We hypothesized that the model would display the auxiliary phrase asymmetry even though the phenomenon was not present in the training data; as mentioned before, the model was not exposed to any code-switched sentences during training.

4.3.2 Simulation 2: “tener model”

In the second simulation we tested explicitly whether adding semantic weight to the Spanish auxiliary verb (i.e., *haber*) would diminish the asymmetry. This was done by taking advantage of the fact that the model’s training set is generated and can therefore be manipulated. To increase the semantic weight of the Spanish auxiliary verb, we modified the Spanish main verb (*tener*) of the model to function both as a main verb and an auxiliary (i.e. similar to English, which uses the verb “to have” both as a main verb and as an auxiliary verb: “the boy has a dog”; “the boy has left”). More specifically, we trained the model (“tener model”) with the same training examples as in the “haber model” simulations, with the only difference that we replaced all instances of *haber* with *tener*. For instance, “el niño ha comido” (i.e., “the boy has eaten”) became “el niño *tiene* comido”. We kept everything else the same as in the “haber model” (the 700 test messages, initialized weights, all the layer sizes, and even the lexicon size, even though *haber* was no longer used, were identical), and we ran 60 networks using the modified target sentences. Because in this simulation *tener* is used both as an independent main verb with semantic content and as an auxiliary verb, we hypothesized that this model would not show the asymmetry.

4.3.3 Simulation 3: “synonym model”

Finally, we ran a third simulation to control for the frequency increase of *tener* which was caused by using it also as an auxiliary verb in the previous model compared to the first model. To make sure that a potential disappearance of the asymmetry in the “tener model” is not simply due to the increase of exposure of *tener*, we made *haber* and *tener* perfect synonyms ("synonym model"). Either could be used as a main verb or an auxiliary verb, whereas we kept the frequency of each verb the same as in the "haber model". Once again, we also kept everything else the same as in the "haber model". If the "synonym model" shows the asymmetry, this will indicate that the asymmetry in the "tener model" disappeared because of the increase in the auxiliary verb frequency and not because of the added semantic weight.

4.4 Results

All three simulations achieved a similar performance on the progressive and perfect test sentences (see Table 4.6). All 1000-sample bootstrapped 95% Confidence Intervals reported in this table and the following figures were calculated over the percentages from each of the 60 model runs.

Table 4.6.: Performance (percentage of test sentences with correct meaning) of the three models at the last epoch. The numbers in the brackets show the bootstrapped 95% Confidence Interval.

	Progressive	Perfect	Total (average)
haber	89.5% [85.4, 91.6]	87.1% [83.1, 89.3]	88.3% [84.2, 90.5]
tener	90.9% [88.0, 92.3]	89.8% [87.9, 91.2]	90.4% [88.1, 91.7]
synonym	90.8% [88.7, 92.3]	87.2% [84.5, 89.1]	89.0% [86.8, 90.7]

For each simulation, we ran a logistic mixed-effects regression analysis comparing the percentage of progressive and perfect switches at the last (40th) training epoch; the analyses include a by-model-run random intercept and a slope of sentence structure (coded as $-.5$ for the perfect structure and $+.5$ for the progressive one), but no by-sentence random effects because test sentences differ between model runs. Comparisons between simulations were performed by including the factor simulation (dummy coded with "haber model" as the reference level) and the interaction with sentence structure; this analysis also included a random slope of simulation[8].

4.4.1 Simulation 1: "haber model"

Importantly, the "haber model" clearly produced the auxiliary phrase asymmetry, as hypothesized. Figure 4.2a shows the average percentage of Spanish-to-English

[8] All regression analyses were performed using the `R` package `lme4` (Bates et al., 2015). The exact script can be found at `https://osf.io/yn8e9/`

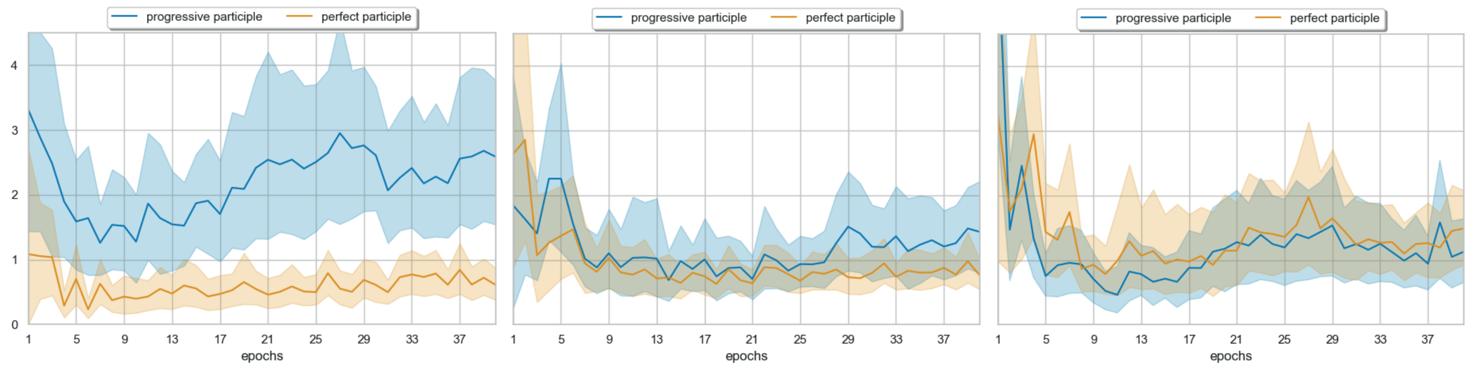

a: haber **b:** tener **c:** synonym

Figure 4.2.: Percentage of Spanish-to-English participle switches (computed over 60 network runs per simulation and over the course of network training) of the correctly produced sentences per structure in the three simulations. Shaded areas show the bootstrapped 95% Confidence Interval.

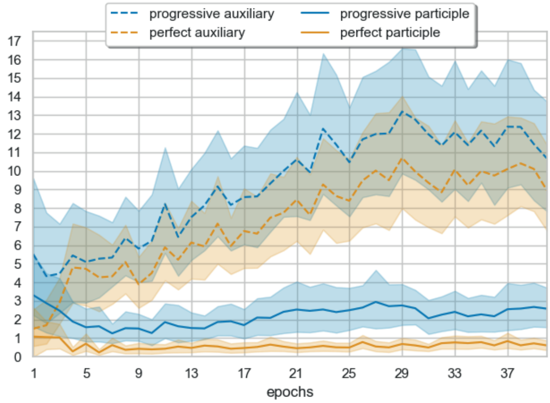

Figure 4.3.: Percentage of code-switches at auxiliary and participle for the progressive and perfect structures in the "haber model", computed over 60 network runs. Shaded areas show the bootstrapped 95% Confidence Interval.

participle switches over 60 model runs for correctly produced sentences per structure (progressive and perfect). The model output showed a strong preference for progressive participle switches over perfect participle switches: at the last training epoch, 2.34% of all correctly produced progressive-form sentences had a switch at the participle whereas only 0.60% of the correctly produced perfect-form sentences had a switch at the participle. The logistic mixed-effects regression analysis showed the difference to be statistically significant[9] $(b = 1.13; z = 3.25; p < .002)$.

Figure 4.3 shows the percentages of code-switches at the auxiliary verb and at the participle for the progressive and perfect structure. A probability of a switch at the auxiliary verb was not significantly different between structures, and a participle switch for the perfect structure was the least preferred switch point. For both structures, the simulations showed a clear preference for a switch at the auxiliary position over the participle one (10.77% progressive-auxiliary switch vs 2.34% for progressive-participle switch and 9.01% perfect-auxiliary switch vs 0.60% for progressive-participle switch). However, the probability of a switch at the auxiliary was not significantly different between the progressive and perfect structures $(b = 0.11; z = 0.93; p > .3)$.[10]

[9]This effect was not caused by the higher frequency of progressive relative to perfect structures; it also appeared when the two structures occurred with equal frequency (Tsoukala, Frank, van den Bosch, et al., 2019).

[10]The switch location (auxiliary vs participle) is not an independent variable; therefore, there is no interaction between location and structure to be tested. However, whether the effect of structure differs between locations can be ascertained by comparing the confidence intervals around structure effect

4.4.2 Simulation 2: "tener model"

When tested on the same 700 messages, the auxiliary phrase asymmetry all but disappeared. The output of the "tener model" (that substituted the original Spanish auxiliary verb from the "haber model" for one with more semantic weight) showed at best a small (non-significant) preference for progressive participle switches over perfect participle switches (1.30% vs 0.77% in the last epoch; Figure 4.2b). This difference is not statistically significant ($b = 0.10; z = 0.37; p > .7$) and is significantly smaller than in the "haber model" (interaction between sentence structure and simulation: $b = -1.05; z = -5.26; p < .0001$).

4.4.3 Simulation 3: "synonym model"

Finally, when tested on the same 700 messages, the "synonym model", which controlled for the frequency increase of the auxiliary verb in the "tener model", did not show a preference for progressive over perfect participle switches either (1.13% progressive-form participle switches vs 1.03% perfect-form participle switches; Figure 4.2c). This difference is not statistically significant ($b = -0.48; z = -1.30; p = 0.19$) and is significantly smaller than in the "haber model" ($b = -1.62; z = -8.75; p < .0001$).

4.5 Discussion

All three simulations support the hypothesis that the asymmetry can be caused by the lack of semantic weight of the Spanish auxiliary verb *haber* "to have". The "haber model" exhibited the auxiliary phrase asymmetry; adding semantic weight to the Spanish perfect-form auxiliary verb ("tener model") was enough to make the asymmetry all but disappear. If the reason that the effect became significantly smaller in the "tener model" was the frequency increase of the auxiliary verb, we would have expected the "synonym model" to produce a similar pattern to the "haber model". In the third simulation, when controlling for the increase in the frequency of the Spanish auxiliary in the second simulation, the "synonym model" did not show a preference for progressive participle switches either, thus further supporting that the lack of semantic weight of *haber* can cause the asymmetry.

The probability of a switch at the auxiliary verb for the progressive structure did not differ significantly from the probability of a switch at the auxiliary for the

sizes. These are (-0.12, 0.35) and (0.45, 1.82) for auxiliary switches and participle switches, respectively, indicating that the effect sizes differ between locations.

perfect structure. Furthermore, a participle switch for the perfect structure was the least preferred switch point; both patterns reflect prior experimental and corpus-based results (Guzzardo Tamargo, Valdés Kroff, and Dussias, 2016). However, unlike these corpus results, in the progressive structure the simulations showed a clear preference for a switch at the auxiliary position. This indicates that the model does not capture the finding that in the corpora (Table 4.1) there is no preference between an auxiliary and a participle switch in the progressive structure. This discrepancy between model results and corpus data could be attributed, for instance, to the limited number of structures used in the simulations.

In this study we have only focused on the grammaticalization account. We speculate, however, that even though in bilinguals the asymmetry is likely driven by grammaticalization, the overall switching patterns are reinforced by exposure to code-switched speech in the community (as claimed by the exposure-based account discussed in the introduction). In principle, the model should be able to simulate exposure-based explanations as well, by running second-generation simulations that receive as target the code-switched sentences of the first simulations. We expect that in this scenario the amount of overall code-switching will increase. Furthermore, we hypothesize that the perfect-form participle-switch will become even less frequent over time.

Previous literature on code-switching (e.g., Pfaff, 1979 and Poplack's 1980 Equivalent Constraint) has argued that code-switching can only occur at points where the surface word order of the two languages is the same. Similarly, grammatical constraints on code-switching from a generative framework (e.g., Functional Head Constraint, Government Constraint) likewise make broad generalizatons on which syntactic junctures code-switches cannot occur (i.e., no code-switches between an auxiliary verb and a main verb). The findings reported here, as well as in Dussias (2003) and Guzzardo Tamargo, Valdés Kroff, and Dussias (2016), have shown that such constraints are not enough to explain the code-switching patterns; the auxiliary phrase asymmetry seems to exist because of differences in the semantic weight of the auxiliary verbs, despite the fact that the structures have the same syntactic patterns. Therefore, it seems that even though syntax and word order play a role in the places where code-switching can occur, they are not the only factors governing code-switching.

As is the case with every research method, using computational cognitive modeling has certain limitations. First, the languages used in the simulations are miniature, and therefore artificial, which could be seen as a disadvantage of this research method. However, using a miniature language is also an advantage as it allows us to remove any confounding factors and to focus on the phenomenon of interest. More importantly, it gives us the unique opportunity to manipulate the languages that the model learns (i.e., the lexicon and/or structures) and to investigate whether changes in the language lead to different code-switching patterns. For instance, in this paper we show that the asymmetry disappears when the Spanish auxiliary verb *haber* is the same as the main verb “to have” and therefore has semantic weight like its progressive-form auxiliary verb equivalent. Second, connectionist models can produce different patterns depending on the generated training sentences and the connection weights that are assigned before the learning starts. It is therefore important not to report on a single model run, nor to hand-pick free parameters that produce the results we would like to see. In this work, we showed that the result is robust by running three separate simulations using 60 networks for each while randomizing (within fixed ranges) all free parameters and the target message-sentence pairs.

The simulations’ goal was to test whether the lexical-syntactic distribution patterns of Spanish and English can lead to the auxiliary phrase asymmetry, and the simulations have provided evidence that it can indeed. We do not claim to (and did not aim to) know what mechanism drives the asymmetry in the models’ output. We speculate, however, that having a higher semantic weight, as in the case of the Spanish and English verbs “to be” (“is”, “está”), leads to more possible upcoming words. Subsequently, this leads to the activation of more output word candidates; the most activated word is less reliably the correct Spanish word, thus increasing the probability that the most activated word is the English translation. The Spanish auxiliary verb *haber*, on the other hand, only has one sense and a very restricted context, as it can only be followed by a participle, thus allowing for fewer options to code-switch.

4.6 Conclusion

We tested whether the auxiliary phrase asymmetry in Spanish–English code-switching could be derived from the properties of the two languages. The “haber model” simulated the attested asymmetry, and the “tener model” showed that this could be attributed to the fact that the Spanish auxiliary *haber* has only a limited, dependent syntactic function (i.e., is more grammaticalized) and is not used as frequently as the English equivalent (“have”). The follow-up model (“synonym model”) used *haber* and *tener* as synonyms, and confirmed that the lack of asymmetry in the “tener model” can be attributed to the syntactic independence of the modified auxiliary verb and not to its increased frequency compared to the “haber model”. The three simulations thus confirm that grammaticalization could be responsible for the asymmetry.

Importantly, we showed that using computational cognitive modeling we can test hypotheses that cannot be experimentally tested in humans, such as changing the function of the Spanish auxiliary verb and observing whether the auxiliary phrase asymmetry persists. We made the grammaticalization hypothesis testable in the model, and showed that indeed the difference in semantic weight between *estar* and *haber* can cause the observed phenomenon, in line with the grammaticalization account.

5

Simulating the effect of cognates on code-switching

Cognates are words that share similar meaning and form between languages, such as 'actor' or 'artificial' in English and Spanish. Clyne (1967; 2003) suggested that cognates are trigger words that facilitate code-switching; this is called the cognate triggering effect. In this chapter, we employ the Bilingual Dual-path model to investigate the cognate triggering effect on balanced and non-balanced (L1 Spanish, L1 English) bilinguals, as well as the effect that the percentage of cognates in a language pair has on code-switching for the Balanced simulation. The simulations show a cognate triggering effect on the L1 English non-balanced simulation, but not on the Balanced and L1 Spanish non-balanced simulations. Additionally, the percentage of cognates in a language pair does not affect the model's code-switching propensity.

5.1 Introduction

Clyne (1967; 2003) suggested that certain words ("triggers") facilitate code-switching. Specifically, he hypothesized that words that have similar meaning and overlapping form, such as cognates[1] are trigger words. Clyne studied the speech patterns of immigrant populations (from the Netherlands, Germany, Croatia, Hungary, Spain, Italy, and Vietnam) in Australia and provided several such anecdotal examples from these bilingual (or multilingual) populations. The code-switching facilitation caused by cognate words is called the triggering (a.k.a. facilitation) effect. In the current chapter, we will investigate the cognate triggering effect in balanced and non-balanced Spanish–English bilinguals, as well as the effect that the percentage of cognates in a language pair has on code-switching, using the Bilingual Dual-path model.

Broersma and De Bot (2006) statistically tested the triggering hypothesis on a small Moroccan Arabic–Dutch corpus, consisting of transcriptions of informal conversations between three bilingual speakers living in the Netherlands. This study found that clauses containing cognates were code-switched 29.6% of the time as opposed to 12.6% of the time for clauses not containing such words. This work was the first to statistically support the triggering hypothesis in a language pair that is not related and, therefore, does not contain many cognate words; only 4.7% of the word tokens (i.e., 104 out of 2,224) were cognates, with most cognates being (proper) nouns. In a later study, Broersma (2009) ran a corpus analysis on an interview conducted with a Dutch–English late bilingual speaker who migrated to New Zealand in her late 30s. Dutch–English is a highly cognate language pair; in the transcription of the interview, 71.4% of word tokens (2,035 out of 2,849) were cognates (proper nouns, content words, and function words). In this corpus analysis, the findings were similar to Broersma and De Bot (2006): 37.5% of the utterances containing a cognate word were code-switched, as opposed to 6.9% of utterances without a cognate word. Broersma, Isurin, et al. (2009) replicated this observation in a corpus study including both a typologically similar and a less related language pair (Dutch–English and Russian–English).

Following Broersma and De Bot (2006), Soto, Cestero, and Hirschberg (2018) ran a corpus analysis on the Bangor Miami corpus (Deuchar et al., 2014), which is the same corpus that we employed in Chapters 3 and 4. The Bangor Miami corpus consists of

[1] Strictly speaking, cognates are words that share a common ancestry, such as father-vader in English and Dutch. However, in psycholinguistic studies the focus is on the *perceived* similarity of two word forms and not on the shared origin. Therefore, proper nouns ("Maria") and loan words (such as "internet") would count as cognates.

