Adaptive Eye Movement Control in a Simple Linguistic Task

by

Michael Shvartsman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Psychology)
in the University of Michigan
2014

Doctoral Committee:

Professor Richard L. Lewis, Co-Chair
Professor Satinder Singh, Co-Chair
Professor Julie Boland
Assistant Professor Jonathan Brennan
Associate Professor John T. Hale, Cornell University

A C K N O W L E D G M E N T S

I might get top billing on this thing by convention, but it took a village to shove me this far. Rather than be mortified for eternity about excluding one person I forgot, I would rather include everyone – but nobody by name. I even got paralyzed by indecision trying to list broad categories. Thanks for bearing with me. Y'all know who you are.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Adaptive Eye Movement Control in a Simple Linguistic Task

by

Michael Shvartsman

Co-Chairs: Richard L. Lewis and Satinder Singh

This dissertation pursues a computationally rational analysis of eye movements in a simple list-reading task. The strength of the computationally rational approach is in the ability to explain why certain phenomena may emerge under the assumption that behavior is an approximately optimal adaptation to the joint constraints of an organism's intrinsic computational constraints and task demands. The provided theory and model integrates a framework of lexical processing as active perception (Norris, 2006) with oculomotor constraints derived from a broad-coverage model of eye movement control in reading (Reichle, Warren & McConnell 2009). The first portion of the thesis provides experimental evidence of adaptation of fixation durations to quantitatively-expressed payoffs in a simple reading task, and adaptation in the model on the same dimension. The second portion explores spillover lexical frequency effects in the same framework and how they may emerge from a model that can adaptively allocate processing resources to information drawn from perception (foveal or parafoveal), or memory. In addition to implications for eye movement control in reading, these findings can be interpreted to bear on task adaptation in reading, as well as the adaptive use of perception and memory in a sequential sampling framework.

CHAPTER 1

Introduction

This dissertation is intended to advance the scientific understanding of human eye movement behavior in reading. It pursues this work from a perspective of computational rationality. In contrast to solely descriptive theories, the computational rationality approach provides explanations of how certain properties of the eye movement process may reflect adaptation to underlying constraints imposed by the reader's cognitive and behavioral architecture, task structure, and task goals. The key contributions are an adaptive model of eye movement control in reading, experimental evidence of adaptation of the same kind that the model exhibits, and an exploration of spillover frequency effects in a rational adaptive control framework. Broadly interpreted, these findings bear on task adaptation in reading, as well as the adaptive use of perception and memory in a sequential sampling framework.

The theoretical work integrates a sequential sampling model of lexical processing (Norris, 2006, 2009) with a cognitive and oculomotor machine architecture based on a state-of-the-art model of eye movement control in reading (Reichle, Warren & McConnell, 2009). The number of sequential samples used can be varied to produce different speed-accuracy tradeoffs. In this way, the sequential sampling component provides a parametric locus of adaptive control. The model generates predictions by finding near-optimal behavior with respect to this adaptive locus, as constrained by a particular machine architecture. These predictions are tested on humans performing the same task as the model, under the same speed-accuracy tradeoffs imposed via quantitatively-expressed payoffs. The task used is the List Lexical Decision Task (Meyer, Schvaneveldt & Ruddy, 1974), a simple widely-spaced list-reading task that has some useful similarities to sentence-reading in a more naturalistic context.

The remainder of this chapter motivates the approach and key theoretical components, provides an overview of the key results, and concludes with a roadmap of the remainder of the dissertation.

1.1 Rational approaches to understanding human behavior

The theoretical work in the thesis assumes that human behavior is an approximately optimal solution to the control problem posed by the combination of the human's environment, underlying cognitive architecture, and task goals. In doing so, it draws on earlier work in related formalisms such as bounded rationality (Simon, 1955), rational analysis (e.g. Anderson, 1991), and most recently computational rationality (Lewis, Howes & Singh, 2013). Outside of psychology, the work draws some of its terminology from reinforcement learning (Sutton & Barto, 1998), and draws on the precise statement of the assumption of *bounded optimality* provided by Russell & Subramanian (1995).

The rational perspective is supported by a body of work on the adaptation to task context and goals, long known to have major effects on human performance in a variety of cognitive tasks (e.g. Green & Swets, 1966). More recent work on visual attention in both linguistic and nonlinguistic contexts in particular indicates that attention strategies are strongly shaped by prevailing task goals (e.g., Rothkopf, Ballard & Hayhoe, 2007; Ballard & Hayhoe, 2009; see Salverda, Brown & Tanenhaus 2011 for a recent review). In general, effects of strategic adaptation penetrate all levels of human performance (Newell, 1973), from the most elementary perceptual decisions (Green & Swets, 1966) to more complex multi-tasking scenarios (Meyer & Kieras, 1997; Howes, Lewis & Vera, 2009).

In psycholinguistic research, task effects likewise have a longstanding history (for an early analysis see the seminal chapter by Forster, 1979). For example, in the area of single-word lexical processing, there are robust differences in how frequency and other important effects are manifest in naming vs. lexical decision tasks (e.g. Grainger, 1990). Task context in the form of experimental list composition and goal manipulation via instructional emphases have significant effects, and have received detailed theoretical treatments (Wagenmakers, Ratcliff, Gomez & McKoon, 2008).

There is also a small but growing line of empirical work demonstrating task effects on eye movements in reading. McConkie, Rayner & Wilson (1973) have shown that participants tend to read longer when anticipating difficult questions (for example, questions of a factual nature), as well as when they were financially incentivized to answer the questions correctly. Rayner & Raney (1996) have shown that the lexical frequency effect is eliminated when subjects read words in search of a target word rather than reading for comprehension. Wotschack (2009) replicated the question-difficulty finding of McConkie and colleagues and also found that increasing the frequency of comprehension questions, as well as instructing the participants to proofread, led to slower reading speeds. Finally, Reichle, Reineberg & Schooler (2010) polled participants to determine when they 'zoned out', and showed that the 'mindless' reading during those periods is significantly different from normal reading, again consistent with eye movements in reading being actively and adaptively controlled.

Rational models have also provided novel explanations for certain phenomena in reading. In particular, *ideal observer* models has shown how saccade targetting decisions might be a consequence of rational perceptual recognition (Legge, Klitz & Tjan, 1997; Legge, Hooven, Klitz, Mansfield & Tjan, 2002; Bicknell & Levy, 2010b), and how regressive eye movements might be a rational adaptation to falling confidence in past inputs (Bicknell & Levy, 2010a). In addition, work in a reinforcement learning context has shown how certain basic properties of eye movements like viewing position effects and parafoveal preview might emerge as an adaptive consequence of basic physiological and psychological constraints on the eye movement system (Reichle & Laurent, 2006; Liu & Reichle, 2010). The thesis adds to this body of work by showing how quantitative variation in speed-accuracy tradeoff imposed by quantiatively varying payoffs rationally yields qualitative and quantitative changes in reading strategy, and also by showing how and why spillover lexical frequency effects may emerge in response to the noise structure of the information processing system.

1.2 Key theoretical components of the thesis model

Lexical processing as sequential sampling. The lexical processing component in the model is implemented as a sequential probability ratio test (SPRT Wald, 1945; Wald & Wolfowitz, 1948) and one of its multihypothesis variants (Baum & Veeravalli, 1994; Draglia, Tartakovsky & Veeravalli, 1999). In doing so, the dissertation follows models of lexical decision and naming by Norris (2006, 2009) and models of reading by Bicknell & Levy (2010a,b). The SPRT is a sequential statistical test that works by computing the log likelihood ratio between hypotheses (i.e. words) as a function of incoming samples (i.e. noisy percepts), stopping when the ratio crosses a threshold. In addition to the successes above, it has also been used to explain neural firing patterns in the basal ganglia associated with perceptual decision-making (Bogacz & Gurney, 2007). Because the changing likelihoods over time form a random walk, the SPRT is also compatible with findings of random walk models like the Drift Diffusion Model

(DDM, Ratcliff, 1978), though Norris (2009) argues that the DDM is less parsimonious than the SPRT.

Oculomotor architecture. The oculomotor architecture used in the model is inspired by a longstanding model of eye movement control in reading, E-Z Reader (Reichle et al., 2009). The architectural constraints are realized as stochastic delays in the oculomotor system. The first delay is on the visual information arriving from the eye to the lexical processor due to preattentive visual processing. This delay is called the *eye-brain lag* (EBL). The second is a delay from when a saccade decision is made to when the eye starts moving (SPT: *saccade planning time*). One of the challenges for eye movement control models is the fact that typical estimates of saccade planning time take up a considerable portion of the total duration of a fixation (Becker & Jürgens, 1979; Rayner, Slowiaczek, Clifton & Bertera, 1983), leaving little time for lexical processing to influence fixation duration in a strictly serial stage model. The third delay is the time it takes the saccade to be executed. These delays are all modeled using gamma distributions. There are additional well-accepted ideas about saccade target and landing site distributions present in E-Z Reader that are not included in the dissertation model because they are of less relevance to the experimental task in the thesis.

The thesis model also follows E-Z Reader in assuming that only one word is processed at a time. This assumption is under debate: Engbert, Nuthmann, Richter & Kliegl (2005) make a strong case for a parallel model based on work in dynamic field theory by Erlhagen & Schöner (2002), as well as their own modeling results. However, the short words and wide word spacing of the LLDT, combined with visual acuity limitations, may limit even a parallel attention model to a near-serial behavior. The model also assumes that saccade plans are directly triggered by lexical processing, another area of active debate. The alternate choice is that saccades are triggered by an autonomous brainstem mechanism and only indirectly modulated by lexical processing (again as advocated by Engbert et al., 2005 following Findlay & Walker, 1999).

Generating predictions by optimization to payoff. The model generates predictions by treating the thresholds of the SPRT as a locus of *adaptive control* and finding the setting of these threshold parameters that maximizes mean payoff in the task given to the humans. This optimization is necessary for the model to make predictions, but is not a part of the theory in the sense that the theory makes no predictions about learning or task-level adaptation, only about its endpoint. The remainder of model parameters are fixed to a priori estimated values whenever possible, either taken from prior literature, or estimated from human participants outside of the model. A few remaining parameters are either fixed to reasonable values and justified in the text, or optimized to some aspect of human data. In this optimization, only near-optimal settings of the adaptive threshold parameters are considered. In this way, the model makes nearly zero-parameter predictions of data.

1.3 An overview of key results

The thesis provides two sets of key results. The first is on the adaptation of human eye movement control to task payoffs, and the second is an exploration of lexical frequency *spillover* effects in reading.

1.3.1 Adaptation to task payoff in reading

While adaptation to task differences in reading has been shown previously (e.g. McConkie et al., 1973; Rayner & Raney, 1996; Wotschack, 2009), the thesis provides the first precise quantitative payoff manipulation in reading. Participants are sensitive to payoff feedback even when the payoff has no intuitive interpretation (i.e. participants are not explicitly told to try to be fast or accurate and nothing about the payoff provides that intuition). In addition, there are models of active perception in reading (e.g. Bicknell & Levy, 2010b) that assume readers attempt to identify words, and models outside of reading that assume that perception is explicitly optimized to the task humans are attempting to perform rather than for more general information acquisition (e.g. Ballard & Hayhoe, 2009; Tatler, Hayhoe, Land & Ballard, 2011). But this is the first model to combine both ideas and do the latter kind of optimization (i.e. model doing the task) in the domain of reading.

1.3.2 Exploring lexical frequency spillover effects in reading

The second set of key result deals with so-called lexical frequency *spillover* effects in reading [1]. Spillover effects in general are effects on words that 'spill over' or appear downstream from the word that drives them. In the case of lexical frequency spillover effects, these are when words that follow high-frequency words are read more quickly

[1] While there are other types of spillover effects, the thesis will always take spillover to mean lexical frequency spillover.

(Inhoff & Rayner, 1986). The thesis provides a number of related results in an effort to understand these effects.

First, the thesis provides empirical evidence that explanations of spillover relying on parafoveal preview (Reichle, Pollatsek & Rayner, 2006) cannot account for spillover effects in the LLDT. This is because the effects persist with parafoveal preview gaze-contingently masked. Second, it shows how a dominant model of eye movement control in reading (E-Z Reader, Reichle et al., 2009) is partially consistent with spillover effects arising from post-perceptual delay as well as parafoveal preview. Third, it explores how an adaptive model that can allocate its processing resources to current perception or post-perceptual memory-based reprocessing might recover spillover frequency effects as a rational consequence of noise in lexical processing, and how this effect arises from a counterintuitive but fundamental fact about sequential samplers used in series. Fourth, it introduces parafoveal preview into this model and shows that using parafoveal preview can be adaptive in spite of it being noisier than foveal processing, and how parafoveal preview may amplify the memory-based spillover effect, yielding spillover effects in a broader portion of the free parameter space of the model. Finally, it attempts a direct fit of the model's free parameters to human data, with mixed results. Some discussion of implications is provided for future iterations of the model.

1.4 The structure of the dissertation

The chapters are structured as follows: Chapter 2 provides definitions and background on the terms and concepts used in the dissertation, including a broader discussion of terms-of-consensus in the field and justification for major assumptions. Chapter 3 discusses in detail the structure of the model that is used for the adaptation results and serves as the baseline for additional modeling later in the thesis.

Chapter 4 provides evidence that human participants adapt to quantitatively-expressed payoffs in the LLDT. Furthermore, it shows that they do so by adjusting individual fixation durations, a fine-grained component of their control strategy (in contrast to coarse strategic adjustments like increasing focus on the task). This chapter also details modeling results that recover adaptation along the same locus and provide a theory of adaptive eye movement control in the LLDT, and shows how the architecture bounds help constrain the model's predictions and in fact improve its predictive power. Finally, some discrepancies between model and humans are discussed.

Chapter 5, is the beginning of the thesis exploration of spillover effects. It shows that participants in the LLDT do show lexical frequency spillover effects even though LLDT was designed to maximize single fixations and minimize word-to-word dependence. The chapter reports on a gaze-contingent masking paradigm used to test whether the spillover effect is driven by parafoveal preview in spite of the wide word spacing in the LLDT, and shows that spillover persists without preview. Modeling results in E-Z Reader are reported showing that it (and models that share its key assumptions) can only recover the qualitative pattern of these results without parafoveal preview by assuming that processing of a word can delay the processing of the next word.

Chapter 6 develops an explanation of the spillover frequency effect in the framework of the model discussed in Chapter 3. Specifically, it develops a model that can adaptively allocate its processing to either current foveal perception, or a memory-based review of past input. Such a model recovers spillover effects when the memory-based review delays the start of processing on the next word. This delay is shown to be sensitive to frequency due to a property of sequenced thresholded samplers. Implications of this property are discussed for a broader theory of the integration of information across decisions. The chapter also compares this model with one that is able to give priority to perception over memory, and shows that the perception-priority policies in such a model do not perform as well over a range of noise settings Parafoveal preview is introduced into the model and shown to improve model performance in some portions of the speed-accuracy tradeoff curve.

Chapter 7 shows attempted quantitative fits of noise settings to human data, and shows that the model with memory sampling but no-preview cannot simultaneously recover spillover effects and single fixation durations close to human values. A number of modifications to the model are undertaken in light of this problem. The same fitting exercise is reported on the preview-capable model, with greater success.

A conclusion addresses some current challenges and future directions, including extending the model to a fuller sentence-reading setting and pursuing a richer policy space for the adaptive control of attention.

CHAPTER 2

Relevant empirical phenomena and current accounts

2.1 Rational approaches to understanding behavior

The work in this thesis explores behavior as an adaptation to both task and the intrinsic computational constraints on the organism. Theories developed from this perspective attempt to explain *why* particular phenomena appear, in addition to sometimes describing how they work. This approach is often contrasted with *mechanistic* approaches (e.g. the review by Chater & Oaksford, 1999) that attempt to describe the structure and sequence of operations that yield phenomena of interest without assuming or attempting to understand why they come about. This section provides an overview of *Rational Analysis* (Anderson, 1990), arguably the most influential of the rationalist approaches to cognition. It then discusses the role that the cognitive architecture or bounds of an organism might play in changing what rational behavior looks like, and an alternate formalism called *Computational Rationality* that explicitly makes agent bounds part of the theory. This is the approach taken in the thesis.

2.1.1 Rational Analysis

An influential formalization of the adaptationist approach to understanding behavior is Rational Analysis (Anderson, 1990, 1991). It attempts to understand the problem human agents are trying to solve by specifying the environment they operate in and their goals, and seeking to capture signature correspondences between modeled and measured behavior. Here is Anderson's procedure:

1. Specify precisely the goals of the cognitive system.

2. Develop a formal model of the environment to which the system is adapted.

3. Make minimal assumptions about computational limitations.
4. Derive the optimal behavior function given 1-3 above.
5. Examine the empirical literature to see whether the predictions of the behavioral function are confirmed.
6. Repeat, iteratively refining the theory defined by 1-3 under the optimality assumption.

Some notable early successes of this approach are work on the power-law relationship in forgetting (Anderson, 1990), on apparent irrationalities in the Wason selection task (Oaksford & Chater, 1994), on patterns in human categorization (Anderson & Matessa, 1990), and on saccade targeting in normal and pathological reading (Legge et al., 1997). Recent work is even more widespread. A few examples from the domain of psycholinguistics and reading bear mentioning: Bicknell & Levy (2010b) extended the model of Legge et al. (2002) by introducing uncertainty in word identity and provided better predictions of saccade targets. They also showed how a rational model might choose to regress as a result of falling confidence in past inputs (Bicknell & Levy, 2010a). In addition, Hale (2011) provided a general framework for rational analyses of the comprehender's parsing choices, and used this framework to explain garden path effects and the subject-object assymetry, phenomena of longstanding interest in psycholinguistics.

2.1.2 From rational analysis to computational rationality

The statement of Anderson's steps 2 and 3 places the explanatory burden on the environment first and mechanism second, which is part of the reason for the contrast drawn between this work and mechanistic approaches. However, architecture bounds carry part of the explanatory burden in rational analysis, even if they are not given credit for doing so. This is the case for a few reasons. Simon (1955) foreshadowed one by remarking that there is substantial room for debate on what aspects of the system are environmental or part of the organism, and foreshadowed another (Simon, 1991), by noting that the notion of minimal assumptions on computational limitations is technically imprecise.

A rational theory must assume costs to information processing, limitations on knowledge availability or acquisition, and other limits, and the minimal limitations needed depend on the empirical phenomenon of interest, which means that the theory partially relies on them for its explanation. Indeed, Anderson (1990) explicitly adds a 'capacity limit' as a cost in information processing, and adds a limitation on short term memory as well. In ideal-observer approaches to vision and perceptual decision-making closer to the thesis work (e.g. Green & Swets, 1966; Legge et al., 1997; Norris, 2006, among many), the computational limitation is the notion of noisy perceptual information. In the case of signal detection theory, one can appeal to this noise as being a part of the stimulus itself, but in a sequential model putting the noise inside the organism (i.e. as a part of the architecture) is more natural, by the following argument: if the stimulus on the page or screen is a single unchanging draw from a noisy distribution, then a rational agent might observe it once to make its inference. But if a single unchanging sample on the page is noisily transmitted through the information processing system, then a rational agent may well use multiple samples to update over its intrinsic noise.

This is not a criticism of the powerful method that rational analysis provides for understanding why behavior is the way it is, only of the emphasis it places on a the task and environment portion of the theory. The computational assumptions may truly be the minimally needed computational limitations needed for emergence of particular behavior. And even if the assumptions made were not the *minimally* required ones, the theories may still yield plausible predictive explanations of human behavior. Rather, thde argument is that rational models are already organism-bounded and that 'minimal assumptions about computational limitations' serves as an important but underexplored theoretical degree of freedom. Therefore, it is better to have a formal place in the theory for computational constraints (i.e. cognitive architecture) in addition to task constraints. The parsimony requirement for minimal assumptions then applies to both the computational limitations, and the agent's task environment.

Howes, Lewis, Vera and Singh (Howes, Vera & Lewis, 2006; Howes et al., 2009; Lewis et al., 2013) have championed such an approach explicitly making room as needed for both task and cognitive architectural constraints in the theory, and explicitly making a distinction between constraints imposed by the organism, constraints imposed by the task environment being performed, and constraints imposed by the broader environment the organism operates in outside of the task. Most recently termed Computational Rationality, this is the approach taken in the thesis.

2.1.3 Other approaches to rationality under bounds

Computational rationality is not the first rationalist approach explicitly making room for agent bounds in the theory. Two other theoretical frameworks explicitly tackle the agent bounds question: the framework of Bounded Rationality (Simon, 1955) and the cognitive model ACT-R (Anderson, 1990) ACT-R. They are briefly contrasted with computational rationality here.

Bounded Rationality. Simon (e.g. 1955; 1991), makes the same observation that leads Lewis et al. to their proposal of computational rationality: that behavior is jointly determined by the task constraints, and the internal constraints on the organism. But the focus of bounded rationality is in using rational behavior as an upper limit of performance, and in understanding which good-enough *satisficing* solutions humans can more easily find and therefore choose over the exactly optimal ones. The emphasis on optimization bounds is in contrast to the approaches above, which implicitly treat optimization as free (having been done over time by evolution or extensive practice).

In addition, in bounded rationality there is no precise notion of optimality under bounds, and reasonable suboptimal solutions are preferred and justified under the organism's adaptive constraints. The notion of bounds in the optimization process itself may be useful, but the computationally rational approach can explicitly take that into account, deriving the complex interaction between bounds on the agent imposed by the task, internal computational constraints, and the optimization process.

ACT-R. Rather than consider cognitive architectural and mechanistic bounds on rationality, ACT-R looks in the other direction and implements mechanisms and cognitive architectural features that are justified by rational analyses. Such a formulation can be taken as a strong modularity claim (Fodor, 1983), in the sense that the adaptation of each module of the architecture happens independently. It also implies a view in which the adaptation of interest is to the challenges the organism faces in its environment and not to the task in a particular experiment. This is problematic given evidence of within-task adaptation discussed previously. To reiterate, such evidence seems to exist across the task complexity spectrum: for example, adaptation to payoff in a simple psycophysics discrimination task (Green & Swets, 1966) and adaptation to question frequency and difficulty in a reading task (Wotschack, 2009).

2.2 Empirical effects in eye movements in reading

There is a substantial body of empirical work on what determines when and where the eyes move during reading – far beyond the scope of this chapter. Instead, this section will review the basic facts about eye movements, and then use lexical frequency effects as a case study for how a single property of words affects reading times. This latter portion serves double duty, as foveal and spillover frequency effects are of empirical and theoretical interest in the latter parts of the dissertation.

2.2.1 Eye movements in reading: basic definitions and phenomena

The subjective perspective on reading is that of slow, methodical scanning across the page. In reality, slow eye movements are rare in reading. Rather, the eye makes rapid movements from position to position called *saccades*. Saccades are fast, with peak velocities of as much as 500 degrees of visual angle per second (Rayner, 1998). Due to this speed, very little information is recovered during saccades, in a phenomenon called *saccadic suppression* (Matin, 1974). Instead, information is recovered when the eye is relatively still[1] in stops called *fixations* mostly lasting betwen about 250 and 350ms in length in normal reading.

The duration of fixations is determined by many factors correlated with the difficulty of individual words, ranging from typographic variability, to linguistic properties at higher levels of abstraction. Some core word-level effects are those of lexical frequency (e.g. Inhoff & Rayner, 1986), length (e.g. Just & Carpenter, 1980), lexical neighborhood size (Andrews, 1997), all of which lead to slower reading times. Many others are reviewed in detail by Rayner (1998). Other effects reflect processing above the word level, either as omnibus effects of predictability given context (e.g. Balota, Pollatsek & Rayner, 1985) or as narrower effects of sentence processing (e.g. garden path effects, Frazier & Rayner, 1982).

A major reason that the eye needs to move at all in reading is that the *acuity* function of the eye, i.e. its effective receptive resolution, is far higher in the center than in the periphery. The area of highest acuity, the *fovea*, covers approximately the central 2 degrees of visual angle, or about 6-8 characters of typically-sized text at typical reading distances (Rayner & Bertera, 1979). Within this area, letters and words are identified with very high accuracy (Bouma, 1973), falling off rapidly outside

[1] The eye experiences a slight tremor even during fixations, called *nystagmus*.

of it. This fact, and the fact that typical saccades are about 7-9 characters in length (Rayner, 1978), both suggest that one of the primary functions of eye movements in reading is to bring information into the fovea.

More precisely, the eyes typically target a position just left-of-center of words in left-to-right writing systems (e.g. Vitu, O'Regan & Mittau, 1990), and just right-of-center in right-to-left systems (Deutsch & Rayner, 1999). This *preferred viewing position* seems to yield faster word identification and a lower probability of *refixation* (another fixation on the same word, usually due to error in the saccade landing location). Given that the fovea is symmetric around the center of fixation, it is surprising that the eye tends to fixate off-center in words. Likewise surprising is the fact that the word identification span is asymmetric, extending farther in the direction of reading than away from it (McConkie & Rayner, 1975).

At least part of the explanation might rely on the fact that readers tend to preview words in their *parafovea* (the area between about 2 and 10 degrees of visual angle from center). There are two pieces of evidence for this effect: first, time viewing words in the parafovea shortens their eventual foveal vieweing; second, short words are occasionally skipped altogether rather than being fixated (e.g. Balota et al., 1985). Both interact with the predictability of the word (i.e. highly predictable words are likelier to be skipped, and receive a greater preview benefit). Finally, the eyes also occasionally return to previously read text, in movements called *regressions*. They are considered different from refixations because they are too distant to be due to saccadic motor error alone, and have recently been compellingly explained as a consequence of dropping confidence about past inputs (Bicknell & Levy, 2010a).

2.2.2 Basic constraints on the oculomotor system in reading

The section above already alluded to the most substantial constraint on the oculomotor system: that the acuity of the eye is limited, and falls off rapidly outside the center of vision. A second constraint is that eye movement planning takes time. Becker & Jürgens (1979) used the double-step paradigm (discussed in greater depth later) to place a lower limit of about 100ms on how long it takes to plan the motor saccadic movement when the intended target is known, though the actual delays between sequential fixations seems to be longer (Rayner et al., 1983). A third constraint is that there is a delay from when the eye lands on a word to when cortical activity begins to be correlated with perceptual information from that word. The lower bound on these delays comes from event-related potential (ERP) studies of the visual system, with the earliest deflection at 50-75ms (Clark, Fan & Hillyard, 1994; Mouchetant-Rostaing, Giard, Bentin, Aguera & Pernier, 2000), but higher values have been found depending on the type of visual stimulus (e.g. Tarkiainen, Helenius, Hansen, Cornelissen & Salmelin, 1999; VanRullen & Thorpe, 2001). These two mandatory delays add up to at least 150ms. Considering that fixations only last about 250-350ms, theories of eye movement control typically assume that some sort of processing continues during saccade planning, if not during the preattentive visual delay as well.

