Domain-general inhibitory control in bilingual language switching: item-specific vs whole-language inhibition

Di Zhu, BE(Hons), PGDipTransInter

Supervisor: Associate Professor Paul Sowman

A thesis submitted in partial fulfilment of the requirements for the degree of

MASTER OF RESEARCH

Department of Cognitive Science, Macquarie University

ARC Centre of Excellence in Cognition and its Disorders

October 10, 2016

Table of Contents

Abstract

This thesis examines the role of domain-general inhibitory control in bilingual speech production. It has been suggested that correct language selection in bilingual production relies on inhibitory control, with recent evidence pointing towards the involvement of domain-general mechanisms. In particular, functional neuroimaging studies report activation of the pre-supplementary motor area (pre-SMA), a brain area responsible for domain-general inhibitory control, during language switching tasks. However, it remains unclear whether the pre-SMA plays an essential role as part of the language control network or simply co-activates with it. In this thesis, I investigate the causal relationship between neuronal activity in the pre-SMA and behavioural performance in language switching, by transiently disrupting this brain area using transcranial magnetic stimulation (TMS). It has been proposed that there may be two levels of inhibition in bilingual control: item-specific inhibition and whole-language inhibition. I start with the hypothesis that these are both at work in language switching and one or both may rely on domain-general inhibitory control.

Two experiments were carried out to test this hypothesis. The first was a behavioural experiment which established the presence of these two levels of inhibition in language switching. Mandarin-English bilinguals performed a picture-naming task involving univalent items (always named in a particular language) and bivalent items (same picture requiring responses in different languages on different trials). In this design, the effect of whole-language inhibition was reflected in the performance decrement on switch trials compared to stay trials, while the effect of item-specific inhibition was reflected in the reduced performance for bivalent items compared to univalent items. In the second experiment, I investigated the causal involvement of the pre-SMA in language control. A repetitive TMS protocol was used to achieve transient disruption of this brain region while the same picture-naming task was performed. The impact on behavioural performance was assessed with respect to the two levels of inhibition. Disruption of the pre-SMA was found to modulate item-specific but not whole-language inhibition, suggesting that only item-specific control recruits this domain-general inhibitory mechanism.

Declaration

I declare that this work, either in part or in whole, has not been submitted elsewhere as either part of a degree or for publication, in any other university or institution. The sources of information for this work were wholly from my own research, except where referenced. The writing of this thesis was independent and free from plagiarism. This research was approved by the Macquarie University Ethics Committee (protocol numbers: 5201200035, 5201400585).

_____

Di Zhu

Acknowledgements

This thesis would not have been possible without all the wonderful people who have supported me throughout the past nine months. It has been a great adventure and I owe my gratitude to many.

First and foremost, I would like to thank my principal supervisor Paul Sowman for giving me the opportunity to embark on an academic path and encouraging me to pursue the topic I am interested in. I am extremely grateful for the generous amount of time you have invested into supervising this project, and the help and expertise you provided me with every step along the way - I have benefitted from these immensely. Thank you also for believing in my abilities and keeping me motivated with your sense of humour. I would like to thank my associate supervisor Matthew Finkbeiner for the helpful discussions on project ideas and for making important improvements to my experimental design.

Many others have contributed to this thesis and made it a really enjoyable experience. I am grateful to Yu and Andy, who were excellent sources of advice on specific technical issues as well as doing research in general. I would like to thank Wei, Yanan, David, Hui, and Xuejing, for valuable feedback on the experiment and help with participant recruitment. In addition, I thank Luan, Marion, Dani, Robert, Chris, and Lydia for their friendship and company. I would also like to say thank you to Edmond Lam, who assisted me in running the TMS experiment and in preliminary data checking, and to all my participants, who endured the experimental sessions even when they found the procedure less than pleasant.

Last but not least, I want to thank my family for supporting me in pursuing my dream. Thank you James for your unconditional love and care, for all your seafood dishes, and your patience and understanding for all the late nights and lost weekends. Thanks to my parents for teaching me to think critically and work diligently, and always giving me the best in everything.

Chapter 1: Introduction

Bilingualism is becoming an increasingly widespread phenomenon. It is reported that at least half of the world's population today speak more than one language (French & Jacquet, 2004). Paralleling the importance of this developing phenomenon is a surge of research interest in bilingualism. One key topic that has attracted much attention is how bilinguals masterfully control their two languages during speech production. When communicating with other bilinguals who share both their languages, they naturally adopt a bilingual mode and switch between these languages. Yet, in a monolingual setting, they are able to keep the languages separate, and converse in the required language without intrusions from the other. Previous research suggests that when a bilingual speaker intends to produce a word in one language, the concept activates relevant lexical nodes in both languages (e.g. Hermans, Bongaerts, De Bot, & Schreuder, 1998; Costa, Caramazza, & Sebastian-Galles, 2000; Colomé, 2001). This gives rise to the question of how bilinguals can ensure that only words in the correct language are selected for output. Finkbeiner, Gollan, and Caramazza (2006) called this the "hard problem" in bilingual lexical selection. It has been suggested that correct language selection in bilingual production relies on a control mechanism which is external to the language system and closely related to executive function (see Bobb, Wodniecka, & Kroll, 2013). In this thesis, the above proposal is evaluated through an investigation of the possible role of executive function in two types of control underlying language selection.

This chapter gives an overview of the current state of research on bilingual language control. Section 1.1 introduces the language switching paradigm, a commonly used experimental paradigm in studying language control in bilinguals. Section 1.2 focuses on a prominent model of bilingual speech production, which is built upon a central hypothesis that correct language selection is achieved through inhibition of the non-target language. Supporting evidence for this hypothesis and arguments against it are presented. Section 1.3 reviews evidence relating to the possible neural mechanisms underlying inhibitory processes in language switching, which points

particularly towards the involvement of domain-general inhibitory control. Section 1.4 discusses the proposal that there are two levels of control in language switching, which may be implemented by different neural mechanisms. Section 1.5 gives an outline of the present study, including the research questions it aims to address and the rationale behind the experimental design.

1.1 The language switching paradigm

One of the most commonly used experimental paradigms in studies of bilingual language control is the language switching paradigm. In this paradigm, bilingual participants are asked to name pictures or numerals in their first language (L1) or second language (L2) according to task instructions. The participant's performance on the task is examined in terms of naming latencies and number of production errors. These studies typically include single-language and/or mixed-language blocks. In a single-language block, the same language is required on all trials. In a mixed-language block, the language requirement may vary from one trial to the next. When there is a language change (i.e. a trial requiring a different language than its preceding trial), this trial is referred to as a *switch trial*; when the language requirement stays the same, it is called a *repetition trial* or *stay trial*. Various types of costs have been identified in language switching: *switch cost* is defined as the difference in naming latencies between switch trials and stay trials, and *mixing cost* refers to the difference between stay trials in mixed-language blocks and those in single-language blocks. While single-language blocks are more commonly used as a baseline (e.g. for calculation of mixing cost), some authors have also examined the cost of changing language between single-language blocks (Misra, Guo, Bobb, & Kroll, 2012; Guo, Liu, Misra, & Kroll, 2011).

The majority of language switching studies employ the *cued switching* design (e.g. Meuter & Allport, 1999; Jackson, Swainson, Cunnington, & Jackson, 2001; Costa & Santesteban, 2004; Verhoef, Roelofs, & Chwilla, 2009). In this design, the language requirement is specified using a language cue displayed on screen (e.g. background colours, national flags), and participants are instructed to respond in the language indicated by the cue on each trial. Studies adopting this design usually report significant switch costs (i.e. switch trials being slower than stay trials) and mixing costs (i.e. stay trials in mixed-language blocks being slower than those in single-language blocks). These findings are generally taken as evidence that language switching (or mixing) takes extra time compared to speaking in a single language. However, it is important to note that these differences in naming latencies should not simply be interpreted as the time it takes to switch (or mix) between languages. For example, longer naming latencies on switch trials can be attributed to processes

such as cue encoding[1], instead of (or in addition to) purely switching between languages.

When performing *cued switching*, participants make "forced" language selections according to the cue given. In the *voluntary switching* design, however, participants have the freedom to choose which language to use (and therefore, to switch language or not) on each trial. This means that they can use "whatever language comes to mind first" for each item (Gollan & Ferreira, 2009). As bilinguals may be more familiar with certain names in one language and other names in another language, this design aims to elicit responses according to lexical accessibility. Switch cost reduction and a mixing facilitation effect (on the non-dominant language) have been observed in voluntary-switching studies (Gollan, Kleinman, & Wierenga, 2014, Exp. 2; Zhang et al., 2015; Gollan & Ferreira, 2009). A variation of the voluntary switching design, called *bottom-up switching*, was recently developed by Kleinman and Gollan (2016). In this design, bilinguals are instructed to use whatever language seems easier the first time they see a picture, but to keep using that same language for every subsequent presentation of the same picture. Kleinman and Gollan found that switch cost was eliminated and mixing cost was substantially reduced in bottom-up switching.

1.2 Language control in bilingual speech production

As introduced earlier, the "hard problem" in bilingual language control (Finkbeiner, Gollan, et al., 2006) concerns how bilinguals ensure correct language selection during speech production. Existing accounts of language control propose three types of solutions to address this problem. The first postulates that only target-language lexical nodes are considered by the lexical selection mechanism, while nodes belonging to the non-target language are ignored (Costa & Caramazza, 1999; Costa, Miozzo, & Caramazza, 1999). In other words, the problem of language selection is solved by assuming that lexical selection is language-specific. The second type of solution states that the intention to speak in one language naturally leads to higher

[1] See Logan and Bundesen (2003, Exp. 3 & 4), for evidence on cue change, not task change, being the source of task switching costs; but see also Heikoop, Declerck, Los, and Koch (in press), who recently examined this in the context of language switching and found significant costs for both cue change and language change.

activation of lexical representations in that language, compared to the non-target language (La Heij, 2005; Finkbeiner, Gollan, et al., 2006). In this view, language selection occurs at the conceptual level, in the same way that monolinguals choose between words with similar meanings. The third solution holds that correct language selection is achieved through inhibition of the non-target language (i.e. the "inhibition hypothesis"). This last view is the most prominent in the literature, and will be discussed in detail below.

1.2.1 The inhibition hypothesis of bilingual language control

The inhibition hypothesis was first proposed by Green (1998), in his inhibitory control model (ICM) of bilingual speech production. This hypothesis derives from Norman and Shallice's (1986) work on the control of actions, as language is "a form of communicative action" (Green, 1998, p. 68). The inhibition account of language control holds that correct language selection is achieved through suppression of lexical nodes in the non-target language. In the ICM, each lexical node is associated with a language tag, which identifies its language membership. Green suggests that activation levels of lexical representations are regulated by language task schemas (e.g. L1 output, L2 output), which are in turn controlled by a supervisory attentional system. When speech output is required in a particular language, say L2, the supervisory attentional system activates the L2 language task schema, which then increases the activation levels of lexical nodes in L2 accordingly. At the same time, the L1 task schema is suppressed, and this suppression also gets passed down to all lexical representations with the L1 language tag. Additionally, the active language task schema (in this case, L2) serves a checking role, by catching any highly activated lexical nodes in the non-target language (i.e. L1) and reactively inhibiting them to prevent them from reaching speech output. According to Green's assumption that more active lexical nodes require stronger suppression, the inhibition hypothesis makes the prediction that lexical nodes in the dominant language will be more suppressed when production takes place in the non-dominant language than vice versa. Green further predicts that this stronger suppression placed on the dominant language will take extra time to overcome when this

language is subsequently required for output. This prediction has been confirmed in many language switching studies, as I will elaborate upon in the next section.

1.2.2 Behavioural evidence for the inhibition hypothesis

In language switching studies, the signature behavioural evidence for inhibition of the non-target language is the asymmetrical switch cost (seminal study by Meuter & Allport, 1999) and reversal of dominance effects (e.g. Costa & Santesteban, 2004; Christoffels, Firk, & Schiller, 2007; Gollan & Ferreira, 2009; Verhoef et al., 2009).

Asymmetrical switch cost refers to the (paradoxically) larger switch cost observed when switching into the dominant language, compared to switching into the non-dominant language. Most of the time, this results in slower responses in the dominant language than in the non-dominant language on switch trials. The original interpretation for this effect, given by Meuter and Allport (1999), is that production in the non-dominant language requires stronger suppression of the dominant language which takes more time to overcome when switching back to the dominant language. This account aligns well with Green's (1998) inhibition hypothesis of language control. The switch cost asymmetry has since been replicated in many studies (e.g. Jackson et al., 2001; Campbell, 2005; Philipp, Gade, & Koch, 2007; Schwieter & Sunderman, 2008), and has become known as the hallmark of inhibitory control in bilingual speech production.

The *language dominance effect* is a robust proficiency effect in the bilingual literature (Hanulová, Davidson, & Indefrey, 2011). It refers to the faster naming time observed when bilinguals name objects in their dominant language compared to naming in their non-dominant language. In the language switching paradigm, however, reversal of this dominance effect (i.e. "*reversed dominance*") has been observed. Specifically, in mixed-language blocks, naming is sometimes slower in the dominant language (on both stay and switch trials), even though single-language blocks are faster in the dominant language. Such a pattern is not universally found in language switching, but it has been observed in a number of studies (e.g. Costa & Santesteban, 2004; Christoffels et al., 2007; Gollan & Ferreira, 2009). This reversed dominance effect is often considered as an indication of sustained inhibition of the dominant language to

facilitate speech production in the non-dominant language (e.g. Gollan et al., 2014; Bobb & Wodniecka, 2013).

1.2.3 Does switch cost asymmetry mean inhibition?

Several authors have put forward alternative views on the interpretation of the switch cost asymmetry. Finkbeiner, Almeida, Janssen, and Caramazza (2006) rejected the inhibition account by showing that switch cost asymmetry was only applicable to bivalent stimuli (to be named in both languages) and not to univalent stimuli (to be named always in the same language). To account for the asymmetrical switch cost on bivalent stimuli, the authors demonstrated, in a single-language task, that the same paradoxical asymmetry can be obtained between easy and difficult responses afforded by the same stimuli. Specifically, it took longer to switch into the easy response, just like the larger cost observed when switching into the dominant language. On the basis of this finding, an alternative account was provided for the switch cost asymmetry. This account argues that responses in the dominant language become available too fast (before the response selection system can finish updating task goal, which is required on switch trials) and are therefore rejected (on switch trials) until the said system is ready, leading to the apparently slower response (see also Finkbeiner & Caramazza, 2006). Verhoef et al. (2009) examined the effect of having longer preparation time for the upcoming language switch, by manipulating the cue-stimulus interval (CSI). They found that switch costs were asymmetrical with short CSI (500ms), but became symmetrical when the CSI was long (1250ms). They propose that the switch cost asymmetry results from the fact that stay trials in the dominant language were exceptionally fast compared to all other trial types (what they term “L1-repeat-benefit”), rather than from inhibition. Philipp et al. (2007) challenged the inhibition hypothesis and provided a persistent activation account as an alternative. They argue that the larger switch cost in the dominant language is due to interference from the strongly activated non-dominant language.