242,475 words, transcribed from informal conversations between 84 Spanish–English speakers. The authors collected a list of (not necessarily identical[2]) cognate words and identified 1,305 unique cognate pairs (types) in the Miami corpus. They then compared utterances with and without cognates with respect to whether they were code-switched or not. A statistical analysis confirmed that code-switched sentences were more likely to contain a cognate, thus supporting the triggering hypothesis.

Broersma, Carter, Donnelly, et al. (2020) explored the triggering effect in a different language pair (Welsh–English) in the Bangor Siarad corpus (Deuchar et al., 2014). The Bangor Siarad corpus is even larger than the Bangor Miami corpus: it contains 447,507 words from 69 informal conversations and a total of 151 speakers. Broersma, Carter, Donnelly, et al. confirmed the triggering hypothesis in this corpus as well: the odds of producing a code-switch were 65% higher when a cognate was present in the sentence compared to sentences that did not contain a cognate.

A number of psycholinguistic experiments have also been conducted to explore the triggering hypothesis, though the results are less conclusive than in the corpus analyses. Broersma (2011) ran a picture-naming task, with color cues indicating the intended language (Dutch or English). Preliminary results suggested that the response time for code-switches following a cognate word were shorter than switches after a non-cognate word, thus supporting the triggering hypothesis. However, further analysis of the data found no such evidence (Broersma, Carter, Donnelly, et al., 2020). In a Dutch–English structure priming experiment, Kootstra, van Hell, and Dijkstra (2012) found that high-proficient bilinguals tended to produce more L1-Dutch to L2-English intra-sentential switches when the prime sentences contained a cognate compared to non-cognate prime sentences; this effect was not found in low-proficient bilinguals. The authors did not investigate the L2-to-L1 switch direction.

There are studies, on the other hand, that found no evidence for the triggering hypothesis. For instance, Bultena and colleagues tested whether the presence of a cognate verb facilitates the comprehension (Bultena, Dijkstra, and van Hell, 2014) and production (Bultena, Dijkstra, and van Hell, 2015) of intra-sentential language switching, both from L1 (Dutch) to L2 (English) and from L2 to L1. The presence of a cognate verb did not have an effect in either experiment or switch direction. Santesteban and Costa (2016) ran a picture naming task with a color cue. They tested two groups of bilinguals based on two proficiency levels: L1 Spanish L2 Catalan (low-proficient) bilinguals and balanced (high-proficient) Spanish-Catalan bilinguals. The results showed a non-significant trend of cognate triggering in the high-proficient

[2]Non-identical cognates do not have the same form; for instance, direction and dirección in English and Spanish.

group, and no effect nor trend in the low-proficient group. The authors concluded that the cognate status of the target or preceding word does not affect the response times in either bilingual group. Last, Broersma, Carter, and Acheson (2016) ran a picture naming study with Welsh–English bilinguals, who were either balanced, slightly English-dominant, or slightly Welsh-dominant. The balanced bilingual and Welsh-dominant participants showed cognate triggering effects, whereas the English-dominant participants showed cognate inhibition when naming in Welsh, and no difference between cognates and non-cognates when naming in English.

In the current chapter, we will investigate the cognate triggering effect of noun concepts using the Bilingual Dual-path model. We decided to focus only on nouns because it is the most documented word type in the aforementioned studies. Because the model simulates sentence production, rather than a specific task such as picture naming[3], or forced code-switching in picture description, we hypothesize that the results will be similar to those of the corpus analyses (Broersma and De Bot, 2006; Broersma, 2009; Broersma, Isurin, et al., 2009; Broersma, Carter, Donnelly, et al., 2020; Soto, Cestero, and Hirschberg, 2018) and we expect that code-switching will occur more often in sentences that contain a cognate. Additionally, we will explore i) whether there is a different triggering effect in balanced versus non-balanced Spanish–English bilingual models, and ii) whether the cognate percentage of a language pair affects the overall probability of code-switching or the effect of cognates in the sentence production of balanced bilinguals. The latter goal has both a methodological motivation, i.e., to ensure that the proportion of cognates we selected for the simulations did not significantly affect the results, and an exploratory theoretical one; to our knowledge, no study has investigated the effect that the proportion of cognate words in a language pair has on code-switching.

5.2 Cognate comparison setup

5.2.1 Bilingual Dual-path model: languages

For this study, we trained the Bilingual Dual-path model using structures similar to the ones used in previous chapters, especially Chapter 3 (see Table 3.1). We were interested in the cognate status of the words and not in the interaction between specific syntactic structures, therefore we only kept one variation of a structure expressing

[3]In fact, because the semantics are given to the model, the process is more similar to a picture-description task. However, because each sentence is independent from the previous one, there is no forced language (switching) cue or any other task-related cognitive control needed in the model.

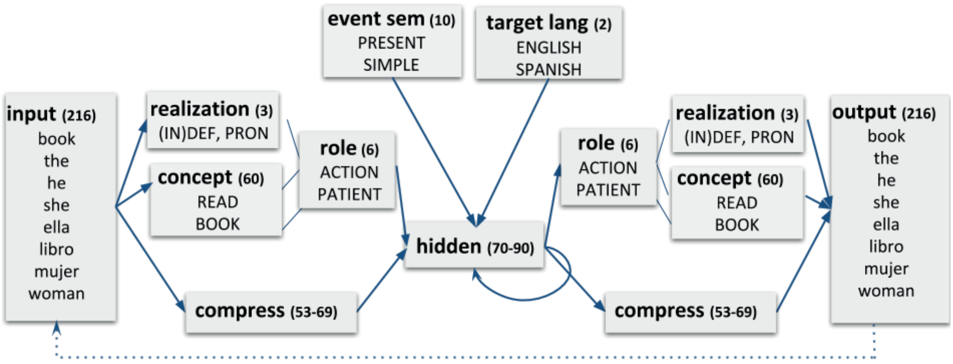

Figure 5.1.: The Bilingual Dual-path model with the layer size information within parentheses.

a RECIPIENT: the prepositional dative in which the RECIPIENT is expressed at the end of the sentence (e.g., "the waitress gave a toy to the girl" instead of "the waitress gave the girl a toy"). We also edited the lexicon by adding 18 more nouns (9 in each language) compared to the lexicon used in Chapter 3 because we are focusing on noun cognates only and we wanted to ensure that there is a large number of nouns in the language pairs. The exact layer sizes can be seen in Figure 5.1; the structure and lexicon files can be found at `https://osf.io/svb8d/`.

As in the previous chapters, we simulated a miniature version of Spanish and English. The percentage of cognates in this language pair depends on the domain, but is expected to be at least 30% (Moss, 1992). For this reason, in the following simulations the default percentage of cognates we used was 30%. However, to explore the effect of the percentage of cognates in a language pair, we ran additional experiments including simulations that had been trained with up to 70% of cognates (see subsection 5.3.2).

5.2.2 Pairwise cognate experiment setup

The artificial languages allow for a clean pairwise comparison between two identical models that have been trained with different cognate sets (hereby: "cognate1" and "cognate2") that contain the same percentage of cognates. In cognate1, 30% of the concepts referring to noun entities (across both languages) were randomly selected and converted to cognates by assigning the English label to both the Spanish and the English word (e.g., the word for the concept DOG became "dog" in Spanish and

English). In cognate2, 30% of non-overlapping (with cognate1) noun concepts were selected and converted to identical cognates.

As we were interested in the effect of cognates on code-switching, using different cognate sets we compared the probability of code-switching in a sentence in which a noun is a cognate versus an identical sentence in which the same concept is not a cognate. The pairwise comparisons were done using a test set containing messages with exactly one cognate, that was a cognate for either cognate1 or cognate2. For a test set sentence with a concept that was a cognate in cognate1, cognate1 was marked as "cognate" and it was compared against cognate2's identical sentence ("non-cognate"), where the same concept was not a cognate. Similarly, if a test set concept was a cognate in cognate2, this network run was marked as "cognate" and was compared against the "non-cognate" cognate1. For instance, consider the following message[4]:

AGENT=ROOSTER, DEF; ACTION=PUSH; PATIENT=BOTTLE, INDEF; EVENT-SEM=SIMPLE, PAST; TARGET-LANG=ES

If ROOSTER is a cognate in cognate1 the sentence would be expressed in (non code-switched) Spanish as "el rooster empujó una botella" whereas in cognate2 (in which ROOSTER is not a cognate) as "el gallo empujó una botella" ("the rooster pushed a bottle" in English).

When generating test messages for either a pairwise or non-pairwise comparison, we imposed two restrictions: i) the cognate should not be expressed as a pronoun and ii) the message should only contain one cognate; it should not contain any other concept that is a cognate in either cognate set. For instance, if in the sentence "the waitress gave the man a cup" the cognate word is "waitress", but another word (e.g., "cup") is also a cognate in either cognate1 or cognate2, this sentence is excluded from the test set.

5.3 Experiments

5.3.1 Balanced versus non-balanced models

As in the previous chapters, the balanced bilingual model was trained simultaneously (for 20 epochs in the current study) on both languages, i.e., with Spanish input randomly sampled using a normal distribution with a mean of 50% and a standard deviation of 8, and the rest in English. Before the training started, 30% of the concepts were converted to cognates (i.e., the Spanish and English form of a concept were

[4]See previous chapters for a thorough explanation of messages; e.g., in section 3.2.2. of Chapter 3.

identical). To run a pairwise comparison, for each cognate set (cognate1 and cognate2) we trained 40 model runs with randomly initialized weights (which were the same for cognate1 and cognate2) and training message-sentence pairs, as done in the previous chapters, and randomly generated cognate sets.

The non-balanced models were first trained for 10 epochs on the L1 only, with the same cognates as in the Balanced model, before it was exposed for another 20 epochs to the bilingual Spanish–English input that was generated to train the balanced bilingual model. The non-balanced models also consisted of 40 model runs per cognate set (cognate1 and cognate2) with the same randomly initialized weights as the Balanced model.

We compared the balanced and non-balanced models using a test set with 600 messages, containing one cognate per message, that had the same distribution of syntactic structures as in the training set. We analyze the effect of cognates on code-switching, both in total and per switch direction for the non-balanced simulations (L1-to-L2 and L2-to-L2).

5.3.2 Effect of percentage of cognates in language pair

In the previous simulations we assigned 30% of noun cognates, because this percentage is realistic for the Spanish–English language pair. To explore the effect of the percentage of cognates in a language pair, and to ascertain that the percentage we selected for the previous simulations did not significantly affect the results, we ran two sets of experiments (a pairwise and a non-pairwise one) on Balanced bilingual models with different percentages of cognates.

Pairwise experiment

First, we ran three pairs of cognate models containing 10, 15, 20, 25, and 30% of nouns that are cognates. We evaluated the three pairs using a single test set: the set consisted of 600 messages, in the same distribution of structures as the training set, containing one cognate per messages from a pool of cognates seen by all pairs (i.e., the cognates selected for the 10% models, minus one cognate that had not been seen by the other four model pairs). The concept could belong to either cognate1 or cognate2.

Non-pairwise experiment

With the pairwise setup we could only test up to 30% of cognates, because for each percentage two cognate sets had to be generated (e.g., for the 30% cognate model, 60% of concepts were converted to cognates; half in cognate1 and the other half in cognate2). Given the additional restriction that each test sentence should contain only one cognate concept, generating 600 test sentences for each of the 40 model runs for a larger percentage of cognates became unfeasible.

To investigate the effect of the percentage of cognates up to 70%, we also compared cognate groups consisting of up to 70% of cognates (0, 10, 20, 30, 40, 50, 60, and 70% of cognates). The difference between this setup and the pairwise comparison is that we did not compare the current cognate set to a "non-cognate" equivalent.

We tested all models using two test sets. First, a test set ("cognate sentences") that was similar to the pairwise-comparison test set, i.e., that contained exactly one cognate per sentence from the ones experienced by the 10% model. The only difference was that, when excluding cognates (that were unique to the 20-70% cognate sets) from the test set to only keep one cognate per message, the excluded cognates were slightly different than the ones excluded in the pairwise experiment generation, thus resulting in different sentences. Testing on the 0% model was not possible because, as the name suggests, it has no cognates. Therefore, we tested only 7 models (10-70% cognates) on this test set.

Second, we evaluated the same models on a test set not containing any cognates ("non-cognate sentences"), which means that we generated sentences without any concept that was a cognate in any of the 7 models. For both test sets, we only kept the sentences that were correctly produced by all models (7 for the "cognate sentences" set, 8 for the "non-cognate sentences").

5.4 Results

5.4.1 Balanced versus non-balanced models

Figure 5.2 shows the code-switching percentages out of sentences that have been correctly produced for all three bilingual simulations (Balanced, L1 Spanish, L1 English) for the 30% cognate and non-cognate sets (total: 5,922 per simulation and set). In line with results from Chapter 3, the non-balanced models code-switch much less than the balanced bilingual model.

To investigate whether there is an overall cognate effect, we ran a logistic mixed-effects regression analysis (Bates et al., 2015) comparing the probability of a code-switch at the last training epoch (i.e., the 20th for the balanced simulation and 30th for the non-balanced simulations) between the cognate and non-cognate models. The analysis includes a by-model-run random intercept and a slope of the cognate status and the simulation type of the model, coded as $-.5$ for non-cognate models and $+.5$ for cognate models, and $-.5$ for non-balanced simulations (L1 English, L1 Spanish) and $+.5$ for the balanced simulation. In all statistical analyses reported here there was no by-sentence random effect because each of the 40 model runs is exposed to different test sentences. The analysis showed an overall cognate effect ($b = 0.12; z = 2.07; p = 0.007$). Additionally, it showed a significant interaction between the cognate effect and the simulation type, indicating that the cognate effect is stronger for the non-balanced models ($b = -0.17; z = -2.49; p = 0.013$). To explore the cognate effect per simulation, we ran three separate analyses, one per simulation, including a by-model-run random intercept and a slope of the cognate status (coded $\pm.5$ as before). The analyses showed no cognate effect for the Balanced simulation ($b = 0.03; z = 0.59; p = 0.55$) nor for the L1 Spanish simulation ($b = 0.05; z = 0.46; p = 0.65$). It did, however, reveal significance for the L1 English simulation ($b = 0.26; z = 3.62; p = 0.0003$). This indicates that for the L1 English simulation, there was a higher probability of code-switching with the cognate than the non-cognate set.

In this section as well as the following ones, we chose to analyze only sentences that had been correctly produced by all (six, in this case) simulations and cognate sets. To ensure that the above pattern does not emerge due to the strict sentence selection, we also analyzed the full set of correctly produced sentences (34,484 for the Balanced simulation, 29,606 for the L1 English simulation, and 30,114 for the L1 Spanish one) and the pattern of statistical significance remained the same. Therefore, we only report here the results based on the identical sentences, so as to ensure a perfect comparison between simulations.

Non-balanced models: cognate in the L1

Figure 5.3a shows the percentage of L1-to-L2 code-switches for the non-balanced bilingual models. Note that the percentages are different than in Figure 5.2 that contains all three simulations because the code-switched percentage plotted here is against all correctly produced L1 sentences only. Additionally, because in this case the L1 sentences are different depending on the simulation (English for the L1 English, Spanish for the L1 Spanish simulation) we cannot restrict the comparison

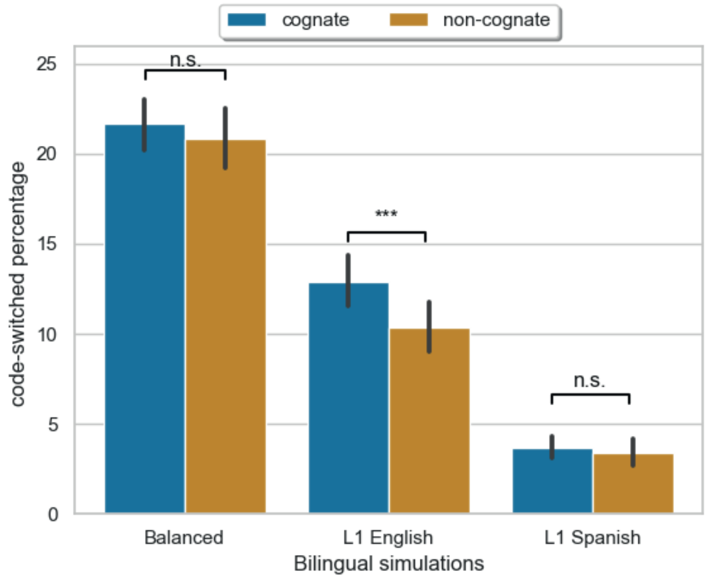

Figure 5.2.: Percentage of code-switches among correctly produced sentences with vs without cognates for the Balanced and Non-balanced bilingual simulations (5,922 sentences per simulation and set). The error bars indicate the 1000-sample bootstrapped 68% Confidence Intervals and the indication of significance comes from the mixed-effects regression analyses per simulation.

to sentences that were correctly produced by all four cases (simulations and cognate sets). Therefore, we report on all correctly produced sentences per simulation: there are 22,036 correctly produced sentences in the L1 English simulations and 24,150 in the L1 Spanish simulations.

To test whether the cognate effect differs depending on the L1, the switch direction (L1-to-L2 vs. L2-to-L1) and the non-balanced bilingual simulation (L1 English vs. L1 Spanish), we ran a logistic mixed-effects regression analysis contrast coding the cognate status as before, and additionally coding the L1 (+.5 if the target language was the L1, −.5 if the target language was the L2) and the simulation (+.5 for the L1 English simulation, −.5 for the L1 Spanish). Like in the previous analyses, this analysis includes a by-model-run random intercept and a slope of the cognate status. The cognate effect differs between the two simulations ($b = 0.39; z = 2.54; p = 0.01$) but not significantly between the two languages ($b = 0.17; z = 1.08; p = 0.28$). Separate analyses per simulation revealed a significant L1-to-L2 cognate effect for the L1 English simulation ($b = 0.73; z = 5.08; p < 0.0001$) but not for the L1 Spanish simulation ($b = -0.03; z = 0.29; p = 0.92$).

