THE EFFECTS OF FLUENCY OF LEXICAL AND PHONOLOGICAL RETRIEVAL PROCESSING ON SPOKEN WORD DURATION

Jason Michael Kahn

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Psychology (Cognitive Program).

Chapel Hill
2015

Approved by:

Jennifer E Arnold

Peter C Gordon

Neil Mulligan

Mark Hollins

Elliott Moreton

ABSTRACT

Jason Michael Kahn: The effects of fluency of lexical and phonological retrieval processing on spoken word duration
(Under the direction of Jennifer E. Arnold)

Spoken words vary considerably in their duration, for a variety of reasons. The current work focused specifically on two of those reasons – the difficulty of lexical retrieval and phonological retrieval and integration – to assess whether they play a significant role in influence the duration of the time to begin speaking and the duration of the words in a simple utterance. Previous work of mine and others has shown that word duration is sensitive to factors that may influence lexical and/or phonological retrieval (Bell et al., 2009; Gahl, 2008; Kahn & Arnold, 2012; Schriefers et al., 1999; *inter alia*), but never in a context that examined the timing of the mechanisms underlying those factors. The current experiments manipulate the semantic and phonological relationship between a prime and a target utterance in order to investigate the mechanisms and their timing. They also manipulate the difficulty of processing other words in a multi-word utterance than the target, to assess whether word duration is affected by that difficulty. The first two experiments focused on the semantic relationship between prime and target, and found that, consistent with prior work, semantic relatedness between prime and target lengthened utterance latencies to begin speaking, while primes that were identical to the target shortened them. Semantic relatedness did not have a great deal of influence on word duration, however, as contrasted with identical primes (which shortened both latencies and durations) and phonologically-related primes. The second two experiments focused on the phonological

relationship between prime and target. Once again consistent with prior work, this relationship elicited shorter latencies. Importantly, it also elicited durational lengthening, most likely as a function of the speaker's language production system having to ensure that it produces the target's phonemes instead of the prime's. Experiments 2 and 4 also included a manipulation of the final word in the utterance (not the target), whether it was consistent between trials or randomly chosen. When it was consistent, it tended to shorten latency to begin speaking, and somewhat less-reliably the duration of the object word and the final word. The findings together show that lexical retrieval is less involved in the determination of duration than phonological retrieval and/or integration. Further, they show that the language production system does not determine the duration of the words in an utterance before it begins speaking that utterance – rather, word duration emerges partly as a function of the ongoing ease or difficulty of completing processing, particularly phonological retrieval.

To my parents and the others who have done so much to support my mind

ACKNOWLEDGEMENTS

If you deign to read this at all, you're probably looking for your own name, so I'll tell you right here that I'm at the end of my rope with this damn thing, and if I forget to name you explicitly, you still have my sincere gratitude, even if that sounds out of character.

I'm sorry, mom and dad, but Jennifer's gotta get first billing on this, because she's been the biggest controller of variance – and I mean that statistically and colloquially. You've always been zealous, and even though it freaked me out a bit to have you follow us to the reception on interview day to ask *more* questions, I suspect that the only reason I managed to do as well as I did in grad school was that zealotry. It took a while for you to get used to me not taking anything seriously, but it's the only way I feel really comfortable with someone, and, personal or professional, our conversations were always challenging and rewarding. You introduced me to everyone, because you *know* everyone, because you do such good and interesting work, and give a damn about other people's – always. It's tiring, sometimes! Thank you, and I'm sorry, and hooray!, and congratulations, and whew.

Second billing, then, but only by a little. I'm lucky enough to have been born into a family with perfectly supportive, perfectly loving, unerringly positive parents, and a sister to match. No one ever questioned this decision except whether I thought it was right for me, and I've never once experienced the doubt and frustration and responsibility that I know some people feel about their home life. You've worked your asses off to keep me free to pursue this crazy track, and I know that it's been hard to keep up the enthusiasm, but know that this (and my life in general) would have been much harder and much worse without you. Thank you, and I'm sorry, and hooray!, and congratulations, and whew.

To Nykoel, for infinite, unending, ceaseless patience during what has possibly been the dullest and hardest part of my life - I promise it'll get better. For adventure, civilizing, fun, discussion, debate, argument, novelty, and love. It was time for me to start doing something new with my life, and you're the only reason I've been able to make it happen. Thank you, and I'm sorry, and hooray!, and congratulations, and whew. What shall we do next?

Special thanks to Joseph, for loving the same bits of language and language research I do, for taking the initiative, always, and for being patient with my weirdnesses. To Erol (and Alisha), for being an enclave of fun, insight, family, friendship, and once in a while serious scholarly discussion. To Elise and Alex, my labmates in arms – sometimes for, sometimes against. To Becky, for Rock Band, tequila, neural synchrony, and a great deal of insight into possibilities. To Yi Ting, for being even more zealous than Jennifer, never looking down on anyone less zealous, and that one time you sent me a message meant for your husband. To DJ (and later Katie), for bottomless inquisitiveness, Yanni, Grungebob, and a home away from academia whenever I needed it. To Milton, for having the best name, and always being able and willing to argue an esoteric point of cognitive psychology at the drop of a hat. To (another) Elise, for saving my life with antidepressants, for saving my soul with rats, and for playing so. many. board games. To Julia, for steak and scotch, for television and a return to high-mindedness, and fanciness. To James, for rocking the fuck out. To Lauren, for a reason to leave Chapel Hill and for saving my rats. To Mike, for addiction.

Thanks also to: Kay, John, Duane, Erica, Agnieszka, Stef, and especially Heidi, for getting me started on all this, although some days I think maybe I should curse you as well. To (another) Erica, for Wittgenstein. To Florian, for statistics and a stern talking-to about givenness, and teaching your students to shotgun beer. To Mara, for solidarity. To Peter, for being utterly inscrutable but scarily good at what he does, and for taking such good care of his students (my friends). To Neil, for being the most impressive scholar I've ever met, even if he picked the wrong thing to study. To Mark, for dauntlessly defending the Psychology department against all challenges, while never losing a modicum of humanity, nor research acumen, nor approachable wisdom. To Elliott, for being almost infuriatingly detail-oriented, a champion for linguistics nerds, and an unimpeachable teacher. To Jen, for forcing me to work up to standard, and for giving me a second chance. To Bruno, for knowing when to be serious and when to have fun, and for asking the hard questions. To Jean-Louis and Joe, for giving a damn about something other than just psychology, and passing on that joy. To Beth, the best lecturer I know, who taught me firm patience in teaching. To all of my research assistants, especially those who helped through the hardest parts of the dissertation, when I couldn't do this by myself even if I'd wanted to. To all of my participants, who never cared a whit about language, even though they use it all day every day.

And to you, the taxpayer, because without funding, how would we know which research was valuable? In all seriousness, my sincere gratitude also goes to the NSF for seeing fit to fund my project (NSF Award # 1251601).

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1: General Introduction

1.1 Spoken word duration.

A simple, but often-overlooked, fact about spoken words is that their duration varies rather a lot. Different speakers produce words slightly variably, and even a single speaker will not uniformly say the same word identically among various conversational circumstances. The age, the gender, and other sociolinguistic factors inherent to speakers account for some of this variability, but this project is concerned instead with cognitive psychological factors. It will focus on how variation in the normal processing of words, and thus the time it takes to plan and to articulate them, leads to corresponding variation in word duration. More specifically, it will focus on how the ease of accessing a word – lexical retrieval - and the ease of accessing and sequence sounds – phonological retrieval and sequencing – leads to variation in both the latency to begin an utterance, and the duration of the words in a multi-word utterance.

A variety of evidence shows that the duration of words varies as a function of what appear to be cognitive psychological factors. Repeated mention of a word, for example, tends to result in shorter duration (Fowler & Housum, 1987). This is frequently attributed to that word's status *as* previously-mentioned, a piece of information that the production system supposedly tracks as it produces the utterance. For similar reasons, a word that becomes available through indirect means in a conversation (e.g. talking about a playground, which makes swings and slides available) sometimes leads to shorter word durations (Gundel, Hedberg & Zacharksi, 1993; Prince, 1992). In this case, participants in such a conversation have a shared cognitive representation of the situation, and rely on this when they produce the word (Brennan & Clark, 1996; Clark, 1996; Gundel et al., 1993; Prince, 1992). In both of these cases, the speaker plans and executes a word differently depending on the context, and this leads to variation in duration.

Simply speaking faster or slower also leads to shorter word durations (Fosler-Lussier & Morgan, 1999; Munson, 2001). Studies that examine the character of databases of running speech, called corpus studies, find that speech rate varies both between speakers and conversational circumstances (Jurafsky, Bell, Gregory & Raymond, 2001). Some people speak faster than others, and some situations lead to faster speech, such as when a speaker is excited or in a hurry. Overall speech rate thus accounts for a significant amount of variation in word duration, and appears to arise from some internal state of the speaker; a personality factor, a state of mind, some environmental factor, or some combination of these.

Corpus studies have also uncovered more-abstract statistical relationships among linguistic variables that affect word duration. How predictable a word is in its surrounding context affects its duration (Bell, Brenier, Gregory, Girand, & Jurafsky, 2009; Jurafsky et al., 2001; *inter alia*). This is true of nouns and verbs (Jurafsky et al., 2001), as well as content and function words (Bell et al., 2009). Content words are what people typically think of as the words that convey the actual content of a spoken utterance – nouns, verbs, adjectives, and other words that carry most of the meaning. Function words, such as “the,” “of,” and “as,” link the content words together, and provide additional linguistic functionality. The ubiquity of the predictability effect, the fact that it appears on all types of words, suggests that it emerges from the structure of the language production system, and is a general property of the system’s status as a cognitive psychological entity.

All this evidence leads to the general hypothesis that variations in representation and/or processing of one sort or another lead either directly or indirectly to durational variation. That hypothesis, and five more-specific versions of it, guides this investigation. The general form is as follows: speakers' utterance plans and the duration of some of the words in their utterances should vary (although not necessarily covary) as a function of the ease of activating, retrieving, and ordering (and possibly others) the representations involved in producing the utterance. More specifically, the experiments below investigate three major subquestions, nested under the general hypothesis: 1) difficulty from resolving a semantic competitor should lead to variation in latency and/or duration; 2) ease or difficulty from resolving a phonological competitor should lead to variation in latency and/or duration; 3) reusing an identical stimulus should lead to variation in latency and/or duration, and this variation should be consistent with prior work on repetition; 4) the ease of retrieving and integrating a verb with earlier sentential material should influence not only its duration, but the preceding noun's duration as well; 5) the precise timing of stimuli that influence the ease each of these processes matters. Each is separate from the others, but part of the guiding question of how the production system specifies the duration of words. The experiments below test these questions using a paradigm which allows for the manipulation of the ease of completing or making use of each process simultaneously with the timing of the stimuli that purportedly modulate that ease. To preview the results, all five hypotheses find some confirmation. The mechanics of the language production system, specifically how the system treats lexical and phonological retrieval and representation, have an influence on latency to begin speaking as well as word duration, albeit only when presented at particular times.

The production system runs through several stages of processing in order to generate an utterance. These include discourse-level, semantic, phonological, and articulatory as the most common major divisions (Bock, 1982; Clark, 1996; Dell, 1986; Garrett, 1980; Goldrick & Rapp, 2002; Levelt, 1989; Levelt, Roelofs & Meyer, 1999; Schmidt, Meyer & Levelt, 1999). Each level of processing may speed or slow speech production in such a way that it influences or contributes to determining the duration of a word, and its neighbors. Previous investigations of the chronometrics of semantic and phonological processing (as well as speech errors) have given rise to computational models of language production (Dell, 1986; Garrett, 1980; Levelt et al., 1999). In particular, the Picture-Word Interference paradigm (PWI) has given rise to specific predictions and computational models of how the production system responds to stimuli that interact with (either by interfering with or facilitating) its normal operation (Levelt et al., 1999).

Although a latter section in this chapter will go into more detail, a few words about PWI will be helpful here. The task has participants name a target object (most often a simple line drawing of a common object) while ignoring various other stimuli. The stimuli interact with the normal process of naming, which allows for an examination of how manipulating the relationship between target and interfering stimulus leads to differences in outcome. The outcome variable is most frequently the latency to begin speaking, which is taken to reflect how long the speaker took to plan the utterance. Planning an utterance typically involves doing at least some basic semantic and/or phonological processing (see below), which means that PWI tasks measure ease of semantic and/or phonological processing. Unlike most PWI tasks, the experiments below also ask about word duration, using similar methods, in order to extend the findings on latency as well as compare with them. The theoretical interest of such a finding rests primarily on the speculation that word duration and utterance planning do not always go hand in hand, because they do not comprise the same set of processes.

Latency findings in PWI show that if the interfering stimulus exhibits a semantic relationship with the target (e.g. naming a pear while hearing "apple"), speakers take longer to begin speaking (Schriefers, Meyer & Levelt, 1990). This is taken to reflect the difficulty of choosing between the target and the interfering stimulus, either because of the difficulty in selecting the lexical item because of the closely-related conceptual representation(s) (Schriefers et al., 1990), or because of the difficulty involved in excluding irrelevant-but-similar lexical items from further processing (Caramazza, 1997). On the other hand, if the interfering stimulus exhibits a phonological relationship with the target (e.g. naming airplane while hearing "error"), speakers take less time to begin speaking (Schriefers et al., 1990). By the same logic, this reduction in latency reflects the ease of retrieving and/or integrating the phonology of the target with the already-available partial phonological information from the other stimulus.

An analogous argument could reasonably apply to word duration, namely that the ease of completing a variety of processes leads to increased or reduced duration. PWI studies claim that latency variation reflects the amount of time speakers take to formulate an utterance plan. For duration the argument would run variation reflects how speakers execute that plan, or possibly how they execute that plan in parallel with planning latter portions of the utterance. To date, no empirical evidence has established the verity of this analogous argument. PWI studies almost exclusively measure latency, and corpus studies of word duration give only a probabilistic demonstration of the relationships among variables. The experiments below measure not only latency, but also duration, in an attempt to show that some forms of variation in production processing lead to differences in both.

Like any cognitive mechanism, the precise timing of processing matters a good deal for language production. For example, semantic processing generally precedes phonological processing. This means that any stimulus that influences semantic processing must precede any stimulus that influences phonological processing, given reasonable assumptions about how cognitive processing works (i.e. that some basic processing must be completed before any processes that rely on the output of those basic processes can begin, and that processes terminate after some finite amount of time). The PWI studies sketched above found that semantic relationships between prime and target display different timecourses of influence than phonological relationships. Semantic primes' influence precedes phonological primes', which reflects theoretical constraints on the course of production processing (Levelt et al., 1999; Schriefers et al., 1990).

More importantly, the time window in which semantic primes influence latency has tended to be fairly narrow (Schriefers et al., 1990; Starreveld & La Heij, 1996). This suggests that in order to have an influence on the ease of lexical retrieval, for example, a prime must not only appear before one that would influence phonological retrieval, it must also appear within a certain time window. A prime that appears too early might have its influence fade before any relevant processing takes place, whereas one that appears too late might simply miss its only opportunity to have an influence. The experiments below control the timing of the primes with this consideration in mind. PWI standardly manipulates the timing of when the primes appear relative to the targets in order to gain insight into the timing of the processing components of the production system. The current work does the same, while both respecting prior findings on when latency effects appear, and widening the measurement range to account for possible longer-lasting effects on the planning and/or execution of word duration. In particular, it is possible that in addition to short-lived effects based on transient relationships among semantic or phonological representations, speakers may consciously and/or automatically preserve potentially relevant information for later processing. The manipulations in the studies below account for this possibility.

The goal of this project is to take a step toward understanding the production system and how it specifies duration, to examine the effects of variation in individual processes on how speakers plan and articulate multi-word utterances. The focusing question of the entire project is whether the difficulty associated with retrieving and planning words leads to durational variation, and if it does, when. The answers to the general question, as well as the specific questions about lexical and phonological retrieval and integration, will be used to further circumscribe existing mechanisms, or, if necessary, suggest alternatives. This project will take the first steps toward this goal in 6 chapters, four of them empirical, with four corresponding pairs of experiments (1a&b, 2a&b, etc.). This chapter will focus on the theoretical background and motivation, and introduce the paradigm. The second will focus on the relationship between lexical retrieval and word duration, and the third will follow up on this question by extending the focus to the utterance's verb. The fourth and fifth will complete the analogy for phonological retrieval and integration, where the fourth will look at the effect on the object word and the fifth will extend the focus once again to the verb. The sixth will summarize and conclude.

1.2 Word duration in the language production system

The importance of understanding why word duration varies – why the production system does not produce the same tokens of words time and time again – lies partly in understanding what duration indicates for comprehenders, and correspondingly what speakers use it to indicate. Word duration tends to mark parts of an utterance that contain information about the speaker's communicative intention (e.g. Ladd, 1996; Pierrehumbert & Hirschberg, 1998; Selkirk, 1995). Most basically, it signals what word the speaker actually produced. The difference between “thyme” and “time,” for example, lies in the duration of the middle vowel, and is a reliable (although subtle) indicator of the difference between the two (Gahl, 2008). Comprehenders can tell very early in an utterance whether a speaker is producing “ham” or “hamster” – before any disambiguating information is present in the speech sample (Dahan, Tanenhaus & Chambers, 2002; Brown, Salverda & Gunlogson, 2015). This suggests that comprehenders use duration to both make predictions about what the speaker is likely going to say, and to confidently settle on an interpretation of the speech signal.

Duration contributes to what is known as prosodic accenting. An accented word carries a pitch accent (Pierrehumbert & Hirschberg, 1990), but also tends to have longer duration and higher intensity than an unaccented word (Ladd, 1996). Accents mark information that a speaker wishes to emphasize for one reason or another (Gussenhoven, 1984; Schwartzchild, 1999; Selkirk, 1995). An accent distinguishes, for example, between *Did you have FUN?* and *Did YOU have fun?*, where in the first case the speaker wants to focus on the addressee’s emotional state, and in the second case instead on the addressee as opposed to some other person. Focus-marking like this changes the meaning of the utterance without changing any of the words or structure. Fundamentally, this means that acoustics plays a role in meaning. Comprehenders understand this function of durational emphasis, and modulate their understanding of a speaker’s intention quite rapidly (Dahan et al., 2002). A lengthened “camel,” for example, tends to mark the camel as new to the discourse, and the sheer fact of lengthening allows the comprehender to make that inference almost as soon as the information becomes available.

Similarly, duration can also index whether the speaker is having difficulty retrieving or otherwise processing the name for an entity, which in turn can influence how comprehenders interpret what the speaker intends to refer to (Arnold, Pancani & Rosa, under revision). When describing a scene filled with objects, for example, longer word duration (as well as disfluency) can signal that the speaker is going to refer to an entity that is not in discourse focus – either something new or something difficult to refer to. Comprehenders use this durational information to make educated (although probabilistic) guesses about the speaker's referential intent.

This line of evidence highlights the importance of duration for language comprehension, and thus also for production. Effective understanding of an utterance depends in part on making use of durational signals in the speech stream. Unfortunately, reliably interpreting these signals may not be as easy as simply following rules such as "if the word is longer than normal, interpret it as emphasizing the corresponding entity." Speakers do not seem to employ such rules, at least not infallibly. Even after controlling for a variety of prosodic factors such as sentence position and accenting, paralinguistic factors such as the conditional probability of one word following another affect word duration (Bell et al., 2009). Within an accent category, too, words exhibit acoustic variation, including in duration (Aylett & Turk, 2004; Watson, Arnold & Tanenhaus, 2008). Prosodic accents themselves do not always reliably correspond with the information categories they are thought to convey (i.e. an accented word does not always mark focused information, and an unaccented word does not always mark unfocused information; Sityaev, 2000), with accented words predominating irrespective of information status. Longer duration can mean any of the things mentioned above (novelty, emphasis, etc.), or it may simply mean that the next word in the utterance is infrequent, or otherwise difficult to process. Durational lengthening is not necessarily *meaningful,* in other words, and may simply reflect fluctuations in the mechanics of the language production system.

Comprehenders thus face a difficult problem. They must determine which acoustic variations constitute meaningful differences – the difference between "ba" and "da," for example

- and which they can safely relegate to "meaningless" processing artifacts – the difference between "pa" and "p-ha" (the latter has a longer plosive portion, where the speaker continues to exhale slightly between the consonant and the vowel), which is not a meaningful distinction in English, but is in Hindi. Ladd (1996), for example, distinguishes between intonational meaning, an intention-based acoustic factor, from paralinguistic prosody, which arises from unintentional and uninteresting deviations from normal processing. A speaker might say "p-ha" when speaking English because of noise in the processing system, but a comprehender would not treat this as a meaningful action, and would instead interpret the acoustics as "pa." However, although the difference between "pa" and "p-ha" does not carry meaning in English, the durational difference between "time" and "thyme" does. When it comes to acoustics, comprehenders cannot simply collapse across every example of acoustic variation. They must make informed decisions about which variations contain useful information about the speaker's intentions, and which they can discard as completely meaningless.

The decision is not as easy as it would be if all variation emerged from intention. Indeed, Arnold and Watson (2013) have argued that fully understanding both language production and comprehension will require understanding the relationship between message-based (i.e. intentional) and processing-based prosodic variation, an attitude the current project shares. If the speech signal contains both meaningful and paralinguistic variation, which it appears to, research that takes acoustics as an outcome variable cannot *a priori* discard an acoustic variant unless comprehenders reliably do as well. As a corollary, studying speech requires studying when speakers introduce systematic acoustic variation into the speech signal. Systematic variation implies a systematic sensitivity to some stimulus or processing operation. Understanding these relationships is a key part of understanding the architecture and function of the production system.

Several lines of evidence suggest that "emergent" processing phenomena, such as disfluency, utterance planning, and especially durational variation, have theoretical relevance to how language production proceeds. Disfluencies in running speech, such as "umm," "uhh," or stuttering and repetition, frequently indicate that the speaker is experiencing some particular type of processing difficulty. Fox Tree and Clark (2002), for example, found reliable differences between the circumstances where speakers used "uhh" and where they used "umm." "Uhh" most commonly indicated a short processing hiccup, which was quickly repaired or overcome. "Umm," by contrast, typically indicated a deeper difficulty, and was commonly followed by either restarting the utterance entirely, or making other significant changes to where the utterance plan originally led.

Comprehension experiments also show that disfluencies tell listeners when a speaker is having processing difficulty (Arnold & Tanenhaus, 2011). In a simple eye-tracking experiment, listeners had to click on an object in one of four corners of a computer screen according to spoken (pre-recorded) instructions from a speaker. Some of the objects were familiar, such as apples or trees, but some were complex, unfamiliar objects with no readily-available name. The experiment manipulated the fluency of the instructions, introducing natural-sounding disfluencies ("Click on the uhh…") before the part of the instruction that told the listener which object to click. After hearing a disfluency, listeners looked more frequently to the unfamiliar, hard-to-name object. This suggests that they had drawn an inference that the speaker was having trouble naming the object, and was using the disfluency to buy time to generate an appropriate referring expressing. If listeners had ignored this information, emerging as it did from the regular operation of the speaker's production system, they would have thrown out an important cue to the speaker's intention, despite its formal lack of meaning.

How speakers structure their utterances is another important cue to how production processes operate that does not necessarily convey meaning. A variety of evidence suggests that speakers modulate how much time they spend planning an utterance as a function of the complexity of its components (Bock, Irwin, Davidson & Levelt, 2003; Gleitman, January, Nappa & Trueswell, 2007; Konopka, 2012). Each of these studies also shows that processing factors, such as ease of access of the objects in a scene, or the availability of a syntactic structure from prior processing (e.g. priming), influences which entities the speaker selects to talk about first. More-available entities tend to appear earlier in the sentence, and primed syntactic forms enhance this effect. The choice of which entity to put first in an utterance is not necessarily meaningful, given that changing from passive to active (e.g., "the mailman was bitten by the dog" to "the dog bit the mailman") does not substantially change the message conveyed. However, in many cases the sentence structure reveals how the speaker processed the information that went into it. Like disfluency and the durational differences in the following section, processing variation can give rise to differences in the speech signal that do not themselves constitute meaningful alternatives, but which can be useful because of what they imply about how production proceeds.

1.3 Givenness and word duration

Durational variation may be another "emergent" property of the language production system, but the degree to which that hypothesis holds depends on the type of explanation offered to account for it. One theoretical distinction that lies at the core of the debate about meaningful and processing-based durational differences is the one between given and new information.

Prince (1992) defines new entities as those that have gone unmentioned or have not featured recently in the conversation, and defines given entities as those that have recently been mentioned or brought to the attention of the conversational participants. Speakers' representations of entities in a conversation are thought include information about their status as given or new (Levelt, 1989), which influences other processing variables such as word order (Bock, 1982) and decision to use a pronoun (Gundel et al., 1993). Given entities tend to appear earlier in an utterance than new entities (Arnold, Wasow, Langosco & Gingstrom, 2000; Chafe, 1994), and to appear more often as pronouns than as fully-realized lexical noun phrases (Gundel et al., 1993).

Givenness also appears to correlate with duration, although this is typically mediated by prosodic factors (Gussenhoven, 1984; Pierrehumbert & Hirschberg, 1990; Ladd, 1996). Referential expressions for given entities often receive deaccenting (Pierrehumbert & Hirschberg, 1990), and therefore shorter duration, than referential expressions for new entities (Brown, 1983; Fowler & Housum, 1987; Kahn & Arnold, 2012; Lam & Watson, 2010; Lieberman, 1963). For example, the duration of *pants* in *That man's pants appear to be missing in action* will most likely be shorter after uttering *Is he wearing pants?* than after *Did you see that?!*.

One message-based explanation for this phenomenon is that speakers' representations of givenness lead to shorter duration in order to manipulate the attention of their listener away from already-known information toward novel information. Information a listener already has access to can be hypoarticulated (i.e. deaccented), because she needs less acoustic specificity to keep that information active, or to reactivate it. Conversely, novel information can be hyperarticulated (i.e. accented), either in order to maximize the chance of the listener comprehending, or to provide emphasis (Lindblom, 1990). Both of these approaches rely on the concept of givenness, or something very similar. Under these theories, speakers – or, more accurately, the production system – make reference to an explicit representation of information status, and use that as the major or sole determiner of emphasis, and therefore duration.

An alternative explanation, based more narrowly on speaker-internal representations, could focus less on the listener, and more on simple repetition. Any entity a speaker has recently referred to should change the state of the system in such a way that subsequent processing of the same entity (i.e. during repetition) should be easier. This might result from more-available representational information, tighter connections among the necessary mechanisms and representations, closer coordination among the various processes, or any combination of these things. The upshot is a production system that can more easily reach an articulatory goal. Notably, this occurs without making explicit use of a representation of a listener, and instead relying on the regular operation of a cognitive system. Speakers do appear to represent the listener to varying degrees depending on the context, sometimes relying largely on their own internal states (Arnold, Kahn & Pancani, 2012; Bard & Aylett, 2004), and sometimes using an explicit representation to inform either lexical (Brennan & Clark, 1996) or acoustic variation (Galati & Brennan, 2010). Although an important theoretical question in its own right, the listener's place in determining word duration will take a background role in subsequent discussion, with the foreground granted to the speaker's processing. Current theorizing suggests that the most fruitful approach to the question of how speakers deal with listeners lies in detailing when they do and do not appear to consider the listener, as opposed to making categorical statements (Brown-Schmidt, 2012).

Earlier work of mine provided evidence of a relationship among givenness, facilitation of the production system, and spoken word duration, and more specifically that processing facilitation influences word duration over and above repeated mention (Kahn & Arnold, 2012). The experiments tested the hypothesis that general processing fluency in the production system increases after recently producing identical material (a hypothesis that the current work also tests), with the behavioral consequence of reduced latency to speak and/or duration of spoken words. In particular, the hypothesis was that priming information which differentially activated different levels of processing should also differentially elicit duration variation.

An instruction-giving task asked speakers to give simple descriptions of objects moving on a computer screen (e.g. “the elephant shrinks”). Before issuing instructions, speakers would either hear the names of the target objects spoken out loud (linguistic priming), see the target objects flash (non-linguistic priming), or receive no priming stimulus. The non-linguistic stimuli were intended to activate only high-level representations, such as conceptual organization of the scene and some semantic information (the fact that the elephant, the airplane, and the balloon flashed), whereas the linguistic stimuli were intended to activate those representations as well as lower-level information associated with phonology (“elephant,” “airplane,” “balloon”). The difference between the linguistic and non-linguistic prime lies primarily in which representations in the production system they should target. Both should activate conceptual representations, but the linguistic condition should also activate semantic and phonological representations directly.

The dependent variable was the duration of the individual word segments of the instructions, and particularly the duration of the target word (e.g. *airplane*). Relative to the control condition, both the linguistic and non-linguistic conditions elicited shorter latencies to begin speaking, suggesting that the primes affected the planning of the utterance. The explanation for this effect is that the prime facilitated the formulation of a conceptual representation. The non-linguistic condition also elicited shorter durations than the control condition on the target word, but more importantly, the linguistic condition elicited even more reduction than the non-linguistic condition, a three-way difference. The most likely explanation for the durational reduction thus resides in the differential effects the primes had on the processing system. The non-linguistic condition sped planning and reduced target word durations relative to the control condition, indicating that speakers had an easier time generating a plan and executing it. The linguistic condition elicited additional duration on the target object word, as well as shorter durations on the determiner and the action word, suggesting an overall processing fluency effect. Importantly, the reduced word durations emerged without a corresponding difference in planning time between the linguistic and non-linguistic conditions, suggesting that the plan did not make any contribution, but rather that later processing was responsible.

Arguably, the primes in Kahn and Arnold (2012) did not constitute a sufficiently rich discourse context to truly establish them as given entities. The name of an object without any additional contextual information may or may not elicit a representation of the object as a given entity, in other words. It does not connect with prior conversation (Chafe, 1994; Clark, 1996), nor does it establish what the speaker is likely thinking about or will continue to talk about (Ariel, 2004; Gundel et al, 1993; Prince, 1992). Even if speakers were not treating the primes as having established an entity as given in a "proper" theoretical sense, the question of why they reduced the object word durations remains. In that paper, we argued that regardless of the primes' discourse status (i.e. whether they were given or not), they still affected processing. A more useful term for primes like that is repeated mention, which does not carry the connotation of a rich discourse representation of givenness, but still retains many of the processing-relevant properties that demonstrably affect speech production. Below, “repeated mention” will be used instead of “givenness,” in order to avoid any objection that the primes do not truly make the entities given.

In the context of repeated mention, the central question for this project is the degree to which ease of processing contributes independently of any information status to variation in word duration. Theories and models of production both leave open the question of how exactly duration emerges from the individual processes in the system, and indeed whether ease of processing contributes to duration variation at all. Similar to Kahn and Arnold (2012), the experiments below will attempt to hold prosody and information status constant while manipulating the difficulty of lexical and phonological retrieval. This will help to circumscribe the role that production processes play in determining duration, which will in turn inform theoretical and computational models of production. The focus on semantics and phonology, and on their role in multi-word utterances, is intended to give both broad and deep answers to questions about how the language production system calculates and produces duration as a function of its recent internal states, as well as its anticipatory processing of upcoming material.

Ultimately, any durational variation will have to be explained in terms of the operation of the language production system, in one way or another. This includes both meaning- and processing-based variation. The most reasonable approach allows each to have their place, and the goal of the current experiments is not to provide an alternative to givenness, but rather to focus on how and when processing influences duration. The next section provides some relevant information about current theories of the organization of the language production system, thereby providing a more-detailed understanding of exactly what is meant by “processing” for the remainder of this discussion. Understanding current thinking on this organization is a necessary part of the current study, relying as it does on theoretical assumptions about how semantic and phonological processing operate.

1.4 Models of language production

Models of language production make a distinction between discourse, semantic, phonological, and articulatory types of processing, in this order. Some models include slightly more or fewer levels, but this list is fairly typical. Many production models take the form of information-processing diagrams, and account for a range of empirical evidence in domains such as speech errors (Garrett, 1980), information status (Bock, 1982), and lower-level processes such as coarticulation and resyllabification (the blending of two nearby syllables, as in *an apple*; Levelt, 1989). More computationally-oriented models (e.g. Dell, 1986; Levelt, Roelofs & Meyer, 1999) make more specific claims about the underlying mechanisms of production, but frequently gain this specificity at the cost of scope. Current computational models (Dell, 1986; Goldrick & Rapp, 2002; Levelt et al., 1999; Starreveld & La Heij, 1996) focus on the production of single words, while non-computational models deal with whole sentences as well as words (Bock, 1982; Garrett, 1980; Levelt, 1989). A more detailed discussion of multi-word processing appears later in the chapter.

Production models generally agree that processing begins with some form of semantic or conceptual representation. Semantic nodes in these models become active to a greater or lesser degree depending on how well they match the speaker's intended concept. A speaker's wanting to talk about a rat, for example, might activate nodes for RESEARCH ANIMAL, CHEESE, and WHISKERS. It might also activate higher-order or lower-order nodes, such as ANIMAL or DUMBO RAT, or other members of similar categories, such as MOUSE, SQUIRREL, or CAT. Part of planning to speak involves a speaker selecting from among the available entities those which best suit the message she wants to convey (Levelt, 1989). Often this involves choosing among alternatives, such as the difference between “cop” and “police.” It may also include a process of deciding whether to include an entity at all – mentioning the murder weapon or not, for example (Brown & Dell, 1987). Still further, she may want to be very precise about her word choice, or may decide to remain vague, for example saying “accomplished” instead of “did.” All of these steps may form part of the message planning portion of language production processing. Some may not occur for every utterance, but the majority of speech acts involve narrowing down the range of possible entities, selecting an appropriate specificity, and finally settling on an appropriate word (Levelt, 1989). Computational models have mechanisms that specify a certain period of time (Dell, 1986) or a requisite level of activation relative to other semantic nodes (Levelt et al., 1999) for selecting a word or words from a set of activated semantic nodes.

Semantic processing comprises the several subprocesses detailed in the paragraph above, including assessing the range of possible entities, accessing the corresponding semantic information, and choosing the lexical entry most appropriate for that information and the intended message, and mental processing takes time. Factors that make this processing easier, such as having recently selected the entity, may reduce the time needed to run from message to utterance. Conversely, factors that make processing more difficult, such as having to select among similar alternatives, may increase the amount of time it takes to reach the actual speech act. The number of plausible competitors for a target concept (e.g. semantically-related words) is sometimes taken as a predictor of the time course of selection (Levelt et al., 1999). The opposite also holds – activated supporting concepts (such as membership in a category) predict faster lexical selection (Glaser & Düngelhoff, 1984). There is currently no evidence for or against the hypothesis that ease of semantic processing has any effect on word duration, however.