Another problem with the switch cost asymmetry serving as evidence for inhibition is that it has not been reliably observed in all language switching studies. Costa and Santesteban (2004) tested highly proficient bilinguals and these participants showed

symmetrical switch costs between the two languages. Now this is not surprising, as the inhibition hypothesis does predict similar amount of suppression on the two languages if they are of similar dominance. However, these authors went on to test a group of multilinguals who were highly proficient in their L1 and L2 (but not their L3), and it was found that switch costs were not only symmetrical between the highly proficient L1 and L2, but also between L1 and the much weaker L3. This finding goes against the prediction of the inhibition hypothesis, as it demonstrates that symmetrical switch costs can exist between two languages even if they are not of a similar level of proficiency. In a later study, Costa, Santesteban, and Ivanova (2006) further showed that switch costs remained symmetrical in highly proficient bilinguals regardless of the degree of similarity between the pair of languages tested; it also did not matter whether the participants were early or late bilinguals. So it seems that switch costs are symmetrical as long as participants are highly proficient in at least two languages. To account for these findings, Costa et al. (2006) propose that highly proficient bilinguals may have developed a more efficient language control mechanism which does not rely on inhibition.

The elimination of switch cost asymmetry has also been observed in voluntary language switching. As introduced earlier, the voluntary switching design gives participants the freedom to use whichever language comes to mind first when naming each item. Gollan and Ferreira (2009) report symmetrical switch cost in voluntary switching, not only for balanced bilinguals but also for those who are clearly dominant in one language. This finding reveals that, aside from language proficiency, task-related factors can also affect the switch cost asymmetry, further questioning the robustness of this signatory evidence for inhibition in bilingual language production.

1.2.4 Summary: to inhibit or not to inhibit

Despite these challenges to the switch cost asymmetry, it is important to note that the absence of such asymmetry does not necessarily mean there is no inhibition (see Bobb & Wodniecka, 2013, for a detailed discussion). What it does mean perhaps, is that the phenomenon of switch cost asymmetry is not quite a reliable indicator of inhibition after all. Furthermore, when switch costs are found to be

symmetrical, reversed dominance effect is usually present - this is true for most of the studies discussed above (Costa & Santesteban, 2004; Costa et al., 2006; Gollan & Ferreira, 2009). It may be the case that the absence of switch cost asymmetry in these studies was precisely due to the dominant language being slowed down so much on both stay and switch trials, such that the "switch cost" (being the difference between switch and stay trials in the same language) was not significantly larger than that in the non-dominant language; the switch trials, nonetheless, were much slower in the dominant language than in the non-dominant language, consistent with the prediction of the inhibition hypothesis of language control. The significant overall slowing of the dominant language is, in itself, also strong evidence showing sustained inhibition of the dominant language when production is required in a mix of the two languages. In this way, reversed dominance complements asymmetrical switch cost as evidence supporting the presence of inhibition in language switching.

Evidence for the inhibition of non-target language in bilingual production has also been found in a range of studies employing other paradigms, such as the language version of the n-2 repetition paradigm (Philipp & Koch, 2009), the picture-word interference paradigm, and paradigms exploiting language-specific properties such as cognate status (for a comprehensive review, see Kroll, Bobb, Misra, & Guo, 2008). More direct evidence for the presence of inhibitory processes comes from neural studies of bilingual language control, which reveal inhibition-related brain activities during language switching. These will be reviewed in the next section.

1.3 Domain-general inhibitory control in language switching

Supporting evidence for the involvement of inhibitory processes in language switching has emerged from a number of electrophysiological and neuroimaging studies. These findings point particularly towards the close relationship between bilingual language control and domain-general inhibitory control.

1.3.1 Evidence from ERP studies

Jackson et al. (2001) recorded ERPs in a digit naming task with cued language switching and showed that switch trials elicited a larger N2 component than stay trials did. As the N2 component is associated with response inhibition (such as in go/no-go tasks; Jackson, Jackson, & Roberts, 1999), this result suggests that inhibition of the non-target language may be based on a similar mechanism as response inhibition. These authors further report an asymmetry in this effect such that the N2 modulation is significant only when switching from L1 into L2. Crucially, even though switching in this direction involves more inhibition (as suggested by the larger N2), the switch cost (contrasting L2 switch trials with L2 stay trials) was smaller compared to the cost of switching into L1 (contrasting L1 switch trials with L1 stay trials). These findings are in line with the ICM's claim that L2 production requires stronger suppression of L1 than vice versa, and that the switch cost asymmetry reflects the longer time required to overcome this suppression when returning to L1 after speaking L2.

Misra et al. (2012) conducted an ERP study with single-language blocks only. In this study, bilinguals named a set of pictures firstly all in one language, and then all in the other language. An asymmetrical pattern was again found between switching from L1 to L2 and switching in the other direction. When the L1 block occurred first, L2 naming speed improved and produced more positive ERP waveforms, a result that is consistent with a priming facilitation effect (Guillaume et al., 2009). In contrast, when the L2 block occurred first, L1 naming latencies suffered and the ERP waveforms showed more negativity, no longer reflecting any facilitation. The authors suggest that any priming facilitation on L1 (when it followed the L2 block) was cancelled out by the inhibition applied on L1 during L2 naming, resulting in an overall cost. These results demonstrate that the asymmetry in switch costs not only exists in fast language switching from trial to trial, but also affects language change between blocks even though production within each block occurs in a single language. This suggests that there may be sustained L1 inhibition during L2 production.

1.3.2 Evidence from neuroimaging studies

If language selection is achieved through inhibition of the non-target language (especially the more dominant language), which areas of the brain might be responsible for carrying out such inhibition? Abutalebi and Green (2008) developed a neurocognitive model of bilingual control, in which they propose a brain network involving cortical and subcortical structures tightly related to executive function. In a recent update to this model (Green & Abutalebi, 2013), more brain regions specifically associated with domain-general inhibitory control, such as the right inferior frontal gyrus (rIFG) and pre-supplementary motor area (pre-SMA), have also been incorporated into the proposed language control network.

Recent neuroimaging studies on language switching report differential activation of the pre-SMA (De Baene, Duyck, Brass, & Carreiras, 2015; de Bruin, Roelofs, Dijkstra, & FitzPatrick, 2014; Abutalebi et al., 2013). This brain area is widely regarded as a core component in domain-general inhibitory control (Fedorenko, Duncan, & Kanwisher, 2013; Aron et al., 2007; Xue, Aron, & Poldrack, 2008). The activation of this brain area during language switching lends support to the idea that the inhibitory processes which enable correct language selection rely on neural mechanisms similar to those underlying action selection in non-linguistic tasks. Importantly, de Bruin et al. (2014) found significant pre-SMA activation (indicating inhibitory control) when trilinguals switched into their L2 and L3 but not when switching into L1, consistent with the ICM's interpretation of the switch cost asymmetry (i.e. production in non-dominant language requires more inhibition of the dominant language).

What could be the possible role of the pre-SMA in language control? The pre-SMA is being increasingly recognised for its role in response selection and conflict resolution across domains, especially in demanding tasks (for an overview, see Nachev, Kennard, & Husain, 2008). Abutalebi et al. (2012) compared brain activities in bilinguals during the performance of a language switching task and a flanker task (a non-linguistic task requiring conflict resolution), and found similar pre-SMA activation in both tasks. They report that the pre-SMA works alongside the dorsal anterior cingulate cortex (ACC) in carrying out conflict monitoring and error detection. Green and Abutalebi's (2013) neurocognitive model of bilingual language control also

assigns the role of conflict resolution to the ACC/pre-SMA complex. In a recent meta-analysis of fMRI studies on language switching (Luk, Green, Abutalebi, & Grady, 2012), significant activation was found in midline pre-SMA but not in the ACC, suggesting that the pre-SMA may be more universally engaged in bilingual language control.

1.4 Two levels of control in language switching

De Groot and Christoffels (2006) reviewed models of bilingual speech production and noted the distinction between two possible types of language control: *global control*, which affects all lexical representations in a language simultaneously, and *local control*, which only targets specific lexical representations. According to De Groot and Christoffels, some models implicate the presence of both types of control (such as the ICM), where proactive regulation of the activation levels of the two languages on the global level is complemented by reactive inhibition operating at the local level. Such local reactive inhibition serves to catch and suppress any highly activated lexical nodes in the non-target language (despite already being suppressed at the global level).

Some authors hold the view that local control is the more fundamental type of control in bilingual production. De Groot (2011) points out that the ICM would work fine without needing global control at all, since local control alone would be sufficient to prevent any non-target-language lexical nodes from being selected. Van Assche, Duyck, and Gollan (2013) investigated local control and global control using a verbal fluency task, in which bilingual participants were presented with specific letter prompts and were asked to produce words beginning with those letters in the required language. When exemplars are produced in one language, this presumably involves inhibition of the other language, making subsequent production in that language more difficult. If inhibition is local (i.e. only affecting specific non-target-language lexical items which compete for output), this inhibition would only need to be overcome when the same letter prompt was given subsequently for production of exemplars in the other language. Conversely, if inhibition is global (i.e. affecting the non-target language as a whole), the language that comes second would always be affected, even when letter prompts are not repeated across the two languages. The

authors found that fluency decreased in the dominant language for repeated letter prompts (i.e. when exemplars for the same letter has been previously produced in the non-dominant language), indicating the presence of local (item-specific) inhibition. However, for non-repeated letter prompts, fluency decrease (which would indicate global, whole-language inhibition) was only observed in one group of bilinguals. Specifically, Mandarin-English bilinguals showed evidence of whole-language inhibition, but Dutch-English bilinguals did not. These findings suggest that local inhibition may be more universally adopted in language control, while global inhibition might only be used by some bilinguals (e.g. only if their two languages are highly dissimilar).

Misra et al. (2012) looked for evidence of global control in a picture-naming study. In this study, bilinguals named a set of pictures firstly all in one language, and then all in the other language. The crucial manipulation was the block order (i.e. L1 block first, or L2 block first). They showed that naming latencies in L2 benefited from having named the same pictures in L1 first, whereas L1 naming suffered from slower responses when it followed the L2 block. Importantly, the L1 slowing was long-lasting rather than transient at the point of language switch (i.e. in the first few trials upon switching to the L1 block). The authors interpreted this as evidence for global inhibition of L1 during L2 production. However, there is an alternative possibility: the "global slowing" of L1 (following L2 naming) may reflect item-specific inhibition instead. Since the same set of pictures were named in both languages, all the individual lexical representations for these pictures in L1 would have been strongly suppressed when these pictures were named in the L2 block, which leads to "global slowing" in the subsequent L1 block. In other words, the item-specific suppression here gives the illusion of global, persistent inhibition of L1 as a whole because it affected all of the items that were tested.

The terms "global" and "local" have been used by several authors to denote two possible types of control in bilingual speech production. However, different authors have used them to refer to different kinds of distinction. For example, Christoffels et al. (2007) used these terms to contrast between sustained control due to language context (single-language vs. mixed-language production), which they call "global", and transient control on a per trial basis (stay vs. switch trial), which they call "local". Guo et al. (2011), on the other hand, used "local" to denote the effect of language

mixing and “global” to denote comparison across single-language blocks. These terms can also cause confusion for our purpose here, since *global* inhibition of the whole language is reflected in trial-to-trial (sometimes called *local*) switch costs. Therefore, from here onwards I will adopt the terminology used by Van Assche et al. (2013), where *item-specific inhibition* is defined as suppression of individual non-target-language lexical nodes that are in competition with the node to be selected, and *whole-language inhibition* is defined as simultaneous suppression of all lexical nodes in the non-target language.

1.5 The present study

The present study has two aims. The first is to investigate whether both whole-language inhibition and item-specific inhibition are at work during language switching. The majority of studies carried out so far in the language switching paradigm have focused on trial-to-trial switching, i.e. comparing switch trials to stay trials. On a switch trial, the response language changes from the preceding trial, which presumably requires inhibition of the language just used and activation of the currently relevant language (Green, 1998); on a stay trial, the response language stays the same, therefore no such inhibition/activation is required. According to the definitions of the two types of inhibition given above (end of *Section 1.4*), this comparison between stay and switch trials reflects *whole-language inhibition* (because the difference between a stay trial and a switch trial is whether there is a language change, regardless of what individual lexical items are involved on these trials).

In order to examine *item-specific inhibition*, a different type of comparison is needed. To explain the approach used in this thesis, it is helpful to firstly understand two types of stimuli commonly used in language-switching studies: univalent stimuli and bivalent stimuli. Recall that in the cued language switching paradigm, participants follow the language cue on each trial and respond in the required language. A *univalent stimulus* always requires a response in the same language every time it appears. A *bivalent stimulus* has no fixed association with a particular language, and may elicit different responses on different trials. When a bivalent stimulus needs to be named in one language after having been named in another, the previously used

name is likely to be highly activated by the concept, and needs to be suppressed by item-specific inhibition (De Groot & Christoffels, 2006). In contrast, a univalent stimulus is only ever named in one language, so no such item-specific inhibition is required. Thus, analogous to examining whole-language inhibition through the comparison of switch trials to stay trials, item-specific inhibition can be studied by comparing bivalent stimuli to univalent stimuli. The opportunity to make the latter type of comparison is rare in the current literature, as language-switching studies generally use either univalent or bivalent stimuli, but not both.

The only study so far that has combined the use of univalent and bivalent stimuli in a language switching paradigm is Finkbeiner, Almeida, et al. (2006). Interestingly, these authors argued against both whole-language inhibition (what they call "language suppression") and item-specific inhibition (what they call "lexical suppression"), based on a lack of switch cost on univalent stimuli in their results. However, as Abutalebi and Green (2007) point out, this study suffers from a major confound of task switching occurring on all univalent trials. The univalent stimuli used in these experiments were of a different type (pictures) from the bivalent stimuli (digits), and participants were explicitly instructed to treat them differently (naming digits according to language cues and naming pictures always in English). Therefore, naming of univalent stimuli was always accompanied by a task switch, which may have masked the effects of inhibition. In order to take their results as evidence for the absence of inhibition, one must assume linear additivity of switch costs (between the language switch and task switch), which is unlikely to be true. In addition, this study did not provide the means for a direct comparison between whole-language and item-specific inhibition, as these were investigated separately.

A few studies employing other paradigms have examined item-specific and whole-language inhibition side-by-side in the same experiment. However, their findings are diverse. For example, Van Assche et al. (2013) compared these two types of inhibition in a verbal fluency task. They found robust evidence for item-specific inhibition (i.e. with repeated stimuli), while whole-language inhibition (non-repeated stimuli) was only observed in some bilinguals. On the other hand, Philipp and Koch (2009), who used the n-2 language-repetition paradigm, arrived at quite the opposite conclusion. They found that inhibition affected a language globally, regardless of whether the individual stimuli were repeated or not. Moreover, the cost of language

inhibition was reduced on repeated stimuli, possibly reflecting a repetition priming effect. The apparent discrepancy between these findings is not so surprising, as the experimental tasks employed in these studies were very different and it can be expected that different types of control are required in each situation.

A systematic examination and comparison of item-specific and whole-language inhibition, like in those studies above, has yet to be conducted with the language switching paradigm. Experiment 1 in this thesis (presented in Chapter 2) aimed to fill this gap in the literature by combining univalent and bivalent stimuli[2] in a picture-naming task, with no confound between them. In this task, half of the stimuli were univalent, each consistently eliciting responses in the same language every time it appeared; the other half were bivalent, each imposing varied language requirements throughout the experiment. Univalent and bivalent stimuli were mixed together and appeared under the same circumstances. Since all stimuli were pictures and were to be named in the exact same way (i.e. by following the language cue), the valence of each stimulus remained implicit in the context of the naming task. With this design, univalent and bivalent stimuli were virtually indistinguishable from each other (from the participant's perspective), thus it was justified to compare them directly. By enabling such comparison between univalent and bivalent stimuli (i.e. a measurement of item-specific inhibition) alongside the comparison between stay and switch trials (i.e. a measurement of whole-language inhibition) within the same experimental task, these two types of inhibition in language switching were able to be examined simultaneously and compared side-by-side.