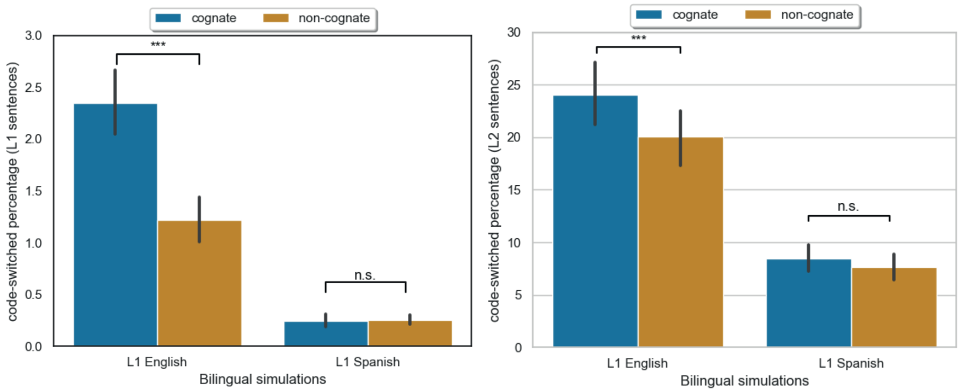

a: L1-to-L2 **b:** L2-to-L1

Figure 5.3.: Percentage of (a) L1-to-L2 and (b) L2-to-L1 code-switches, among correctly produced L1 (22,036 for the L1 English simulation and 24,150 for the L1 Spanish simulation) and L2 sentences (7,570 for the L1 English, 5,964 for the L1 Spanish simulations) respectively, with vs without cognates for the non-balanced bilingual simulations. Note that the percentages are different between the two switch directions and therefore the y-axes differ between the two panels. The error bars indicate the 1000-sample bootstrapped 68% Confidence Intervals and the indication of significance comes from separate mixed-effects regression analyses per simulation and switch direction.

Non-balanced models: cognate in the L2

Figure 5.3b shows the percentage of L2-to-L1 code-switches for the non-balanced bilingual models (7,570 sentences for the L1 English simulations and 5,964 for the L1 Spanish simulations).

In this switch direction, too, the analyses found a significant effect of the cognate status for the L1 English simulation ($b = 0.27; z = 3.82; p < 0.0002$). For the L1 Spanish simulation, the analysis showed no statistical difference between the cognate and non-cognate sets ($b = 0.11; z = 0.99; p = 0.32$).

5.4.2 Effect of percentage of cognates in language pair

Pairwise comparison

Figure 5.4 shows the percentage of code-switches across the five pairs of cognate models tested on the same 600 messages. To ensure a fair analysis, we only kept sentences that were produced correctly among all ten simulations (i.e., cognate vs non-cognate, and 10, 15, 20, 15, 30% cognates). Hence, in these simulations, too, the percentages are against *correctly produced sentences* only, which amount to 6,209 per simulation (out of 24,000 test sentences). To ensure that the criterion is not too strict, we also investigated the model's behavior for all correctly produced sentences (not only the ones that had been correctly produced across all simulations), and the resulting sentences per simulation showed a similar pattern; therefore, these results are not presented here.

A logistic mixed-effects regression analysis (performed similar to the comparison between cognate and non-cognate sets, but this time including the percentage of cognates and their interaction with cognate status) showed no statistically significant effect of percentage of cognates ($b = -0.00; z = -0.39; p = 0.698$), no significant cognate effect ($b = -0.07; z = -0.86; p = 0.39$), nor an interaction between the two ($b = 0.00; z = 1.43; p = 0.15$).

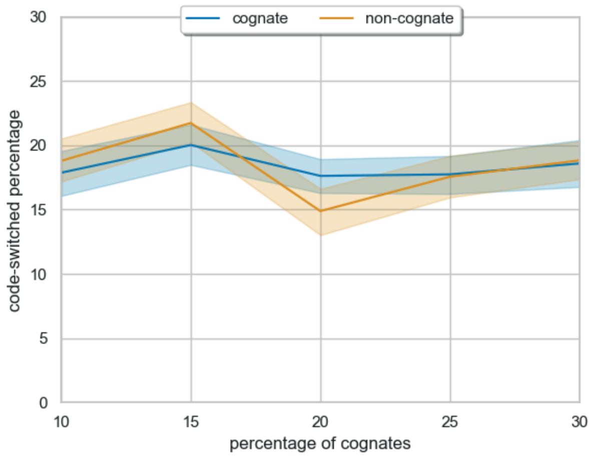

Figure 5.4.: Pairwise comparison of identical sentences with and without cognates; percentage of code-switches among correctly produced sentences at the 20th epoch (6,209 per model) as a function of percentage of cognates in the language. Shaded areas indicate the 1000-sample bootstrapped 68% Confidence Interval.

Non-pairwise comparison

Figure 5.5 shows the percentage of code-switches for the 8 models (containing 0, 10, ..., 70% cognates), tested on two test sets comprising: i) 600 messages with exactly one cognate per sentence ("cognate set") and ii) 600 different messages with no cognate per sentence ("non-cognate set"). As mentioned in the pairwise comparison, we only kept sentences that were correctly produced in all models, resulting in 11,583 test sentences per model for the cognate-containing test set (81,081 sentences in total) and 13,836 for the non-cognate-containing test set (110,688 in total).

Note that the two sets of test sentences are randomly different, because one never contains a cognate whereas the other always does; a comparison between the two sets is possible, but it is not as well controlled as a pairwise comparison. The main focus here is the comparison between models with different percentages of cognates. Additionally, the 0% model was not tested on the test set containing a cognate because, by definition, this model has no cognates. A mixed-effects analysis performed as in the pairwise comparison did not reveal statistical differences between the models containing different percentage of cognate in the language pair ($b = -0.05; z = -1.27; p = 0.21$). It did, however, show a cognate effect ($b = 0.30; z = 9.71; p < 0.00001$) and an interaction between the cognate set and the percentage of cognates ($b = 0.03; z = 2.18; p = 0.0291$).

5.5 Discussion

5.5.1 Balanced versus non-balanced models

The balanced and L1 Spanish simulations show no cognate triggering effect, whereas the L1 English simulations do, in both switch directions (L1-to-L2 and L2-to-L1). It is surprising that only one language pair gets affected by the cognate status of the produced sentences. If we compare the results to Chapter 3 (e.g., Figure 3.3), it is obvious that the L1 Spanish model produces a similar percentage of code-switches compared to that (non-cognate containing) chapter. The lexicon and structures are slightly different between the two chapters, but the difference is not so big that it would prohibit an imperfect comparison between the simulations. Whereas the L1 Spanish simulation stays rather unaffected, the code-switched performance of the L1 English model is boosted overall and even more so when there is a cognate in the sentence, thus supporting the cognate triggering effect.

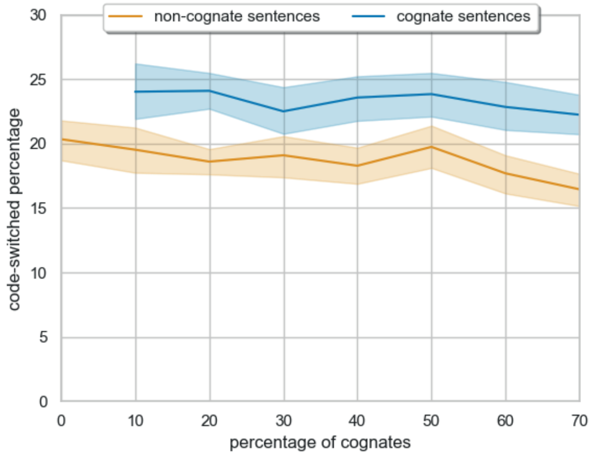

Figure 5.5.: Non-pairwise comparison; percentage of code-switches among correctly produced sentences at the 20th epoch as a function of percentage of cognates in the language, tested on i) sentences containing a cognate (11,583 correct sentences per model) and ii) sentences not containing a cognate (13,836 correct sentences per model). Shaded areas indicate the 1000-sample bootstrapped 68% Confidence Interval.

We can only speculate as to why this effect is only apparent in the L1 English simulation and not in the L1 Spanish and Balanced simulations. First, it may have to do with our methodological choice to convert all cognate words to the English form; 'dog' and the Spanish equivalent 'perro' both became 'dog'[5] when converted to a cognate, whereas it would have been better to randomly assign the Spanish or English form each time a cognate was created. However, it is hard to imagine how this could have affected so dramatically the code-switched performance of the L1 English model. Second, and more importantly, given that Spanish is a gendered language and English is not, it is possible that there is an effect of the grammatical gender. Another difference between the two languages is that the word order differs between NPs that contain adjectives; in English the adjective comes before the noun whereas Spanish prefers post-nominal adjectives ("the small house" - "la casa pequeña"). Given that a switch within an NP between an adjective and a noun is not common (neither among human nor simulated bilinguals), a Spanish noun cognate would most likely be followed by a Spanish adjective, and to produce a code-switch the model would need to switch

[5]In practice, this means that if 'dog' is the 4th node (item) in the lexicon and 'perro' the 12th, the cognate word in a Spanish or English sentence would always activate the 4th node; the 12th would remain unseen throughout training and testing.

after the adjective, at which point the cognate effect might have decreased, whereas for English the code-switch is able to occur right after the cognate. We analyzed the model output to determine whether this is the case with adjective-containing NPs; the results can be seen in Table 5.1, and they indicate that the cognate effect is indeed stronger for English-to-Spanish.

Table 5.1.: Absolute frequencies of Spanish-to-English and English-to-Spanish code-switches in adjective-containing sentences in the non-balanced bilingual models.

	Cognate	
Switch from	Yes	No
Spanish	45	42
English	107	49

5.5.2 Effect of percentage of cognates in language pair

With regards to the percentage of cognates in a language pair and its effect on code-switched production, the simulations do not predict important differences between different cognate percentages. This indicates that the percentage of cognates we used for the Bilingual simulations likely did not affect the overall results. Additionally, this is an interesting prediction that remains to be tested among humans.

The pairwise comparison shows that the cognate and non-cognate models produce a similar percentage of code-switches; there is a small decline in the non-cognate model for the 20% cognate models, but it seems to be a non-significant fluke given that for all other percentages the cognate and non-cognate models produce very similar percentages of code-switched sentences.

The non-pairwise comparison shows a slight decline in the percentage of code-switching as the percentage of cognates in the language pair increases but this trend is not significant either. When inspecting, however, the percentage of code-switches produced in the cognate and non-cognate test sets, the simulations show that there is a larger percentage of code-switches in sentences that contain a cognate compared to sentences that do not contain a cognate (Figure 5.5a). We do not know why this is; one might think that this is due to the difference between the numbers of correctly produced sentences: for the cognate-containing test set, the model has produced 11,583 sentences correctly, whereas for the non-cognate-containing test set the same model produced 13,836 sentences correctly. However, when we counted the number of all code-switches (including incorrect ones) for the two test sets, the cognate-containing test set still produced more code-switches; this indicates that the explanation above does not hold.

The main difference between the pairwise and non-pairwise methods is that in the former different sets of models ("cognate" and "non-cognate") are tested on the same messages, whereas in the non-pairwise comparison the same models (containing 0-70% of cognates) are tested on different messages (one set containing and the other not containing cognate concepts). An additional difference in the non-pairwise comparison is that the non-cognate sentences only include a restricted set of nouns (30% of total noun concepts), because 70% of the noun concepts had been converted to cognates in the 70% cognate model and therefore had to be excluded. Both methods, however, indicate that the percentage of cognates in a language pair does not affect the probability of a code-switch. Given that one method (non-pairwise) exhibited a cognate effect whereas the other method (pairwise comparison) did not, we cannot safely conclude that there is a triggering effect in the Balanced model for any percentage of cognates in the (Spanish–English) language pair. It would be interesting, however, to repeat this work on the L1 English model that exhibited the cognate triggering effect in our experiments.

5.6 Conclusion

In this chapter, we have presented a novel method to investigate the effect of cognates on code-switching, as well as the effect that the percentage of cognates in a language pair has on code-switching probability. Out of the three simulations (Balanced, L1 English and L1 Spanish) only the L1 English simulation produced more code-switches when encountering a cognate. Further work needs to be done to determine why the code-switched production of the L1 English model got enhanced with the mere existence of cognates in the language pair and even more so when presented with cognates. In the future, it would be interesting to investigate the effect that the position of a cognate (e.g., whether it is at the beginning or the end of a sentence) has on code-switched production. More importantly, the model makes predictions on human code-switched behavior that could be attested in future psycholinguistic experiments or corpora analyses. Specifically, one could test whether indeed the percentage of cognates in a language pair does not matter with regards to the probability of code-switching, and whether there are differences between balanced and non-balanced human bilinguals, and between L1 Spanish and L1 English speakers, that match the model's predictions. The latter prediction (that L1 English speakers are more likely than L1 Spanish speakers to code-switch after a noun phrase containing a cognate noun and an adjective) can be tested using the Bangor Miami corpus (Deuchar et al., 2014).

6

"He's pregnant": simulating the confusing case of gender pronoun errors in L2 English

Even advanced Spanish speakers of second language English tend to confuse the pronouns 'he' and 'she', often without even noticing their mistake (Muñoz, 1991). A study by Antón-Méndez (2010) has indicated that a possible reason for this error is the fact that Spanish is a pro-drop language. In order to test this hypothesis, we used an extension of Dual-path (Chang, 2002), a computational cognitive model of sentence production, to simulate two models of bilingual speech production of second language English. One model had Spanish (ES) as a native language, whereas the other learned a Spanish-like language that used the pronoun at all times (non-pro-drop Spanish, NPD_ES). When tested on L2 English sentences, the bilingual pro-drop Spanish model produced significantly more gender pronoun errors, confirming that pronoun dropping could indeed be responsible for the gender confusion in natural language use as well.

This chapter is based on: Tsoukala, Chara, Stefan L. Frank, and Mirjam Broersma (2017). ""He's pregnant": simulating the confusing case of gender pronoun errors in L2". In: *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Ed. by G. Gunzelmann, A. Howes, Tenbrink T., and E.J. Davelaar. London, UK, pp. 3392–3397

6.1 Introduction

Second language (L2) speech errors have been employed in the past as a means to understand bilingual speech production as well as the acquisition process of a foreign language (Antón-Méndez, 2010; Poulisse, 1999). Certain L2 errors are observed more often due to discrepancies between the first language (L1) and the L2. For example, if the expression of a message in the L2 requires the inclusion of a specific feature that would not be necessary in the L1, then speakers of these two languages may produce a speech error in their L2 due to L1 transfer (Odlin, 1989). In this study, we focus on a gender-related L2 pronoun error that has been observed among native speakers of Spanish; namely, errors involving the third person singular nominative pronouns 'he' and 'she'. Even advanced Spanish speakers of L2 English occasionally confuse the two pronouns, referring to an actress as 'he' or a father as 'she', often without even noticing their mistake (Muñoz, 1991). At first, this phenomenon seems surprising because the Spanish language does have two equivalent pronouns ('él' for 'he' and 'ella' for 'she'), and also a very strong separation between the two genders, even more so than in English. For instance, depending on the suffix a word can be feminine or masculine (e.g., maestro - teacher [masculine], maestra - teacher [feminine]; niño - child [masculine, a.k.a. boy], niña - child [feminine, a.k.a. girl]). This means that the gender mistakes that Spanish speakers make in English cannot be attributed to the lack of familiarity with the distinction. Furthermore, the challenge the English pronoun system poses for native speakers of Spanish could not be due to its inherent difficulty, as this would mean that most non-native speakers of English, regardless of their L1, would produce the same mistake. As Muñoz noted, a low proficiency level of the native Spanish speakers is not a reason either. This was also demonstrated in the experiments of Antón-Méndez, where the participants showed an intermediate to upper intermediate knowledge of English. Finally, note that the gender mismatch error cannot be classified as a syntactic error; the produced sentence is grammatically correct, but it conveys the wrong meaning.

A hypothesis which has been put forward (Muñoz, 1991; Antón-Méndez, 2010) regarding the cause of errors in the use of English pronouns is the pro-drop status of the Spanish but not the English language. In pro-drop languages, nominative personal pronouns are often omitted (1b) because the number and person information is conveyed in the conjugated verb (Davidson, 1996), whereas in English the omission of the pronoun would result in an ungrammatical sentence (2b).

1. a) *Él/ella* tiene un perro (Spanish)

b) tiene un perro

2. a) *He/she* has a dog (English)

   b) * has a dog

It is hard to imagine, however, how the pro-drop feature of the L1 might result in a gender pronoun error ("He's walking", when referring to a woman) instead of an omission ("Is walking"), which would be the case in a direct language transfer.

As a matter of fact, native speakers of Spanish have been noted to produce another gender-related pronoun error in English, this time regarding possessive pronouns ('his', 'her'). Due to the high frequency of this type of errors, a lot more emphasis has been given to the misuse of these pronouns than the subject pronouns (White, Muñoz, and Collins, 2007; Antón-Méndez, 2011). The reason that English possessive pronouns pose a challenge for native speakers of Spanish is most likely that in Romance languages the possessive pronoun agrees in gender and number with the possessum, namely the noun that follows, whereas in Germanic languages such as English the possessive pronoun refers to the antecedent. For example:

i His daughters are on vacation. [*his*: 3rd person masculine singular]

ii Sus hijas están de vacaciones. [*sus*: 3rd person (feminine) plural]

Due to the different information encoding Spanish speakers of English may occasionally make gender mistakes such as "He called her mother", where 'her' refers to the antecedent ('he') and not a different female person. This is because 'mother' is female, and a Spanish speaker would use that gender information to construct the possessive pronoun in Spanish. The resulting error in English is, of course, confusing, as a speaker of English would not guess that 'her' in this case refers to the same subject ('he'). The gender error in the case of L2 English possessive pronouns seems clearly due to L1 transfer, because the properties of Spanish are directly applied to English. In the case of the subject pronoun gender errors, on the other hand, it is not evident that the pro-drop feature of one language would lead to a gender error in L2. The present study addresses only the latter type of errors.