2.2.3 Frequency phenomena in reading

Lexical frequency phenomena are a useful case study for the relationship between word properties and reading times. Their robustness has permitted a detailed investigation into their nature in different situations and the extent of their influence. Consider a taxonomy of frequency phenomena, based on the word driving the effect and the word that the effect is seen on. In this taxonomy there are foveal frequency effects (when the foveal word frequency affects foveal reading times), preview effects (when parafoveal frequency affects the eventual reading times on the same word), parafoveal-on-foveal effects (PoF, when parafoveal frequency affects reading times on the foveal word), and spillover effects (when foveal frequency affects the eventual reading time on the parafoveal word).

While the physical parafovea extends up to 10 degrees of visual angle to the right and left of fixation (or about 30-40 characters at typical reading distances), the useful parafovea in reading seems to largely be limited to one word forward. While n+2 preview effects have been documented, they seem to appear in only very particular situations, for example when words n+1 and n+2 are short, high-frequency words (Kliegl, Risse & Laubrock, 2007). Likewise n+2 PoF effects have been recently shown (Bicknell & Levy, 2014) but they seem to only be sensitive to visual features of word n+2 such as length. With this in mind, the remainder of the thesis will assume that the useful parafovea extends to the next word only.

Foveal frequency effects. Words of higher frequency are recognized more quickly in visual presentation (Howes & Solomon, 1951) than words of lower frequency. The relationship appears to logarithmic, with the natural logarithm of lexical frequency counts predicting a variety of reading measures (Inhoff & Rayner, 1986; Rayner & Duffy, 1986; Henderson & Ferreira, 1990; Just & Carpenter, 1980, and many others). This is true independent of both word length, and predictability from context as estimated from the cloze task (a fill-in-the-blank task).

Parafoveal preview frequency effects. Parafoveal preview effects occur when parafoveal information obtained from a word shortens that word's later foveal viewing time (e.g. Balota et al., 1985; McConkie & Rayner, 1975; Rayner, 1975). The way these effects are investigated is by manipulating the quality of the preview of a word and testing the effect of the quality of preview on reading time. For example, the preview could be the true word, a phonologically or orthographically related word, an unrelated word, or a nonword. Recent work has shown that the frequency of previewed words also has some effect on the quality of the preview: high-frequency words only receive preview benefit from identical previews, whereas previewed orthographic neighbors may provide a benefit for lower frequency words, so to this extent parafoveal preview is sensitive to frequency as well (Williams, Perea, Pollatsek & Rayner, 2006).

Parafoveal frequency effects on foveal words. The effect of parafoveal words on foveal processing is investigated in experiments which manipulate the orthographic material available in the fovea. Readers slow down in response to a variety of orthographic irregularities in the parafovea, such as complete masking (Rayner, 1975; Blanchard, Pollatsek & Rayner, 1989), all-capitals strings (Inhoff, Starr & Shindler, 2000) and orthographically irregular nonwords (Starr & Inhoff, 2004). The extent to which lexical information in the parafovea such as frequency affects foveal processing is far less well-established: while some evidence for this has been found in corpus studies (Kennedy & Pynte, 2005) and non-reading tasks (Kennedy, Pynte & Ducrot, 2002), these effects do not replicate in more ordinary reading tasks unless the foveal word is very short (Angele & Rayner, 2011; Kliegl et al., 2007) and may be partially due to mis-targeted saccades (Drieghe, Rayner & Pollatsek, 2008).

Foveal frequency effects on parafoveal words. The frequency of a word affects the reading time of the following word (Rayner & Duffy, 1986; Inhoff & Rayner, 1986). This *spillover* effect is most often investigated in the context of the parafoveal preview phenomenon discussed above. The amount of parafoveal preview benefit is affected by foveal frequency (Inhoff & Rayner, 1986; Henderson & Ferreira, 1990; White, Rayner & Liversedge, 2005). However, the effect of word N foveal frequency on reading times on word N+1 occasionally remains even in the absence of parafoveal preview (Kennison & Clifton, 1995; Schroyens, Vitu, Brysbaert & D'Ydewalle, 1999). This effect is of substantial interest in the latter part of the dissertation, and is discussed in greater detail in §5.1.1.

2.3 Sequential sampling as a model of moment-by-moment decisionmaking

Thresholded random walk models of decisionmaking (e.g. Ratcliff, 1978; Edwards, 1965) have a long and successful history in the cognitive science literature: they recover reaction time and accuracy distributions consistent with human behavioral data, as well as firing patterns of neural populations (e.g. Cook & Maunsell, 2002; Gold & Shadlen, 2007). They have also been applied to lexical processing, both using the random walk Drift Diffusion model (DDM, Ratcliff, Gomez & McKoon, 2004), and using a sequential statistical test called the sequential probability ratio test (SPRT, Norris, 2006, 2009).

The DDM assumes that there is some noisy decision signal available outside of the decision mechanism. In the case of lexical decision this signal is 'wordness', and it is assumed to be sensitive to things like frequency, predictability, length etc. The DDM accumulates this signal to a threshold until a decision is made. Since the signal is noisy, the practical realization of the process is as a random walk in this 'wordness' space.

Norris (2009) makes a strong case for the SPRT or its multihypothesis variant (MSPRT, Baum & Veeravalli, 1994) as a better model of decisionmaking, at least for the case of lexical decision. The claim is that since the SPRT operates explicitly in probability space, there is no need for the ill-defined notion of wordness. Individual samples of the wordness signal in the DDM become likelihoods of hypotheses (i.e. words) in the SPRT, making the random walk a walk in log-likelihood space of the words being perceptually identified. The decision thresholds can now also be interpreted as desired error probabilities, and the starting points of the walk as prior probabilities. If the priors are based on lexical frequency counts, the SPRT recovers the log-frequency effect on reading times. With mild assumptions about the likelihood function, it can also recover neighborhood density effects, and some differences between naming and lexical decision tasks.

2.4 Models of eye movement control in reading

This section will review two current dominant models of eye movement control in reading: E-Z Reader (Reichle et al., 2009) and SWIFT (Richter, Engbert & Kliegl, 2006). These two models are chosen because they claim broad coverage, have seen recent theoretical and empirical treatment, and fall on opposite points of the theoretical debates on attention and directness of oculomotor control in reading. While many other models exist, they either aim for narrower coverage (e.g. regressive eye movements, Bicknell & Levy, 2010a), or are similar variants of those two models: the autonomous-saccade Push-Pull model (Yang & McConkie, 2001) is a less formalized variant of many of the ideas in SWIFT; Glenmore (Reilly & Radach, 2006) is a connectionist model that shares many of the SWIFT assumptions but adds ideas about lexical processing from the connectionist literature (McClelland & Rumelhart, 1981); EMMA (Salvucci, 2001) is a production system model sharing many assumptions of E-Z Reader; and so on. Reichle, Rayner & Pollatsek (2003) provides a more detailed review of some of these models and others. Some narrower-coverage models in a rational perspective will be given additional consideration below.

E-Z Reader E-Z Reader (Reichle et al., 2009) is a serial-attention, direct oculomotor control model of reading that makes minimal assumptions about the mechanistic content of the lexical processing system. It is serial-attention in the sense that only one word can be lexically processed at a given time. It is direct-oculomotor-control in the sense that oculomotor control decisions are directly triggered by the lexical processing system (though lexical processing and oculomotor programming can proceed in parallel). It makes minimal assumptions about the mechanistic content of the lexical processing system in the sense that it does not provide a theory of lexical processing or access. Rather, it includes a lexical processing delay that varies with the logarithm of word frequency, predictability, and length, assuming that any theory of lexical processing will have to at minimum account for those facts. E-Z Reader also makes a number of assumptions about saccade targetting error that let it fit saccade landing site distributions, and assumptions about parafoveal preview that let it fit spillover effects. The total set of fit phenomena in E-Z Reader includes the frequency, predictability, length and spillover effects mentioned above, as well as saccade landing site distributions, refixations, word skipping rates, and both intra- and inter-word regressions.

The E-Z Reader modeling paradigm is entirely descriptive: only training fits to data are provided, and the model is considered to account for phenomena if it can fit the training data with low error. Moreover, fits are done on aggregate and highly colinear data, yielding unrealistically low error (Feng, 2003). E-Z Reader uses over a dozen free parameters to fit 30 means that are correlated with coefficients as high as 0.9. This is not to say that E-Z Reader is necessarily an incorrect model or that its assumptions and claims are unrealistic, but its apparently low-error fits to data may not impl yits ability to strongly constrain the space of possible explanations.

SWIFT In contrast to E-Z Reader, SWIFT (Richter et al., 2006) is a parallel (gradient) attention, indirect oculomotor control model. It is a gradient attention model in the sense that multiple words are processed in parallel, constrained by an activation function. Saccade targets are selected based on ongoing competition between activations of different words, so in that sense the saccade target decision is directly controlled. But the saccade timing is randomly determined by a low-level saccade generation mechanism, so in this sense it is an indirect oculomotor control model. To recover frequency and spillover effects (or in fact any effect of word properties on reading times), SWIFT allows lexical processing to slow down the autonomous saccade generator. This slowdown is itself delayed in what Richter et al. call time-delayed foveal inhibition, allowing SWIFT to recover both word N and word N+1 (spillover) effects using a single mechanism.

The activation function over time forms a random walk in the eyes-are-likely-to-target space, rising to a maximum activation computed from frequency, predictability and length in a similar equation to the one used by E-Z Reader, and falling back to zero. SWIFT has similar empirical coverage to E-Z Reader on word-level effects (frequency, predictability etc) and saccade targetting, and also makes predictions about parafoveal-on-foveal effects (Kennedy & Pynte, 2005). As with E-Z Reader, its modeling paradigm is descriptive, with 13 free parameters. It fits them to minimize total error on 8 different measures on 850 words. This is an improvement over E-Z Reader's fit to means, but not a substantial one since individual word-level measures should be highly correlated across subjects and words.

Rational analyses of eye movements The work in the thesis draws some inspiration from previous work that also considers eye movements as emerging from adaptive constraints on the organism itself. For example, the Mr. Chips model (Legge et al., 1997, 2002) modeled the reading process from an ideal observer perspective, i.e. as a process of rational adaptation to noise. Legge et al. showed how saccade targeting decisions (including viewing position effects, skips, and regressions) can be understood as a rational adaptation to acuity, and moreover how increasing acuity limitations (such as in low-vision participants) change the near-optimal reading pattern. Bicknell & Levy (2010b) extend Mr. Chips to consider uncertainty over incoming word information, and show that this improves the predictive power of the model. They also show how regressions in particular may be a rational adaptation to increasing uncertainty over past inputs (Bicknell & Levy, 2010a).

Reichle and colleagues (Reichle et al., 2006; Liu & Reichle, 2010; Reichle, Liu & Laurent, 2011) likewise show how simple bounds on acuity and attention might yield adaptive behavior like fixating close to the centers of words, anticipating the saccade planning decision before lexical processing is complete, and choosing to preview short words to eliminate costly saccades.

2.5 Summary

This chapter reviewed a few different approaches to understanding behavior as adaptive, and motivated the use of computational rationality rather than one of the other approaches. In addition, it explained the basic terminology and phenomena in eye movement control, and used lexical frequency effects to illustrate the kinds of relationships between cognitive processing and eye movement control that models of eye movement control seek to understand. It provided an overview of recent sequential sampling approaches to lexical recognition that will form the core of the lexical processing component in the theory advanced in the thesis. Finally, it reviewed other rational approaches to eye movement control in reading, as well as dominant non-rational approaches, one of which forms the basis for the oculomotor architecture used in the theory advanced in the thesis.

CHAPTER 3

A Theoretical Model and Empirical Paradigm for Understanding Adaptive Eye Movement Control

This chapter introduces the List Lexical Decision Task (LLDT), a simple empirical paradigm for understanding adaptive eye movement control, as well as the theoretical model of this task, from which inference to reading will be drawn. To introduce it and compare it to the other models reviewed in the previous chapter, consider the following view on theorizing in cognitive science: there is some cognitive skill of interest, supported by aspects of the human cognitive architecture – for example, reading; there is a task or set of tasks that are used to empirically investigate this skill of interest, that hopefully tap the same underlying architecture – for example, a self-paced moving window paradigm; and there is a theory of the skill and architecture built based on data from these tasks. Under this view, there are two inferential steps: one from theory to task data, and one from task data to the actual ecologically interesting task that humans do outside of the lab and the components of cognitive architecture that support it.

With both steps come abstraction choices. E-Z Reader makes the choice to abstract away from sentence-level complexity in its theory, but not in the task measured in the lab. Therefore, there is a challenging inferential step from theory to task, expressed as the criticism that E-Z Reader is a theory of sequential word recognition applied to sentence-reading, or as the criticism that the architecture of E-Z Reader cannot support the task of naturalistic reading. But the benefit is that of a simple inferential step from the laboratory task E-Z Reader uses to the skill of reading in the world outside the lab (i.e. the cognitive architecture components recruited in the lab are more likely those recruited in the naturalistic environment).

The thesis makes a different set of abstraction choices. The advanced theory is a theory of wordlist-reading, but the empirical task used is also a task of wordlist-reading. Therefore, the inferential step from theory to task is much smaller than in the case of E-Z Reader, i.e. the model is doing the same task as the humans, and the architecture constructed is sufficient to support the wordlist-reading task. On the other hand, the inferential step from the list-reading task to the ecologically interesting skill of reading is more challenging: is the list-reading task like sentence-reading? This is similar to the challenge that can be leveled at models like E-Z Reader, except that it is made on the level of empirical data rather than theoretical argument. By gathering data from the simpler task that the theory is modeling, it is possible to address empirically the question of the similarity between the experimental task, tasks like sentence-reading, and reading outside of the lab.

3.1 The list lexical decision task

The dissertation uses the *List Lexical Decision Task* (LLDT). A simple extension of a paradigm first introduced by Meyer et al. (1974), it requires participants to make a single lexical decision on a list of character strings. On each trial of the LLDT, participants fixate on a target on the left side of the screen, and are then presented with a list of character strings, with the first string where the fixation target was. The list of strings is either a set of six English words, or five words with a single pronouncable nonword. Strings are never repeated within a list. The participants are required to indicate by buttonpress, for each trial, which of the two cases it is (all-words or one-nonword). Participants are given a precisely expressed quantitative payoff reflecting some intended speed-accuracy tradeoff (e.g. losing some number of points for inaccurate responses but gaining some number of points as a function of response time independent of response), and are given point feedback after every trial. This is in contrast to more typical instructions to "respond as quickly and accurately as they can," and the precision means that the model can be provided with the same feedback.

Strings are kept short (four characters each) to minimize refixations and maximize single fixations. This will help simplify the modeling effort, as a successful model might only need to perform single fixations. The consistent length should also yield clean estimates of frequency effects unconfounded by length and only minimally confounded by predictability: because of the at-most-single-nonword and non-repetition constraint, previous strings are only slightly predictive of upcoming strings. Strings are also widely spaced (six character spaces) to minimize parafoveal preview and any skips it may drive.

Despite its simplicity, the LLDT has several desirable features:

- The precisely expressed payoff makes it amenable to subtle quantitative payoff manipulations.
- The task requires linguistic processing at the word level (to make the lexical decision), and possibly processing above the word level as well (via the subtle predictability effect noted above).
- Because participants always start viewing the strings fixated on the first one, and because strings are widely spaced, it is expected to yield a left-to-right reading pattern similar to that in normal reading.
- Because information is acquired at each fixation but the task response is global to the trial, the LLDT requires the control of visual attention and the integration of information across multiple saccades to yield a response, and thus poses a joint optimization problem over both sets of decisions.

Later chapters will revisit some of these features and provide some evidence that they are indeed present in the task as performed by human participants.

3.2 The core theory

The theoretical work in the thesis proceeds as follows: it specifies a reading agent that performs the full LLDT – choosing eye movement actions (with some simplifications), as well as the eventual motor response. At each point in time, the agent produces a Bayesian solution to the inference problem posed by the LLDT: it takes in noisy observations of strings in different positions, and updates its beliefs about word and trial identity over time. Summary statistics over this belief distribution form the agent's *state*: the set of observations on which it conditions its *actions*. This inference problem is wrapped in a parametric *machine architecture* that forms the agent's *bounds*: durations of saccade planning and execution, etc. The mapping of states to actions is the agent's *policy*. The agent's policy is specified as a set of threshold parameters defined over the summary statistics in its state.

To generate predictions, the theory fixes the parametric architecture whenever possible to a priori estimated values (taken from prior work or estimated outside of the model), and then finds the set of near-optimal policies with respect to the same payoff given to human participants, i.e. ones that maximize this payoff rather than a fit metric to human data. Some machine architecture parameters cannot be fixed from prior work and are treated as free parameters optimized to data. Values for those parameters are found by only considering near-optimal policies, and then selecting the free parameter settings for which the optimal policy is closest to human data.

Figure 3.1 shows a block sequence diagram of the model reading two words, `chip` and `dive`. A detailed description of each stage is provided below, but first consider the sequence of events in reading a single pair of words in the LLDT model. The model begins fixated on the word `chip`. Before it can begin processing the word and making its decision to saccade, it must first complete some early visual processing, marked as *EBL* on the figure for eye-brain lag. It then begins lexical processing in service of its saccade decision. This process is a sequential Bayesian sampling process, or equivalently a random walk in log probability space, that will be discussed below. Eventually the saccade decision is made (based on an adaptively-set threshold on a summary statistic over the random walk) and saccade planning begins, as the eye is preparing to move on to the next word, `dive`. While the oculomotor plan is being computed, the eye continues receiving information from the current word, annotated as 'free' sampling in the figure because unlike the saccade decision sampling, it is not under adaptive control.

Eventually the saccade planning is complete and the eye begins moving (annotated as *Sac. Exec.* on the figure, i.e. saccade execution). Some perceptual samples are still available for part of the saccade, so 'free' sampling continues for a short time – these are samples that arrived at the eye before the saccade began and are only now done with early visual processing. Eventually the saccade is completed and the eye has now landed on `dive`. Another eye-brain lag delay begins now, after which the saccade decision on `dive` will begin, and the whole process will repeat on the next word. Ongoing throughout and not notated on the figure is the trial-level decision process. The trial level decision is also based on an adaptively-set threshold, but defined over a different summary statistic of the state of the random walk.

The next section will explain and motivate the set of decision processes the model uses as it reads, and then provide more motivation for the other delay stages in the model.

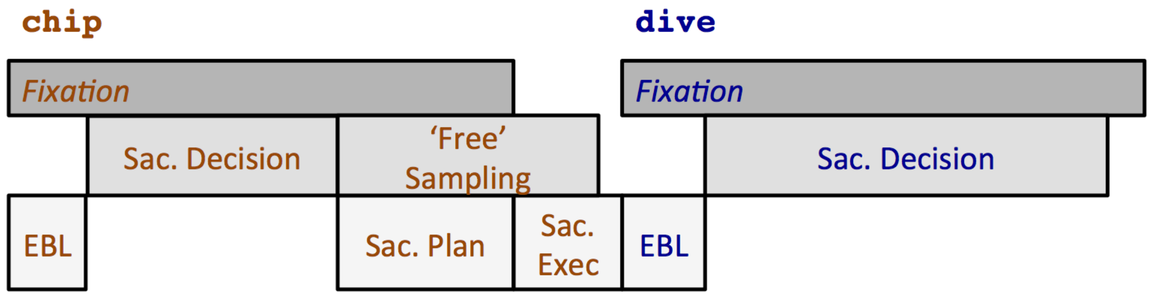

Figure 3.1: The sequence of stages as the model reads one word and moves onto the next. See §3.2 for a discussion.

3.2.1 Lexical processing perceptual identification

The modeling framework in the thesis views lexical processing as a process of perceptual identification: the reader proceeds from not knowing what the word is on the page (but perhaps having a guess) to knowing what it is, up to some level of certainty. A model of such a process should admit at least the following three sources of variability in recognition times:

1. Intrinsic properties of words (e.g. frequency, predictability, length).
2. External properties of the task (e.g. imposed by adaptive pressures).
3. Properties of the recognition process itself (e.g. perceptual noise).

Rise-to-threshold drift models such as the Drift Diffusion Model (DDM; Ratcliff, 1978) admit at least three sources of variability (the start point, the threshold, and the drift rate) and as such can capture these sources of variability. However, and as noted by Norris (2009), such models suffer in interpretability: their random walk is in abstract decision units in an abstract decision space (e.g. 'wordness') rather than in units and space directly corresponding to task. A related alternative is the Sequential Probability Ratio Test (SPRT, Wald, 1945).

The SPRT can be thought of as a dynamic version of Signal Detection Theory (Green & Swets, 1966). Green & Swets were interested in the basic building blocks of visual psychophysics: understanding how much light needs to reach the retina for perceptual recognition to be made. What they found is that there was no fixed amount needed and that the decision participants made depended on what the distributions of signal and noise looked like, and what value was assigned to hits (identifying signal trials) and correct rejections (identifying noise trials). They modeled this behavior using a threshold test: given a signal and noise distribution one can place a threshold somewhere over their combined support. If a given observation is on the signal side a 'yes' response is made; otherwise a 'no' response. If the signal and noise distributions overlap, the model generates false alarms and misses for highly unusual signal or noise trials. The threshold precisely specifies a tradeoff between misses and false alarms that is necessary for any solution to the decision problem with overlapping signal and noise distributions.

The SPRT is analogous to the SDT model, with the addition of the option to wait and receive another sample. At every time step, the log likelihood of the sample observed is computed given both hypotheses (signal and noise). The sum of the log likelihood ratios of samples seen so far is compared to two thresholds. If it is above the upper threshold, one hypothesis is selected. If it is below the lower threshold, the other hypothesis is selected. If it is between the thresholds, another sample is drawn. It is the optimal sequential test between two hypotheses, as long as one of the hypotheses corresponds to the true sample generation scheme. It is sequential in the sense that samples are observed iteratively, optimal in the sense that it provides the lowest error rate for any desired sample size, or the smallest sample size for any desired error rate (Wald & Wolfowitz, 1948), and the condition on the sample generation scheme means that the test may do arbitrarily poorly if the samples are drawn from some distribution that is not one of the hypotheses (Anderson, 1960).

In such a formulation, the SPRT is equivalent to a random walk in summed log probability space. It also has an equivalent Bayesian formulation, where a Bayes update is performed at each step on the incoming evidence samples. In this formulation it is easy to map the SPRT to the sources of variability above: the initial prior (or first log likelihood) reflects the distribution of words in the environment, the thresholds reflect desired sample counts and error probabilities dictated by task constraints, and the likelihood function reflects properties of the recognition process itself.

The multihypothesis variant of the SPRT, the MSPRT (Baum & Veeravalli, 1994; Draglia et al., 1999) only approximates the optimal multihypothesis sequential test. But unlike the optimal test (Tartakovsky, 1988) it has the same threshold form as the SPRT, and is optimal in the lower limit of noise (and equivalently upper limit of sample rate). It has been successfully applied to model lexical decision and single-word recognition data, recovering effects of frequency and neighborhood density, among others (Norris, 2006, 2009). The model variants in the dissertation all share Norris' choices of word representation and evidence likelihood function, though the Bayes updates and decisions are rendered more complicated because of the trial-level statistical structure in the LLDT.

The Bayes update used in the thesis has a correct generative model of the trial (i.e. probability of word and nonword trials, probability of generating words and nonwords in each position, the sample generation noise, etc) with one exception: the update computation does not include the constraint that words in the real LLDT are never repeated within a trial. At each unit of time, the model's belief representation contains the correct probability distribution over strings in each position in light of the samples received so far under this simplification.

Specifically, the value computed is $P(S^k = s_i|e^k, T = t)$, that is, the probability that the string in position K (S^k) is the ith string in the lexicon (s_i) conditioned on an incoming evidence sample from position k (e^k) and the current belief about the trial identity ($T = t$, with T a multinomial distribution consisting of outcomes "all words" and "nonword in position k" for all k). Then, an update of the posterior distribution over trial types is also computed, specifically $P(T = t|e^k)$. Mathematical detail of the update is provided in the appendix.

An adaptively set threshold. The model makes its decisions when particular summary statistics over the belief distribution cross an adaptively set threshold. This summary statistic for the trial-level motor decision is the probability that the trial is an all-words trial (and its complement, the probability that the trial contains a nonword). For the position-level eye movement decision two summary statistics are used in different models. In the first model in the thesis, it is the probability of a word (and its complement, probability of a nonword) in a given position. In the remaining model it is the probability of the most likely string in the current position. Reasoning for these choices is discussed in the next section.

The threshold is is adaptively set in the sense that it is set to maximize a payoff function – the same one given to the humans in the task. The model's predictions for behavior are drawn only from those parameter settings that are near-optimal in the task, in contrast to models that fit all free parameters to data. There is no learning claim here, however: while there is some work on how to set decision thresholds rapidly in an online setting (Simen, Cohen & Holmes, 2006), the model in the thesis is agnostic on this front and finds near optimal thresholds by brute-force offline search. This also permits the exploration of contrasts between optimal and suboptimal policies under the model's assumptions, i.e. explanations that it includes or excludes under the bounded optimality assumption. The particular choice of summary statistics is discussed later.

Sample Rate and Sample Noise. Besides the threshold, the sequential sampler has parameters reflecting the sample rate, the noise terms in the sample generation function, and the functional terms of the likelihood computations. Under the assumption that the model has the correct noise model of the word, the likelihood function is taken to be the inverse of the sample generation function, with the same parameters (i.e. if the sample generation function is a random draw from the gaussian distribution, the likelihood function is the gaussian probability density function with the same parameters). Higher noise means the belief probability will move less at each time step, and lower will make it move more.

The sample rate is a what allows the model to make timing predictions in milliseconds rather than samples. If processing a single sample takes less time, the belief will move more in the same amount of time. Sample noise and sample rate in the model can therefore be thought of as jointly determining a scaling factor between samples and milliseconds: higher noise has the same effect as slower samples, that of making the belief probability move less per millisecond. Making the sampling duration interpretable in millisecond-space is also important for understanding how sampling durations interact with millisecond-level estimates of architectural delays.