The second aim of the present study was to investigate the involvement of domain-general inhibitory control in language switching. The motivation behind this was to verify theoretical accounts in brain models of bilingual language control (e.g. Green & Abutalebi, 2013), and to provide more empirical basis for or against the view that language selection relies on executive function. An increasing amount of neural evidence (reviewed in *Section 1.3*) now suggests that inhibition in language control is accomplished via domain-general mechanisms, and chief among these is the brain area called pre-SMA. While neuroimaging evidence can only reveal an association

[2] When I talk about "univalent" and "bivalent" stimuli here, it is in regards to the number of possible responses they are associated with, in the context of this experiment. Technically, all stimuli are bivalent to the bilinguals, as they can name each picture in both languages. The definitions used here are consistent with Finkbeiner, Almeida, et al. (2006).

between pre-SMA activity and inhibitory control in language switching, the role of this brain region in language control can be further confirmed if a causal relationship is established. Experiment 2 (presented in Chapter 3) explored whether such a causal relationship existed, by externally disrupting the excitability of the pre-SMA and examining the consequence on language switching performance. This disruption was achieved using a non-invasive brain stimulation technique called transcranial magnetic stimulation (TMS).

In this thesis, the involvement of domain-general inhibitory control in language switching was examined with respect to the proposed distinction between whole-language and item-specific inhibition. As explained earlier, the present experimental design affords the ability to inspect both types of inhibition within the same task, thus making it possible to assess whether they share the same underlying neural mechanism. After Experiment 1 established the presence of item-specific inhibition (as indexed by the comparison between univalent and bivalent items) and whole-language inhibition (as indexed by the comparison between stay and switch trials), Experiment 2 then used the same picture-naming task to investigate whether either or both types of inhibition were causally dependent on domain-general inhibitory control. Distinguishing between these two types of inhibition in language switching and examining the role of the pre-SMA in each of them provides more fine-grained information as to what the pre-SMA is responsible for (and what it is not responsible for) in bilingual language control. This might help shed light on the exact role of the pre-SMA and inform future updates to neurocognitive models of bilingual speech production.

1.6 Research questions

To summarise, my research questions in this thesis are as follows:

(a) Does language control in bilingual speech production involve both item-specific inhibition and whole-language inhibition?

(b) If so, do both types of inhibition rely on the same domain-general mechanism?

Chapter 2: Whole-language and item-specific inhibition in language switching (Experiment 1)

This chapter presents a behavioural experiment aimed at identifying whole-language and item-specific inhibition in language switching, as outlined in Chapter 1.

Bilingual participants performed a picture-naming task, in which a small set of pictures were presented repeatedly, and each trial required naming in either English or Mandarin according to a cue. Half of the pictures had consistent language requirement throughout the experiment (univalent items), and the other half had changing language requirements (bivalent items). In this design, whole-language inhibition can be examined when the language requirement changes from one trial to the next (i.e. switch vs. stay trials), and item-specific inhibition can be assessed when the same picture elicits a response in one language after having been named in the other (i.e. bivalent vs. univalent items). Each type of inhibition is indexed by a "cost", reflected in longer naming latencies or higher error rates on the trials affected.

I started with the hypothesis that bilingual language control involves both item-specific inhibition and whole-language inhibition, so the expectations were that switch trials should incur a cost compared to stay trials, and bivalent items should incur a cost compared to univalent items.

2.1 Method

2.1.1 Participants

Sixteen healthy adult Mandarin-English bilinguals (7 males; mean age = 28.2 years) participated for course credit or monetary compensation. Bilinguals were required to be at least moderately proficient in both languages (a minimum self-rating of 4 out of 7, for each language). Participants were free from speech or language impairments, and all had normal or corrected-to-normal vision. One participant was excluded from all analyses due to voice key issues during the experiment (see *2.1.5* for more details). Informed consent was obtained from all participants. The study was approved by the human ethics committee of Macquarie University.

Demographic information and language proficiency self-ratings were collected from all participants using a language history questionnaire (either completed at the end of the experiment, or online in their own time). The multilingual naming test (MINT; Gollan, Weissberger, Runnqvist, Montoya, & Cera, 2012), a 68-item picture-naming test available in both English and Mandarin, was administered to each participant to obtain a more objective measurement of their language proficiency. The naming test was always given after the participant had completed the experimental task, to avoid any possible influence on their performance. Table 2.1 provides a summary of characteristics for the included participants.

Most participants acquired Mandarin at an early age in a home setting (except three who learned the language later at primary school), and they started learning English half way through primary school or from the beginning of high school. In general, the bilinguals were either fairly balanced between English and Mandarin, or slightly more dominant in Mandarin. They switch between languages quite regularly in everyday life. It may be worth noting that the participants were not strictly Mandarin-English bilinguals. Since most Mandarin speakers also speak another variant of Chinese, it is difficult to find such pure Mandarin-English bilinguals. However, it was ensured that participants were only included in the study if Mandarin and English were their two strongest languages.

	Mean	*SD*
Age	28.2	5.8
Age of first exposure to Mandarin	1.9	2.9
Age of first exposure to English	10.0	3.0
Mandarin MINT score [a]	60.9	4.4
English MINT score [a]	53.4	6.3
Mandarin listening ability [b]	6.6	0.9
Mandarin speaking ability [b]	6.4	0.9
Mandarin reading ability [b]	6.7	0.8
Mandarin writing ability [b]	6.5	0.9
English listening ability [b]	5.5	0.9
English speaking ability [b]	5.0	0.8
English reading ability [b]	5.9	0.5
English writing ability [b]	5.0	0.6
Percent Mandarin use currently [c]	47.1	24.9
Percent English use currently [c]	48.9	20.6
Percent Mandarin use during childhood [c]	75.2	35.4
Percent English use during childhood [c]	8.4	9.6
Switching frequency currently [d]	4.2	1.2
Switching frequency in childhood [d]	2.1	1.5

Table 2.1. Characteristics of included participants in Experiment 1.

[a] Maximum possible score in the MINT test is 68 for each language.
[b] Language proficiency based on self-ratings on a 7-point scale: 1 = little to no knowledge, 7 = like a native speaker.
[c] Percentages for Mandarin and English use do not add up to 100 percent, as some participants reported also speaking another variant of Chinese.
[d] Based on a 6-point scale: 1 = never, 2 = very infrequently, 3 = occasionally, 4 = two to three times per conversation, 5 = several times per conversation, 6 = constantly.

2.1.2 Materials

Eight black-and-white line drawings were selected from the stimuli used by Kleinman and Gollan (2016) in their picture-naming study. The pictures were [English/Mandarin(hanyu pinyin[3])]: *hand-shou, door-men, tree-shu, horse-ma, pencil-qianbi, bone-gutou, king-guowang, grapes-putao*. Each picture was to be named in English, Mandarin or both in the experiment. Pictures were selected such that

[3] The romanisation system for spelling out Mandarin sounds.

naming ambiguity (i.e. more than one possible name for a picture) was minimised in both languages, and no within-language or cross-language homophones existed among the sixteen possible target names. All target names in English were either one- or two-syllable words that were 4-6 letters long, and all target names in Mandarin were one- or two-character words (in Mandarin, one character is one syllable). It was ensured that there was minimal semantic relatedness between any two pictures, so that the sequence of pictures could be fully randomised without the risk of any semantic interference effects on naming latencies.

2.1.3 Design and procedure

As explained in Chapter 1, the picture-naming task was designed to allow a direct comparison between univalent and bivalent items (to examine item-specific inhibition), and between stay and switch trials (to examine whole-language inhibition). The task consisted of a training block and a testing block. Each univalent item maintained consistent language requirement throughout the two blocks, while each bivalent item was trained on one language and tested in the other. Item-language pairings were randomly generated for each participant when the experiment started, such that four out of the eight pictures were associated with English and the other four with Mandarin. Next, out of the four pictures associated with each language, two were randomly selected to be univalent and the other two were assigned to be bivalent. In the training block, the original item-language pairings were followed. In the testing block, those pictures that were assigned to be bivalent changed their language requirement (i.e. if it was originally trained in English, it now required naming in Mandarin, and vice versa), while the univalent pictures stayed in their original language (see Figure 2.1). The language requirement on each trial was specified using a language cue, which appeared simultaneously with the picture stimulus. The language cue was either "*What is this?*", indicating the response was to be given in English, or the Chinese equivalent "*这是什么?*", indicating a response in Mandarin was required. These language cues were designed to elicit responses in each language more naturally (compared to the commonly used cues, such as background colours or national flags), so as to minimise any cue-processing and related costs.

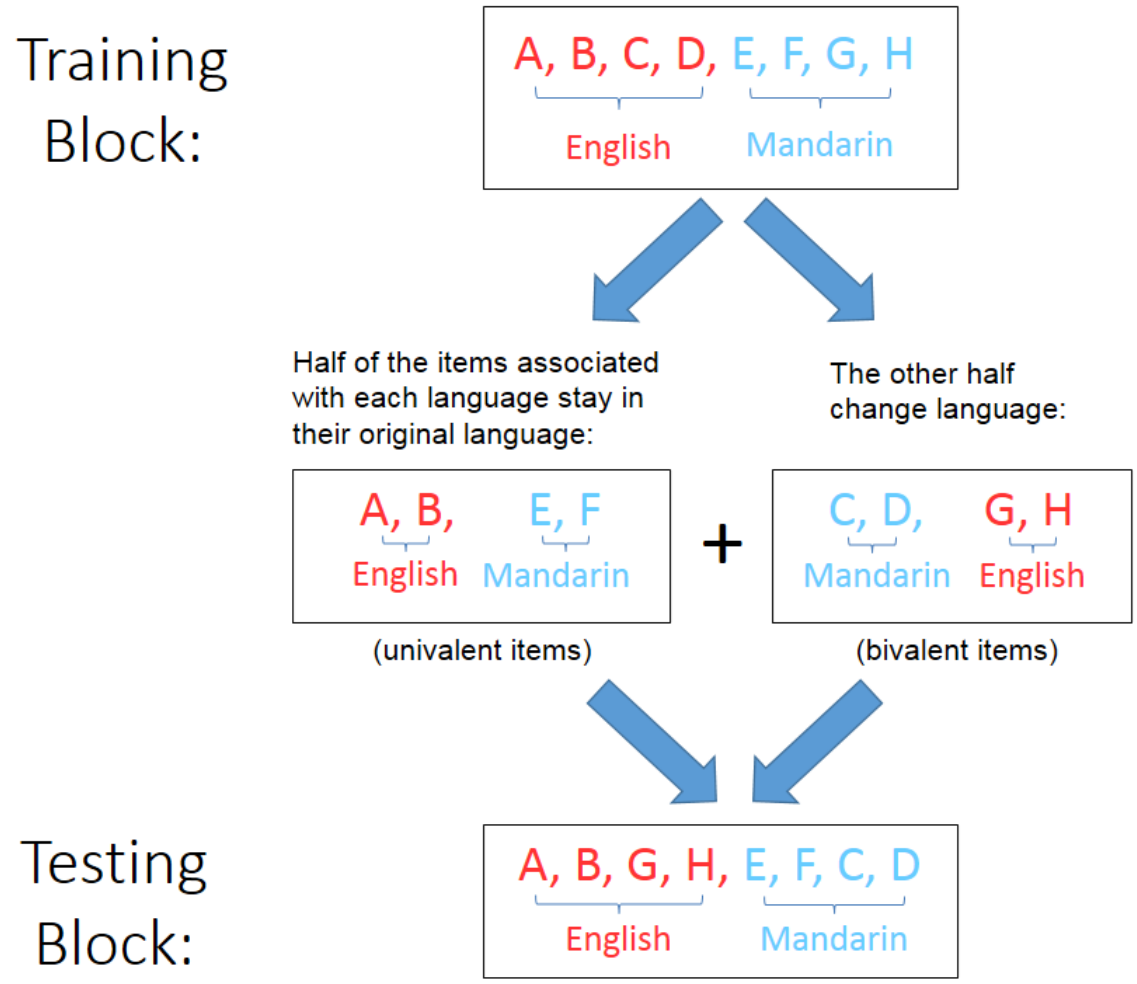

Figure 2.1. Illustration of the procedure used in Experiment 1 to achieve balanced assignments of language and valence to the picture stimuli. A total of eight pictures were used in this experiment. Here each letter (e.g. 'A') represents one picture item. Items associated with English are shown in red; items associated with Mandarin are shown in blue. Univalent items maintained consistent language requirement in the two blocks, while bivalent items were trained and tested in opposite languages. Item-language pairings for the training block were randomly generated for each participant, such that four out of the eight pictures were associated with English and the other four with Mandarin (*top row*). Next, out of the four pictures associated with each language, two were randomly selected to be univalent and the other two were assigned to be bivalent. The language requirement for each bivalent item was changed (*middle row*). This produced the set of item-language associations to be used in the testing block (*bottom row*).

Participants were tested individually in a soundproof room. Each session lasted 35-45 minutes. The experiment was programmed in, and controlled by the Presentation software (Neurobehavioral Systems, Version 18.3). Stimuli were displayed on a Samsung SyncMaster SA950 (27 inch) monitor, connected to a Dell Optiplex 9010 PC (3.2GHz Intel i5-3470 CPU, 8GB RAM). Participants were seated comfortably in a chair 80cm away from the monitor. Vocal responses were recorded through a microphone, and a voice key was set up in Presentation to detect response onset. The microphone amplifier volume was adjusted individually for each participant to optimise the functioning of the voice key. Before the picture-naming task commenced, participants were given verbal and onscreen instructions, which asked them to name the pictures as quickly and accurately as possible according to the language cue on each trial. Instructions were followed by a short practice block, which consisted of the same stimuli used in the experiment proper. Each stimulus appeared twice in the practice block. The purpose was to allow participants to familiarise themselves with the task as well as to make sure they had no trouble naming each picture. After a short break, participants initiated the training block themselves by pressing a key when they were ready. A short break was given after the training block was completed, and then participants initiated the testing block, again by pressing a key themselves.

In the training block, each picture stimulus appeared 12 times. Pictures were presented in a random order for each participant, with the constraints that each picture appeared an equal number of times on stay trials and switch trials, and that no two consecutive trials had the same picture. In the testing block, trials were presented in the form of triads (i.e. groups of three), similar to the quartet structure used by Finkbeiner, Almeida, et al. (2006). In the triad structure, each (critical) trial was preceded by two filler trials. These fillers served a setup purpose (i.e. they were not included in data analysis), but appeared no different to critical trials from the participants' perspective. The two filler trials in a triad always required responses in the same language, to ensure that each critical trial had a run-length of two (i.e. a switch trial would not directly follow another switch trial, which could result in a "stacked" effect). Thus, an example of a stay trial could be English -> English -> English, and a switch trial could be Mandarin -> Mandarin -> English. Each target picture stimulus appeared 12 times on critical trials (six stay trials and six switch

trials), resulting in a total of 96 critical trials. As language and valence were already assigned earlier to all picture items in a random and balanced manner, this created critical trials that were fully balanced across language, valence, and trial type (stay vs. switch), eliminating possible bias due to factors other than the variables of interest. In addition, the same eight picture stimuli were used on the filler trials, so each picture appeared 24 times as a filler. The triads were constructed in such a way that there was no repetition of pictures within each triad, and then all the triads were presented to the participant in a random sequence. The use of filler trials allowed dynamic sequences to be generated for each participant on the fly, and further ensured participants would not be able to make predictions about the upcoming trial (as fillers were indistinguishable from critical trials). To avoid the naming of bivalent items on filler trials potentially overriding the training effect (and to maintain the bivalency throughout the testing block), filler trials used the original item-language pairings consistent with the training block. Thus, opposite languages were required on filler and critical trials for bivalent items[4].