Antón-Méndez (2010) has investigated the hypothesis that the pro-drop feature of Spanish is responsible for the gender pronoun errors in L2 English ("pro-drop hypothesis"). She conducted an experiment eliciting semi-spontaneous speech in English, where she compared native Spanish and native French speakers of L2 English with respect to the pronoun errors they produced. French was chosen as it is a Romance language that is similar to Spanish in several aspects, but which, in contrast to the Spanish language, is not a pro-drop language. Each test group consisted of 20 participants who were comparable in terms of education, age of English acquisition, frequency of use and proficiency. The participants were shown 43 illustrations and were asked questions designed to elicit pronoun production. The subjects were instructed to respond freely, and the pronoun errors they produced were recorded. The types of reported errors fall in the following categories: person errors (e.g., 'I' instead of 'you'), number errors (e.g., 'I' instead of 'we'), gender errors ('he' instead of 'she' and vice-versa), animacy errors (e.g., 'he' instead of 'it'), omission errors (e.g., 'is swimming'), insertion errors (e.g., 'the boy he played' instead of 'the boy played') and other errors (e.g., 'it' instead of 'there' in 'there is').

Spanish speakers of L2 English indeed made significantly more gender errors (4.30%) compared to other types of pronoun errors and to the French group (0.68%)[1]. The pronoun errors recorded were not due to erroneous transfer of the Spanish L1 grammar, as the Spanish speakers made no omission errors ('is swimming'); thus, in none of the items of Antón-Méndez's experiment did the subjects omit a pronoun, which would have been the case in a grammatical transfer. Importantly, even though there were slightly more 'he' than 'she' errors (he: 5.68%, she: 2.98%), the difference is not statistically significant. Therefore, the Spanish speakers were not using a default pronoun (e.g., always 'he' instead of 'she'). The use of 'he' as the default pronoun would have suggested that another factor might underlie the error, for instance, the difficulty that the English phonology poses for speakers of Spanish. The Spanish phonology does not contain the phonemes /ʃ/ in 'she' and /h/ in 'he', therefore one explanation for the gender pronoun issue could be at the phonological level. In the present study we focused only on the pro-drop feature, not because we disregard the potential role of the phonology, but because we wanted to investigate whether the pro-drop feature has the capacity of causing this type of gender errors in L2.

In order to focus on the pro-drop feature, we simulated bilingual sentence production using computational cognitive modeling. The pro-drop feature is not the sole difference between the French and Spanish languages, and one could argue that the differences in the error patterns between the two groups could have been partially attributed to confounding factors, for instance, to a different L2 English teaching system in Spain and France.

Using computational modeling we can remove all possible confounds and therefore minimize the variance by focusing only on the phenomenon of interest, which in

[1]The percentages are calculated by Antón-Méndez (p. 129, Table 6) and they represent the frequency of the gender pronoun mistake (68 and 10, respectively) with respect to the total number of pronouns produced where this particular mistake could have occurred.

this case is the pro-drop feature and its possible effect on L2 English pronouns. For this reason, we modified Dual-path (Chang, 2002), a computational cognitive model of sentence production, to account for bilingualism. We then compared L2 English speech production of simulated native speakers of Spanish (ES) on the one hand, to L2 production of simulated native speakers of a Spanish-like language ('non-pro-drop Spanish', NPD_ES) on the other hand. The latter contained all the features of the Spanish language (lexicon, allowed structures) except the pro-drop feature; therefore, pronouns needed to be used at all times. All input languages (ES, NPD_ES and EN) were artificially generated and based on the Spanish and English language, using a subset of their lexica and syntactic structures. If the bilingual Spanish-English (ES-EN) Dual-path model produces significantly more subject pronoun errors in English than its Spanish-like non-pro-drop equivalent (NPD_ES-EN), it will be clear that the pro-drop feature of the Spanish language is the reason for this particular L2 error in the simulation, as the two simulated languages differ only in their pro-dropness. If this is the case, we will have confirmed that the pro-drop feature has the capacity to lead to gender pronoun errors in L2 English.

6.2 Method

In order to simulate Spanish speakers of L2 English, we developed two bilingual models using a modified version of Dual-path which is a connectionist model based on the Simple Recurrent Network (SRN; Elman 1990) architecture (Chang, 2002).

6.2.1 Bilingual Dual-path model

Dual-path (Figure 6.1) learns to convert a message into a sentence by predicting the sentence word by word ("next word prediction"). It has two pathways that influence the production of each word; the meaning system which learns concepts, roles and event semantics, and the sequencing system which is an SRN that learns to abstract syntactic patterns. Both paths influence the word output layer. The sequencing system consists of one recurrent hidden layer (of 30 units in our simulations) and two "compress" layers (of 12 units each) that are placed between the input word, the hidden layer and the output word.

The meaning system learns to map the input word onto a concept, which is linked to a specific thematic role (that is given for each sentence through fixed connections). The fixed connections allow the separation between concepts and roles, which, in turn,

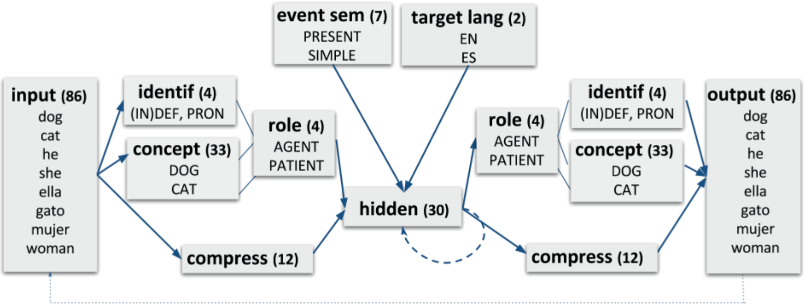

Figure 6.1.: Bilingual Dual-path model

enables the model to generalize and to produce words in novel places. The thematic role is connected to the hidden layer, and so is the “event-semantics" layer. The hidden layer spreads the activation to the next thematic role (in the meaning path, and the “compress" unit in the syntactic path), which is in turn linked to a specific predicted concept that is used as input to the output word layer, along with the “compress" unit.

In the original model, all layers use the tanh activation function, except the output layer and the (comprehended/input) role layer that use softmax. In the modified version of the model, we also employed softmax for the predicted role layer. This led to a stricter selection of the upcoming thematic role which helped overcome a difficulty that the model had with learning the correct articles regarding gender and definiteness (e.g., ‘a’ vs ‘the’). Furthermore, our version has a “target language" layer in the meaning path that is used as an additional input to the hidden layer, along with the “event-semantics" layer. The “target language" denotes the intended spoken language and helps the model handle more than one language. The modified model can be found at `https://github.com/xtsoukala/dual_path` .

6.2.2 Input languages

Message

Dual-path is trained using randomly generated sentences paired with their meaning (Chang, Dell, and Bock, 2006). The meaning (message) contained information regarding four thematic roles (AGENT, PATIENT, ACTION, RECIPIENT). A concept (e.g., ‘WOMAN’ for the English word ‘woman’ or Spanish word ‘mujer’) was assigned to each thematic role depending on the meaning that needed to be expressed (e.g., in the sentence “the woman run -s" the message would include AGENT= WOMAN, DEF). Furthermore, the message contained event-semantic information (denoted as ‘EVENT-SEM’), which gave information regarding the tense (PRESENT or PAST) and aspect (SIMPLE or PROGRESSIVE). The message contained information about the target language (ES or EN) as well. This information was given at the beginning of the sentence along with the roles and the event-semantics, so that the model knew whether it was supposed to produce an English or Spanish sentence.

Structures

The allowed structures for all languages were the following (where ‘S’, the subject, is omitted in the pro-drop case):

1. (S)V: (Subject) - Verb, e.g., “He runs"
2. (S)VO: (Subject) - Verb - Object, e.g., “She kicked the ball"
3. (S)VIODO: (Subject) - Verb - Indirect Object - Direct Object, e.g., “He gave the girl a book"
4. (S)VDOIO: (Subject) - Verb - Indirect Object - Direct Object, e.g., “He gave a book to the girl"

The sentences in English and in non-pro-drop Spanish always started with a pronoun, and the sentences in pro-drop Spanish never started with a pronoun but always with a verb.

Lexicon

The total lexicon consisted of 34 nouns: 11 male (‘man’, ‘boy’, ‘father’, ‘brother’, ‘dog’, ‘hombre’, ‘niño’, ‘padre’, ‘hermano’, ‘perro’, ‘maestro’), 11 female (‘woman’, ‘girl’, ‘mother’, ‘sister’, ‘cat’, ‘mujer’, ‘niña’, ‘madre’, ‘hermana’, ‘gata’, ‘enfermera’) and 12 inanimate (‘ball’, ‘stick’, ‘toy’, ‘kite’, ‘key’, ‘bag’, ‘pelota’, ‘palo’, ‘juguete’, ‘cometa’, ‘llave’, ‘bolso’), 24 verbs (e.g., ‘give’, ‘show’, ‘walk’, ‘throw’, ‘present’, ‘dar’, ‘lanzar’, ‘presentar’, ‘nadar’, ‘caminar’) and 26 function words (e.g., articles (‘a’, ‘the’, ‘un’, ‘una’, ‘el’, ‘la’), pronouns (‘he’, ‘she’, ‘él’, ‘ella’) and auxiliary verbs (‘is’, ‘was’, ‘está’, ‘estaba’)).

The model treats the verb lemma (‘give’) and the suffix (‘-s’) as two different units. Syntactic information (such as ‘verb’, ‘noun’) is not given explicitly, but is learned by the model during training through the syntactic path. The syntactic gender was also learned implicitly during training through the article of Noun Phrases (NP) and pronouns. Semantic gender (e.g., 'ACTRESS, F', 'ACTOR, M') was not included in the model.

Thematic roles could be expressed using either an NP with definite (DEF) or indefinite (INDEF) articles (e.g., 'the woman', 'a woman') or the pronoun (PRON) equivalent ('she').

Example

The following message would be the same across languages:

AGENT=WOMAN, PRON;

ACTION=GIVE;

PATIENT=INDEF, KEY;

RECIPIENT=DEF, GIRL;

EVENT-SEM=SIMPLE, PRESENT, AGENT, PATIENT, RECIPIENT

and it would be expressed linguistically in the following manner for the three languages:

1. she give -s the girl a key . [EN]
2. d -a a la niña una llave . [ES]
3. ella d -a a la niña una llave . [NPD_ES]

6.2.3 Training

The two models were trained on 2000 randomly generated sentences (training set) and tested on 500 unseen sentences (test set). The models contained almost identical sets, with the only difference that the NPD_ES model expressed the subject pronoun at all times, whereas the ES model never did and always started with a verb. For each model we ran 100 simulations using the same input, but different random initial weights per simulation, as the input and the weights are the only non-deterministic parts of the model. The models were trained for 20 epochs, where 1 epoch corresponds to a full iteration of the training set (2000 sentences). At the beginning of each epoch, the training set was shuffled. In order to simulate late L2 acquisition, we first trained the models for 20 epochs using Spanish input only, and then used the fully trained weights as initial weights for the bilingual models. The bilingual input consisted of newly

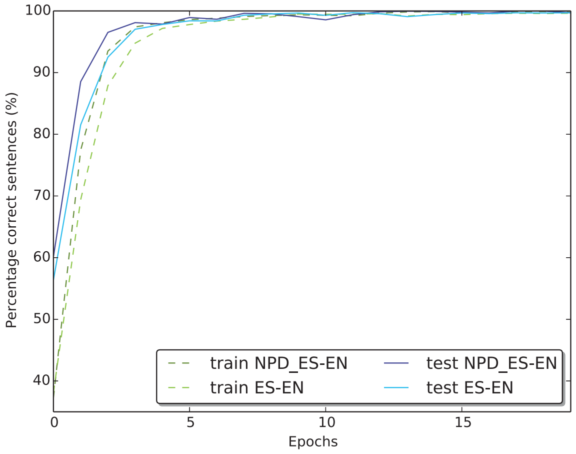

Figure 6.2.: Performance on the training and test sets over the training period (20 epochs) averaged over 93 simulations for the two bilingual models. Performance is measured in percentage of correctly produced Spanish and English sentences.

generated (2000 training and 500 test) sentences, this time using 50% (pro-drop or non-pro-drop according to the model) Spanish and 50% English. We excluded from the analysis 7 simulations that did not manage to learn at least 75% of the test set by the end of the training in one of the two models, leading to a total of 93 simulations. Both bilingual models were able to perform equally well by the end of the training, reaching 99.69% correct for ES-EN and 99.70% correct for NPD_ES-EN (Figure 6.2) on the test set that contained English and Spanish sentences.

6.3 Results

In order to assess the performance of the two bilingual models on L2 pronouns, we focused only on the English sentences (50% of the test set). If a pronoun error was detected and the sentence was grammatical, it was classified as a gender pronoun error. We compared the performance of the two bilingual models with regard to the gender pronoun error production. If the models had a comparable performance we would not be able to confirm that the pro-drop feature has the capacity to lead to gender pronoun errors in L2 English. If, on the other hand, the NPD_ES model made fewer gender pronoun errors than the ES model it would indicate that the pro-drop feature is a possible explanation.

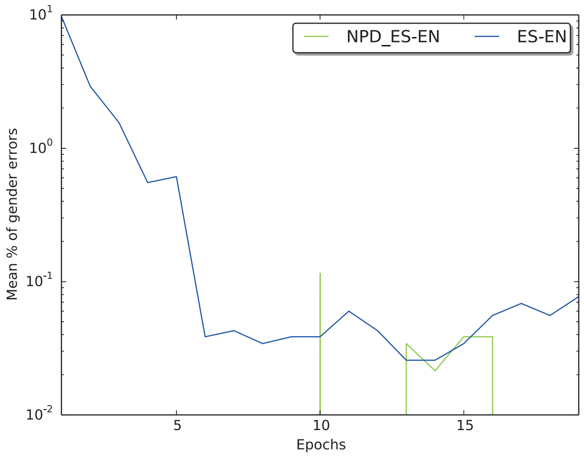

Figure 6.3.: Production of English gender errors (in log scale) averaged over 93 simulations for the two bilingual models. Note that the gaps in the non-pro-drop model are because log(0), for 0% error rate, is not defined, and therefore not plotted.

The non-pro-drop Spanish-English (NPD_ES-EN) bilingual model (Figure 6.3) produced almost no gender pronoun errors (maximum percentage: 0.11%) whereas the bilingual model based on pro-drop Spanish (ES-EN) initially produced 9.75% pronoun errors, gradually dropping to 0.05%.

Crucially, the ES-EN model never reached 0% (minimum error rate: 0.02%) whereas the NPD_ES-EN model did. Following visual inspection, we ran a z-test for proportions from epoch 5 onwards to test for a difference in error rate between the models. The difference is significant (z=7; p<.001).

6.4 Discussion

Our simulations showed that a bilingual model with L1 pro-drop Spanish and L2 English produced significantly more gender pronoun errors than a similar model with L1 non-pro-drop Spanish. These sentences were grammatically correct: the only error they contained was a pronoun with incorrect gender. Given that the only difference between the two L1s was the pro-drop feature, we have demonstrated that the pro-drop nature of Spanish can indeed cause the gender pronoun error as observed in L1 Spanish speakers of L2 English.

Why the pro-drop feature does not lead to a direct language transfer (“is walking") in either the model or humans remains to be investigated, as the current simulations and results do not explain how pro-dropness in L1 could lead to gender errors in L2. Nevertheless, having a computational model that simulates the gender pronoun errors in L2 English can point us in the right direction. Our hypothesis for the occurrence of the gender error is that the gender information is not as crucial for the message planning, at least in the subject position, of a pro-drop language, and is therefore weaker or omitted, even when producing sentences in a non-pro-drop L2.

It is important to point out that the Dual-path model does not contain a phonological level (Garrett, 1988). One might have thought that the reason Spanish speakers confuse the words ‘he’ and ‘she’ is because of the difficulty the English phonology poses for native speakers of Spanish. However, our simulations have produced gender errors without having any phonological representations. This does not mean that phonology could not play a role, but rather that it is not the only possible explanation.

It is also crucial to note two simplifying assumptions in these simulations. First, as mentioned in the Method section, the input for all three languages (EN, ES, NPD_ES) was artificially generated and it only represented a subset of the actual languages. In general, using natural input would be preferable as it would increase the validity and naturalness of the results. However, the benefit of miniature languages that are typically used in cognitive modeling is that they can be easily manipulated. For instance, in the simulations described here we were able to add and remove the pro-drop feature at will, leaving everything else the same, and thus to isolate this important feature from confounding factors.

Second, a crucial simplifying assumption in the miniature language is the absence of full NP subjects. We therefore repeated the simulations using new input for all languages, this time including 50% pronouns at the subject position and 50% noun phrases. Preliminary simulations show no gender errors in either model, which means that further research is needed using more natural language input, starting with a more naturalistic proportion of pronouns and NPs in the subject position based on English and Spanish corpora.

6.5 Conclusion

Computational modeling can be used to validate or generate linguistic hypotheses while focusing on specific factors of interest and minimizing the variance. In this study, we have addressed the question as to whether the pro-drop feature of the Spanish language has the capacity to cause the gender pronoun errors that Spanish speakers of L2 English have been shown to produce (Muñoz, 1991; Antón-Méndez, 2010). The reported simulations showed that the model with L1 pro-drop Spanish produced more gender pronoun errors in L2 English than the model with L1 non-pro-drop Spanish, which is a necessary but not sufficient condition for the pro-drop hypothesis.

7 General discussion

7.1 Summary of results

In this dissertation, I have presented a novel method for investigating code-switching and cross-linguistic influence in bilingual sentence production using computational modeling. Specifically, in **Chapter 2**, I introduced the Bilingual Dual-path model and showed that a computational cognitive model of monolingual sentence production can be extended to bilingualism by adding a language control node and by training it with bilingual input. Furthermore, by allowing the bilingual model to produce a sentence in either language, I demonstrated that the model produces code-switched sentences, even though it has not been exposed to code-switched input. This indicates that i) some aspects of code-switching can be attributed to internal factors and to the distribution of the languages involved, and therefore ii) a bilingual speaker does not need (extensive) exposure to code-switching in order to code-switch. In this chapter I simulated balanced bilinguals, i.e., bilinguals who have been exposed to both languages simultaneously and to a similar extent, and showed that several produced patterns are in line with those observed among human (balanced) bilinguals, as reported in psycholinguistic and sociolinguistic studies.