After semantic processing comes phonological. Many models of language production assume that the lexical information associated with selecting a particular word (e.g. “caldera”) from a semantic set contains a link to the phonological form for that word (Levelt, 1989). The phonological form includes information about the actual phonemes (the smallest units of meaningful sound), their ordering (“cat” has the same phonemes as “tack,” just in a different order), how they are organized into syllables, their prosodic emphasis (i.e. which part of the word gets emphasized, as in the difference between “CONsole” and “conSOLE”), and possibly some basic information about how to combine with other words. Once the lexical item has been selected, as long as nothing goes wrong, all of this phonological information becomes accessible. Each part is important for delivering a fluent utterance.

Once the word form becomes active, all of the information combines to form a set of articulatory routines (Levelt & Wheeldon, 1994), which then feed out to the articulatory motor system. Generally speaking, syllable structure emerges at this level. This includes both standalone as well as mixed phonemic information both within and between words. What this means is that before delivering phonological information to the articulatory stage of processing, the production system sometimes reorganizes the structure of the phonemes in the utterance. In some cases this means simply modifying how the articulatory system should pronounce a vowel or a consonant given the need to pronounce other nearby vowels or consonants. This is called coarticulation, and like word lengthening, happens more frequently than most speakers realize. A similar phenomenon, termed “resyllabification,” involves the movement and recombination of phonemes in order to respect articulatory constraints and prosodic boundaries. Individual phonemes and collections of phonemes sometimes continue to stand alone in their standard form, but other times come together in fairly standard alternative ways, for example in the difference between “an apple” and *a napple*. Both of these phonemic movement processes happen during phonological processing, and constitute a significant portion of the processing that occurs there.

The field’s current understanding of how language production works comes in large part from catalogs and analyses of speech errors. Both computational and non-computational production models make use of speech error data (e.g. Dell 1986, Garrett, 1980), because they license inferences about the organization of the system. Error distribution patterns suggest that semantic information is processed separately from phonological. Semantic errors, for example, occur primarily within syntactic categories – nouns exchange with nouns, verbs exchange with verbs. Phonological errors, by contrast, cross syntactic categories but tend to respect typical phonological boundaries – word onsets exchange with word onsets. A similar set of “errors,” tip-of-the-tongue states, suggests that semantic processing comes before phonological. In such a state, the speaker knows the semantics of a word, but cannot access its phonological form.

Recently, computational models have also begun to incorporate chronometric data (Levelt et al., 1999). Similar to error-based model-construction, this process combines hypotheses about how the timing of production should work with empirical data about how it actually does work (Roelofs, 2002). The separation of semantic and phonological processing stages, for example, is partly based on the error data outlined above. It is also partly based on the timing of when certain types of primes, either semantic or phonological, appear to interfere with the normal processing of a linguistic stimulus. The experimental paradigm that gives rise to this set of inferences is called the Picture-Word Interference (PWI) paradigm, and is the inspiration for the experiments in this project. It allows for fairly direct manipulation of the ease of selected types of production processing, which permits the type of inference generation used to construct the Levelt et al. model. PWI will be discussed in more detail below.

Semantic processing occurring earlier and separately from phonology means that any effects on fluency must pass either through or around phonological processing. Increased fluency in semantic processing may, in other words, have downstream effects on phonological processing. Similarly, increased fluency in phonological processing, such as pre-activated phonological representations, may feed backward to have an influence on semantic processing. The degree to which the stages of processing in production are interactive with each other is hotly debated (see Goldrick and Rapp, (2002), for a recent review). Feedback and discreteness of information transfer between levels in particular are still open questions. Although these questions are relevant to the interpretation of the data presented below, the experiments do not test between them, and examine the influence of each processing level separately.

This section outlined contemporary theorizing about the production of single words. Success for these theories is typically measured in terms of accuracy, namely whether the model produced an error and of what type, as well as latency, namely how long the model took to produce the correct word or some other. Relying on error and chronometric data allows these models to approximate many of the accuracy and processing difficulties encountered during regular speech, such as perseverations and frequency effects. To date, however, models do not make specific predictions about word duration, even at the single-word level. This is by design, given that existing models focus on explaining speech error patterns and speech onset latencies.

As stated above, the long-term goal of this project is to supply empirical evidence for the relationship between ease of processing and the duration of words in running speech, so that models may begin to use duration as an outcome measure. In order draw on the same theoretical background and empirical evidence, the manipulations below will use a slightly-modified PWI paradigm. They will also expand the range of outcomes-of-interest to include word duration.

This will further increase compatibility between this and prior work, because it will allow for a simultaneous comparison of how the ease of semantic and phonological processing affect latency and how they affect word duration, while still allowing me to ask novel questions.

Like the models described in this section, typical PWI focuses on single-word production. Most speech, however, consists of many words in rapid succession – a standard conversational utterance. An effective investigation of word duration, therefore, should take place in the context of multi-word utterances. Including multiple words in the outcome measurements raises other questions about how speakers produce fluent running speech.

1.5 Multi-word utterances, planning, and incrementality

Making a useful conversational contribution most often entails uttering multiple words in a row. Anything more than a simple affirmative or interjection requires simultaneously planning and then linearly executing several words in order to maintain fluency, and thus hold the conversational floor (Clark, 1996). To accomplish this, speakers could mentally plan out an entire utterance before beginning to speak. In fact, they do seem to plan up to a syntactic or prosodic phrase in advance (e.g. *the green apple rotates*; Clark & Wasow, 1998; Jescheniak, Schriefers & Hantsch, 2003, Schnur, Costa & Caramazza, 2006). However, speech does not consist of long pauses punctuated by extremely fluent delivery, even in short sentences. Rather, speakers tend to alternate between relatively short pauses and relatively fluent delivery (Beattie, 1983), which suggests that at least some essential processing occurs simultaneously with the actual articulation of the utterance.

Multi-word utterance production involves more than simply appending unrelated chains of individual words' processes. The major difficulty in moving beyond the production of a single word lies in striking a balance between the processing of upcoming words with the processing of the current word. This is no simple task, requiring the speaker to make decisions about syntactic form, word order, lexical items, and prosody, among other things (Levelt, 1989). More goes into creating an utterance than planning and executing a series of unrelated words, like beads on a string. Theoretical and computational models account for this complexity by allowing, and in fact requiring, parallel and simultaneous processing of several different pieces of information at once (Bock, 1982; Dell, Oppenheim & Kittredge, 2008; Eberhard, Cutting & Bock, 2005). For example, information status interacts with the ease of accessing a lexical entry to influence word order choices in some models (Bock, 1982). The important question in the context of word duration is the extent to which the processing of one word's duration influences the processing of the surrounding words' durations. The major question guiding the discussion in the previous sections was whether the fluency of semantic and phonological processing of a single word influences that word's duration. The major question of this section is whether the duration of some target word varies as a function of the fluency of processing the words that surround it.

Theoretically speaking, planning and articulation of some words in a sentence necessarily bears on the planning and articulation of other words, if for no other reason than syntactic agreement. In English, any nouns that correspond with a particular verb in a sentence must agree with that verb in number and tense (e.g. *the boy plays*, *the boys played*). Other languages exhibit similar agreement requirements, both within and between syntactic categories (Corbett, 2006). Sometimes this creates what are called long-distance dependencies, which refers to situations where many otherwise irrelevant words intervene between two words that agree grammatically (e.g. *the fox who was quick and brown jumps*, where *fox* must agree with *jumps*). These are quintessential cases of linguistic phenomena that require planning, either in the form of pre-planning (noting the number of *fox* and marking *jumps* then and there with the appropriate inflection) or pro-active long-term retention (noting the number of fox and relying on memory to ensure that *jumps* gets marked at the time it gets produced). In either case, the system formulates a plan for grammatical constraints, and then proceeds with the rest of the utterance (Eberhard et al., 2005). Any effect of syntactic planning should thus also remain constant.

Multi-word utterance processing necessarily involves interactions between words not only from syntactic dependency, but also because at least some semantic and phonological processing must occur in parallel, and often does. In order to plan a full phrase, speakers must activate each lemma in the phrase, and do so in the proper order. To achieve full syllabification of an utterance, they must also activate at least some phonological information. In a short utterance like *the airplane rotates*, a speaker might plan the entire utterance before beginning to speak, which would entail activating lemmas for THE, AIRPLANE, and ROTATE, as well as the corresponding phonemes. Within a word, too, especially a long one, speakers must access each phoneme, and make sure that they appear in order in the utterance, or they risk confusing *tack* with *cat*.

The possibility of planning an entire utterance in advance raises questions about how much planning speakers actually do in normal situations. One possibility is that they plan an entire utterance, perhaps as much as a whole paragraph, before uttering a single word. In fact, periods of fluent delivery are often punctuated by periods of planning (Beattie, 1983), suggesting that although speakers do not plan entire paragraphs, they do a significant amount of preprocessing. Frequently, speakers plan approximately a single clause in advance (Allum & Wheeldon, 2007; Smith & Wheeldon, 1999), with much of the planning tied to processing the verb (Ferriera, 1993; Schnur et al., 2006). When talking about scenes, they can quickly extract the gist and begin planning a syntactic structure accordingly (Bock et al., 2003; Griffin & Bock, 1998). Even phonological information that becomes available early reduces planning time (Abdel Rahman & Melinger, 2008; Jescheniak et al., 2003; Schnur et al., 2006), suggesting that planning can extend to all levels of production.

On the other hand, processing and planning during production also sometimes proceed incrementally. Speakers need not formulate an entire paragraph before beginning an utterance, and can in that case offload some processing until after they begin speaking. In simple two-picture scenes, for example, onset to speak latencies for the first-named object are unaffected by the difficulty of accessing the second-named object (Meyer, Sleiderink & Levelt, 1998). This suggests that not only do speakers begin an utterance before accessing all the necessary information, they simultaneously plan and execute the latter portion. In other simple scenes, such as ones depicting objects that differ only in color, speakers include color word modifiers only when they have noticed the ambiguity in time (Brown-Schmidt & Tanenhaus, 2008). For example, in the context of a scene with two butterflies, one blue and one orange, *the butterfly* is an ambiguous reference. When giving an instruction to a listener, however, speakers will include the color modifier only if they fixate on both butterflies sufficiently before they begin speaking. Otherwise, they simply use *the butterfly*, and repair the utterance post-hoc, as in *the butter- uh, the blue butterfly*. Speakers appear capable of quickly modifying utterances, and of incorporating relevant information into their plans as necessary. The speed with which they do so suggests that the behavioral output, including factors like word duration, may vary in a time-locked fashion according to the ease of planning or executing individual processing steps.

Other studies show that planning time is affected by the difficulty of accessing information about a scene. In a task where they have to name two pictured objects while minimizing pauses between the two names, participants begin speaking sooner if they see a picture of an object with a two-syllable name than an object with a one-syllable name (Griffin, 2003). This suggests that they intentionally begin speaking early in the two-syllable case under that assumption that while speaking the first word they will finish planning the second. More direct evidence of this effect on duration comes from a task where participants named two objects where the second had its frequency manipulated (Christodoulou, 2012). Low-frequency second objects elicited longer durations than high-frequency objects, even after controlling for frequency. More importantly, the latency to begin speaking did not differ between conditions, suggesting that in both cases speakers planned to the same degree and modulated the duration of executing articulation. Once again, the evidence points toward a modulation of duration based on processing fluency.

The incrementality and planning/execution data lead to several unanswered questions about how ease of processing influences word duration. All of these questions follow from the idea that speakers can partly begin or continue processing one component of an utterance (a lexical entry, some phonological information, a word's place in the ongoing syntactic and prosodic form, etc.) while completing processing of some other component. For example, does the ease of processing word *n+1* influence duration the duration of word *n*, similar to its influence on latency in Schnur et al., (2006), and duration in Christodoulou (2012)? It might if speakers process word *n* different depending on the production system's assessment of the difficulty of word *n+1.* The same applies in reverse – does processing word *n* influence processing on word *n+1* in a way that leads to correlated durational variation? Are there certain time windows during processing in which it is possible to influence word duration in these ways? The experiments below test these questions, by extending the PWI paradigm to multi-word utterances, and manipulating the ease of processing each of the content words in those utterances.

Up to this point, the discussion has centered around the hypothesis that ease and fluency of processing might have an influence on word duration. The next two sections will instead discuss the experimental evidence that gives reason to think that this might be the case. The following section will focus on word duration in particular, while the one after that will go into detail about PWI, including the evidence for a processing effect on latency to begin speaking. The final section will highlight some questions that the other sections leave unanswered, and show how the current study attempts to answer them.

1.6 Ease and fluency of processing

Ample evidence shows that word duration varies in a fine-grained way in response to a wide range of processing factors. Perhaps the most well-known and robust predictor of word duration in this literature is simply word frequency. Low-frequency words tend to be produced with longer duration than high-frequency words (Zipf, 1929), across a wide range of word types and categories (Bell et al., 2009). In fact, frequency controls significant variance in nearly every study mentioned in this section, and more of it than nearly any other factor (Jaeger, 2010). Conceptually, many of the other measures rely on frequency, as well, such as transitional probably, which is the frequency of word *w+1* following word *w*. The locus of the effect of word frequency in the production system is debated, with some claims that it has its primary effect on phonological form (Levelt et al., 1999), and competing evidence that lexical frequency also plays a role (Gahl, 2008). Regardless, frequency effects provide evidence that fluency of processing leads to chronometric variation, both on latency and duration.

Most demonstrations of the effect of frequency on word duration take the form of corpus analyses that examine the probabilistic distribution of language, both spoken and written. Several different measures of probabilistic relationships between words show that words which precede high-frequency words generally have shorter durations than those which precede low-frequency words (Bell et al., 2009). The conditional probability of one word given some prior word, for example, is one such measure. This suggests that the production system either tracks these relationships in a vast statistical space, or has some other mechanism for modulating duration based on fluency of processing. Bell et al. propose a mechanism at the interface between semantic and phonological processing to explain the effect of surrounding words on word duration. When ongoing processing is delayed by extended word retrieval, as in the case of a low-frequency word, or having to activate a normally-unrelated word, the production system delays by extending the duration of the current word in order to avoid slipping into disfluency. Word lengthening, on this account, indexes the degree to which the system experiences processing difficulty, and then smoothes word duration to accommodate fluctuations.

A similar proposal holds that prosodic variables such as phrase-final lengthening and accent placement explain most of the variance in word duration, even after accounting for frequency effects (Aylett & Turk, 2004). Unlike the Bell et al., (2009) study, Aylett and Turk represented prosody explicitly in their model alongside frequency effects. They found that prosodic variables did indeed explain a significant portion of the variance, but that frequency also contributed significantly. This suggests that the fluency of processing plays a role in word duration independently of its contribution to prosody (cf. Arnold & Watson, 2012).

Other frequency-as-fluency effects pervade the production system. The size of a word's neighborhood affects its duration. Neighbors are typically defined as words that can be formed by changing, adding, or deleting phonemes from some target word. *Cat*'s neighbors, for example, include words like *bat*, *scat*, and *at*. Words with many high-frequency neighbors also tend to have shorter durations than those with few, or low-frequency neighbors (Gahl, Yao & Johnson, 2012). Low-frequency words with high-frequency homophones also tend to have shorter pronunciation durations than those without (Bybee, 1999). Words with low phonotactic probability (the transitional probability between two or more phonemes) tend to have longer durations than those with high probability (Vitevich, Armbrüster & Chu, 2004). Syntactic regularity (i.e. subcategorization frequency) affects the duration of words that appear at syntactic junctures (Gahl & Garnsey, 2004). All of these effects suggest that language production is sensitive to fluctuations in processing fluency, and in fact tracks several types of probabilistic relationships when determining exactly how long a duration to give a particular word. Generally speaking, easy processing leads to shorter word durations, and more difficult processing leads to longer word durations. This could happen for several reasons. Most directly, a speaker may elongate or shorten a word in order to promote or maintain speaking fluency. In a multi-word utterance, some words may be more difficult to process than others, which means that in order to speak fluently, speakers may have to manage the precise duration of each word. One possible explanation is that more difficult-to-process words may lead to increased duration on either themselves or earlier words, depending on how the system manages duration. Although an explanation like this has been offered to explain variation in spoken word duration (Bell et al., 2009), the exact nature of the relationship between ease of processing, multi-word utterances, and word duration remains an open empirical question.

Another possible explanation for the putative relationship between ease of processing and word duration lies in the production system's assessment of the ease of processing itself. The system may know approximately how long semantic processing should take for a given word, a simple quantity or range of quantities to track. When selecting some semantic entity, the production system could easily compare the expected difficulty of processing to the observed, and send this information along as a signal to latter sub-processes that it found a disparity. A lower observed value than expected might lead the system to say, “Hurry up, we’ve got this covered!” Conversely, a higher-than-expected value might cause the system to throw the brakes, slowing things down to prevent disfluency. Simple comparators like this could exist throughout the language production system (and cognition generally), micromanaging subsequent processes in order to achieve a desired level of fluency.

Both the explanation in terms of coordinating between words and the explanation in terms of processing difficulty on a particular word rely on the assumption that speakers desire fluency and work to maintain it. Shannon (1948) characterizes this process as the smoothing of the mathematically-defined amount of information transmission over a noisy channel. The information a word contributes to a conversation is measured in terms of its probability of appearing in a context. In a noisy information channel like speech, maintaining uniform information density over time maximizes communicative efficiency. Speakers do appear to modulate their utterances based on information density, even over and above the other probabilistic measures mentioned so far (Jaeger, 2010). This includes syntactic variation (Jaeger, 2006), use of contractions (Levy & Jaeger, 2007), and even overall sentence complexity (Genzel & Charniak, 2002). The noisy channel is one of two things: either the processing from message to speech, or the interaction between speech and the listener(s). In both cases, it makes sense to characterize duration variation in terms of ensuring a balance between least effort and reliable transmission. Modulating word duration, in addition to engaging in some non-negligible amount of pre-planning, is a speaker's primary means of remaining fluent.

All of the evidence reviewed here points to ease of processing having some influence on word duration. However, the majority of these studies make use of probabilistic methods, without any direct manipulation of variables. This allows for a more targeted look at which variables (e.g. neighborhood density or mutual information) have an influence, but is less conclusive than experimental methods, because a causal relationship is more difficult to establish. The experimental methods employed below take a step further toward establishing a causal relationship between sub-processes like lexical or phonological retrieval and variation in latency and duration. They ask similar questions about the relationships among the variables, but introduce more-direct manipulation. This represents what I consider the most productive union between statistical/corpus-based methods and experimental methods, with one supporting the other.

1.7 The Picture-Word Interference Paradigm

An established method for investigating the time course and ease of processing an utterance is the Picture-Word Interference (PWI) paradigm. The current studies use a variant of this method, which not only permits investigation of fluency questions in general, but also provides a theoretical and empirical baseline for comparison with previous general patterns. The paradigm generalizes on the Stroop paradigm, in which a participant names a patch of color while ignoring a distractor (Glaser & Glaser, 1989). In a PWI task, a participant sits in front of a computer screen and names simple picture targets while ignoring a written or spoken distractor word (e.g. Schriefers et al., 2000). This additional stimulus, typically referred to as the interfering stimulus (IS), is used to influence how easily the speaker can process the target. Note that "interfering stimulus" is a slight misnomer, because it may also facilitate processing, but I will continue to use the label to remain consistent with PWI usage. The two major independent variables in these studies are the information relationship between the IS and the picture (i.e. related, unrelated, identical, control), and the SOA (stimulus onset asynchrony – the time between the presentation of a stimulus and a prime).

Manipulating the information relationship permits a comparison among facilitatory, inhibitory, and neutral effects of the IS. This allows for a general investigation of the range of fluency-of-processing effects. Facilitation appears as decreased latency to name the picture, and inhibition as longer latency. Neutral or control stimuli frequently take the form of silence in spoken distractor studies, or a row of upper-case X's in written distractor studies, in order to establish a baseline of naming latency in the presence of an otherwise informationally-identical prime. Unsurprisingly, an IS that appears as the name of the target ("airplane" for a picture of an airplane) elicits speeded processing, and thus shorter latencies to begin speaking (Glaser & Glaser, 1989). However, the degree to which this exceeds the facilitation from other stimuli varies according to the task. Unrelated-but-still-linguistic interfering stimuli (such as "plant" for an airplane) tend to elicit longer latencies than a null or control condition. This is taken to stem from the additional interference brought on by having to ignore a potentially-relevant linguistic stimulus.

A simple match in overall stimulus quality or category results in longer latencies to begin speaking, but the most telling relationship between interfering and target stimulus is one where the IS actually has some processing-relevant property in common with the target. Suppose that the speaker's target on a trial is a picture of an apple. If the interfering stimulus is "orange," not only is there an overall match insofar as both stimuli are linguistic, apples and oranges have a semantic relationship as well. This makes an orange a better candidate for selection than an interfering stimulus like “airplane,” which means the production system has slightly more difficulty selecting the target than in the case where the interfering stimulus has nothing with the target except the sheer fact of being a word. Semantic relatedness between interfering and target stimulus creates even more processing difficulty than mere wordhood, and thus increases latencies even further. The opposite tends to be true for phonological relatedness. Hearing “error” just before seeing an airplane target makes naming the airplane slightly easier. The initial phonemes are the same in “error” and “airplane,” so the production system has activated that phoneme by the time it begins to run through the processes necessary to name the target. Naming the target is easier, and latencies tend to go down in this situation.

Manipulating SOA permits a mapping of the time course of the effects listed above. In principle, external stimuli may influence some processes within a certain time window of their execution (i.e. sometime slightly before and/or during the actual process). Varying SOA either within or between subjects can reveal the time course of processing, or at least place constraints on when it can be influenced. Computational models of language production make use of this property of the PWI paradigm, and introduce or remove temporal processing constraints appropriately (Levelt et al., 1999; Starreveld & La Heij, 1996).

In particular, Levelt et al. (1999) draw on a temporal distinction between the effects of semantic and phonological primes reported in Schriefers et al., (1990). Their manipulation made use of an auditory prime, and varied SOA from -150ms to 150ms in 150ms increments. They observed semantic inhibition at -150ms, reflected in longer latencies to begin speaking, although not longer word durations. A summary of these results, as well as others outlined below, is shown in Figure 1. Inhibition did not appear at the other SOAs, nor were there durational effects.

This pattern suggests an early lexical competition effect that tapers off or disappears after lexical selection occurs. In the model, semantic activation accrues to the target word from seeing the picture stimulus. In a control condition, with no competing semantic information, lexical selection proceeds as normal. In a condition with a related distractor, however, activation also accrues to sematically-related conceptual representations. The system must decide between the two, and this decision takes time. Corroborating evidence from similar tasks shows that his effect is not readily explained in terms of conceptual conflict or visual confusability (Damian, Vigliocco & Levelt, 2001), suggesting that the effect is localized within the language production system.

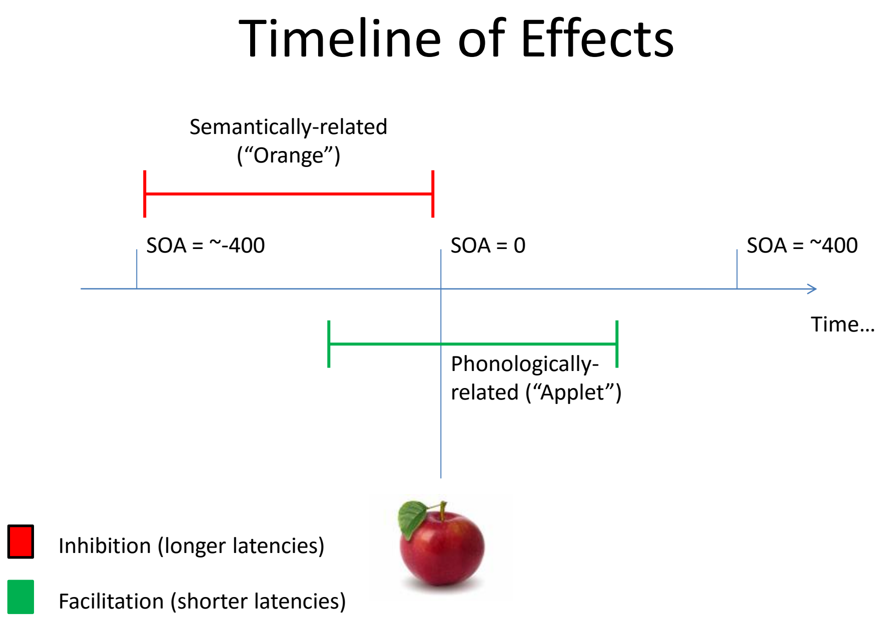

Figure 1: A summary of Picture Word Interference effects, with Schriefers et al., (1990) and Starrveled & La Heij (1996) used as examples. A picture of an apple appears at SOA = 0, with semantic inhibitions effects ranging from approximately -400SOA to 0SOA, and phonological facilitation effects ranging from approximately -200SOA to 200SOA.

By contrast to the timing of semantic inhibition, Schriefers et al. (1990) also observed a phonological facilitation effect at 150ms, but not at earlier SOAs. Phonological facilitation appeared as decreased latencies to speak, although the effect did not also carry over to word duration, similar to the effect of semantic inhibition. In the model, selection of the picture's lemma begins to activate the phonemes associated with the name. The appearance of the auditory stimulus 150ms later appears to coincide with this activation, and produces activation of its own on some or all of the target's phonemes. The additional activation speeds the utterance plan, and speakers begin speaking earlier. The later appearance of this effect relative to the semantic inhibition effect appears to reflect the later processing of phonological information, an assumption of essentially all models of production. Levelt et al (1999) interpret the lack of overlap between phonological facilitation and semantic inhibition as evidence for a discrete, feedforward network, with little-to-no interaction between semantic and phonological processing.

Not all PWI studies report the same pattern of results as Schriefers et al., however. Starreveld and La Heij (1996), for example, make use of written primes instead of auditory ones, and vary semantic and orthographic relatedness. Semantic inhibition appears once again at small negative SOAs. Orthographic facilitation, akin to phonological facilitation, spans a longer timeframe, from -200ms to 100ms, in 100ms increments. They attribute this long-lasting facilitatory effect to the ease of selecting phonological representations, brought on by the prime. Importantly, they also include an orthographically-and-semantically related condition (e.g. *cat* and *rat*), which induced latencies between the interference brought on by semantic inhibition and the facilitation brought on by orthographic facilitation. This they take as evidence that semantic and phonological processing interact, contrary to the assumptions and argumentation in Levelt et al. (1999). The findings are not directly comparable, given the difference in prime type, but any production model would have to account for both. Starreveld and La Heij offer a production model that localizes competition at word form selection (i.e. the word's phonological information) instead of at the lexical level. This simultaneously explains both semantic interference and orthographic facilitation. It could also explain phonological facilitation by a similar argument.

Still other studies report semantic facilitation instead of semantic inhibition. Many PWI experiments make use of semantic relations most accurately described as fellow category members. The word *cat* interferes with processing a picture of a dog because members of the same category compete readily with each other for lexical selection. A subset of experiments instead use semantic associates, such as *honeycomb* for a picture of a bee. Although still related to each other, researchers argue that the relationship is indirect, or mediated by a conceptual category that is less readily apparent to the participants (Belke, Meyer & Damian, 2005). Making the category apparent eliminates the behavioral difference between category members and associates (Abdel Rahman & Melinger, 2009), suggesting a conceptual mediation that leads to lexical competition. The difference between the behavioral effects for these two types of stimuli lies primarily in how a particular experiment emphasizes the conceptual relationship(s).

Stroop-like experiments, including PWI, rarely report word duration, because they focus primarily on the processes involved in planning. Notable exceptions to this trend include Experiment 2 in Schriefers and Teruel (1999) and Schnur et al., (2006), both of which report null effects of their priming manipulations on duration. To my knowledge, the only study to report an effect of semantic relatedness on word duration is not a standard PWI study, in that it makes use of both written targets and written primes (Balota, Boland & Shields, 1989). Although the interference does not take the form of picture-word, the manipulation and questions are very similar. Participants read a target word aloud, either before or after seeing a printed prime word. When the prime word came before the target, participants were told to ignore the prime and simply read the target quickly and accurately. Another set of participants saw the opposite order, with target before prime. In that case, they were told to read the target word aloud only after the prime word had appeared, once again ignoring the prime itself except insofar as it cued the response. Across two experiments, participants thus saw both negative (prime before target) and positive (target before prime) SOAs, and in greater ranges than normally appear in PWI studies, from approximately -1200ms to 1200ms.

The results for latency exhibited a semantic facilitation effect at small-to-medium negative SOAs, consistent with other orthographic facilitation effects (Starreveld & La Heij, 1996). This effect disappeared at positive SOAs. Instead, latencies decreased as SOA increased, reflecting the fact that speakers needed less time to plan, and simply executed. Importantly, a semantic facilitation effect on duration also appeared at small-to-medium positive SOAs, although not at any of the negative SOAs. A summary of these results is shown in Figure 2. Balota el al. suggest that this reflects the influence of early activation of phonological codes, itself a consequence of preactivation or early selection of a lemma. Semantic priming, in this case, led to shorter word duration. The restriction in time to “post-recognition processes” suggests that “relatedness is influencing *how* the subject produces the word, not simply *when* (Balota et al., 1989, pp. 22).” This coheres well with the arguments for an effect of processing fluency outlined above. The Balota et al. manipulation differs significantly from standard PWI studies, in that it uses written primes and targets. The results may thus apply only within the written modality, but this is an open question.

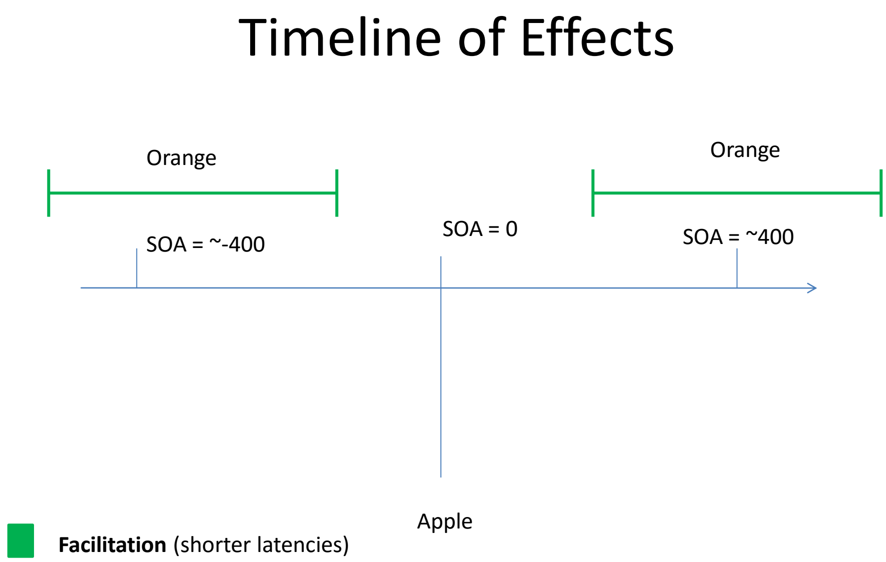

Figure 2: A summary of the effects of the word-naming studies in Balota et al., (1989). Facilitation of latencies is shown on the left, from -600SOA to -300SOA, and facilitation of durations is shown on the right, from 300SOA to 600SOA.

1.8 Experimental overview

The foregoing discussion, particularly the evidence reviewed in the previous section, leads to the hypothesis-driven question that underlies all of these experiments: How does the ease of completing the processes necessary for planning and producing an utterance affect the latency to begin speaking that utterance and the durations of its words? Implicit in this question, but worth noting explicitly, is the possibility that the duration of each word is not necessarily fully determined by the time a speaker begins an utterance, and thus that post-planning processes also have an influence. This is especially important for the two questions about retrieval, namely whether lexical retrieval and phonological retrieval are influenced by competition. Speeding or slowing the processing involved in retrieval may influence planning latency, but it may also influence word duration independently, if the production system does not always complete each of these steps before beginning an utterance. The question about identity priming also falls under this umbrella, because repeatedly processing the information associated with a target word will most likely ease planning and duration, but depending on the timing of presentation the two outcomes may vary independently (and, in fact, do). Finally, the action word follows the same pattern. Repeatedly producing the same word as part of a simple instruction will most likely make that whole instruction easier to plan. The more important question is whether that repetition also influences the duration of the words in the utterance as a function of making lexical and phonological retrieval easier, and what the timing of this influence is.

The PWI task operationalizes fluency of processing in terms of the effect of distractor relatedness on latency, and in this set of experiments, duration. This allows for an investigation of fluency both in terms of facilitation and inhibition, by comparing the related and unrelated conditions with the control and identical conditions. PWI operationalizes timing in terms of when a priming stimulus appears relative to when the target stimulus appears. This allows the experiments to simultaneously investigate facilitation and inhibition with a focus on when exactly those differences emerge, which sheds light on when the system engages in certain types of processing. To best capitalize on these operationalizations, the experiments below draw on a well-established literature on the timing of priming effects of fluency of lexical and phonological retrieval on latency to begin speaking.

Work in the domain of word production, and existing theories about language production, suggests that lexical retrieval, and all of the processing necessary for it, occurs in a certain time window during speech. Prior PWI findings have circumscribed this window by showing that semantic interference ends mostly or entirely before the time window where phonological processing begins (Schriefers et al., 1990; Starreveld & La Heij, 1996). This coheres well with theories of word production that say lexical retrieval must occur before phonological processing (Bock, 1982; Dell, 1986; Garrett, 1980; Levelt, 1996), and has informed at least one computational model of word production that capitalizes on the chronometric data (Level et al., 1999). The same PWI findings are somewhat mixed on exactly when phonological facilitation occurs, but generally agree that lexical retrieval difficulty from semantic inhibition and phonological retrieval ease from phonological facilitation take place at separate times. Two strong conclusions from these findings guide both the manipulations and the predictions for the current experiments: the effect of the semantically-related primes on latency should be inhibitory and should occur primarily if not exclusively when the primes appear before the targets (instead of the reverse), and the effect of phonologically-related primes on latency should be facilitatory and should appear later than the effects of semantic primes.