The theoretical interpretation of these parameters depends on whether the theory treats the SPRT as a convenient discrete approximation of a continuous-time decision process, or as a discrete model of a discrete process. Treating the discreteness of the SPRT as a theoretical claim rather than a modeling convenience aligns with evidence for discrete oscillations in human perceptual processing (visual, Dehaene, 1993; and auditory, Giraud & Poeppel, 2012) as well as short-term memory (Elliott & Müller, 2000). Likewise, recent work implementing the MPSRT using Poisson spiking neurons (Zhang & Bogacz, 2010), together with a way of estimating such Poisson spike train parameters from neural sequence data (Lehky, 2010) may pave the way for a priori noise estimates in the future, as well.

The thesis makes the simpler choice of treating these parameters as free. Since they exhibit perfect parameter trading, sample rate is fixed and noise used as the fit parameter. The fixed value is 10ms, which is coarse enough to keep simulation tractable while not being so coarse it wipes out effects of interest.

3.2.2 Decision dynamics

Recall that the model's belief representation contains the probability distribution over strings in each position. It must condition decisions on this probability distribution. It does so by computing summary statistics corresponding to hypotheses of interest and comparing those probabilities to adaptively-set decision thresholds, as in the SPRT or MSPRT. Under the rationality assumption, these summary statistics should be selected such that the decision over them optimizes payoff. But optimality explanations implicitly admit optimization over payoffs in two different kinds of environments. One is the actual task environment – in this case, the LLDT – and the other is the ecological environment that the humans are adapted to – in this case, ordinary reading. Humans could be adapted to the latter either if the adaptation cost is high, and therefore it is not worthwhile to adapt to the experimental task, or if the adaptation is on the level of populations rather than organisms. Lewis et al. (2013) call the latter case *ecological optimality* to emphasize the role of the organism's out-of-the-laboratory environment in the explanation. Explanations that do not appeal to ecological bounds are stronger because they require fewer assumptions, so they should be preferred. The decision dynamics in the model are discussed in context of these two options.

3.2.2.1 Motor response decision and action.

In the case of motor responses, there are two decisions: respond-yes and respond-no. The simple optimality assumption is that participants map the decisions onto the two possibilities indicated in the task instructions; all-words, and not-all-words. That is, assuming w.l.o.g. that $T = t_1$ reflects the outcome that the trial is an all-words trial ($T = w$, and that $T = t_{k+1}$ reflects the outcome that the trial has a nonword in the kth position ($T = n^k$), the model computes $P(T = w|e_k) = P(T = t_1|e_k)$ and $P(T = t_n|e^k) = \sum_k P(T = n^k|e^k) = \sum_k P(T = t_{k+1}|e^k)$. The maximum of those two values is the decision variable, such that when it is equal or greater than the threshold, the motor response decision is made.

The two-choice case might appear to map the decision onto Wald's original SPRT, inheriting its optimality. However, the analytical form of the expected sample size in Wald's work relies on samples being i.i.d. which is not the case in the LLDT because they arrive from different word positions, so it may be the case that the optimality result does not apply in the LLDT hypothesis set.

The ecological optimality alternative is adding the assumption that participants are relying on preexisting, highly-practiced reading strategies – for example, responding when either all words are identified or when a nonword conclusively fails to be identified as a known word. This is the back-off choice if the non-ecological option is not successful. As it will turn out, the simpler option above seems to work well and the additional assumption is not added.

3.2.2.2 Eye movement decision and action.

In the case of eye movements, there is no mapping from probabilities to eye movement decisions provided in the task instructions. The full space of actions might include targeting a saccade on any word, or even on any character (as is in the case in related models, Legge et al., 1997, 2002; Bicknell & Levy, 2010b). But the LLDT was designed to maximize sequential left-to-right eye movements in part to simplify the eye movement action space, so by assumption the action the model will take will always be to saccade to the next word. Such a model will fail to recover word-skipping and regressions, and will therefore likely behave suboptimally (Bicknell & Levy 2010a showed that regressive policies outperform non-regressive ones). On the other hand, it will simplify the rest of the modeling effort.

This leaves the model with one eye movement action: *plan-saccade-forward*. The choice that remains is the decision that the action is conditioned on. A simple option from the perspective of task optimality is the same choice as in the buttonpress decision: to condition the eye movement decision on the probability that the string being fixated is a word (or nonword). That is, assuming without loss of generality that the words are $s_i, i = 1..n$ and nonwords are $s_i, i = n..m$ and m is the full size of the lexicon, the model computes $\sum_{i=1}^{n} P(S^k = s_i|e^k, T = t)$ and $\sum_{i=n}^{m} P(S^k = s_i|e^k, T = t)$. At each time step the higher of the two is compared to the threshold. If the threshold is crossed, eye movement planning begins.

This is the choice used in successfully modeling the lexical decision task by Norris (2006), and is also the choice taken in the modeling of payoff adaptation in Chapter 4.

An alternative under the ecological optimality argument is that participants plan a saccade when the current word is sufficiently identified, rather than just sufficiently discriminated between word and nonword, again with the caveat that conditioning on probabilities over past words, the nature of the trial, or something else might perform better. That is, make the decision variable becomes $\max_i P(S^k = s_i|e^k, T = t)$. When the posterior probability of some specific string is high enough, eye movement planning begins. This back-off choice will be made in the modeling of spillover, and motivated in Chapter 6.

3.2.3 Oculomotor architecture: delays and distributions

The oculomotor architecture in the model takes the form of stochastic delays. These are an abstract way of modeling processes that are not lexical processing, the primary adaptive component of the model. In particular, the computational details of motor planning and execution, saccade planning and execution, and early visual processing are all abstracted away and modeled as delays with fixed gamma distributions. The specific values chosen and the justification for the distributional form are detailed below.

3.2.3.1 Gamma-distributed delays

The gamma distribution is continuous and zero-bounded, making it convenient to use for delays that cannot take on negative values. It has been used to model delays on saccade planning and related values with some success in prior work (Reichle et al., 2009). In addition, one of its interpretations is as the distribution of waiting times on poissson-distributed events (such as neural spikes), making it a natural choice for simulating durations of neural processes of interest that are not given a more detailed treatment.

For ease of interpretation, one can recast gamma distribution parameters in terms of desired mean and variance. In particular, the scale α and rate β of the gamma distribution are computed from means and variances as follows:

$$\alpha = \frac{\mu^2}{\sigma^2} \qquad\qquad \beta = \frac{\mu}{\sigma^2} \tag{3.1}$$

For all the stochastic delays, σ is set to $0.3 \times \mu$, which is close to values used in prior work (Engbert et al., 2005; Reichle et al., 2006)

3.2.3.2 Manual motor delays.

A delay on planning and executing a buttonpress response is needed to be able to model the full LLDT. For this delay, the dissertation draws upon the EPIC cognitive architecture (Kieras, Wood & Meyer, 1997), which provides a model of motor planning.

In EPIC, motor planning and execution take 50ms per motor feature activated. Kieras et al. (1997) provide an example of a simple 'peck' keypress that involves five features: the motor style used, the hand, the finger, the direction of the motion, and the extent of the motion. A mean of 250ms for motor planning and execution is therefore used.

3.2.3.3 Eye-brain lag: early visual processing

The eye-brain lag reflects the processing that happens from when the eye lands on a word to when linguistic processing begins. Typical estimates are taken from the electrophysiology literature, by looking for the first appearance of stimulus-correlated or task-related activity in the electrophysiological record. The time of the first stimulus-related activity and first task-related activity, however, depend on stimulus and task. Early visual evoked potentials in response to checker stimuli peak as early as 65ms and start earlier (C1; Clark et al., 1994). Stimulus-related potentials in more complex tasks tend to come later: 40-90ms in face discrimination (Mouchetant-Rostaing et al., 2000), 75ms in visual category discrimination (VanRullen & Thorpe, 2001), and as late as 125ms in letter discrimination Tarkiainen et al. (1999), in MEG, though these are peak measurements and the first deflection is earlier. Task-related activity comes even later, at about 150ms for category discrimination (VanRullen & Thorpe, 2001) and 143ms for letter discrimination (Tarkiainen et al., 1999). There is one outlier in these findings: Seeck, Michel, Mainwaring, Cosgrove, Blume, Ives, Landis & Schomer (1997), showed a task-related deflection as early as 50ms, in both ERP and intracranial EEG in a face recognition task.

If eye-brain lag is a shorthand way of modeling the processing between the retina and the lexical recognition process, then the duration of eye-brain lag might reflect an implicit claim about the kind of features available to the linguistic process, and the location where it occurs. One might address this theoretically by making a strong theoretical commitment to a particular lexical representation, or empirically by investigating first deflections of the kind discussed above in a task closer to the LLDT.

A simpler option, taken in the dissertation, is to use the value of 50ms motivated by prior modeling work (Reichle et al., 2009), which is also consistent with some of the shorter estimates discussed above.

3.2.3.4 Saccade planning time

Saccade planning time in the model is the duration from when the decision to make a saccade is made, to when the oculomotor action can actually begin execution. Previous estimates used in eye movement control modeling have ranged widely, from around 125ms (Engbert et al., 2005; Reichle et al., 2009) to 240ms (Reichle et al., 2003). Estimating this value is challenging because saccades might be partially pre-planned if the decisions of where and when to move are not made at the same time, and because the planning duration might be proportional to the distance the eye travel, or some properties of the content of the visual field.

Rayner et al. (1983) tackled this question by trying to create a saccade task with a visual field and goals somewhat similar to reading: targets arranged in a horizontal sequence. Their specific concern was with whether there is a minimal refractory period between saccades, and whether advance targeting or timing information (i.e. a predictable time-to-move or predictable location) can shorten saccades by pre-planning them. But the results can also be interpreted to understand the general question of how long saccade planning takes.

Most similar to reading is the second experiment in (Rayner et al., 1983): participants see a series of fixation crosses, but only move from one to the next in response to a cue (presented foveally or parafoveally). By delaying the timing of the cue, it is possible to estimate the time to plan a saccade when the saccade target is known, but the time to start planning it is not known in advance. This is similar to ordinary reading where most saccades are targeted in a consistent left-of-center position in words, but the timing is variable.

Rayner et al. found that with cues delayed 50ms or less, saccade planning took about 220-240ms (depending on whether the cue was foveal or parafoveal), dropping closer to 200ms with cues delayed more than 50ms and up to the limit tested of 200ms. There is one problem with this estimate, however: Rayner and colleagues removed all fixations 100ms or less after the cue, treating them as anticipatory. Treating such ultra-fast saccades as anticipatory is largely consistent with what is known about saccadic reaction times, but will bias the data if there are more anticipatory saccades at longer delays.

A different way to estimate saccade planning duration is to use the double-step paradigm (Becker & Jürgens, 1979). In this paradigm, participants are asked to saccade from a fixation cross to a target – but a second target is occasionally added after the first, in which case participants attempt to saccade to both in sequence. The second saccade is usually faster, suggesting that it is partially programmed in parallel with the first after the second target appears. Becker & Jürgens fit a linear model to their saccade latencies as a function of second target delay and use its slope to estimate a theoretical minimum amount of time needed to plan a saccade when its target and timing are both known, in parallel with other saccade planning processes. More recent modeling efforts (Reichle et al., 2006, 2009; Engbert et al., 2005) use the lower values estimated from this task.

This lower value (between 112ms and 148ms depending on where the second target appears) is the minimum amount of time saccade planning takes, even assuming that the reader can do as much parallel or anticipatory saccade planning as possible. The dissertation model uses a mean of 125ms, following the latest version of E-Z Reader (Reichle et al., 2009) and close to the mean value of 115ms used by SWIFT (Richter et al., 2006). Chapter 4 will briefly discuss what the consequences might have been of choosing the higher estimate from the cued saccade task rather than the lower one from the countermanding paradigm.

3.2.3.5 Saccade Execution time.

The duration that the saccade takes can be estimated from the eye movement task without using the model. The only challenge is determining the boundaries between fixations and saccades in the eye movement record, and standard methods for this are provided in the software shipped with eyetracking systems. A mean value of 40ms was estimated from participants in the LLDT (full experiment reported in Chapter 4).

3.2.4 Direct control and serial attention: current debates

Two of the major debates in the current literature are on the allocation of attention and the directness of control in reading. The thesis model is a serial-attention model: samples from only one word are processed at a time. It is a direct-oculomotor-control model, meaning that saccade decisions are conditioned on the lexical processing stream, rather than some other autonomous process. Justification for these choices is provided below.

3.2.4.1 Serial Attention

Inheriting a debate from the attention community, models of eye movement control are split based on whether they view attention as essentially serial (i.e. focused on one word at a time), or parallel. Serial models are defended on several grounds: they are easy to formulate and interpret, because they always make it clear what is being processed; they are more strongly constrained than the parallel alternative, in the sense that a parallel attention model reduces to serial attention as a special case; and they provide word-order information to the sentence processing system ‘for free’.

Reichle, Liversedge, Pollatsek & Rayner (2009) make additional challenges to parallel models on the grounds that mental representations of words as separate units might be difficult to create if they are not processed ont at a time, and that word order might be difficult to compute if the perceptual input is not over word-size units. These latter arguments are, however, not compelling: Bayesian theories that assume a generative model of sentences and the availability of letter information (Legge et al., 1997, 2002; Bicknell & Levy, 2010b) can rationally update word-level beliefs from information of multiple partial words. In addition, as Reichle and colleagues note, the focus on word-level units is if anything a disadvantage for understanding reading in non-alphabetic languages like Chinese, or languages like Thai where word boundaries are not marked.

The challenge to purely serial models like E-Z Reader is accounting for spillover, preview, and parafoveal-on-foveal effects: all effects where the current word reading time is affected by properties of words that precede or follow. E-Z Reader addresses this challenge by imposing parallelism between oculomotor and lexical-attentional processing rather than in attention. By decoupling attention from eye position, the theory allows attention (and therefore processing) of one word extend into fixations on words that precede or follow. This explanation is not sufficient to cover parafoveal-on-foveal effects (when words forward of the fixated word affect fixation durations on the fixated word), but there is still an ongoing debate in the literature on the existence and robustness of these effects (Kennedy et al., 2002; Kennedy & Pynte, 2005; Kliegl et al., 2007; Drieghe et al., 2008).

The advantage of parallel attention models empirically is in their ability to account for the aforementioned parafoveal-on-foveal effects, as well as the full set of empirical facts serial models cover. The challenge to them is on grounds of predictive flexibility: theoretically, the space of parallel models can in principle span the range from full parallelism, through word-level serialism, down to phoneme- or grapheme-level serial models. Such models therefore require additional a priori estimates or constraints as to the extent of parallelism.

When it comes to the LLDT, this distinction may not be important, however. Given the short words and wide spaces used in the LLDT, parafoveal words may well be at far enough eccentricity for their processing to be minimal even in a fully-parallel model. A serial model is simpler to build, faster to simulate, and easier to interpret because it is easy to understand which words and fixations trigger each action in the model's near-optimal policies. It also greatly simplifies the choice of policy space for the model: a model that has information coming from multiple words at a time might also condition its behavior on multiple words at a time, whereas it is natural for a model that only takes in information from one word position to also condition its behavior on a single position. Therefore, a serial attention scheme is used for the dissertation model.

3.2.4.2 Direct oculomotor control

Another debate in the literature on eye movement control in reading is on the directness of the link between lexical and oculomotor processing. At one extreme of the spectrum are models that assume saccade targets and timings are primarily governed by visual features like the location of letters and spaces rather than ongoing lexical processing (e.g. Reilly & O'Regan, 1998). At the other are models that assume both saccade targets and timings are directly controlled by the lexical processing stream, though with the delays noted above (e.g. Reichle et al., 2009). In between are models that make intermediate choices, for example keeping saccade targets controlled by lexical processing but letting timings be largely autonomous (e.g Engbert et al., 2005). See Reichle et al. (2003, Table 1) for a more detailed review and taxonomy of models on this dimension.

This debate is driven by two fundamental facts about eye movements: on the one hand, standard estimates of saccade planning durations for sequential eye movements range as high as 200ms (Rayner et al., 1983), and typical fixations in reading are only slightly longer. Even lower bounds on how long a saccade takes to plan in a simple psychophysics task is at about 100ms (Becker & Jürgens, 1979). On the other hand, word-level properties do affect reading times (as reviewed in §2.2). The time left for cognitive processing to affect a direct eye movement control decision and yield these word-level effects is therefore only about 50-150ms, with at least about 50ms of this time taken up by eye-brain lag.

In light of these duration estimates, the reader appears to spend more time viewing a word *after* deciding to move on to the next one than he or she takes to makes the decision. There are two primary ways of understanding this fact. The first is taken by so-called primary oculomotor-control models. These models argue that the above facts mean that decisions of when and where to saccade are driven primarily by the same low-level perceptual considerations as the decisions in the sequential cued saccade tasks. Some example proposals are that the reader always attempts to fixate on the so-called optimal viewing position, just left-of-center (O'Regan & Lévy-Schoen, 1987), or to the longest word within some reasonable range to the right (Reilly & O'Regan, 1998), and to do so after a fixed amount of time. Other proposals claim that fixation targets are determined by the lexical processing system, but the timing is nonetheless controlled by an autonomous saccade clock (Engbert et al., 2005). In practice, all of these models must allow some higher level information to make it indirectly into the oculomotor stream, otherwise they will fail to capture the known effects of word- and sentence-level properties on reading times.

The second way of handling this empirical puzzle comes from from a key insight by Morrison (1984): that if saccade planning can occur in parallel with lexical processing, and given that saccades can be canceled after planning starts but before execution (Becker & Jürgens, 1979), then lexical processing can play a significant role in the where and when decisions in eye movements by operating in parallel with saccade planning. Among the most prominent models of this type is E-Z Reader (Reichle et al., 2009), discussed in greater detail in chapter 5.

The thesis model assumes that saccade timing is directly controlled by lexical processing (as delayed by saccade planning). When the lexical processing threshold is reached, saccade planning begins. The theory remains agnostic on the question of saccade targeting by assuming that in the LLDT both direct and indirect control of saccade targets will look quite similar and can be approximated by assuming saccades sequentially target each word. The choice is made in part in the interest of simplicity: understanding and interpreting what the model is doing is far easier when the actions of theoretical interest are under active control. But direct control was also used in ideal-observer approaches to eye movements similar to the thesis work (e.g. Bicknell & Levy, 2010a,b) so it has some empirical support.

3.2.5 How the model makes predictions

The model generates predictions by finding the optimal *policy* (i.e. a mapping from its probabilistic belief space to eye and hand actions), optimal in the sense that it maximizes the same payoff given to the human participants. To the extent possible, it makes these predictions without making direct contact with data from the LLDT. In the ideal case, there would be no 'free' model parameters: all architecture constraints would be determined a priori from previous research, and the remaining parameters would all be policy parameters optimized to maximize payoff in the same task given to human participants. This is in contrast to typical computational modeling outside a rational framework, where many or even all model parameters are 'free' parameters fit to data, and their values are part of the descriptive power of the model (with respect to the dataset it is fit against) and its predictive power (with respect to new datasets other than the one it is fit against).

In most practical cases, it is hard or impossible to eliminate all free parameters. In the dissertation model, the free parameters are the sample rate, and the noise in the sample generation and likelihood function. The former is fixed to 10ms as discussed above, and the latter estimated. To correctly estimate it while maintaining bounded optimality assumption, a nested fit is undertaken. First, near-optimal models with respect to payoff are found by varying policy parameters, over a range of noise parameters. Then, the noise parameters are selected for which near-optimal models best match human data.

3.2.6 Putting it all together: the full architecture

Figure 3.2 shows the diagram of the full model. On the right is a schematic view of the (M)SPRT lexical processor and its decision dynamics, indicating the flow of information (green arrows). On the left is the oculomotor architecture, indicating the sequence of events (blue arrows). The red arrows indicate the adaptive decisions. Note in particular the blue arrow between saccade planning and sample draws: consistent with Morrison's assumption, samples continue to be drawn during saccade planning.

This is the architecture used in the treatment of payoff adaptation in chapter 4. The decision dynamics and lexical processing components will change in the treatment of spillover effects, as memory and parafoveal processing force some changes in model assumptions. The task, oculomotor architecture and payoff all remain fixed throughout the dissertation.

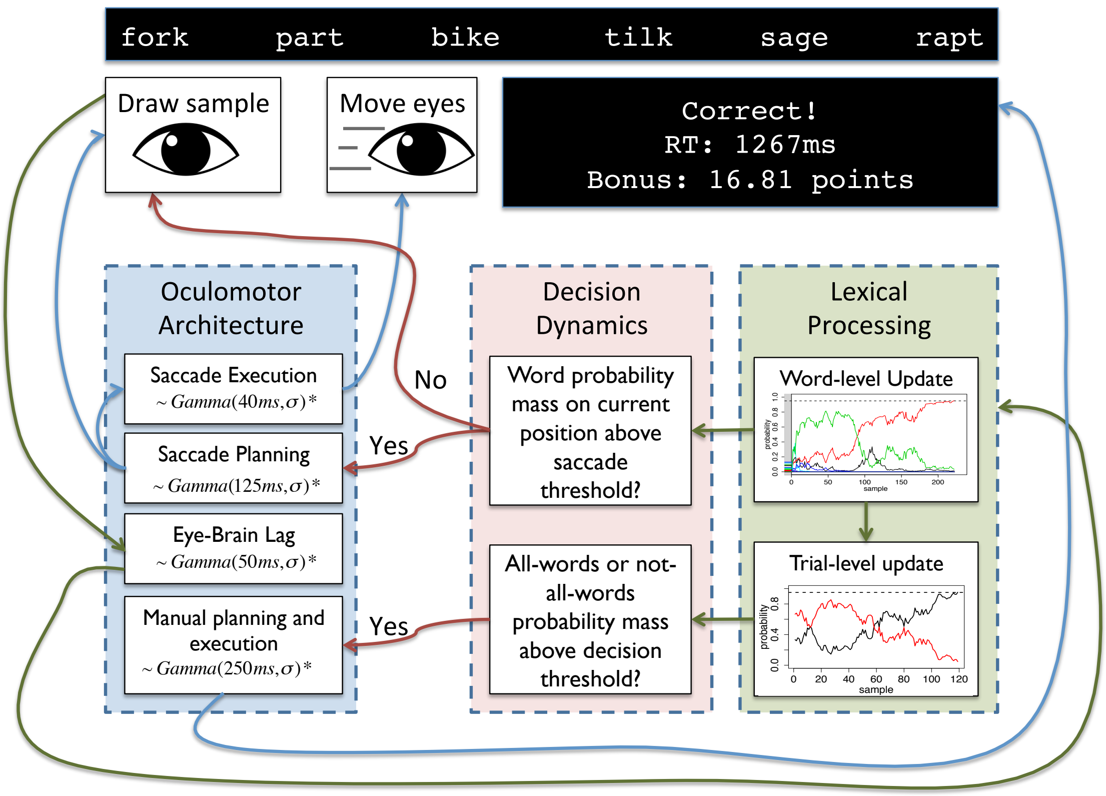

Figure 3.2: Architecture for the base model. The blue arrows reflect the sequence of events: samples are drawn during saccade planning or if the eye movement decision has not been made; saccade execution follows saccade planning, and is itself followed by the eye arriving at the next word; manual (buttonpress) planning and execution leads to the end of the trial and feedback. The green arrows reflect the flow of information: samples from the eye are delayed by the eye-brain lag before arriving at the word-level and trial-level updates discussed in the text; information from the content of the updates arrives at the decision mechanism. Red arrows reflect actions of the decision mechanism: motor planning starts if the word probability mass exceeds the threshold and does not otherwise; motor response planning begins if the task decision belief crosses threshold.

CHAPTER 4

Evidence for Adaptive Eye Movement Behavior in the LLDT

This chapter discusses results of a human experiment in the LLDT showing that participants do adapt their eye movement decisions on the single fixation level to a quantitative payoff manipulation. The same manipulation is undertaken in the model introduced in the previous chapter, and model results show a behavioral signature of adaptation similar to that in the human participants. These results are also reported in Lewis, Shvartsman & Singh 2013.

Task goals and context have long been known to have major effects on human performance in psycholinguistic experiments (for an early analysis see the seminal chapter by Forster, 1979). For example, in the area of single-word lexical processing, there are robust differences in how frequency and other important effects manifest in *naming* vs. *lexical decision* tasks (e.g. Grainger, 1990). Task context in the form of experimental list composition and goal manipulation via instructional emphases have significant effects, and have received detailed theoretical treatments (Wagenmakers et al., 2008).

There is also some work investigating task effects on eye movements in reading. For example, McConkie et al. (1973) have shown that participants tend to read longer when anticipating more difficult questions (for example, questions of a factual nature), as well as when they were financially incentivized to answer the questions correctly. Rayner & Raney (1996) have shown that the lexical frequency effect is eliminated when subjects read words in search of a target word rather than reading for comprehension. Finally, Wotschack (2009) replicated the question difficulty finding of McConkie et al. and also found that increasing the frequency of comprehension questions or instructing the participants to proofread led to slower reading speeds.

Recent work on visual attention in both linguistic and nonlinguistic contexts indicates that attention strategies are affected by task goals (e.g., Rothkopf et al., 2007;

	Accuracy	Balanced	Speed
Incorrect penalty	-150	-50	-25
Speed bonus (per second under 5s)	8	6.7	5.7

Table 4.1: Quantitative payoffs given to both model and human participants.

Ballard & Hayhoe, 2009), and more broadly there is a history of research showing that effects of strategic adaptation appear on multiple levels of complexity, from simple perceptual decisions (Green & Swets, 1966) to complex multi-tasking scenarios (Meyer & Kieras, 1997; Howes et al., 2009).

But while this prior work manipulates task type and difficulty, there has not been a manipulation of quantitative speed-accuracy tradeoffs of the kind pursued here. The advantage of this quantitative manipulation comes when adding the assumptions that humans are attempting to maximize the quantitative payoffs given. Under this assumption one can map model behavior under different payoffs to behavior of different participant groups under those same payoffs, and attempt to understand the graded nature of the adaptive process.

This is exactly what is done in this chapter. Three payoff were designed to impose different speed-accuracy tradeoffs for a given level of success, and were all defined in terms of a bonus for speed and penalty for incorrect responses. The bonus was continuous at the millisecond level, starting at zero points for responses longer than 5s and rising by a different number of points per second for each payoff. The three are detailed in Table 4.1. The labels of 'accuracy', 'balanced' and 'speed' are intended to reflect the intended emphasis of the payoff. For example, the speed payoff gives a smaller bonus for going fast, requiring participants to go faster than they would for the same time bonus in the accuracy payoff, but has a smaller inaccuracy penalty to yield an approximately equal payoff to the accuracy payoff. The accuracy payoff in turn gives a bigger bonus for going fast while imposing a larger inaccuracy penalty, allowing participants to spend longer and achieve more correct responses without sacrificing as many points as they would in the speed payoff. The payoffs were designed to be approximately equal in expectation based on an earlier model (Shvartsman, Lewis, Singh, Smith & Bartek, 2011).