In **Chapter 3** I tested the robustness of the model by randomizing almost all free parameters of the simulations. Additionally, I simulated non-balanced bilinguals, i.e., speakers who have first acquired a native language (L1) before becoming exposed to their second language (L2). At the early stages of L2 acquisition, the non-balanced models, that had not been exposed enough to the L2, preferred to switch back into the L1 frequently. However, when these models became somewhat proficient in the L2, they produced very few code-switches. The balanced models, on the other hand, code-switched much more frequently than the non-balanced ones. Both the L2-acquisition phenomenon and the code-switching differences between balanced and non-balanced speakers have been observed in human bilinguals. Furthermore, I analyzed the code-switching patterns in the Miami corpus, which contains spontaneous dialogues between Spanish–English speakers living in Miami, Florida. I compared the simulated patterns of balanced and non-balanced speakers to the ones produced by the Miami corpus speakers and to a corpus analysis reported by Poplack (1980). As mentioned in this chapter, it would be impossible to observe a perfect match between the model's code-switching patterns and the corpus data because the simulations use *artificial miniature languages*, which have features of English and Spanish. Nevertheless, some patterns were similar to what human bilinguals produced in the two corpora.

Having shown that the model reliably code-switches, and that the patterns sometimes reflect those produced by human bilinguals, in **Chapter 4** I employed the model to shed light on the auxiliary phrase asymmetry observed in Spanish–English bilingual speech. Through manipulations of the miniature languages, I showed that this asymmetry can be attributed to the distributional patterns of the two languages, and more specifically to the differences in the semantic weight of the auxiliary verbs, even though the structures have the same syntactic patterns. This indicates that even though syntax and word order play an important role in (dis)allowing code-switching in certain places, they are not the only factors affecting code-switching. Through this experiment, I also demonstrated the advantage of using artificial languages; they allow us to modify the grammar and frequency of the languages involved and to explicitly test the effect these have on speech patterns. For instance, in this chapter I changed the function of a Spanish auxiliary verb and consequently observed that the auxiliary phrase asymmetry disappeared; the role of the auxiliary verb cannot be experimentally tested in humans, whereas it is easily tested using computational modeling.

In **Chapter 5**, I tested the effect of cognates on code-switching among balanced and non-balanced bilinguals, as well as the effect that the percentage of cognates in a language pair has on code-switching. The artificial languages allow for a clean set-up, i.e., to compare the production of a non-cognate-containing sentence with that of an identical sentence in which one of the words is a cognate. The model predicts no cognate effect for the balanced bilingual model, whereas it does show the so-called cognate triggering effect for the L1 English model, but not for the L1 Spanish one. As, to the best of my knowledge, there are no available data comparing the cognate triggering effect in the sentence production of balanced bilinguals and non-balanced bilinguals with different L1s, these predictions remain to be tested in humans.

Last, in **Chapter 6** I showed that the model can also be employed to explain other cross-linguistic phenomena, such as L2 speech errors. I used the model to show that a subject pronoun error, which is sometimes produced by Spanish speakers of L2 English, can be attributed to the pro-drop status of the L1 Spanish. Additionally, in work not included in this chapter, Khoe, Tsoukala, Kootstra, and Frank (2020) employed the model to simulate cross-linguistic structural priming between Spanish and English, and showed that cross-language priming does occur in the model; this study opened a new research possibility to test how syntax is shared across languages in the bilingual mind.

7.2 Overall conclusions

All the above results show that via computational cognitive modeling of bilingual sentence production one can remove all extra-linguistic influences in language use and investigate which patterns are caused by the interaction of the statistical properties of the languages involved in combination with the cognitive system, rather than by community norms. However, it is important to acknowledge that no explanation of the process of code-switching, or any speech pattern, can be complete without taking into consideration the impact of community practices. Community norms have been shown to influence code-switching patterns, to the extent that there can be opposite preferences between communities that use the same language pair; for instance, Blokzijl, Deuchar, and Parafita Couto (2017) analyzed the Miami corpus and a Nicaraguan Spanish–English creole corpus and found that speakers in Miami preferred to use a Spanish determiner in a mixed Determiner Phrase, whereas in Nicaragua only the English creole determiner was used. Similarly, Balam, Parafita Couto, and Stadthagen-González (2020) examined three Spanish–English communities (from Northern Belize, New Mexico, and Puerto Rico) with respect to their preference for two code-switched compound verb constructions; the U.S. bilinguals showed a different preference to the Northern Belize community, once again confirming that there are aspects other than grammar that affect code-switching across communities that speak the same language pair.

Modeling can help disentangle the linguistic from the extra-linguistic factors. For example, the simulations in **Chapter 3** produced certain patterns and in certain switch directions (e.g., more Spanish determiners than English ones) even though there were no extra-linguistic influences available to the model. More importantly, in **Chapter 4** the model helped answer the question as to whether the auxiliary phrase asymmetry exhibited by Spanish–English speakers arises due to community practices or whether it can be attributed to the statistical properties of the two languages. The model was able to simulate the asymmetry without exposure to this phenomenon, thus showing that at least part of it is due to the properties of the Spanish auxiliary verb.

Note, however, that, as is the case with all psycholinguistic experiments, testing a hypothesis and showing that a phenomenon *can* occur because of a specific linguistic attribute (e.g., the pro-drop feature of Spanish in Chapter 6) does not exclude other possible explanations for the same phenomenon; it merely supports the plausibility of the hypothesis under consideration.

Finally, from a cognitive perspective, it is interesting that the model contains no complex code-switching cognitive control (such as active inhibition). Several cognitive models of code-switching, such as the Control Process Model of code-switching (CPM; Green and Wei, 2014), suggest temporary inhibition during the production of alternations and insertions; however, the minimal computational model presented here can produce such patterns without any added restrictions. This is also in contrast with Green's (2018) statement that "varying language activation levels is an insufficient mechanism to explain the variety of language use" (p. 1). This is not to say that the CPM is not a correct model: it aims to explain how different types of output arise in situations with a different context, something that I have not attempted to do here. It is interesting, however, that the Bilingual Dual-path model exhibited alternations and insertions without employing any complex mechanisms.

7.3 Future directions

In this dissertation I focused on Spanish–English code-switching and cross-linguistic processing because there are corpora, such as the Miami corpus, and several studies based on communities that code-switch between English and Spanish; this makes the model output more easily attestable compared to using a language pair that is less documented. However, the model can be trained on any language pair. More interestingly, one could use this model to investigate code-switching between (more than) three languages. The Bilingual Dual-path model should, in principle, be easily **extended to multilingualism** by being exposed to several languages simultaneously during training, and by using the corresponding target language nodes.

As mentioned in several chapters, the model is trained on miniature languages. The main advantage of this is that it helps us remove the complexity of language and focus only on the phenomena of interest (such as the auxiliary phrase asymmetry in Chapter 4, the cognate status of a word in Chapter 5, and the subject pronoun error in Chapter 6). It could be fruitful, however, to **extend the language in a (more) naturalistic manner**. The first step towards this action would be to replace the 'concept' layer, which is currently a one-hot vector, with a cross-lingual word embedding. One-hot vector means that if the concept of interest (e.g., DOG) is the 10th unit in the 'concept' layer, only the 10th unit will be activated; each concept is either on (1) or off (0). This also means that to add one concept one would need to extend the layer vector by one value, which could become computationally costly for a large lexicon and set of concepts. A cross-lingual word embedding (e.g., Klementiev, Titov, and Bhattarai, 2012), on the other hand, is a dense and low-dimensional way to encode bilingual lexical semantics, thus making it a scalable way of representing a large number of concepts. In a preliminary experiment, I replaced the concept layer with a pre-trained English embedding[1] and trained the model with the miniature English sentences. The performance was initially lower (during the first epochs) compared to using a one-hot vector, but the Bilingual Dual-path quickly learned to use the embedding. This indicates that only a small architectural change needs to be made for the model to receive a larger input set. Apart from this change in the representation of concepts in the computational model, one would presumably want to use corpus data for training. As is done now with the artificial languages, any new input sentences would need to be semantically annotated in order to be used as training input to the model. The model receives as input a message with thematic roles and the corresponding concepts, therefore, one would need to automatically assign semantic role labels to the training sentences.

An alternative future direction would be to **add phonological information to the model**. As mentioned in several chapters and in Chapter 6 in particular, the phonology also plays an important role in bilingual production and code-switching. In this dissertation I have focused on the role of semantics and the syntactic distribution of the languages involved. In the future it would be interesting to add phonological information as well to enrich the model and to investigate the role of phonology on code-switching and cross-linguistic phenomena.

Last, as shown in the previous chapters, the Bilingual Dual-path model can simulate certain phenomena and test specific hypotheses regarding the interaction of two languages. However, in this dissertation I did not analyze specific processes of the model to investigate *why* it produces certain patterns, and to **explain how the specific cognitive processes of bilingual production and code-switching work**.

[1] Specifically, a word2vec (Mikolov et al., 2013) trained on the text8 corpus, which is a small subset of Wikipedia data and can be downloaded through Gensim (Řehůřek and Sojka, 2010).

Nevertheless, it is important to point out that the model is able to code-switch merely by being exposed to two languages and by having (and manipulating) a language control node that sets the conversational setting and allows the model to produce in either language; no other cognitive control is required for the model to code-switch.

There are two straightforward ways to explore the processes of the model: First, to test the **average activation of specific layer units** for certain structures, concepts, or around a code-switch. Second, to **deactivate (a.k.a. "lesion") the syntactic and semantic streams** and observe the code-switching behavior of the models. Following Chang (2002), who lesioned the two streams to simulate aphasic production, in two preliminary experiments I lesioned two trained models on either the syntactic or the semantic path. Both lesioned models produced code-switches; it would be interesting to quantitatively analyze the produced patterns so as to make predictions about code-switching in bilingual aphasia, and to better understand the role of semantics versus syntactics in code-switching.

7.4 Concluding remarks

In this dissertation I have argued that computational cognitive modeling is an interesting and novel method to research bilingualism and code-switching. The Bilingual Dual-path model in particular, with its straightforward architecture that consists of i) a semantic stream, ii) a syntactic stream, iii) a language control node that sets the conversational setting, shows that no complex code-switching cognitive control is needed for a model to code-switch. Additionally, the fact that Dual-path, a model of monolingual sentence production, was altered only minimally to be able to handle two languages simultaneously, shows that the cognitive mechanisms of bilingualism may not be that different from those of monolingualism. As far as the syntactic and semantic levels are concerned, the bilingual cognitive architecture and the monolingual architecture may be indistinguishable (Frank, 2021).

Furthermore, the model learned to code-switch without exposure to code-switched input, which indicates that part of code-switching can be attributed to internal factors and the distribution of the languages involved, and not only on community practices. This is important when answering questions regarding whether a pattern emerges due to internal factors or community-based influence, as mentioned above.

Last, in Chapter 1 (Section 1.2) I discussed several grammatical models that have attempted to explain code-switching patterns. The EC model, for instance, has argued that code-switching can only occur if the word order of the two languages is identical. However, the Bilingual Dual-path had no hard-coded constraints. Even so, the model did produce certain constraint-like patterns, such as the dispreference in Chapter 4 to produce a Spanish auxiliary verb together with an English participle. Therefore, deriving hard constraints from observed code-switched patterns may be too limiting or not sufficient, because, as discussed in the same chapter, word order and syntax are not the only factors governing code-switching. Instead, grammatical constraints can emerge from the interaction between properties of the production system and the languages involved.

Bibliography

Adel, Heike, Ngoc Thang Vu, and Tanja Schultz (2013). "Combination of recurrent neural networks and factored language models for code-switching language modeling". In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 206–211.

Antón-Méndez, Inés (2010). "Gender bender: Gender errors in L2 pronoun production". In: *Journal of Psycholinguistic Research* 39, pp. 119–139.

— (2011). "Whose? L2-English speakers' possessive pronoun gender errors". In: *Bilingualism: Language and Cognition* 14, pp. 318–331.

Balam, Osmer, María del Carmen Parafita Couto, and Hans Stadthagen-González (2020). "Bilingual verbs in three Spanish/English code-switching communities". In: *International Journal of Bilingualism* 24.5–6, pp. 952–967.

Bates, Douglas, Martin Mächler, Ben Bolker, and Steve Walker (2015). "Fitting Linear Mixed-Effects Models Using lme4". In: *Journal of Statistical Software* 67.1, pp. 1–48.

Belazi, Hedi M., Edward J. Rubin, and Almeida Jacqueline Toribio (1994). "Code switching and X-bar theory: The functional head constraint". In: *Linguistic inquiry*, pp. 221–237.

Bentahila, Abdelali (1983). "Motivations for code-switching among Arabic-French bilinguals in Morocco". In: *Language & Communication* 3.3, pp. 233–243.

Bhat, Gayatri, Monojit Choudhury, and Kalika Bali (2016). "Grammatical constraints on intra-sentential code-switching: From theories to working models". In: *arXiv preprint arXiv:1612.04538*.

Blokzijl, Jeffrey, Margaret Deuchar, and María del Carmen Parafita Couto (2017). "Determiner asymmetry in mixed nominal constructions: The role of grammatical factors in data from Miami and Nicaragua". In: *Languages* 2.4, p. 20.

Bock, Kathryn (1986). "Syntactic persistence in language production". In: *Cognitive Psychology* 18.3, pp. 355–387.

Bock, Kathryn and Helga Loebell (1990). “Framing sentences”. In: *Cognition* 35.1, pp. 1–39.

Branigan, Holly P., Martin J. Pickering, and Mikihiro Tanaka (2008). “Contributions of animacy to grammatical function assignment and word order during production”. In: *Lingua* 118, pp. 172–189.

Broersma, Mirjam (2009). “Triggered codeswitching between cognate languages”. In: *Bilingualism: Language and Cognition* 12, pp. 447–462.

— (2011). “Triggered code-switching: Evidence from picture naming experiments”. In: *Modeling bilingualism from structure to chaos: In honor of Kees de Bot*, pp. 37–57.

Broersma, Mirjam, Diana Carter, and Daniel J. Acheson (2016). “Cognate costs in bilingual speech production: Evidence from language switching”. In: *Frontiers in Psychology* 7.1461.

Broersma, Mirjam, Diana Carter, Kevin Donnelly, and Agnieszka Konopka (2020). “Triggered codeswitching: Lexical processing and conversational dynamics”. In: *Bilingualism: Language and Cognition* 23, pp. 295–308.

Broersma, Mirjam and Kees De Bot (2006). “Triggered codeswitching: A corpus-based evaluation of the original triggering hypothesis and a new alternative”. In: *Bilingualism: Language and Cognition* 9, pp. 1–13.

Broersma, Mirjam, Ludmila Isurin, Sybrine Bultena, and Kees de Bot (2009). “Triggered code switching: Evidence from Dutch–English and Russian–English bilinguals”. In: *Multidisciplinary approaches to code switching*. John Benjamins. Chap. 5, pp. 103–128.

Bullock, Barbara E. and Almeida Jacqueline Toribio, eds. (2009a). *The Cambridge handbook of linguistic code-switching*. Cambridge University Press.

— (2009b). *Themes in the study of code-switching*. Cambridge University Press.

— (2009c). “Trying to hit a moving target: On the sociophonetics of code-switching”. In: *Multidisciplinary Approaches to Code Switching* 41, pp. 189–206.

Bultena, Sybrine, Ton Dijkstra, and Janet G. van Hell (2014). “Cognate effects in sentence context depend on word class, L2 proficiency, and task”. In: *The Quarterly Journal of Experimental Psychology* 67, pp. 1214–1241.

— (2015). “Switch cost modulations in bilingual sentence processing: Evidence from shadowing”. In: *Language, Cognition and Neuroscience* 30, pp. 586–605.

Callahan, Laura (2002). “The matrix language frame model and Spanish/English codeswitching in fiction”. In: *Language & Communication* 22.1, pp. 1–16.

Calvillo, Jesús, Harm Brouwer, and Matthew W Crocker (2016). "Connectionist semantic systematicity in language production". In: *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX, pp. 2555–2560.

Chang, Franklin (2002). "Symbolically speaking: A connectionist model of sentence production". In: *Cognitive Science* 26, pp. 609–651.

— (2009). "Learning to order words: A connectionist model of heavy NP shift and accessibility effects in Japanese and English". In: *Journal of Memory and Language* 61.3, pp. 374–397.

Chang, Franklin, Michael Baumann, Sandra Pappert, and Hartmut Fitz (2015). "Do lemmas speak German? A verb position effect in German structural priming". In: *Cognitive Science* 39.5, pp. 1113–1130.

Chang, Franklin, Gary S. Dell, and Kathryn Bock (2006). "Becoming syntactic". In: *Psychological Review* 113.2, p. 234.

Chang, Franklin, Gary S. Dell, Kathryn Bock, and Zenzi M. Griffin (2000). "Structural priming as implicit learning: A comparison of models of sentence production". In: *Journal of psycholinguistic research* 29.2, pp. 217–230.

Clyne, Michael G. (1967). *Transference and triggering: Observations on the language assimilation of postwar German-speaking migrants in Australia*. Martinus Nijhoff The Hague.

— (1980). "Triggering and language processing". In: *Canadian Journal of Psychology/Revue canadienne de psychologie* 34.4, p. 400.

— (2003). *Dynamics of language contact: English and immigrant languages*. Cambridge University Press.

Cohen, Jonathan D., Kevin Dunbar, and James L. McClelland (1990). "On the control of automatic processes: a parallel distributed processing account of the Stroop effect". In: *Psychological review* 97.3, p. 332.

Davidson, Brad (1996). "'Pragmatic weight' and Spanish subject pronouns: The pragmatic and discourse uses of 'tú' and 'yo' in spoken Madrid Spanish". In: *Journal of Pragmatics* 26, pp. 543–565.

Dell, Gary S. and Franklin Chang (2014). "The P-chain: Relating sentence production and its disorders to comprehension and acquisition". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1634, pp. 1–9.