In order to preserve the theoretical approach while still asking novel questions, all of the experiments in the current study will take the same general form, based on previous manipulations of timing and primes (Balota et al., 1989; Schriefers et al., 1990; Starreveld & La Heij, 1996). Participants' basic task will be to see a picture move on a computer screen, and to describe this motion with a simple description. In response to seeing a picture of an airplane rotate, for example, the participant might say *The airplane rotates.* Two major manipulations surround this basic task – an interfering stimulus (IS) to interact with the processing of the target, and a timing manipulation, to assess when interaction between the IS and the target utterance can occur. The IS on each trial is a spoken word (or, in the control condition for two of the experiments, a non-linguistic cue to respond) which is mean to interact with processing, for example to interfere with lexical retrieval, or facilitate phonological retrieval, or to generate repetition.

The relationship between the IS and the target is what allows the manipulations to answer the first three questions, about lexical and phonological retrieval, and identity priming. The simplest contrast will hold between a silent no-prime control condition and an identity condition, where the name of the picture itself serves as the prime. The control condition establishes the baseline of processing difficulty for a particular item, without any specific external interference. The identical condition represents a higher boundary of facilitation, as seen in studies investigating givenness and/or repetition (e.g. Fowler & Housum, 1987; Kahn & Arnold, 2012). Including both of these allows for a comparison with prior work on repetition priming (Glaser & Glaser, 1989; Kahn & Arnold, 2012), as well as an independent comparison of whether identical primes induce ease of lexical and phonological retrieval, and thus shorter latencies and most likely shorter durations, compared the control condition. The related condition is the core prime type, as in most PWI studies. Where participants' responses fall in this condition relative to both the control and identical conditions will determine whether the related prime facilitated or inhibited processing. The most important comparison is between the related and unrelated condition, which will determine whether relatedness *per se* caused any differences in outcome. This is the key to answering the first and second questions, about the fluency of lexical and phonological retrieval in the face of competition.

My experiments depart more significantly in timing from prior work than they do in the IS manipulation. Most studies in the PWI format introduce only a small temporal gap between prime and target, sometimes even having them appear simultaneously. They also do not make a strong distinction between primes being presented before versus after the appearance of the target, because participants are told to ignore the primes in favor of simply naming the target. The experiments in this dissertation have fairly large temporal gaps between primes and targets, and although the participants were instructed to ignore the content of the primes, the manipulations make a larger difference than previous studies between primes that appear before targets and primes that appear after targets. The experiments were run in pairs (a & b), with one including prime stimuli that occur only on one side or the other of the prime. In other words, Experiment 1a had only primes that appeared before targets, and Experiment 1b had only primes that appeared after targets. Figure 3 illustrates this difference between the two experiment types – one that has primes that appear before the targets (at negative temporal offset), and one that has primes that appear after the targets (at positive temporal offset). PWI studies typically block SOA, presenting only one for a set of trials, then randomly selecting another. The separation of negative and positive SOAs also sidesteps the complication of having to make participants keep two sets of instructions in mind simultaneously, one for positive and one for negative.

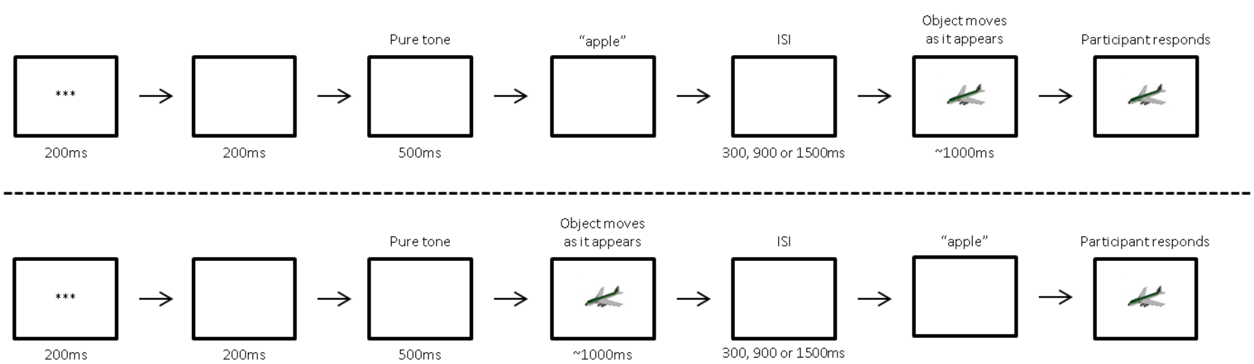

Figure 3: (Top) – Order of presentation of the events in Experiment 1a, with prime before target. (Bottom) – Order of presentation of the events in Experiment 1b, with target before prime.

One important difference between the experiments below and previous PWI studies is the use of Interstimulus Intervals (ISIs) instead of SOAs. Previous PWI studies calculated the distance between primes and targets based on when the prime initially appeared or was spoken to the participant and when the target appeared, which is how SOA is defined. The studies below use an ISI instead, which is slightly different for the two experiment types. In the a types, where the primes appear before the target, the ISI spans the temporal distance between the offset of the spoken prime and the onset of the target object's motion. In the b types, where the primes appear after the target, the ISI spans the temporal distance between the end of the object's motion and the beginning of the spoken prime. As contrasted with SOAs, the ISI allows participants to hear the whole prime, even in the positive ISI case (because utterances that were initiated earlier were excluded from analysis). This is especially important for the phonological primes, which share onsets but not offsets. In that case, the ISI manipulation ensures that they hear the offset phoneme(s) as well before they initiate their response. One downside of using an ISI instead of an SOA is that it makes comparing the results below to previous PWI studies more difficult, because they use slightly different timing. On the other hand, one advantage of the current timing scheme is that it targets a new set of questions. SOA manipulations are frequently intended to capture transient effects, and have done so. The use of ISI allows me to test whether previously ephemeral effects of semantic interference and phonological facilitation also persist through the full presentation of a several-hundred millisecond long priming stimulus. From the perspective of cognition, this means looking at longer-lasting consequences of processing of the prime, in the form of longer-term retention, rehearsal, or other less-transient shifts in the state of the language production system.

Note that this does not entail that ISIs should unwaveringly exhibit similar timing patterns as the SOAs used in previous studies, because the longer interval between prime onset and target onset means that the primes in this experiment appear well beyond the point at which previous PWI work has measured ease of semantic processing. In all likelihood, the longer interval between prime and target onsets will lead to the ISIs tapping into at least some processes that SOAs do not (with the reverse also being true). For example, the longer interval may allow and/or require speakers to retain more information about the prime, through elaboration or rehearsal. Such a routine would still allow semantic interference to occur, but at ISIs that exceed the magnitude of previously-investigated SOAs. Another possibility is that ISIs are simply too long to tap into production-relevant processes. Speakers may hear the primes in this task and simply discard them, which would result in a null effect on latency. This is particularly likely (although not ultimately the case) at the longest ISI (1500ms), given that speakers have no reason to retain information about the prime or rehearse it in this task.

Requiring multi-word utterance responses also creates an opportunity to include an additional manipulation, to investigate facilitatory effects on other places in the utterance. Like the experiments, the chapters are paired, with Chapters 2 and 4 making use of a prime that is meant to interact with the processing of the object word, and Chapters 3 and 5 making use the object prime manipulation as well as of a manipulation (not a prime) that is meant to interact the processing of the action word. Chapter 2, for example, will present a picture of an airplane rotating, and ask speakers to describe this event while ignoring a spoken prime word such as "helicopter." On other trials, the objects will perform different actions, such as expanding, or shrinking, all randomly selected. In Chapter 3, another set of participants might see the same pairing of target and prime (the airplane and "helicopter"), but in a block of trials where all of the targets rotate. This orthogonal contrast between consistent and random action words allows for a simultaneous investigation of the first three questions (semantic and phonological relatedness, and identity priming) and the fourth question, whether facilitation of the processing of the action word influences duration as well.

The questions that these experiments address are interrelated, and the data are dense and complex. I have chosen to organize the presentation in a way that addresses each question more or less in order. Chapters 2 and 3 will thus focus on the question of whether fluency of semantic processing influences word duration. They will also include a partial discussion of identity priming, which is investigated in parallel with questions about relatedness. Chapters 4 and 5 will shift the focus to the question of whether fluency of phonological processing influences word duration, and will also include the remainder of the discussion about the effects of identity priming. Discussion of the fourth question, about how the fluency of processing the action word affects duration, will appear at the end of Chapter 5, after all of the data have been presented. The final chapter will address some other minor issues, tie up loose ends, and suggest some future directions for work on word duration.

Chapter 2: Experiments 1a&b

2.1 Introduction

The central question of this project is what affects the duration of spoken words. Chapter 2 focuses particularly on whether, how, and when the difficulty of lexical retrieval affects their duration, as well as whether identity priming exhibits repetition-based reduction. Of the processes involved in semantic processing generally, the lexical retrieval manipulation here focuses on the selection of a particular word when other semantically-related words are also available for processing, and the identity prime manipulation focuses on the reprocessing of a word after repetition. The previous chapter established the project's broad purpose of looking at variation in production processes, as well as the importance of the timing of interference and the appearance of these effects on particular words and times in a speech stream. This chapter makes these questions more specific. First of all, does the ease or difficulty (i.e. overall fluency) of lexical retrieval during multi-word utterance production affect the way in which that utterance is produced? Second, does the timing of the influence matter? In other words, does any interfering stimulus have to appear at, before, or after a certain time point? And if it appears at one time point, how long does its influence persist? Each of these questions is relevant under two major realms of inquiry.

2.2 Questions about competition

PWI studies find that only within a fairly narrow range of milliseconds does a related stimulus have an effect on latency, and thus on planning (Schriefers et al., 1990), although they disagree on when precisely this window begins and ends (Starreveld & La Heij, 1996). One major open question here is whether relatedness might also have an effect on duration, either in addition to or in place of its effect on latency. The time window may be narrow or wide, and it may overlap partly or entirely with the window for when relatedness affects latency. This overlap is particularly likely if the planning and articulation of an utterance both have an effect on the duration of the words in that utterance. A quickly-planned utterance may be executed quickly, leading to overall speeding (Kello, Plaut & MacWhinney, 2000). The coupling between planning and execution may also lead to shortened latencies at the same time as lengthened durations, as the system compensates for speeded planning by drawing out more-difficult words.

Two other minor questions arise from the departure from single-word investigation to multi-word utterances. The first is whether the additional planning time necessary for a multi-word utterance leads to effects appearing at longer ISIs, and the second is whether duration effects appear on words other than the target.

The greatest magnitude of SOA used in most PWI experiments, either positive or negative, is approximately 500ms. This stems from the assumption that many of the effects of interest are short-lived, and appear in narrow time windows (Balota et al., 1989; Levelt et al., 1999; Schriefers et al., 1990). Under that assumption, it does not make sense to investigate large SOAs nor ISIs. For a multi-word utterance, however, the assumption of transient effects does not necessarily hold. Planning to begin speaking often takes on the order of 800 milliseconds for simple utterances (Griffin & Bock, 1998), meaning that the plan may begin before the prime even appears, in the case of the positive primes. In the case of the negative primes, the target word is not available until after the prime appears, so formulation of a plan could consist only of generating a syntactic frame and other pre-content-word operations. To cover both possibilities, the tasks in the experiments in this chapter include ISIs -1500, -900, and -300 (or 300, 900, and 1500, depending on the experiment).

In all of the experiments, a moving picture is used as the target instead of a static image. The primary function of the movement is to elicit a multi-word utterance from the participant, including a determiner (viz., “the”), the target word, and a verb or action word. This allows for the measurement of four regions, as opposed to only one, which is typical. The use of a multi-word utterance maximizes the chances of observing when behavioral differences, if any, appear in the participants’ responses. A standard PWI task would measure only onset latency, but the present task has four outcome variables. With this paradigm, it is possible to ask whether the influence of lexical retrieval appears on the targeted word, or others in the utterance also. The focus for this chapter is on the onset latency and the object word, because the primes are focused on eliciting processing differences on only those two regions. The next chapter (as well as Chapter 5), will extend the investigation to the duration of the action word. The other region, the determiner in the simple utterance, will not receive any attention.

2.3 Questions about identity priming

Another open question is whether and at which time points identity effects will appear, both on latency and duration. Prior studies using this or similar paradigms find fairly ubiquitous effects of an identity prime on latency, but do not report data on duration (Glaser & Glaser, 1989; Starrveld & La Heij, 1996). The only study I am aware of that asks about the relationship between latency and duration effects with an identity prime did not vary SOA (Kahn & Arnold, 2012). The current study thus represents a significant step forward in understanding the timing of identity priming and its effects on word production.

Two patterns of results are likely here, depending on how the production system deals with repetition. If production must incorporate repetition into the plan, treating the repeated word as part of an ongoing discourse or conversation, only the negative ISIs should exhibit the full range of effects of repetition, namely reduced latencies as well as durations. At the positive ISIs, participants should not have enough time to fully make use of the repetition, because they will have planned their utterance almost entirely by the time the prime cues their response. However, if the production system is also sensitive to post-planning repetition, the positive ISIs may exhibit durational differences without latency differences. This would reflect a more online, less plan-driven form of durational determination.

2.4 Experiments 1a&b

Methods

Participants. 36 native English-speaking undergraduates from the University of North Carolina at Chapel Hill participated in each experiment (72 participants total) for pay ($7.50 for half an hour).

Materials & Methods.

Visual materials. 100 colored line drawings of simple objects served as picture stimuli (Rossion & Pourtois, 2001, drawing from Snodgrass & Vanderwart, 1980). These drawings were normed by native French speakers for imageability, visual complexity, familiarity, and name agreement in French. The latter measure was not used, because the speakers here responded in English.

Recording equipment. The auditory primes were recorded on a Marantz professional solid-state recorder over an Audio Technica professional microphone headset**,** using citation format (i.e., no distinguishing inflection, and clear pronunciation).

Trial Structure. The general character of the trials is the presentation of a target stimulus, which moves, and an interfering stimulus presented either before or after the target. The first of two major differences between the two experiments' trials was their temporal order. Primes appeared before targets in 1a, and targets appeared before primes in 1b – otherwise everything was the same. Trials in 1a consisted of 1) a string of 3 asterisks appearing for 200ms; 2) a blank screen for 200ms; 3) a 400Hz pure tone for 500ms; 4) the spoken prime; 5) the ISI gap; 6) the picture appearing and moving. Movements took four forms: rotating, shrinking, expanding, or fading. Trials in 1b consisted of 1) a string of 3 asterisks appearing for 200ms; 2) a blank screen for 200ms; 3) a 400Hz pure tone for 500ms; 4) the picture appearing and moving; 5) the ISI gap; 6) the spoken prime.

Prime types. Three types of prime were used, in addition to a control condition: identical, related, unrelated. For the identical condition, the English name of each picture as reported in Rossion and Pourtois (2001) was recorded. The control condition consisted of silence instead of any object name. For the related condition, primes were chosen so that they formed a close association with the target object without being fully synonymous (e.g. *flashlight* as a prime for *lamp*). The related primes also did not exhibit a phonological relationship with the picture, or a semantic relationship with another picture. The appendix contains a list of primes for both this and the phonological experiments.

Table 1: An example set of interfering stimuli for the target picture object, airplane. These are actual stimuli used in the experiments.

Target Example	Name (Identical)	Related Prime	Example unrelated Prime	Control Prime
	"Airplane"	"Rocket"	"Apple"	[Silence]

Timing of prime. Each item was also presented at a particular ISI, with the prime appearing at either 1500, 900, or 300ms before the target (hereafter referred to as the -1500 ISI, -900 ISI, and -300 ISI conditions) in Experiment 1a, or the same times after the target (hereafter referred to as the 300ISI, 900ISI, and 1500ISI conditions).

List Design and Counterbalancing. Participants saw each item only once, at a particular condition and ISI, both of which were counterbalanced across participants. Condition had four levels, and ISI had three, for a total of twelve possible combinations of condition and ISI. Twelve groups of items consisting of 8 items each were created and matched with one of the combinations of condition and ISI (e.g. unrelated-300 or control-900 or identical-1500). Twelve lists were created as well, each with all 96 items (the 12 groups of 8) matched to a condition and ISI pair, with counterbalancing of item and condition-ISI across the lists. This ensured that items appeared once per list (and thus only once per participant) and equally often in each and every condition-ISI pairing across all participants. See the appendix for example lists.

The unrelated prime word for an item in the unrelated condition on a particular list was formed by randomly pairing pictures in this condition with the related primes from items appearing in the control condition, without allowing semantic or phonological similarity. Each list was seen by 3 participants, and each subject saw each item once, in one ISI-condition pairing (a total of 96 items). This contrasts with PWI practice, where participants frequently see each item multiple times. A single presentation was used in order to reduce repeated-mention reduction effects (Fowler & Housum, 1987; Kahn & Arnold, 2012). ISI was blocked, again in keeping with standard PWI practice, such that a participant saw all of the -1500 ISI items at once, for example, before seeing all of the -900 ISI items, and then seeing the -300 ISI items. ISI block order was counterbalanced among participants to control for order effects (i.e. if one participant saw -900ISI first, the next participant saw -300ISI first, and the next saw -1500ISI first, and so on).

Object movements. A rotating object would quickly make a 360 degree spin about its center. A shrinking object would shrink to roughly half its original size, while an expanding object would roughly double in size. A fading object would be mixed with a blank image, so that it appeared to fade into the white background. The movements all took approximately the same amount of time (i.e. rotating took approximately the same amount of time as fading, and rotation actions on two arbitrary trials took approximately the same amount of time, 1000ms in all cases). Movement began immediately when the object appeared on the screen.

Procedure. Participants sat in front of a laptop computer, approximately two feet away from the screen. They were told by the experimenter that they would see simple objects appear on the screen and move in one of four ways, and that subsequently they would have to describe what they saw, speaking as quickly and accurately as they could. They were shown examples of each of the four movement types (rotate, shrink, expand, or fade), and told how they should describe each movement. For example, after seeing a picture of a lobster rotate, the participants would be told to describe that movement as "the lobster rotates." Participants were corrected at any point during the instruction and practice phase if they used the wrong sentence frame, particularly if they dropped the determiner ("the") or used the past tense.

Participants then saw 20 practice stimuli with the ISIs all randomly chosen from -1500, -900, and -300 ms, and stimuli drawn from the Rossion & Portois (2001) picture database not used in the testing phase. Five pictures were presented in each condition, always in the same order across participants, first from the control condition, then identical, then unrelated, and finally related condition. Here participants in Experiment 1b were corrected if they began naming before the prime (i.e. did not wait until the prime had appeared). Neither the pictures nor primes for the practice movements nor the practice stimuli appeared during the remainder of the experiment. Participants were asked whether they understood the task, and the experiment began after addressing any questions they had. A customized Python script controlled the experiment. The entire experiment took approximately 25 minutes to complete, with 10 minutes for instructions and 15 for the script to run.

Analysis. Individual utterance recordings were submitted to forced alignment software that generated time-aligned textfiles for each segment of interest (latency, determiner, object and action word; Yuan & Liberman, 2008). The textfiles were inspected by research assistants both visually and acoustically for accuracy using acoustic analysis software (Praat; Boersma & Weenink 2009), and these modified textfiles were finalized by the author. The duration of each segment was then automatically extracted from the textfiles with Praat, and submitted individually to the analyses below. The onset to begin speaking was measured from onset of the target movement. Each other word duration was defined as closely as possible as the onset to the offset of the word. Utterances were excluded from the model if they were disfluent in any way (i.e., pauses longer than 250ms, coughs, false starts, disfluent 'the') or if the speaker used the wrong object or action word. The means and standard deviations reported below come from this trimmed dataset, as do the models.

The logarithm of each duration was taken to increase its linear relationship with the predictors. Outlier exclusion was not necessary after this transformation. This variable was then submitted to a multi-level model, built in two steps. Models were constructed with the lmer function from R's lme4 package, which uses maximum likelihood estimation to obtain parameter estimates. Baseline models were constructed in order to specify the random effects structure and to identify the significant control variables. Each baseline model included random intercepts for subject and item, as well as random slopes for condition by subject and item when possible. If the slope for the contrast correlated with the intercept, or the individual contrasts correlated with each other, greater than .9, the redundant slope was removed to avoid over-parametrizing the model. This latter condition frequently (although not always) obtained, with no particular pattern in the relationship among DV, random factor, and inclusion.

After adding the random effects, each model was built stepwise to include the following control variables: log frequency (based on the original Snodgrass & Vanderwart (1980) English names), imageability, visual complexity, and familiarity of the experiment target (taken from Rossion & Purtois (2001)), the number of syllables of the experimental target, and the trial number (to control for practice effects). The durations of the other segments of the utterance were also included in order to control for speech rate. The model for object noun duration, for example, included the onset, article, and action word duration. Any predictor that did not approach significance during this stage (a *t* value of approximately 1.5) was removed. A typical baseline model included the three random effects and approximately 4 fixed effect predictors. Detailed accounts of which predictors were included in each model can be found in the appendix.

One of these baseline models was constructed for each combination of ISI and utterance segment (e.g. a single model for the onset to begin speaking for stimuli presented at -1500ISI), for a total of twelve models. Condition was entered as a treatment-coded predictor to these baseline models, and then submitted to Westfall tests for multiple comparison (Bretz, Hothorn & Westfall, 2011). This allowed for tests of all pairwise comparisons (each level of condition against each other level, within ISI) while correcting for simultaneous inference. The Westfall correction procedure is slightly less conservative than Tukey comparisons while preserving a low

Type I error (Bretz et al., 2011). The final analyses are reported with the Westfall estimates of the differences between conditions, and the associated significance. The report below focuses on the theoretically-relevant comparisons: related vs. unrelated, and identical vs. control.

Results

Summary statistics for each condition and ISI appear in Tables 2 through 5, broken down by segment. The effects of relatedness show up on latency at negative ISIs, and on duration only at -900ISI, with no effects at all at positive ISIs. The overall effect of identical primes was to facilitate both latency and duration, although with slightly different temporal profiles. Parameter estimates for each of these effects appear in Tables 6 and 7.

Table 2: Onset latencies and standard deviations (in parentheses) for Experiments 1 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	2293 (433)	2269 (351)	2214 (299)	1480 (563)	1368 (436)	1378 (408)
Identical	1853 (390)	1875 (414)	1834 (389)	1702 (448)	1719 (432)	1795 (525)
Related	1939 (395)	1921 (378)	1882 (409)	1809 (423)	1821 (430)	1897 (524)
Unrelated	1813 (332)	1886 (414)	1818 (332)	1769 (418)	1815 (431)	1820 (445)

Table 3: Determiner durations and standard deviations (in parentheses) for Experiments 1 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	89 (31)	83 (25)	83 (29)	101 (35)	102 (34)	104 (33)
Identical	81 (27)	78 (27)	74 (26)	101 (33)	102 (35)	102 (38)
Related	89 (32)	88 (29)	86 (30)	103 (33)	108 (37)	108 (34)
Unrelated	87 (32)	87 (32)	84 (29)	104 (35)	108 (39)	108 (37)

Table 4: Object word durations and standard deviations (in parentheses) for Experiments 1 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	387 (101)	388 (106)	387 (111)	416 (111)	422 (114)	419 (111)
Identical	370 (103)	367 (104)	363 (105)	419 (113)	414 (109)	417 (114)
Related	403 (107)	410 (110)	400 (112)	429 (108)	434 (122)	430 (121)
Unrelated	403 (114)	398 (109)	398 (118)	430 (118)	434 (109)	429 (113)

Table 5: Action word durations and standard deviations (in parentheses) for Experiments 1 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	661 (131)	652 (122)	656 (130)	691 (131)	680 (125)	691 (129)
Identical	640 (132)	637 (121)	628 (117)	693 (130)	698 (128)	697 (131)
Related	667 (113)	648 (121)	642 (124)	697 (124)	706 (124)	693 (134)
Unrelated	651 (122)	658 (108)	641 (119)	692 (124)	700 (129)	708 (141)

Table 6: Westfall comparisons for Experiment 1a. Dashes indicate non-significant effects (p>.05)

		Critical Comparisons				Non-critical comparisons							
Segment	ISI	Identical vs. control		Unrelated vs. related		Identical vs. Unrelated		Identical vs. Related		Unrelated vs. Control		Related vs. Control	
		z	p	z	p	z	p	Z	p	z	p	z	p
Latency	-1500	-2.86	0.01	-3.92	0.001	-2.2	0.05	-1.84	0.07	-4.98	0.001	--	--
	-900	--	--	--	--	--	--	--	--	--	--	--	--
	-300	--	--	2.34	0.038	-2.82	0.025	--	--	-2.52	0.03	--	--
Determiner	-1500	-2.58	0.05	--	--	-2.43	0.05	-2.42	0.05	--	--	--	--
	-900	-2.36	0.04	--	--	-3.53	0.002	-3.59	0.002	--	--	--	--
	-300	-3.98	0.0001	--	--	-5.43	0.0001	-4.89	0.0001	--	--	--	--
Object	-1500	-2.6	0.02	--	--	-5.45	0.001	-5.15	0.001	2.78	0.01	2.68	0.01
	-900	-3.85	0.001	1.91	0.05	-5	0.001	-6.77	0.0001	1.91	0.05	3.14	0.005
	-300	3.06	0.004	--	--	-5.98	0.001	-5.25	0.001	-2.84	0.01	-2.28	0.02
Action	-1500	--	--	--	--	--	--	--	--	--	--	--	--
	-900	--	--	--	--	--	--	--	--	--	--	--	--
	-300	-2.84	0.02	--	--	--	--	--	--	--	--	--	--

Table 7: Westfall comparisons for Experiment 1b. Dashes indicate non-significant effects (p>.05)

		Critical Comparisons				Non-critical comparisons							
Segment	ISI	Identical vs. control		Unrelated vs. related		Identical vs. Unrelated		Identical vs. Related		Unrelated vs. Control		Related vs. Control	
		z	p	z	p	z	p	z	p	z	p	z	p
Latency	-1500	14.32	0.001	--	--	--	--	2.4	0.043	15.41	0.001	16.39	0.001
	-900	13.39	0.001	--	--	3.13	0.004	3.22	0.004	16.33	0.001	16.3	0.001
	-300	8.4	0.001	--	--	2.28	0.023	3.54	0.001	10.43	0.001	11.62	0.001
Determiner	-1500	--	--	--	--	2.12	0.085	2.62	0.04	--	--	--	--
	-900	--	--	--	--	3.03	0.013	2.44	0.039	--	--	--	--
	-300	--	--	--	--	--	--	--	--	--	--	--	--
Object	-1500	--	--	--	--	--	--	--	--	--	--	--	--
	-900	--	--	--	--	3.48	0.003	3.13	0.005	--	--	--	--
	-300	--	--	--	--	2.794	0.014	3.08	0.011	2.82	0.014	3.07	0.011
Action	-1500	6.84	0.0001	--	--	--	--	--	--	7.83	0.0001	6.71	0.0001
	-900	7.28	0.0001	--	--	2.42	0.041	2.39	0.041	9.2	0.0001	9.12	0.0001
	-300	3.26	0.002	--	--	--	--	--	--	4.48	0.001	5	0.001

Latency. Comparisons between the related and unrelated conditions revealed significant lengthening after related primes at -1500ISI and -300ISI, and a marginal effect at -900ISI. There were no significant differences at the positive ISIs. These results are summarized in Figure 4.

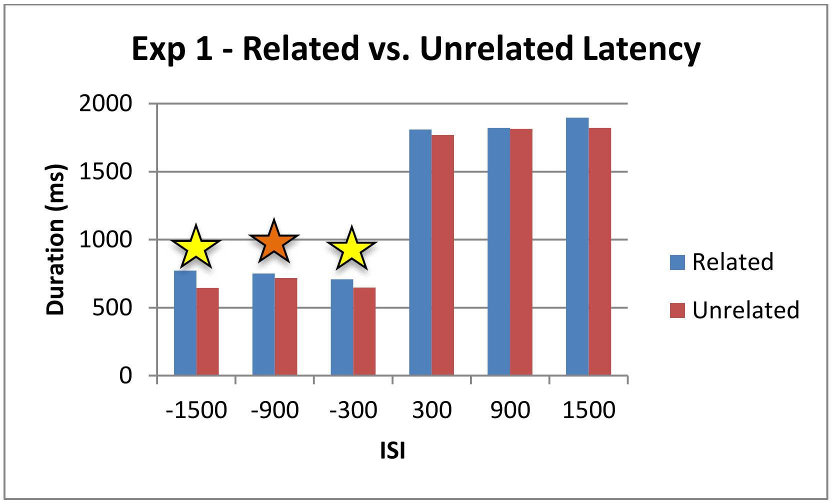

Figure 4: Latencies for the related vs. unrelated comparisons, with yellow stars to indicate comparisons that were significant (*p* < .05), and orange stars to indicate comparisons that were marginally significant (*p* < .01).

Comparisons between the control and identical conditions revealed two widespread effects in opposite directions. At the negative ISIs, the identical condition was shorter than control. At the positive ISIs, the reverse was true. This is contrary to expectations, and is discussed further below. These results are summarized in Figure 5.

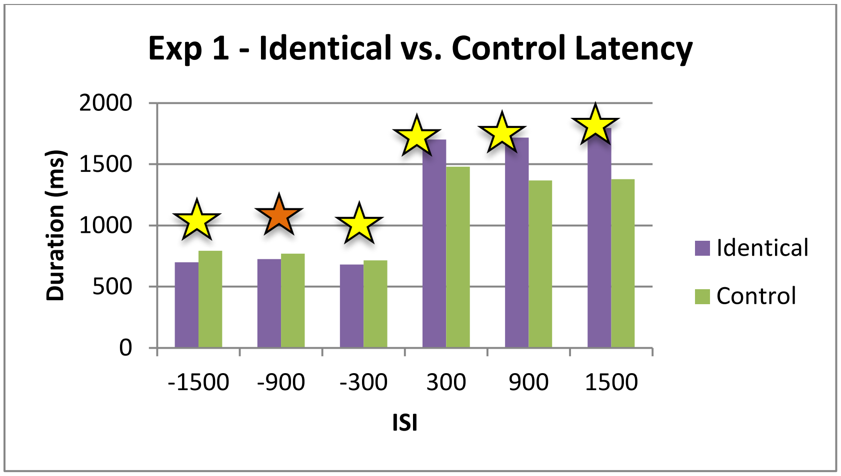

Figure 5: Latencies for identical vs. control comparisons, with yellow stars to indicate comparisons that were significant (*p* < .05), and orange stars to indicate comparisons that were marginally significant (.05 < *p* < .01).

Object word. For the related vs. unrelated comparisons, the only (marginally) significant difference ($p < .056$) was at -900ISI. No other comparisons were significant, as shown in Figure 6.

As with latency, the differences between the control and identical condition were slightly more robust. All negative ISIs exhibited significant differences, with identical shorter than control, while none of the positive ISIs did, as shown in Figure 7.

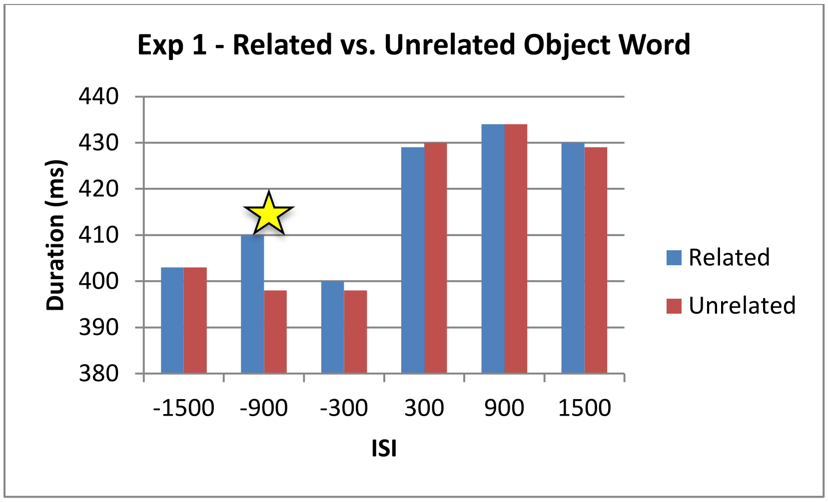

Figure 6: Object word durations for the related vs. unrelated comparisons, with yellow stars to indicate comparisons that were significant (*p* < .05).

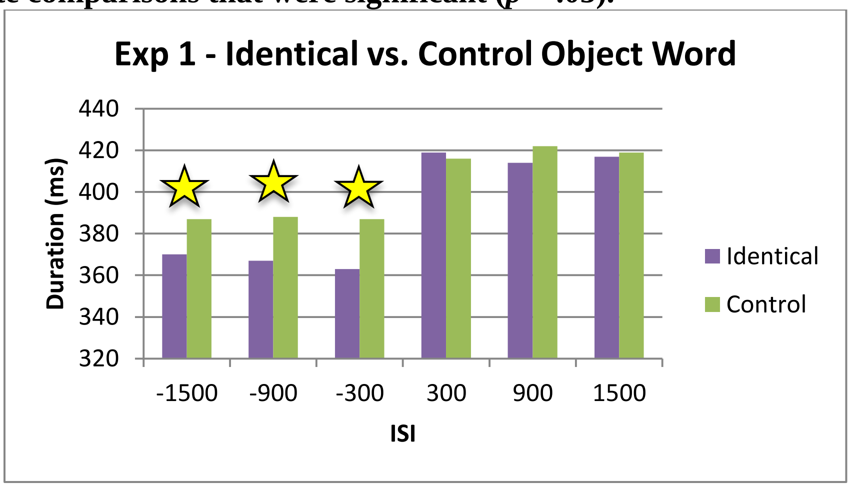

Figure 7: Object word durations for the identical vs. control comparisons, with yellow stars to indicate comparisons that were significant (*p* < .05).