Parameter	Mean	Std Deviation
Eye-Brain Lag	50ms	15ms
Saccade programming time	125ms	37.5ms
Saccade execution time	40ms	12ms
Motor preparation and execution time	100ms	30ms
Trial onset detection and refixation	150ms	45ms
Sample duration	10ms	0
Gaussian sample noise	0	1.2

Table 4.2: Oculomotor architecture parameters; all are fixed in advance as described in Chapter 1, except sample noise, which is fit as described in the text. Standard deviation for all but the sample duration and sample noise is fixed at 0.3 times the mean.

4.1 Model and simulation specification

The core of the model is as described in Chapter 3: a sequential sampler embedded in an oculomotor architecture. The key fixed parameters are summarized in Table 4.2.

The first four items in the table are architectural delays described in the previous chapter. The next item is a trial onset detection and refixation delay: participants in the LLDT see a gaze-contingent trial start, and it may take some time for them to detect the appearance of the list of strings and potentially refixate in a more optimal viewing position on the first string. This delay was simulated in the model by adding another Gamma deviate with a mean of 150ms and SD of 45ms in the beginning of each trial. While it provided a slight improvement of the overall quantitative fit, it was not of substantial theoretical interest and was dropped from the models in the remaining chapters. The next two items are sample rate and sample noise. Sample rate is fixed at 10ms as discussed previously, and sample noise was fit to single fixation durations, with only near-optimal policies considered.

4.1.1 Simulation details

To implement the simultaneous optimization of threshold parameters to payoff and noise parameters to fixation duraions, a grid search was undertaken, with values described in Table 4.3. The threshold parameter grid ranged from low enough to yield near-chance performance to high enough to be unreasonably accurate (and correspondingly slow). The noise parameter grid ranged widely enough to identify a clear minimum of error against single fixation duration. Each point in this grid was evaluated for 300,000 monte carlo simulation trials. Expected values were computed for each of the measures of interest (payoffs, reaction times, single fixation duration, etc.) by taking means of these values over the simulated model trials.

The words and nonwords in the trials were drawn from the experimental list used for human participants according to the same drawing rules as in the human experiment: the probability of a word trial was 0.50 (and thus a nonword trial 0.50), the nonword position was uniform in nonword trials, and words and nonwords were drawn uniformly from the lexicon. Across the 300,000 trials there were 50 different word and nonword lexicons of approximately 500 strings each. The word lexicons always included the experimental words and an additional set of words drawn uniformly randomly from the set of 1,500 English four letter words represented in Kučera & Francis (1967); the nonword lexicons always included the experimental nonwords and an additional set of nonwords constructed from letter bigrams that appeared in the word list. The model's performance is always evaluated on the words and nonwords from the human experiments, but for the model these strings are not distinguished in any way from the rest of the model's lexicon. Aggregating results across different model lexicons ensures that the results are not driven by a particular lexicon choice (though preliminary explorations indicated the results are robust against this choice, as well as the choice of lexicon size).

Statistical tests are not reported on the empirical measures that the model produces: at 300,000 simulated trials for each noise and policy setting, the confidence intervals around the reported means are negligible. From the expected values thus estimated, mean trial payoffs were computed for each architecture (determined solely by the noise parameter; the remaining architecture parameters were fixed as described in Chapter 1). A set of near-optimal policies[1] was found with respect to each of the three payoffs. Among these optimal policies, one was selected to yield the smallest root mean square error (RMSE) against the mean human SFDs, across all payoff conditions.

Parameter	Search range	Final value	Source
Saccade Threshold	0.80, 0.85, 0.86, 0.87, 0.88, 0.89, 0.90, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99, 0.995, 0.999, 0.9999, 0.99999	0.99 (Acc), 0.97 (Bal), 0.92 (Speed)	Maximizing payoff given task and architecture
Decision Threshold	0.80, 0.85, 0.90, 0.92, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99, 0.999, 0.9999, 0.99999	0.999 (Acc), 0.999 (Bal), 0.99 (Speed)	Maximizing payoff given task and architecture
Sample noise	1.05, 1.1, 1.15, 1.2, 1.25, 1.3	1.2	Minimize deviation from human SFD, **within optimal policies**

Table 4.3: Fit model parameters. Only the sample noise parameter is 'free' in the conventional sense (i.e. it is estimated by maximizing fit to human data), and even in this case it is only fit to a single empirical measure and evaluated against many. The remaining parameters (the thresholds) are fit by maximizing payoff in a particular condition (just as the humans are being asked to do).

4.1.2 The relationship between policy and payoff

Figure 4.1 provides views of the payoff surface in the two-dimensional policy space. In the first three panels, payoff is plotted against saccade threshold and each separate line corresponds to a separate decision threshold. In the fourth panel, payoff is plotted against decision threshold and each line corresponds to a separate saccade threshold—thus these are different views of the same 2-D payoff surface. Recall that a policy is simply a pair of threshold values (*saccade-threshold, decision-threshold*). The circled point at the top of each payoff plot represents the policy that yields the maximum expected payoff under this formulation of the policy space, and its value is given as the pair of numbers to the left of the point. The colored points represent policy points that are within 0.2 payoff units of the optimal. Values of each point are computed from means of 300,000 Monte Carlo trials.

Consider the Accuracy payoff graph at the left. This graph indicates that there is a flat region of the payoff surface when saccade thresholds are below about 0.85; this

[1] The thesis will always refer to these as *near*-optimal because they are not formally guaranteed to be the absolute best policies, due to the fact that the dissertation uses brute-force discretized search over the payoff function made noisy by randomness in architectural delays and sample generation. More accurate estimates of the value of the payoff function found by running more trials may help further separate these near-optimal policies into better and worse performers. Likewise finer discretization of the space may find slightly better policies close to the near-optimal ones found.

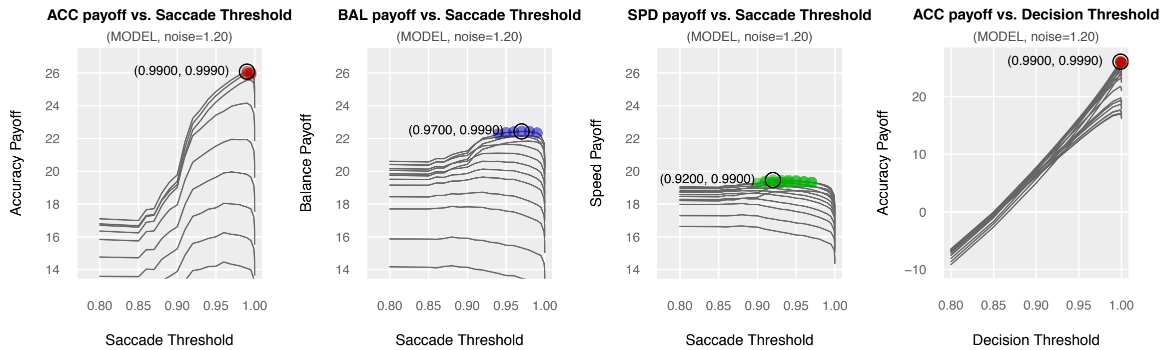

Figure 4.1: Expected payoffs generated by the model over the 2-D policy space defined by decision and saccade thresholds. The first three panels relate expected Accuracy, Balanced, and Speed payoff to saccade thresholds; each separate line corresponds to a separate decision threshold. The fourth (rightmost) panel relates the expected Accuracy payoff to decision thresholds; each separate line corresponds to a separate saccade threshold. (The Balanced and Speed payoffs, not shown, have rising but shallower slopes). The circled points represent the optimal policies, whose value is indicated at the left of the point. The colored points (*red* for Accuracy, *blue* for Balanced, and *green* for Speed) represent policies that are within 0.2 expected payoff units of the optimal point; thus the spread of these points in the Balanced and Speed payoffs reflects the flatter surfaces for those payoffs. Expected payoff values were computed over 300,000 simulated trials.

corresponds to thresholds where the saccade program is initiated almost immediately upon fixation. There is a steep rise in payoff as saccade thresholds increase—because more samples are obtained and accuracies are increasing significantly—up to a maximum point near a threshold of 0.99, followed by a steep decline as the additional gain from increased accuracy diminishes and the time cost begins to dominate. This relationship holds for most of the good performing decision thresholds. The relationship between decision threshold and payoff has a similar but simpler profile over the range explored: a steady increase in payoff as the decision threshold increases, followed by a steep drop as the time cost begins to dominate.

The Balanced and Speed payoffs have a similar profile as the Accuracy payoff—for the saccade thresholds, a flat region, a rise, and a sharp drop. But the qualitative shape differs considerably in the region of the maximum; payoff surface is considerably flatter for the Balance payoff, and very flat for Speed. Thus there is considerably more spread in the range of thresholds that perform within some close threshold of the optimum. The visualization of the payoff space suggests more success in separating the Accuracy condition from the other two than in separating Balanced and Speed from each other, a fact somewhat consistent with the human data (Figure 4.2.

The graphs in Figure 4.1 provide a visualization of the fundamental basis of the model's link between task payoff and predicted behavior: they depict the nature of the adaptive space that the human participants must navigate according to the model. The overall differences in payoff levels suggest that at least in the case of the model, the attempt to make the payoffs approximately equal in attainable points was not successful. While slightly unfortunate, this is not very important as long as the model does show differences between policies that are near-optimal with respect to different payoffs.

The next step is to relate these same policy points to predicted behavior directly – to do this, the next section will detail the human experimental methods, so that the model and human results can be reported in parallel subsequently.

4.2 Human experiment methods

4.2.1 Participants

Sixty-one members of the University of Michigan community participated in the experiments. Data from thirteen were unusable due to calibration problems, failure to complete the experiment, or equipment malfunctions, leaving a total of 48. Participants were given a baseline of $10 for participation, plus a bonus of $1 for each 1000 points they earned in the task.

4.2.2 Stimuli

Participants responded to 200 trials of the LLDT divided into 10 blocks, preceded by a 10-item practice block. There were two types of trials in each block: half of trials contained all-words lists, and the other half contained 5 words and one non-word in a randomly drawn position. Words were all 4 characters long and drawn from a 234-word subset of the Brown Corpus (Kučera & Francis, 1967), containing 117 high-frequency words (mean frequency count 239.2, SD 186.0) and 117 low frequency words (mean frequency count 5.6, SD 12.8). Nonwords were also all 4 characters long, and were drawn from a list of 53 nonwords pronounceable according to English phonotactics. While this means that participants saw each string more than once, the number of times a string was seen had no significant effect on fixation durations (effect of -0.64 ms per time seen, $p=0.61$).

4.2.3 Procedure

Each participant was assigned to one of the three payoff conditions used to make predictions from the model. They were not told the name of their payoff, only a quantitative description of the requirements (e.g. "You will receive a point for each 125 milliseconds by which your response is faster than 5000 ms (5 s). You will lose 150 points if your response is incorrect. You will get a $1 bonus for each 1000 points.").

Items were presented on a CRT monitor in a 20pt Courier font, separated by 6 characters of whitespace. This resulted in each word covering 0.7 inches or 1.6 degrees of visual angle, and whitespace covering 1.48 inches or 3.4 degrees of visual angle at a distance of 25 inches from the screen. Each trial started with a fixation dot at the location of the first string. The entire six-string list would appear once subjects fixated, and the trial ended after subjects responded using a Cedrus response box. Eye movements were measured using an SR-Research Eyelink II head-mounted eye-tracker operating at 500Hz. Single-point drift correction was performed before every trial.

Statistical Methods Data analysis on the human data was carried out using mixed effects regression (Pinheiro & Bates, 2000) using the `lme4` package (Bates, Mächler, Bolker & Walker, 2014) for the R environment for statistical computing (R Core Team, 2014). For inference, models with maximal random-effects structures were fit: in trial-level analyses of condition this included by-participant and by-trial random intercepts, and by-trial random slopes (by-subject random slopes are not necessary because ours was a between-subjects design). In string-level analysis of condition this additionally included random slopes and intercepts of word and list position. In string-level analysis of frequency this included random slopes and intercepts of position but only random intercept of word (since frequency is a between-word factor). In string-level analysis of position this included a random slope and intercept of word. Linear models were fit to all timing measures, and a logit model was fit to accuracy.

Response times and single fixation durations (SFDs) that were farther than 3 standard deviations from the mean of those respective measures were removed. Some additional response times had to be removed due to a bug in response collection code. To mitigate against the possibility that some participants were quicker to adapt to the payoff than others, all empirical measures were tested both for the full set of participants and for only the top half of participants (in mean trial payoff) for each condition. There was no difference in the significance or direction of reported effects between the full group and the top-half group; however, numerically some of the results from the top group appear closer to the model's performance. Fixation durations from the first and last strings in the list were excluded to eliminate noise associated with response selection and gaze-contingent refixation. In addition, words that appeared after a nonword in a trial were also excluded from the analysis. Over 17,000 single fixations remained in the dataset.

Hypothesis tests were conducted using a single pair of normalized orthogonal contrasts. The first contrast, and the one of theoretical interest, is the contrast between the accuracy and speed payoffs (i.e. coding accuracy and speed as ± 0.5 and balanced as 0). The second contrast is included for orthogonality (coding accuracy and speed as 0.408 and balanced as -0.815) but is not theoretically informative, so its results are not reported. This contrast design allowed the use of the balanced condition for purposes of improving error estimates, increasing overall power. Log frequency and position were treated as numeric linear predictions for the purposes of those hypothesis tests. The effect of each contrast is reported as the p value of a likelihood ratio test (using the Chi squared distribution) comparing two models identical except for the presence of the contrast set of theoretical interest.

Reported condition means are estimates of fixed effect of condition in a model containing random intercepts as described above, but no random slopes. This allowed the use of Markov-Chain Monte Carlo (MCMC) sampling to estimate 68.2% (i.e. 1 standard error on each side) highest posterior density (HPD) intervals around these estimates. While HPD intervals derived from such models may be overly small (i.e. anti-conservative), they are useful due to being more interpretable than equivalent confidence intervals: the HPD interval is simply the interval within which the mean is likely to lie with this probability. Table 4.6 details the specific contrasts tested, and the sections below walk through them in more detail.

4.3 Human and model results

4.3.1 Trial level effects

Figure 4.2 shows the results for key trial-level measures: response times (on correct and incorrect trials), and percentage correct. The top row is the set of human results, and the bottom row is the set of model results from a set of policies at or near optimal (within 0.2 payoff units) plotted. Table 4.4 shows condition means for human participants and model.

	Response Time (ms)		% Correct		Payoff	
Condition	Human	Model	Human	Model	Human	Model
Accuracy	1667	1644	92	98	12.91	23.44
Balanced	1548	1546	87	97	16.08	21.52
Speed	1494	1455	88	95	16.65	18.89

Table 4.4: Trial-level measures; payoff reported is mean payoff per trial. See Table 4.6 for significance tests.

The key empirical result here is that the human participants show marginally decreased accuracies and response times across the Accuracy-Balanced-Speed payoff conditions: the accuracy condition results in slower response times (Table 4.6, contrast set (a)) and higher percentage correct (Table 4.6, contrast set (b)). The model predicts this trend because the optimal thresholds for Accuracy are higher (Table 4.3) than Speed, leading to slower but more accurate responses. This is also consistent with the design goals of the payoff scheme, which is designed to encourage on average slower but more correct responses in the accuracy as compared to the speed condition. There is a significant discrepancy in predicted accuracies, addressed later in the chapter.

The model correctly predicts that correct word trials will show slower responses than incorrect trials, with the converse holding in nonword trials (a reliable cross-over interaction, Table 4.6, contrast set (c)). This result is a consequence of the fact that 'all-words' responses (correct or incorrect) tend to come after reading the full list, whereas 'not-all-words' responses tend to come after only reading a subset of words. Although this is not surprising behavior for the humans, the model need not have behaved this way: there are suboptimal strategies in the explored space that set the decision threshold low enough that an 'all-words' response are made before all strings are read, and ones that set it high enough that 'not-all-words' responses are not made until after the sixth string.

4.3.2 String level effects

Figure 4.3 shows the results of key string-level measures: single fixation duration across fixation types and by frequency class. Single fixation durations are reported because this is the measure that the model is able to quantitatively predict (refixations, regressions and skips are all not in the model). 71% of the strings were fixated only

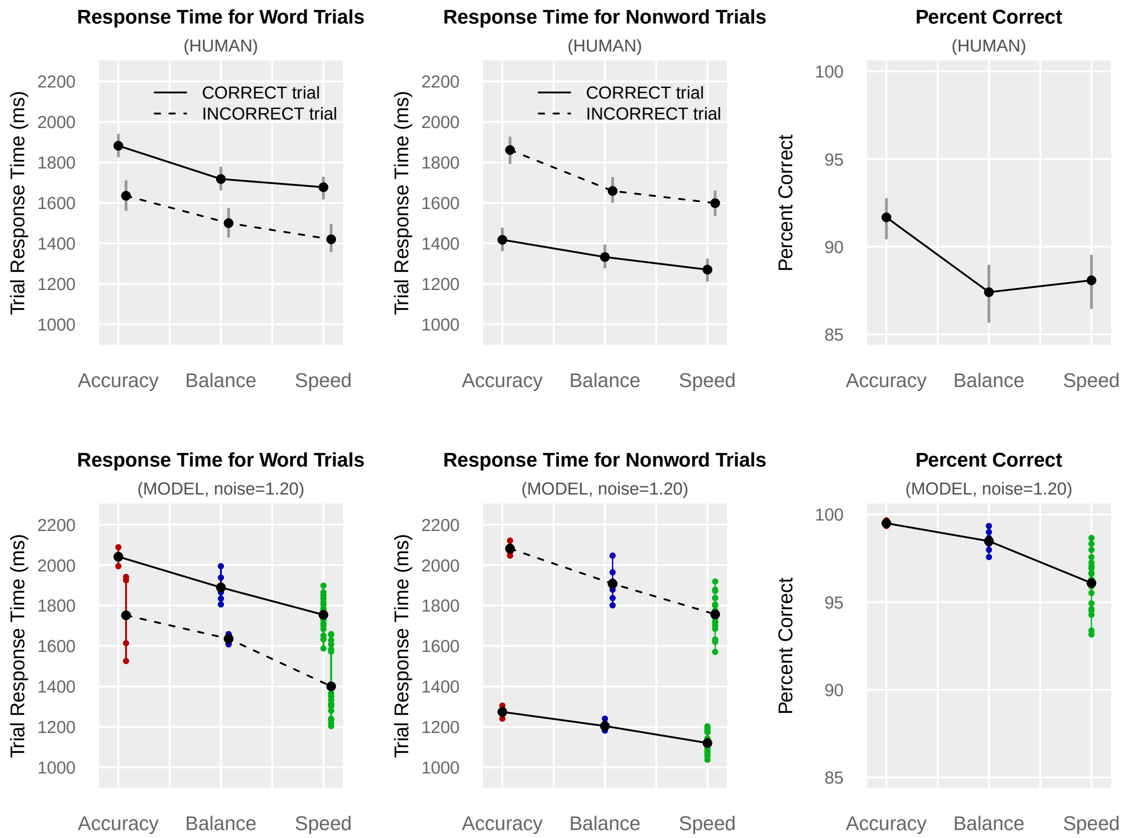

Figure 4.2: Empirical measures at the level of the trial for the full set of human participants and the computational model. The colored points represent predictions corresponding to the best-performing policies identified in Figure 4.1, the lines connect the means of this set. The error bars on the human data correspond to one standard error estimated from posterior densities of the mixed effects models.

once, and other fixation-level measures (e.g. first fixation and total fixation times) do show the same patterns[2].

The key result here is that the human data is consistent with the model prediction of slower fixation durations in the Accuracy as compared to the Speed condition (Table 4.6, contrast set (d)). The model provides a straightforward explanation of this effect: when a payoff provides pressure to respond more correctly, a higher saccade

[2]One interesting pattern the human results show that the model cannot recover is occasional regressions from the final string to suspected nonword positions (both true nonwords and low-frequency words). This pattern is broadly consistent with the rational explanation for regressions provided by Bicknell & Levy (2010a): as more words are read, the probability of an all-words trial increases, reducing the probability of a nonword in that past position, and resulting in a regression.

threshold will increase the amount of information acquired, increasing the likelihood of a correct response.

		Single Fixation Duration	
Condition	Frequency	Human	Model
Accuracy	High	253.57	255.59
	Low	274.47	294.24
Balanced	High	227.23	227.37
	Low	244.46	256.71
Speed	High	236.34	205.40
	Low	251.29	224.30

Table 4.5: Single Fixation Durations for words in the high and low frequency bins, indicating an effect of lexical frequency. See Table 4.6 for significance tests.

The model also correctly predicts an effect of log frequency on fixation durations: fixations on highly frequent words are faster (Table 4.6, contrast set (d)); see Table 4.5 for numeric estimates). Norris (2006) has already shown that this is a consequence of an otherwise unconstrained ideal observer model making decisions on single words, but the comparable magnitude of effects between the human was not a necessary outcome of the model embedded in an architecture (see below for similar architectures yielding the wrong effect magnitude).

The model also predicts a larger frequency effect in the Accuracy condition as compared to the Speed condition, for the same reason: thresholds are set lower in the Speed condition and fewer samples are obtained during the adaptively controlled, pre-saccade-programming stage of sampling that affects fixation durations and thus frequency effects. The human data are numerically consistent with this effect, but the interaction is not reliable.

The model makes other interesting and more subtle predictions attested in the human data (Figure 4.4, left two columns). First, nonwords are read more slowly than words. In the model, this is a consequence of the fact that the prior probability that any given string is a word is much higher than a nonword. It therefore takes more evidence (more sampling time) to reach the nonword threshold. Note that these predictions of the model are valid in spite of the simplistic nonword representation chosen: a reasonable belief representation that does not explicitly assign probabilities to nonwords will still be far likelier to expect known words than otherwise, and the

SPRT provides a mechanism that explains why less expected visual items might take longer to identify.

Furthermore, the word-nonword difference is predicted to be larger in the Accuracy condition than Balance and Speed, an effect that appears in the human data as a reliable interaction between condition and string type. This reflects a nonlinear effect of distance-to-threshold on the expected number of samples required for decision in the SPRT, the same mechanism that yields the logarithmic frequency effect on single fixation durations and its interaction with payoff condition.

In addition, the effect of trial accuracy is different for words and nonwords (another interaction): fixation durations on words are about the same in correct and incorrect trials, but the model predicts that nonwords are read more *quickly* in incorrect trials, and the human data shows this pattern (Table 4.6, set (g)). In the model this arises because the prior is closer to the word threshold than the nonword threshold, and the expected number of samples (i.e. time) to reach threshold is a function of the distance from prior to threshold. Therefore, incorrect random walks to the word threshold take fewer samples than correct random walks to the nonword threshold. This latter insight will prove helpful in understanding spillover frequency effects in the latter portions of the thesis.

Finally, the model predicts that strings in later positions are read somewhat more slowly than strings in early positions for all three payoff conditions; this effect, though tiny, is also reliable in the human data (Table 4.6, contrast set (e)). This is another consequence of the list-level Bayesian update. The reason is a somewhat counterintuitive property of the probabilistic structure of the task: as evidence is accumulated identifying strings as words in the list, the probability of an all-words list increases—but the probability that any one of the individual strings remaining is a nonword *increases* slightly. Thus, the prior belief that each remaining string is a word is slightly lower, and in expectation additional samples are required to hit the word threshold. This shows that participants are sensitive to the subtle inter-word dependency in the LLDT, one of the task's desired properties. Unfortunately this suggests a misalignment between the probabilistic structure of the LLDT and normal reading, which is not reported to show such a systematic slowdown.

Contrast set	Effect	Estimate	p
(a)	Condition on RT	-180.36	0.09
(b)	Condition on % Correct (logit)	-0.40	0.08
(c)	Trial Type (word vs nonword) on RT	**-355**	<**0.001**
	Correctness (correct vs incorrect) on RT	**-91**	<**0.001**
	Correctness x trialtype interaction on RT	**-596**	<**0.001**
(d)	Condition on SFD	**-21**	**0.04**
	Frequency on SFD	**-4.45**	<**0.001**
	Frequency x Condition interaction on SFD	1.64	0.29
(e)	Position on SFD	**3.26**	**0.007**
	Position x Condition interaction on SFD	-0.65	0.44
(f)	String (word vs nonword) type on SFD	**97**	<**0.001**
	String type x Condition interaction on SFD	**-61**	**0.05**
(g)	String (word vs nonword) x Correctness (correct vs incorrect) on SFD	**71.2**	<**0.001**

Table 4.6: Coefficient estimates and p-values calculated using a likelihood ratio test between two linear models identical except for the presence of the tested predictor. Lines separate different linear models. Condition was coded as a set of orthogonal contrasts; reported here is the Accuracy vs. Speed contrast.

4.3.3 How could it be otherwise? Why a match in adaptation and its locus is important

The previous two sections showed a two-way correspondence between humans in model both in the fact that they adapt to payoff, and in the fact that they adapt along the same locus (individual fixation durations). This two-way correspondence need not have been the case on either the trial payoff or the fixation duration level, and there are reasonable alternative possibilities for both.

In the case of trial-level payoff adaptation, Both the model and the humans could have failed to respond to the payoff. The humans could have failed to respond to payoff if mapping the abstract numeric payoff to moment-by-moment strategy differences was too difficult in the short duration of the task – because the cost of creating a new strategy was higher than using an old one, because the granularity of humans' internal strategic representation is not sufficient to distinguish between the three designed payoffs, or for some other reasons. They also could have failed to respond to payoff because of other independent pressures like overall task difficulty (putting them at ceiling or floor), or conflicting task goals (such as finishing the experiment and leaving).

Likewise the model could have failed to respond to payoff if its cognitive architectural constraints overconstrain the policy space. For example, if saccade planning duration is long enough, any word will be (nearly) perfectly identified during saccade planning, and active (threshold-controlled) sampling is not needed. In such an architecture the best threshold would be as low as possible, regardless of payoff, and optimal performance for all payoffs would look identical.

Even given that there was an a correspondence between model and humans in the effect of payoff on trial response times, there need not have been that correspondence in the case of fixation durations. The model could have behaved in other ways to yield the same response time differences, for example by keeping the saccade threshold the same across payoffs and only varying the trial decision threshold. This would yield an unnatural (for humans) behavior where more or fewer words are read as a result of different speed-accuracy tradeoffs. The humans, too, have this unusual strategy available and do not choose it. They could have also adapted via other strategies such as minimizing mind wandering, increasing their attentional resources, etc. Rather than (or more likely, in addition to) these other adaptations, humans show evidence of adapting their moment-to-moment saccade timing to the differential speed-accuracy pressures.