Dell, Gary S., Franklin Chang, and Zenzi M. Griffin (1999). "Connectionist models of language production: Lexical access and grammatical encoding". In: *Cognitive Science* 23.4, pp. 517–542.

Deuchar, Margaret, Peredur Davies, Jon Herring, María del Carmen Parafita Couto, and Diana Carter (2014). “Building bilingual corpora”. In: *Advances in the Study of Bilingualism*. Ed. by E. Thomas and I. Mennen. Bristol, UK: Multilingual Matters, pp. 93–111.

Dewaele, Jean-Marc and Li Wei (2014a). “Attitudes towards code-switching among adult mono-and multilingual language users”. In: *Journal of Multilingual and Multicultural Development* 35.3, pp. 235–251.

— (2014b). “Intra-and inter-individual variation in self-reported code-switching patterns of adult multilinguals”. In: *International Journal of Multilingualism* 11.2, pp. 225–246.

Dijkstra, Ton and Walter J.B. van Heuven (1998). “The BIA model and bilingual word recognition”. In: *Localist connectionist approaches to human cognition*, pp. 189–225.

— (2002). “The architecture of the bilingual word recognition system: From identification to decision”. In: *Bilingualism. Language and Cognition* 5, pp. 175–197.

Dijkstra, Ton, Alexander Wahl, Franka Buytenhuijs, Nino van Halem, Zina Al-Jibouri, Marcel De Korte, and Steven Rekké (2019). “Multilink: a computational model for bilingual word recognition and word translation”. In: *Bilingualism: Language and Cognition* 22.4, pp. 657–679.

Dussias, Paola E. (2003). “Spanish–English code mixing at the auxiliary phrase: evidence from eye-movement data”. In: *Revista Internacional de Lingüistica Iberoamericana* 1.2, pp. 7–34.

Elman, Jeffrey L. (1990). “Finding structure in time”. In: *Cognitive Science* 14.2, pp. 179–211.

Fernandez, Carla B., Kaitlyn A. Litcofsky, and Janet G. van Hell (2019). “Neural correlates of intra-sentential code-switching in the auditory modality”. In: *Journal of Neurolinguistics* 51, pp. 17–41.

Filippi, Roberto, Themis Karaminis, and Michael S.C. Thomas (2014). “Language switching in bilingual production: Empirical data and computational modelling”. In: *Bilingualism: Language and Cognition* 17.2, pp. 294–315.

Fishman, Joshua A., Robert Leon Cooper, and Roxana Ma Newman (1971). *Bilingualism in the Barrio*. Vol. 1. Indiana University.

Fitz, Hartmut and Franklin Chang (2017). “Meaningful questions: The acquisition of auxiliary inversion in a connectionist model of sentence production”. In: *Cognition* 166, pp. 225–250.

— (2019). “Language ERPs reflect learning through prediction error propagation”. In: *Cognitive Psychology* 111, pp. 15–52.

Fodor, Jerry A. and Zenon W. Pylyshyn (1988). “Connectionism and cognitive architecture: A critical analysis”. In: *Cognition* 28.1-2, pp. 3–71.

Frank, Stefan L. (2021). “Toward computational models of multilingual sentence processing”. In: *Language Learning* 71.S1, pp. 193–218.

Frank, Stefan L., Padraic Monaghan, and Chara Tsoukala (2019). “Neural network models of language acquisition and processing”. In: *Human Language: from Genes and Brains to Behavior*. Ed. by Peter Hagoort. The MIT Press, pp. 277–281.

French, Robert M. (1998). “A simple recurrent network model of bilingual memory”. In: *Proceedings of the 20th Annual Conference of the Cognitive Science Society*, pp. 368–373.

Fricke, Melinda and Gerrit Jan Kootstra (2016). “Primed codeswitching in spontaneous bilingual dialogue”. In: *Journal of Memory and Language* 91, pp. 181–201.

Garrett, Merrill F. (1988). “Processes in language production”. In: *Linguistics: The Cambridge survey* 3, pp. 69–96.

Genesee, Fred, Elena Nicoladis, and Johanne Paradis (1995). “Language differentiation in early bilingual development”. In: *Journal of child language* 22.3, pp. 611–631.

Giancaspro, David (2015). “Code-switching at the auxiliary-VP boundary: A comparison of heritage speakers and L2 learners”. In: *Linguistic Approaches to Bilingualism* 5.3, pp. 379–407.

Gilbert, Sam J. and Tim Shallice (2002). “Task switching: A PDP model”. In: *Cognitive Psychology* 44.3, pp. 297–337.

Glorot, Xavier and Yoshua Bengio (2010). “Understanding the difficulty of training deep feedforward neural networks”. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 249–256.

Goldrick, Matthew, Michael Putnam, and Lara Schwarz (2016). “Coactivation in bilingual grammars: A computational account of code mixing”. In: *Bilingualism: Language and Cognition* 19, pp. 857–876.

Gollan, Tamar H. and Victor S. Ferreira (2009). “Should I stay or should I switch? A cost–benefit analysis of voluntary language switching in young and aging bilinguals”. In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35.3, p. 640.

Gonen, Hila and Yoav Goldberg (2019). “Language Modeling for Code-Switching: Evaluation, Integration of Monolingual Data, and Discriminative Training”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Pro-*

cessing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 4166–4176.
Gordon, Jean K. and Gary S. Dell (2003). “Learning to divide the labor: An account of deficits in light and heavy verb production”. In: *Cognitive Science* 27.1, pp. 1–40.
Green, David W. (2018). “Language control and code-switching”. In: *Languages* 3.2, p. 8.
Green, David W. and Li Wei (2014). “A control process model of code-switching”. In: *Language, Cognition and Neuroscience* 29.4, pp. 499–511.
Grosjean, François (1997). “Processing mixed language: Issues, findings, and models”. In: *Tutorials in Bilingualism: Psycholinguistic Perspectives*, pp. 225–254.
— (2001). “The bilingual’s language modes”. In: *One mind, two languages: Bilingual language processing*. Ed. by Janet L. Nicol. Oxford: Blackwell, pp. 1–22.
Guzmán, Gualberto A., Joseph Ricard, Jacqueline Serigos, Barbara E. Bullock, and Almeida Jacqueline Toribio (2017). “Metrics for Modeling Code-Switching Across Corpora”. In: *INTERSPEECH*, pp. 67–71.
Guzzardo Tamargo, Rosa E. and Paola E. Dussias (2013). “Processing of Spanish-English code-switches by late bilinguals”. In: *Proceedings of the Annual Boston University Conference on Language Development. Boston University Conference on Language Development*. Vol. 37. NIH Public Access, pp. 134–146.
Guzzardo Tamargo, Rosa E., Jorge R. Valdés Kroff, and Paola E. Dussias (2016). “Examining the relationship between comprehension and production processes in code-switched language”. In: *Journal of Memory and Language* 89, pp. 138–161.
Hartsuiker, Robert J., Sarah Bernolet, Sofie Schoonbaert, Sara Speybroeck, and Dieter Vanderelst (2008). “Syntactic priming persists while the lexical boost decays: Evidence from written and spoken dialogue”. In: *Journal of Memory and Language* 58, pp. 214–238.
Heath, Jeffrey G. (1984). “Language contact and language change”. In: *Annual Review of Anthropology* 13.1, pp. 367–384.
Hinaut, Xavier, Florian Lance, Colas Droin, Maxime Petit, Gregoire Pointeau, and Peter Ford Dominey (2015). “Corticostriatal response selection in sentence production: Insights from neural network simulation with reservoir computing”. In: *Brain and Language* 150, pp. 54–68.
Isurin, Ludmila, Donald Winford, and Kees De Bot (2009). *Multidisciplinary approaches to code switching*. Vol. 41. John Benjamins Publishing.

Janciauskas, Marius and Franklin Chang (2018). “Input and age-dependent variation in second language learning: A connectionist account”. In: *Cognitive Science* 42, pp. 519–554.

Johnson, Melvin, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. (2017). “Google’s multilingual neural machine translation system: Enabling zero-shot translation”. In: *Transactions of the Association for Computational Linguistics* 5, pp. 339–351.

Joshi, Aravind (1982). “Processing of sentences with intra-sentential code-switching”. In: *Coling 1982: Proceedings of the Ninth International Conference on Computational Linguistics*.

Karaminis, Themis and Michael S.C. Thomas (2010). “A cross-linguistic model of the acquisition of inflectional morphology in English and Modern Greek”. In: *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*. Ed. by S. Ohlsson and R. Catrambone. Austin, TX, pp. 730–735.

Khoe, Yung Han, Chara Tsoukala, Gerrit Jan Kootstra, and Stefan L. Frank (2020). “Modeling cross-language structural priming in sentence production”. In: *Proceedings of the 18th Annual Meeting of the International Conference on Cognitive Modeling*.

Klementiev, Alexandre, Ivan Titov, and Binod Bhattarai (2012). “Inducing crosslingual distributed representations of words”. In: *Proceedings of COLING 2012*, pp. 1459–1474.

Kootstra, Gerrit Jan, Janet G. van Hell, and Ton Dijkstra (2010). “Syntactic alignment and shared word order in code-switched sentence production: Evidence from bilingual monologue and dialogue”. In: *Journal of Memory and Language* 63.2, pp. 210–231.

— (2012). “Priming of code-switches in sentences: The role of lexical repetition, cognates, and language proficiency”. In: *Bilingualism: Language and Cognition* 15, pp. 797–819.

Lantto, Hanna (2012). “Grammatical code-switching patterns of early and late Basque-Spanish bilinguals”. In: *Sociolinguistic Studies* 6.1, p. 21.

Lederberg, Amy R. and Cesáreo Morales (1985). “Code switching by bilinguals: Evidence against a third grammar”. In: *Journal of Psycholinguistic Research* 14.2, pp. 113–136.

Li, Ping and Igor Farkaš (2002). “A self-organizing connectionist model of bilingual processing”. In: *Advances in Psychology*. Vol. 134. Elsevier, pp. 59–85.

Li, Ping, Xiaowei Zhao, and Brian MacWhinney (2007). “Dynamic self-organization and early lexical development in children”. In: *Cognitive Science* 31.4, pp. 581–612.

Li, Ying and Pascale Fung (2014). “Language modeling with functional head constraint for code switching speech recognition”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 907–916.

Lipski, John M. (1978). “Code-switching and the problem of bilingual competence”. In: *Aspects of Bilingualism*. Ed. by Michel Paradis, pp. 250–264.

— (2019). “Field-testing code-switching constraints: A report on a strategic languages project”. In: *Languages* 4.1, p. 7.

Litcofsky, Kaitlyn A. and Janet G. van Hell (2017). “Neural correlates of intra-sentential codeswitching: Switching direction affects switching costs”. In: *Neuropsychologia* 97, pp. 112–139.

MacDonald, Maryellen C. (2013). “How language production shapes language form and comprehension”. In: *Frontiers in Psychology* 4, p. 226.

MacSwan, Jeff (2005). “Précis of a minimalist approach to intrasentential code switching”. In: *Italian Journal of Linguistics* 17.1, p. 55.

— (2014). *Grammatical theory and bilingual codeswitching*. MIT Press.

Malhotra, Sunita (1980). “Hindi-English, Code Switching and Language Choice in Urban, Uppermiddle-class Indian Families”. In: *Kansas Working Papers in Linguistics* 5, pp. 39–46.

Matras, Yaron (2013). “Languages in contact in a world marked by change and mobility”. In: *Revue française de linguistique appliquée* 18.2, pp. 7–13.

Menacer, Mohamed Amine, David Langlois, Denis Jouvet, Dominique Fohr, Odile Mella, and Kamel Smaıli (2019). “Machine translation on a parallel code-switched corpus”. In: *Canadian Conference on Artificial Intelligence*. Springer, pp. 426–432.

Meuter, Renata F.I. and Alan Allport (1999). “Bilingual language switching in naming: Asymmetrical costs of language selection”. In: *Journal of Memory and Language* 40.1, pp. 25–40.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean (2013). “Efficient estimation of word representations in vector space”. In: *arXiv preprint arXiv:1301.3781*.

Moss, Gillian (1992). “Cognate recognition: Its importance in the teaching of ESP reading courses to Spanish speakers”. In: *English for specific purposes* 11, pp. 141–158.

Munarriz-Ibarrola, Amaia, María del Carmen Parafita Couto, and Emma Vanden Wyngaerd (2018). "Methodologies for intra-sentential code-switching research". In: *Linguistic Approaches to Bilingualism* 8, pp. 1–4.

Muñoz, Carmen (1991). "Why are HE and SHE a problem for Spanish learners of English?" In: *Revista Española de Lingüıstica Aplicada*, pp. 129–136.

Muysken, Pieter (2000). *Bilingual Speech: A Typology of Code-Mixing*. Cambridge University Press.

Myers-Scotton, Carol (1993). "Common and uncommon ground: Social and structural factors in codeswitching". In: *Language in Society* 22, pp. 475–503.

Odlin, Terence (1989). *Language transfer: Cross-linguistic influence in language learning*. Cambridge: CUP.

Petersen, Jennifer (1988). "Word-internal code-switching constraints in a bilingual child's grammar". In: *Linguistics* 26, pp. 479–493.

Pfaff, Carol W. (1979). "Constraints on language mixing: intrasentential code-switching and borrowing in Spanish/English". In: *Language* 55, pp. 291–318.

Poplack, Shana (1978). *Syntactic structure and social function of code-switching*. Vol. 2. Centro de Estudios Puertorriqueños, City University of New York.

— (1980). "Sometimes I'll start a sentence in Spanish y termino en Español: toward a typology of code-switching". In: *Linguistics* 18.7-8, pp. 581–618.

Poulisse, Nanda (1999). *Slips of the tongue: Speech errors in first and second language production*. Vol. 20. Amsterdam: John Benjamins.

Pratapa, Adithya, Gayatri Bhat, Monojit Choudhury, Sunayana Sitaram, Sandipan Dandapat, and Kalika Bali (2018). "Language modeling for code-mixing: The role of linguistic theory based synthetic data". In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1543–1553.

Řehůřek, Radim and Petr Sojka (May 2010). "Software Framework for Topic Modelling with Large Corpora". In: *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. Valletta, Malta: ELRA, pp. 45–50.

Romaine, Suzanne (1986). "The syntax and semantics of the code-mixed compound verb in Panjabi/English bilingual discourse". In: *Language and Linguistics: The Interdependence of Theory, Data and Application*. Ed. by Deborah Tannen and J.E. Alatis, pp. 35–49.

Ross, John Robert (1967). "Constraints on variables in syntax". PhD thesis. Cambridge, MA: Massachusetts Institute of Technology.

Rumelhart, D. E., G. E. Hinton, and R. J. Williams (1986). "Learning internal representations by error propagation". In: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1*. Ed. by David E. Rumelhart, James L. McClelland, and The PDP Research Group. Cambridge, MA: MIT Press, pp. 318–362. ISBN: 0-262-68053-X.

Sankoff, David and Sylvie Mainville (1986). "Code-switching of context-free grammars". In: *Theoretical Linguistics* 13.1-2, pp. 75–90.

Santesteban, Mikel and Albert Costa (2016). "Are cognate words "special"?" In: *Cognitive control and consequences of multilingualism* 2, p. 97.

Sert, Olcay (2005). "The functions of code-switching in ELT classrooms". In: *The Internet TESL Journal* 11.8.

Shan, Changhao, Chao Weng, Guangsen Wang, Dan Su, Min Luo, Dong Yu, and Lei Xie (2019). "Investigating end-to-end speech recognition for Mandarin-English code-switching". In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 6056–6060.

Smolensky, Paul, Matthew Goldrick, and Donald Mathis (2014). "Optimization and quantization in gradient symbol systems: a framework for integrating the continuous and the discrete in cognition". In: *Cognitive Science* 38.6, pp. 1102–1138.

Solorio, Thamar, Elizabeth Blair, Suraj Maharjan, Steven Bethard, Mona Diab, Mahmoud Ghoneim, Abdelati Hawwari, Fahad AlGhamdi, Julia Hirschberg, Alison Chang, et al. (2014). "Overview for the first shared task on language identification in code-switched data". In: *Proceedings of the First Workshop on Computational Approaches to Code Switching*, pp. 62–72.

Solorio, Thamar and Yang Liu (2008). "Learning to predict code-switching points". In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 973–981.

Soto, Victor, Nishmar Cestero, and Julia Hirschberg (2018). "The role of cognate words, POS tags and entrainment in code-switching". In: *Interspeech*, pp. 1938–1942.

Thomas, Michael S.C. and Walter J.B. van Heuven (2005). "Computational models of bilingual comprehension". In: *Handbook of bilingualism: Psycholinguistic approaches*, pp. 202–225.

Timm, Leonora (1975). "Spanish–English code-mixing: el porqué [sic] and how-not-to". In: *Romance Philology* 28, pp. 473–482.

Toribio, Almeida Jacqueline (2001). "On the emergence of bilingual code-mixing competence". In: *Bilingualism: Language and Cognition* 4, pp. 203–231.

Tsoukala, Chara, Mirjam Broersma, Antal van den Bosch, and Stefan L. Frank (2021). "Simulating code-switching using a neural network model of bilingual sentence production". In: *Computational Brain & Behavior* 4, pp. 87–100.

Tsoukala, Chara, Stefan L. Frank, and Mirjam Broersma (2017). ""He's pregnant": simulating the confusing case of gender pronoun errors in L2". In: *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Ed. by G. Gunzelmann, A. Howes, Tenbrink T., and E.J. Davelaar. London, UK, pp. 3392–3397.

Tsoukala, Chara, Stefan L. Frank, Antal van den Bosch, Jorge Valdés Kroff, and Mirjam Broersma (2019). "Simulating Spanish-English code-switching: El modelo está generating code-switches". In: *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*. Minneapolis, Minnesota: Association for Computational Linguistics, pp. 20–29.

— (2020). "Modeling the auxiliary phrase asymmetry in code-switched Spanish–English". In: *Bilingualism: Language and Cognition*. DOI: 10.1017/S1366728920000449.

Twomey, Katherine E., Franklin Chang, and Ben Ambridge (2014). "Do as I say, not as I do: A lexical distributional account of English locative verb class acquisition". In: *Cognitive Psychology* 73, pp. 41–71.