Discussion

Questions About Competition. Theoretically similar to previous findings in PWI, the related condition elicited significantly longer latencies than the unrelated condition all of the negative ISIs. Prior studies have shown that semantically-related primes elicit longer latencies to begin speaking than semantically-unrelated primes (Balota et al., 1989; Schriefers et al., 1990; Starreveld & La Heij, 1996), but only at short, negative SOAs. This has typically been explained as a difference in planning difficulty or timing, where it takes longer to activate or select the appropriate item in the presence of semantically-related competitors. That explanation seems appropriate here as well, although with the caveat that ISIs are not the same as SOAs. This difference results in speakers experiencing a larger temporal gap between prime and stimulus presentation than participants in earlier studies, but they still exhibit the inhibitory effect. To compare, the average duration of the primes in these experiments is approximately 700ms, which effectively makes the ISIs into SOAs of 1000, 1600, and 2300ms. To my knowledge, a semantic priming effect has not been reported at SOAs nor ISIs of this magnitude, the earliest effect being the one reported for -600ms in Balota et al., (1989).

The presence of inhibition in the current experiments likely reflects two characteristics of the paradigm: the sentential frame and the blocking design. The sentential frame requires participants to engage in more-complex mental processing than most PWI studies, which typically require only a single-word response. Bock (1982), for example, puts sentential frame processing temporally prior to lexical access in some cases. Under this scheme, participants would have begun, if not also completed, formulating a syntactic frame before beginning lexical access. It is thus not surprising that effects might appear at slightly larger temporal differences (viz. ISIs) than normal. Assuming that some amount of information processing is necessary to determine how the activated semantic information will fit into the inchoate sentential frame (Bock, 1982; Levelt, 1989), having to prepare a more-complex sentential frame (i.e. one with a determiner and verb, instead of a single word) may commit speakers more firmly to their initial formulation. In other words, participants might experience semantic inhibition effects of the same type as in single-word naming, but instead as a function of having committed early to a sentence plan that they must then revise. The blocking design also may serve as an aid to participants, who may prepare and hold a response in abeyance until the target appears. The two explanations may interact, as well – if participants anticipated that they would (sometimes) need the prime word, they might activate and have it ready earlier in the course of processing than in a less-predictive design. However, a pilot experiment showed that even with an unblocked, fully randomized SOA presentation, participants exhibited the semantic priming effect at -1500ms, suggesting that they may have learned to treat 1500ms as the cut-off point.

Positive ISIs did not exhibit any semantic inhibition, which is theoretically consistent with prior work, respecting the caveats already discussed. Previous PWI studies have found that semantic interference tends to be restricted to negative SOAs (Schriefers et al., 1990), or at most, very short positive SOAs (Starreveld & La Heij, 1996). According to these accounts, this is because lexical retrieval has effectively finished by the time the prime appears, meaning that it cannot create competition. The same argument applies here – the primes appeared far too late to elicit competition between them and the target, meaning speakers did not experience any competition-based difficulty in selecting the appropriate lexical item.

The other novel relatedness outcome was the semantic inhibition effect on the object word duration at -900ISI. Two things are notable about this effect: 1) its appearance at all – to my knowledge the first evidence of durational variance of any sort in a PWI paradigm; 2) its absence in the other ISIs, especially the other negative ISIs. As covered in the first chapter, PWI studies do not typically measure duration. The notable exceptions (Jescheniak et al., 2003; Schriefers et al., 1990) found null effects. I postpone a more detailed discussion of this effect until all of the semantic relatedness data have been reported, at the end of the next chapter.

Questions About Identity Priming. There were reduced latencies and object word durations in the identical compared to the control condition at negative ISIs, which coheres well with the theory underlying previous PWI findings (Glaser & Glaser, 1989), as well as previous work of mine (Kahn & Arnold, 2012). Here identity primes appear to have made processing easier overall, both the planning of the utterance and its execution. However, the reduced durations could reflect either meaning-based (information status) or facilitation effects. Even in a sparse conversational context such as this, speakers could have used the prime as a reason to trigger reduction (similar to Schmidt, et al., 1999), as opposed to reducing because of facilitated processing. In the Schmidt et al. model, speakers select a pronoun instead of a full lexical phrase in situations where the conversational context licenses a pronoun (e.g. a repetition). In the current experiment, speakers' production systems might only have needed to recognize the fact of repetition in order to choose a reduced form, akin to choosing a pronoun. On a more facilitation-based account, by contrast, the identity prime could have activated representations ahead of time, or lowered the threshold for selection, or eased the process of integrating the representations with other parts of the utterance. Both of these explanations account for the experimental outcome well, and the current data cannot distinguish between them.

At the positive ISIs, identical latencies were longer than control latencies. Unfortunately, the manipulation in these experiments limits the ability to interpret this difference. In Experiment 1b, participants had to respond to a visual cue (the row of x's) instead of an auditory one in the control condition (i.e. the spoken prime). This meant that the control condition was slightly easier to respond to than the identical, and participants could begin their responses sooner. Comparing the identical condition with the related and unrelated conditions shows facilitation from identity priming, lending support to this hypothesis. Experiments 2 and 4 address the disparity between identical and control prime modalities by introducing a spoken control prime instead of a silence control.

The identity prime led to shorter object word durations only at the negative ISIs, which is somewhat surprising, given that it should have a similar effect on processing regardless of the timing. The result is open to two interpretations. One is that the difference between the positive and negative ISIs appears here as well, more precisely as a consequence of it. The shorter planning time in the control condition may have allowed speakers to reduce their overall production time, particularly on the latency and object word duration. On this interpretation, the difference between the control and identical conditions' effects on latency lies not in the ease of planning, but the timing of the cue to begin speaking. It takes less time to finish processing a line of printed x's than it does to listen to a word, meaning that speakers could begin their response more quickly in the control condition. The lack of a difference in duration between the control and identical conditions at positive ISIs would be explained in terms of the overall ease of completing the naming task in the control condition, which did not exceed nor lag behind the ease of completing it in the identical condition. This explanation finds support in Chapters 3 and 5, where the durational results more closely reflect identity-based facilitation.

Chapter 2 has shown that PWI can be extended to multi-word utterances and make use of ISIs and still show the same general patterns of semantic inhibition as seen in earlier studies. Identity priming also appeared, but not in all of the expected places, due to a limitation of the manipulation. The next chapter will continue the investigation of both of these questions (whether semantic relatedness and identity priming affect word duration), while adding a new layer of complexity.

Chapter 3: Experiments 2a&b

3.1 Introduction

The previous chapter showed that the difficulty of lexical retrieval has an influence on latency, and also object word duration in a limited time window. In multi-word utterances, however, it is possible that the ease of retrieving each lexical item, not just the target word, may have an influence on the duration of not only that word, but others in the same utterance. It is also possible that integrating several words together has an influence on word duration, depending on the relative difficulty of retrieving each of them, and of ensuring that each fits into the appropriate place in the utterance. In addition to the questions from the previous chapter, this chapter will thus also investigate how the ease of retrieving and integrating the action word (including lexical and phonological retrieval) influences the duration of each of the segments of the same multi-word utterance. Along with the manipulations of prime and ISI, the experiments here (and in Chapter 5) will include blocks where participants use the same action word throughout (viz. “rotates”). Relative to the randomized actions in the first two experiments, this manipulation is designed to facilitate retrieval and integration of the action word. Modulating the ease of processing the action word may have an influence on how speakers plan and execute the entire utterance, if they engage in any planning of the action word before they begin speaking. Even if they do not, the ease of processing may influence the duration of the action word itself, or the object word which immediately precedes it, because speakers will have less difficulty creating an utterance plan and executing it.

3.2 Relationship with previous work

Chapter 3 will be able to draw not only on prior PWI work, but also the previous chapter, to increase confidence that the method is consistently tapping into the same or similar processes. Chapter 2 showed a semantic inhibition effect at both -300 and -1500ISI; the experiments here should show the same pattern. Experiment 1a also showed significant lengthening of the object word after related primes, but only at -900ISI. The null effect at -300 and -1500 may have been a lack of power, or a genuine null effect. It was judged to be more important to investigate whether the addition of the consistency condition would lead to durational variation at -300 and -1500ISI than to continue pursuing the -900 effect. If the absence of an effect at -300 and -1500ISI was truly a null effect, object word duration should not vary in this experiment as a function of relatedness.

3.3 Questions about competition

This chapter further investigates the same questions about relatedness as in Chapter 2, namely whether ease of lexical retrieval affects latency and word duration, and whether these effects are restricted in time. The previous chapter established that latency is readily subject to semantic inhibition effects, but duration is not. That serves as preliminary evidence that the latency the speaker takes to begin speaking is not the only factor that influences the length of the words in an utterance, however much of the latency is taken up by pre-planning. Only at -900ISI did the previous experiments show any difference in word duration. This was coupled, however, with a lengthening effect on latency. It may be the case that a longer latency does not necessarily result in a longer target word duration, but that a longer target word duration necessitates a longer plan. The manipulations in Experiment 3 restrict the range of ISIs to the 300 and 1500 magnitude, positive and negative, to focus on assessing the plausibility of that causal relationship. This restriction also serves as further investigation of whether speakers experience semantic inhibition difficulty at ISIs that exceed the magnitude of previous investigations' SOAs. A replication of the inhibition findings from Chapter 2 would further strengthen the idea that speakers retain or otherwise continue to experience semantic competition when planning their utterances, whether this interference emerges primarily during pre-planning or whether it occurs as part of incremental planning-while-executing.

3.4 Questions about identity priming

The identity prime here carries the same theoretical and practical importance as it did in the previous chapter. The questions are whether latency and duration reduction will appear, and if so, whether they appear together. Chapter 2 showed consistent reduction effects at the negative latencies, but the interpretation of the positive latencies was complicated by the use of a printed control prime instead of an auditory one. This chapter (and Chapter 5) makes use of an auditory control prime (the word "nothing") to make the control and identical conditions comparable. There is no reason to expect the negative ISIs to differ significantly from the results in Chapter 2. The positive ISIs will reflect whether the identity prime affects planning even when the speaker does not have an opportunity to treat the prime as a repeated mention while planning the utterance. If it does, latency should be reduced. Duration will likely be reduced regardless, because of repetition priming.

3.5 Questions about consistency

Apart from the new control prime, the other major difference between this manipulation and the previous chapter is the use of consistent action blocks. In these, participants will have 48 trials in a row where the target object always moves in the same way (*viz*. rotates). This makes the task of describing the action easier, in a way similar to the identity prime. Although nothing during the trial itself indicates that the object will move identically, the manipulation and instructions still make the action fully predictable. The participant can keep the action word readied throughout the consistent block, and thus does not have to spend any additional time retrieving or otherwise processing it, beyond what is necessary to insert it alongside the target object word.

This manipulation gives rise to two new questions: 1) will the consistency manipulation affect latency and word duration at all? It should affect both, because speakers will have access to both the semantic and phonological information associated with the action word before they even begin to plan. This should make planning easier, and, like the identity prime, serves as a locus of repetition on the action word; 2) which segments of the utterance will it affect – latency, object word duration, action word duration, or some combination of these three? If action word consistency only facilitates the processing of the repeated word, through sheer repetition priming, only the action word should exhibit reduction. However, if the effect is more general, speakers may speed up the object word with the knowledge that they do not need to do as much processing on the upcoming action word.

3.6 Experiments 2a&b

Methods

Participants. 40 native English-speaking undergraduates from the University of North Carolina at Chapel Hill participated in each experiment for pay ($7.50 for half an hour), or participation credit in an introductory psychology class.

Materials. The picture stimuli, prime words, trial structure, and object movements were identical to those used in Experiments 1a&b.

Prime types. The identical, related, and unrelated conditions were the same as in previous experiments. The control condition was changed in this and subsequent experiments from a series of printed x's to a spoken word: "nothing." This brought the control prime in line with the other primes' modality.

Timing of prime. In these and the experiments in Chapter 5, only the 300 and 1500ms magnitude ISIs were used, not the 900ms. The 900 magnitude ISI was not retained for two reasons: the only novel result it showed was the -900ISI result on duration in Experiment 1, and including it would reduce the power that is instead going to be used to detect effects at the other ISIs.

List Design and Counterbalancing. In these experiments, participants saw each target picture a total of twice – once at one ISI and once at the other. Unbeknownst to participants, the experiment was divided into two halves, with all 96 items appearing once in the first half, with no repetition until the second half. Each item (i.e. object word) was paired with an ISI and condition in the same manner as in the previous experiments, for a total of 8 lists (two ISIs x four conditions). Also as in the previous experiments, ISIs were presented in blocks, with all 1500ISIs within a list (48 items) appearing before all 300ISIs in the same list (another 48 items), or the reverse in half of the lists. This marked the first half of the experiment. In the second half, items were presented in a list that paired them with the other ISI and a different condition. Parallel to all of this, half of each set of items that appeared at a particular ISI (24 from each for a total of 48) was presented in a block that consisted only of rotates actions, and the other half of each set was presented in a block that consisted of all four actions, as in the previous experiments. In the second half of the experiment, the halves were switched, so that each item appeared once in the consistent condition and once in the random condition for each participant. From the participants' perspective, the experiment consisted of 8 blocks of trials, each of which was homogeneous within itself with respect to consistency and ISI, but not with respect to condition, which was randomized (per the lists).

Procedure. The procedure was identical to the earlier experiments', up until the point where the practice trials had ended. At that point, participants were told by the experimenter that the experiment was divided into 8 blocks. They were told that each block would either be a rotates-only block or a random action block, and that they would be told by the program before each block began which type it was. This set of experiments took a total of 45 minutes, with 10 minutes for instructions and between 25 and 30 minutes for actual presentation. The program stopped after each block to allow the participant a moment to rest, and to give the action instructions for the upcoming block.

Analysis. The analysis was the same as in previous chapters, with the following exceptions. In this pair of experiments, only the first half of each participant's data was entered into the model, to avoid repetition effects. Further, instead of constructing a model for each ISI, one baseline model was constructed for each utterance segment (latency, object word duration, action word duration), each of which then had three simple comparisons and their interactions added: 300ISI vs. 1500ISI (ISI variable), consistent action vs. random action (consistency variable), and then either related vs. unrelated, or identical vs. control (condition variable). This allowed a single control model for each segment to test all of the relevant questions, including both of the most important condition comparisons (related vs. unrelated, and identical vs. control). When an interaction involved the ISI variable, planned comparisons were conducted within each ISI. These therefore included only half of the possible datapoints, but directly compared whether condition and consistency (as well as their interaction) were significant at each ISI. Note as well that just like in the previous chapters, the ISI variable in the model distinguishes only between the 1500 and the 300 ISI – positive and negative models were constructed separately.

Results

Summary statistics for each condition and ISI appear in Tables 8 through 11, broken down by segment. Parameter estimates for each of these effects appear in the Appendix, with the other parameter estimate tables. Main effects in this chapter (and Chapter 5) are reported as being in the absence of a significant interaction, unless the interaction is explicitly mentioned.

Related vs. Unrelated Models. The related condition elicited longer latencies than the unrelated condition, but only at negative ISIs (Figure 8). There were no other effects of relatedness on latency to begin speaking, including any interactions. Consistency (i.e. all actions were “rotates”) led to shorter latencies than inconsistency at negative ISIs, but not positive ISIs. See Figure 9.

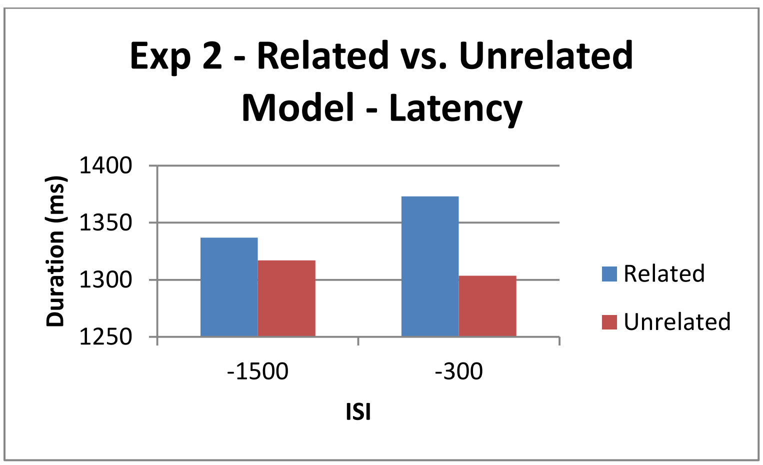

Figure 8: Graph showing the durations of the latency, broken down by ISI and condition. This illustrates a main effect of relatedness.

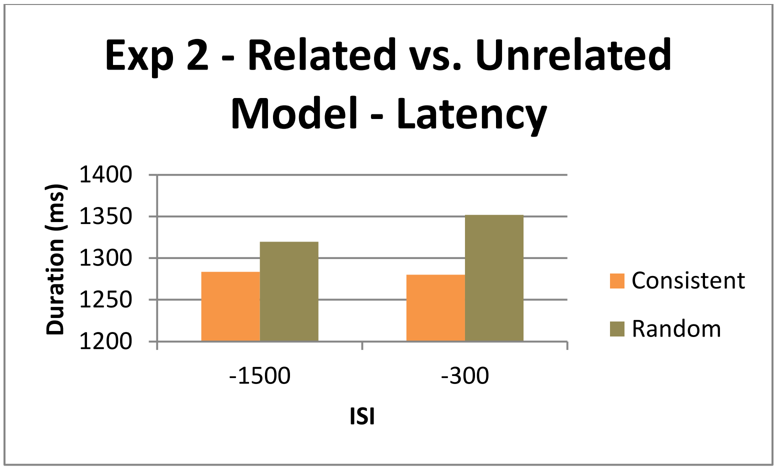

Figure 9: Graph showing the durations of the latency, broken down by ISI and consistency. This illustrates a main effect of consistency.

Object word durations did not exhibit any main effect of relatedness. Consistency led to shorter object word durations at negative but not positive ISIs (Figure 10). This was especially true in the related condition, as evidenced by a significant interaction between condition and consistency, with the shortest object word durations in the related and consistent levels. See Figure 11.

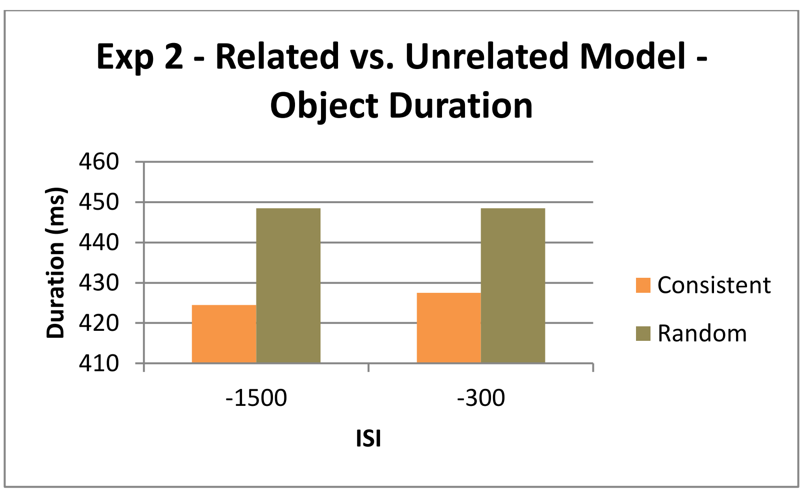

Figure 10: Graph showing the durations of the object word, broken down by ISI and consistency. This illustrates a main effect of consistency.

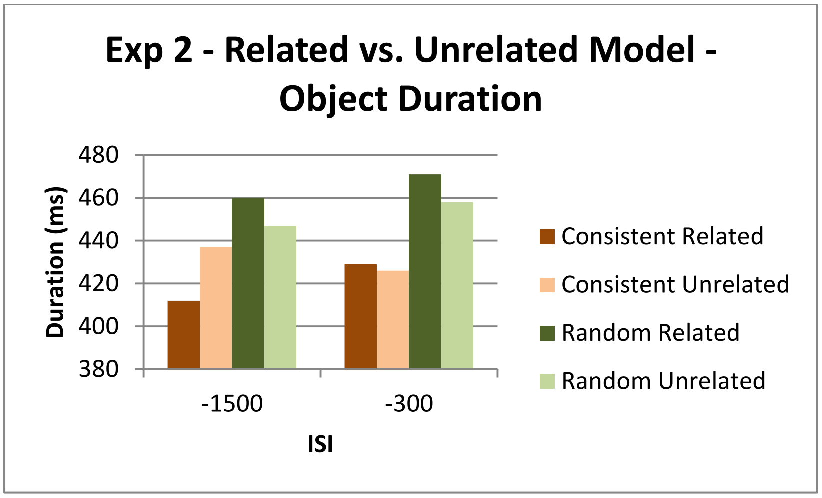

Figure 11: Graph showing the interaction between consistency and condition, with the shortest durations in the combined consistent and related condition (the orange-tinted bars).

Like object words, action words were unaffected by the relatedness manipulation. Action word durations were significantly shorter in consistent than inconsistent blocks at positive ISIs, but not negative ISIs. There was a significant interaction between consistency and ISI for the negative ISIs, but simple effects analyses did not reveal any significant consistency differences in either ISI alone (Figure 12).

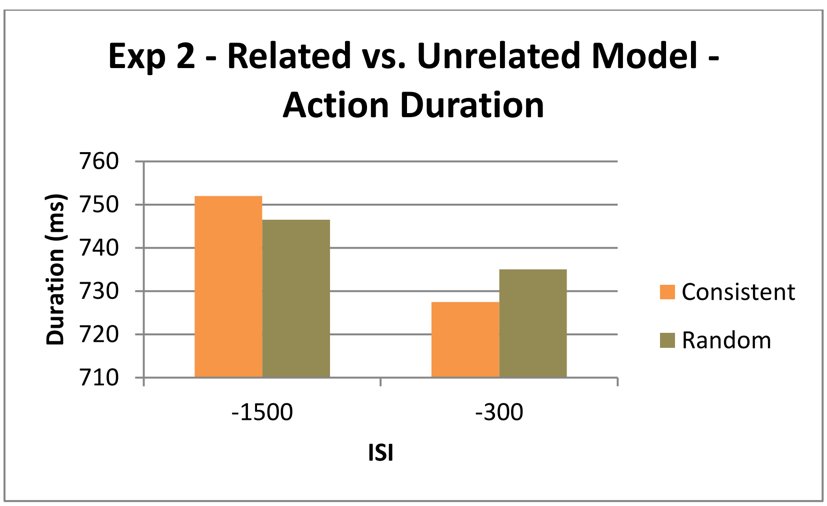

Figure 12: Graph showing the durations of the action word, broken down by ISI and consistency. This apparent interaction did not result in a significant simple effect of consistency at either ISI.

Identity vs. Control Models. The identical condition elicited shorter latencies than control at the negative ISIs but not positive. Consistency led to shorter latencies at all ISIs (Figure 13). At the positive ISIs, 1500ISI was shorter than 300ISI. At the negative ISIs, condition and consistency both interacted with ISI. A simple effects analysis is shown in Figure 14. It revealed significantly longer control latencies at -300 but not -1500, and consistency effects in both ISIs.

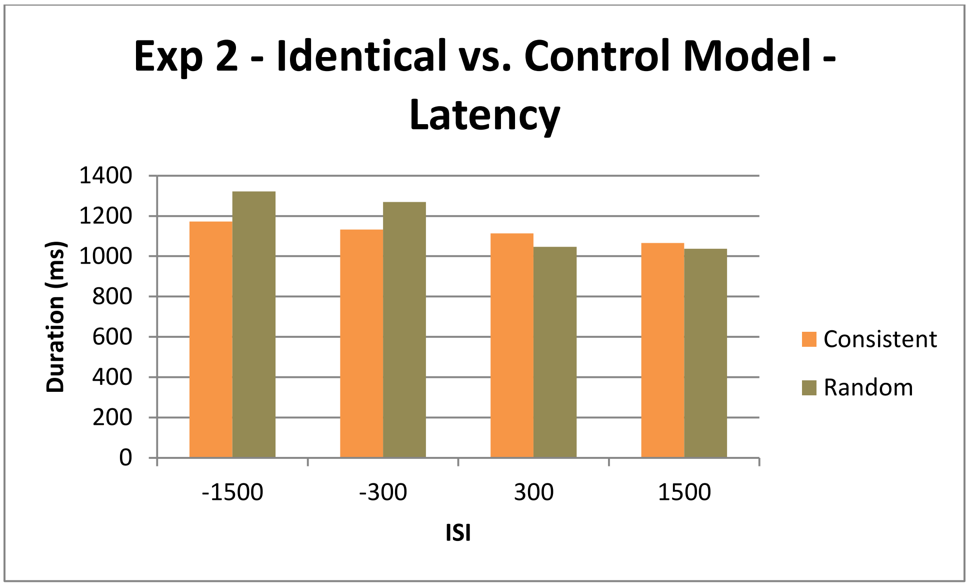

Figure 13: Graph showing the latency to begin speaking, broken down by ISI and consistency. This illustrates two main effects of consistency, one in the negative ISI model and one in the positive ISI model.

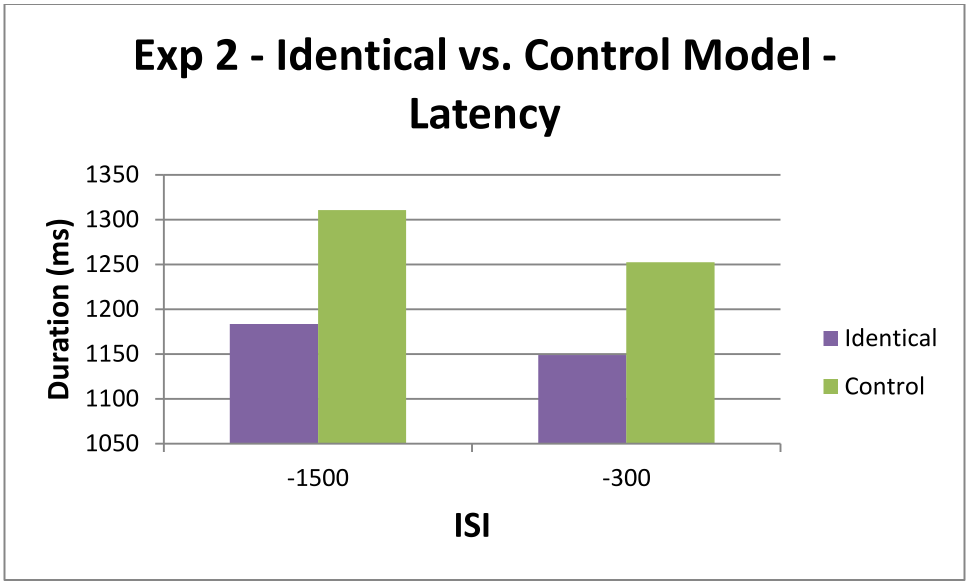

Figure 14: Graph showing the latency to begin speaking at -300ISI, broken down by condition. Here only the difference at -300ISI was significant, as explained in the text.

For the object words, identity priming led to shorter durations than control at all ISIs. This effect disappeared in a simple effects analysis that was conditioned on a significant interaction between condition and ISI, presumably due to halved power (Figure 15). Consistency led to shorter durations at negative ISIs, but longer durations at positive ISIs. At negative ISIs there was also a significant interaction between consistency and ISI. The simple effects analysis showed consistency effects at both -1500 and -300ISI, but also a marginally-significant interaction between condition and consistency at 300, which showed that the shortest durations were in the consistent and identical levels. A significant three-way interaction in the positive ISI model corroborated the shorter durations from identical priming that appeared in the general model. It also showed that consistency led to longer durations only at 300ISI. See Figure 16.

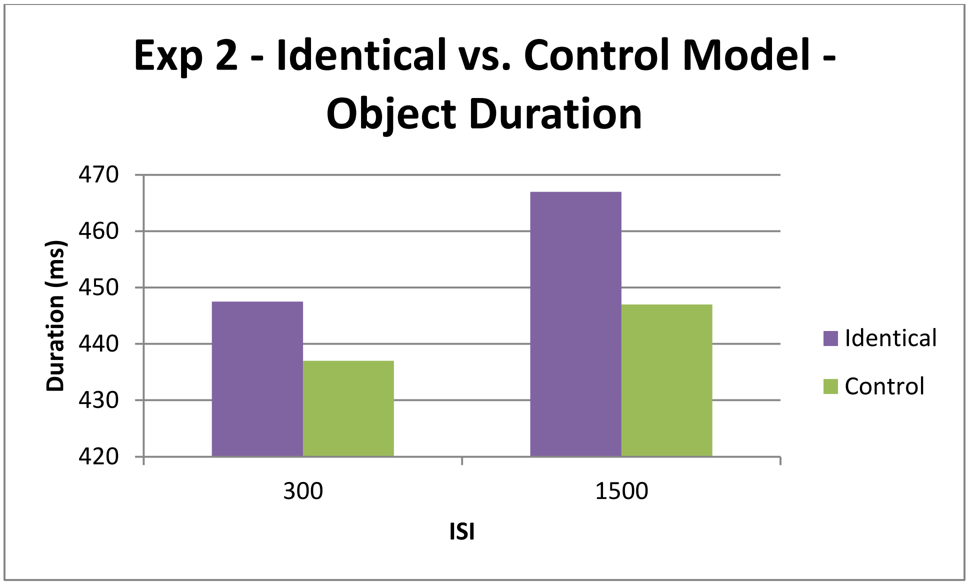

Figure 15: Graph showing the durations of the object word, broken down by ISI and condition. This illustrates a main effect of relatedness (identity).

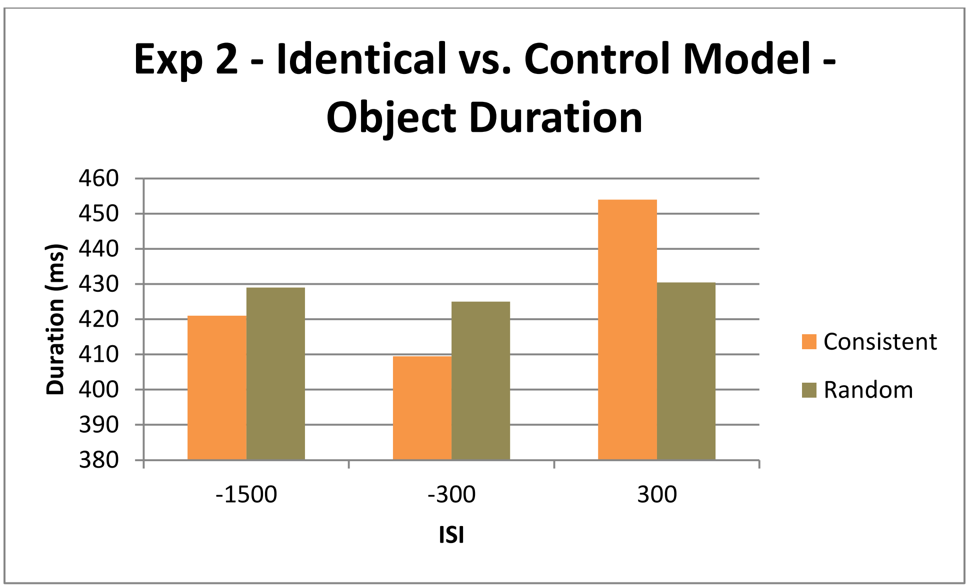

Figure 16: Graph showing the durations of the object word, broken down by ISI and consistency. All of these effects are significant, although not in the same direction. The negative ISIs had shorter duration in consistent than random trials, while 300ISI had longer duration in consistent than random trials.

For action words, the effect of consistency was significant in both the negative and positive models. ISI was a significant predictor for the positive ISIs, with longer durations at

1500ISI. For the negative ISIs, the three-way interaction among condition, consistency, and ISI was significant, and mitigated the effect of consistency at -1500ISI in the identical condition (see Figure 17). A simple effects analysis broken down by ISI showed that the consistency effect did not fully reach significance at either ISI. This is not conclusive, however, given the halved power, and caution is necessary in interpreting this difference.

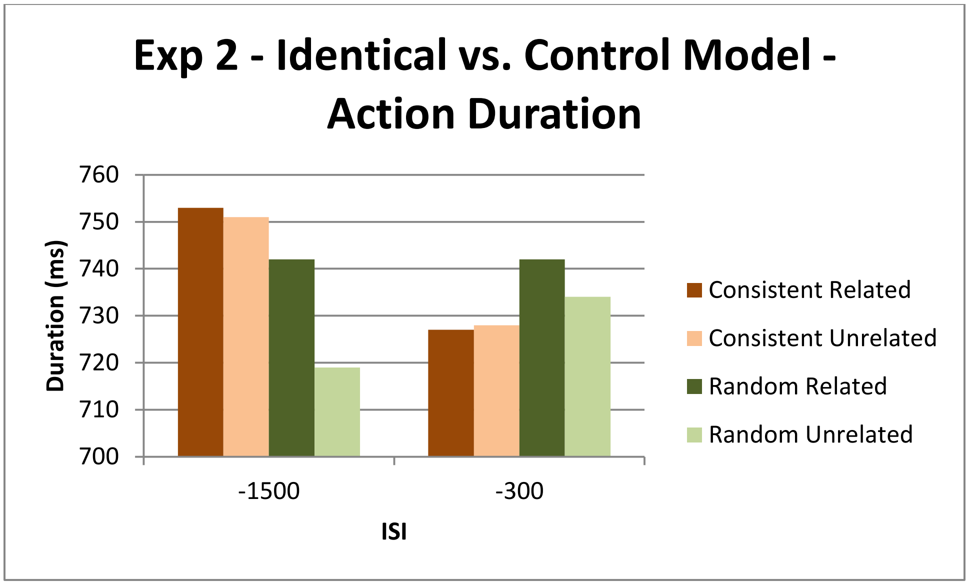

Figure 17: Graph showing the durations of the action word, broken down by condition, consistency, and ISI. The overall effect of consistency is significant, but the effect is mitigated at -1500ISI, as shown by the difference between the green and orange bars.