The trial structure is shown in Figure 2.2. Each trial started with a fixation cross which appeared at the centre of the screen for 350 ms. This was followed by a blank screen for 150 ms, before the language cue and picture stimulus appeared simultaneously on screen. The picture was displayed at the centre of the screen, while the language cue was located above it. Sound recording started as soon as the stimulus appeared. The trial was terminated upon the voice key being triggered by a response, or 3 seconds after stimulus onset if no response was detected. The inter-trial interval lasted 850 ms, during which a blank screen was displayed, and then the next trial started. The vocal response on each trial was saved as an individual wave file for later verification.

[4] Note that the filler trials alone may be sufficient to create the bivalency, which means the training block may be redundant.

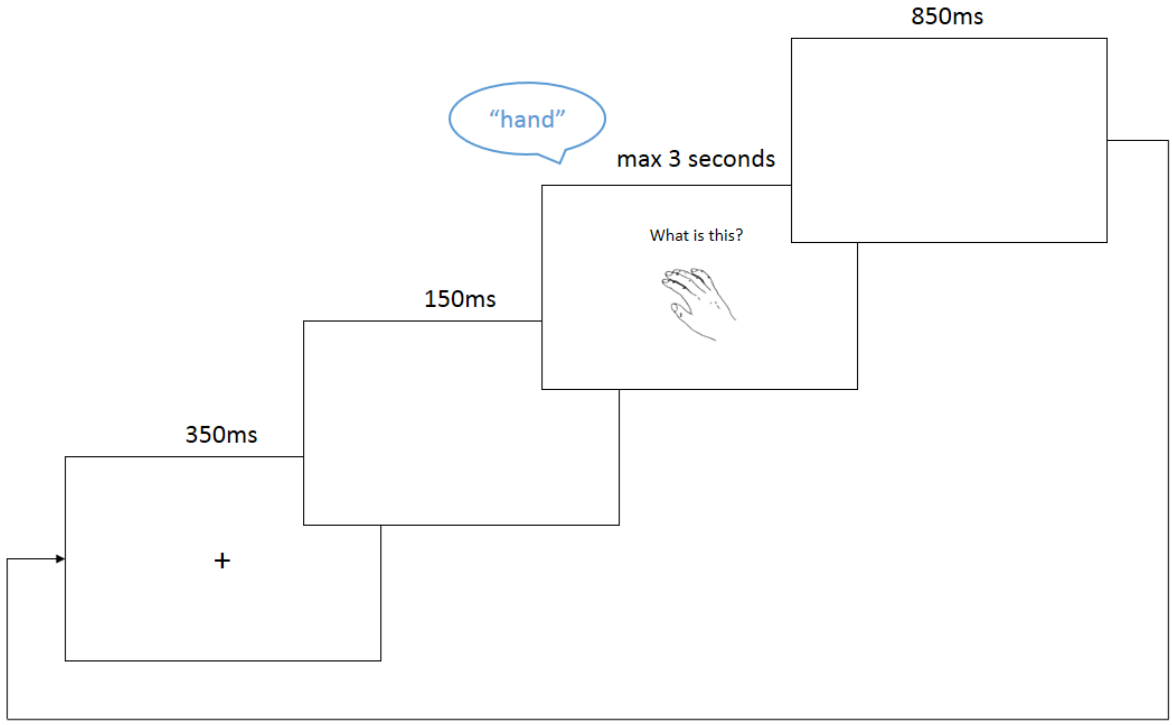

Figure 2.2 An example of a naming trial, showing the sequence of frames and the display of language cue ("What is this?") and target stimulus (the picture of the hand). This example trial requires the response "hand" in English.

2.1.4 Post-processing and trial exclusions

The voice key in Presentation is triggered when the input speech volume from the microphone reaches a certain threshold. This was intended to serve two purposes in the experiment: ending the current trial when a response is detected, and automatically reporting a reaction time (RT) value for each trial. While the speech detection was good enough for ending trials, the RT output (in milliseconds) did not reach the expected level of accuracy (i.e. the detected RTs did not consistently align with response onset across all trials). In order to obtain more accurate RT values, all of the wave files were processed offline using in-house software for speech onset detection. The detection output for each wave file was visualised as a graph and visually inspected to ensure accuracy. Any inaccurately detected RTs were identified and those trials were subsequently excluded from the RT analysis.

Error coding was performed manually for all trials by checking the sound recording against the target response. The definition of “error” used here was a broad one, which included incorrect responses as well as all verbal disfluencies (e.g. partial responses, stuttering, and utterance repairs). If the participant started giving the correct response but hesitated before having sounded out the complete word, or if they started to make a mistake but quickly corrected themselves, these were all counted as error trials, or perhaps more accurately named “high conflict trials”. In other words, only straightforward correct responses were scored as correct. The reasoning is that those disfluencies represent cases of high conflict (which we are interested in for the same reason that we are interested in error trials), and determining which of these trials should be classified as correct and which as error often must involve subjective interpretation of the response given.

One participant was excluded from all data analyses due to heavy breathing triggering the voice key on a large number of trials. Even though this did not affect the RT values (as speech onsets were correctly detected by the post-processing procedure described above), the early triggering of voice key meant that the trial ended (and stimulus disappeared) before an actual response was produced. This could affect the RT for the current trial in unknown ways. Moreover, the early ending of trials resulted in shortened inter-trial interval (which started as soon as each trial ended), and it appeared that there was insufficient time following these trials for the participant to get ready for the upcoming trial.

2.1.5 Data analysis

Two within-subjects variables were examined in the statistical analysis: valence (univalent items vs. bivalent items) and trial type (stay trials vs. switch trials). Each variable had two levels, making a total of four conditions. Mean reaction times (RT) and error rates (ER) in picture naming were calculated for each participant in each condition. Following previous convention in language switching studies, RTs less than 250ms or greater than 3000ms were treated as outliers (Kleinman & Gollan, 2016; Gollan et al., 2014). Outlier trials, as well as trials where non-response noise (e.g. coughing) was present before the actual response, were counted as bad trials and excluded from both RT and ER analysis. The error rate is, therefore, the number

of error trials out of the total number of good trials. All error trials were excluded from the RT analysis, as were the trials identified earlier with inaccurately detected RTs.

Response language (L1 vs. L2) was not included as a factor in data analysis. This was based on the consideration that the training block in the experimental design may have created a temporary "dominant language" for each picture item by training it in that language[5]. In this case, the participants' natural language dominance may not be meaningful, and it may provide misleading information instead (due to a confound with valence)[6].

A two-way repeated measures analysis of variance (ANOVA) was conducted in SPSS (IBM SPSS Statistics, Version 22). The factors in the analysis were valence and trial type. In addition, follow-up t-tests were performed separately on univalent items and bivalent items to determine whether the effect of trial type was significant for both. All effects were categorised as significant at $p < .05$ and marginally significant at $p < .10$.

[5] Bilinguals often have one dominant language in certain domains while the other language is dominant in other subject areas, depending on which words are used more frequently. This means that "language dominance" is not tied to a particular language, but rather based on the usage frequency of individual lexical items.

[6] In future experiments, I plan to remove the training block and see if bivalency can be achieved using the filler trials in the testing block alone. If this new design can still successfully differentiate between univalent and bivalent items, it would resolve the confound between valance and language dominance, and allow both factors to be directly examined in data analysis.

2.2 Results

Following the trial exclusion procedure described above, approximately 1.0% of trials were classified as bad trials and excluded from both the ER and RT analyses. An additional 8.5% were error trials and these were also excluded from the RT analysis. Mean reaction times in each of the four conditions are shown in Figure 2.3.

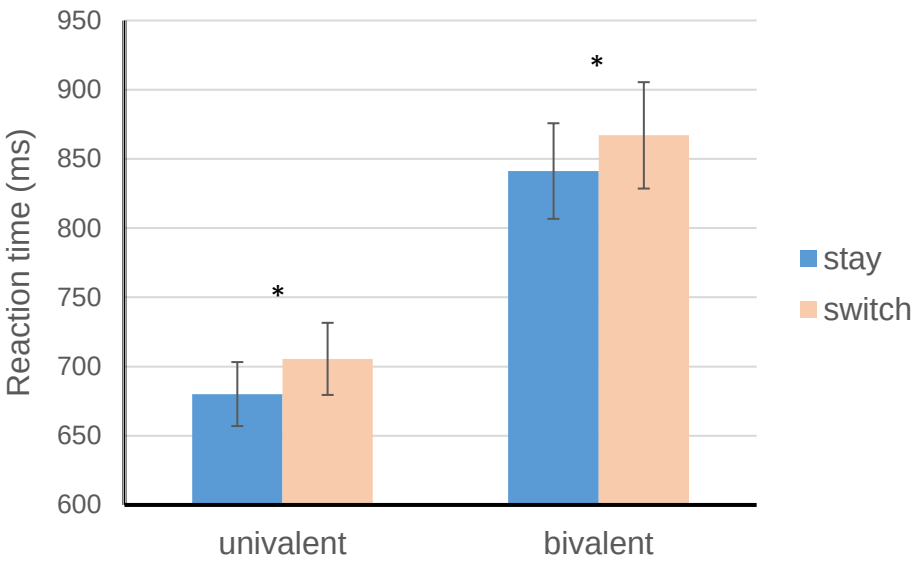

Figure 2.3. Reaction times (in milliseconds) as a function of valence (univalent vs. bivalent items) and trial type (stay vs. switch trials). Error bars indicate one standard error above and below the mean values. Planned follow-up t-tests on the effect of trial type were performed separately for univalent and bivalent items. *: $p < .05$.

The RT analysis revealed main effects of both valence and trial type on picture naming latencies (see Figure 2.4). Bivalent items were named more slowly (M ± SEM = 854 ± 36ms) than univalent items (M = 693 ± 24ms): F (1, 14) = 39.152, p < .001, η_p^2 = .737. Naming on switch trials took longer (M = 786 ± 30ms) than on stay trials (M = 761 ± 26ms): F (1, 14) = 6.468, p = .023, η_p^2 = .316. There was no significant interaction between valence and trial type. Planned follow-up comparisons (see Figure 2.1) revealed numerically similar switch costs for univalent items (β = 25ms) and bivalent items (β = 26ms), and both were statistically significant (univalent: t = 2.273, p = .039; bivalent: t = 2.282, p = .039).

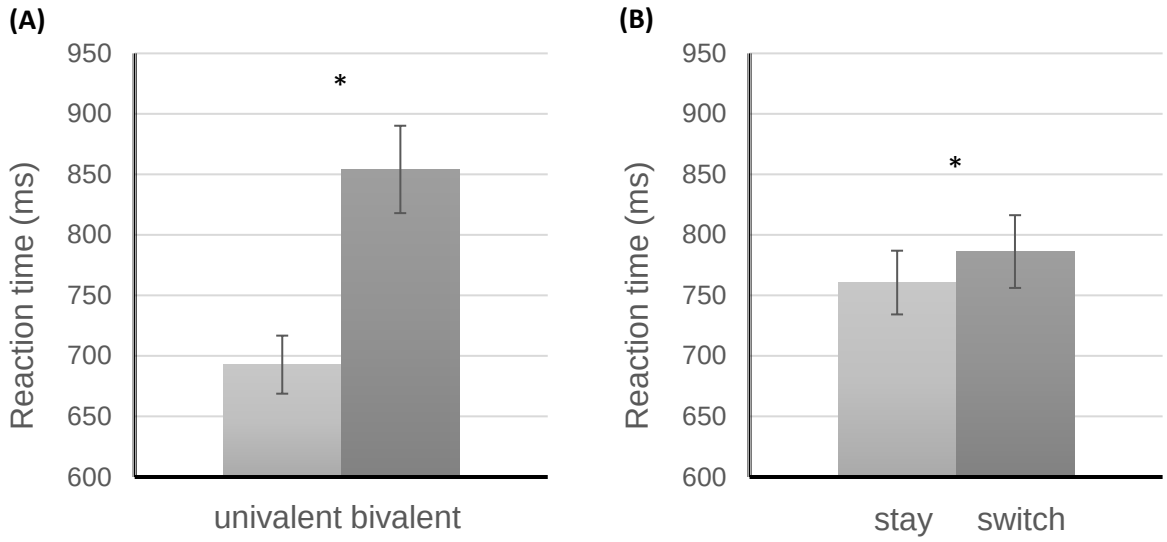

Figure 2.4. Main effects of valence and trial type on picture naming latencies. Both were statistically significant, but the effect size of valence was much larger.
(A) Mean reaction times for univalent and bivalent items, collapsed across trial types.
(B) Mean reaction times on stay and switch trials, collapsed across valence.
Error bars indicate one standard error above and below the mean values. *: p < .05.

The error rate analysis showed a main effect of valence, with more errors occurring on bivalent items (M = 15.4 ± 1.7%) than on univalent items (M = 1.9 ± 0.7%): F (1, 14) = 50.895, p < .001, η_p^2 = .784. No other effects or interactions were significant in the ER analysis.

Effect tested	Measure	df	*F*	*p*	Partial Eta Squared
Valence	RT	1, 14	39.152	**<.001**	.737
	ER	1, 14	50.895	**<.001**	.784
Trial type	RT	1, 14	6.468	**.023**	.316
	ER	1, 14	2.961	.107	.175
Valence * trial type	RT	1, 14	.001	.972	.000
	ER	1, 14	.441	.518	.031

Table 2.2. Statistical analysis results for reaction times (RT) and error rates (ER) in Experiment 1. Factors in the analysis were valence (univalent items vs. bivalent items) and trial type (stay trials vs. switch trials). Statistically significant effects (p < .05) are shown in **bold**.

2.3 Discussion

2.3.1 Whole-language and item-specific inhibition

The results of this experiment reveal two important effects. Firstly, there is a large difference between univalent and bivalent items in terms of both naming latencies and error rates (i.e. main effect of valence). Secondly, there is a difference in naming latencies between stay and switch trials (i.e. main effect of trial type). These findings will be interpreted below in regards to the distinction between whole-language and item-specific inhibition, advanced by De Groot and Christoffels (2006).

According to De Groot and Christoffels (see also De Groot, 2011, Chapter 6), whole-language inhibition in bilingual speech production occurs during the early stage of lexical selection and entails inhibition of the non-target language on a global level. For example, if an English-Mandarin bilingual is required to name a picture in English on one trial, the correct selection of the English word involves globally increasing the activation levels of all lexical nodes in English and globally inhibiting all lexical nodes in Mandarin. Now if Mandarin is required on the next trial, the previous inhibition of Mandarin must be overcome to enable production in this language. In other words, whenever a language change is required (i.e. switch trials), the current language in use must be suppressed and the other language, which was suppressed on the previous trial, must be reactivated. This brings about a cost in RT, presumably because it takes time to overcome prior inhibition: the amount of time required to reactivate the currently required language is a function of how strongly it was suppressed previously (Green, 1998). In contrast, stay trials have the same language requirement as their preceding trial, and therefore do not require such resolution of inhibition. It follows that stay trials should not incur such an RT cost. Thus, whole-language inhibition predicts longer reaction times on switch trials than on stay trials. This prediction is confirmed by the main effect of trial type in this experiment. The RT difference between the two trial types (switch RT - stay RT) is usually referred to as the "switch cost". Since this cost indexes the amount of time to recover from whole-language inhibition[7], I will call it the "whole-language inhibition cost" here, for ease of comparison with the "item-specific inhibition cost", which is introduced next.