Even finding the benchmark frequency effect in the model is not a trivial finding, as any frequency effect whatsoever was not a necessary consequence of the architecture-bounded model: the effect can disappear with sufficiently low saccade thresholds, even though the model can still perform the task. Indeed, in the Speed condition, many good-performing thresholds nearly make the frequency effect disappear. The reason is that thresholds can be set that are near the prior belief in expectation, so a saccade program is initiated immediately, before samples are taken. There is therefore little or no opportunity for frequency to affect fixation duration, but samples continue to arrive during the saccade programming time so the task can be performed. A different set of architecture bounds would also eliminate the effect: with a saccade planning time long enough to identify the word and perform the task, low saccade thresholds will perform increasingly better and again eliminate the frequency effect.

4.4 How does architecture shape adaptation?

The reported model predicts a detailed empirical signature of adaptation across task payoffs that is attested in the human subjects, at the level of single fixation duration. It does so by optimizing its policy (and therefore behavior) under the joint constraints of task and oculomotor architecture. The previous sections and human empirical design test its varied predictions under different task constraints. What follows next are two explorations of how the oculomotor architecture might affect the model's predictions, the first analytical and the second computational. The goal here is not to compare qualitative changes in architecture like changes from serial to parallel attention. Rather, the intent is to take advantage of the parametric nature of the machine architecture and see how the quantitative detail of the machine might make a difference in predictions.

The first argument is made analytically: recall the discussion in Chapter 1 of eye-brain lag and saccade-planning time. The eventual values selected were from a recent version of E-Z Reader (Reichle et al., 2009) and consistent with shorter estimates in the literature (EBL 50ms: Clark et al., 1994; Seeck et al., 1997; SPT 125ms: Becker & Jürgens, 1979). But imagine that the values selected were instead the higher ones used in earlier E-Z Reader variants (Reichle et al., 2003) and consistent with higher estimates in the literature (EBL 90ms: Mouchetant-Rostaing et al., 2000; SPT 240ms: Rayner et al., 1983). These estimates would yield a lower bound on expected fixation durations of 330ms, since each fixation must include both saccade planning and eye-brain lag, but does not have to include any sampling before saccade planning begins.

Now recall the optimization procedure: first, optimal policies are found for each noise setting, and then the best noise setting is found based on correspondence between SFD in the optimal policy for that setting and the human SFDs. Human SFDs in the LLDT are approximately 250ms, so the noise setting that will be selected is the one at which SFD is as close as possible to 250ms, which will be at the expected minimum of 330ms. There will be some noise settings for which such a policy is optimal and saccade planning begins immediately, and these settings and policies will be selected. At the very least, these policies will show neither frequency effects nor effects of payoff on fixation durations.

Without simulation it is not clear whether they will also fail to recover other effects, but the loss of these two effects is sufficient for the argument that bounds play a strong role in adaptation. It is important to note that this extreme model is still an adaptive, bounded optimal control model – it just so happens that the same adaptive solution is optimal under any task settings. By this argument one can conclude that long saccade planning and eye-brain lag duration are not compatible with some aspect of the rest of the architecture in being able to yield human-like behavior under the bounded optimality assumption.

The second argument is made by simulation. Consider a set of models with increasingly simple and short oculomotor dynamics. As in the analytical argument above, these are not 'lesioned' variants of the model. They are each valid models in their own right, each with new best-fitting noise parameters and new optimal policies. The simplest of these is a *minimal model* that dispenses with eye-brain lag, saccade programming time, and saccade execution times. In that sense, it is the simplest kind of ideal observer model that can accomplish the task by controlling the timing of sequential saccades. The next two architectural variants rely on the insight above that the duration of saccade planning (and therefore 'free' post-threshold sampling) is the aspect of the architecture that most strongly interacts with the policy and therefore the interesting adaptive component of the model. The first maintains only the saccade programming time (thus eliminating eye-brain lag and saccade execution duration), and the second keeps eye-brain lag and saccade execution and eliminates saccade programming.

Figure 4.5 shows the best achievable fit available to each of those models, against SFD (the metric to which the noise parameter is fit). The model with only saccade planning performs nearly as well (in terms of recovering human data patterns) as the model with the full architecture. The minimal model does least well, consistent with the notion that the architecture is improving predictive power, but – remarkably – the architecture lacking only a saccade planning delay performs nearly as poorly. This is consistent with the insight above that saccade planning is driving most of the impact of the architecture.

These differences are driven by the following sequence of changes: First, as expected, less 'free' sampling means that optimal policies would require higher thresholds to achieve similar accuracies. The sharper, taller peaks in figure 4.6 show this for the most striking case, the minimal model. Notice that the optimal policy in the speed payoff in the minimal model is the same as the optimal policy in the accuracy payoff in the full architecture, and thresholds go higher from there.

Next, with a bigger role for threshold in the model's predictions, all effects that rely on distance from prior to threshold are magnified. See Figure 4.7, which shows single fixation durations by frequency class, and compare it to the bottom right panel on Figure 4.3, noting the difference in y-axis scales. Both the effect of payoff and the effect of frequency (high vs. low) are magnified. Recall from Figure 4.5 that the fit to human data cannot be made better by adjusting the noise parameter; it is already maximizing fit to single fixation durations possible within the space of optimal policies. This also illustrates again that with a single free parameter, the model (whether with a full oculomotor policy or a minimal one) is substantially restricted in its space of predictions.

4.5 Discussion

4.5.1 Summary of major results

The key challenge in this chapter was demonstrating payoff adaptation in humans in the LLDT, and showing that the model can correctly recover this adaptation along the same adaptive locus chosen by the humans. The between-subjects payoff manipulation in the LLDT was successful, and the detailed empirical predictions of the model were largely supported in the human results, including the key result on the locus of adaptation: single fixation durations in the human data were modulated by payoff in ways predicted by the optimal control model.

Next, the adaptation to architecture as well as payoff was explored by manipulating the model architecture. The first key finding from that exploration is that optimal policies in different architectures clearly differ. The second is that the modified architectures without saccade programming time provide poorer fits to the human data, even if the single free noise parameter is allowed to be refitted. In particular, they overestimate the size of the payoff effects and frequency effects.

4.5.2 Comparing the LLDT to normal reading

The goal in designing the LLDT was to create a task that was similar to sentence-reading while maximizing the number of single fixations, minimizing regressions and word skipping, and keeping the overall reading patterns amenable to simple analysis. The scanning pattern was primarily left-to-right as designed: 71.2% of strings in the analysis were fixated once. This appears to be well above typical rates of SFDs in reading experiments, likely driven in large part by the short length and wide spacing of the strings in the LLDT.

The proportion of regressions originating in strings in the analysis (i.e. excluding first, last, and post-nonword strings) was 7.5%, consistent with recent reading experiments: Engbert et al. (2005) shows regression out probability between 1% and 6%; Reichle et al. (2009) shows regression out probability between about 8 and 11%, and Levy, Bicknell, Slattery & Rayner (2009) shows regression out probability between about 10 and 20%. The proportion of regressions originating from the last string, however, was far above: as many as 50% of final strings were the source of a regression, often to a nonword or low-frequency word (Sanders, 2013). This is likely higher than in ordinary reading, but it is difficult to make a direct comparison, since sentence-final words are usually excluded from sentence-reading studies. However, the regressions to nonwords or low-frequency words (likely suspected nonwords) are consistent with sentence-reading behavior where participants regress to points of particular difficulty like the beginning of ambiguous regions (e.g. Frazier & Rayner, 1982). One point of departure from ordinary reading that LLDT behavior does make is the position effect: in the LLDT, reading times get slower farther into the list. This is in contrast to the commonly accepted pattern in sentence-reading.

4.5.3 The Accuracy discrepancy

There are two key discrepancies between model predictions and human performance: first, that the model achieves higher accuracy overall than the human participants. Second, and possibly related to the first, the maximum achievable mean trial payoff in the model in the accuracy payoff condition is higher than that in the speed payoff condition, whereas humans empirically achieve a higher mean trial payoff in the speed payoff condition than in the accuracy payoff condition. These discrepancies cannot be straightforwardly addressed by simply increasing the noise parameter. The model would adjust its thresholds to maintain higher accuracy levels (and in doing so increase fixation durations and response times) to maximize payoff, illustrating the constrained nature of the predictions. Thus there is a genuine discrepancy that cannot be explained by the present model. There are a few possible reasons for this discrepancy, likely acting in concert:

Fragility of Good Performance in Accuracy payoff Recall from Figure 4.1 that the peak in the Accuracy payoff is far more fragile than the one in the speed and balanced payoffs. Under any payoff, the optimal policy is balanced between two extremes: one where the model takes enough time to rarely receive the large inaccuracy penalty, and one where the model is so fast that getting the inaccuracy penalty is mostly compensated for by the RT bonus. Moving from the optimum in either direction worsens performance, but this drop is far faster in the accuracy payoff than the other two. The model is not sensitive to this fact in any way: the optimization happens entirely offline and the best policy can generally be found, no matter how fragile or unusual it might be. However, the humans have a much shorter time to adapt to the task, so one can imagine that this narrow range is lost amidst the variance present in the task and the noise in the subjects' own internal processing. This idea is consistent recent work by Zacksenhouse, Bogacz & Holmes (2010), who show that the behavior of lower-scoring participants is consistent with maximin rather than reward maximization, i.e. maximizing some degree of guaranteed performance under a presumed level of uncertainty rather than finding absolute maxima. Exploring this possibility would require a theory of adaptation or learning in the task.

A second possibility is that even if participants can find the right threshold and are attempting to maximize reward, their ability of setting thresholds is not fine-grained enough to target this policy even if they are able to surmise its existence. This latter possibility can be explored in the model by introducing noise in the threshold setting. This would only the reversed best achievable mean trial payoff issue, not the overall accuracy discrepancy.

Nondecision Error Human participants may occasionally make errors not strictly related to their decision process: they can engage in mind wandering, they can execute an erroneous or involuntary manual response, or make similar non-decision-related actions, contributing to a higher error rate. To estimate the nondecision error, recall that participants (and model) tend to respond on the nonword to nonword trials, and on the last word on word trials. Presumably then, 'all-words' responses on the nonword in nonword trials should be quite rare, and mostly a signature of 'nonword' responses that were accidentally realized as 'all-words' responses. The percentage of those (as a proportion of total errors) should be on the upper end of possible estimates for non-decision error, for nonword trials. For all-words trials, it is likewise possible to exploit the fact that participants tend to respond on the end of the trial. Therefore, the proportion of 'nonword' responses on the final word of all-words trials should be an estimate of nondecision error on all-words trials.

The overall proportion of such errors in the LLDT (in both directions) is approximately 10% of total errors (varying between 6% and 25% depending on payoff condition and trial type). Even taking the highest error rate in the model (about 5%, in the Speed payoff condition) and boosting it by 25% would increase the error rate only to 6.25%, which is still lower than the error in the *Accuracy* payoff condition in the humans. Therefore, nondecision error cannot be the sole reason for the accuracy discrepancy, though it can contribute to perhaps a percentage point of the difference.

Reward rate or time optimization Unlike the model, participants do have other goals and plans in life besides the LLDT, and they know that when they are done with the experiment they can go and pursue them. With the exception of a very few participants genuinely excited about helping science or with enough competitive spirit to attempt to set a performance record in the task, LLDT task participants are typical undergraduate volunteers who are far more interested in getting their \$10 baseline quickly than working hard to receive their bonus (which will vary between \$2 and \$4 for most participants). Even without this large baseline, one can imagine the additional opportunity cost imposed on the participant by spending additional time in the testing room as compared to going about their day. In light of this, it is surprising that the payoff manipulation worked as well as it has in this sample, and even more surprising that it worked in unpaid pilot groups.

The overall increase in time pressure that reward rate optimization would engender should yield faster but less accurate performance overall, and would disproportionately affect the Accuracy payoff. It can be explored in the model by optimizing reward rate rather than mean trial reward, or empirically in humans by changing the proportion of the baseline payment to the achievable bonus.

Uncertainty over nonword identity The model has veridical knowledge of which strings are classified as words and which are nonwords. But some of the very low-frequency words on the experiment list (e.g. *bard* and *nigh*) may simply not be in the participant's lexicons, and no amount of additional sampling will overcome such errors. These therefore represent a class of incorrect responses which are not possible for the model to make. One can address this empirically in the humans, by explicitly testing the full list of strings without time pressure post-experiment, and excluding trials with words not known by the participant.

Another way of tackling this is addressing it in the model, in one of a few different ways. One is giving the model some uncertainty over which strings are words and which are nonwords. The simplest way of doing that is giving each string a 'wordness' percentage based on its phonotactics or some other rating, and scale the posteriors accordingly (vs. each string being 100% word or nonword). It is unclear, however, what this 'wordness' rating would be, and how it would be independent of frequency and similarity to other words, both of which are already in the model (the latter via neighborhood density). It also has the undesirable property of bringing a 'wordness' signal into the MSPRT, but as noted by Norris (2009), one of the advantages of the (M)SPRT over related models like Ratcliff's DDM is exactly the replacement of a poorly-explained 'wordness' signal with a well-motivated and easily explainable appeal to prior experience.

A solution to this problem is to generate this wordness signal from a generative model of nonwords that does not explicitly represent each nonword. Such a model would provide an interpretable continuous 'wordness' signal of exactly the right kind: the probability that a string is a (phonotactically and otherwise) valid English string. But it is unclear what exactly such a model looks like: it seems like the main distinguishable property of well-designed, pronouncable, phonotactically valid nonwords (as all the nonwords in the LLDT are) as compared to words is precisely the fact they are nonwords. A related solution is to find a sequential test with H_0='samples are drawn from distribution X' and H_1='samples are *not* drawn from distribution X'. In the context of the (M)SPRT, this would require a way of computing the likelihood that the sample is drawn from none of the hypotheses.

Finally, one can consider a related set of tests that can reject all hypotheses. Consider an SPRT choosing between hypotheses H_1 and H_2 when evidence is actually generated according to some other hypothesis H_3 – for example, the case where the test is determining between Gaussians with two means μ_1 and μ_2 when the actual sampling mean is some $\mu_1 > \mu_3 > \mu_2$. In a degenerate case, one can imagine that μ_3 yields samples equally likely under both H_1 and H_2, and therefore the SPRT never converges. Even in more reasonable cases, the test can still do worse than the fixed-sample test, as noted by Anderson (1960), even though the SPRT in general does far better (Wald, 1947). A truncated test with a time-dependent threshold is provided by Anderson that is robust against the possibility of a third hypothesis and has a better bound on the expected maximum sample size than the SPRT in this condition. But this test still accepts one of the two original hypotheses. Perhaps a test can be found that is robust against this property but actually makes the rejection decision like a one-sided test. Such a test could be used to select one of the words, or reject the possibility of any of them, consistent with the intuition of what happens in lexical decision.

4.5.4 The speed of payoff adaptation

The payoffs used in the experiments reported here are fairly subtle: while the penalty for incorrect trials is fairly large, especially in the Accuracy payoff condition, the distinction between e.g. 6.7 and 5.7 points per second is not substantial. One might therefore expect that the adaptation reported earlier in the chapter appears gradually as participants get more familiar with a particular payoff scheme, and that behavior would look similar across payoffs during the first few trials of the experiment, diverging as time goes by. This is not the case, however. Instead, the gap between payoffs exists starting at the first test block (following a half-size practice block), as indicated in figure 4.8. This is somewhat surprising but not entirely inconsistent with previous results. For example, Simen et al. (2006) provide an algorithm for threshold-setting in a diffusion model for a simpler two-alternative forced choice task that can adapt to a new threshold in as few as 20-30 trials. One might expect a similar mechanism to be in play here.

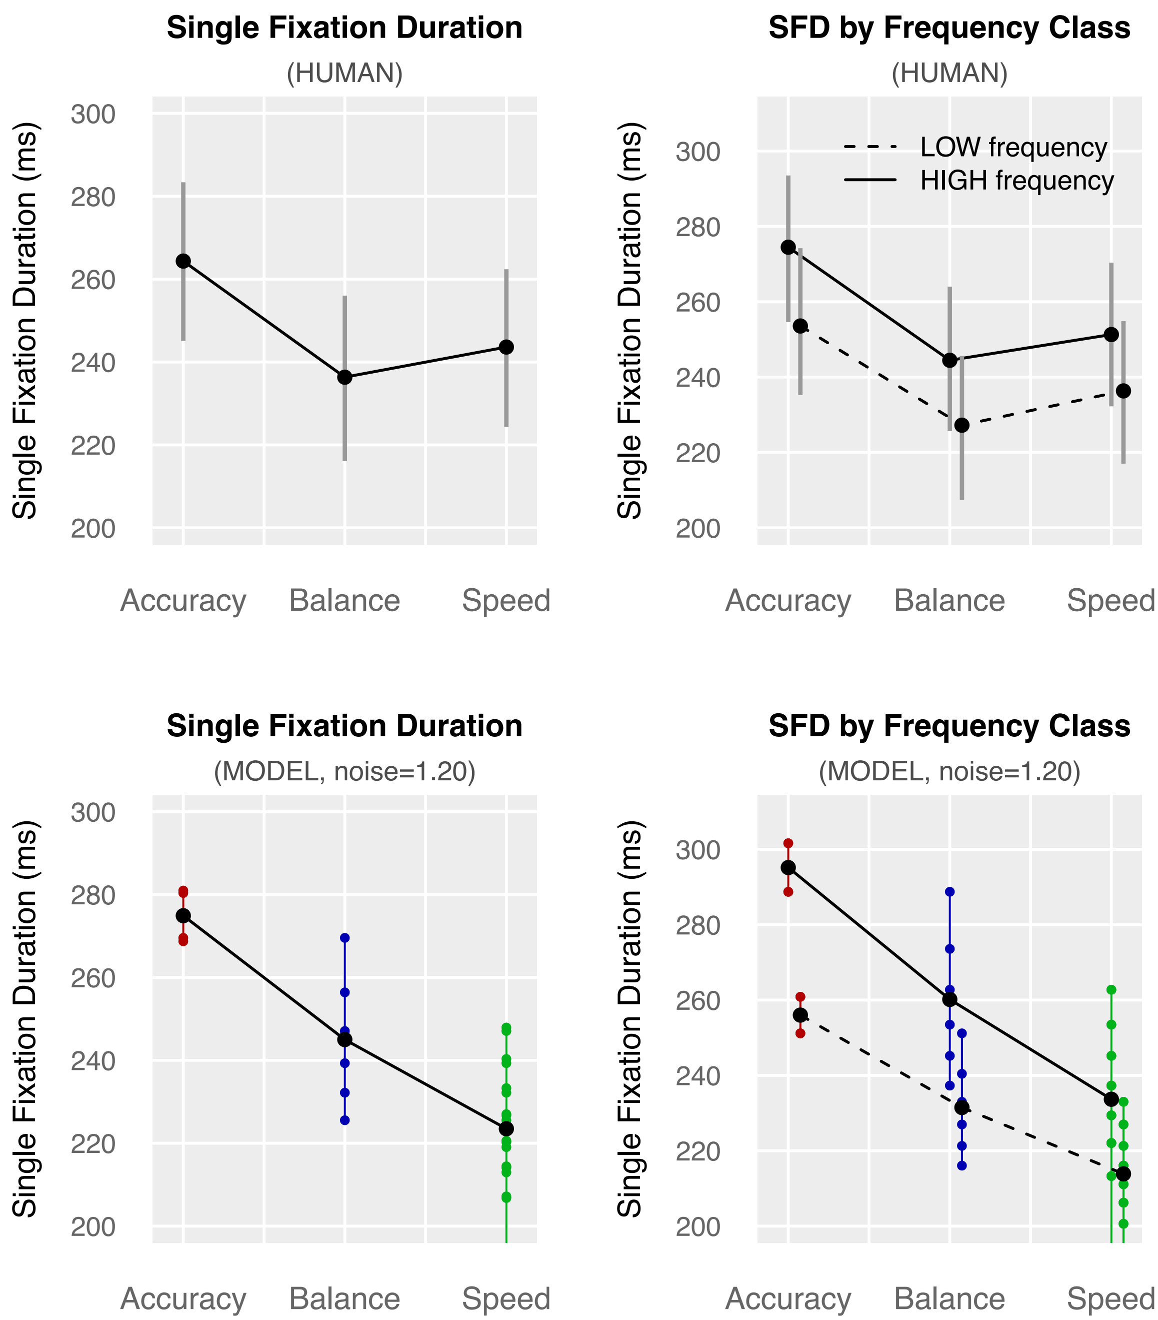

Figure 4.3: The colored points represent predictions corresponding to the best-performing policies identified in Figure 4.1, the lines connect the means of this set. The frequency effect is shown here here as means of low and high frequency bins (median-split) but all statistical models used continuous predictors. The error bars on the human data correspond to one standard error estimated from posterior densities of the mixed effects models.

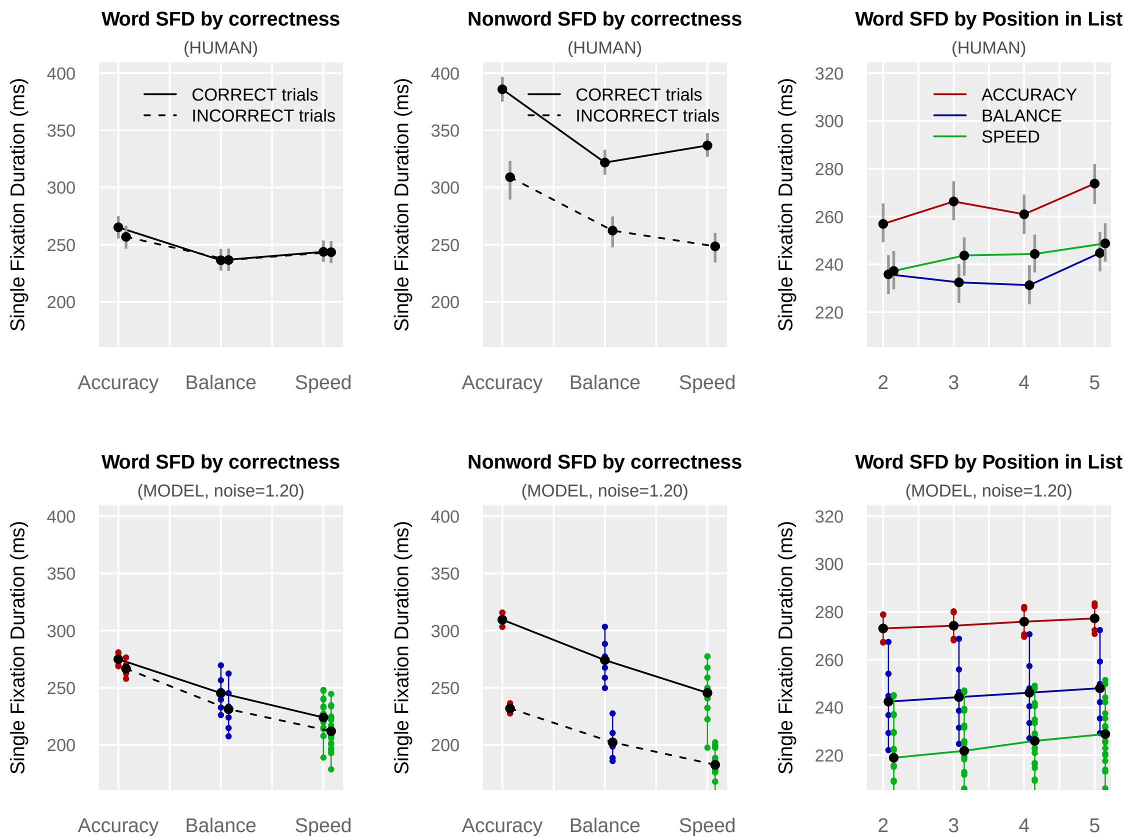

Figure 4.4: Single fixation durations for word and nonword, by correctness (left two columns), and single fixation durations for words by position in the list (rightmost column), for the full set of human participants and the computational model. The colored points represent predictions corresponding to the best-performing policies identified in Figure 4.1, the lines connect the means of this set. The error bars on the human data correspond to one standard error estimated from posterior densities of the mixed effects models.

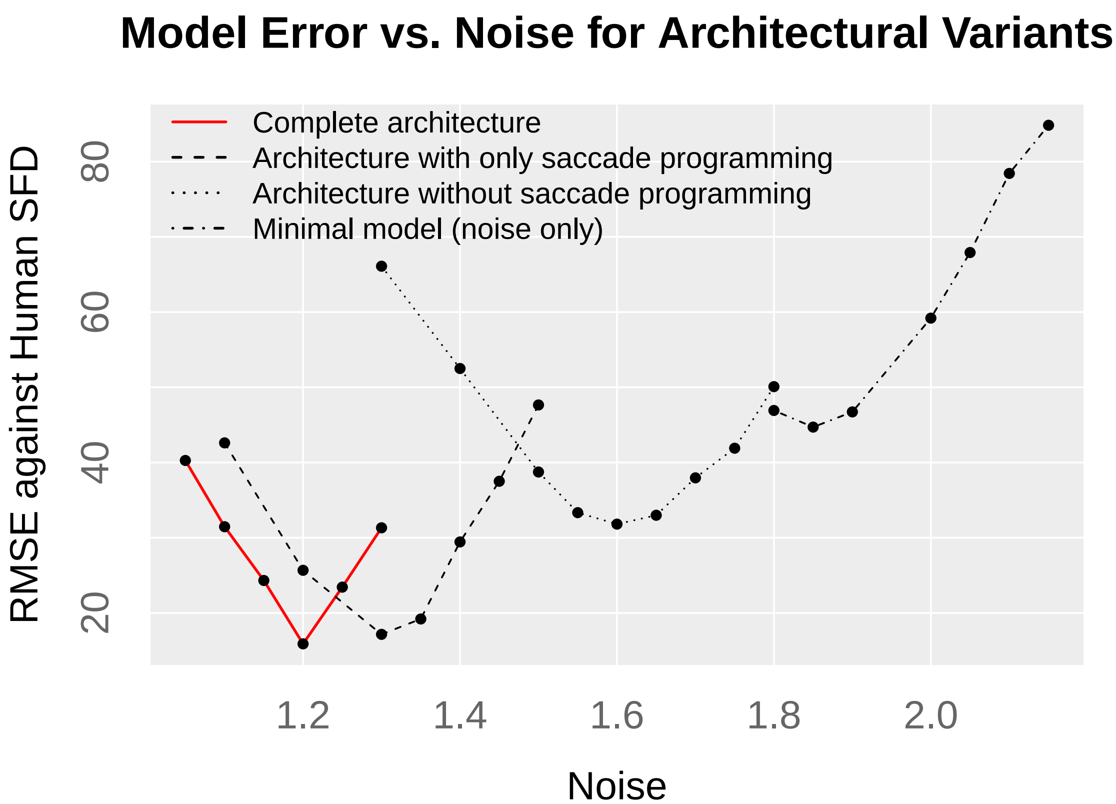

Figure 4.5: Root mean squared error of model predictions (against mean single fixation duration for the three payoff conditions) for four architectural variants. For each architecture, new optimal control policies are derived. In red is the complete architecture explored above and includes saccade programming, eye-brain-lag, and saccade execution. The minimal model dispenses with these oculomotor constraints. The other two models explore the effect of including excluding the saccade programming.