Wang, Zhongqing, Yue Zhang, Sophia Lee, Shoushan Li, and Guodong Zhou (2016). "A bilingual attention network for code-switched emotion prediction". In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 1624–1634.

White, Joanna, Carmen Muñoz, and Laura Collins (2007). "The his/her challenge: Making progress in a 'regular' L2 programme". In: *Language Awareness* 16, pp. 278–299.

Yılmaz, Emre, Henk van den Heuvel, and David van Leeuwen (2016). "Investigating bilingual deep neural networks for automatic recognition of code-switching Frisian speech". In: *Procedia Computer Science* 81, pp. 159–166.

A

Appendix

A.1 English summary

Code-switching occurs when a multilingual speaker alternates (switches) between their languages in a single conversation context (for instance, in a single sentence). Code-switching is done by many people who speak more than one language, and not only by people who grew up in bilingual communities; it is an important aspect of human language. Despite that, we still do not fully understand the underlying processes of this phenomenon. The goal of this dissertation was to come closer to understanding how code-switching works; I aimed to do so by simulating sentence production using computational cognitive modeling.

Specifically, In Chapter 2 I worked with a computational cognitive model, which is based on the so-called recurrent neural network architecture, to simulate sentence production in simultaneous Spanish-English bilinguals (namely, people who grew up with both languages). The goal was to investigate whether the model would code-switch without being exposed to code-switched sentences. The model indeed produced code-switches even without any exposure to such input and the patterns of code-switches are in line with what human bilinguals produce according to earlier linguistic work. To my knowledge, this is the first computational cognitive model that aims to simulate code-switched sentence production.

In Chapter 3, apart from early balanced Spanish–English bilinguals, I simulated late speakers of English who have Spanish as a dominant native language, and late speakers of Spanish who have English as a dominant native language. The simulations predicted how code-switching patterns differ between early balanced and late non-balanced bilinguals; the balanced bilingual simulation code-switched considerably more frequently, which is in line with what human bilinguals do. Additionally, I compared the patterns produced by the simulations to those that human bilinguals produce and identified noticeable commonalities and differences. Having shown that the model produces patterns that are somewhat realistic, in Chapter 4 I used this model to understand why Spanish–English bilinguals produce a code-switch between an auxiliary verb (“to have”, “to be”) and a verb in a certain grammatical structure (the progressive structure, e.g., “Maria está voting”; “Maria is voting”) but not so much in a similar structure called the perfect structure (“Maria ha voted”; “Maria has voted”). This is called the “auxiliary phrase asymmetry”. Using the model and the code-switches it produced, I was able to identify a possible reason for this asymmetry, which has to do with the way the Spanish auxiliary verb *haber* (“to have”) is used in Spanish.

In Chapter 5 I focused on cognates, which are words that share similar meaning and form between languages, such as 'actor' or 'artificial' in English and Spanish. Cognates are said to be "trigger words", which means that if a bilingual speaker hears or produces a cognate word, they are more likely to code-switch; this is called the cognate triggering effect. In this chapter, I investigated the cognate triggering effect on balanced and non-balanced bilinguals, as well as the effect that the percentage of cognates in a language pair has on code-switching. The model predicted no cognate triggering effect for the balanced bilingual model, whereas it did show the effect for the L1 English model, but not for the L1 Spanish one. As, to the best of my knowledge, there are no available data comparing the cognate triggering effect in the sentence production of balanced bilinguals and non-balanced bilinguals with different L1s, these predictions remain to be tested in humans. As for the amount of cognate words, according to the model the frequency of cognates in a language pair does not affect the amount of code-switching.

Last, in Chapter 6 I showed that the model can also be employed to explain other bilingual phenomena, such as mistakes that non-native speakers produce. I used the model to show that a subject pronoun gender error (e.g., 'he' instead of 'she'), which is sometimes produced by Spanish speakers when speaking English, can be attributed to the fact that in Spanish the equivalent of 'he' and 'she' (él/ella) can be omitted in a sentence.

Overall, in this dissertation I have argued that computational cognitive modeling is a useful and novel method to research bilingualism and code-switching. Furthermore, the model learned to code-switch without exposure to code-switched input, which indicates that part of code-switching can be attributed to the distribution of the languages involved, and not only on community-based influence. This is important when answering questions regarding whether a pattern emerges due to internal factors or community practices.

A.2 Nederlandse samenvatting

Codewisseling vindt plaats wanneer een meertalige spreker in één gesprekscontext (bijvoorbeeld in één zin) afwisselt tussen zijn of haar talen. Veel mensen die meer dan één taal spreken, en niet alleen mensen die in tweetalige gemeenschappen zijn opgegroeid, doen aan codewisselen; het is een belangrijk aspect van menselijke taal. Desondanks begrijpen we de onderliggende processen van dit fenomeen nog steeds niet volledig. Het doel van dit proefschrift was om beter te begrijpen hoe codewisseling werkt; ik heb geprobeerd dit te doen door zinproductie te simuleren met behulp van computationele cognitieve modellen.

Om preciezer te zijn, heb ik in Hoofdstuk 2 gewerkt met een computationeel cognitief model dat gebaseerd is op de zogenaamde recurrent neuraal netwerk architectuur, om zinsproductie te simuleren in simultaan Spaans-Engels tweetaligen (dat wil zeggen, mensen die zijn opgegroeid met beide talen). Het doel was te onderzoeken of het model codewisselingen zou produceren zonder te zijn blootgesteld aan zinnen met codewisseling. Het model produceerde inderdaad codewisselingen, zelfs zonder blootstelling aan dergelijke input, en de patronen van codewisseling zijn vergelijkbaar met de producties van menselijke tweetaligen volgens eerder taalkundig onderzoek. Voor zover ik weet, is dit het eerste computationeel cognitieve model dat als doel heeft om codewisseling in zinnen te simuleren.

In Hoofdstuk 3 heb ik niet alleen gebalanceerde, vroege Spaans-Engels tweetaligen gesimuleerd, maar ook late sprekers van het Engels met Spaans als dominante moedertaal, en late sprekers van het Spaans die Engels als dominante moedertaal hebben. De simulaties voorspelden hoe codewisselingspatronen verschillen tussen gebalanceerde vroege en niet-gebalanceerde late tweetaligen; de gebalanceerde tweetalige simulaties codewisselden aanzienlijk vaker, wat overeenkomt met wat menselijke tweetaligen doen. Bovendien heb ik de patronen die de simulaties produceerden vergeleken met die van menselijke tweetaligen, en heb ik opmerkelijke overeenkomsten en verschillen geïdentificeerd. Nadat ik heb laten zien dat het model patronen produceert die enigszins realistisch zijn, heb ik in Hoofdstuk 4 dit model gebruikt om te begrijpen waarom Spaans-Engels tweetaligen codewisselen tussen een hulpwerkwoord ("to have", "to be") en een hoofdwerkwoord in een bepaalde grammaticale structuur (de "progressive" structuur, zoals "Maria está voting"; "Maria is voting") maar niet zozeer in een vergelijkbare structuur (de voltooide tijd: "Maria ha voted"; "Maria has voted"). Dit wordt de "hulpwerkwoord-asymmetrie" genoemd. Met behulp van het model en de codewisselingen die het produceerde, kon ik een mogelijke reden voor deze asymmetrie identificeren, die te maken heeft met de manier waarop het Spaanse hulpwerkwoord *haber* ("hebben") wordt gebruikt in het Spaans.

In Hoofdstuk 5 concentreerde ik me op cognaten, woorden die eenzelfde betekenis en vorm hebben in verschillende talen, zoals 'actor' of 'artificial' in het Engels en Spaans, of 'effect' in het Engels en Nederlands. Cognaten zouden "triggerwoorden" zijn, wat betekent dat als een tweetalige spreker een cognaat hoort of produceert, hij/zij eerder geneigd is te codewisselen; dit wordt het cognaat-triggeringseffect genoemd. In dit hoofdstuk heb ik het cognaat-triggeringseffect onderzocht bij gebalanceerde en niet-gebalanceerde tweetaligen, evenals het effect dat het percentage cognaten in een talenpaar heeft op codewisseling. Het model voorspelde geen cognaat-triggeringseffect voor het gebalanceerde tweetalige model, terwijl het wel het effect liet zien voor het L1 Engelse model, maar niet voor het L1 Spaanse model. Aangezien er, voor zover ik weet, geen data beschikbaar zijn die het cognaat-triggeringseffect vergelijken in zinsproductie van gebalanceerde tweetaligen en niet-gebalanceerde tweetaligen met verschillende L1's, moeten deze voorspellingen nog getest worden bij mensen. Wat betreft het aantal cognaten, heeft volgens het model de frequentie van cognaten in een talenpaar geen invloed op de hoeveelheid codewisseling.

Ten slotte heb ik in Hoofdstuk 6 laten zien dat het model ook kan worden gebruikt om andere tweetalige fenomenen te verklaren, zoals fouten die niet-moedertaalsprekers maken. Ik gebruikte het model om aan te tonen dat een fout in het geslacht van het voornaamwoord (bijv. 'he' ['hij'] in plaats van 'she' ['zij']), die soms door Spaanstaligen wordt gemaakt wanneer ze Engels spreken, kan worden toegeschreven aan het feit dat in het Spaans het equivalent van 'hij' en 'zij' (él/ella) in een zin kan worden weggelaten.

In dit proefschrift heb ik betoogd dat computationele cognitieve modellering een nuttige en nieuwe methode vormt om tweetaligheid en codewisseling te onderzoeken. Bovendien leerde het model codewisselingen zonder blootstelling aan input met codewisseling, wat aangeeft dat codewisseling deels kan worden toegeschreven aan de patronen in de betrokken talen, en niet alleen aan invloed vanuit de gemeenschap. Dit is belangrijk bij het beantwoorden van de vraag of een patroon ontstaat door interne factoren of door praktijken in de gemeenschap.

A.3 Publications

Thesis-related publications

Journal papers

- Tsoukala, Chara, Mirjam Broersma, Antal van den Bosch, and Stefan L. Frank (2021). "Simulating code-switching using a neural network model of bilingual sentence production". In: *Computational Brain & Behavior* 4, pp. 87–100
- Tsoukala, Chara, Stefan L. Frank, Antal van den Bosch, Jorge Valdés Kroff, and Mirjam Broersma (2020). "Modeling the auxiliary phrase asymmetry in code-switched Spanish–English". In: *Bilingualism: Language and Cognition*. DOI: 10.1017/S1366728920000449

Conference proceedings

- Tsoukala, Chara, Stefan L. Frank, and Mirjam Broersma (2017). ""He's pregnant": simulating the confusing case of gender pronoun errors in L2". In: *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Ed. by G. Gunzelmann, A. Howes, Tenbrink T., and E.J. Davelaar. London, UK, pp. 3392–3397
- Tsoukala, Chara, Stefan L. Frank, Antal van den Bosch, Jorge Valdés Kroff, and Mirjam Broersma (2019). "Simulating Spanish-English code-switching: El modelo está generating code-switches". In: *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*. Minneapolis, Minnesota: Association for Computational Linguistics, pp. 20–29
- Khoe, Yung Han, Chara Tsoukala, Gerrit Jan Kootstra, and Stefan L. Frank (2020). "Modeling cross-language structural priming in sentence production". In: *Proceedings of the 18th Annual Meeting of the International Conference on Cognitive Modeling*

Book chapter

- Frank, Stefan L., Padraic Monaghan, and Chara Tsoukala (2019). "Neural network models of language acquisition and processing". In: *Human Language: from Genes and Brains to Behavior*. Ed. by Peter Hagoort. The MIT Press, pp. 277–281

Interactive Machine Translation (IMT-)related publications

- Sanchis-Trilles, G., Alabau, V., Buck, C., Carl, M., Casacuberta, F., García-Martínez, M., ... , Tsoukala, C. & Vidal, E. (2014). Interactive translation prediction versus conventional post-editing in practice: a study with the CasMaCat workbench. *Machine Translation*, 28(3-4), 217–235.

- Koehn, P., Tsoukala, C., & Saint-Amand, H. (2014, June). Refinements to interactive translation prediction based on search graphs. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* (pp. 574–578).

- Alabau, V., Buck, C., Carl, M., Casacuberta, F., García-Martínez, M., Germann, U., ... & Tsoukala, C. (2014, April). CASMACAT: A computer-assisted translation workbench. In *Proceedings of the Demonstrations at the 14th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 25–28).

- Alabau, V., Bonk, R., Buck, C., Carl, M., Casacuberta, F., García-Martínez, M., ..., & Tsoukala, C. (2013). CASMACAT: An open source workbench for advanced computer aided translation. *The Prague Bulletin of Mathematical Linguistics*, 100(1), 101–112.

- Alabau, V., Bonk, R., Buck, C., Carl, M., Casacuberta, F., García-Martínez, M., ... & Tsoukala, C. (2013). Advanced computer aided translation with a web-based workbench. In *2nd Workshop on Post-Editing Technologies and Practice* (pp. 55-62).

A.4 Acknowledgements

Welcome to the most popular part of a thesis!

First of all, I would like to thank my supervisors, Stefan Frank, Mirjam Broersma, and Antal van den Bosch, without whom this dissertation would not have been the same. Our meetings were "interdisciplinary work in action"; your diverse backgrounds led to stimulating conversations and taught me to consider different points of view. But, most importantly, the meetings were fun: I truly enjoyed getting to know you over the years and I feel lucky you were my supervisors!

I am grateful to the manuscript committee and doctoral examination board for agreeing to carefully read my thesis and evaluate my work. I look forward to talking to you at the public defense!

For the research input, I would like to thank the Language and Speech Technology (LST) and the "Cognitive and developmental aspects of multilingualism" groups, also for encouraging me to present my research at various stages. Additionally, I would like to thank Prof. Franklin Chang and Dr. Hartmut Fitz for answering any Dual-path-related questions I had when I first started. I'd also like to thank the cognitive modeling coffee meeting (Stefan, Yung, Danny, Jinbiao and Alessandro) for the occasional modeling talks and great coffee. Stefan, thank you for your help with the *nederlandse samenvatting*, which would have remained a sad edit of automatically translated text without your contribution. Yung, thank you for your interest in the model, it feels nice that someone other than me is using the scripts!

Jorge, thank you for supervising me at the University of Florida and for being such a great host! You made me feel at home at the department and the lab meetings, and on top of that you (and Kevin and Sasha) showed me around to make sure I see manatees, birds and alligators. I never thought I'd end up thinking that alligators are cute (of all adjectives)! Florida was a truly unique experience.

In Nijmegen I was lucky enough to be part of the International Max Plank Research school (IMPRS) and the Language in Interaction (LiI) consortium, both of which introduced me to PhD students and research from all sorts of (psycho)linguistic and neuroscientific domains. Kevin, you are a great IMPRS coordinator; thank you for the educational content and the pizza nights! As for the fellow PhD students, I am especially grateful to Alessandro, Marvin (also for the modeling chats and for sharing this LaTeX template with me), Lara, Xiaochen and Izabela for the social part! Also to Joe, even though we are not in touch, for helping me get going (i.e., not panic) on the

zipline at the Klimpark in Amsterdam (Marvin, that *heidag* was your brilliant idea and you didn't even join!)

From the department I'd like to thank Ingeborg, Hongling, Antje, Laura, Afrooz for the international lunch/coffee breaks (we should have met more often!), and Susanne for the corridor chats. Special thanks to Ingeborg for sharing not only an office with me for almost 4 years (well, when I wasn't away) but also the ups and downs of the PhD life.

Christina and Alessandro, thank you for being my paranymphs and my closest friends! You both made my life in Nijmegen so much better, and I'm glad you will get to know each other even now. Alessandro, traveling with you to two conferences made the experience so much more fun! I'm looking forward to an actual trip or to hosting you and Panos in Athens. Χριστίνα, το Donders (Discussions) έκανε όντως wonders! Σ'ευχαριστώ για τις βόλτες και τα βιβλιαράκια, ανυπομονώ να βρεθούμε απο κοντά για boxing-ο-aerial και για να γνωρίσω τη Maze σας!

My life in Nijmegen became much more enjoyable due to four more people: Christonikos, Patricia, Aaron, and Calle. Thank you all for the walks and long chats!

Απο τους μη-Ναιμεχενιώτες θα ήθελα καταρχάς να ευχαριστήσω τους γονείς μου. Ο,τι και να γράψω δε θα καλύψει την ευγνωμοσύνη μου για τη στήριξή σας όλα αυτα τα χρόνια. Είστε υπέροχοι, και σας ευχαριστώ που με στηρίζετε ο,τι και να κάνω, ακόμα κι αν (ή ακριβώς επειδή) αυτό περιελάμβανε το να αλλάζω χώρα κάθε λίγα χρόνια με αποτέλεσμα να μη μπορώ να μιλήσω "σωστά" ελληνικά (αν και ελπίζω να σας έχω πείσει πως και το code-switching είναι σωστό!)

Ευχαριστώ επίσης τους "βελγοσυγγενείς" για τις βολτούλες στην Ολλανδία και το Βέλγιο. Ήταν πολύ όμορφο που μέναμε κοντά και μπορούσα απλά να πεταχτώ κάποιο ΣΚ ή τις γιορτές να σας δω! Χαίρομαι που ήρθαμε πιο κοντά και που κάνουμε πολύωρες κλήσεις με τη γιαγιά κάθε Σάββατο.

Amalia (μανίτσι), thank you for the amazing cover, for being a fan of code-switching and for your friendship all these years–you and Aggeliki really make me feel like το τρίτο δίδυμο (ακόμα και χωρίς τις φράτζες). Aggeliki and Petroutso, thanks for all the comté cheese you provide each time I visit you in Grenoble. Ανυπομονώ να γνωρίσω και τη Μάγια απο κοντά και να ακούσω τα γαλλοελληνικά της!

Mike, as Nish said when I told him what my thesis is about, you're one of the reasons I ended up studying code-switching! Πέρα απο αυτό, ευχαριστώ για όλα τα ταξίδια και τα meltdown (που μειώνονται επικίνδυνα όσο περνάνε τα χρόνια). I can't wait to visit you and Nish in Singapore when this is over!