Table 8: Summary statistics for the onset latency, with standard deviations in parentheses. Experiment 2a is on the left, and Experiment 2b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	1234 (353)	1387 (343)	1195 (301)	1310 (312)	1075 (605)	1013 (573)	1022 (401)	975 (489)
Identical	1110 (283)	1257 (297)	1069 (289)	1229 (395)	1151 (287)	1079 (510)	1111 (299)	1099 (477)
Related	1301 (405)	1373 (393)	1301 (404)	1445 (468)	1223 (330)	1126 (473)	1214 (388)	1153 (524)
Unrelated	1266 (402)	1368 (401)	1259 (361)	1348 (407)	1192 (419)	1114 (509)	1218 (279)	1079 (532)

Table 9: Summary statistics for the determiner, with standard deviations in parentheses. Experiment 2a is on the left, and Experiment 2b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	132 (62)	132 (43)	134 (55)	134 (53)	156 (61)	130 (37)	139 (52)	148 (67)
Identical	126 (43)	124 (37)	123 (46)	128 (52)	132 (42)	128 (35)	129 (45)	123 (37)
Related	144 (65)	142 (63)	138 (75)	135 (60	128 (43)	125 (33)	136 (44)	125 (42)
Unrelated	141 (72)	144 (51)	133 (62)	132 (63)	132 (41)	131 (46)	135 (42)	134 (45)

Table 9: Summary statistics for the object word durations, with standard deviations in parentheses. Experiment 2a is on the left, and Experiment 2b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	434 (114)	450 (132)	398 (114)	433 (132)	445 (125)	429 (169)	454 (125)	440 (138)
Identical	408 (119)	437 (133)	421 (124)	413 (123)	463 (128)	432 (137)	471 (128)	463 (136)
Related	412 (111)	460 (155)	429 (130)	471 (128)	479 (135)	464 (142)	477 (135)	455 (155)
Unrelated	437 (123)	447 (123)	426 (129)	458 (133)	468 (128)	484 (144)	496 (128)	448 (144)

Table 10: Summary statistics for the action word durations, with standard deviations in parentheses. Experiment 2a is on the left, and Experiment 2b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	719 (109)	736 (127)	738 (104)	727 (123)	771 (141)	795 (163)	758 (147)	770 (138)
Identical	735 (104)	732 (128)	731 (105)	724 (128)	798 (117)	778 (145)	806 (112)	786 (133)
Related	753 (104)	742 (134)	727 (116)	742 (137)	780 (136)	815 (137)	804 (109)	796 (136)
Unrelated	751 (106)	719 (134)	728 (101)	734 (135)	797 (113)	786 (136)	805 (120)	785 (141)

Discussion

Questions About Relatedness. The same latency effects are on display here as in Chapter 2, namely longer latencies after related primes than unrelated. The difference appears at the same ISIs as well, both -300 and -1500ISI, but neither 300 nor 1500ISI. This shows continuity, and the same interpretation applies as before as well – speakers spent longer completing their utterance plan when they had difficulty accessing or selecting the relevant semantic information. The real question is why this happens, in terms of language production mechanisms.

Related primes increased latency more than unrelated primes at only negative ISIs, which is the same pattern as in Experiment 1. Together, the results from this experiment give good reason to think that speakers are engaged in similar processing at -300 and -1500ISI. The lengthening of planning time at -300ISI was most likely due to the association between the prime and target. Schriefers et al. (1990) claim that this association leads to increased difficulty in selecting the correct lexical item, and therefore increased planning time. Mahon, Costa, Peterson, Vargas and Caramazza (2007) offer an alternative, namely that speakers experience difficulty in selecting a response from a list of already-processed word representations. According to this explanation, speakers select and process the semantic information for both the target and the prime, and then engage in a response-exclusion process that discards the unusable prime. Both the associative-difficulty and the response-exclusion explanations could account for the increased latency after related primes here, and these experiments cannot distinguish between the two. Regardless of the details of the underlying process(es), semantically-related primes increase latency to begin speaking even for multi-word utterances. The relationship between the prime and the target is sufficient to lengthen planning time independently of the other processes that are involved in creating a fluent utterance.

Related primes also increased latency when they occurred at -1500ISI, an unprecedentedly large amount of time. More importantly, both negative ISIs produced the same pattern of results, suggesting that speakers engaged in similar processing at both intervals. The nearest precedent for this effect comes from Balota et al., (1989) who found lengthened latencies to begin speaking at -600ms while measuring out to 1200ms, both positive and negative. The lack of effect at larger SOAs may simply be a function of experimental design – their participants uttered a single word, and in response to a written as opposed to pictured stimulus. Starreveld and La Heij (1996) argued that part of the difference between their own findings and the outcome of Schriefers et al., (1990) stemmed from their written as opposed to pictured stimuli. They attributed this to the comprehension system activating phonological information through the orthographic information, and the production system capitalizing on the already-available phonological information. The difference between the current results and Balota et al., (1989) could similarly lie in the fact that written stimuli activate orthographic and possibly also phonological information. If that were the case, the production system would require less involvement of semantic processing, with the consequence that semantic relatedness had a greater impact on processing in the current study than it had in the others.

Three types of explanation are possible here for why ISIs elicited responses similar in character to the ones that small negative SOAs have elicited in other studies. All rely on the core processing assumption that the prime and target's semantic representations compete with each other, which lengthens the time it takes speakers to complete the utterance plan. With that assumption, the data here tell the same story as data from previous PWI work. Although an unlikely explanation because of the low predictability of the relationship between prime and target, one of the reasons that speakers might still experience processing difficulty at -1500ISI is that they actively kept the prime around by rehearsing it until the target appeared. If they engaged in such a strategy, speakers could maintain enough activation on the prime that it would continue to interfere with selection and/or retrieval of the target during utterance planning. The utility of this approach would also extend indefinitely backward in time, which could explain why speakers experienced semantic inhibition difficulty while planning the targets in the -900 and -1500ISI conditions. Similarly, it could explain the positive ISI data – speakers had no opportunity to rehearse anything when the prime appeared after the target, and thus did not experience any competition. However, there is no good reason for speakers to engage in an explicit rehearsal process, because even though the primes sometimes exhibited relationships with the targets, on only 25% of trials (the identical trials) was that relationship predictive of the required response.

Another straightforward explanation, and a similar one, is that multi-word utterances require enough additional planning, above and beyond single-word utterances. For example, following Bock (1982) and Levelt (1989), additional syntactic processing is necessary for multi-word utterances compared to single-word ones. Speakers must also select the rhetorical and information status roles of each of the entities in the utterance, and generate connecting grammatical material, such as the determiner and inflectional information (whether to pluralize the action word). The production system has to make sure that each activated/selected lexical entry appears in the correct syntactic position, with the correct prosody, and appropriately integrated phonologically with its neighbors. All of this processing has to be completed within a reasonable timeframe in order to maintain fluency and/or hold the conversational floor. In these experiments with ISIs, speakers would try to complete as many of the requisite processes as possible ahead of time, including activating the syntactic frame, and slotting in the prime as a potential candidate for the target word. When the target finally appears, because the prime has been slotted in as a candidate for the object word, the speaker must take some time to remove it. The greater semantic similarity between the prime and target in the related condition makes this slightly more difficult, and therefore more time-consuming.

Another possible explanation for the semantic inhibition at the negative ISIs is that the semantic system automatically retained information about the primes until the targets appeared, without any conscious effort on the speakers' part. This explanation would require that the production system held on to semantic information from the prime for on the order of 2 seconds in the -1500ISI. Semantic priming has been shown to persist across hundreds of trials (Munson & Solomon, 2004), and the same pattern is true for syntactic priming (Bock & Griffin, 2000). Importantly, both of those studies included many filler trials, which would presumably have introduced enough noise into the production system that any activation which would have led to priming would have to have been significant and very persistent. That contrasts with the situation in the current study, which presented no information between the prime and target. It is thus entirely possible that the production system simply let the semantic information from the prime reverberate, perhaps on the off chance that it would be helpful for identical trials. Persistent activation of this sort could easily explain semantic inhibition, via the same competition mechanism as in prior work (Schriefers et al., 1990).

Regardless of whether one of the explanations offered here is correct, or whether some other mechanism explains these effects, the manipulations in Experiments 1 and 2 elicited increased latencies after related primes at negative ISIs. The most plausible account for this effect is semantic competition, and I have attempted to also account for the latency between prime and target introduced by the ISIs. Assuming one or more of these explanations holds, the current results corroborate prior PWI work that has shown semantic inhibition, and provide additional evidence in favor of lengthened latencies to begin speaking as a function of difficulty of semantic processing.

Unlike latency, and Experiment 1, Experiment 2 showed no effect of relatedness on duration. However, this result is consistent with the results from the ISIs that were also present in the earlier experiments, namely -300 and -1500. The null outcome here strengthens the conclusion from Chapter 2 that durational lengthening failed to appear not due to lack of power, but that the effect simply does not exist at these ISIs. The durational lengthening at -900ISI in Experiment 1 raises the question of whether 900ms is a special number, or whether some time window in that region would allow semantic relatedness to affect duration reliably. If so, it would rule out any of the explanations offered above that treated -300 and -1500 as performing similar processes because they work along the same continuum. There is no good reason that speakers should have engaged in the same processing at a relatively early time point and a relatively late time point, but not at a time point in between, with the exception that 900ms is special in some way.

One possible explanation is that 900ms is long enough that the speaker processes the prime to the point that it competes all the way to the point of articulation. In other words, the prime might arrive during a time window when the speaker is seriously entertaining it as a candidate for insertion into the utterance, and that candidacy persists through to articulation. The system would have to assess or retain the fact that it experienced difficulty during the planning stage of processing and use that as evidence that it needed to slow down a difficult-to-process word (cf. Bell et al., 2009). In that case, the production system might slow down its production of the word, either because it needs to buy time or because it recognizes that it had serious difficulty with completing processing the competition between the two. Levelt et al., (1999), in their model of lexical access, note that occasionally it appears to be the case that two lexical items get selected and fully processed during production. The competition to determine what actually gets said is resolved only at the final stage(s) of production, where a final monitoring process catches slips of the tongue, curse words, or in this case, words with serious semantic competitors.

The other notable property of the semantic inhibition effect on word duration at -900ISI is that it appears with a marginally-significant latency lengthening, but the other lengthened latencies at negative ISIs do not exhibit similar (or any) duration effects. This suggests that the length of the plan that precedes utterance onset does not entirely determine utterance duration, and vice versa. If it did, all of the negative ISIs would have led to durational lengthening as well. This observation has the implication that at least some of the processing that influences a word's duration occurs after utterance initiation, according to the logic of the PWI paradigm. According to Levelt et al., (1999), selection of the lexical item signals the end of the planning phase in single-word production. This is why they attribute the latency differences in Schriefers et al., (1990) to differences in planning, namely competition at the pre-selection portion of semantic processing. If the results here are not fully attributable to differences in this part of processing, that means that at least some of the post-selection processes of production influence word duration.

There are no existing provisions in the model for specifying the duration of a word. Nor is the model designed to deal with the complexity of producing a multi-word utterance, which is not the same as stringing together several independent single-word models. The conclusion from the previous paragraph may thus be too strong. If the planning of a multi-word utterance differs significantly from the planning of a single-word utterance (for example, that it requires multiple/parallel lexical access, or storing number agreement information about a word, or resyllabification/coarticulation), the planning phase could still include some process(es) that influence word duration. However, it is not entirely clear why the difference in duration would appear in tandem with a difference in latency at -900ISI, but not the other negative ISIs. If some process in the planning phase exerted an influence on duration in this dataset, that influence is obscured.

On the other hand, other work in a PWI-like paradigm has offered a post-planning processing explanation for the influence of semantic processing fluency on word duration. Like Experiment 1, Balota et al., (1989) found differences in word duration after semantically-related primes, which they attributed to the speed of phoneme activation, referencing Dell (1986)'s connectionist model of production. The model itself makes no specific claims about chronometrics, so it is not fully clear how to delineate its processing into planning and post-planning stages, and thus this explanation is subject to the same ambiguity described in the previous paragraph. The notable property of the Dell (1986) model that distinguishes it from the Levelt et al., (1999) model, however, is that semantic processing variation (e.g. how quickly or strongly a lexical item is selected compared to its competitors) is allowed to influence phonological processing. This means that semantic competition could lead directly to speed of phonological selection, and at a post-lexical-selection stage of processing. Per the definition of lexical selection as the planning stage (Levelt et al., 1999), post-planning processes have an influence on word duration. As Balota et al. note, however, the Dell model does not include the assumption of speeded phonological selection as part of its architecture – they added it to explain their data. In a reasonable range of connectionist-style architecture, however, it is a justifiable assumption.

Whatever the processes, the latency outcome shows a clear semantic inhibition effect. The duration outcome will require further investigation. I take this issue up again in the final chapter, after presenting the phonological results.

Questions about identity priming. The use of an auditory control prime here licenses a more conclusive interpretation of the effects of the identity prime, especially on latency. Once again, identity elicited shorter latencies and object words at negative ISIs. Importantly, at positive ISIs it had the same effect on duration, but no significant effect on latency. This suggests two things: 1) in the previous experiment, participants were indeed using the control condition as a cue to respond, making the comparison between identical and control conditions in those experiments infelicitous for current purposes; 2) only at negative ISIs is the identical prime easing planning to a greater degree than the control prime, given the null effect on latency at positive ISIs. The fact that durational reduction occurred without a corresponding latency reduction, however, suggests that planning variation is not a necessary precursor to durational variation. The phonological experiments have useful data to contribute to this discussion, and so I put off further speculation about the underlying processes until Chapter 5.

Questions about consistency. The consistency manipulation tended, although not uniformly, to reduce latency and action word duration in these experiments. This makes sense, because similar to an identity prime, knowing that the same action would repeat on a trial should have made it easier for speakers to plan the utterance. It also made processing the action word itself easier, as evidenced by the shorter durations. Planning the utterance and reducing the duration of its words seemed to operate on different processes, however. At the negative ISIs, consistency reduced planning latency and object word duration, suggesting that speakers used the processing facilitation they experienced on the action word to speed their production of the object word. This is consistent with corpus analyses that show that speakers lengthen words when upcoming production is difficult, and shorten them when upcoming production is easy (Bell et al., 2009; Christodoulou, 2012; Jaeger, 2010). Having the action word ahead of time meant that in spite of the semantic relatedness, speakers had an easier time integrating the action word into the ongoing production of the utterance.

Semantic relatedness did not have much of an influence on word duration in this experiment. Semantic competition did not propagate forward to word duration – the difficulty in selecting the appropriate lexical item ended there, and did not influence how the phonetic information was encoded. However, being able to select the lexical information for “rotates,” and to also prepare the phonological information, eased the production of the action word. The production system appears to have looked ahead, saw that the remainder of the utterance would prove simple to finish, and reduced the duration of the object word. This provides experimental corroboration of the statistical corpus findings. It also coheres well with prior findings that suggest that speakers engage in some planning of upcoming words while they are already producing speech (Christodoulou, 2012). Further, it suggests that speakers completed some processing of every word in the utterance before they began speaking, but also that they modulated the duration of individual words as a function of the difficulty of processing the surrounding words.

The same pattern of results did not appear at the positive ISIs. Latencies were not reduced, and neither were object word durations. Action words, however, were, suggesting that some other pattern of processing transpired. Speakers knew for certain in consistent blocks that the target objects would rotate, and so could ready the representations associated with the action word confident that that pre-processing would pay off. They did not know whether the prime would have any relationship with the target, and so had no need to integrate that representation into the full formulation of the utterance. The exception to this was the model for the identical and control primes, which had a significant effect of consistency on the latency to begin speaking. In other words, speakers had an easy time planning when the prime and target matched, but only then, and only when the action word allowed for the planning of the entire utterance to be easy.

Consistency will also be investigated in Chapter 5, and has slightly different effects there. Further discussion is thus saved until then. The next chapter returns to the experimental design of the previous chapter, with no consistency manipulation.

Chapter 4: Experiments 3a&b

4.1 Introduction

The central question has the same general form in this chapter, namely how ease of processing affects word duration, but the focus shifts from lexical to phonological-centered processing. In Chapters 2 and 3, the specific interest was whether difficulty experienced during lexical retrieval translated to lengthening of word duration in the subsequently-produced utterance. Here, instead of a semantic alternative, the manipulation introduces partly- but not completely-overlapping phonological material. Similar to the semantic competitors in the previous two chapters, this manipulation is designed to interact with ongoing phonological processing. Unlike those chapters, however, the processing relationship between the prime and target depends less obviously on competition, because they share some but not all information. Will the shared information make it easier to plan and execute the utterance, or will the unshared information at the end of the word make planning and execution more difficult? Does the timing of any influencing information matter? The PWI paradigm once again provides a useful method for answering these questions, and the experimental design is identical to Chapter 2.

4.2 Relationship with prior work

Like the semantic interference effect, prior work shows that in single-word naming, phonological facilitation appears only at certain time points. However, there is ongoing debate about exactly what those time points are. Schriefers et al., (1999) found a strict separation between the semantic interference effect and the phonological facilitation effect. The former occurred only at small negative SOAs, whereas the latter occurred only at small positive SOAs. However, Starreveld and La Heij (1996) used a similar paradigm and found that the phonological facilitation effect was slightly more widespread, extending into small negative SOAs as well. They explained the difference in terms of the prime stimulus, which was written instead of auditory. The prime was taken to activate orthographic representations, and to spread this activation upward to the phonological system. The speed of orthographic uptake, and activation spread to phonology, widened the window of phonological facilitation.

As in Chapter 2, the discussion here will benefit from confirming that this manipulation also elicits phonological facilitation on latency, because that strengthens my assumption that it taps into some of the same processes as prior studies. One guiding question for this chapter is thus the extent to which the results in the experiments below cohere theoretically with the results of previous studies. Despite the slight differences in timing, both of the PWI studies discussed in the previous paragraph found facilitation from the phonological primes, in the form of reduced latencies to begin speaking, and in general, phonological effects appeared after semantic effects. Because of the closer similarity between the current paradigm's auditory primes and the primes in the Schriefers et al. study (as opposed to the written primes in Starreveld and La Heij, 1996), the results here should more closely resemble Schriefers et al., and thus only positive ISIs should exhibit an effect of phonological facilitation on latency.

4.3 Questions About Relatedness

Like semantic relatedness, there are no PWI studies that systematically investigate the relationship between phonological processing and word duration, to my knowledge (although see Baese-berk & Goldrick, 2009; Watson & Buxo-Lugo, to appear). The question for this chapter is how the production system will react to phonologically-related primes in how it produces duration. It is possible, for example, that having already-activated phonemes will lead to shorter object word durations, analogous to the logic that explains why having already-activated phonemes leads to reduction in the identical condition. It is also possible that having part of the word form activated or processed ahead of time, but not the whole, ultimately-conflicting word form, will lead the system down a garden path. In this case, it would need additional processing time to ensure that the correct word form was actually produced, instead of the related one. This would lead to extended word duration, and is similar (i.e., follows a similar logic) to the observed longer duration of words with dense phonological neighborhoods (Gahl, 2008; Munson, 2001, *inter alia*).

The theoretical mini-structure in the previous paragraph leads to the expectation that latency and duration results should appear at the same SOAs. Reduced planning time should correlate with reduced duration, because the system should respond to the already-active phonemes in both outcome variables. The extent to which the results deviate from this prediction would show that duration is a phenomenon that does not emerge from a single process or time point, but rather that it emerges as part of several of the standard operations of the system.

4.4 Questions About identity priming

Questions about the relationship between the identical and control condition remain the same as well, with the additional question of whether the effects in this pair of experiments appear at the same time as the effects did in Chapter 2. The related condition differs between this chapter and that one, but all three of the other conditions – control, identical, and unrelated – are the same. Because the relationship between the two conditions is entirely identical between the two experiments, there is no reason to expect any difference.

4.5 Experiments 3a&b

Methods

Participants. 36 native English-speaking undergraduates from the University of North Carolina at Chapel Hill participated in each experiment for pay ($7.50 for half an hour), or participation credit in an introductory psychology class. 1 participant was excluded from Experiment 3b due to technical difficulties with the recording.

Materials. The picture stimuli and list structure were identical to those used in Experiments 1a&b.

For the related condition, phonologically-related primes were chosen based on three major criteria: 1) they had an identical (as near as was possible) relationship with the onset phoneme of the target word; 2) they had the same prosodic structure as the target word; 3) they did not overlap phonologically with any part of the word after the onset. Analogously to Experiment 1, the related primes also did not exhibit a semantic relationship with the picture, nor a phonological relationship with another picture in the stimulus set, to the extent that this was possible.

Procedure, and Analysis. These were the same as in Experiments 1a&b.

Results

Summary statistics for each condition and ISI appear in Tables 12 through 15, broken down by segment. Relatedness led to longer object word durations at every ISI, at least marginally. Similarly, identity led to shorter object word durations at all ISIs except -1500ISI. The only effects of relatedness on the latency to begin speaking were at -1500 and -300ISI, with no effects at the positive ISIs. The effects of identity on latency were presented at all ISIs, although in the opposite direction from expectation at positive ISIs. Parameter estimates for each of these effects appear in Tables 5 and 6, and detailed information about the models appears in the Appendix.

Table 12: Onset latencies and standard deviations (in parentheses) for Experiments 3 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	2326 (356)	2346 (433)	2318 (435)	1472 (525)	1389 (438)	1418 (425)
Identical	1870 (322)	1864 (327)	1842 (299)	1781 (421)	1799 (445)	1829 (474)
Related	1877 (335)	1914 (397)	1898 (349)	1783 (450)	1840 (480)	1853 (457)
Unrelated	1962 (418)	1938 (384)	1941 (362)	1780 (462)	1780 (462)	1872 (463)

Table 13: Determiner durations and standard deviations (in parentheses) for Experiments 3 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500

Control	111 (33)	106 (34)	109 (38)	102 (32)	106 (31)	112 (32)
Identical	105 (34)	101 (36)	101 (31)	108 (31)	110 (30)	110 (30)
Related	111 (39)	108 (33)	107 (30)	112 (34)	114 (34)	115 (33)
Unrelated	110 (37)	108 (35)	111 (40)	111 (31)	111 (33)	116 (34)

Table 11: Object word durations and standard deviations (in parentheses) for Experiments 3 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	431 (127)	423 (122)	435 (126)	438 (120)	441 (112)	439 (117)
Identical	417 (131)	404 (114)	414 (121)	428 (109)	429 (121)	436 (111)
Related	464 (142)	473 (127)	467 (132)	461 (133)	463 (117)	473 (122)

Unrelated	437 (127)	455 (130)	451 (131)	446 (115)	455 (130)	459 (122)

Table 12: Action word durations and standard deviations (in parentheses) for Experiments 3 a&b, broken down by condition and ISI.

	-1500	-900	-300	300	900	1500
Control	702 (138)	692 (138)	702 (136)	713 (121)	724 (120)	725 (115)
Identical	675 (149)	678 (136)	688 (149)	724 (122)	729 (121)	749 (123)
Related	692 (139)	702 (131)	698 (139)	719 (106)	724 (113)	740 (119)
Unrelated	700 (139)	695 (127)	716 (139)	729 (122)	728 (121)	736 (114)

Latency. The effect of relatedness appeared only at ISIs -300 and -1500, with shorter latencies in response to related primes. Note that this is the opposite pattern from Chapter 2, which had longer latencies after related primes. Identity priming appeared at every ISI, but with shorter identity than control latencies at negative ISIs, and the reverse at positive ISIs. Like in Chapter 2, the interpretation of these differences is compromised by the disparate modalities of the control and identical primes.

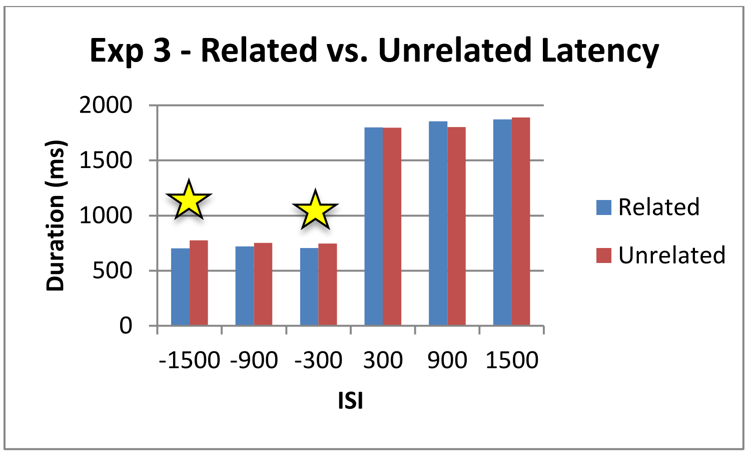

Figure 18: Onset latencies for related vs. unrelated comparisons, with yellow stars to indicate comparisons that were significant (p < .05).

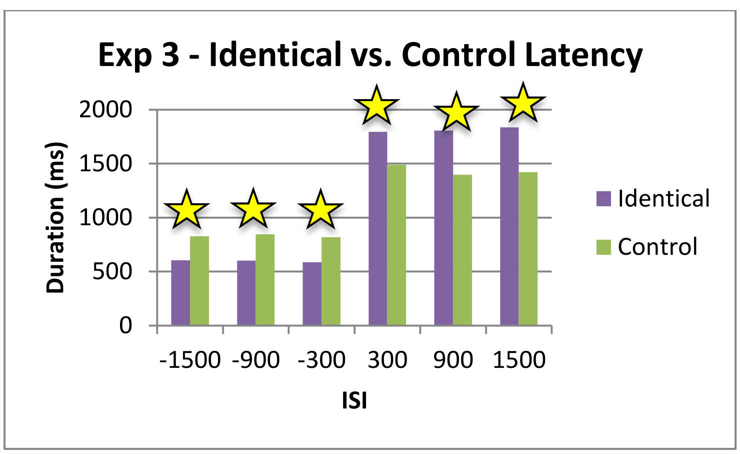

Figure 19: Onset latencies for identical vs. control comparisons, with yellow stars to indicate comparisons that were significant (p < .05).

Object Word Duration. Of particular interest is the breadth and direction of the relatedness effects on duration. The related condition was longer or marginally longer at every ISI, with significant effects at all ISIs except 300 and 900. The effect of identity priming was slightly less widespread, with a null effect at -1500ISI, but significantly shorter durations in the identical condition at all other ISIs except 900 (which was marginal).

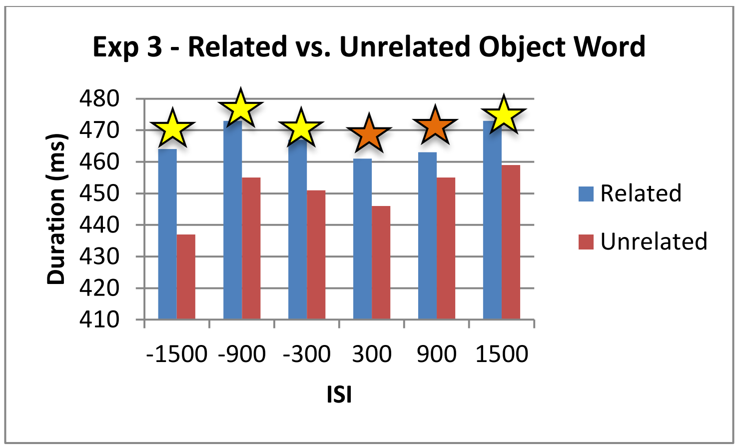

Figure 20: Object word durations for the related vs. unrelated comparisons, with yellow stars to indicate comparisons that were significant (*p* < .05), and orange stars to indicate comparisons that were marginally significant (.05 < *p* < .01).

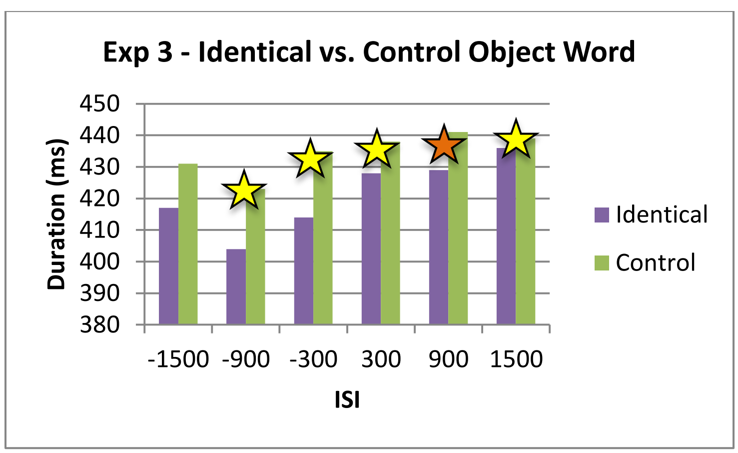

Figure 21: Object word durations for identical vs. control comparisons, with yellow stars to indicate comparisons that were significant (*p* < .05), and orange stars to indicate comparisons that were marginally significant (.05 < *p* < .01).

Table 13: The collection of parameter estimates, including *p*-values, for Experiment 3a. Critical comparisons are in the leftmost four columns, and non-critical comparisons in all others. Non-significant estimates are indicated with two dashes (--).

		Critical Comparisons				Non-critical comparisons							
Segment	ISI	Identical vs. control		Unrelated vs. related		Identical vs. Unrelated		Identical vs. Related		Unrelated vs. Control		Related vs. Control	
		z	p	z	p	Z	p	z	p	z	p	z	p
Latency	-1500	-19.434	0.001	3.449	0.002	2.72	0.007	--	--	-15.413	0.001	-19.267	0.001
	-900	-19.5	0.0001	--	--	--	--	--	--	-16.63	0.0001	-17.677	0.0001
	-300	-19.569	0.001	2.522	0.015	2.779	0.015	--	--	-16	0.001	-18.413	0.001
Determiner	-1500	--	--	--	--	--	--	--	--	--	--	--	--
	-900	--	--	--	--	2.89	0.02	2.084	0.09	2.178	0.07	--	--
	-300	--	--	--	--	--	--	--	--	--	--	--	--
Object	-1500	--	--	-2.384	0.003	3.38	0.002	6.027	0.001	2.72	0.007	4.732	0.001
	-900	--	--	-5.737	0.001	2.814	0.014	8.967	0.001	--	--	6.38	0.001
	-300	--	--	2.341	0.038	4.675	0.001	7.167	0.001	3.124	0.0017	5.039	0.001
Action	-1500	-2.965	0.016	--	--	2.154	0.079	--	--	--	--	-2.286	0.058
	-900	--	--	--	--	--	--	--	--	--	--	--	--
	-300	--	--	--	--	2.569	0.05	--	--	--	--	--	--

Table 14: The collection of parameter estimates, including *p*-values, for Experiment 3b. Critical comparisons are in the leftmost four columns, and non-critical comparisons in all others. Non-significant estimates are indicated with two dashes (--).

		Critical Comparisons				Non-critical comparisons							
Segment	ISI	Identical vs. control		Unrelated vs. related		Identical vs. Unrelated		Identical vs. Related		Unrelated vs. Control		Related vs. Control	
		z	p	z	p	z	p	z	p	z	p	z	p
Latency	-1500	14.534	0.0001	--	--	--	--	--	--	14.304	0.0001	13.727	0.0001
	-900	13.598	0.0001	--	--	--	--	--	--	12.077	0.0001	13.548	0.0001
	-300	11.065	0.0001	--	--	--	--	--	--	10.185	0.0001	9.464	0.0001
Determiner	-1500	--	--	--	--	--	--	--	--	--	--	--	--
	-900	--	--	--	--	--	--	2.535	0.054	--	--	2.313	0.054
	-300	2.15	0.062	--	--	1.914	0.062	2.345	0.05	3.875	0.001	4.271	0.001
Object	-1500	1.93	0.05	-2.69	0.01	6.64	0.001	3.83	0.001	4.18	0.001	--	--
	-900	-2.18	0.05	-2.23	0.05	4.592	0.001	7	0.001	2.1	0.05	4.15	0.001
	-300	-2.717	0.013	--	--	4.809	0.001	3.3	0.003	--	--	--	--
Action	-1500	7.655	0.0001	--	--	--	--	--	--	6.564	0.0001	7.004	0.0001
	-900	5.058	1.00E-05	--	--	--	--	--	--	5.467	1.00E-05	5.556	1.00E-05
	-300	5.861	0.0001	--	--	--	--	--	--	6.406	0.0001	5.429	0.0001

Discussion

Relationship with Prior Work. Two types of results would have been consistent with prior studies using the PWI methodology: latency effects at only positive ISIs (following Schriefers et al., 1990), or latency effects at positive and very slightly negative ISIs (following Starreveld & La Heij, 1996). Neither pattern appeared in these data, nor in fact were there any latency effects at positive ISIs from the relatedness comparison at all. However, -300 and -1500ISI exhibited latency effects, and in the expected direction, namely with shorter latencies after related than unrelated primes. In a multi-word naming study, Jescheniak et al., (2003) found facilitation, then a null effect, then inhibition on latency at 0, 150, and 300SOA, respectively, so the current results are not without precedent. Phonological primes here appeared to affect latency in a graded fashion depending on ISI. It is worth noting that like both the Schriefers et al., (1990) and the Starreveld & La Heij (1996) studies, the interpretation of the results here lies partly in relating them explicitly to the timing of the semantic prime results, not the absolute value of the SOA or ISI. In that sense, the results are more consistent with Jescheniak et al., (2003), but can still serve as relevant data for looking at phonological processing alongside other PWI work.

At least two explanations could account for the difference between this set of results and Schriefers et al., and Starreveld & La Heij's findings. The first is that the use of a multi-word utterance changes production processes sufficiently that positive SOAs were not ever going to exhibit any relatedness differences. Single-word utterances do not require any form of phonological integration, like coarticulation or resyllabification, where adjacent words take on phonological aspects of each other as part of pre-articulation processing. Speakers' having to engage in integration here may have preempted the facilitation effects seen in single-word studies, because they may have needed additional processing time to perform the integration or assess its appropriateness/success. On this explanation, this outcome represents an alternative but still true pattern of effects, and the difference from earlier work arises from the use of a more-complex utterance. The second explanation is that in a multi-word utterance, the object word duration exhibits the phonological processing difficulty instead of the latency, as evidenced by the marginal-to-significant effects on duration. Chapter 5 shows a similar pattern of effects to those presented above, and so I postpone further discussion until then.

Questions about Competition. There are two important competition findings in this pair of experiments. First, the related prime had a nearly ubiquitous lengthening effect on duration, compared to the unrelated prime. Two aspects of this effect are novel: how widespread it was, and its direction. Compared to the latency results, the duration results are almost uniform. This suggests that the speakers were responding to the prime words slightly differently in latencies than in durations when there was partial phonological overlap with the target. One possible explanation is that the overlapping portion served to highlight the non-overlapping portion (i.e. the shared onsets highlighted the unshared offsets), which called for a more emphatic pronunciation of the target, in order to distinguish it from the prime. Interestingly, the corresponding lack of a latency effect where many of the duration effects appear suggests that the effect was not a pre-planned or intentional emphasis. Rather, it emerged as the speaker engaged in whichever processes were responsible for producing the target word online. Another related possibility is that the early activation of the prime's phoneme created confusion during phonological integration (e.g. coarticulation or syllable preparation; Levelt, 1989; Levelt & Wheeldon, 1998). Having the onset phoneme should have facilitated processing, but having it activated in conjunction with a competing phoneme (the offset phoneme) would require restructuring the phonemic output that was intended for the articulators. The experiments in the next chapter show similar effects at the positive SOAs, and the discussion there will go into further detail about possible processing mechanisms.