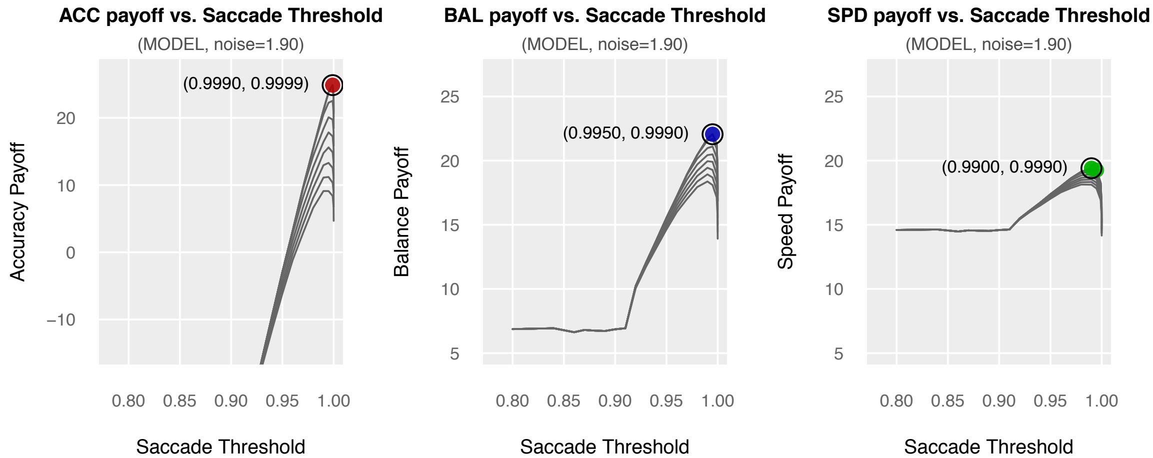

Figure 4.6: The payoff surface and optimal policies for the three payofs for the *minimal model* that results from eliminating the oculomotor constraints; compare to Figure 4.1.

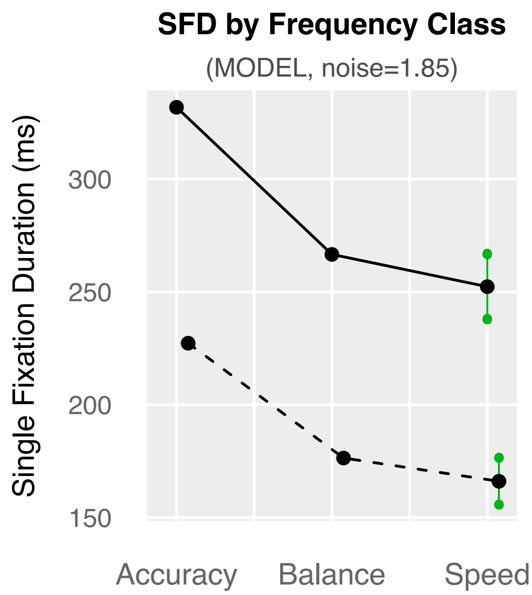

Figure 4.7: Single fixation durations by payoff and frequency class for the minimal model. Dashed line is high frequency, solid line is low frequency (binned by median split). The predicted magnitude of both payoff effects and frequency effects increase, yielding poorer fit to human data (compare to the bottom right panel on Figure 4.3). See Figure 4.5 for source of the noise parameter.

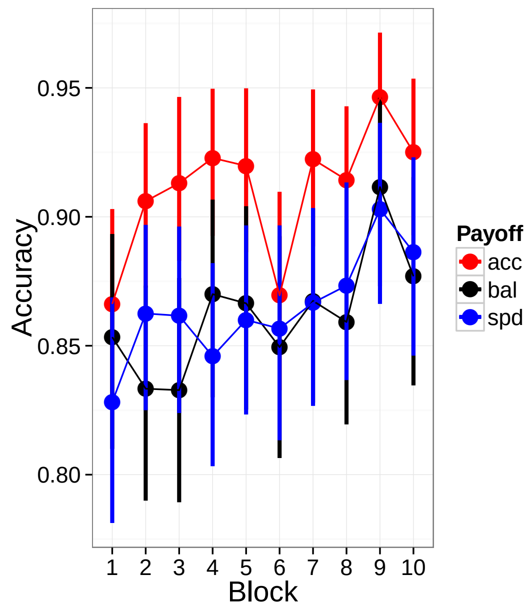

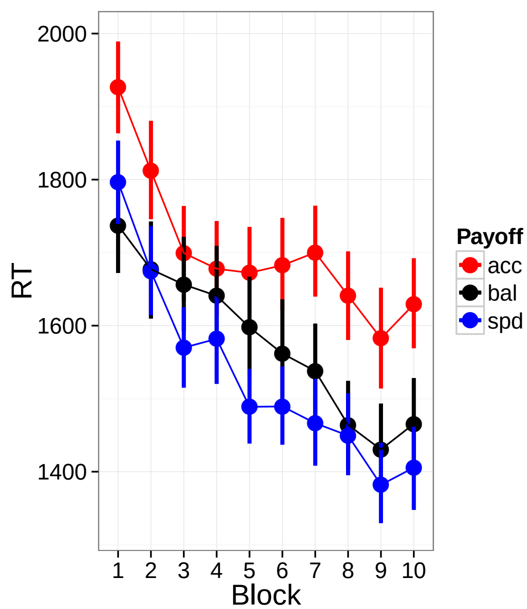

Figure 4.8: Percentage correct (top) and response time (bottom) by block, by payoff condition. Both figures indicate that adaptation happens early, by the first block. Most of the learning that happens over time yields overall improvements in both speed and accuracy, independent of payoff

CHAPTER 5

Spillover frequency effects in the masked LLDT

5.1 Introduction

This chapter adds to the body of work on spillover frequency effects in reading by showing that they persist in the LLDT in spite of its wide word spacing and even when parafoveal preview is masked. As such, it provides evidence for spillover as a cognitive rather than perceptual phenomenon and puts it in the explanatory domain of models such as the one advanced in the thesis. In addition, this chapter embarks on a modeling exercise in E-Z Reader, the model of eye movement control on which the architecture of the thesis model is based. It shows that there are in fact two different mechanisms for spillover effects possible under the E-Z Reader assumptions: one relying on parafoveal preview and another using a post-perceptual delay (though only some reported parameter fits use both). This suggests a key revision to the model that guides remainder of modeling work in the dissertation.

5.1.1 Spillover frequency effects

As discussed in Chapter 2, spillover frequency effects are defined as the speedup of the reading time of a word as a function of the frequency of the preceding word. One explanation for these effects is that they are a result of parafoveal preview in a resource-constrained model. Such a model must make the tradeoff between foveal and parafoveal processing. Easier foveal processing (e.g. on higher frequency words) makes more resources available for parafoveal processing, speeding up the eventual recognition of the next word.

This explanation is supported by evidence that spillover effects disappear when parafoveal preview is incorrect (Inhoff & Rayner, 1986; Henderson & Ferreira, 1990).

This evidence comes from experiments using gaze-contingent boundary paradigms (McConkie & Rayner, 1975): a target word is replaced with a non-useful preview (typically a set of consonants), and then changed to the correct preview when the eye crosses an invisible boundary placed before the target word. The measurement is of the fixation time on the target word as a function of the frequency of the pre-target word.

However, the disappearance of spillover effects with incorrect preview is not always attested. In particular, Kennison & Clifton (1995) show only a reduction but not a disappearance of the effect, and Schroyens et al. (1999) show little effect of masking on spillover. Most interestingly, White et al. (2005) show no disappearance of spillover effect at all in participants who were aware that something unusual was going on in how the words were displayed, but an effect about twice the size as the ones shown in the other paper in unaware participants.

One explanation for this contrast, in light of the findings by White et al., is that participants who are not aware that the display change occasionally provides inaccurate preview attempt to take advantage of preview, whereas participants aware of the preview's occasional uselessness adapt to this constraint and find another way to extend processing one word into the next. This explanation also provides a new way of making sense of a finding by Morrison & Rayner (1981), who showed that saccade distances depend on characters skipped independent of the visual angle traversed. If participants target saccades the same character distance forward regardless of where in their acuity function the target lies, they might be able to vary the amount of processing they apportion to different parts of their perceptual span and therefore choose to use or ignore parafoveal information.

5.1.2 Spillover in the unmasked LLDT

The LLDT was explicitly designed to minimize parafoveal preview and other word-to-word dependencies. Under the explanation for preview provided above, one might therefore expect not to see spillover effects in the LLDT. However, this is not the case: figure 5.1 shows single fixation durations by condition, split into bins based on frequency of the current and previous word, with both foveal and spillover frequency effects readily apparent. Statistical analysis (using the methods described in section 4.2.3) bears this out (foveal $\beta = 4.6633, \chi^2 = 219.85, p < 0.0001$, spillover $\beta = 3.9639, \chi^2 = 5998.9, p < 0.0001$).

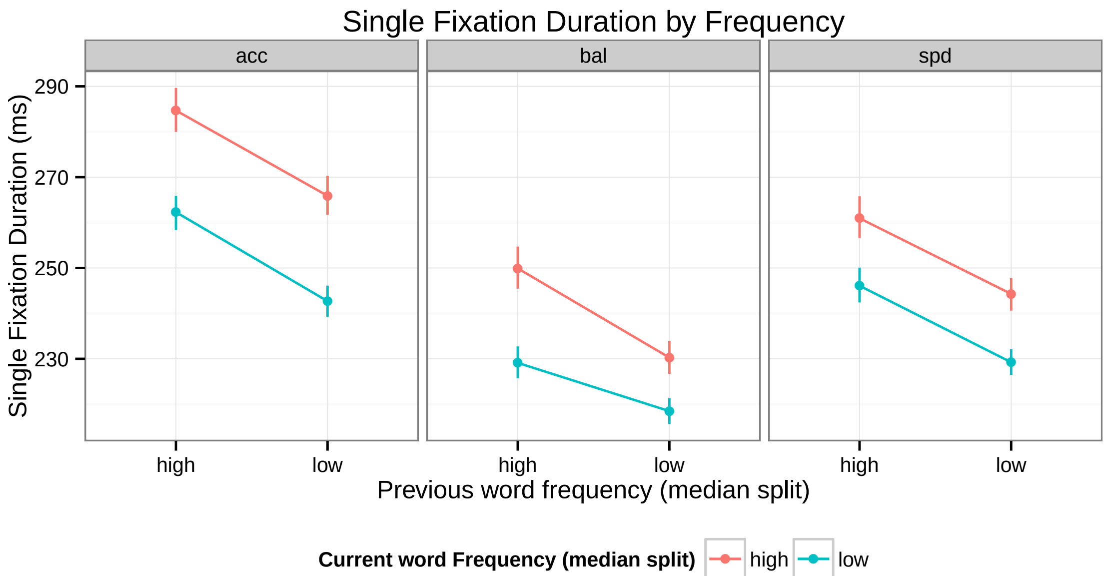

Figure 5.1: Plot shows single fixation durations by condition in the unmasked LLDT, by foveal and spillover (previous) frequency bin. Both foveal frequency (indicated by separation between blue and red lines) and parafoveal frequency (indicated by line slopes) have a significant effect on single fixation durations.

This evidence is consistent with one of two possibilities: the first is that this is a parafoveal preview effect and that participants are able to preview more distant words than expected. In the LLDT experiment in the previous chapter, participants' heads were not fixed, so they could have been sitting farther from the screen than 25 inches and therefore fitting adjacent words into their reading span in spite of the wide spacing. Even if their screen distance was as intended, perhaps the reading span estimates in the literature are variable or incorrect, again allowing for unintended preview.

To mitigate against both of those possibilities, another LLDT experiment was run, this time in a gaze-contingent moving window paradigm. Two different implementations of the gaze-contingent paradigm were implemented, one a novel predictive unmasking design. The experiment is described in the next section, with both a comparison of the two masking variants and an investigation of spillover.

5.2 Spillover in the masked LLDT

5.2.1 Methods

Participants 46 members of the University of Michigan community participated in the predictive unmasking experiment. Data from 14 was unusable due to calibration problems, failure to complete the experiment, or equipment malfunctions, leaving a total of 32. The data loss disproportionately affected the Speed payoff condition, with only 8 participants retained in that condition (compared to 11 and 13 in the other conditions, and 15-16 per condition in the experiment in chapter 4). Participants were paid a $10 baseline for participation, plus $1 per 1000 points as in the previous experiment. 51 participants participated in the boundary paradigm experiment, all in the Balanced payoff condition. Data from 18 was unusable due to calibration problems, failure to complete the experiment, or equipment malfunctions, leaving a total N of 33. Participants in one of the masking variants (the boundary paradigm) participated for course credit.

Stimuli The stimulus generation mechanism was exactly the same as in chapter 4: participants responded to 10 blocks of 20 trials of the masked LLDT, with half of the trials in each block containing all-words lists, and the other half containing a single nonword per list. Words were all 4 characters long and drawn from a bimodal distribution containing high and low frequency words. Nonwords were all pronouncable. The mask was four hashmark characters (####).

Procedure Items were presented on an LCD monitor in a 20pt Courier font, separated by 6 characters of whitespace. This resulted in each word covering 0.7 inches or 1.6 degrees of visual angle, and whitespace covering 1.48 inches or 3.4 degrees of visual angle at a distance of 25 inches from the screen. Each trial started with a drift correction, and then a fixation dot at the location of the first string. The entire six-string list would appear once subjects fixated (with a mask on words 2-6), and the trial ended after subjects responded using the keyboard. Eye movements were measured using an SR-Research Eyelink 1000 table-mounted eye-tracker operating at 1000Hz (preditive unmasking) or a Eyelink-II head-mounted eye-tracker operating at 250Hz (boundary paradigm). Each participant was given an explanation of the task, as well as a quantitative description of the payoff requirements (i.e. "You will receive a point for each 150 milliseconds by which your response is faster than 5000 ms (5 s). You will lose 50 points if your response is incorrect."). Participants were told that words other than the one fixated will be masked.

Statistical methods Data analysis was carried out using mixed effects regression (Pinheiro & Bates, 2000) using the `lme4` package (Bates et al., 2014) for the R environment for statistical computing (R Core Team, 2014). Barr, Levy, Scheepers & Tily (2013) recommend using random intercepts for all grouping factors and additionally random slopes for grouping factors for which the manipulation of interest is within-factor. In the LLDT, the payoff and masking manipulation is between-subject and the frequency manipulations are between-word, so in those cases intercepts are sufficient. The payoff and masking manipulation are within-word so random slopes might be indicated for those tests (consistent with the assumption that some words may reflect the payoff manipulation more than others). Additionally, previous-string frequency is a within-word factor, so a random slope there is recommended as well. Such 'maximal' (in Barr and colleagues' terminology) fits were attempted, backing off to a model with uncorrelated slopes and intercepts first and then to an intercepts-only model as needed. Fixed effect covariates of trial type and trial number were additionally used in trial-level analyses. In string-level analysis covariates of position and string type (word or nonword) were additionally used. The first and last strings in each trial were excluded from analysis, as were all nonwords and all strings that appeared after a nonword. Though the plots show binned means, centered and log-transformed frequency were used as predictors in the statistical models. Statistical significance was tested by likelihood ratio test between the model with the key contrast of interest and one otherwise-identical one without it.

5.2.2 A new masking method

Typically, gaze-contingent masking designs are implemented using either *moving window* (McConkie & Rayner, 1975) or *boundary* (Rayner, 1975) paradigms. In the former, the display is redrawn after each saccade to maintain a consistent viewing window around the fixation point, with the surrounding information masked. In the latter, the display is only redrawn if an invisible boundary (usually between words) is crossed, with the information on the fixation side of the boundary unmasked and the remainder masked. The challenge with such methods is to redraw the display quickly enough for a seamless reading experience, considering that saccades last only about 30-40ms. In the moving window paradigm, the redraw can only commence when the eye lands, so the eye will keep fixating on the mask for the duration of the display change, which is reported to be as short as 8-17ms (Slattery, Angele & Rayner, 2011) on what must be highly specialized hardware, but can be as long as 30-40ms in the hardware used for the dissertation experiments. The boundary paradigm can subtract about 20ms from this amount (assuming a 40ms saccade and a boundary placed at the midpoint between strings) but creates a new problem: longer-distance saccades that cross multiple boundaries will create noticeable flicker as multiple redraws are queued in the system and just-unmasked strings are re-masked.

In spite of far faster overall processors, noticeable flicker appears to be a bigger problem on modern commodity hardware than it was on those earlier specialized systems. To create a masked variant of the LLDT maximally similar to unmasked reading, such delays or flickers are undesirable. Consequently, a novel masking method was developed that is faster than the boundary paradigm and does not create flicker for long-distance saccades[1]. The masking was performed by detecting saccades (based on a velocity and accelleration cutoff), at which point the word to the immediate right of the current word was unmasked. After the command to unmask the rightward word is sent to the display, the approximate midpoint of the saccade is detected by waiting for a deceleration of the eye, and used to predict the intended landing position. If the intended landing position is not the next word to the right, the word that was just unmasked is masked again, and the predicted landing position unmasked. This complicated scheme was necessary to keep redraw delays short: short saccades tended to conclude before their deceleration could be detected, whereas long ones could be accurately extrapolated. The double-redraw was less noticeable than in the boundary paradigm because in this case the re-masked word was farther in the parafovea, whereas re-masked words in long distance saccades in the boundary paradigm are always adjacent. The subjective impression of the swapping was overall seamless, except on rare occasion when the saccade extrapolation is incorrect and the wrong word is unmasked.

The mask always consisted of four hashmark characters (####) which replaced each string except for the one being fixated. Figure 5.2 shows single fixation durations in the unmasked experiment reported in the previous chapter, as well as the boundary paradigm (the theoretically faster of the two types of masking in prior work) and the new predictive unmasking method. Only the balanced payoff was used for this comparison because it was the only payoff used in all three masking types. Total N was 15 for unmasked, 11 for predictive unmasking, and 33 for the boundary paradigm.

[1] Many thanks to Craig Sanders for experimenting with different methods of gaze-contingent masking and for the eventual implementation of this method.

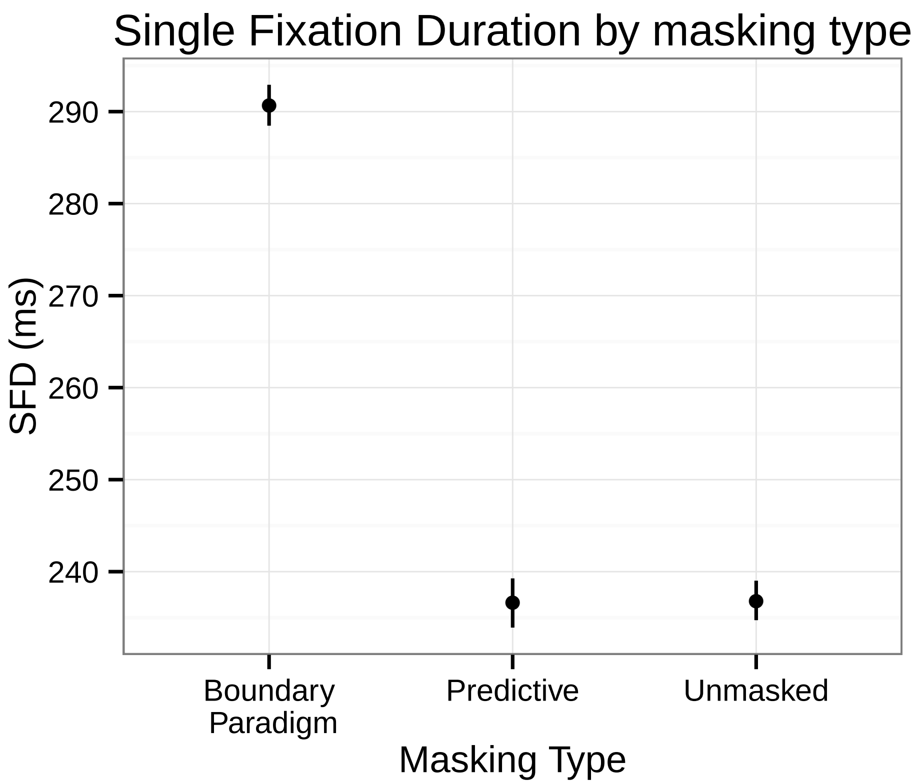

Figure 5.2: The predictive masking algorithm yields fixation differences similar to those in unmasked reading, unlike an otherwise-equivalent conventional boundary paradigm design.

The comparison is not entirely well-matched: the three datasets were gathered on three different semesters, the boundary paradigm participants participated for credit rather than for pay, and the predictive unmasking manipulation took advantage of a new 1000Hz eyetracker. However, these mismatches are not sufficient in explaining the differences: moving from 250Hz to 1000Hz should only provide a benefit of a handful of milliseconds for detecting the eye position, and even unpaid participants optimize points when they are provided (e.g. Schumacher, Lauber, Glass, Zurbriggen, Gmeindl, Kieras & Meyer, 1999).

Two sets of mixed effects models were fit to the single fixation durations, differing by how the experiment contrast was coded. In the first pair of models, the experiment was coded as a zero-sum contrast between the boundary and unmasked conditions, with the orthogonal contrast to it used as a covariate. In the second pair, the contrast was coded between the predictive and unmasked contrast, with the orthogonal contrast to this contrast used as a fixed covariate instead. Backoff to uncorrelated slope and intercept of the random effect by word was needed for convergence of the boundary paradigm models. A significant slowdown was found for the boundary paradigm as compared to the unmasked dataset ($\beta = 53.5ms, \chi^2 = 17.835, p < 0.0001$) but not for the predictive unmasking paradigm ($\beta = 2.9ms, \chi^2 = 0.0388, p = 0.8438$). This is consistent with the claim that the predictive unmasking paradigm provides a more natural reading experience (or at least more similar to reading in the unmasked LLDT).

5.2.3 Adaptation to payoff

Only the predictive masking manipulation was performed on all three payoffs. Figure 5.3 shows response time by payoff overall, as well as separated by block. Unfortunately, the payoff manipulation in this masked experiment appears less clean than in the unmasked one: there is in fact no significant effect of the contrast between the speed and accuracy payoffs on RT, though the direction is correct ($\beta = -16.7, \chi^2 = 0.0168, p = 0.8968$). There are two apparent reasons for this. First, the masked experiment is likely underpowered to estimate the payoff condition effect (especially in the speed condition, where it has half as many participants as the unmasked experiment). Second, the masking manipulation may have made adaptation to payoff more difficult overall. The block-level plots support this argument somewhat, with the speed payoff condition speeding up more overall as a function of block than the others. Statistical analysis bears this out, with a significant payoff-trial index interaction: participants speed up overall as a function of trial, but speed up more in the speed payoff ($\beta = -1.41, \chi^2 = 28.3423, p < 0.0001$). The situation in single fixation durations (plotted in figure 5.4) is similar: directionally correct but non-significant effect of payoff ($\beta = -1.3, \chi^2 = 0.011, p = 0.9163$) and a small but highly significant interaction with trial number ($\beta = -0.08474, \chi^2 = 10.555, p = 0.0012$). The lack of a significant payoff adaptation finding here is of concern, but should not be viewed as a full nonreplication of the results in Chapter 4 due to the apparent lack of power.

5.2.4 Frequency and spillover effects

For analysis of frequency and spillover (previous-word frequency) effects, models were fit with uncorrelated intercept and slope of the previous-word frequency effect by word, because the model with correlated slope and intercept failed to converge. To

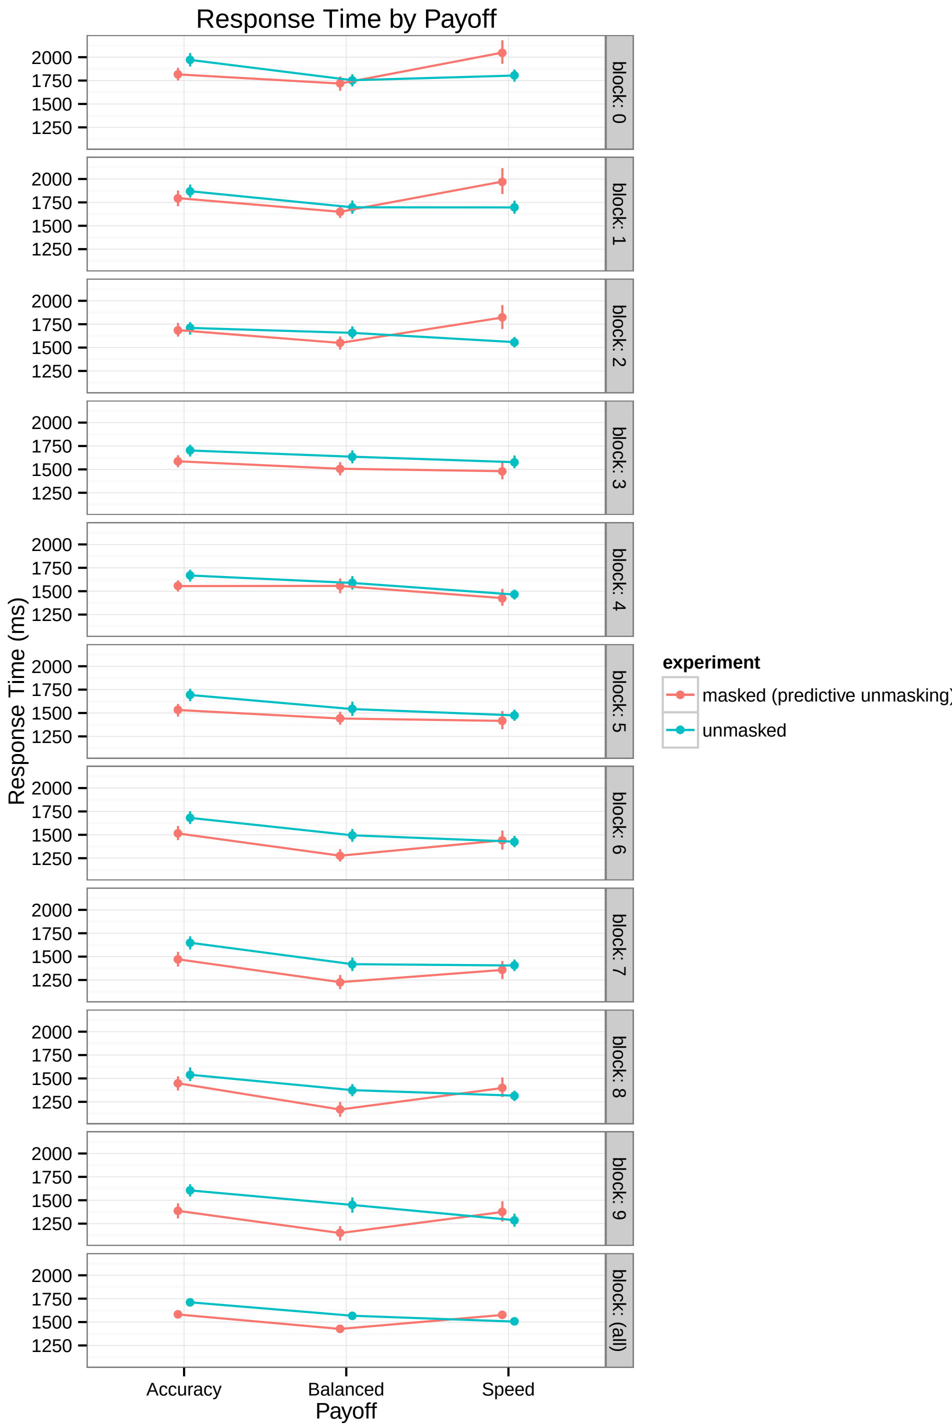

Figure 5.3: RT by condition by block, both masked and unmasked. The effect of payoff on RT in the unmasked experiment disappears in the masked case (bottom panel), though it is present in a sequence of blocks towards the middle of the experiment.