Άγγελε και Μυρτώ, ευχαριστώ για όλα τα Ουτρεχτιανά ΣΚ και για τη φιλοξενία, ειδικά τον πρώτο μήνα (έναν ολόκληρο μήνα!) που δεν έβρισκα σπίτι! Μακάρι να το ανταποδώσω. In retrospect συνειδητοποιώ πως θα ήθελα να βρισκόμαστε πιο συχνά, μου έχετε ήδη λείψει!

Nicole, thank you for all the *uitjes* around the Netherlands! After all these years, I can safely say you were the best roommate I've ever had! I'm looking forward to temporarily becoming roommates again at that long-promised trip (I'll make sure we rent a place without a gas stove, sure!)

Σοφία, Όλγα, Μπρέγιαννη, σας ευχαριστώ για την υπέροχη παρέα στην Αθήνα και την Κεφαλονιά, ειδικά το τελευταίο καραντινοκορονοέτος (επίσης ευχαριστώ που δε ρωτάγατε πότε θα παραδώσω τη διατριβή!)

Και, τέλος, Σπύρο ("αυτό το τραγούδι το'γραψα πριν από 45 χρόνια"), σε ευχαριστώ παραπάνω απ'ο,τι μπορώ να εκφράσω με λόγια για τη στήριξη σου και που είσαι ο τέλειος συνταξιδιώτης (sputnik, όπως μάθαμε στο Kennedy Space Center) στη ζωή και τα ταξίδια.

A.5 Biography

Chara Tsoukala was born in Athens, Greece. She studied Computer Science in Athens, with an exchange semester in Darmstadt, Germany, and she got a Master's degree in "Human-Machine Communication" from the University of Groningen, the Netherlands. For her MSc thesis she moved to Edinburgh, U.K., to work with the Statistical Machine Translation group on Interactive Translation Prediction, namely a tool for translators. She also worked there as a research assistant for the EU project CASMACAT (Cognitive Analysis and Statistical Methods for Advanced Computer Aided Translation) and was later employed at a startup in Lisbon, Portugal, that provides translations using crowd-sourcing techniques. Through translation she got interested in multilingualism, particularly in code-switching and the way multilinguals control their several languages. Following Dr. Stefan Frank's open call for a PhD traineeship on computational cognitive models of psycholinguistic processes, she explored code-switching using cognitive modeling, which led to the current dissertation. During her time as a PhD candidate, she briefly visited Stanford University and she spent three months at the University of Florida under the supervision of Dr. Jorge Valdés Kroff. Chara is currently working as a Machine Learning developer at Sciling.

A.6 MPI Series in Psycholinguistics

1. The electrophysiology of speaking: Investigations on the time course of semantic, syntactic, and phonological processing. *Miranda van Turennout*
2. The role of the syllable in speech production: Evidence from lexical statistics, metalinguistics, masked priming, and electromagnetic midsagittal articulography. *Niels O. Schiller*
3. Lexical access in the production of ellipsis and pronouns. *Bernadette M. Schmitt*
4. The open-/closed-class distinction in spoken-word recognition. *Alette Haveman*
5. The acquisition of phonetic categories in young infants: A self-organising artificial neural network approach. *Kay Behnke*
6. Gesture and speech production. *Jan-Peter de Ruiter*
7. Comparative intonational phonology: English and German. *Esther Grabe*
8. Finiteness in adult and child German. *Ingeborg Lasser*
9. Language input for word discovery. *Joost van de Weijer*
10. Inherent complement verbs revisited: Towards an understanding of argument structure in Ewe. *James Essegbey*
11. Producing past and plural inflections. *Dirk Janssen*
12. Valence and transitivity in Saliba: An Oceanic language of Papua New Guinea. *Anna Margetts*
13. From speech to words. *Arie van der Lugt*
14. Simple and complex verbs in Jaminjung: A study of event categorisation in an Australian language. *Eva Schultze-Berndt*
15. Interpreting indefinites: An experimental study of children's language comprehension. *Irene Krämer*
16. Language-specific listening: The case of phonetic sequences. *Andrea Weber*
17. Moving eyes and naming objects. *Femke van der Meulen*
18. Analogy in morphology: The selection of linking elements in Dutch compounds. *Andrea Krott*
19. Morphology in speech comprehension. *Kerstin Mauth*
20. Morphological families in the mental lexicon. *Nivja H. de Jong*
21. Fixed expressions and the production of idioms. *Simone A. Sprenger*
22. The grammatical coding of postural semantics in Goemai (a West Chadic language of Nigeria). *Birgit Hellwig*
23. Paradigmatic structures in morphological processing: Computational and cross-linguistic experimental studies. *Fermín Moscoso del Prado Martín*

24. Contextual influences on spoken-word processing: An electrophysiological approach. *Daniëlle van den Brink*
25. Perceptual relevance of prevoicing in Dutch. *Petra M. van Alphen*
26. Syllables in speech production: Effects of syllable preparation and syllable frequency. *Joana Cholin*
27. Producing complex spoken numerals for time and space. *Marjolein Meeuwissen*
28. Morphology in auditory lexical processing: Sensitivity to fine phonetic detail and insensitivity to suffix reduction. *Rachèl J. J. K. Kemps*
29. At the same time...: The expression of simultaneity in learner varieties. *Barbara Schmiedtová*
30. A grammar of Jalonke argument structure. *Friederike Lüpke*
31. Agrammatic comprehension: An electrophysiological approach. *Marlies Wassenaar*
32. The structure and use of shape-based noun classes in Miraña (North West Amazon). *Frank Seifart*
33. Prosodically-conditioned detail in the recognition of spoken words. *Anne Pier Salverda*
34. Phonetic and lexical processing in a second language. *Mirjam Broersma*
35. Retrieving semantic and syntactic word properties. *Oliver Müller*
36. Lexically-guided perceptual learning in speech processing. *Frank Eisner*
37. Sensitivity to detailed acoustic information in word recognition. *Keren B. Shatzman*
38. The relationship between spoken word production and comprehension. *Rebecca Özdemir*
39. Disfluency: Interrupting speech and gesture. *Mandana Seyfeddinipur*
40. The acquisition of phonological structure: Distinguishing contrastive from non-contrastive variation. *Christiane Dietrich*
41. Cognitive cladistics and the relativity of spatial cognition. *Daniel B.M. Haun*
42. The acquisition of auditory categories. *Martijn Goudbeek*
43. Affix reduction in spoken Dutch. *Mark Pluymaekers*
44. Continuous-speech segmentation at the beginning of language acquisition: Electrophysiological evidence. *Valesca Kooijman*
45. Space and iconicity in German Sign Language (DGS). *Pamela Perniss*
46. On the production of morphologically complex words with special attention to effects of frequency. *Heidrun Bien*

47. Crosslinguistic influence in first and second languages: Convergence in speech and gesture. *Amanda Brown*
48. The acquisition of verb compounding in Mandarin Chinese. *Jidong Chen*
49. Phoneme inventories and patterns of speech sound perception. *Anita Wagner*
50. Lexical processing of morphologically complex words: An information-theoretical perspective. *Victor Kuperman*
51. A grammar of Savosavo, a Papuan language of the Solomon Islands. *Claudia Wegener*
52. Prosodic structure in speech production and perception. *Claudia Kuzla*
53. The acquisition of finiteness by Turkish learners of German and Turkish learners of French: Investigating knowledge of forms and functions in production and comprehension. *Sarah Schimke*
54. Studies on intonation and information structure in child and adult German. *Laura de Ruiter*
55. Processing the fine temporal structure of spoken words. *Eva Reinisch*
56. Semantics and (ir)regular inflection in morphological processing. *Wieke Tabak*
57. Processing strongly reduced forms in casual speech. *Susanne Brouwer*
58. Ambiguous pronoun resolution in L1 and L2 German and Dutch. *Miriam Ellert*
59. Lexical interactions in non-native speech comprehension: Evidence from electro-encephalography, eye-tracking, and functional magnetic resonance imaging. *Ian FitzPatrick*
60. Processing casual speech in native and non-native language. *Annelie Tuinman*
61. Split intransitivity in Rotokas, a Papuan language of Bougainville. *Stuart Robinson*
62. Evidentiality and intersubjectivity in Yurakaré: An interactional account. *Sonja Gipper*
63. The influence of information structure on language comprehension: A neurocognitive perspective. *Lin Wang*
64. The meaning and use of ideophones in Siwu. *Mark Dingemanse*
65. The role of acoustic detail and context in the comprehension of reduced pronunciation variants. *Marco van de Ven*
66. Speech reduction in spontaneous French and Spanish. *Francisco Torreira*
67. The relevance of early word recognition: Insights from the infant brain. *Caroline Junge*
68. Adjusting to different speakers: Extrinsic normalization in vowel perception. *Matthias J. Sjerps*

69. Structuring language. Contributions to the neurocognition of syntax. *Katrien R. Segaert*
70. Infants' appreciation of others' mental states in prelinguistic communication: A second person approach to mindreading. *Birgit Knudsen*
71. Gaze behavior in face-to-face interaction. *Federico Rossano*
72. Sign-spatiality in Kata Kolok: how a village sign language of Bali inscribes its signing space. *Conny de Vos*
73. Who is talking? Behavioural and neural evidence for norm-based coding in voice identity learning. *Attila Andics*
74. Lexical processing of foreign-accented speech: Rapid and flexible adaptation. *Marijt Witteman*
75. The use of deictic versus representational gestures in infancy. *Daniel Puccini*
76. Territories of knowledge in Japanese conversation. *Kaoru Hayano*
77. Family and neighbourhood relations in the mental lexicon: A cross-language perspective. *Kimberley Mulder*
78. Contributions of executive control to individual differences in word production. *Zeshu Shao*
79. Hearing speech and seeing speech: Perceptual adjustments in auditory-visual processing. *Patrick van der Zande*
80. High pitches and thick voices: The role of language in space-pitch associations. *Sarah Dolscheid*
81. Seeing what's next: Processing and anticipating language referring to objects. *Joost Rommers*
82. Mental representation and processing of reduced words in casual speech. *Iris Hanique*
83. The many ways listeners adapt to reductions in casual speech. *Katja Poellmann*
84. Contrasting opposite polarity in Germanic and Romance languages: Verum Focus and affirmative particles in native speakers and advanced L2 learners. *Giuseppina Turco*
85. Morphological processing in younger and older people: Evidence for flexible dual-route access. *Jana Reifegerste*
86. Semantic and syntactic constraints on the production of subject-verb agreement. *Alma Veenstra*
87. The acquisition of morphophonological alternations across languages. *Helen Buckler*
88. The evolutionary dynamics of motion event encoding. *Annemarie Verkerk*

89. Rediscovering a forgotten language. *Jiyoun Choi*
90. The road to native listening: Language-general perception, language-specific input. *Sho Tsuji*
91. Infants'understanding of communication as participants and observers. *Gudmundur Bjarki Thorgrímsson*
92. Information structure in Avatime. *Saskia van Putten*
93. Switch reference in Whitesands. *Jeremy Hammond*
94. Machine learning for gesture recognition from videos. *Binyam Gebrekidan Gebre*
95. Acquisition of spatial language by signing and speaking children: a comparison of Turkish sign language (TID) and Turkish. *Beyza Sümer*
96. An ear for pitch: on the effects of experience and aptitude in processing pitch in language and music. *Salomi Savvatia Asaridou*
97. lncrementality and Flexibility in Sentence Production. *Maartje van de Velde*
98. Social learning dynamics in chimpanzees: Reflections on (nonhuman) animal culture. *Edwin van Leeuwen*
99. The request system in Italian interaction. *Giovanni Rossi*
100. Timing turns in conversation: A temporal preparation account. *Lilla Magyari*
101. Assessing birth language memory in young adoptees. *Wencui Zhou*
102. A social and neurobiological approach to pointing in speech and gesture. *David Peeters*
103. Investigating the genetic basis of reading and language skills. *Alessandro Gialluisi*
104. Conversation Electrified: The Electrophysiology of Spoken Speech Act Recognition. *Rósa Signý Gisladottir*
105. Modelling Multimodal Language Processing. *Alastair Smith*
106. Predicting language in different contexts: The nature and limits of mechanisms in anticipatory language processing. *Florian Hintz*
107. Situational variation in non-native communication. *Huib Kouwenhoven*
108. Sustained attention in language production. *Suzanne Jongman*
109. Acoustic reduction in spoken-word processing: Distributional, syntactic, morphosyntatic, and orthographic effects. *Malte Viebahn*
110. Nativeness, dominance, and the flexibility of listening to spoken language. *Laurence Bruggeman*
111. Semantic specificity of perception verbs in Maniq. *Ewelina Wnuk*
112. On the identification of FOXP2 gene enhancers and their role in brain development. *Martin Becker*

113. Events in language and thought: The case of serial verb constructions in Avatime. *Rebecca Defina*
114. Deciphering common and rare genetic effects on reading ability. *Amaia Carrión Castillo*
115. Music and language comprehension in the brain. *Richard Kunert*
116. Comprehending Comprehension: Insights from neuronal oscillations on the neuronal basis of language. *Nietzsche H.L. Lam*
117. The biology of variation in anatomical brain asymmetries. *Tulio Guadalupe*
118. Language processing in a conversation context. *Lotte Schoot*
119. Achieving mutual understanding in Argentine Sign Language. *Elizabeth Manrique*
120. Talking Sense: the behavioural and neural correlates of sound symbolism. *Gwilym Lockwood*
121. Getting under your skin: The role of perspective and simulation of experience in narrative comprehension. *Franziska Hartung*
122. Sensorimotor experience in speech perception. *Will Schuerman*
123. Explorations of beta-band neural oscillations during language comprehension: Sentence processing and beyond. *Ashley Lewis*
124. Influences on the magnitude of syntactic priming. *Evelien Heyselaar*
125. Lapse organization in interaction. *Elliott Hoey*
126. The processing of reduced word pronunciation variants by natives and foreign language learners: Evidence from French casual speech. *Sophie Brand*
127. The neighbors will tell you what to expect: Effects of aging and predictability on language processing. *Cornelia Moers*
128. The role of voice and word order in incremental sentence processing. *Sebastian Sauppe*
129. Learning from the (un)expected: Age and individual differences in statistical learning and perceptual learning in speech. *Thordis Neger*
130. Mental representations of Dutch regular morphologically complex neologisms. *Laura de Vaan*
131. Speech production, perception, and input of simultaneous bilingual preschoolers: Evidence from voice onset time. *Antje Stoehr*
132. A holistic approach to understanding pre-history. *Vishnupriya Kolipakam*
133. Characterization of transcription factors in monogenic disorders of speech and language. *Sara Busquets Estruch*
134. Indirect request comprehension in different contexts. *Johanne Tromp*

135. Envisioning Language - An Exploration of Perceptual Processes in Language Comprehension. *Markus Ostarek*
136. Listening for the WHAT and the HOW: Older adults' processing of semantic and affective information in speech. *Juliane Kirsch*
137. Let the agents do the talking: on the influence of vocal tract anatomy on speech during ontogeny and glossogeny. *Rick Janssen*
138. Age and hearing loss effects on speech processing. *Xaver Koch*
139. Vocabulary knowledge and learning: Individual differences in adult native speakers. *Nina Mainz*
140. The face in face-to-face communication: Signals of understanding and non-understanding. *Paul Hömke*
141. Person reference and interaction in Umpila/Kuuku Ya'u narrative. *Clair Hill*
142. Beyond the language given: The neurobiological infrastructure for pragmatic inferencing. *Jana Bašnáková*
143. From Kawapanan to Shawi: Topics in language variation and change. *Luis Miguel Rojas-Berscia*
144. On the oscillatory dynamics underlying speech-gesture integration in clear and adverse listening conditions. *Linda Drijvers*
145. Understanding temporal overlap between production and comprehension. *Amie Fairs*
146. The role of exemplars in speech comprehension. *Annika Nijveld*
147. A network of interacting proteins disrupted in language-related disorders. *Elliot Sollis*
148. Fast speech can sound slow: Effects of contextual speech rate on word recognition. *Merel Maslowski*
149. Reason-giving in everyday activities. *Julija Baranova*
150. Speech planning in dialogue - Psycholinguistic studies of the timing of turn taking. *Mathias Barthel*
151. The role of neural feedback in language unification: How awareness affects combinatorial processing. *Valeria Mongelli*
152. Exploring social biases in language processing. *Sara Iacozza*
153. Vocal learning in the pale spear-nosed bat, Phyllostomus discolor. *Ella Lattenkamp*
154. The effect of language contact on speech and gesture: The case of Turkish-Dutch bilinguals in the Netherlands. *Elif Zeynep Azar*

155. Language and society: How social pressures shape grammatical structure. *Limor Raviv*
156. The moment in between: Planning speech while listening. *Svetlana-Lito Gerakaki*
157. How speaking fast is like running: Modelling control of speaking rate. *Joe Rodd*
158. The power of context: How linguistic contextual information shapes brain dynamics during sentence processing. *René Terporten*
159. Neurobiological models of sentence processing. *Marvin Uhlmann*
160. Individual differences in syntactic knowledge and processing: The role of literacy experience. *Saoradh Favier*
161. Memory for speaking and listening. *Eirini Zormpa*
162. Masculine generic pronouns: Investigating the processing of an unintended gender cue. *Theresa Redl*
163. Properties, structures and operations: Studies on language processing in the brain using computational linguistics and naturalistic stimuli. *Alessandro Lopopolo*
164. Investigating spoken language comprehension as perceptual inference. *Greta Kaufeld*
165. What was that Spanish word again? Investigations into the cognitive mechanisms underlying foreign language attrition. *Anne Mickan*
166. A tale of two modalities: How modality shapes language production and visual attention. *Francie Manhardt*
167. Why do we change how we speak? Multivariate genetic analyses of language and related traits across development and disorder. *Ellen Verhoef*
168. Variation in form and meaning across the Japonic language family with a focus on the Ryukyuan languages. *John Huisman*
169. Bilingual sentence production and code-switching: Neural network simulations. *Chara Tsoukala*