The related primes also shortened latency relative to the unrelated primes, but only at -1500 and -300ISI. In Chapter 2, the semantic inhibition effect on latency appeared not only at -300ISI, but also -1500ISI, and marginally at -900ISI. I argued there that speakers may have rehearsed their responses, keeping some of the prime's information active, or it might have remained active simply from being slotted in as a potential candidate for production. These explanations could also account for the phonological data, but with slightly more difficulty. Phonological codes do not stay active as long as semantic representations, due to their having to pass more rapidly through the production system. A single word can have many phonemes, each of which has to become active at and only at the correct time. To accomplish this, the production system has to rapidly activate then rapidly deactivate phonemes, meaning that any phonological information that becomes active is not likely to remain that way for long. To affect latency or duration at these ISIs, the phonological information would have had to remain active for longer than would be expected under normal conversational circumstances. I am not aware of any data similar to Munson & Solomon (2004) that looked at phonological priming (instead of semantic) at long delays, but if the production system exhibited that behavior, it would explain these data as well.

An alternative explanation is that the shared phonological portion of the prime makes planning it slightly easier, but the unshared portion causes the speaker to lengthen the target object word's duration, for one or two reasons. The first part of this explanation is familiar by now from earlier studies – phonological overlap facilitates utterance plans. The second part has its root in either intentional prosodic emphasis based in planning or monitoring, processing-based delay, or both. Speakers can intend to emphasize words in order to distinguish them from competitors (Selkirk, 1995), which is a decision that occurs at the message-planning phase of processing (Levelt, 1989). Assuming that this decision takes time, it is a somewhat unlikely explanation, given the phonological facilitation effects. However, the production system also appears to have a monitoring mechanism or mechanisms that come into play immediately before articulation and ensure that the speaker does not accidentally produce unwarranted swear words, or that earlier processing hasn't accidentally mixed up a word or a phoneme (Levelt, 1989). Such a mechanism could see the only-partial resemblance between the prime and target and lengthen the duration of the object word to highlight the difference. A similar but slightly mechanically different way of reaching the same behavioral outcome is if the system has to take some additional time sorting out whether it has combined the onset and offset phonemes properly. This could directly result in delay – in other words lengthening – on the object word, or the lengthening could result from the system assessing the difficulty and slowing down to accommodate it, akin to the lengthening mechanism described by Bell et al., (2009). The current data do not distinguish between these two explanations, and were not designed to. Importantly, though, neither relies on planning to explain lengthening. This issue is taken up again in the next chapter.

Questions About Identity Priming. Surprisingly, the identity priming results were slightly different from Chapter 2's. The effect on duration was more robust, with identity primes leading to at least marginally shorter durations than the control condition at every ISI except -1500ISI. It is not clear why seeing phonologically-related primes as part of the experiment, instead of semantically-related ones, should lead to a more significant difference between identical and control priming. One explanation is that the presence of the phonologically-related primes, which partly but not fully overlapped with the target, created the need to draw a phonetic contrast between the target and the phonologically-related prime. This would distinguish trials with related primes from trials with identical primes for the speakers, making them more likely to reduce after identical primes.

Just as in Experiment 1, the control condition latencies make interpreting the identical vs. control comparison difficult. The control condition was a series of printed x's, as it was in Chapter 2. The printed x's as a control stimulus appear to have made the speaker's task easier, and latencies were shorter overall. While the negative ISIs with identity primes elicited shorter latencies, the opposite was true for the positive ISIs. Because these experiments were run concurrently with the experiments in Chapter 2, it was not possible to correct this disparity. Just as Experiment 2b fixed this problem for the semantic stimuli, Experiment 4 (reported in Chapter 5) will fix it for the phonological stimuli.

Chapter 5: Experiments 4a&b

5.1 Introduction

The previous chapter showed that the effects of phonological primes were notably more widespread than the effects of semantic primes. Duration was lengthened at nearly every SOA after a phonologically-related prime, which was explained in terms of late-in-processing emphasis. Latency cohered with prior PWI work, and was shorter after related primes instead of longer, like durations were. This suggests that even if speakers are lengthening in order to emphasize and/or mark difficulty, the processes responsible take place after planning has ended, and also in spite of a shorter planning phase. Chapter 5 continues to ask questions about the relationship between the fluency of phonological processing and word duration, and completes the analogy among the chapters. Like Chapter 3, the manipulations below also include blocks with consistent actions, and ask how the ease of retrieving and integration the action word influences the duration of each of the segments of the multi-word utterance frame.

5.2 Questions about competition

Effectively, the questions here are the same as in Chapter 4. Does phonological relatedness of an interfering stimulus have an effect on word duration, and if so, at which ISIs? Does phonological relatedness of an interfering stimulus have an effect on latency to begin speaking, and if so, at which ISIs? Like Chapter 3, this chapter relies both on previous PWI work and earlier experiments in this project. Chapter 4 showed a phonological facilitation effect on latency, which the experiments below should also show. Object duration also varied almost uniformly as a function of phonological relatedness, with longer durations after related primes. Evidence of durational lengthening would corroborate these results.

5.3 Questions about consistency

The guiding questions here are similar to the ones in Chapter 3: 1) will the consistency manipulation affect word duration at all? It should – Chapter 3 showed broad consistency effects on each of the segments of the utterance, and the manipulation is the same here. 2) which segments of the utterance will it affect – latency, object word duration, action word duration, or some combination of these three? In Chapter 3 consistency had an effect on all three segments, although with different time profiles for each. Once again, the expectation is the same here. The more interesting question is whether consistency will interact with condition at all. If speakers were simply emphasizing the object word in Chapter 4, the phonological information from the consistency manipulation will most likely have no additional effect on word duration. This is because speakers will have set the word for emphasis, and so will not benefit from the additional ease of processing that having early access to phonemes grants. However, if the production system has an easier time completing the coarticulation process because of the consistency manipulation, object word duration will be affected by consistency.

5.4 Experiments 4a&b

Methods

Participants. 40 native English-speaking undergraduates from the University of North Carolina at Chapel Hill participated in each experiment for pay ($7.50 for half an hour), or participation credit in an introductory psychology class.

Materials. All of the materials were the same as in Experiments 2 a&b, with the exception of the prime words, which were the same as Experiments 3 a&b (the phonologically-related primes).

Procedure. The procedure was the same as in Experiments 2 a&b.

Analysis. Due to a technical error, the data from the first half of Experiment 4a did not include the relevant primes. As such, the latter half of these participants' data was used in the model for this experiment. Experiment 4b used the first half of participants' data, and all other parts of the analysis proceeded the same as in the previous chapter.

Results

Summary statistics for each condition and ISI appear in Tables 18 through 21, broken down by segment. Parameter estimates for each of these effects appear in the Appendix with the other parameter estimate tables.

Table 15: Summary statistics for the latency to begin speaking, with standard deviations in parentheses. Experiment 4a is on the left, and Experiment 4b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	1176 (276)	1262 (317)	1177 (286)	1194 (237)	1106 (288)	994 (473)	1125 (305)	932 (491)
Identical	959 (198)	1103 (235)	1004 (288)	1121 (367)	1096 (256)	1040 (507)	1139 (278)	1028 (499)
Related	1259 (329)	1344 (370)	1269 (359)	1402 (369)	1149 (348)	1041 (466)	1156 (303)	1072 (482)
Unrelated	1206 (334)	1276 (335)	1249 (348)	1344 (344)	1153 (307)	1139 (537)	1203 (316)	1077 (496)

Table 16: Summary statistics for the determiner durations, with standard deviations in parentheses. Experiment 4a is on the left, and Experiment 4b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	118 (47)	126 (47)	123 (48)	128 (69)	131 (38)	125 (44)	130 (45)	120 (37)
Identical	119 (43)	119 (39)	117 (41)	123 (43)	123 (44)	126 (41)	125 (40)	125 (33)
Related	135 (66)	137 (67)	124 (53)	145 (83)	130 (53)	130 (45)	137 (45)	125 (39)
Unrelated	124 (51)	131 (49)	129 (67)	135 (65)	133 (56)	128 (44)	130 (45)	137 (41)

Table 17: Summary statistics for the object word durations, with standard deviations in parentheses. Experiment 4a is on the left, and Experiment 4b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	406 (112)	425 (137)	384 (103)	421 (135)	415 (113)	409 (121)	440 (122)	425 (126)
Identical	379 (101)	419 (117)	394 (101)	390 (125)	420 (127)	393 (117)	418 (120)	422 (121)
Related	387 (99)	444 (137)	415 (107)	462 (138)	453 (127)	459 (132)	433 (125)	435 (133)
Unrelated	415 (115)	415 (117)	402 (111)	447 (134)	421 (117)	441 (132)	451 (133)	420 (135)

Table 18: Summary statistics for the action word durations, with standard deviations in parentheses. Experiment 4a is on the left, and Experiment 4b on the right, broken down by ISI, condition, and then consistency.

	-1500		-300		300		1500	
	Consistent	Random	Consistent	Random	Consistent	Random	Consistent	Random
Control	692 (116)	706 (142)	696 (113)	703 (148)	781 (109)	785 (146)	767 (130)	773 (131)
Identical	711 (132)	683 (124)	692 (112)	691 (143)	771 (107)	767 (147)	785 (102)	787 (149)
Related	708 (111)	702 (142)	688 (124)	725 (141)	758 (117)	788 (138)	781 (111)	792 (136)
Unrelated	699 (112)	685 (144)	693 (107)	703 (138)	770 (113)	782 (139)	784 (114)	788 (169)

Related vs. Unrelated Models. The condition manipulation elicited longer latencies after related primes than unrelated, at the negative but not positive ISIs (Figure 22). Consistency shortened latencies significantly at negative ISIs but lengthened them marginally at positive ISIs (Figure 23).

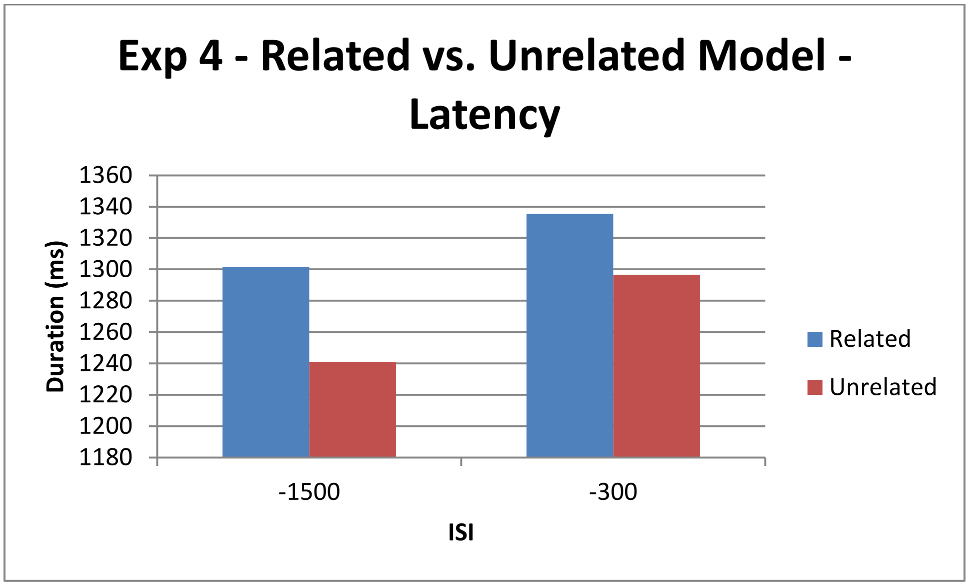

Figure 22: Graph showing the durations of latency, broken down by ISI and condition. This illustrates a main effect of relatedness.

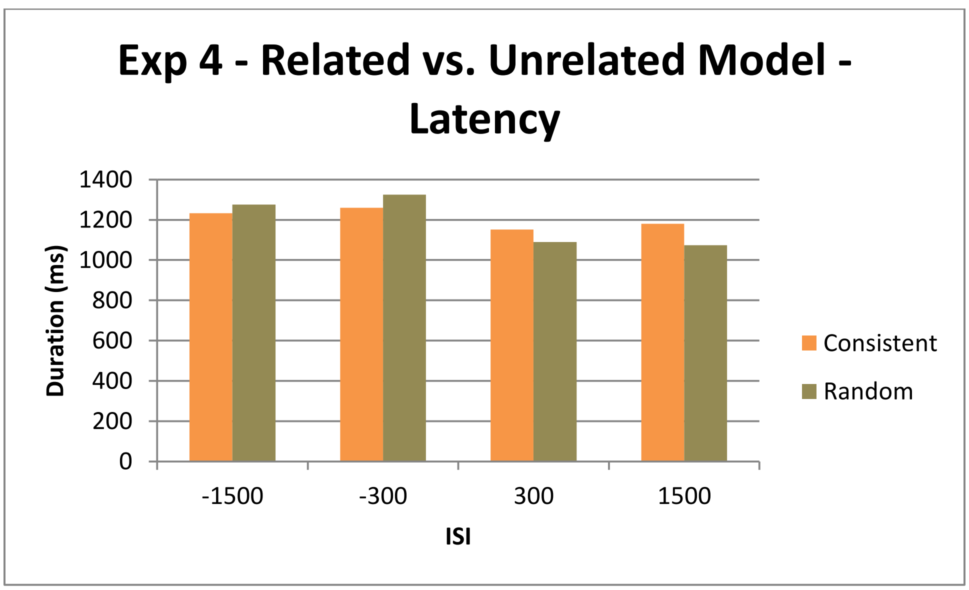

Figure 23: Graph showing the durations of latency, broken down by ISI and consistency. This illustrates two different effects of consistency: a main effect with consistent shorter than random at negative ISIs, and a marginal main effect in the opposite direction at positive ISIs.

Conversely, relatedness led to longer object word durations at positive ISIs but not negative (Figure 24). Despite the appearance of the bars in the figure, the interaction between condition and ISI did not reach significance (t = -1.8). Consistency led to shorter object word durations, but only at negative ISIs (Figure 25).

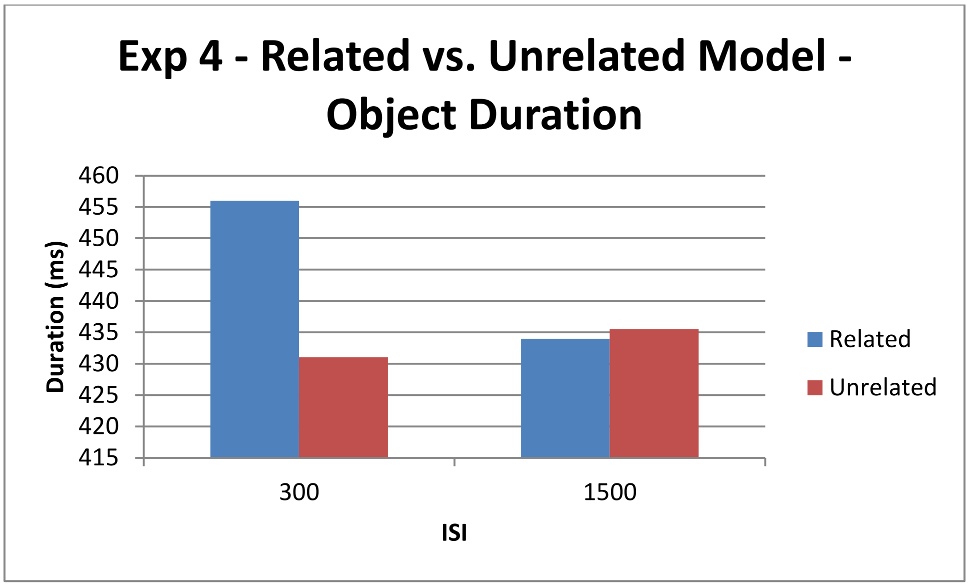

Figure 24: Graph showing the durations of the object words, broken down by ISI and condition. This illustrates a main effect of relatedness (related longer than unrelated).

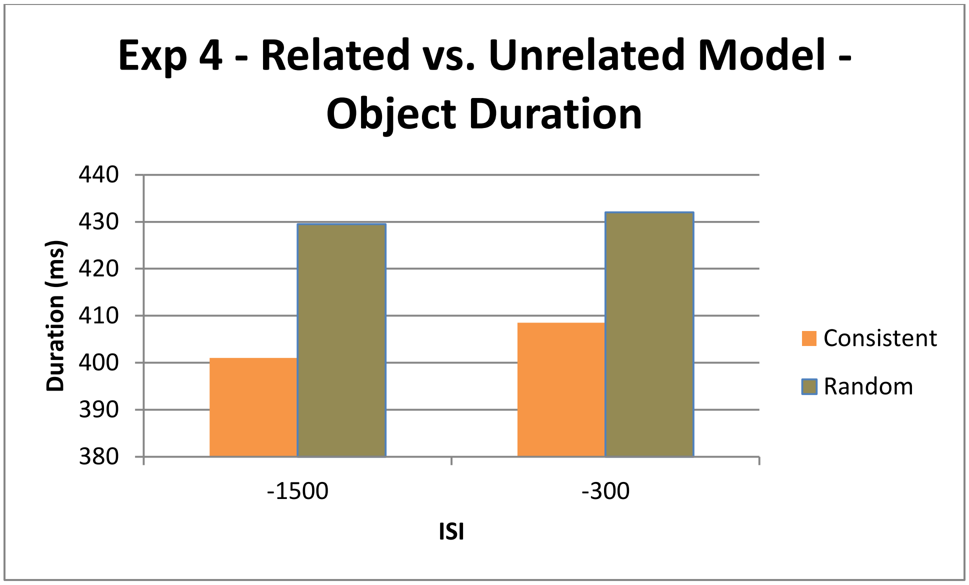

Figure 25: Graph showing the duration of the object words, broken down by ISI and consistency. This illustrates a main effect of consistency.

Action word durations, however, were shorter in consistent blocks at both positive and negative ISIs (Figure 26). At the negative ISIs only, there was a significant effect of condition and of ISI, with shorter durations in related blocks and at -1500ISI.

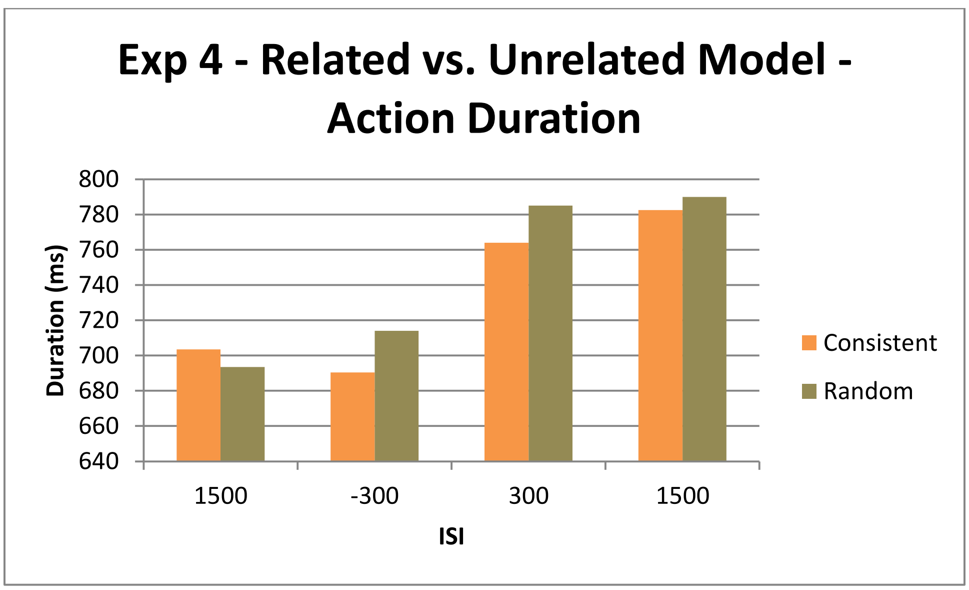

Figure 26: Graph showing the duration of the action words, broken down by ISI and consistency. This illustrates two main effects of consistency, one in the negative ISI model

and one in the positive. Both results are in the same direction, with consistent shorter than random.

Identity vs. Control Models. The identical condition elicited shorter latencies than control at the negative ISIs but not positive. Consistency had the same effect, with shorter latencies in the consistent condition. Consistency also interacted with the other two variables, which is shown in Figure 1. Condition and consistency were both significant at both ISIs, with a stronger effect at -1500ISI for both variables. There was also a significant interaction between condition and consistency, with even shorter durations in the consistent and related condition, but only at -1500ISI. At the positive ISIs, consistency interacted with ISI, with consistency eliciting longer latencies at 1500ISI only. These simple effects are illustrated in Figures 27 and 28.

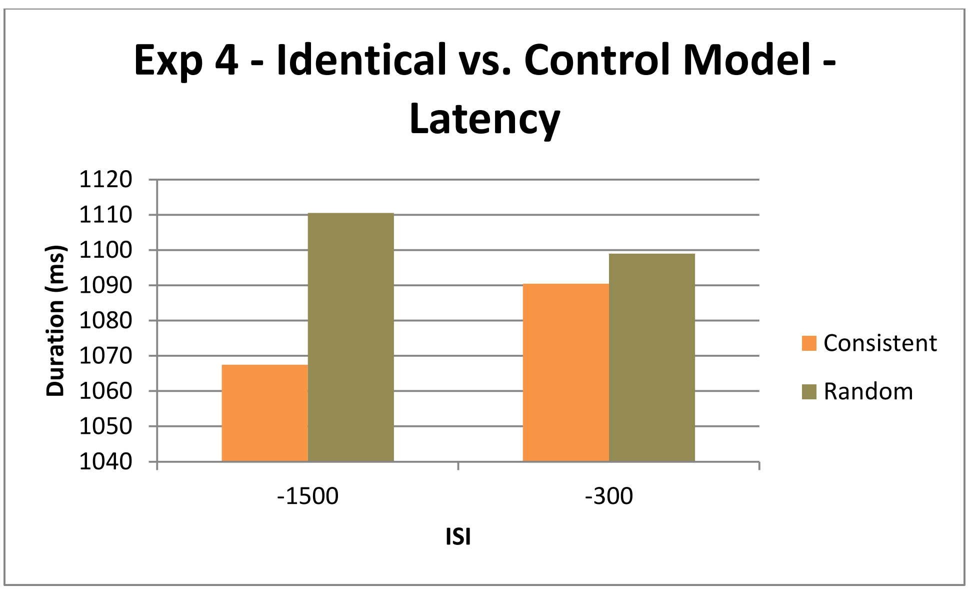

Figure 27: Graph showing the durations of the latency, broken down by ISI and consistency. This illustrates a main effect of consistency.

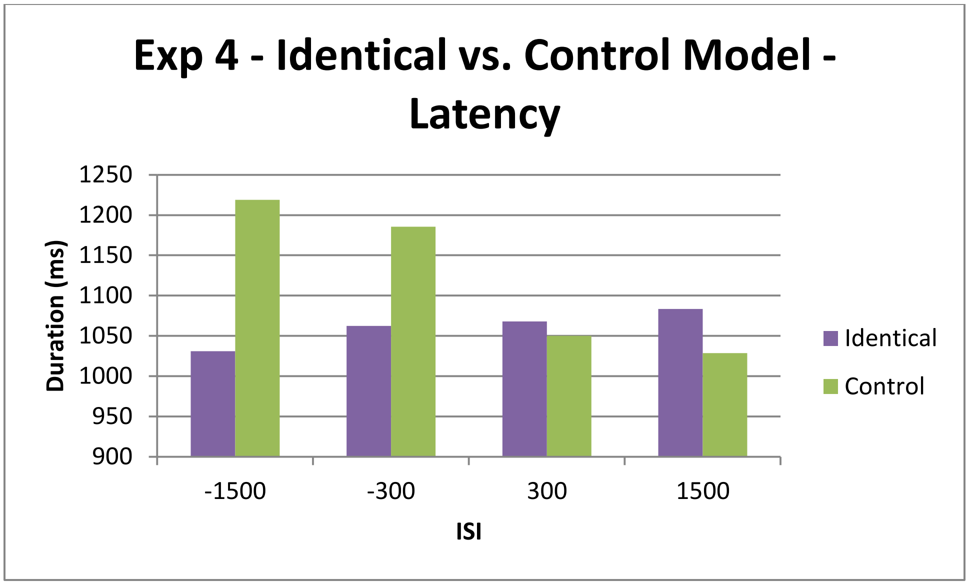

Figure 28: Graph showing the durations of latency, broken down by ISI and condition. The negative ISIs show a main effect of relatedness (identity), and the positive ISIs show a simple effect of relatedness (identity) at only 1500ISI, in the opposite direction.

For the object words, the only effect was decreased object word durations at the positive ISIs.

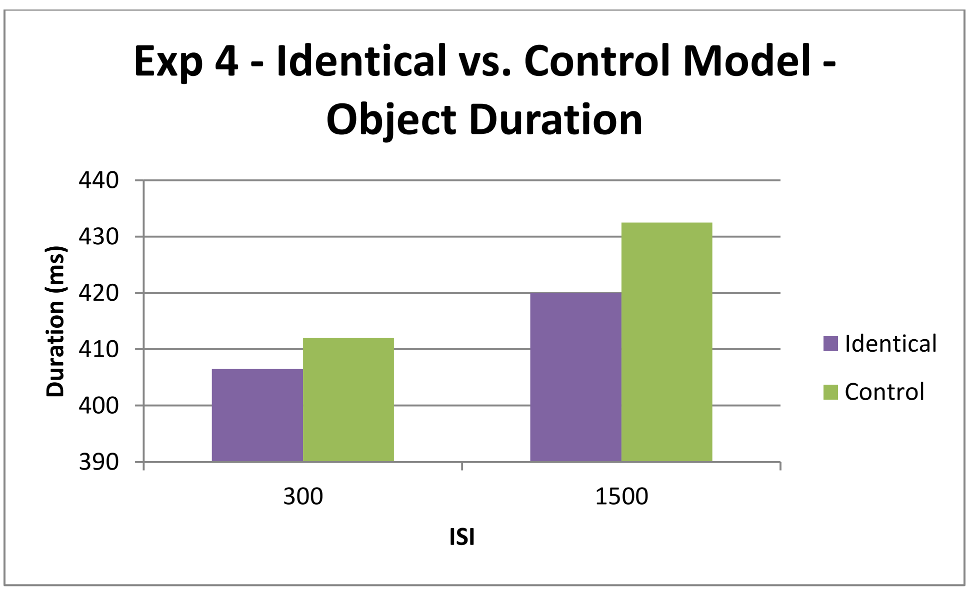

Figure 29: Graph showing the durations of the action word, broken down by ISI and condition. This illustrates a main effect of relatedness (identity).

For action words, consistency led to shorter durations at both positive and negative ISIs.

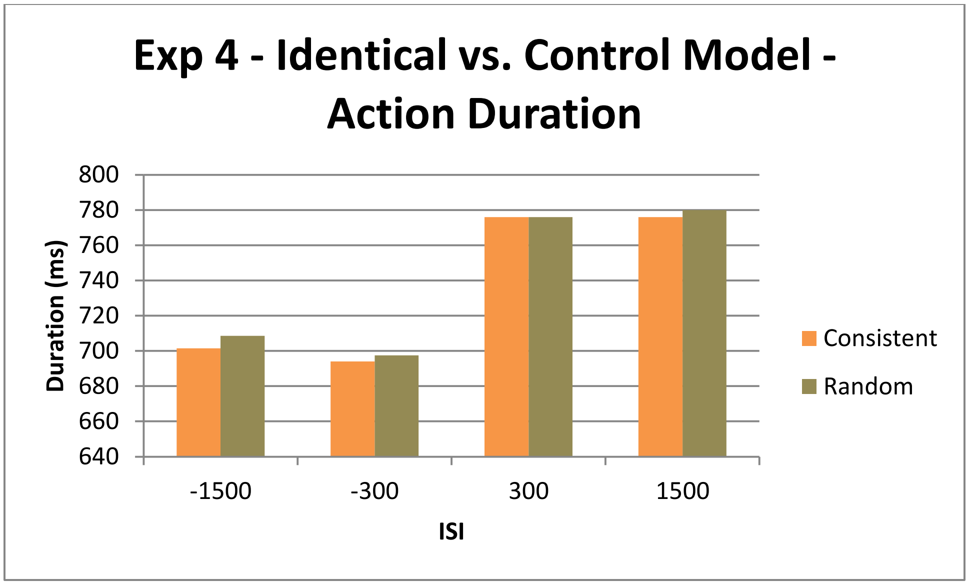

Figure 30: Graph showing the durations of the action word, broken down by ISI and consistency. This illustrates two main effects of consistency, one in the negative ISI model and one in the positive. Both results are in the same direction, with consistent shorter than random.

Discussion

Questions About Relatedness. The relatedness effect on the object word ran in the same direction as before, with related primes leading to longer durations than unrelated primes, providing additional evidence for the similar outcome from Experiment 3. Once again it seems likely that the production system is modulating length because of phonological retrieval and/or integration, either from the additional time needed to ensure that the correct phonemes get produced, or from contrastive emphasis. Notably, lengthening appeared in spite of an absence of a latency effect, which differs from the facilitatory effect on latency found in Chapter 4. The null latency effect is somewhat surprising, and I will return to it momentarily. The noteworthy point here is that object word duration still differed, meaning that emphasis occurred without any lengthening of the time speakers spent planning. It is possible that adding emphasis to a word requires no additional planning time, but this seems like a dubious assumption. Prosodic theory treats emphasis as a decision to be made, if not also the marked form (Gussenhoven, 1984; Ladd, 1996; Pierrehumbert & Hirschberg, 1990; Selkirk, 1995), which means that assessing whether to emphasize a word should take time. The durational lengthening without a latency lengthening in this experiment (and the same in the previous experiment) suggests that speakers emphasized here without going through any additional processing during the planning phase.

Here, as opposed to Chapter 4, however, object word lengthening appeared only at positive ISIs. This raises questions about what differences in the manipulation made the negative ISI effect not appear in this experiment. The negative ISIs also differed from the positive in that they exhibited a lengthening effect on latency as opposed to a null effect. One possibility is that the planning of emphasis appeared at the negative ISIs, whereas at the positive ISIs emphasis was a post-planning effect. This hypothesis is supported by the consistency data. Consistency had a reduction effect on all segments of the utterance at the negative ISIs, but only on latency and the action word duration at the positive ISIs. One of the predictions for this experiment was that consistency might interact with relatedness if speakers could complete at least some of the coarticulation processes necessary for the utterance before beginning the action word. This appears to have been the case at the negative ISIs but not the positive, given that the consistency manipulation affected the object word duration only at the negative ISIs.

Having the action word ahead of time allowed speakers to speed up the whole utterance, because they had nearly all of the phonological information they needed. But this was only possible when the prime word appeared before the target, because the prime words were designed to share not only onset, but also prosodic structure and number of syllables. This allowed the speakers to plan the majority of the structure of the utterance ahead of time, setting out the syntax, prosody, and 2 of the 3 lexical items (“the,” and “rotates”). The only difficulty they encountered was in deciding whether they needed to emphasize the object word, which occurred during the planning phase. Then, any actual durational emphasis was washed out by the reduction from consistency.

When the prime word appeared after the target, in the positive ISIs, speakers could not plan whether they would emphasize the object word or not. This meant that the plan was not delayed, but emphasis still had to be processed. In that case, speakers did not delay the onset of speaking, but instead spent a little bit of extra time on the object word. The effects of consistency still appeared on latency and the action word duration, because it was still easier to process those portions of the utterance, both in planning and execution.

The take-home message from the consistency and condition manipulations in this experiment is that the timing of processing-relevant information matters a great deal when assessing whether that information will affect word duration. Although the two variables did not interact statistically, they did make processing-relevant information available to the system at different times. The condition manipulation eased the processing of the onset phoneme, but made it necessary to emphasize the object word. This required additional planning time if the consistency manipulation affected the object word duration (as in the negative ISIs), but resulted in actual durational emphasis when the consistency manipulation did not affect processing early enough to integrate with the processing of the object word (as in the positive ISIs).

Questions about identity priming. Taken together, the evidence from the identity prime comparisons shows that reduction in latency does not correlate perfectly with reduction in word duration. This suggests that speakers need not always plan to reduce the duration of words in an utterance, and that the production system can decide or otherwise determine the duration of a word or words as a function of online processing. It is worth mentioning that this conclusion is weakened slightly by the lack of direct evidence for the comparison between the control and identity conditions at the positive ISIs in Chapters 2 and 4. As in Chapter 3, the spoken control prime in Chapter 5's experiments allowed for a direct comparison with the identical condition, in contrast with the visual prime in the other chapters. The general conclusion about the relationship between control and identity priming is that in the experiments with printed control primes, participants used the prime merely as a cue to respond. Although that comparison may have theoretical implications as well, the more relevant set of results for current purposes comes from the auditory control condition, and so I limit my focus somewhat to there.

Throughout all of the studies reported here, identity priming reduced latencies to begin speaking. Apart from the limitations on interpreting Experiments 1 and 3, the reduction in response to identical primes was very widespread. This is expected – speakers should find processing the same information again after a short interval easier. The same representations are active or partly active, the relevant processing has been recently completed and so should experience little noise or interference from related processes that might otherwise have to be ignored, and in general the word is readily available to be inserted into the ongoing utterance or utterance plan. The two sub-questions that have guided the investigation of the effect of identity primes have been their timing and which segments they affect. These questions bear on the more general question of how facilitation of processing affects word duration, because identity primes are the paradigm case of facilitated processing. If a time point does not exhibit reduced latency or duration after an identity prime, or a segment of the utterance only responds to identity priming under certain circumstances, that suggests that facilitation has a limited capacity to influence processing during those time windows or on those utterance segments – the production system is not set up to entertain or capitalize on that relationship. The corollary of such a finding is that weaker forms of facilitation (e.g. only phonological, or only semantic) are less likely to have an influence under the same circumstances. Further, the influence of these weaker forms of facilitation shows how much each contributes to determining word duration by itself. This makes the identity prime's effects not only the baseline for interpreting facilitation, it also has the potential to show whether having the entire system primed has synergistic effects.