Following this global balancing of activation levels of the two languages, item-specific inhibition operates at a later stage in lexical selection and involves suppression of specific lexical nodes which are in competition with the target node (De Groot & Christoffels, 2006). This type of control is reactive and ensures correct selection of the desired output by catching any highly activated lexical nodes in the non-target language (despite early global adjustments) and suppressing them. According to Green (1998), the strength of this suppression is proportional to the

[7] Note that overcoming prior inhibition is just one possible interpretation of this cost. Following Meuter and Allport (1999) and Green (1998), this is the predominant interpretation given in the language switching literature. However, there are other processes which may contribute to the "switch cost"; those possibilities will be discussed in more detail in Chapter 4.

activation level of the individual non-target lexical nodes, such that more activated lexical nodes receive stronger inhibition. Since each bivalent stimulus in this experiment is associated with two responses, fierce competition can be expected each time a bivalent picture appears. Therefore, item-specific inhibition is needed to actively suppress the response from occurring in the wrong language (according to the language requirement on the current trial) and to ensure selection of the correct response. This item-specific inhibition should be particularly strong here as the two competing responses for each bivalent picture are translation-equivalents of each other, and both should be highly activated by the picture stimulus. According to the same assumption outlined above for whole-language inhibition, once a particular response is suppressed, it takes time to overcome this suppression when the same response is required for output on a later trial. As the item-specific inhibition here is supposedly very strong, resolving the inhibition can be expected to incur a large cost in RT. On the other hand, each univalent picture is always associated with the same response and does not require such item-specific inhibition. Thus, the naming of bivalent items should be much slower than that of univalent items. The significant main effect of valence in this experiment confirms this prediction. Since the RT cost (bivalent RT - univalent RT) reflects the amount of time needed to recover from item-specific inhibition, I will call it the "item-specific inhibition cost".

Both the "whole-language inhibition cost" and the "item-specific inhibition cost" showed up in the results (as main effect of trial type and main effect of valence, respectively), supporting the hypothesis that both types of inhibition are at work during language switching. The whole-language inhibition cost obtained here was smaller than the switch cost found in most language switching studies (for a summary, see Bobb & Wodniecka, 2013). This suggests that the language cues used in this experiment, which were designed to elicit responses naturally and minimise cue-processing costs, may have successfully achieved their purpose. Thus, we may consider this whole-language inhibition cost as truly reflecting the cost of switching language[8]. As expected, the effect size of valence is notably larger than the effect of trial type, suggesting that item-specific inhibition is of much greater strength and takes longer to overcome. Numerically, it takes 161 ms to resolve item-

[8] Of course, I cannot say so conclusively until a direct comparison is carried out between the different types of language cues.

specific inhibition applied previously, while whole-language inhibition only takes 25 ms to overcome. Since the present study is the first to examine item-specific and whole-language inhibition side-by-side in the language switching paradigm, comparison of results can only be made with similar studies conducted using other paradigms. Van Assche et al. (2013) investigated item-specific and whole-language control in Mandarin-English bilinguals using a verbal fluency task. They found that both types of inhibition exist, but item-specific inhibition is much stronger. These findings are perfectly consistent with the present study.

To summarise the results of Experiment 1, evidence was found for both item-specific and whole-language inhibition in language switching. However, it is unclear whether these two types of inhibition are of the same nature (i.e. both suppressing non-target nodes at the level of lexical representations), with simply a strength difference between them. Given the proposed different timing at which whole-language and item-specific inhibition operate and the distinct purpose they serve, as well as the vast difference in magnitude between the costs associated with them in this experiment (note also that while the RT cost was significant for both, the increase in error rate was significant for item-specific inhibition only), it is not unreasonable to suspect that the underlying neural mechanism for these two types of inhibition may be dissociable. This will be investigated further in Experiment 2, where TMS will be used to examine whether both types of inhibition recruit the same domain-general inhibitory mechanism.

2.3.2 Whole-language inhibition in univalent items

An additional objective of this experiment, in the case that whole-language inhibition was found to exist, was to examine whether both univalent and bivalent items are similarly affected by this type of inhibition. While significant switch costs have been consistently reported for bivalent stimuli in language switching studies (e.g. Costa & Santesteban, 2004; Gollan et al., 2014, Exp. 2), elimination of switch cost has been observed when univalent stimuli were used (Finkbeiner, Almeida, et al., 2006; Kleinman & Gollan, 2016). This is a very interesting finding; however, neither of these studies directly compares univalent and bivalent stimuli which differ in nothing but valence, to demonstrate a difference in switch cost purely related to stimulus

valence. Finkbeiner, Almeida, et al. (2006) included a confound of task switching on all their univalent stimuli, such that any effect of valence could also be attributed to task switching (see previous discussions in *Section 1.5*). Kleinman and Gollan (2016) examined univalent stimuli in a separate block (i.e. not mixed together with bivalent stimuli), making it difficult to draw any comparison; in addition, this study adopted the "bottom-up design" (see *Section 1.1*), which leaves the freedom of language choice to the participants, so their findings may or may not be transferrable to cued language switching. The present experiment provides an ideal opportunity to revalidate the findings discussed above, as its design allows a direct comparison of univalent and bivalent stimuli without confounds.

Finding out whether the naming of both univalent and bivalent items were affected similarly by whole-language inhibition is important for the purpose of providing find-grained information to guide the next experiment. If we want to examine how whole-language inhibition is modulated by the disruption of a brain area, we should know when exactly this type of inhibition is applicable in the first place (e.g. on all trials or only certain trials). Kleinman and Gollan (2016) eliminated switch cost in their bottom-up block (where all items are univalent and each named in whichever language the participant preferred for this item). These authors argued that inhibitory control may be suspended when processing univalent items, as lexical selection can be driven purely by accessibility: a more efficient mechanism. Thus, I wanted to see whether the elimination of switch cost on univalent items would be reproduced under forced language selection (with the same training procedure to create temporary accessibility preference towards the cued language). If a similar pattern was obtained (i.e. bivalent items being associated with switch costs while univalent items are not), this would provide more robust support for the claim that the control mechanism behind whole-language inhibition is flexible and can be engaged when required and disengaged at other times to allow more efficient processing.

However, the results from this experiment showed otherwise. There was no interaction between trial type and valence in the ANOVA, suggesting a lack of significant difference in switch costs between univalent and bivalent stimuli. In other words, there was no evidence for a flexible whole-language control mechanism which was selectively engaged in high-conflict situations (i.e. bivalent items) only. The numerical values of the univalent and bivalent switch costs only differed by 1ms.

Planned follow-up t-tests, which were performed separately for univalent and bivalent items, confirmed that both had statistically reliable switch costs. These results provide strong evidence that both univalent and bivalent items in this experiment were consistently affected by whole-language inhibition.

This finding is important for two reasons. Firstly, it provides guidance for the next experiment in regards to what kind of comparison is appropriate in assessing the (differential) effect of TMS on whole-language inhibition and item-specific inhibition. Hypothetically, if the whole-language inhibition cost was found to occur on bivalent items only, it would not be appropriate to expect its cost to be modulated to the same degree as the item-specific inhibition cost even if both types of inhibition were equally affected by TMS; in this case, univalent items (which were not subject to whole-language inhibition in the first place) should be left out of the comparison. Now the actual finding is that the whole-language inhibition cost was consistently found across univalent and bivalent items in this experiment; therefore, it is justified to perform a direct comparison of TMS effect on the whole-language and item-specific inhibition costs, regardless of valence. Secondly, the finding above confirms that the vast difference in magnitude between item-specific and whole-language inhibition costs (as discussed earlier in *2.3.1*) was real rather than apparent (i.e. it was not due to univalent items having a reduced switch cost and masking an otherwise large switch cost on bivalent items). This strengthens the supposition that the two types of inhibition may involve different underlying mechanisms.

Chapter 3: Domain-general inhibitory control in language switching (Experiment 2)

In the behavioural experiment presented in Chapter 2, two types of inhibition cost were identified. The difference between univalent and bivalent items constituted the "item-specific inhibition cost", while the comparison of stay trials to switch trials yielded the "whole-language inhibition cost". Both costs were found in the language-switching task, indicating the presence of both types of inhibition. However, given the proposal that whole-language and item-specific inhibition operate at different timing and serve distinct purposes in lexical selection (De Groot & Christoffels, 2006), as well as the vast difference in magnitude between the two costs found in the behavioural experiment, it is reasonable to consider that these two types of control may operate via different inhibitory mechanisms. Specifically, the question of interest here is whether one or both of them engages executive control.

The current experiment examines a particular brain area responsible for domain-general inhibitory control: the pre-SMA. This brain area has been found to be active during language switching, suggesting its possible involvement in carrying out inhibition in language control. However, it remains unclear whether the pre-SMA has a causal role in language inhibition, and if so, what its precise function is. These questions are investigated in this experiment, by perturbing the area using a repetitive TMS protocol and observing the effect on each of the two types of language inhibition identified earlier. If TMS modulates one type of inhibition cost but not the other, we can then pinpoint exactly when (i.e. under what situations) the pre-SMA is recruited and infer more precisely what role it plays in language control.

I hypothesised that the pre-SMA serves to facilitate conflict resolution in bilingual language control, and therefore it should have a more prominent role in item-specific inhibition, where a high level of conflict is present between two responses.

3.1 Method

3.1.1 Participants

Ten healthy right-handed adult Mandarin-English bilinguals (5 males; mean age = 26.8 years) participated for monetary compensation. Four of them were returning participants who had previously been tested in Experiment 1. All bilinguals had native-like proficiency in Mandarin (which they learned at home) and were moderately proficient in English (which they learned at school). All participants were free from neurological disorders and met the safety requirements for undergoing MRI and TMS. None of the participants were taking any psychiatric medication, and all had normal or corrected-to-normal vision. All participants gave informed consent before taking part in the experiment. The study was approved by the human ethics committee of Macquarie University.

Individual high-resolution T1-weighted brain MRI images were obtained for each participant for the purpose of localising the target area for TMS. Each participant was then tested in two separate TMS sessions at least one week apart (for eight participants, the interval was exactly one week; for the remaining two, the interval was two weeks). The TMS sessions were all scheduled in the afternoon, and the two sessions for the same participant always took place at the same time of day (with at most one-hour difference between the starting time) to minimise possible circadian effects (Sale, Ridding, & Nordstrom, 2008). Testing order was fully counterbalanced in regards to TMS order (pre-SMA stimulation in first session, or control site in first session), including counterbalancing within the two participants who had two-week interval between the first and second session. One participant was excluded due to technical issues with the MRI scans obtained. A new participant was recruited as replacement in order to maintain the counterbalancing of TMS order.

Demographic and language proficiency information was collected from the participants as in Experiment 1, and a summary of this information (for the included participants) is presented in Table 3.1.

	Mean	*SD*
Age	26.8	3.6
Age of first exposure to Mandarin	1.0	2.3
Age of first exposure to English	11.0	1.4
Mandarin MINT score [a]	63.1	2.1
English MINT score [a]	42.8	5.8
Mandarin listening ability [b]	7.0	0.0
Mandarin speaking ability [b]	7.0	0.0
Mandarin reading ability [b]	7.0	0.0
Mandarin writing ability [b]	7.0	0.0
English listening ability [b]	4.9	0.6
English speaking ability [b]	4.3	0.7
English reading ability [b]	5.1	0.9
English writing ability [b]	4.5	0.8
Percent Mandarin use currently [c]	73.9	10.5
Percent English use currently [c]	25.1	10.5
Percent Mandarin use during childhood [c]	90.1	21.3
Percent English use during childhood [c]	2.9	3.3
Switching frequency currently [d]	3.7	1.3
Switching frequency in childhood [d]	1.5	0.5

Table 3.1. Characteristics of included participants in Experiment 2.

[a] Maximum possible score in the MINT test is 68 for each language.
[b] Language proficiency based on self-ratings on a 7-point scale: 1 = little to no knowledge, 7 = like a native speaker.
[c] Percentages for Mandarin and English use do not add up to 100 percent, as some participants reported also speaking another variant of Chinese.
[d] Based on a 6-point scale: 1 = never, 2 = very infrequently, 3 = occasionally, 4 = two to three times per conversation, 5 = several times per conversation, 6 = constantly.

3.1.2 Target localisation

The pre-SMA is a small cortical region located on the medial frontal cortex (very close to the midline between the two hemispheres of the brain). The fMRI studies that identified activation of this area in language switching simply referred to it as “pre-SMA” (without stating whether it is the left or right side) or “midline pre-SMA” (Luk et al., 2012). However, a precise target location is required for TMS, as stimulating on the midline (i.e. on top of the medial longitudinal fissure) would likely

result in ineffective stimulation of either the left or right pre-SMA. The right pre-SMA was chosen in this study because it is more commonly accepted as part of the inhibitory control network (Cai, George, Verbruggen, Chambers, & Aron, 2012).

A high-resolution T1-weighted structural brain MRI scan (slice thickness: 1 x 1 x 1 mm) was obtained for each participant (Macquarie Medical Imaging, Macquarie University Hospital, Sydney). The images were firstly reoriented as necessary such that the head was upright and the anterior commissure (AC) and posterior commissure (PC) were on the same horizontal line. The pre-SMA was then located anatomically, adapting the procedure described by Tremblay and Gracco (2009) to locate the right rather than the left pre-SMA, viz. a vertical line was drawn 10mm anterior to the AC, forming a coronal plane which intersects the cerebral cortex at the top. The right pre-SMA was identified as a point along the intersection on the medial most portion of the right superior frontal gyrus (SFG). The coordinates of this point were noted and a white spherical blob was drawn onto the MRI at this position using an in-house Matlab script.

Localisation of TMS target on the participant was guided by a frameless stereotaxic system (Visor2, ANT Neuro, Enschede, Netherlands; http://www.ant-neuro.com). The MRI images for each individual participant were loaded into the navigation system and a 3D model of the head and brain was reconstructed from these images. The target location was then marked in the system at the location of the white blob drawn earlier. During each TMS session, a MRI coregistration procedure was performed to link the 3D model to the participant's head in real space. The participant wore a headband with reflective spherical markers attached, the positions of which were tracked by the navigation system, which then guided the placement of the coil over the predefined target location.

The vertex, which served as the control site, was defined as the halfway point between the nasion and inion (Cai et al., 2012). This location was determined with tape measurement and the desired coil position was marked for later use. For both stimulation sites, the coil was held with the handle pointing to a posterior direction. The same MRI coregistration procedure and tape measurement were carried out during both the experimental session and control session to make the two sessions appear identical from the participant's perspective, and participants were told that

two areas of interest were being investigated. During the debriefing at the end of the entire experiment, participants reported similar sensations from TMS during both sessions, and some expressed surprise upon learning that one of these sessions was the control. When asked to guess which session was experimental and which was control, they were unable to tell (more than half gave the incorrect answer).

3.1.3 TMS Procedure

Magnetic stimulation was delivered using a Magstim Rapid2 stimulator (Magstim Co., Whitland, UK), with a hand-held 70-mm figure-of-eight coil. Resting motor threshold (RMT) was determined individually for each participant. The RMT was defined as the minimum intensity applied on the right primary motor cortex (M1) to elicit three visible twitches on the contralateral first dorsal interosseous (FDI) muscle out of five consecutive stimuli. Participants were instructed to keep their hand muscles relaxed while the RMT was determined.

Continuous theta-burst stimulation (cTBS; Huang, Edwards, Rounis, Bhatia, & Rothwell, 2005) was used to achieve transient suppression of the right pre-SMA. This is a repetitive TMS protocol capable of inducing a reduction of cortical excitability thought to be mediated by long-term-depression-like mechanisms (Huang, Chen, Rothwell, & Wen, 2007). Suppressive effects of cTBS on pre-SMA excitability has previously been demonstrated (e.g. Dietrich, Hertrich, Ackermann, Ziemann, & Müller-Dahlhaus, 2015). In the cTBS protocol, each burst consisted of 3 pulses delivered at 50Hz, and the bursts were repeated at 5Hz. As such, a total of 600 pulses were delivered over a period of 40 seconds.