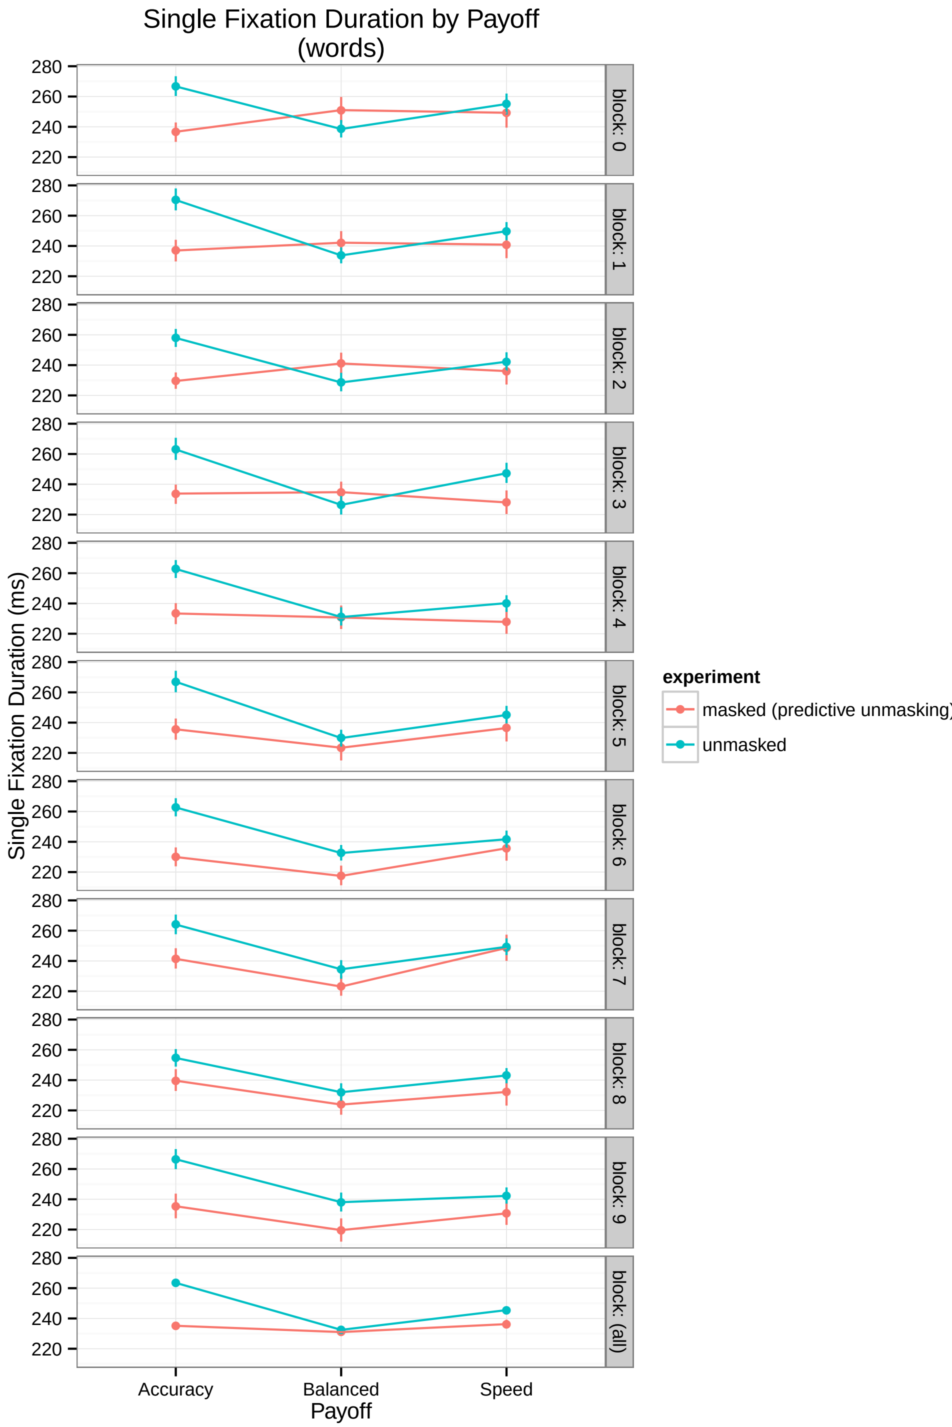

Figure 5.4: Single fixation duration by condition by block, both masked and unmasked. The effect of payoff on RT in the unmasked experiment, disappears entirely in the masked case (bottom panel), though it is present towards the end of the experiment.

test the effect of foveal frequency, a current-word frequency predictor was added to a model already containing previous-word frequency. Likewise to test spillover (previous-word) frequency, a previous-word frequency predictor was added to a model already including a current frequency effect. Both effects were highly significant in all three datasets (summarized in table 5.1). If anything, the spillover effect appears larger in the predictive unmasking manipulation and the frequency effect appears smaller. Both of these differences are reflected in significant interactions (frequency by experiment interaction $\beta = -0.96, \chi^2 = 5.7351, p = 0.017$; spillover by experiment interaction $\beta = 1.573, \chi^2 = 15.418, p < 0.0001$). These two interactions suggest that perhaps not exactly the same process is going on in both masked and unmasked experiments, and chapter 6 explores this possibility in greater depth. But certainly it is the case that spillover effects persist even in the absence of parafoveal preview in the LLDT.

	Frequency effects			Spillover effects		
Dataset	β	χ^2	p	β	χ^2	p
Unmasked	-4.66	179.84	$<$0.0001	-3.89	99.008	$<$0.0001
Masked (boundary paradigm)	-4.63	113.64	$<$0.0001	-4.88	86.384	$<$0.0001
Masked (predictive unmasking)	-3.65	193.35	$<$0.0001	-5.46	68.102	$<$0.0001

Table 5.1: Both current-word (foveal) and previous-word (spillover) frequency effects are highly significant regardless of masking manipulation.

5.3 Exploring other explanations for spillover effects

The experiments discussed above are problematic for explanations that rely solely on parafoveal preview to explain spillover effects. One solution is to take the avenue of models like SWIFT (Engbert et al., 2005), which include spillover effects as a fundamental property of the oculomotor architecture. In the case of SWIFT, this is because saccades are autonomously generated and can only be delayed by difficult cognitive processing. This delay is itself delayed, so that occasionally the 'wrong' (next) saccade is delayed instead of the current one, producing a spillover effect.

This solution is not the simplest move, however. To use the explanation presented and previously motivated by SWIFT would require fairly dramatic changes to the oculomotor architecture that was quite successful in recovering adaptive behavior in the previous chapter. To diverge from the SWIFT attention assumption while retaining its notion of time-delayed foveal inhibition has the net effect is 'building in' the effect rather than understanding it as an adaptive process. This is not to say that all effects must be understood as fundamentally emergent – some properties of the output behavior might well be fixed by the architecture or task at hand. But especially given the empirical fact that spillover effect sizes change as a function of parafoveal masking, it is important to understand if the spillover effect could be explained as a consequence of rational saccadic control.

Rather than build the effect in, the remainder of the chapter is devoted to exploring alternative explanations possible within a serial attention, direct-oculomotor control model like the one used in the thesis. A partial implementation of E-Z Reader is used instead of the adaptive model for this purpose. This has a few benefits: first, as a highly successful descriptive model of the oculomotor control process in reading, the test of E-Z Reader against a new dataset is of interest for its own sake. Second, by relaxing the bounded optimality assumption (which E-Z Reader does not make) one achieves additional flexibility in modeling which, while undesirable in general, is useful for understanding the type of explanations possible within the constraints of a direct-oculomotor control, serial-attention model. Finally, using E-Z Reader's analytic form for mean fixation duration provides for substantially faster simulation, again facilitating easier exploration.

5.3.1 An overview of E-Z Reader

E-Z Reader (Reichle et al., 2009) is one of the major broad-coverage models of eye movement control in reading. It makes a number of key assumptions. First, that attention is allocated serially, such that only one word can be lexically processed at a given time. Second, that oculomotor control decisions are directly triggered by the lexical processing system. And third, that oculomotor and lexical processing can proceed in parallel (i.e. that saccade programming and lexical processing can run concurrently). In these ways, it is a direct architectural progenitor of the model in the dissertation. figure 5.5 shows the sequence of stages in a typical trial of E-Z Reader (reprinted from Reichle et al., 2009). L_1 is the initial stage of lexical processing, at the end of which the process splits into two streams: a motor stream in which a saccade is programmed to the next word, and the continuation of the lexical processing stream (L_2). This is the sense in which the eye-control in E-Z Reader is direct: eye movement control is directly initiated by a completion of a stage in the lexical processor.

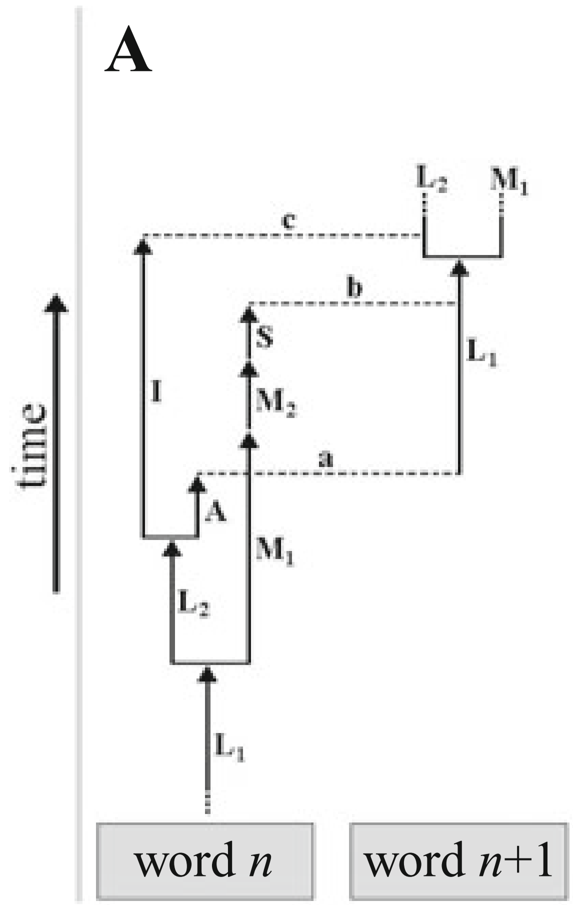

Figure 5.5: A timecourse diagram for E-Z Reader. Note that the first stage of lexical processing L_1 on word $n+1$ cannot begin until the second stage of lexical processing L_2 and the attention shift A is complete on word n. However, the motor processing stream, consisting of a labile phase M_1, non-labile phase M_2 of saccade programming and actual saccade execution S, can proceed in parallel with lexical processing. Postlexical integration I stands in for higher-level language processing.

Saccade programming has two phases: a *labile* (i.e. cancelable) M_1 phase and non-labile M_2 phase. It is followed by the saccade itself (S) and an additional delay imposed by pre-attentive visual processing (the eye brain lag, V, not displayed in figure 5.5), after which the lexical processor begins receiving information from the foveated word. In parallel with the eye stream, lexical processing (L_2) completes, and is followed by a postlexical integration phase as well as an attention shift (A) to the following word. In this sense the model allows solely for serial attention allocation –

the model can only process one word at a time, so if the attention shift occurs before the eyes have foveated on a word, that word begins to be previewed parafoveally.

In a given model run, all of the above stages have durations drawn from Gamma distributions, with fixed means in the case of all parameters except the lexical processing stages. Mean durations for the delays are motivated by prior empirical work (though not directly set to prior work estimates) or fit to data. The data used to fit the model are means of six different eye movement measures computed for five word frequency classes from the 48-sentence Schilling corpus (Schilling, Rayner & Chumbley, 1998). This method of fitting a large number of free parameter to a small number of correlated empirical measures generated from a small dataset has been criticized for potentially overfitting the target training set (Feng, 2003).

The mean duration of the lexical processing time L_1 is calculated as a function of frequency and predictability, using the following formula[2]:

$$t(L_1) = \alpha_1 - \alpha_2 \ln(f_n) - \alpha_3 p_n \tag{5.1}$$

Here, f_n is the frequency of word n, p_n is the predictability of that word given context computed from a cloze task, and the alpha parameters are all free parameters fit to minimize deviation from the training set. The mean of the distribution of $t(L_2)$ is constrained to be a fixed proportion of $t(L_1)$, though the proportion is a fit free parameter in the model (Δ). The model also takes into account the fact that visual information is of highest fidelity in the fovea by penalizing the L_1 duration as a function of the mean distance from the fixation point to the letters in the processed word according to the following formula:

$$t(L_1) \leftarrow t(L_1) \epsilon^{\sum \frac{|fixation - letter|}{N}} \tag{5.2}$$

Here ϵ is a free parameter. The way that the adjustment is computed does take into account the fact that L_1 could be going on during multiple fixations (i.e. a word could be previewed and then viewed foveally).

Frequency effects in E-Z Reader The duration of the lexical processing stages in E-Z Reader, L_1 and L_2, are defined to be a function of lexical frequency. Given that in model fits the value of parameter α_2 in equation 5.1 is constrained to be positive,

[2]For conciseness the word-skipping case is omitted, in which $t(L_1)$ is set to zero with probability equal to the word's probability given context.

the duration of these stages will be shorter for higher frequency words. The decision of when to move the eyes is a function of when L_1 completes, so in this sense E-Z Reader provides a descriptive account of the frequency effect.

Spillover effects in E-Z Reader In E-Z Reader, the total time the eyes spend on a word is a function of the time of the two oculomotor programming stages M_1 and M_2, plus the time for the initial lexical processing stage L_1, minus the proportion of this stage that is completed during parafoveal preview from the previous word. Given that attention allocation is assumed to be serial in E-Z Reader, L_1 can only begin parafoveally when lexical processing on the previous word is complete and attention is shifted. Since lexical processing time (the time of both L_1 and L_2) is in part a function of frequency, higher frequency previous words will complete processing quicker, allow for more of L_1 to be completed parafoveally, and shorten the fixation time on the current word. This is a mathematical formalization of the preview-driven spillover effect explanation provided in the beginning on the chapter.

There is an alternate explanation for spillover effects in E-Z Reader: that L_2 extends past the start of L_1 on the next string if an unusually long L_2 is drawn. The reports on E-Z Reader do not discuss this possibility, but the E-Z Reader codebase does not exclude it. This explanation will come into play when L_2 is longer than $V + M_1 + M_2$ (saccade planning and eye-brain lag), or 200ms. Since L_2 is a fixed proportion of L_1, L_1 needs to be longer than $\frac{200}{\Delta}$ for L_2 to cause a spillover delay. The mean duration of a fixation will be $L_1 + M_1 + M_2 = L_1 + 150$; for typical single fixation durations of 250ms on short words, this implies a value of L_1 of about 100ms, and therefore a Δ of about 2. With typical reported fit values for Δ of about 0.5 and upper limits of 1.0 in searches (Reichle et al., 2006, 2009; Reichle, Pollatsek & Rayner, 2012) the delay possibility seems underexplored in the literature, and is of primary interest here.

5.3.2 Modeling the LLDT in E-Z Reader

A simplified version of E-Z Reader 10 was implemented to model the LLDT and explore different possibilities for the source of the spillover effect.

Simplification Details. The simplified codebase implements all of the stages detailed in Figure 5.5 except lexical integration, and penalizes lexical processing as a function of eccentricity in exactly the way the full E-Z Reader model specifies. The main simplification made is to simplify saccade targeting and timing in the model:

all saccades depart from and arrive on the second letter of the string in question. All strings in the LLDT are four letters long, so this approximates the preferred viewing position (Vitu et al., 1990). In addition, all fixations in the simplified implementation are single fixations. There are no refixations or regressions, hence lexical integration is not necessary, so it is also not included. This is justified in that the short, widely separated words in the LLDT are likely to encourage a high proportion of single fixations. Finally, string predictability in the LLDT is very minimal and its effects subtle (i.e. the position effect in Chapter 4), so the predictability of all strings was set to zero for simplicity.

This leaves three free parameters in the model: α_1 (the baseline duration of the time from processing start to saccade planning start), α_2 (the additional effect of log frequency on the time from processing start to saccade planning start), and Δ (the ratio of the duration from processing start to saccade planning start to the duration from saccade planning start to processing end). Masking in the simulation is implemented by setting the preview benefit to zero: when the model completes L_2 on one word, lexical processing simply stops until preattentive visual processing (V) on the following word is complete.

Parameter search settings. Parameter ranges were chosen that both subsumed the ranges previously used in the E-Z Reader literature and contained minima with respect to the fit metric chosen. Specifically, α_1 ranged from 25 to 175, α_2 from 0 to 10, and Δ was initially swept from 0 to 5, discretized evenly into 100K parameter combinations. The Δ ranges was then extended to 10 with another 100K combinations, since fit minima for Δ were not found in the first range. The resultant step size on α_1 was about 3, the step size on α_2 was about 0.2, and on Δ was 0.1.

For all simulations, 10,000 trials were used for each of four different word pairs: a high-frequency word followed by another high-frequency word, a low-frequency word followed by another low frequency word, and then two mixed types with a high frequency word followed by a low frequency word and vice versa. This is consistent with the E-Z Reader practice of fitting the model on means of binned frequencies. The frequencies were set to the means of those bins in the human data (240 per million for high-frequency and 7 per million for low). Reading times were only measured on the second word. The large number of trials reduced standard errors in the fixation duration estimates to under 1ms in most parameter combinations, and below 3.5ms in all of them. Error bars for such small errors are nearly invisible on the scale of the plots provided, so they are not plotted. To provide some indication that the simplified implementation arrives at similar results to those of the full E-Z Reader model, Figure 5.6 shows a plot of the cumulative durations of the various stages in E-Z Reader 10 as a function of word lexical frequency (Reichle et al., 2012), and below it the same plot generated from the simplified implementation.

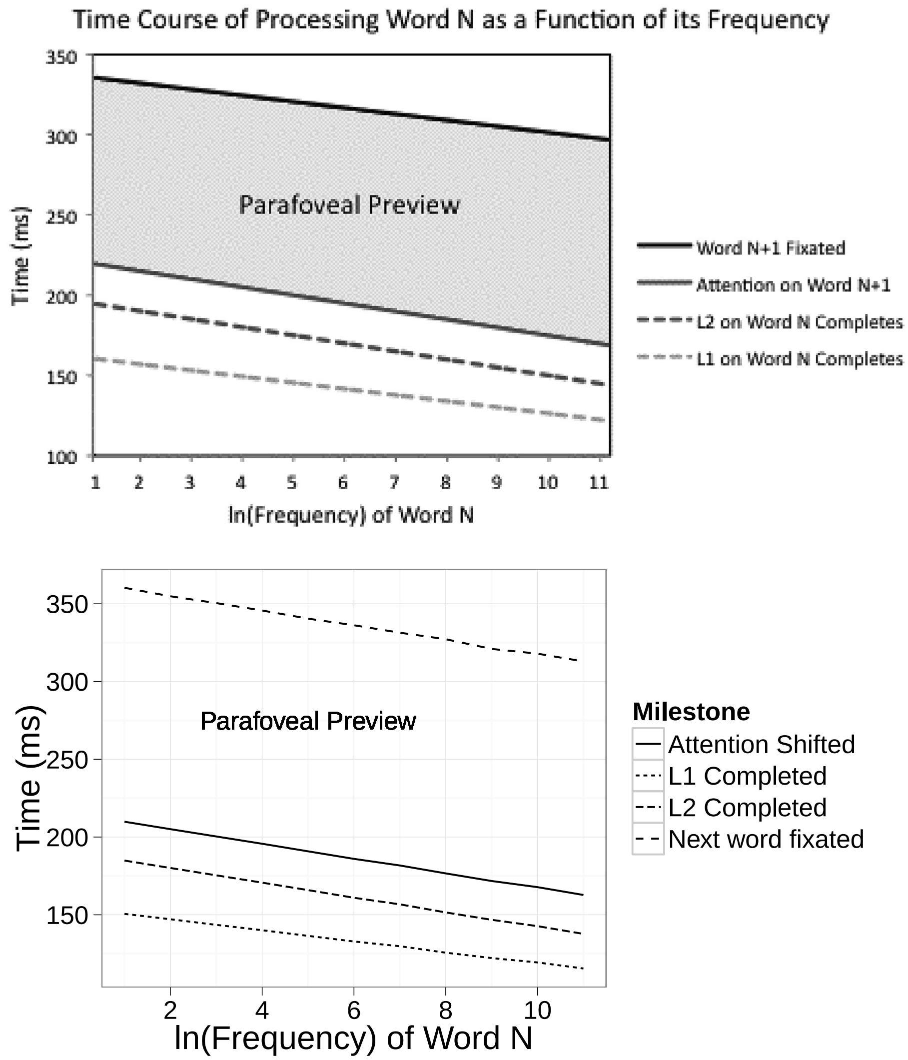

Figure 5.6: Timecourse of stages in full E-Z Reader and the simplified model as a function of frequency, suggesting that the simplified variant recovers key timing properties of E-Z Reader. Top plot reprinted from Reichle et al. 2012; bottom plot generated from the simplified model.

Model training and evaluation. Standard practice in the psychological modeling literature is to only report descriptive training set fits rather than predictive test-set fits. This is of concern because of the potential to *overfit* the training set and capture random noise in it rather than systematic regularity (see Roberts & Pashler, 2000; Pitt, Myung & Zhang, 2002; Pitt & Myung, 2002, for broader discussion of overfitting issues in psychology). This chapter provides test-set fits using two different published sets of parameter values: those from fitting E-Z Reader 10 on the Schilling corpus (Schilling et al., 1998), and those from fitting on a sentence-reading dataset from Rayner & Fischer (1996). The former is selected because the Schilling corpus has been the standard training set for E-Z Reader for many years. The latter is used because it uses high values of α_1, α_2 and Δ (167, 7.5 and 0.97, respectively) and might therefore yield high enough values of L_2 to show delay-spillover rather than preview-spillover. Both use the parameter values provided by Reichle et al. (2012).

For training on the LLDT, the loss function was the sum of the root-mean-squared error of fixation durations in the four frequency classes tested (the interaction between high and low foveal frequency, and high and low previous word frequency). Separate fits were produced based on minimizing error on the unmasked and predictive unmasking datasets, and then tested on the remaining datasets in each case. Fits on the boundary paradigm behave similarly to the predictive unmasking fits, except with slower fixations overall (as one might expect from the overall slower fixations in that dataset), and are not reported.

The test sets are not necessarily a fair comparison: if participants are choosing different strategies in response to the masking manipulation, then it is not reasonable to expect the same parameters of E-Z Reader to reflect both strategies. Therefore, this chapter also discusses individual training fits to the three different LLDT datsets, as a descriptive assessment of what might be going on there. Chapter 6 revisits the question of preview and delay in the adaptive framework of the rest of the thesis.

Predictions from previous fits Figure 5.7 shows predicted fixation durations for the four word classes tested. A few conclusions can be drawn from these predictions: first, E-Z Reader predicts quite different fixation durations in the LLDT depending on the training set use. This is somewhat surprising given that both datasets are similar sentence-reading sets, and speaks to the overfitting concern noted above. It is also puzzling that such different parameter values (e.g. scaling the frequency effect by a factor of nearly 4) are needed to fit those two similar sentence-reading datasets. The fit from the Schilling corpus significantly underestimates both the foveal and spillover frequency effects in the LLDT. Second, E-Z fails to recover the spillover frequency effect in the presence of masking using the Schilling corpus fit. Third, as suspected given the relatively high values of α_1 and Δ, the fit on the Rayner and Fischer corpus does predict spillover effects in the LLDT, though it overestimates fixation durations overall.