The effect of identity priming was very robust at negative ISIs, with only a single deviation from significance into marginal significance at -900ISI in Experiment 1. All other latencies were shorter after identity primes than control primes at negative ISIs. The conceptual explanation for this is simple – having early access to a significant portion of the linguistic information associated with the eventual target made planning to utter the name of that target easier, and thus the latencies shorter. The processing explanation is perhaps not so straightforward. Two courses are possible here, which were discussed in more detail in Kahn and Arnold (2012): 1) an information-status representation such as givenness played a role; 2) the production system was facilitated by the already-available information from the prime. The two are not mutually exclusive, and it is almost certainly the case that both types of processing will be necessary to fully account for the range of phenomena presented both here as well as in more-complex discourse situations. The two pieces of evidence from this set of experiments that distinguish between an information-status explanation and a facilitation explanation are discussed in more detail below, because they are object word reductions.

The positive ISIs' effects are also worth discussing before turning to object word durations, however. Experiments 1b and 3b did not have identical and control primes in the same modality, making a direct comparison between them problematic. Experiment 3b, for example, displayed significant differences between the identical and control conditions, but in the opposite direction from expectation (control shorter than identical). In the discussion for that experiment, I noted that speakers seemed to have used the printed x's in the control condition as a cue to respond, whereas they actually processed the identical stimulus as more than simply a cue. However, Experiments 2b and 4b did not have this same limitation. Somewhat surprisingly, given the robust results at the negative ISIs, the identical condition did not differ from the control condition in the positive ISIs in Experiments 2b and 4b.

Assuming that speakers were still processing the identical primes in approximately the same way as they did in Experiments 1b and 3b, this means that the control condition was taking longer to process than it did in those experiments. It seems that the word "nothing" was more difficult to process than printed x's, but not so difficult to process that it elicited longer latencies than the identical condition (but only at the positive ISIs). Note that this is more likely than the alternative, which is that the identical condition was reduced more while the control condition stayed the same, because the negative ISIs exhibited the same pattern through all 4 experiments. The null effect in Experiments 2b and 4b suggests that even after seeing the identical prime, speakers did not have any more or less difficulty processing it during the planning portion of production than they did in the control condition. The ease in processing difficulty frequently appeared instead on the duration of the object word. Another possibility is that at the positive ISIs, speakers had already planned most of the utterance, and so did not have time to integrate the identical prime.

The reduction effect of the identical prime on the object words' duration was widespread, similar to the effect on latency. In fact, there were only major two places where this effect did not appear – the positive ISIs in Experiment 1, and the negative ISIs in Experiment 4. These gaps complement the latency results – every ISI had either a latency reduction or a duration reduction, and many had both. This suggests that the identical primes had a very reliable effect on processing, but that something about the manipulations made that effect manifest in slightly different ways. The timing of the prime, the presence of either semantic or phonological relatedness, and the presence of consistency are the three sources of variation. All three of these variables interacted, and not in a simple way. Unfortunately, the problems with the modality of the control condition may partly explain this interaction. The timing of planning, which was apparently affected by the difference between the printed x's and the word "nothing," may also have affected the duration of the words in the planned utterance. Because of this ambiguity, a satisfying interpretation of the object word durations in the positive ISIs from Experiments 1 and 3 will have to wait on additional data.

The remaining data points suggest that object word duration is easily affected by the identical primes, in slight contrast with latency. The most straightforward explanation of this effect is one based on information status. The prime made the prime word slightly predictable, a reasonable candidate for future reference, and readily available to the speaker. All of these properties are important, although neither necessary nor sufficient, for an information status categorization of repeated-mention. Without a larger conversational context, it is less likely that the speaker treated the prime as a repeated-mention in the same way as s/he might while speaking with a friend. However, reduction on the object word can frequently and felicitously mark an entity as having been previously-mentioned (Pierrehumbert & Hirschberg, 1990; Prince, 1992), which means these data could indicate that the speakers treated the primes as markers of information status.

An information status explanation has a more difficult time explaining the conjunction of the latency and duration data than the duration data alone. Experiments 2 and 4 were both split between positive and negative ISIs on whether the object word duration was affected or the latency was affected or both. The positive ISIs for both experiments showed durational reduction, but not latency reduction, whereas the negative ISIs showed latency reduction in both experiments, but durational reduction only in Experiment 2. This shows that speakers did not treat the identical primes the same between the experiments, and there does not seem to be much of a reason to treat them differently from an information status perspective. In all cases, the prime made the prime word predictable, a good candidate for future reference, and readily available to the speaker. It did this without affecting the latency in the positive ISIs. Information status is typically conceived as a variable that gets processed very early in language production (Bock, 1982; Levelt, 1986; Schmidt et al., 1999), typically before lexical selection takes place. By the logic of PWI, the null effect on latency between the identical and control conditions signifies that speakers took a similar amount of time to plan the utterance. In all likelihood, then, information status was not considered, or at least not in a way that affected the length of time it took speakers to plan.

Explaining the reduction on duration in absence of a similar reduction in latency without making reference to information status requires looking at processes that occur after planning is completed, or mostly completed. The identical primes should have called attention to the split between overlap and non-overlap with the related primes. That is, because the onsets in the related primes were shared but the offsets were not, speakers may have simultaneously experienced eased processing on the onset and a need to emphasize the offsets' difference. This is the same sort of explanation offered for the difference between "the BLACK car" and "the black CAR," each of which highlights something different about the conversational situation. Emphasis can appear as durational lengthening (Ladd, 1996; Pierrehumbert & Hirschberg, 1990), which would explain the durational lengthening here. Interestingly, it appears even alongside the latency phonological facilitation effects. This suggests that speakers do not set out to emphasize in their plan, but actually have an easier time initiating the plan due to the activated onset phoneme, and then still lengthen the duration of the object word. A post-planning explanation for this phenomenon requires that some process or processes in the production system that take place after lexical selection has an influence on duration. Levelt et al. (1999) include a post-phonological processing, pre-articulation monitor in their model, which gives the system an opportunity to censor or modify its output before actually sending it to be articulated. The monitor is necessary for other reasons (e.g. censoring swearing, or avoiding slips of the tongue), which means that this explanation does not require an ad-hoc additional assumption.

Unlike the relatedness section, the monitoring explanation is not the only plausible facilitation-based explanation for the identity priming data. The identity primes most likely initiated or completed many of the processes associated with processing the targets, such as activation of the semantic or phonological information, selection of the same, incorporation of these selections into the sentence frame, and so on. Both an explanation in terms of speed of phonological activation or selection (ala Balota et al., 1989), or an explanation in terms of coordination between levels of processing (ala Bell et al., 2009), would also suffice. The pre-activation or selection or retention of phonological information from the prime would facilitate subsequent processing of the target word, on either explanation.

Of course, all of the explanations mentioned so far (information status, monitoring, coordination, or activation speed) are not mutually exclusive. Information status seems to be necessary as an explanatory construct, as mentioned above. Monitoring, too, seems necessary to explain other phenomena in production, also outlined above. It is less obvious that processing-based explanations are necessary, although data such as the ones in this study, and in corpus studies like Bell et al., (2009), suggest that they explain reduction phenomena that the others cannot.

Questions about consistency. This second pair of experiments provides additional information about how the timing of information relative to actual production matters. The consistency manipulation had a universal reduction effect on action word duration, and a nearly universal reduction effect on latency. This shows once again that the consistency manipulation has a fairly targeted effect, appearing primarily on the word that was relevant to the repeated action. Knowing which action was going to appear ahead of time allowed the participant to begin speaking earlier in most cases, and to reduce the duration of the expected word. Unlike Chapter 3, object word duration was affected only in the related-unrelated model at the negative ISIs, suggesting that in this case, the consistency manipulation did not also allow speakers to reduce the duration of the preceding word.

However, a similar tradeoff to the one that appeared in Chapter 3 appears here as well. There, at the negative ISIs object word duration was reduced while action word duration showed a null effect. Conversely, at the positive ISIs, action word duration was reduced while object word duration differed. This was explained in terms of the production system reducing either based on anticipating the ease of producing the action word, or reducing the action word itself when it couldn't anticipate. A similar phenomenon seems to have occurred here. At the negative ISIs, speakers managed to anticipate the action word and integrate processing it in such a way that the duration of the object word was reduced. The same explanation is likely – the representations and processes for the action word had recently been completed, and so when speakers went to slot the word into the inchoate utterance, they found that slotting easy, and could speed the duration of the preceding word(s; the determiner was significantly shorter in the consistent blocks as well). At the positive ISIs, anticipation of the action word was not sufficient – speakers could not integrate the eased processing into the ongoing processing of the utterance as whole, and so reduced only the predictable word itself while spending more time planning. However, the null effect of consistency on the duration of the object word freed up resources to consider the object word relative to the prime. Here, as in Chapter 4, the related prime elicited longer object word durations than the unrelated prime. Once again, this was not a statistical interaction between prime type and consistency, but rather a contingency relationship – if speakers found themselves able to consider the object word independently of the processing of the action word, they emphasized the object word and reduced the action word. But this was only possible at the positive ISIs, where integrating the processing of the action word into the ongoing processing of the utterance was less possible.

The difference between the positive and negative ISIs both here and in Experiment 2 highlights the fact that speakers sometimes engage in long-term planning of utterances, and therefore word duration, and sometimes rely more on incremental processing to determine duration. Consistency did not interact with condition significantly at any point, meaning that consistency had its effects over and above any difference based on relatedness or identity priming. Not only does this show that speakers planned to the end of these utterances before beginning to speak (consistent with Jescheniak et al., 2003, who found similar effects), but that each part of the plan completed without much reference to the other parts. This was especially apparent at the negative ISIs, after related stimuli. Latencies in those levels of the independent variables were longer because of relatedness, but shorter because of consistency.

Chapter 6: General Discussion

6.1 Summary of results

All of the relevant comparisons investigated here revealed at least some significant and theoretically-interesting findings. Semantically-related primes elicited longer latencies at negative ISIs, and some hint of longer durations at -900ISI. Phonologically-related primes elicited shorter latencies at negative ISIs, and had widespread durational lengthening effects across the timepoints. Identity primes, too, shortened latencies, but also durations, and importantly the differences emerged at non-overlapping time points. The consistency manipulation shortened planning latencies and action word durations, but more importantly, it also shortened object word durations. Together, these provide answers to all of the questions the project proposed to investigate, including whether ease of processing has different effects at different time points.

6.2 Implications for Processing

The majority of the findings here point to a role for lexical and/or phonological retrieval and integration in affecting the latency and duration of words in multi-word utterances. Latencies, as well as object word durations and (when manipulated) action word durations varied as a function of manipulations that were intended to facilitate or inhibit language production processing. Semantic inhibition and phonological facilitation both appeared on latencies in the appropriate places, suggesting that speakers had a more difficult time retrieving the correct item in the face of semantically-related competitors, and an easier time retrieving and/or integrating in the face of phonologically-related ones. Object word durations were longer in one instance after semantically-related primes, but more robustly after phonologically-related primes. This suggests that semantic relatedness has a limited capacity to affect duration, compared with phonological relatedness. Phonological relatedness, however, is not as limited, and had a wide-ranging lengthening effect. The coordination of phonological processing, either between it and the semantic level of processing, or between the phonological processing of one word and another, appears to have a greater capacity to influence duration.

The ease of phonological retrieval and integration appears to have a greater effect on duration than the ease of lexical retrieval. I have previously suggested that facilitation throughout the production system (i.e. at any level) has the capacity to affect word duration (Kahn & Arnold, 2012). This claim does not stand up to scrutiny, both in follow-up work in a similar paradigm (Kahn & Arnold, 2013) and here. Kahn and Arnold (2013) tested whether facilitating the articulatory level of processing would lead to reduction, and found null results. The current results are not null, but the effect of semantic processing on duration is clearly restricted. The claim in Bell et al., (2009), that the coordination in the production system between semantic and phonological processing controls much of the variance in word duration, seems closer to the truth. The manipulations in Experiments 2 and 4 targeted phonological retrieval and integration in particular, and found widespread lengthening effects. Although this effect is potentially explainable as emphasis, it is important to note that in some cases the duration differenced appeared in the absence of a latency difference. This disparity suggests that the phonological manipulation had an influence after the planning of the utterance had completed, which coheres well with Bell et al.'s explanation. The production system appears to have modulated word duration online, post-planning, in order to respond to fluctuations in processing fluency.

The ease of phonological retrieval and/or integration having a greater influence on word duration also makes sense given the assumed architecture of the production system, where phonological processing necessarily interacts more closely with articulatory processing. Any stimulus which influences processes that directly precede the ones that are ultimately responsible for the mechanical articulation of an utterance (i.e. phonological processes) will nearly always have a greater influence on duration than a stimulus that influences processes further upstream. In other words, the influence of a semantically-related prime still has to pass through the parts of the system that are responsible for phonological processing, making its ability to influence duration indirect at best. A phonologically-related stimulus has to pass through fewer downstream processes, if it passes through any others at all, before its influence reaches articulatory processing.

The link between latency variation and duration variation is relatively weak here. The two measures frequently did not vary together, or if they did, in opposite directions. This suggests that the length of the plan that precedes utterance onset does not entirely determine utterance duration, and vice versa. As an example, if the length of the plan did determine word duration, all of the negative SOAs in Experiment 1 – which were longer after related than unrelated primes – would have led to durational lengthening as well. They did not – only -900SOA exhibited any durational lengthening. A similar point was made in the previous paragraph about the phonological results. The processing implication of this finding is that at least some of the processing that influences a word's duration occurs after utterance initiation, according to the logic of the PWI paradigm.

There are no existing provisions in production models for specifying the duration of a word. Nor are extant computational models designed to deal with the complexity of producing a multi-word utterance, which is not the same as stringing together several independent single-word models. Note that this is intentional and not a flaw or oversight – models of production to date have focused on phenomena other than word duration. The conclusion from the previous paragraph may thus be too strong. If the planning of a multi-word utterance differs significantly from the planning of a single-word utterance (for example, that it requires multiple lexical access, or storing number information about a word, or resyllabification), the planning phase could still include some process(es) that influence word duration. However, it is not entirely clear why the difference in duration would appear in tandem with a difference in latency at -900SOA, but not the other negative SOAs. If the influence comes from some process in the planning phase, that influence is obscured.

On the other hand, prior work in a PWI-like paradigm offered a post-planning processing explanation for the influence of semantic processing fluency on word duration. Like Experiment 1, Balota et al., (1989) found duration differences, which they attributed to the speed of phoneme activation, referencing Dell (1986)'s model of production. The model itself makes no specific claims about chronometrics, so it is not fully clear how to delineate its processing into planning and post-planning stages. The notable property of this model that distinguishes it from the Levelt et al., (1999) model, however, is that semantic processing variation (e.g. how quickly or strongly a lexical item is selected compared to its competitors) is allowed to influence phonological processing. This means that semantic competition could lead almost directly to speed of phonological selection, and at a post-lexical-selection stage of processing. As Balota et al. note, however, the model does not include this assumption as part of its architecture – they added it to explain their data.

Identity priming had a strong effect throughout the experiment, particularly on object word duration. This is not surprising either, given previous work in this paradigm (e.g. Glaser & Glaser, 1989) and others (e.g. Fowler & Housum, 1987; Kahn & Arnold, 2012). Identity primes should have made it easier to plan the utterance by making important information about its contents available during the planning stage. Whether this information took the form of already-active representations at the semantic and/or phonological level, lowered selection thresholds, ease of integration with other words, or some other facilitated processing is unclear. It may also be the case, as discussed in Chapters 2 and 3, that speakers selected a reduced form of the object word during the planning phase. Regardless, identity primes had the expected facilitating effect. One final implication for processing concerns the role of identity primes. The motivation for including them in the manipulation was twofold: 1) to establish a baseline for facilitated processing, against the background of prior PWI findings; 2) to investigate the effects of repeated-mention on word duration. As to the first, apart from the difficulties with Experiments 1b and 3b, the manipulation successfully elicited reduction. Latencies and object word durations were generally reduced, confirming that planning and executing the utterance was easier after identity primes. As to the second, the real question is why planning and executing were easier. Two explanations have been offered; 1) a trigger-based mechanism, where speakers saw an identity prime and planned to reduce the object word from the outset; 2) a facilitation-based account, where some portion of the post-planning processes were easier to complete, and thus allowed the speaker to reduce the object word.

The earlier chapters remained somewhat agnostic about which of these explanations was more likely. In many cases throughout these experiments, latency and duration were shorter at the same time, but this was not always true. Experiments 2 and 4 in particular showed divergence, where the negative SOAs showed reduced latencies but the positive SOAs showed only reduced duration. One possible explanation for the disparity is that in one case speakers had an easier time planning and in the other case an easier time executing. This would most likely involve different processes – perhaps a trigger mechanism at the negative SOAs and a more online set of facilitation-based reduction processes at the positive SOAs. In fact, in order to explain the reduction on object word duration in the absence of a planning effect, a trigger mechanism would have to take no time at all, and have no impact on processing. Like the monitoring mechanism discussed earlier, the trigger mechanism for reduced forms appears to be necessary for reasons other than duration management. But facilitation-based mechanisms that don't make use of the trigger also appear to be necessary to explain the rest of the data in the current studies.

The consistency manipulation, where speakers saw the same action word repeated throughout a block of trials, generally resulted in shorter action word durations. This is not surprising, given that speakers should have found processing that word easier in consistent blocks than inconsistent blocks. It shows that identity priming of the sort used in the condition manipulations is not the only way to facilitate the processing of a word – priming implicitly also made processing easier. The consistency manipulation also shortened object word durations in some cases. The presence of this effect depended on SOA, or more accurately, whether the action word was integrable with the ongoing processing of a fluent utterance. At negative SOAs, speakers had the ability to modulate the duration of the object word – by reducing it – when they knew that the upcoming action word required little planning. At the positive SOAs, they restricted reduction to the word that they knew would reliably appear every time – the action word – instead of predictively reducing the duration of the object word.

6.3 Implications for methodology

Previous studies have shown both a semantic inhibition effect and a phonological facilitation effect on latency (Schriefers et al., 1990; Starreveld & La Heij, 1996, *inter alia*). Semantic inhibition effects tend to appear at negative SOAs, in the form of increased latencies to begin naming objects in response to related primes than unrelated primes. Phonological facilitation effects tend to appear at slightly negative or positive SOAs, with shorter latencies after related primes. The results of the current study generally support these findings, particularly semantic inhibition. The timing of the current phonological facilitation effects was slightly more restricted, but is also slightly more varied in prior studies.

The conceptual replication shows that the non-standard manipulations in the current study did not compromise its ability to find standard effects. This implies that it tapped into at least some of the same processes, primarily semantic- and phonological-level processing. More importantly, it shows that the other significant effects also most likely tapped into the processing system in some of the same ways, taking into account the fact that speakers might have had to hold responses in abeyance until the experiment signaled them to respond. The duration effect at -900SOA in Experiment 1, for example, most likely emerged because of inhibited semantic processing, and the consistency effects in Experiments 2 and 4 because of facilitated action word processing. The logic that leads from reduced latency to reduced planning time via facilitated processing also leads from reduced duration to reduced processing difficulty.

Looking forward, the finding of standard PWI effects alongside novel processing-relevant effects suggests that duration as an outcome measure could play a larger role in PWI investigations than it has in the past. Models of language production are beginning to look to chronometric evidence (Levelt et al., 1999), and the field is becoming increasingly interested in the implications that durational variation has for language production (Bell et al., 2009, Gahl, 2008; Jaeger, 2010). As mentioned in the introduction, one of the goals of the current study is to provide evidence for such investigations, by showing how duration varies in theoretically-relevant ways. Another related goal is to show that duration as an outcome variable has the potential to offer additional useful data to studies that are already extracting chronometric data from spoken utterances.

Other PWI studies that have measured duration found null results (Schriefers et al., 1990; Jescheniak et al., 2003), but PWI studies tend to use single-word naming as their behavioral response. As I argued in Chapter 1, a multi-word utterance has more space for expansion, both in terms of the sheer length of the utterance and the amount of processing necessary to produce it. Duration is thus not only a potentially useful tool to investigate processing, but may ultimately be necessary to fully understand the production system. Multi-word utterances involve processes that single-word naming does not, including coarticulation, resyllabification, syntactic agreement, morphological agreement, and others (Levelt, 1989). The duration results here and in other studies (Bell et al., 2009; Christodoulou, 2012) suggest that speakers perform some of this processing online – that is, while they are speaking other words, not before they begin speaking an utterance. Measuring duration, and specifically measuring it in multi-word utterances, is likely to be necessary for a fuller understanding of everyday language production.

To my knowledge, the greatest magnitude of SOA used in a PWI study was 1200ms, in Balota et al., (1989). The results from the current study suggest that this magnitude may not be sufficiently large to capture some of the processes involved in production. The 1500-magnitude ISIs here quite reliably showed significant effects, and were, more often than not, almost identical to the 300-magnitude effects. Prior chapters have discussed the possibility that speakers were merely holding a response in anticipation of the appearance of the target or prime. This is indeed possible, but the evidence is not conclusive. 1500 and 300 magnitude ISIs did not behave fully identically, suggesting some differences in processing. It is not entirely clear what these differences might be, because this is the first evidence that production-relevant processing might be occurring so far in advance of the initiation of an utterance. Future work in PWI, and particularly studies of multi-word utterances, may need to consider using large- and small-magnitude SOAs or ISIs to ensure that it measures all of the relevant effects.

6.4 Limitations and Future Directions

The most obvious limitation in these experiments is in the mismatch between the modalities of the primes in Experiments 1 and 3. This prevented a reliable interpretation of the effects of the control prime that spanned all four experiments. Experiments 2 and 4 standardized the modality of presentation, but did not directly replicate the experimental conditions of Experiments 1 and 3, making a comparison somewhat problematic. Although Schriefers et al., (1990) conducted a systematic investigation of the use of a variety of control primes, and found no difference between silence and a null auditory condition, that did not generalize to these experiments. The major significant difference was the use of a multi-word utterance response. It may be necessary in the future to conduct another systematic investigation of how different priming modalities affect latency, but especially duration, to more confidently establish an appropriate baseline.

Another desirable comparison that these data do not permit is a more-direct one between semantic and phonological information. Because participants only saw one type of relatedness, it was not possible to draw any conclusions about how speakers responded differently to the two, particularly in the timing of any variation. The current experiments thus have less to say about the ordering of particular processes – their strength is in showing that timing does matter. Future work will have to begin putting the processes in order, and showing when they interact, and whether they proceed simultaneously or serially. This would bring PWI investigations of multi-word utterances more closely in line with Schriefers et al., (1990), who used a within-subjects design to look at how individual speakers were affected by both semantic and phonological processing.

Using only 3 ISIs in the first four experiments, and 2 in the latter four, was necessary to reduce the complexity of presentation. Unfortunately, it did not allow for a fine-grained exploration of exactly when the primes had their effect. A less exploratory experiment or set of experiments that was more focused on the precise timing of these effects (on the order of a hundred milliseconds, instead of hundreds), could stagger the ISIs more closely. Like the paragraph above, this would bring PWI investigations of multi-word utterances more closely in line with prior work (Schriefers et al., 1990, *inter alia*) that used closely-spaced ISIs. Pilot work for the current studies showed few if any effects at -1200 and -600ISI, but additional tests are probably warranted.

The presentation of the primes at positive SOAs in previous PWI experiments has mostly taken a form where the SOA is small enough that the prime appears somewhat contemporaneously with the target. The current experiments did not use temporal differences that small, because of the need to explore the longer-term planning necessary for multi-word utterances. One consequence of this is that the positive and negative ISIs may have created a very different response pattern, one that was more different than it was in previous studies. This is especially true because the primes did not pertain to the whole response, in the sense that they primed only a single word in a three-word utterance. Previous studies of single-word production in PWI did not have this difficulty. One possibility for a future experiment that would bring the two SOA or ISI valences in closer comparison would be to use a multi-word prime. Unfortunately, this raises difficulties of its own, namely controlling the duration of the primes to make them reasonably comparable.

6.5 Conclusion

I began with the question of whether processing facilitation affects the latency and duration of multi-word utterances. The results suggest that the answer is yes, with some limitations. Lexical selection and retrieval affect planning, but have an effect on duration that is restricted in time. Phonological retrieval and integration also affect planning, but have a widespread effect on duration compared to the semantic effects. Latency and duration did not always pattern together, suggesting that the determination of duration is not always made before speakers begin their utterances, and that they sometimes modulate duration online, as they speak. Identity priming had more complex effects than straightforward reduction, and, like the other relatedness effects, sometimes affected latency and duration differently. The takeaway is thus similar for all types of relatedness – facilitated processing affects word duration, but only when the relevant information is available while the production system can use it to advance the current state of processing. If the prime appeared too late to affect the relevant process(es), it did not influence latency or duration. Consistency had a similar effect as well, reducing action word durations but also affect object word durations, but only when the information was available at the right time.

APPENDIX

Part A: Example lists

The following are three of the 12 lists used in Experiment 1. The lefthand column is the name of the target picture. The middle column is the condition in which the target appeared, and the rightmost is the ISI at which it appeared. Note, for example, that the first row has airplane appearing in three different conditions, and the same ISI.

airplane	identical	900		airplane	control	900		airplane	unrelated	900
alligator	control	1500		alligator	unrelated	1500		alligator	related	1500
axe	related	1500		axe	identical	1500		axe	control	1500
ball	related	900		ball	identical	900		ball	control	900
balloon	related	900		balloon	identical	900		balloon	control	900
banana	unrelated	300		banana	related	300		banana	identical	300
barn	control	900		barn	unrelated	900		barn	related	900
barrel	control	300		barrel	unrelated	300		barrel	related	300
bear	unrelated	1500		bear	related	1500		bear	identical	1500
bed	unrelated	1500		bed	related	1500		bed	identical	1500
bell	control	900		bell	unrelated	900		bell	related	900
book	identical	1500		book	control	1500		book	unrelated	1500
boot	related	1500		boot	identical	1500		boot	control	1500
bottle	unrelated	900		bottle	related	900		bottle	identical	900
bowl	unrelated	300		bowl	related	300		bowl	identical	300
bread	related	1500		bread	identical	1500		bread	control	1500
broom	related	900		broom	identical	900		broom	control	900
bus	identical	1500		bus	control	1500		bus	unrelated	1500
cannon	unrelated	900		cannon	related	900		cannon	identical	900
car	identical	300		car	control	300		car	unrelated	300
cat	identical	900		cat	control	900		cat	unrelated	900
chain	related	300		chain	identical	300		chain	control	300
chair	control	900		chair	unrelated	900		chair	related	900
chicken	unrelated	300		chicken	related	300		chicken	identical	300
clock	unrelated	1500		clock	related	1500		clock	identical	1500
clown	control	1500		clown	unrelated	1500		clown	related	1500
comb	related	300		comb	identical	300		comb	control	300
corn	identical	300		corn	control	300		corn	unrelated	300
cow	unrelated	1500		cow	related	1500		cow	identical	1500
crown	unrelated	900		crown	related	900		crown	identical	900
cup	related	900		cup	identical	900		cup	control	900
deer	control	300		deer	unrelated	300		deer	related	300
dog	related	300		dog	identical	300		dog	control	300
door	identical	1500		door	control	1500		door	unrelated	1500
dress	related	1500		dress	identical	1500		dress	control	1500
drum	related	1500		drum	identical	1500		drum	control	1500
fish	control	900		fish	unrelated	900		fish	related	900

flag	unrelated	300		**flag**	related	300		**flag**	identical	300
flower	control	1500		**flower**	unrelated	1500		**flower**	related	1500
foot	control	1500		**foot**	unrelated	1500		**foot**	related	1500
football	identical	900		**football**	control	900		**football**	unrelated	900
frog	control	1500		frog	unrelated	1500		frog	related	1500
giraffe	unrelated	900		giraffe	related	900		giraffe	identical	900
glasses	related	1500		**glasses**	identical	1500		**glasses**	control	1500
glove	related	300		glove	identical	300		glove	control	300
grapes	identical	300		grapes	control	300		grapes	unrelated	300
guitar	control	900		**guitar**	unrelated	900		**guitar**	related	900
gun	identical	1500		**gun**	control	1500		**gun**	unrelated	1500
hair	identical	300		**hair**	control	300		**hair**	unrelated	300
hammer	unrelated	300		hammer	related	300		hammer	identical	300
heart	unrelated	300		**heart**	related	300		**heart**	identical	300
horse	related	300		**horse**	identical	300		**horse**	control	300
house	control	1500		**house**	unrelated	1500		**house**	related	1500
knife	unrelated	900		**knife**	related	900		**knife**	identical	900
lamp	identical	900		**lamp**	control	900		**lamp**	unrelated	900
leaf	control	900		**leaf**	unrelated	900		**leaf**	related	900
leg	control	300		**leg**	unrelated	300		**leg**	related	300
lettuce	identical	1500		lettuce	control	1500		lettuce	unrelated	1500
lips	unrelated	900		**lips**	related	900		**lips**	identical	900
monkey	identical	900		monkey	control	900		monkey	unrelated	900
mountain	unrelated	1500		**mountain**	related	1500		**mountain**	identical	1500
mushroom	unrelated	900		mushroom	related	900		mushroom	identical	900
nose	control	900		**nose**	unrelated	900		**nose**	related	900
owl	related	300		owl	identical	300		owl	control	300
pants	control	300		pants	unrelated	300		pants	related	300
peacock	related	300		peacock	identical	300		peacock	control	300
pencil	identical	1500		**pencil**	control	1500		**pencil**	unrelated	1500
penguin	unrelated	300		penguin	related	300		penguin	identical	300
piano	related	900		**piano**	identical	900		**piano**	control	900
pig	control	1500		pig	unrelated	1500		pig	related	1500
pumpkin	identical	300		pumpkin	control	300		pumpkin	unrelated	300
ring	related	1500		**ring**	identical	1500		**ring**	control	1500
saw	unrelated	900		**saw**	related	900		**saw**	identical	900
screw	related	1500		**screw**	identical	1500		**screw**	control	1500
sheep	unrelated	1500		**sheep**	related	1500		**sheep**	identical	1500
shoe	identical	300		**shoe**	control	300		**shoe**	unrelated	300
skunk	unrelated	1500		skunk	related	1500		skunk	identical	1500
snake	control	1500		**snake**	unrelated	1500		**snake**	related	1500
snowman	identical	900		snowman	control	900		snowman	unrelated	900
spider	unrelated	1500		spider	related	1500		spider	identical	1500
spoon	control	300		spoon	unrelated	300		spoon	related	300
squirrel	related	900		squirrel	identical	900		squirrel	control	900
star	identical	900		**star**	control	900		**star**	unrelated	900
suitcase	identical	300		**suitcase**	control	300		**suitcase**	unrelated	300
swing	related	300		**swing**	identical	300		**swing**	control	300
table	control	300		**table**	unrelated	300		**table**	related	300
tie	unrelated	300		**tie**	related	300		**tie**	identical	300
tiger	control	900		tiger	unrelated	900		tiger	related	900
train	identical	300		**train**	control	300		**train**	unrelated	300
truck	related	900		**truck**	identical	900		**truck**	control	900
turtle	identical	1500		turtle	control	1500		turtle	unrelated	1500
umbrella	identical	900		umbrella	control	900		umbrella	unrelated	900
wagon	identical	1500		**wagon**	control	1500		**wagon**	unrelated	1500
well	control	300		**well**	unrelated	300		**well**	related	300
wheel	related	900		**wheel**	identical	900		**wheel**	control	900
windmill	control	300		windmill	unrelated	300		windmill	related	300

<u>Part B: Blocking design for Experiments 2 and 4</u>

These experiments manipulated consistency, ISI, and condition relative to a single target word. Targets appeared twice in each experiment, but changed ISIs and consistencies. If a target appeared in the first half of the experiment in the random-1500 pair of levels, it appeared in the second half of the experiment at the 300-consistent pair of levels. Condition was pseudo-randomized, as can be seen with "airplane" below. The top two lists are lists 1 and 2, and the bottom are lists 7 and 8. Participants who saw list 1 in the first half of the experiment would see list 7 in the second half, and vice versa. Participants who saw list 2 in the first half would see list 8 in the second half, continuing through the pairing of list 6 and list 12. This ensured that no participant saw an item more than twice, and saw each item in a different set of levels of the independent variables.

airplane	identical	1500	random		airplane	control	1500	random
alligator	control	300	consistent		alligator	control	300	consistent
axe	related	1500	random		axe	identical	1500	random
ball	related	300	consistent		ball	related	300	consistent
balloon	related	300	consistent		balloon	related	300	consistent
banana	unrelated	1500	consistent		banana	related	1500	consistent
barn	control	1500	random		barn	unrelated	1500	random
barrel	control	1500	consistent		barrel	unrelated	1500	consistent
bear	unrelated	1500	random		bear	related	1500	random
bed	unrelated	1500	random		bed	related	1500	random
bell	control	1500	random		bell	unrelated	1500	random
airplane	related	300	consistent		airplane	identical	300	consistent
alligator	identical	1500	random		alligator	identical	1500	random
axe	unrelated	300	consistent		axe	related	300	consistent
ball	unrelated	1500	random		ball	unrelated	1500	random
balloon	unrelated	1500	random		balloon	unrelated	1500	random
banana	control	300	random		banana	unrelated	300	random
barn	identical	300	consistent		barn	control	300	consistent
barrel	identical	300	random		barrel	control	300	random
bear	control	300	consistent		bear	unrelated	300	consistent
bed	control	300	consistent		bed	unrelated	300	consistent
bell	identical	300	consistent		bell	control	300	consistent

Part C: Semantic and Phonological primes for each item

TARGET	SEMANTIC PRIME	PHONOLOGICAL PRIME
airplane	rocket	error
alligator	crocodile	alley
axe	hatchet	act
ball	globe	bomb
balloon	streamer	believe
banana	pear	bungalow
barn	stable	bard
barrel	keg	berry
bear	lion	bale
bed	futon	bend
bell	gong	bet
book	newspaper	bun
boot	slipper	boon
bottle	jar	bottom
bowl	saucer	bone
bread	cracker	breath
broom	mop	brood
bus	trolley	dome
cannon	crossbow	candle

car	moped	cap
cat	ferret	cab
chain	necklace	chase
chair	desk	change
chicken	rooster	chitter
clock	alarm	cloth
clown	mime	clod
comb	brush	code
corn	carrot	court
cow	bison	cob
crown	tiara	crop
cup	mug	cub
deer	elk	deep
dog	wolf	doll
door	drawer	gut
dress	skirt	drain
drum	bongo	droop
fish	dolphin	fist
flag	banner	flip
flower	vine	flouted
foot	ankle	four

football	helmet	futon
frog	toad	frost
giraffe	cheetah	jerk
glasses	contacts	glamour
glove	mitten	glum
grapes	plum	grain
guitar	banjo	giddy
gun	taser	lended
hair	mustache	hate
hammer	wrench	hamster
heart	lung	hard
horse	buffalo	horn
house	apartment	howl
knife	plate	night
lamp	flashlight	land
leaf	grass	leap
leg	fin	left
lettuce	cucumber	pending
lips	teeth	lid
monkey	baboon	Monday
mountain	hill	mouthing

mushroom	fern	mustard
nose	snout	note
owl	swallow	hour
pants	trousers	path
peacock	cardinal	people
pencil	paper	turnstile
penguin	seal	peddler
piano	organ	peering
pig	boar	pit
pumpkin	gourd	pummel
ring	bracelet	rim
saw	drill	sod
screw	nail	scrap
sheep	ram	sheen
shoe	sock	show
skunk	fox	scuttle
snake	gecko	snail
snowman	scarecrow	stall
spider	roach	spying
spoon	ladle	spool
squirrel	rat	squirted

star	planet	umpire
suitcase	backpack	super
swing	slide	sweet
table	bench	taper
tie	scarf	type
tiger	panther	tightly
train	subway	trail
truck	motorcycle	trust
turtle	lizard	bull
umbrella	poncho	snorer
wagon	caravan	waxy
well	canal	west
wheel	tire	weed
windmill	pinwheel	winter

Part D: Parameter estimate tables

The following are the tables for the parameter estimates of the models reported in the main text. Experiments 1 and 3 have six tables each, while Experiments 2 and 4 have four tables each. The same control variables were used in each model, and so the majority of each table is identical among the experiments. Experiments 2 and 4 differ slightly in that the parameter estimates for the variables of interest are also reported, unlike the variables of interest in Experiments 1 and 3. This is because of the different methods of assessing significance used between the experiments. The condition parameter estimates in Experiments 1 and 3 were not so important because the significance of the condition was assessed using multiple comparisons. In Experiments 2 and 4, *t*-values were treated as significant if they exceeded a magnitude of 2, following Baayen (2008). The values reported below are *t*-values. Values greater than 2 are printed in bold. Estimates that did not exceed a magnitude of 1.5 are indicated with two dashes (--) to reduce clutter.