In accordance with previous studies (Chiou, Sowman, Etchell, & Rich, 2014), stimulation intensity for each individual was calculated as 80% of their RMT. The average RMT for the participants in this experiment was 69% (range 60~76%), meaning that the intensities to apply should be ranging from 48~61%. However, due to the capacity limit on the stimulator available, the maximum output intensity achievable in the cTBS protocol was 51%. Therefore, the highest stimulation intensity given to any participant was capped at 51%.

3.1.4 Behavioural task

The behavioural task was the same picture-naming task used in Experiment 1, with identical materials and procedure. After the RMT was determined and MRI coregistration was performed, the participant was given verbal and onscreen instructions for the task and completed the first part of picture naming (i.e. the training block). The coregistration accuracy was checked (by validating the nasion position) immediately before TMS to ensure the navigation markers worn on the participant's head did not move relative to the head (in one case where the validation failed, the coregistration procedure was carried out again before TMS). Then, cTBS was delivered for 40 seconds and the participant was instructed to rest for 5 minutes without talking. This waiting time was based on observations on the after-effects of cTBS over M1, where the modulation of motor evoked potentials (MEP) was found to be most reliable at 5 minutes post-stimulation (Vernet et al., 2014). After the 5-minute waiting time, the participant was instructed to proceed to the second part of picture naming (i.e. the testing block).

To make the results from the two TMS sessions more comparable and to ensure there was no contradicting training effects during the two sessions, the same item-language pairings and item-valence assignment were always maintained for each individual. During the first TMS session, the pairings were randomly generated just as in Experiment 1. However, for the second session, the previously generated pairings were used instead of new pairings being created. For those who had already participated in the behavioural experiment before, the item-language pairings generated for each of these individuals during that experiment were copied over for use in this experiment. Additionally, for these participants, there was an interval of minimum three weeks between the behavioural session and the first TMS session to allow potential practice effects to dissipate.

3.1.5 Data analysis

Offline RT detection, trial exclusions, and error coding were carried out as in Experiment 1. A new variable was introduced in this experiment, as each participant received TMS stimulation to two different locations: pre-SMA and control site. Thus, three within-subjects variables were examined in the statistical analysis: valence

(univalent items vs. bivalent items), trial type (stay trials vs. switch trials), and TMS condition (control vs. pre-SMA). Each variable had two levels, making a total of eight conditions. Mean reaction times (RT) and error rates (ER) in picture naming were calculated for each participant in each condition. A 2x2x2 repeated measures ANOVA was conducted with the three factors above. All effects were categorised as significant at $p < .05$ and marginally significant at $p < .10$.

3.2 Results

Following the trial exclusion procedure described in Experiment 1, approximately 1.0% of trials were classified as bad trials and excluded from both the ER and RT analyses. An additional 13.4% were error trials and these were also excluded from the RT analysis. Mean reaction times in each of the eight conditions are shown in Figure 3.1.

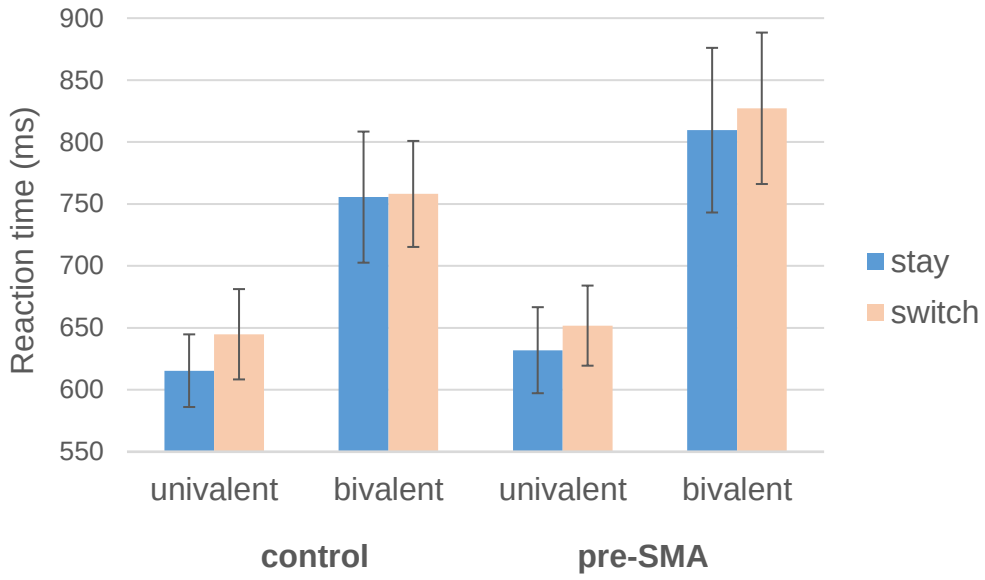

Figure 3.1. Reaction times (in milliseconds) as a function of TMS condition (control site vs. pre-SMA), valence (univalent vs. bivalent items), and trial type (stay vs. switch trials). Error bars indicate one standard error above and below the mean values.

The RT analysis revealed a main effect of valence on picture naming latencies, with bivalent items taking much longer to name (*M* ± *SEM* = 788 ± 54ms) compared to univalent items (*M* = 636 ± 31ms): *F* (1, 9) = 16.988, *p* = .003, η_p^2 = .654. Trial type had a marginally significant effect, with longer RTs on switch trials (*M* = 720 ± 37ms) than on stay trials (*M* = 703 ± 42ms): *F* (1, 9) = 3.982, *p* = .077, η_p^2 = .307. Crucially, there was a marginally significant interaction between TMS condition and valence: *F* (1, 9) = 4.532, *p* = .062, η_p^2 = .335. Post hoc analyses performed separately on univalent and bivalent items revealed that TMS prolonged RTs on bivalent items (pre-SMA: *M* = 818 ± 63ms; control site: *M* = 757 ± 47ms): *F* (1, 9) = 5.025, *p* = .052, η_p^2 = .358, while having no significant effect on univalent items (pre-SMA: *M* = 642 ± 33ms; control site: *M* = 630 ± 33ms): *F* (1, 9) = 0.216, *p* = .653, η_p^2 = .023. Mean reaction times in the post hoc analyses are shown in Figure 3.2. No other effects or interactions were significant in the RT analysis.

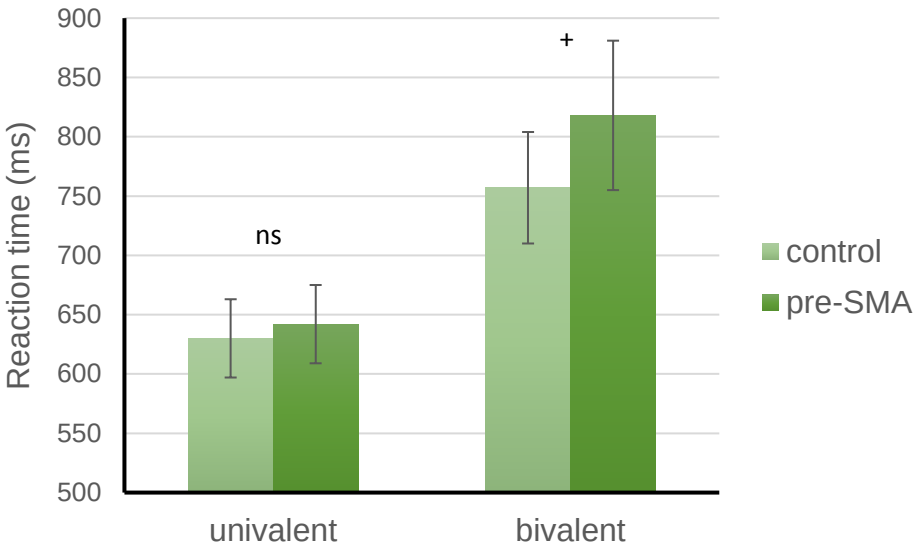

Figure 3.2. Post hoc analyses of TMS effect (pre-SMA vs. control site stimulation), performed separately for univalent items and bivalent items. Error bars indicate one standard error above and below the mean values. *: p < .05; +: *p* < .1; ns: *p* > .1.

Similar to the effects on naming latencies, the ER analysis showed main effects of valence and trial type on error rates. More errors occurred on bivalent items (*M* = 24.8 ± 5.5%) compared to univalent items (*M* = 2.5 ± 1.0%): *F* (1, 9) = 19.945, *p*

= .002, η_p^2 = .689. Switch trials induced more errors (*M* = 15.4 ± 3.7%) than stay trials did (*M* = 11.9 ± 2.4%): *F* (1, 9) = 5.301, *p* = .047, η_p^2 = .371. No other effects or interactions were significant in the ER analysis.

Effect tested	Measure	df	*F*	*p*	Partial Eta Squared
TMS	RT	1, 9	2.405	.155	.211
	ER	1, 9	2.468	.151	.215
Valence	RT	1, 9	16.988	**.003**	.654
	ER	1, 9	19.945	**.002**	.689
Trial type	RT	1, 9	3.982	*.077*	.307
	ER	1, 9	5.301	**.047**	.371
TMS * valence	RT	1, 9	4.532	*.062*	.335
	ER	1, 9	.925	.361	.093
TMS * trial type	RT	1, 9	.032	.862	.004
	ER	1, 9	.309	.592	.033
Valence * trial type	RT	1, 9	.464	.513	.049
	ER	1, 9	.455	.517	.048
TMS * valence * trial type	RT	1, 9	.705	.423	.073
	ER	1, 9	1.263	.290	.123

Table 3.2. Statistical analysis results for reaction times (RT) and error rates (ER) in Experiment 2. Factors in the analysis were valence (univalent items vs. bivalent items), trial type (stay trials vs. switch trials), and TMS condition (control site vs. pre-SMA). Statistically significant effects ($p < .05$) are shown in **bold**, and marginally significant effects are shown in *italics* ($p < .10$).

3.3 Discussion

3.3.1 Essential role of pre-SMA in item-specific inhibition

The purpose of this experiment was to investigate whether the pre-SMA plays an essential role in whole-language inhibition and/or item-specific inhibition, by examining how these two types of inhibition are affected when this brain region is disrupted. Firstly, the index for each type of inhibition, as introduced in Experiment 1, was successfully replicated here. Recall that bivalent items are subject to item-specific inhibition while univalent items are not, and naming on switch trials requires whole-language inhibition while naming on stay trials does not. Therefore, the main effect of valence in the RT and ER analyses represent the “item-specific inhibition cost”, and the main effect of trial type (marginally significant in the RT analysis and significant in ER analysis) correspond to the “whole-language inhibition cost”. As in Experiment 1, the effect of valence and that of trial type were once again shown to be vastly different in magnitude, suggesting possibly different underlying mechanisms for the two types of inhibition they represent.

The main result of interest in this experiment is the marginally significant interaction between TMS condition and valence. This interaction shows that TMS affected univalent and bivalent items differently. The TMS protocol used (cTBS) was intended to induce a reduction of cortical excitability at the stimulation site, resulting in an inhibitory effect on the target brain region. Using a control site (vertex) as baseline in comparison allows a direct examination of the consequence of target site (pre-SMA) stimulation, without the risk of the observed effect being merely a generic effect of applying TMS. Post hoc analyses revealed that TMS on pre-SMA (compared to control site) delayed naming latencies, and this delay affected bivalent items but not univalent items. This shows that inhibition of pre-SMA activity has an impact on certain processes that are peculiar to bivalent items. As the difference between bivalent and univalent items lies in whether item-specific inhibition is involved, this result offers strong evidence for a causal relationship between pre-SMA activity and item-specific inhibition. Another way to look at this is that TMS increased the existing RT difference between bivalent and univalent items (i.e. the item-specific inhibition cost). This larger cost following pre-SMA disruption indicates an essential role of this brain region in carrying out item-specific inhibition. This finding is in agreement with

Branzi, Della Rosa, Canini, Costa, and Abutalebi (2015), who suggest that item-specific inhibition in bilingual production primarily relies on control processes mediated by the pre-SMA. In contrast to valence, trial type did not have a significant interaction with TMS condition ($F < 1$), meaning that the whole-language inhibition cost was not modulated by pre-SMA perturbation. This suggests a lack of pre-SMA involvement in whole-language inhibition. Further, there was no three-way interaction between TMS condition, valence, and trial type ($F < 1$), revealing that bivalent items were affected by TMS in similar ways whether they appeared on a stay or switch trial.

One may argue that the lack of TMS modulation on the whole-language inhibition cost in this experiment was due to the latter being only marginally significant in the first place. However, even if a variable is not significant as a main effect itself, this does not prevent it from having a significant interaction with another variable. If we assume that whole-language inhibition is of similar nature to item-specific inhibition, then we should expect the whole-language inhibition cost to become larger following pre-SMA disruption (just like the item-specific inhibition cost did), regardless of whether the cost is significant to begin with. As for why the whole-language inhibition cost was only marginally significant here (whereas it was significant in Experiment 1), a possible explanation is that bilinguals became more efficient at switching through practice, as they each attended two sessions in this experiment (compared to a single session in Experiment 1). It is reasonable to consider that the behavioural task used in these experiments may attract a practice effect. Although care has been taken to ensure any practice effect would not confound or influence the TMS effect (as TMS order was fully counterbalanced between participants), the practice may have nonetheless resulted in a reduction in switch costs overall when averaging across the two sessions. In support of this, the item-specific inhibition cost in this experiment was also smaller than that in Experiment 1, showing a possible improvement in switching efficiency.

It should be acknowledged that the key finding - an interaction between TMS condition and valence - was only marginally significant in this experiment. However, this was likely a result of the small sample size used. Due to limited timeframe for this Master's project, only ten participants were included in the TMS experiment (each participant had to undergo one MRI session plus two TMS sessions). Note that the F-value for the interaction was large, so it is likely that, with more power (i.e. a

larger sample), this effect would be statistically robust[9]. In the following discussions, I will assume this effect is real.

3.3.2 Subcomponents of item-specific inhibition: what is the consequence of pre-SMA disruption?

According to the present assumption (as explained in Experiment 1 discussion, see *2.3.1*), the item-specific inhibition cost on bivalent items arises due to the time required to overcome prior inhibition applied on the current target response when it was the non-target competitor on a previous trial (where its translation-equivalent was the target response). Following this, it should be expected that disruption of inhibitory control would reduce the strength of that prior inhibition, making it easier to overcome on the current trial. Assuming that the pre-SMA is engaged to carry out inhibition in this task, the disruption of this brain region should lead to weaker inhibition and therefore faster responses for bivalent items (i.e. a reduction in item-specific inhibition cost). However, the results show the opposite. Disruption of pre-SMA increased naming latencies on bivalent trials, instead of reducing them.