Table 19: Parameter estimates for Experiment 1, at -1500SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	***-3.63***	**--**	**--**	**--**
Syllables	*-1.74*	***-2.18***	***6.59***	**--**
Imageability	**--**	**--**	**--**	**--**
Familiarity	**--**	**--**	**--**	**--**
Visual Complexity	**--**	**--**	**--**	**--**
Log Frequency	**--**	*1.56*	***-3.59***	**--**
Log Onset Duration	***N/A***	***3.94***	*1.76*	**--**
Log *the* Duration	***4.19***	***N/A***	***2.52***	**--**
Log Object Duration	***2.14***	*1.99*	***N/A***	**--**
Log Action Duration	**--**	***3.09***	***2.65***	***N/A***

Table 20: Parameter estimates for Experiment 1, at -900SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	--	***-2.2***	*-1.7*	--
Syllables	***-3.28***	*-1.79*	***6.46***	--
Imageability	--	--	--	--
Familiarity	***-2.23***	--	--	--
Visual Complexity	***-2.07***	--	--	--
Log Frequency	--	***2.27***	***-3.73***	--
Log Onset Duration	***N/A***	***2.88***	***2.89***	--
Log *the* Duration	***3.15***	***N/A***	***3.51***	--
Log Object Duration	***3.65***	***3.3***	***N/A***	--
Log Action Duration	--	--	--	***N/A***

Table 21 Parameter estimates for Experiment 1, at -300SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	***-2.06***	**--**	**--**	**--**
Syllables	*-1.99*	**--**	***6.99***	**--**
Imageability	**--**	**--**	**--**	**--**
Familiarity	***-2.19***	**--**	**--**	**--**
Visual Complexity	***-3.78***	**--**	**--**	**--**
Log Frequency	**--**	**--**	***-4.2***	**--**
Log Onset Duration	***N/A***	***3.05***	***3.7***	**--**
Log *the* Duration	***2.8***	***N/A***	**--**	**--**
Log Object Duration	***3.67***	**--**	***N/A***	**--**
Log Action Duration	**--**	**--**	**--**	***N/A***

Table 22: Parameter estimates for Experiment 1, at 300SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	***-2.95***	**--**	*1.88*	*-1.48*
Syllables	**--**	**--**	***6.72***	**--**
Imageability	**--**	**--**	**--**	**--**
Familiarity	***-2.28***	*-1.69*	**--**	***-2.1***
Visual Complexity	**--**	**--**	**--**	**--**
Log Frequency	***2.17***	**--**	***-3.48***	*1.95*
Log Onset Duration	***N/A***	***3.2***	**--**	***-10.81***
Log *the* Duration	***3.26***	***N/A***	*-1.9*	***2.02***
Log Object Duration	*1.66*	***-2.03***	***N/A***	**--**
Log Action Duration	***-10.05***	***2.49***	**--**	***N/A***

Table 23: Parameter estimates for Experiment 1, at 900SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	**--**	**--**	***3.53***	**--**
Syllables	**--**	**--**	***6.68***	**--**
Imageability	**--**	**--**	**--**	**--**
Familiarity	**--**	**--**	**--**	***-2.36***
Visual Complexity	*-1.76*	**--**	**--**	**--**
Log Frequency	**--**	**--**	***-4.02***	**--**
Log Onset Duration	***N/A***	***3.25***	***2.3***	***-10.77***
Log *the* Duration	***3.92***	***N/A***	***-3.74***	**--**
Log Object Duration	***2.25***	***-3.53***	***N/A***	***-3.05***
Log Action Duration	***-9.76***	**--**	***-2.41***	***N/A***

Table 24: Parameter estimates for Experiment 1, at 1500SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	--	--	--	--
Syllables	--	--	***6.89***	--
Imageability	--	--	--	--
Familiarity	--	--	--	*-1.78*
Visual Complexity	--	--	--	--
Log Frequency	--	--	***-4.3***	***2.02***
Log Onset Duration	***N/A***	***2.11***	--	***-13.13***
Log *the* Duration	***2.04***	***N/A***	--	*1.7*
Log Object Duration	--	--	***N/A***	--
Log Action Duration	***-12.18***	***2.32***	--	***N/A***

Table 25: Parameter estimates for the related vs. unrelated model for Experiment 2, at the negative SOAs (-1500 and -300)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	***4.46***	*-2.3*	--	--
Syllables	***-1.57***	*-2*	*6.13*	--
Imageability	--	--	--	--
Familiarity	--	--	--	-2.28
Visual Complexity	-2.05	--	--	--
Log Frequency	--	--	***-3.42***	*2.77*
Log Onset Duration	N/A	4.14	*1.9*	*2.15*
Log *the* Duration	***3.57***	N/A	***2.4***	--
Log Object Duration	***1.99***	***2.03***	*N/A*	***2.15***
Log Action Duration	***1.85***	--	***2.06***	*N/A*
Condition	*3.82*	--	--	--
Consistency	*-6.6*	--	*-4.6*	--
SOA	--	***3.34***	--	--
Condition * Consistency	--	--	-2.08	--
Condition * SOA	--	--	--	--
Consistency * SOA	--	-1.7	--	2.92
Three-way Interaction	-1.6	--	1.55	-1.89

Table 26: Parameter estimates for the identical vs. control model for Experiment 2, at the negative SOAs (-1500 and -300)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	**--**	*-3.08*	--	--
Syllables	***-2.53***	--	*7.77*	--
Imageability	--	--	--	--
Familiarity	*-1.81*	--	--	-2
Visual Complexity	-3.35	--	--	--
Log Frequency	--	--	***-4.62***	*2.36*
Log Onset Duration	N/A	2.93	*3.35*	--
Log *the* Duration	***2.89***	N/A	***-2.63***	***1.83***
Log Object Duration	***3.55***	***-2.91***	*N/A*	***2.01***
Log Action Duration	**--**	--	***1.77***	*N/A*
Condition	*-11.81*	**--**	*-2.63*	--
Consistency	*-10.19*	--	*-2.08*	*-4*
SOA	***1.83***	--	**--**	--
Condition * Consistency	--	--	--	--
Condition * SOA	2.05	--	-2.05	--
Consistency * SOA	-2.01	--	--	--
Three-way Interaction	--	--	--	2.47

Table 27: Parameter estimates for the related vs. unrelated model in Experiment 2, at the positive SOAs (300 and 1500)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	--	*2.08*	*5.48*	*3.91*
Syllables	--	*-1.64*	5.42	--
Imageability	--	--	--	--
Familiarity	--	--	--	--
Visual Complexity	-2.13	--	--	--
Log Frequency	--	--	*-4.23*	*2.17*
Log Onset Duration	N/A	-1.94	*4.3*	--
Log *the* Duration	***-2.02***	N/A	--	***3.87***
Log Object Duration	***4.25***	--	*N/A*	***2.45***
Log Action Duration	--	*4.32*	***2.17***	--
Condition	--	***-2.89***	--	--
Consistency	--	*3.45*	*1.65*	*-3.08*
SOA	--	***2.39***	--	--
Condition * Consistency	--	--	--	*-1.56*
Condition * SOA	--	--	--	--
Consistency * SOA	--	--	2.06	--
Three-way Interaction	--	--	-1.82	--

Table 28: Parameter estimates for the identical vs. control model in Experiment 2, at the positive SOAs (300 and 1500)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	***-3.64***	--	*5.97*	*3.26*
Syllables	--	--	3.04	*1.82*
Imageability	--	--	--	*-1.77*
Familiarity	--	--	--	-2.99
Visual Complexity	--	--	--	--
Log Frequency	--	1.78	*-2.62*	*1.8*
Log Onset Duration	N/A	--	*2.62*	--
Log *the* Duration	--	N/A	***1.64***	--
Log Object Duration	***5.38***	--	*N/A*	--
Log Action Duration	--	*3.56*	--	*N/A*
Condition	--	--	*-3.14*	*0.83*
Consistency	*-2.04*	*2.28*	*2.86*	*-2.83*
SOA	*-2.27*	--	--	*2.05*
Condition * Consistency	--	--	--	--
Condition * SOA	--	--	--	--
Consistency * SOA	*-1.5*	--	--	--
Three-way Interaction	--	1.66	2.63	--

Table 29: Parameter estimates for Experiment 3, at -1500SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	--	--	--	--
Syllables	--	--	***6.96***	--
Imageability	--	--	--	--
Familiarity	--	--	*-1.54*	***-2.58***
Visual Complexity	***-2.72***	--	--	--
Log Frequency	--	--	***-2.84***	*1.75*
Log Onset Duration	***N/A***	*1.55*	***2.54***	--
Log *the* Duration	--	***N/A***	***2.02***	***2.99***
Log Object Duration	***3.82***	***2.65***	***N/A***	--
Log Action Duration	--	***3.82***	--	***N/A***

Table 30: Parameter estimates for Experiment 3, at -900SOA

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	***-2.27***	*1.65*	--	--
Syllables	--	--	***5.8***	--
Imageability	--	--	--	--
Familiarity	***-2.02***	--	--	*-1.75*
Visual Complexity	*-1.54*	--	--	--
Log Frequency	--	--	***-4.11***	*1.82*
Log Onset Duration	***N/A***	***2.21***	*1.69*	--
Log *the* Duration	***2.45***	***N/A***	--	***5.2***
Log Object Duration	*1.96*	--	***N/A***	--
Log Action Duration	--	***5.79***	--	***N/A***

Table 31: Parameter estimates for Experiment 3, at -300SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	***--***	*-1.68*	***--***	*1.61*
Syllables	*-1.92*	***--***	***5.94***	***--***
Imageability	***--***	***--***	***--***	***--***
Familiarity	***-2.05***	***--***	***--***	***-3.24***
Visual Complexity	***--***	***--***	***--***	***--***
Log Frequency	***--***	***2.33***	***-3.7***	***--***
Log Onset Duration	***N/A***	***3.35***	***2.64***	***--***
Log *the* Duration	***2.48***	***N/A***	***4.88***	***3.07***
Log Object Duration	***4.05***	***5.88***	***N/A***	***-3.23***
Log Action Duration	***--***	***3.91***	*-2*	***N/A***

Table 32: Parameter estimates for Experiment 3, at 300SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	**--**	**--**	***2.51***	**--**
Syllables	**--**	***-2.23***	***5.4***	**--**
Imageability	***2.1***	**--**	**--**	**--**
Familiarity	**--**	**--**	**--**	***-3.04***
Visual Complexity	**--**	**--**	**--**	**--**
Log Frequency	*1.67*	**--**	***-4.12***	***2.16***
Log Onset Duration	***N/A***	***4.15***	***5.17***	***-11.51***
Log *the* Duration	***4.21***	***N/A***	**--**	***2.59***
Log Object Duration	***5***	**--**	***N/A***	**--**
Log Action Duration	***-10.5***	***2.78***	**--**	***N/A***

Table 33: Parameter estimates for Experiment 3, at 900SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	**--**	**--**	**--**	***2.01***
Syllables	**--**	**--**	***6.44***	**--**
Imageability	**--**	**--**	**--**	**--**
Familiarity	**--**	**--**	**--**	*-1.84*
Visual Complexity	**--**	**--**	**--**	**--**
Log Frequency	**--**	**--**	***-4.19***	**--**
Log Onset Duration	***N/A***	***6.12***	*1.89*	***-10.1***
Log *the* Duration	***5.18***	***N/A***	**--**	***3.24***
Log Object Duration	***3.21***	**--**	***N/A***	***-2.15***
Log Action Duration	***-9.34***	***3.32***	*-2*	***N/A***

Table 34: Parameter estimates for Experiment 3, at 1500SOA

	Latency to Speak	**Determiner (*the*)**	**Target word**	**Action word**
Trial	***-2.15***	**--**	***2.82***	**--**
Syllables	**--**	**--**	***5.69***	**--**
Imageability	**--**	**--**	**--**	**--**
Familiarity	**--**	**--**	**--**	***-3.57***
Visual Complexity	***-2.14***	**--**	**--**	**--**
Log Frequency	**--**	**--**	***-3.59***	***2.03***
Log Onset Duration	***N/A***	***3.44***	***3.52***	***-11.03***
Log *the* Duration	***3.46***	***N/A***	*1.68*	***3.18***
Log Object Duration	***4.5***	*1.5*	***N/A***	**--**
Log Action Duration	***-10***	***3.28***	**--**	***N/A***

Table 35: Parameter estimates for the related vs. unrelated model in Experiment 4, at the negative SOAs (-1500 and -300)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	***1.8***	--	--	--
Syllables	--	*-2.17*	*7.64*	--
Imageability	--	--	--	--
Familiarity	*-2.21*	--	-1.68	--
Visual Complexity	-4.5	--	--	--
Log Frequency	--	--	--	--
Log Onset Duration	N/A	5.44	--	--
Log *the* Duration	***5.38***	N/A	--	***2.88***
Log Object Duration	--	--	*N/A*	***3.03***
Log Action Duration	***1.74***	*3.27*	***3.46***	*N/A*
Condition	*3.91*	--	--	*1.76*
Consistency	*-7.36*	*-2.89*	*-4.26*	*-4.7*
SOA	***-2.75***	--	***-1.54***	--
Condition * Consistency	--	--	-1.74	--
Condition * SOA	--	--	--	--
Consistency * SOA	--	--	--	2.35
Three-way Interaction	--	--	--	1.72

Table 36: Parameter estimates for the identical vs. control model in Experiment 4, at the negative SOAs (-1500 and -300)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	--	--	--	*-3.51*
Syllables	--	--	*7.77*	--
Imageability	--	--	--	--
Familiarity	--	--	--	-1.9
Visual Complexity	-2.37	--	--	--
Log Frequency	--	1.8	***-3.72***	*2.83*
Log Onset Duration	N/A	--	*2.85*	--
Log *the* Duration	--	N/A	--	***4.2***
Log Object Duration	***3.34***	--	*N/A*	***5.27***
Log Action Duration	--	*4.46*	***6.12***	*N/A*
Condition	*-18.34*	***-1.58***	--	--
Consistency	*-10.14*	--	--	*-2.3*
SOA	***-1.92***	--	--	--
Condition * Consistency	-2.97	--	--	*1.77*
Condition * SOA	--	--	--	--
Consistency * SOA	-3.84	--	--	--
Three-way Interaction	--	--	-1.67	--

Table 37: Parameter estimates for the related vs. unrelated model in Experiment 4, at the positive SOAs (300 and 1500)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	--	*1.71*	--	*5.47*
Syllables	--	--	*5.83*	--
Imageability	--	--	--	--
Familiarity	--	--	--	-3.17
Visual Complexity	--	--	--	--
Log Frequency	--	1.98	***-4.24***	*2.63*
Log Onset Duration	N/A	--	*2.3*	--
Log *the* Duration	--	N/A	***1.87***	***1.64***
Log Object Duration	***3***	***2.08***	*N/A*	***2.58***
Log Action Duration	--	*2.33*	***3.04***	*N/A*
Condition	*-1.7*	--	*3.61*	--
Consistency	*-2.46*	--	--	*-6*
SOA	--	--	***-1.8***	*1.7*
Condition * Consistency	--	--	--	--
Condition * SOA	--	--	-1.76	--
Consistency * SOA	--	--	--	1.79
Three-way Interaction	--	--	--	--

Table 38: Parameter estimates for the identical vs. control model in Experiment 4, at the positive SOAs (300 and 1500)

	Latency to Speak	Determiner (*the*)	Target word	Action word
Trial	--	--	*6.57*	*2.97*
Syllables	***2.61***	--	*6.28*	--
Imageability	--	--	--	--
Familiarity	--	--	--	-2.78
Visual Complexity	--	--	--	--
Log Frequency	--	--	***-4.05***	*2.05*
Log Onset Duration	N/A	--	--	--
Log *the* Duration	***-1.5***	N/A	--	--
Log Object Duration	--	--	*N/A*	--
Log Action Duration	--	*1.84*	--	*N/A*
Condition	*1.63*	***-1.66***	*-3.41*	--
Consistency	--	*1.58*	--	*-5.15*
SOA	--	--	--	--
Condition * Consistency	--	--	--	--
Condition * SOA	1.78	--	--	--
Consistency * SOA	2.52	--	--	--
Three-way Interaction	--	--	--	--

REFERENCES

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

Ariel, M. (2004). Accessibility Marking: Discourse Functions, Discourse Profiles, and Processing Cues. Discourse Processes, 37(2), 91-116. doi:10.1207/s15326950dp3702_2

Arnold, J.E., Kahn, J.M. & Pancani, G. (2012) Audience Design Affects Acoustic Reduction Via Production Facilitation. Psychological Bulletin & Review, 19, 505-512

Arnold, J. E., Pancani, G. P., & Rosa, E. C.. (under revision). Listeners perceive acoustic prominence differently for distracted and fluent speakers. Ms., University of North Carolina at Chapel Hill.

Arnold, J. E. & Watson, D. G. (2013). Synthesizing meaning and processing approaches to prosody: performance matters. Language, Cognition, and Neuroscience.

Arnold, J. E., Wasow, T., Losongco, T., & Ginstrom, R. (2000). Heaviness vs. newness: The effects of structural complexity and discourse status on constituent ordering. Language, 76, 28-55.

Arnold, J.E. & Watson, D. (under review). Synthesizing meaning and processing approaches to prosody: performance matters.

Aylett, M. & Turk, A. (2004). The Smooth Signal Redundancy Hypothesis: A Functional Explanation for Relationships between Redundancy, Prosodic Prominence, and Duration in Spontaneous Speech. Language and Speech, 47 (1), 31-56.

Baayen, R. H. (2008). Analyzing Linguistic Data. A Practical Introduction to Statistics Using R. Cambridge University Press.

Baese-Berk, M., & Goldrick, M. (2009). Mechanisms of interaction in speech production. Language and cognitive processes, 24(4), 527-554. doi:10.1080/01690960802299378

Balota, D.A., Boland, J.E. & Shields, L.W. (1989). Priming in pronunciation: Beyond pattern recognition and onset latency. Journal of Memory and Language, 28(1), 14-36.

Beattie, G. (1983). Talk: An analysis of speech and non-verbal behavior in conversation. Milton Keyes: Open University Press.

Belke, E., Meyer, A. S., & Damian, M. F. (2005). Refractory effects in picture naming as assessed in a semantic blocking paradigm. The Quarterly Journal of Experimental Psychology, 58(4), 667-692.

Bell, A., Brenier, J. M., Gregory, M., Girand, C., & Jurafsky, D. (2009). Predictability effects on durations of content and function words in conversational English. Journal of Memory and Language, 60, 92-111.

Bell, A., Jurafsky, D., Fosler-Lussier, E., Girand, C., Gregory, M., & Gildea, D. (2003). Effects of disfluencies, predictability, and utterance position on word form variation in English conversation. Journal of the Acoustical Society of America, 113, 1001–1024.

Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. Psychological Review, 89, 1–47.

Bock, J.K., Irwin, D. E. (1980). Syntactic Effects of Information Availability in Sentence Production. Journal of Verbal Learning and Verbal Behavior, 19, 467-484.

Bock, K., Irwin, D. E., Davidson, D. J., & Levelt, W. J. (2003). Minding the clock. Journal of Memory and Language, 48(4), 653-685.Gleitman

Boersma, P. & Weenink, D. (2009). Praat: doing phonetics by computer (Version 5.1.15) [Computer program]. Retrieved August 30, 2009, from http://www.praat.org/

Brennan, S. E. & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. Journal of Experimental Psychology: Learning, Memory and Cognition, 22(6),482-1493. doi: 10.1.1.121.3930

Brown, G. (1983). Prosodic Structure and the Given/New Distinction. A. Cutler & D. R. Ladd (Eds.), Prosody: Models and Measurements (pp. 67-77).

Brown-Schmidt, S. (2012). Beyond common and privileged: Gradient representations of common ground in real-time language use. Language and Cognitive Processes 27(1), 62-89.

Brown-Schmidt, S., & Tanenhaus, M. K. (2008). Real-time investigation of referential domains in unscripted conversation: A targeted language game approach. Cognitive science, 32(4), 643-684.

Bybee, Joan. 1999. Use impacts morphological representation. Behavioral and Brain Sciences 22: 1016-1017.

Chafe, W. (1976) Givenness, contrastiveness, definiteness, subjects, topics and point of view. In Li, C. (ed.). Subject and Topic. New York: Academic Press.

Chafe, W. L. (1976). Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In C. N. Li (Ed.), Subject and topic (pp. 25-56). New York: Academic Press.

Clark, H. H. & Fox Tree, J. E. (2002). Using uh and um in spontaneous speaking. Cognition, 84, 73-111.

Clark, H. H., & Wasow, T. (1998). Repeating words in spontaneous speech.Cognitive psychology, 37(3), 201-242.

Corbett, G. G. (2006). Agreement (Vol. 109). Cambridge University Press.

Cutler, A., Dahan, D., & Van Donselaar, W. A. (1997). Prosody in the comprehension of spoken language: A literature review. Language and Speech, 40, 141-202.

Dahan, D., Tanenhaus, M. K., & Chambers, C. G. (2002). Accent and reference resolution in spoken-language comprehension. Journal of Memory and Language, 47, 292-314.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. Psychological review, 93(3), 283-321.

Dell, G. S., Oppenheim, G. M., & Kittredge, A. K. (2008). Saying the right word at the right time: Syntagmatic and paradigmatic interference in sentence production. Invited paper for Language and Cognitive Processes, 23, 583-608.

Eberhard, K. M., Cutting, J. C., & Bock, K. (2005). Making syntax of sense: number agreement in sentence production. Psychological Review, 112(3), 531.

Ferreira, F. (1993). The creation of prosody during sentence production. Psychological Review, 100, 233-253.

Ferreira, F. (2000). Syntax in language production: An approach using tree-adjoining grammars. In L. Wheeldon (Ed.), Aspects of Language Production. Cambridge, MA: MIT Press.

Fosler-Lussier, E., & Morgan, N. (1999). Effects of speaking rate and word predictability on conversational pronunciations. Speech Communication, 29, 137–158.

Fowler, C. A., & Housum, J. (1987). Talkers' signaling of "new" and "old" words in speech and listeners' perception and use of the distinction. Journal of Memory and Language, 26, 489-504.

Gahl, S. (2008). Time and thyme are not homophones: The effect of lemma frequency on word durations in spontaneous speech. Language, 84(3), 474-496.

Gahl, S. & Garnsey, S. M. (2004). Knowledge of grammar, knowledge of usage: syntactic probabilities affect pronunciation variation. Language, 80 (4), 748-775.

Gahl, S., Yao, Y., Johnson, K. (2012). Why reduce? Phonological neighborhood density and phonetic reduction in spontaneous speech. Journal of Memory and Language.

Galati, A. & Brennan, S.E. (2010). Attenuating information in spoken communication: For the speaker, or for the addressee? Journal of Memory and Language, 62(1), 35-51. doi:10.1016/j.jml.2009.09.002

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

Genzel, D., & Charniak, E. (2002, July). Entropy rate constancy in text. InProceedings of the 40th annual meeting on association for computational linguistics (pp. 199-206). Association for Computational Linguistics.

Glaser, W. R., & Glaser, M. O. (1989). Context effects in stroop-like word and picture processing. Journal of Experimental Psychology: General, 118(1), 13.

Glaser, W. R., & Düngelhoff, F. J. (1984). The time course of picture-word interference. Journal of Experimental Psychology: Human Perception and Performance, 10(5), 640.

Gleitman, L. R., January, D., Nappa, R., & Trueswell, J. C. (2007). On the give and take between event apprehension and utterance formulation. Journal of memory and language, 57(4), 544-569.

Griffin, Z. M., & Bock, K. (1998). Constraint, word frequency, and levels of processing in spoken word production. Journal of Memory and Language, 38, 313-338.

Goldrick, M., & Rapp, B. (2002). A restricted interaction account (RIA) of spoken word production: The best of both worlds. Aphasiology, 16(1), 20-55. doi:10.1080/02687040143000203

Gundel, J., N. Hedberg & R. Zacharski (1993) Cognitive status and the form of referring expressions in discourse. Language, 69(2), 274-307. doi:10.2307/416535

Gussenhoven, C. (1984). On the grammar and semantics of sentence accents(Vol. 16). Walter de Gruyter.Schwartzchild

Hothorn, T., Bretz, F., Westfall, P., Heiberger, R. M., & Schuetzenmeister, A. (2014). Multcomp: simultaneous inference in general parametric models. R package version, 1-3.

Jaeger, T.F. (2010). Redundancy and reduction: Speakers manage syntactic information density. Cognitive Psychology, 61(1), 23-62.

Jescheniak, J. D., Schriefers, H., & Hantsch, A. (2003). Utterance format effects phonological priming in the picture-word task: Implications for models of phonological encoding in speech production. Journal of Experimental Psychology: Human Perception and Performance, 29(2), 441-454. doi:10.1037/0096-1523.29.2.441

Jurafsky, D., Bell, A., Gregory, M., & Raymond, W. D. (2001). Probabilistic relations between words: Evidence from reduction in lexical production. In J. Bybee & P. Hopper (Eds.), Frequency and the emergence of linguistic structure (pp. 229-254). Amsterdam: Benjamin.

Kahn, J.M. & Arnold, J. E. (2012) A Processing-Centered Look at the Contribution of Givenness to Durational Reduction. Journal of Memory & Language, 67, 311-325

Kello, C. T., Plaut, D. C., & MacWhinney, B. (2000). The task dependence of staged versus cascaded processing: An empirical and computational study of Stroop interference in speech perception. Journal of Experimental Psychology: General, 129(3), 340-360. doi:10.1037//0096-3445.129.3.340

Konopka, A. E. (2012). Planning ahead: How recent experience with structures and words changes the scope of linguistic planning. Journal of Memory and Language, 66(1), 143-162.Chafe, W. L. (1994). Discourse, Consciousness, and Time. Chicago: Chicago University Press.

Ladd, D.R., (1996). Intonational Phonology: Cambridge Studies in Linguistics 79. Cambridge: Cambridge University Press.

Lam, T.Q. & Watson, D.G. (2010). Repetition is easy: Why Repeated Referents Have Reduced Prominence. Memory & Cognition 38(8), 1137-1146. doi: 10.3758/MC.38.8.1137

Levelt, W. J. (1989). Speaking: From intention to articulation. Cambridge, MA. MIT Press.

Levelt, W.J.M., Roelofs, A., & Meyer, A.S. (1999). A theory of lexical access in speech production. Behavioral and Brain Sciences, 22, 1-38.

Levelt, W. J., & Wheeldon, L. (1994). Do speakers have access to a mental syllabary?. Cognition, 50(1), 239-269.

Levy, R., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. In B. Schlokopf, J. Platt, & T. Homan (Eds.), Advances in neural information processing systems (NIPS) 19 (pp. 849-856). Cambridge, MA: MIT Press.

Lieberman, P. (1963). Some effects of the semantic and grammatical context on the production and preparation of speech. Language and Speech, 6, 172-175.

Lindblom, B. (1990). Explaining variation: A sketch of the H and H theory. In W. Hardcastle & A. Marchal (Eds.), Speech production and speech modeling (pp. 403-439). Dordrecht: Kluwer Academic.

Mahon, B. Z., Costa, A., Peterson, R., Vargas, K. A., & Caramazza, A. (2007). Lexical selection is not by competition: a reinterpretation of semantic interference and facilitation effects in the picture-word interference paradigm.Journal of Experimental Psychology: Learning, Memory, and Cognition, 33(3), 503.

Meyer, A. S., & Schriefers, H. (1991). Phonological facilitation in picture/word interference experiments: Effects of stimulus onset asynchrony and types of interfering stimuli. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17(6), 1146-1160. doi:10.1037//0278-7393.17.6.1146

Meyer, A. S., Sleiderink, A. M., & Levelt, W. J. (1998). Viewing and naming objects: Eye movements during noun phrase production. Cognition, 66(2), B25-B33.

Munson, B. (2001). Phonological pattern frequency and speech production in adults and children. Journal of Speech, Language, and Hearing Research, 44, 778-792.

Munson, B., & Solomon, N. P. (2004). The Effect of Phonological Neighborhood Density on Vowel Articulation. Hearing Research, 47, 1048-1058.

Pierrehumbert, J. B., & Hirschberg, J. (1990). The meaning of intonational contours in the interpretation of discourse. In P. R. Coehn et al. (Eds.), Intentions in Communication (pp. 271-311). Cambridge, MA: MIT Press.

Pluymaekers, M., Ernestus, M., & Baayen, R. H. (2005a). Articulatory planning is continuous and sensitive to informational redundancy. Phonetica, 62, 146–159.

Prince, E. (1992). The ZPG letter: Subjects, definiteness, and information-status. In S. Thomspon & W. Mann (Eds.), Discourse description: Diverse analyses of a fund raising text (pp. 295-325). Philadelphia: John Benjamins B.V.

Rahman, R. A., & Melinger, A. (2008). Enhanced phonological facilitation and traces of concurrent word form activation in speech production: An object-naming study with multiple distractors. The Quarterly Journal of Experimental Psychology, 61(9), 1410-1440.

Roelofs, A. (2002). Spoken language planning and the initiation of articulation. The Quarterly Journal of Experimental Psychology, 55A(2), 465-483. doi:10.1080/0272498014300048

Rossion, B., & Pourtois, G. (2001). Revisiting Snodgrass and Vanderwart's object database: Color and texture improve object recognition [On-line]. Journal of Vision, 1, 413a.

Schmitt, B. M., Meyer, A.S., & Levelt, W. J. (1999). Lexical access in the production of pronouns. Cognition, 69(3), 313-35.

Schnur, T. T., Costa, A., & Caramazza, A. (2006). Planning at the phonological level during sentence production. Journal of Psycholinguistic Research, 35(2), 189-213.

Schriefers, H., & Teruel, E. (1999). Phonological facilitation in the production of two-word utterances. European Journal of Cognitive Psychology, 11, 17–50.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, familiarity and visual complexity. Journal of Experimental Psychology: Human Learning & Memory, 6, 174-215.

Sityaev, D. (2000). The relationship between accentuation and information status of discourse referents : A corpus study. UCL Working Papers in Linguistics, 12, 285-303.

Shannon, C. E. (1948). A mathematical theory of communication. Bell Technical Systems Journal, 27, 379-423.

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

Starreveld, P. A., & La Heij, W. (1996). Time-course analysis of semantic and orthographic context effects in picture naming. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22(4), 896.Selkirk 1995

Vitevitch, M. S., Armbrüster, J., & Chu, S. (2004). Sublexical and lexical representations in speech production: effects of phonotactic probability and onset density. Journal of Experimental Psychology: Learning, Memory, and Cognition, 30(2), 514.

Watson, D. G. (2010). The many roads to prominence: Understanding emphasis in conversation. Psychology of Learning and Motivation, 163–183.

Watson, D.G., Arnold, J.E., & Tanenhaus, M.K. (2008). Tic tac TOE: Effects of predictability and importance on acoustic prominence in language production. Cognition, 106, 156-1557.

Watson, D. G., Buxó-Lugo, A., & Simmons, D. C. (to appear). The effect of phonological encoding on word duration: Selection takes time. In E. Gibson and L. Frazier, Explicit and Implicit Prosody in Sentence Processing.

Wheeldon, L. R. & Lahiri, A. (1997) Prosodic units in speech production. Journal of Memory and Language 37:356–81.

Yuan, J., & Liberman, M. (2008). Speaker identification on the SCOTUS corpus.Journal of the Acoustical Society of America, 123(5), 3878.

Zipf, G. (1929). Relative frequency as a determinant of phonetic change. Harvard Studies in Classical Philology, 40, 1-95.