We could speculate that the disruption of pre-SMA led to other compensatory mechanism kicking in to assist in order to successfully inhibit the non-target response, ultimately resulting in stronger inhibition, which then took longer to overcome subsequently. However, there is a more parsimonious explanation if we are willing to put aside the assumption that the item-specific inhibition cost simply reflects the time required to recover from prior inhibition of the currently relevant response. While the process of overcoming inhibition may certainly be a component of this cost, it is likely that the time required to suppress the currently irrelevant response (i.e. interference suppression) is also a component. Here I consider the possibility that the "interference suppression" component is what gets prolonged due to pre-SMA disruption, leading to the observed RT increase on bivalent items. It has been shown that disrupting the activity of pre-SMA can slow down the inhibition process such that it takes more time to complete successfully (Obeso, Robles, Marron, & Redolar-Ripoll, 2013). Therefore, after TMS was delivered over this brain region, it might take longer to achieve the appropriate level of item-specific inhibition

[9] I will be testing more participants in the near future to confirm this.

in order to prevent the prepotent (but incorrect) response from being selected and reaching speech output. This process is required for bivalent items only, therefore bivalent items (and not univalent items) suffered longer RTs following pre-SMA disruption. Note that both components mentioned above (suppressing currently irrelevant response and overcoming prior inhibition of the current target response) are distinguishing features of bivalent trials (compared to univalent trials), so it is possible that either or both of them are impacted by pre-SMA disruption and contribute to the RT modulation. In other words, pre-SMA disruption may have indeed resulted in weaker item-specific inhibition, which took less time to overcome (i.e. RT facilitation for bivalent items), but at the same time the disruption also led to increased difficulty in carrying out inhibition on the current trial, which made these RTs longer. If the prolonging effect was large enough, the RT facilitation could be masked (and reversed) such that the overall observable effect was longer RT. The present results cannot discern whether this hidden RT facilitation in one component exists alongside the RT increase in another component; however, if the facilitation exists, it can only mean that the actual amount of RT increase in the latter component must have been even greater than what was observed. Therefore, either way there is robust evidence supporting the role of pre-SMA in the inhibition of interfering responses (i.e. words in the non-target language).

It is interesting to note that, unlike in the RT analysis, there is no interaction between TMS condition and valence in the ER analysis. This is contrary to my original prediction. I predicted that TMS would significantly increase error rates on bivalent items but not on univalent items, because disruption of pre-SMA function would affect item-specific inhibition (which is only applicable to bivalent items) and cause more selection errors where this type of inhibition is required. However, there was no evidence of such interaction. Bivalent items induced a lot more errors overall compared to univalent items (i.e. main effect of valence), but this effect was constant regardless of the TMS condition applied. The most likely reason for this lack of TMS modulation on error rates is that the slight disruption of the pre-SMA achieved with this protocol was not enough to cause bilinguals to actually output the wrong lexical item; instead, it just hindered the process of inhibiting the competing non-target response and selecting the correct one. In other words, the pre-SMA disruption may have resulted in a reduced ability to inhibit the undesired response, but the inhibition

was nonetheless successful (just taking longer to accomplish), and the correct target word was still produced in the end. Therefore, the effect of TMS disruption appeared in the form of longer reaction times, not higher error rates.

3.3.3 Lack of pre-SMA involvement in whole-language inhibition

The results of this experiment show that pre-SMA plays an essential role in item-specific inhibition, but no evidence was found for its involvement in whole-language inhibition. Given that the presumed role of the pre-SMA is to carry out inhibition, why is this brain area not uniformly engaged in both types of inhibition? To answer this question, it may be helpful to consider what the differences are between whole-language and item-specific inhibition (other than the scope of influence).

Firstly, De Groot and Christoffels (2006) state that whole-language and item-specific inhibition occur at different time points during lexical selection. Whole-language control occurs in the early stage to adjust the activation levels of the two languages according to task demand, such that all elements of the target language are made more available for selection and at the same time all elements of the non-target language are made less available. Item-specific control, on the other hand, acts at a later stage and serves to catch any non-target language lexical nodes that are highly activated despite the early global adjustments. In other words, whole-language inhibition creates a preparatory setting in the language system to facilitate correct selection, while item-specific inhibition performs the checking procedure just before output from the system. The different timing at which the two types of inhibition are in action and the distinct functions they serve may demand different neural mechanisms to be engaged. While the pre-SMA was found to be essential in carrying out item-specific inhibition in this experiment, whole-language inhibition may be accomplished using a different mechanism (which may or may not be domain-general), and thus was not modulated by the disruption of pre-SMA activity.

Secondly, De Groot (2011) suggests that the two types of inhibition may exert control on different targets: while whole-language control acts on the language system proper, item-specific control targets imminent outputs of the system to prevent undesired words from appearing in speech. In this sense, item-specific inhibition may

operate via a mechanism similar to response inhibition. Since pre-SMA is important in carrying out response inhibition (Aron & Poldrack, 2006), this may explain why pre-SMA is engaged for item-specific but not whole-language inhibition. Going one step further, what was observed as “item-specific inhibition” in this experiment could simply be response inhibition (to hold back the prepotent response). While this may suggest that the pre-SMA does not really have a role in language control after all (since it merely carries out vocal response inhibition in this task), the fact remains that this type of language switching (i.e. changing between two languages in naming the same item) requires the pre-SMA, whether its role is to suppress lexical representations during selection or to inhibit undesired responses at the output stage.

3.3.4 Summary

There was an important trend showing that the pre-SMA plays an essential role in item-specific inhibition but not whole-language inhibition. While this finding confirms the involvement of domain-general inhibitory mechanism in language switching, it also suggests that this mechanism may be engaged under certain circumstances only. Specifically, language switching is unlikely to recruit domain-general control when each concept remains uniquely associated with a particular language (i.e. univalent items in this experiment), even if language switches are performed from one concept to the next (i.e. switch trials). A real life example of this would be using a mixture of two languages in a conversation, but consistently referring to each concept using the same word. On the other hand, alternating between translation-equivalents that correspond to the same concept (i.e. bivalent items) is likely to engage domain-general mechanism. Examples of this include translating after oneself (i.e. repeating the same message in the other language), or switching between two languages in such a way that the same concepts are mentioned in both languages.

Chapter 4: General Discussion

This thesis aimed to answer two questions about bilingual language control. The first question was in regards to whether both item-specific and whole-language inhibition are involved in language switching. This question arose out of the recent debate on the distinction between two levels of control in bilingual production. Experiment 1 examined this issue by incorporating both item-specific and whole-language inhibition into a cued language switching paradigm and comparing them directly. The findings indicate that both types of inhibition are at work during language switching, and they are each associated with a "cost". Item-specific inhibition was indexed by the effect of valence (i.e. naming latencies for bivalent items were significantly slower compared to univalent items), and whole-language inhibition was indexed by the effect of trial type (i.e. naming latencies on switch trials were slower than on stay trials). The cost of item-specific inhibition was found to be much greater in magnitude compared to whole-language inhibition.

The second question in this thesis concerned whether both types of inhibition operate via domain-general mechanisms. This question arose from the growing amount of empirical evidence suggesting a close relationship between language control and executive function in bilinguals (Bobb et al., 2013). Experiment 2 investigated the involvement of executive control in language switching by perturbing the pre-SMA, a brain area responsible for domain-general inhibitory control, with non-invasive brain stimulation. A repetitive TMS protocol was used to disrupt the functioning of this brain area, and the consequence on language switching performance was examined. The results showed that the naming of univalent and bivalent items were differentially impacted by the perturbation of pre-SMA, while no such difference was found between stay trials and switch trials. This revealed an essential role of the pre-SMA in item-specific, but not whole-language, inhibition.

In this chapter, I would like to discuss two particular issues. In section 4.1, I will talk about the role of pre-SMA in language control and how this relates to its general role in executive function. In section 4.2, I will take a closer look at whole-language

inhibition and explore an alternative explanation on why it was not affected by the disruption of the pre-SMA.

4.1 The role of pre-SMA in language control

In the TMS experiment, disruption of pre-SMA was shown to increase naming latencies for bivalent stimuli, regardless of whether they appeared on a stay or switch trial. Univalent stimuli, on the other hand, were unaffected by this disruption. This highlights a role of the pre-SMA in resolving conflict when there are two highly accessible names for the same stimulus competing for selection, but not when switching between languages in the absence of such fierce competition. This aligns well with the theoretical account in the latest neurocognitive model of language control (Green & Abutalebi, 2013), which proposes a role for the pre-SMA in conflict monitoring and resolution. As discussed in *3.3.2*, the pre-SMA may perform such conflict resolution by means of interference suppression, i.e. inhibiting the currently irrelevant, prepotent response to allow successful selection of the correct (but possibly less automatic) response. Outside the linguistic domain, the pre-SMA has also been implicated in resolving conflicts in action selection. For example, Nachev, Rees, Parton, Kennard, and Husain (2005) found pre-SMA involvement in situations of response conflict, and they suggest that its function may be related to resolving competition between different action plans to allow the desired action to be performed. In linking the functioning of this brain region across domains, Abutalebi et al. (2012) compared brain activities in bilinguals who performed a language switching task and a non-linguistic conflict resolution task (flanker task), and they found the pre-SMA to be similarly involved across these tasks.

All together these findings point to a role of the pre-SMA in response selection under high-conflict situations, and suggest that the involvement of the pre-SMA in language control is indeed related to its general role in executive function. Crucially, the necessity of pre-SMA in item-specific inhibition but not whole-language inhibition suggests that this brain region resolves conflicts at the level of responses or stimulus-response associations (i.e. switching between different possible responses for the same stimulus), rather than at the task level (i.e. switching between languages from trial to trial). As Branzi et al. (2015) point out, the pre-SMA is

engaged in cases where stimulus-response remapping is required[10] (i.e. when a picture has previously been named in the other language), and that is why this brain region is essential in the naming of bivalent but not univalent items.

4.2 An alternative view on whole-language inhibition

Allow me to consider an alternative answer to the question of why item-specific inhibition engages the pre-SMA but whole-language inhibition does not. In the discussions so far, I have interpreted this as evidence for dissociable inhibitory mechanisms underlying item-specific and whole-language inhibition, due to the distinct functions they serve in language control. Here I would like to embark on a slightly different path and take a closer look at what we have been referring to as "whole-language inhibition".

According to Green's (1998) formulation of whole-language inhibition, successful naming on switch trials requires inhibition of one language task schema (the currently active one) and activation of the other. The inhibition of the currently active language task schema then passes on the inhibition to all lexical representations in that language (De Groot & Christoffels, 2006). Following this account, suppression of lexical representations takes place on every switch trial, similar to the inhibitory process encountered on bivalent items. This is what led to the perplexing question of what makes one type of inhibition different from another. Here I propose a small modification to the above account to solve this problem: whole-language control operates on the level of language task schemas only, and does not extend suppression down to the level of lexical representations. In this way, early whole-language control serves the purpose of setting task goal (i.e. which language to output in) without regulating activation levels of lexical representations, leaving the latter to be the responsibility of the late reactive control. As De Groot (2011, p. 309) points out, selection of non-target language lexical nodes could be successfully

[10] Although the design of the present study did not strictly enforce that every bivalent trial required a stimulus-response remapping (i.e. between two critical trials with the same picture stimulus, there was always a filler with that picture, to be named in the non-target response language), it is very likely that most of them did (as there were twice as many fillers as there were critical trials). This can be examined more carefully in the future by controlling whether there was a change of language from the last time that same picture was named, thus allowing a direct examination of the consequence of stimulus-response remapping.

prevented by reactive control alone. Thus, the so-called “whole-language inhibition cost” does not reflect the process of carrying out inhibition on lexical representations or overcoming such inhibition exerted previously; instead, the cost relates to switching between language task schemas only (i.e. for one schema to take over the other).

A piece of evidence that offers support for this proposal rests in the very similar switch costs found on univalent and bivalent items. In the following discussion, I will refer to the results from the behavioural experiment (presented in Chapter 2), as they are not subject to any potential contamination from TMS-related effects. Recall from Chapter 2 that whole-language inhibition affects switch trials but not stay trials, while item-specific control affects bivalent items but not univalent items. If we assume whole-language and item-specific inhibition both alter activation levels of lexical representations (De Groot & Christoffels, 2006), then a bivalent item appearing on a switch trial should experience the combined effect from the two types of inhibition. It is unlikely that these two sources of inhibition acting on the same lexical node would simply stack their effects on top of each other, resulting in an overall cost that is the linear sum of the two inhibition costs. Rather, the smaller cost (from whole-language inhibition) would likely be masked by the larger cost (from item-specific inhibition). In other words, since bivalent items are already strongly suppressed by item-specific inhibition on both stay and switch trials, whole-language inhibition (which occurs on switch trials) should not be able to suppress them much further, and therefore should have little additional effect. It follows that bivalent items should have similar naming latencies on stay and switch trials, i.e. minimal or drastically reduced “switch cost”. However, robust switch costs were obtained for both univalent and bivalent stimuli, and the magnitude of univalent and bivalent switch costs only differed by 1ms. This finding challenges the idea that the switch cost arises from inhibition on the level of lexical representations, and supports my proposal that whole-language control acts on the level of language task schemas only.

Now if the switch cost resides at the level of language task schemas, it most likely represents one or more of the processes involved in switching between these task schemas. As Green (1998) constructed his account of language control based on the control of actions, here I will generalise the notion of language task schemas and assume they are like any other task sets. Thus, we can say the switch cost relates to

switching between task sets (L1 output or L2 output). This cost may or may not be related to inhibitory processes acting on the task-set level. Although the proposal under consideration stipulates whole-language control does not exert inhibition on lexical representations, processes such as suppressing the currently irrelevant task (i.e. the language task schema to switch away from) and overcoming prior inhibition of the currently relevant task (i.e. the language task schema to switch into) may nonetheless be applicable to task sets. However, even if these processes exist on the task-set level, it would not be surprising if the neural mechanism of inhibiting task sets is different from that of inhibiting individual lexical nodes (or stimulus-response associations), as one can expect the mental representations of tasks to be quite different in nature from mental representations of words (i.e. lexical nodes).

Furthermore, evidence suggests an absence of such inhibition even at the task-set level. In the task switching literature, it is generally accepted that the switch cost comprises of at least two components. One component relates to endogenous control processes that can be completed before stimulus onset; this part of the switch cost has been shown to be avoidable with advance preparation for the upcoming switch (Meiran, 1996; Rogers & Monsell, 1995; Karayanidis, Coltheart, Michie, & Murphy, 2003). The other component (often called “residual switch cost”) remains even if there is ample preparation time, and is thought to reflect the cost of resolving lingering activation of previously executed task or lingering inhibition of currently relevant task (Arbuthnott & Frank, 2000). The switch cost asymmetry in bilingual production, which is often considered the hallmark of inhibitory control in the language switching literature, has been shown to disappear when ample preparation time is given before stimulus onset. Verhoef et al. (2009) compared the effect of short and long cue-stimulus intervals, and found that the asymmetry in switch costs is eliminated with long interval. This suggests that the asymmetry resides in the first component of switch cost mentioned above, rather than being part of the “residual switch cost” and resulting from the different amount of time required to overcome prior inhibition applied at varied strengths. It should be noted though that such elimination of switch cost asymmetry was not universally observed when long preparation time was given (see Philipp et al., 2007).

In summary, the so-called “whole-language inhibition” may not involve inhibition of lexical representations at all (and possibly not even inhibition of language task

schemas). Hence, the “whole-language inhibition cost” was not affected by the disruption of inhibitory control in the TMS experiment, while the item-specific inhibition cost, which was truly due to inhibition, was affected. This possibility stands as an alternative explanation for the findings from Experiment 2.

Chapter 5: Summary

This thesis makes two important contributions to the language switching literature. Firstly, most language-switching studies conducted so far have focussed on examining trial-to-trial switching (which is thought to reflect inhibitory control of complete language systems, i.e. whole-language inhibition), with a lack of attention on how this differs from the kind of control required when the same concept must be named in two languages alternately (i.e. item-specific inhibition). Although many studies have used bivalent stimuli (repeated stimuli which required naming in both languages), only the trial-to-trial switch costs were examined, while the effect of being bivalent was left unexplored (as there were usually no univalent stimuli to compare to). To my knowledge, Experiment 1 in this thesis is the first behavioural experiment to make a direct comparison (without any confounding factors) between whole-language and item-specific inhibition in the cued language-switching paradigm. Such comparison is important as it provides the unique opportunity to explore different types of inhibitory control in language switching and the different mechanisms underlying these control processes.

The second major contribution of this thesis is the establishment of a causal relationship between pre-SMA excitability and performance in language switching. While several neuroimaging studies have already reported activation of this brain region during language switching, it remained unclear whether the pre-SMA played an essential role or simply co-activated with the language control network. Using non-invasive brain stimulation, Experiment 2 in this thesis demonstrated a performance impact when the pre-SMA was disrupted, uncovering an essential role of this brain region in bilingual language control. Thus, the present findings complement the neuroimaging evidence and support the proposal that language control in bilinguals relies on executive function. Crucially, the pre-SMA was found to be essentially engaged in item-specific inhibition but not whole-language inhibition, suggesting that not all types of switching effectively exercise executive control. For example, the pre-SMA might not be recruited for simply switching between languages, but only in particular situations which involve a high degree of response

conflict, e.g. when naming the same item in different languages and switching between these possible names.

These are exciting findings which have potentially very practical implications. Given the considerable current interest in enhancing executive function and combating dementia and cognitive decline, language switching represents a simple exercise people can do in daily life[11] (without going out of their way or spending any extra time). The findings in this thesis provide some useful direction on what type of language switching might be effective in exercising executive control and gaining cognitive benefits. As suggested above, merely using more than one language may not be effective - the key is to talk about the same topics in both languages. For example, if you speak Mandarin with your family and communicate in English with your friends, describing the same events to both groups of people would constitute a good exercise for your executive control system. On the other hand, if your two languages are kept completely separate for use in different domains of your life, being bilingual might not give you any advantage in cognitive function.

[11] Note that this is not necessarily restricted to bilinguals. It has previously been suggested that lexical selection in bilinguals is qualitatively no different than that in monolinguals. Given that bilingual language control is implemented via domain-general inhibitory mechanism (as the present findings suggest), it is well within reason to suppose that monolingual speakers might use this same inhibitory mechanism to select between words which have similar meaning but are used in different context (e.g. register selection). If this is the case, then monolingual speakers can practice executive control in a similar fashion.

References

Abutalebi, Della Rosa, P. A., Ding, G., Weekes, B., Costa, A., & Green, D. W. (2013). Language proficiency modulates the engagement of cognitive control areas in multilinguals. *Cortex, 49*(3), 905-911.

Abutalebi, & Green, D. (2007). Bilingual language production: The neurocognition of language representation and control. *Journal of Neurolinguistics, 20*(3), 242-275. doi:10.1016/j.jneuroling.2006.10.003

Abutalebi, & Green, D. W. (2008). Control mechanisms in bilingual language production: Neural evidence from language switching studies. *Language and cognitive processes, 23*(4), 557-582.

Abutalebi, J., Della Rosa, P. A., Green, D. W., Hernandez, M., Scifo, P., Keim, R., . . . Costa, A. (2012). Bilingualism Tunes the Anterior Cingulate Cortex for Conflict Monitoring. *Cerebral Cortex, 22*(9), 2076-2086.

Arbuthnott, K., & Frank, J. (2000). Executive control in set switching: Residual switch cost and task-set inhibition. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 54*(1), 33.

Aron, A. R., Durston, S., Eagle, D. M., Logan, G. D., Stinear, C. M., & Stuphorn, V. (2007). Converging evidence for a fronto-basal-ganglia network for inhibitory control of action and cognition. *The Journal of Neuroscience, 27*(44), 11860-11864.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *The Journal of Neuroscience, 26*(9), 2424-2433.

Bobb, S. C., & Wodniecka, Z. (2013). Language switching in picture naming: What asymmetric switch costs (do not) tell us about inhibition in bilingual speech planning. *Journal of Cognitive Psychology, 25*(5), 568-585. doi:10.1080/20445911.2013.792822

Bobb, S. C., Wodniecka, Z., & Kroll, J. F. (2013). What bilinguals tell us about cognitive control: Overview to the special issue. *Journal of Cognitive Psychology, 25*(5), 493-496.

Branzi, F. M., Della Rosa, P. A., Canini, M., Costa, A., & Abutalebi, J. (2015). Language control in bilinguals: Monitoring and response selection. *Cerebral Cortex*, bhv052.

Cai, W., George, J. S., Verbruggen, F., Chambers, C. D., & Aron, A. R. (2012). The role of the right presupplementary motor area in stopping action: two studies with event-related transcranial magnetic stimulation. *Journal of neurophysiology, 108*(2), 380-389.

Campbell, J. I. (2005). Asymmetrical language switching costs in Chinese–English bilinguals' number naming and simple arithmetic. *Bilingualism: Language and Cognition, 8*(01), 85-91.

Chiou, R., Sowman, P. F., Etchell, A. C., & Rich, A. N. (2014). A conceptual lemon: Theta burst stimulation to the left anterior temporal lobe untangles object representation and its canonical color. *Journal of Cognitive Neuroscience, 26*(5), 1066-1074.

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

Colomé, À. (2001). Lexical activation in bilinguals' speech production: Language-specific or language-independent? *Journal of Memory and Language, 45*(4), 721-736.

Costa, A., & Caramazza, A. (1999). Is lexical selection in bilingual speech production language-specific? Further evidence from Spanish–English and English–Spanish bilinguals. *Bilingualism: Language and Cognition, 2*(03), 231-244.

Costa, A., Caramazza, A., & Sebastian-Galles, N. (2000). The cognate facilitation effect: implications for models of lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(5), 1283.

Costa, A., Miozzo, M., & Caramazza, A. (1999). Lexical selection in bilinguals: Do words in the bilingual's two lexicons compete for selection? *Journal of Memory and Language, 41*(3), 365-397.

Costa, A., & Santesteban, M. (2004). Lexical access in bilingual speech production: Evidence from language switching in highly proficient bilinguals and L2 learners. *Journal of Memory and Language, 50*(4), 491-511.

Costa, A., Santesteban, M., & Ivanova, I. (2006). How do highly proficient bilinguals control their lexicalization process? Inhibitory and language-specific selection mechanisms are both functional. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*(5), 1057.

De Baene, W., Duyck, W., Brass, M., & Carreiras, M. (2015). Brain circuit for cognitive control is shared by task and language switching. *Journal of Cognitive Neuroscience*.

de Bruin, A., Roelofs, A., Dijkstra, T., & FitzPatrick, I. (2014). Domain-general inhibition areas of the brain are involved in language switching: FMRI evidence from trilingual speakers. *Neuroimage, 90*, 348-359.

De Groot, A. M. (2011). *Language and cognition in bilinguals and multilinguals: An introduction*: Psychology Press.

De Groot, A. M., & Christoffels, I. K. (2006). Language control in bilinguals: Monolingual tasks and simultaneous interpreting. *Bilingualism: Language and Cognition, 9*(02), 189-201.

Dietrich, S., Hertrich, I., Ackermann, H., Ziemann, U., & Müller-Dahlhaus, F. (2015). P76. Transient suppression of speech comprehension by continuous theta-burst magnetic stimulation of the pre-SMA. *Clinical Neurophysiology, 126*(8), e131-e132.

Fedorenko, E., Duncan, J., & Kanwisher, N. (2013). Broad domain generality in focal regions of frontal and parietal cortex. doi:10.1073/pnas.1315235110

Finkbeiner, M., Almeida, J., Janssen, N., & Caramazza, A. (2006). Lexical selection in bilingual speech production does not involve language suppression. *Journal of experimental psychology. Learning, memory, and cognition, 32*(5), 1075-1089. doi:10.1037/0278-7393.32.5.1075

Finkbeiner, M., & Caramazza, A. (2006). Now you see it, now you don't: On turning semantic interference into facilitation in a Stroop-like task. *Cortex, 42*(6), 790-796.

Finkbeiner, M., Gollan, T. H., & Caramazza, A. (2006). Lexical access in bilingual speakers: What's the (hard) problem? *Bilingualism: Language and Cognition, 9*(02), 153-153. doi:10.1017/S1366728906002501

French, R. M., & Jacquet, M. (2004). Understanding bilingual memory: models and data. *Trends in cognitive sciences, 8*(2), 87-93.

Gollan, T. H., & Ferreira, V. S. (2009). Should I stay or should I switch? A cost-benefit analysis of voluntary language switching in young and aging bilinguals. *J Exp Psychol Learn Mem Cogn, 35*(3), 640-665. doi:10.1037/a0014981

Gollan, T. H., Kleinman, D., & Wierenga, C. E. (2014). What's easier: doing what you want, or being told what to do? Cued versus voluntary language and task switching. *Journal of experimental psychology. General, 143*(6), 2167-2195. doi:10.1037/a0038006

Gollan, T. H., Weissberger, G. H., Runnqvist, E., Montoya, R. I., & Cera, C. M. (2012). Self-ratings of Spoken Language Dominance: A Multi-Lingual Naming Test (MINT) and Preliminary Norms for Young and Aging Spanish-English Bilinguals. *Biling (Camb Engl), 15*(3), 594-615. doi:10.1017/S1366728911000332

Green, D. W. (1998). Mental control of the bilingual lexico-semantic system. *Bilingualism: Language and Cognition, 1*(02), 67. doi:10.1017/S1366728998000133

Green, D. W., & Abutalebi, J. (2013). Language control in bilinguals: The adaptive control hypothesis. *Journal of Cognitive Psychology, 25*(5), 515-530. doi:10.1080/20445911.2013.796377

Guillaume, C., Guillery-Girard, B., Chaby, L., Lebreton, K., Hugueville, L., Eustache, F., & Fiori, N. (2009). The time course of repetition effects for familiar faces and objects: An ERP study. *Brain research, 1248*, 149-161.

Guo, T., Liu, H., Misra, M., & Kroll, J. F. (2011). Local and global inhibition in bilingual word production: fMRI evidence from Chinese-English bilinguals. *Neuroimage, 56*(4), 2300-2309. doi:10.1016/j.neuroimage.2011.03.049

Hanulová, J., Davidson, D. J., & Indefrey, P. (2011). Where does the delay in L2 picture naming come from? Psycholinguistic and neurocognitive evidence on second language word production. *Language and cognitive processes, 26*(7), 902-934.

Heikoop, K. W., Declerck, M., Los, S. A., & Koch, I. (in press). Dissociating language-switch costs from cue-switch costs in bilingual language switching.

Hermans, D., Bongaerts, T., De Bot, K., & Schreuder, R. (1998). Producing words in a foreign language: Can speakers prevent interference from their first language? *Bilingualism: Language and Cognition, 1*(03), 213-229.

Huang, Y.-Z., Chen, R.-S., Rothwell, J. C., & Wen, H.-Y. (2007). The after-effect of human theta burst stimulation is NMDA receptor dependent. *Clinical Neurophysiology, 118*(5), 1028-1032.

Huang, Y.-Z., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta burst stimulation of the human motor cortex. *Neuron, 45*(2), 201-206.

Jackson, G. M., Swainson, R., Cunnington, R., & Jackson, S. R. (2001). ERP correlates of executive control during repeated language switching. *Bilingualism: Language and Cognition, 4*(02), 169-178.

Jackson, S. R., Jackson, G. M., & Roberts, M. (1999). The selection and suppression of action: ERP correlates of executive control in humans. *Neuroreport, 10*(4), 861-865.

Karayanidis, F., Coltheart, M., Michie, P. T., & Murphy, K. (2003). Electrophysiological correlates of anticipatory and poststimulus components of task switching. *Psychophysiology, 40*(3), 329-348.

Kleinman, D., & Gollan, T. H. (2016). Speaking two languages for the price of one: Bypassing language control mechanisms via accessibility-driven switches. *Psychological Science.*, 1-39. doi:10.1177/0956797616634633

Kroll, J. F., Bobb, S. C., Misra, M., & Guo, T. (2008). Language selection in bilingual speech: Evidence for inhibitory processes. *Acta psychologica, 128*(3), 416-430.

La Heij, W. (2005). Selection processes in monolingual and bilingual lexical access. *Handbook of bilingualism*, 289.

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance, 29*(3), 575.

Luk, G., Green, D. W., Abutalebi, J., & Grady, C. (2012). Cognitive control for language switching in bilinguals: A quantitative meta-analysis of functional neuroimaging studies. *Language and cognitive processes, 27*(10), 1479-1488.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*(6), 1423.

Meuter, R. F., & Allport, A. (1999). Bilingual language switching in naming: Asymmetrical costs of language selection. *Journal of Memory and Language, 40*(1), 25-40.

Misra, M., Guo, T., Bobb, S. C., & Kroll, J. F. (2012). When bilinguals choose a single word to speak: Electrophysiological evidence for inhibition of the native language. *Journal of Memory and Language, 67*(1), 224-237.

Nachev, P., Kennard, C., & Husain, M. (2008). Functional role of the supplementary and pre-supplementary motor areas. *Nature Reviews Neuroscience, 9*(11), 856-869.

Nachev, P., Rees, G., Parton, A., Kennard, C., & Husain, M. (2005). Volition and conflict in human medial frontal cortex. *Current Biology, 15*(2), 122-128.

Norman, D. A., & Shallice, T. (1986). Attention to action *Consciousness and self-regulation* (pp. 1-18): Springer.

Obeso, I., Robles, N., Marron, E. M., & Redolar-Ripoll, D. (2013). Dissociating the Role of the pre-SMA in Response Inhibition and Switching: A Combined Online and Offline TMS Approach. *Front Hum Neurosci, 7*, 150. doi:10.3389/fnhum.2013.00150

Philipp, A. M., Gade, M., & Koch, I. (2007). Inhibitory processes in language switching: Evidence from switching language-defined response sets. *European Journal of Cognitive Psychology, 19*(3), 395-416.

Philipp, A. M., & Koch, I. (2009). Inhibition in language switching: what is inhibited when switching between languages in naming tasks? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*(5), 1187.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of experimental psychology: General, 124*(2), 207.

Sale, M. V., Ridding, M. C., & Nordstrom, M. A. (2008). Cortisol inhibits neuroplasticity induction in human motor cortex. *The Journal of Neuroscience, 28*(33), 8285-8293.

Schwieter, J. W., & Sunderman, G. (2008). Language switching in bilingual speech production: In search of the language-specific selection mechanism. *The Mental Lexicon, 3*(2), 214-238.

Tremblay, P., & Gracco, V. L. (2009). Contribution of the pre-SMA to the production of words and non-speech oral motor gestures, as revealed by repetitive transcranial magnetic stimulation (rTMS). *Brain research, 1268*, 112-124.

Van Assche, E., Duyck, W., & Gollan, T. H. (2013). Whole-language and item-specific control in bilingual language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 39*(6), 1781.

Verhoef, K., Roelofs, A., & Chwilla, D. J. (2009). Role of inhibition in language switching: Evidence from event-related brain potentials in overt picture naming. *Cognition, 110*(1), 84-99.

Vernet, M., Bashir, S., Yoo, W. K., Oberman, L., Mizrahi, I., Ifert-Miller, F., . . . Pascual-Leone, A. (2014). Reproducibility of the effects of theta burst stimulation on motor cortical plasticity in healthy participants. *Clin Neurophysiol, 125*(2), 320-326. doi:10.1016/j.clinph.2013.07.004

Xue, G., Aron, A. R., & Poldrack, R. A. (2008). Common neural substrates for inhibition of spoken and manual responses. *Cerebral Cortex, 18*(8), 1923-1932.

Zhang, Y., Wang, T., Huang, P., Li, D., Qiu, J., Shen, T., & Xie, P. (2015). Free Language Selection in the Bilingual Brain: An Event-Related fMRI Study. *Scientific reports, 5*.