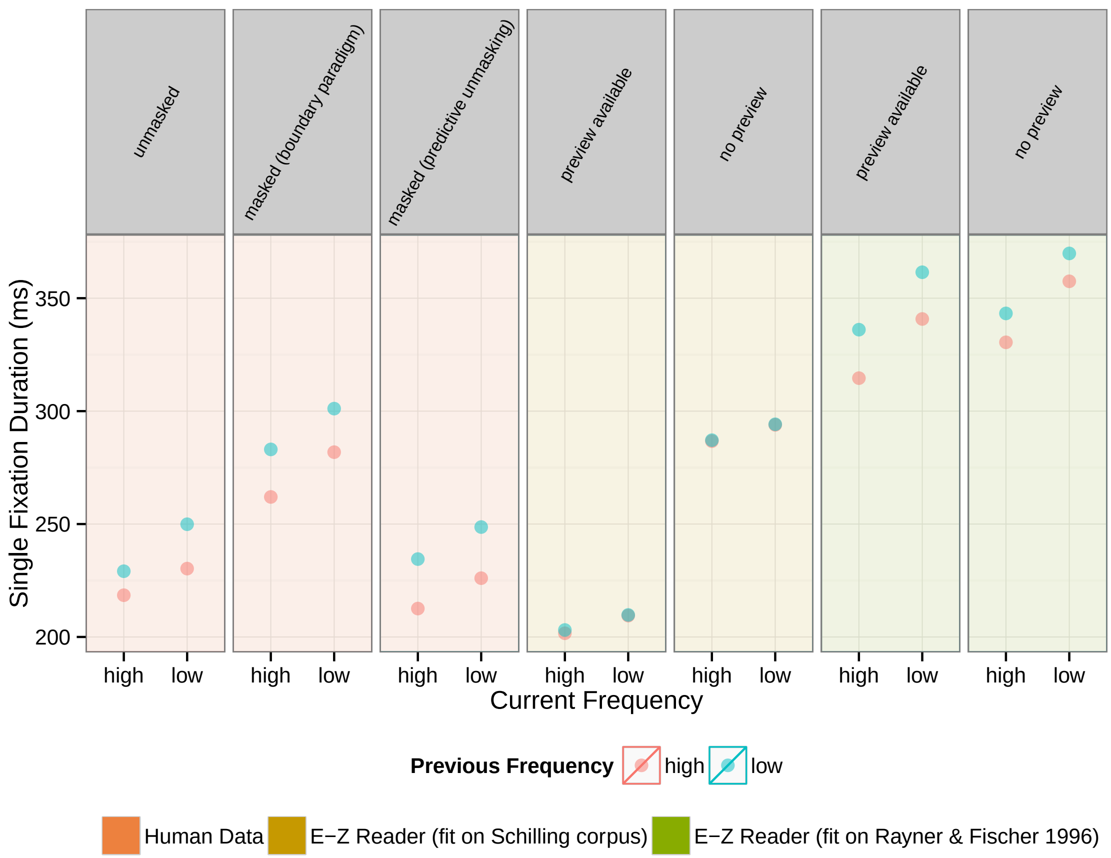

Figure 5.7: E-Z Reader provides poor predictive fits on the LLDT dataset when trained on the Schilling corpus, but somewhat better ones using the fits from the Rayner and Fischer experiment. Schilling corpus fit parameters: $\alpha_1 = 98, \alpha_2 = 2, \Delta = 0.25$. Rayner and Fischer fit parameters: $\alpha_1 = 162, \alpha_2 = 7.5, \Delta = 0.97$

Predictions from fits on the LLDT Figure 5.8 shows predicted fixation durations using the parameters estimated to minimize error on the unmasked LLDT dataset. The training fit is good (left green panel compared to left red panel), recovering both frequency and spillover effect sizes, though overestimating the fixation durations on high frequency words following low frequency words (left blue point). The prediction for the unmasked fit (right green panel compared to the two rightmost red panels) underestimates the spillover frequency effect, consistent with the argument that the spillover effect is primarily driven by preview in that fit.

Figure 5.9 shows the predicted fixation durations with the predictive unmasking data as the training set, with essentially the opposite effect. The training fits are again good (right green panel as compared to right red panel), but the large delay times in the masked condition turn into preview times when preview is available, creating substantial overestimates of the spillover effect (left green panel compared to left red panel).

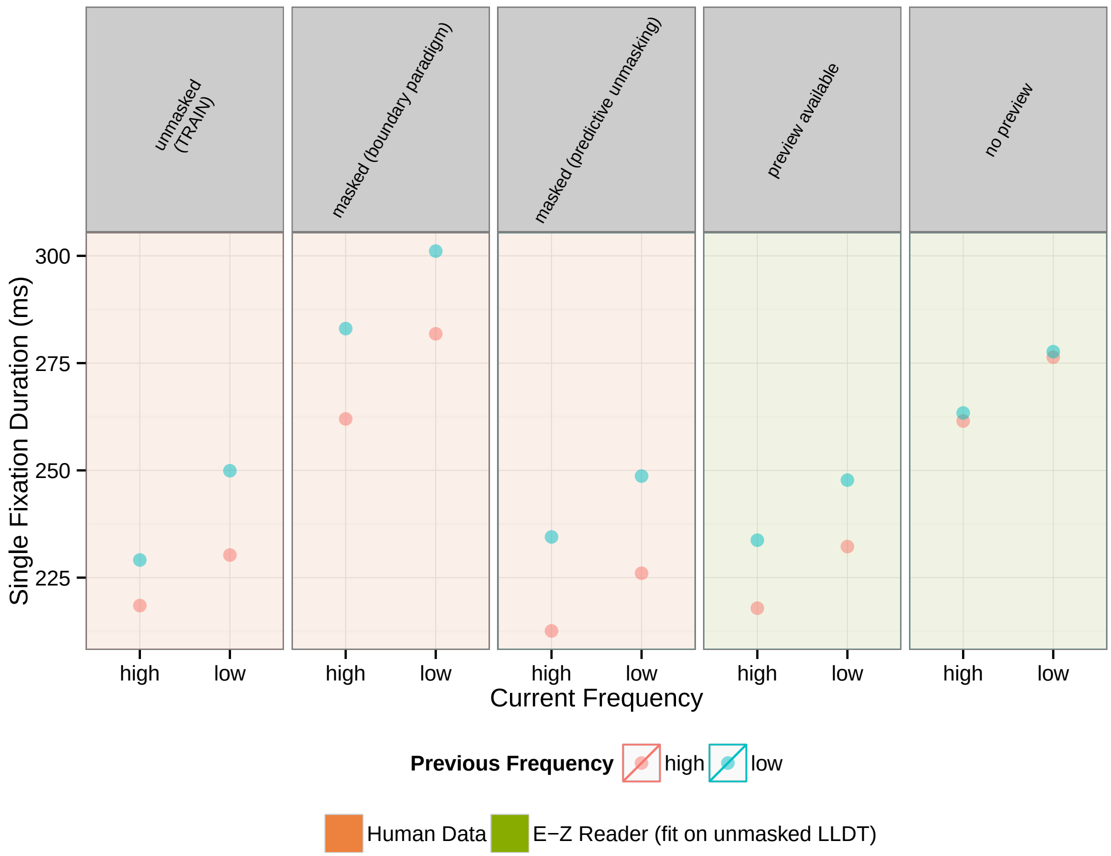

Figure 5.8: E-Z Reader captures the frequency effect in both masked and unmasked data, but greatly underestimates the spillover effect in the masked case when trained on unmasked data. Parameters: $\alpha_1 = 83.7, \alpha_2 = 4.13, \Delta = 1.2$

Understanding parameter values at best fits As noted above, the training/test split between masked and unmasked data is somewhat unfair because something strategically different is likely be going on in the two task variants. Therefore, it may be instructive to look at what best training fits to the three different LLDT variants might look like and what they imply descriptively about what might be

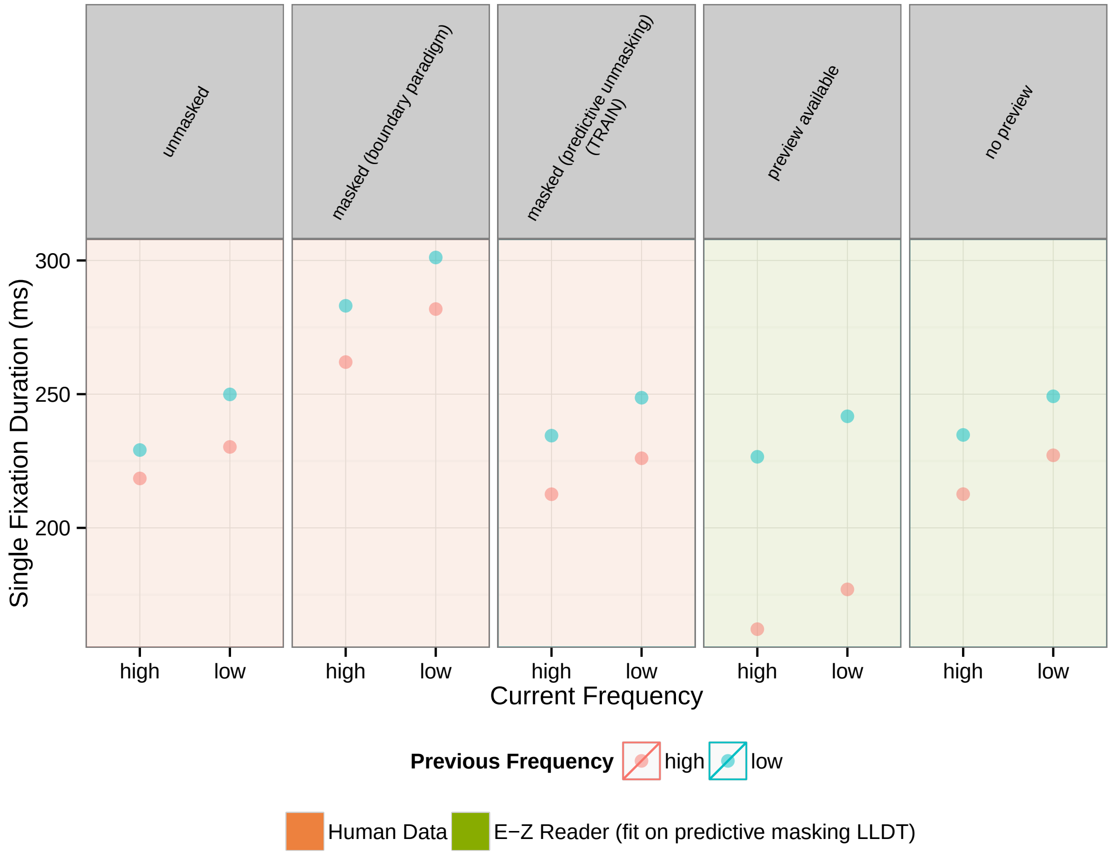

Figure 5.9: E-Z Reader trained on maksed data greatly overestimates the spillover effect in the unmasked case. Parameters: $\alpha_1 = 34.78, \alpha_2 = 4.13, \Delta = 5.33$

going on. Figure 5.10 shows the SFD error at different values of α_2, at the best values of the other parameters. α_2 is the parameter that adds a frequency-based delay to fixations, and as one might expect the minima map nearly directly onto foveal frequency effect sizes in the data. They are nearly in the same location for the unmasked and predictive-unmasking masked sets, and slightly larger for the boundary paradigm.

Of greater interest are the other two parameters. Figure 5.11 shows the SFD error at different values of α_1, at the best value of the other parameters in that range, for the three task variants. α_1 is the baseline duration of the pre-saccade-planning portion of processing (L_1). The error on the unmasked task is not extremely sensitive to α, with a fairly wide range of values yielding relatively low errors. This is a positive result for E-Z Reader, in the sense that its fits do not seem strongly sensitive to this free parameter. The loss minima for the other two datasets, however, are far sharper. Both lie to the right of the minimum for the unmasked dataset, consistent with the argument that to fit the masked dataset E-Z Reader chooses shorter times before saccade planning and (in order to fit mean SFD) longer processing times after. Bolstering this argument further is the loss surface for Δ (Figure 5.12), which shows the minima in the exact opposite order. To fit the short fixation durations and larger spillover effects in the predictive-unmasking dataset, the fit selects very short pre-saccade processing times, and then very long lexical recognition times that extend far into the following word.

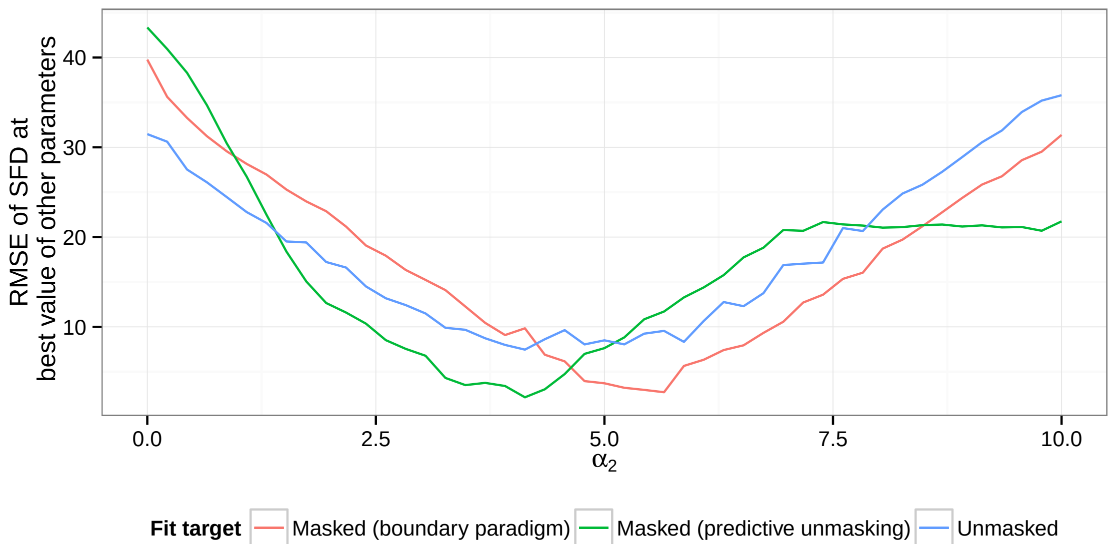

Figure 5.10: E-Z Reader RMSE as a function of α_2. Minima are nearly in the same location independent of dataset fit because this parameter primarily governs the foveal frequency effect.

5.3.3 The implications for spillover-as-delay for models of reading

The previous section illustrated how a serial, direct-oculomotor control model of reading like E-Z Reader can recover frequency and spillover effects in two different parafoveal masking tasks by combining parafoveal preview and delay. The delay processing cannot be perceptual, both because the word recognition span is strongly asymmetric forward, and because the effect persists when previous strings are masked. This is again consistent with spillover being not solely a perceptual phenomenon. In the context of E-Z Reader, this is a fairly natural move due to L_2 already not being

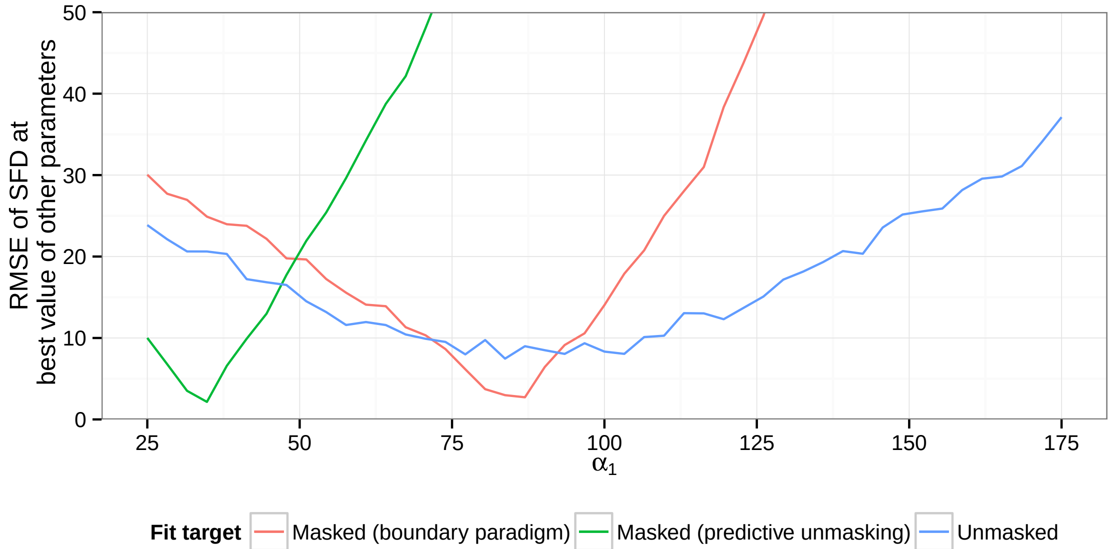

Figure 5.11: E-Z Reader RMSE as a function of α_1. While the minimum for the unmasked data is fairly shallow (suggesting many reasonably-good fits available), those in the masked datasets are far sharper, and at lower values of α_1. These shorter baselines (and the correspondingly longer post-saccade-planning portions of processing seen in the values of Δ in figure 5.12) are consistent with a delay-based explanation for spillover effects.

perceptually bound. It is also consistent with response hysteresis accounts such as those informally discussed in the buttonpress moving-window paradigm (i.e. that it is hard to get out of a buttonpress rhythm) or formally embodied with the time-delayed foveal inhibition mechanism of autonomous-saccade models like SWIFT (Engbert et al., 2005).

In addition, the previous section illustrated a challenge of non-adaptive models of behavior: their free parameters provide no locus of adaptation and their fitting procedures cannot distinguish between differences in strategy and differences in the architecture of the reading agent. The remainder of the thesis will extend the model presented in the second chapter to include both delay and preview. To do so, it will need to provide a rational explanation for why a delay might extend into the processing of the next word.

This is necessary because the move of simply adding a lexical access stage is undesirable: a significant portion of the elegance of the SPRT as a model of lexical access is that it provides a mechanistic explanation for lexical access making contact between the lexicon (via the priors) and perception (via incoming noisy samples).

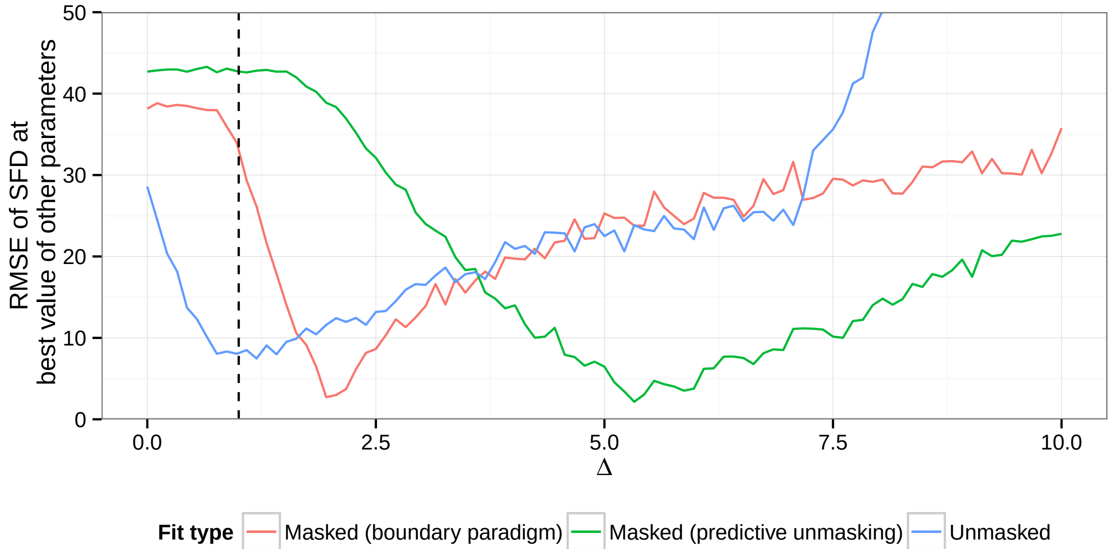

Figure 5.12: E-Z Reader RMSE as a function of Δ. The upper of 1.0 limit used in typical fits is marked with a dashed vertical line. The higher values needed to fit unmasked data are consistent with delay-based rather than preview-based explanations of spillover.

To maintain this elegance, the delay should consist of continued sampling, but with the eye moved on this continued sampling cannot be from perception. A plausible extension is to assume that this sampling continues from memory, but this introduces the question of why a rational reader might choose to sample from memory instead of either spending longer on the word in the first place, or moving on and sampling from perception on the following word.

The next chapter explores this notion, and introduces a model in which it is rational to reprocess past samples after the eye has moved on, and therefore yield spillover effects. The final data chapter adds parafoveal preview into the model, in an attempt to build a model that can adaptively change its behavior to yield predictions on both masked and unmasked paradigms under a single architecture.

CHAPTER 6

Computationally rational spillover

6.1 Introduction and overview

The previous chapter provided evidence that spillover effects persist in the absence of parafoveal preview in the LLDT and reviewed evidence of spillover persisting with incorrect preview in other tasks (Schroyens et al., 1999; Kennison & Clifton, 1995). This evidence is used to make a case for a *delay*-based explanation for spillover effects. This chapter provides such an explanation in the context of the adaptive model presented earlier in the thesis, and shows why delaying perception in favor of processing past inputs may be rational. The specific idea explored is that delay is due to memory-based reprocessing of past inputs, which may be rational as a way to mitigate against noise in the update process itself.

The model in this chapter extends the model described in Chapter 3 by adding a second MSPRT that draws its samples from a simple short-term memory buffer. A few results are provided using this model: first, that frequency spillover emerges in top-performing policies, where performance is evaluated on the same task and payoff given to human participants as in Chapter 4. Second, that a model capable of spillover does no worse than an otherwise identical model that can eliminate spillover by always attending to perception when it can. Third, that the spillover-capable policies in such a model perform no worse than spillover-incapable ones across the speed-accuracy tradeoff curve (in the sense that they achieve a no worse a speed at the same accuracy or no worse an accuracy at the same speed), and in fact perform better in some portions of the noise parameter space. And fourth, that the origin of the effect is in a counter-intuitive but fundamental property of the dynamics of sequenced thresholded samplers.

Building on this memory model, a parafoveal preview model is constructed. This model is capable of navigating a tradeoff between foveal, parafoveal and memory

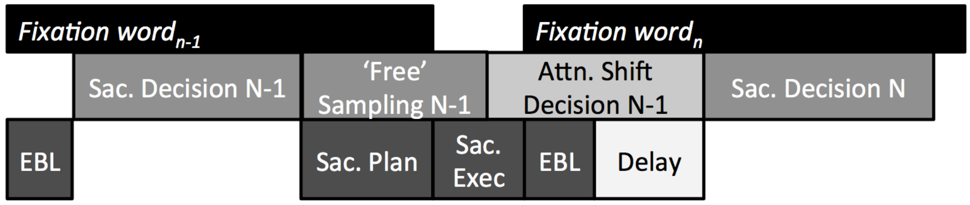

Figure 6.1: Example dynamics of a decision to saccade from word N-1 to word N. The memory-driven attention shift decision can delay the start of perceptual sampling on the next word even though eye-brain lag (EBL) is complete, potentially creating spillover. A detailed description of the dynamics depicted in this figure is in §6.2.

sampling. The preview-capable model performs better (in speed-accuracy tradeoff) at high RTs and equally or only slightly worse at lower RTs. The practical impact on spillover is shown to be minimal, however, in the widely-spaced setting of the LLDT, with spillover effects primarily driven by the memory delay even in the preview-capable model. However, parafoveal preview is shown to increase the range of free parameter values in the model that show rational spillover, due to a property of the memory component of the model.

6.2 A model of saccadic control with noisy memory for recent perception

The key extension to the model in Chapter 3 is a noisy memory that buffers perceptual input. It is most easily understood by first considering the dynamics of a single decision to saccade from one word to the next, as presented in Figure 6.1. As in the earlier model variant, there is an eye-brain lag (EBL) delay from when the eye first fixates on word N to when information from the retina becomes available for perceptual processing. A sequence of noisy perceptual samples then arrive and are integrated via an incremental and noisy Bayesian update of a probability distribution over lexical hypotheses as in the previous model variant. In addition, the perceptual samples are also buffered by storing them in a memory that contains samples from only the most recent word.

As in the previous model variant, saccade planning is initiated when a *saccade threshold* is reached – but the decision variable is different, as will be described below. Then, perceptual sampling continues in parallel with saccade planning until the fixation ends, and then for another EBL amount longer (these are samples received at the retina during the fixation and only now arriving at the lexical processor). This perceptual sampling period is marked as *free sampling* in Figure 6.1 because its length is not under adaptive control.

The model then switches to sampling from its memory, continuing the sampling process until an *attention shift threshold* is reached. If this threshold had already been reached during the earlier perceptual sampling stages, attention shifts instantly. Otherwise attention remains on word N-1 even if the eye has saccaded to word N and the eye-brain lag on word N is completed[1]. Perceptual samples from word N will not be processed until attention is shifted away from the memory-based processing of word N-1. Thus the memory processing on word N-1 may delay processing of perceptual samples from word N and produce delays of the kind explored using E-Z Reader in the previous chapter (though note that these delays must be frequency-sensitive for spillover frequency effects to arise). Perceptual samples arriving during this time are buffered in the memory. In this way the posterior update is a limited computational resource and its relative allocation to perception or memory is determined by the saccade and attention shift thresholds. If the time to reach the attention shift threshold is sensitive to the frequency of word N-1, the model may exhibit a spillover frequency effect.

One more departure made from the previous model is the decision variable. In the model in Chapter 4, the eye movement decision was conditioned on the word or nonword belief in a given position. §3.2.2.2 discusses the possibility of making the decision conditioned on either the specific task at hand or on the task most similar to it in humans' ordinary life. The move in this model is from using the former to the latter, after preliminary

Preliminary simulations showed that with that prior and the substantial amount of 'free' sampling that saccade planning time imposes after a first threshold is crossed, nearly any second threshold would be crossed during this portion of sampling, eliminating all opportunity for delay.

: the model is assumed to have a correct model of the trial generation procedure and set the prior probability of words and nonwords accordingly.

In particular, that the probability that the trial is an all-words trial is 0.5 and nonwords are distributed in nonword trials uniformly. Therefore the prior probability for the word hypothesis at each position is $0.5 + (\frac{5}{6} * 0.5) = \frac{11}{12} = 0.91\overline{6}$.

[1] Because there is a fixed set of memory samples available, the attention shift decision is not guaranteed to converge, unlike the saccade decision. It nearly always converges, but a 30-sample deadline is used to prevent infinite sequences.

The theoretical interpretation of this change is that participants might choose to rely on their highly-honed eye movement strategy tuned to word identification rather than a new strategy for lexical decision. One can attempt to address this empirically: facilitatory effects of lexical neighborhood density on response times tend to appear in lexical decision whereas inhibitory ones appear in word recognition (see Norris, 2006; Andrews, 1997 for reviews). The explanation provided by Norris in the context of the SPRT is that if the hypotheses are individual words (as in perceptual identification), they will compete, leading to inhibition on word identification times. If the hypotheses are the classes of words and nonwords (as in lexical decision), one word will facilitate other words and lead to facilitation in word identification times.

This was investigated in the LLDT using lexical neighborhood norms computed from the MCWord orthographic wordform database (Medler & Binder, 2005). The dataset used was the full set of LLDT experiments (masked and unmasked) with covariates of experiment type and condition in addition to the ones reported in §4.2.3. There was no significant effect of lexical neighborhood density size on single fixation duration for words in either direction ($\beta = -0.12, \chi^2 = 0.9174, p = 0.3382$). With neither facilitatory nor inhibitory effect of lexical neighborhood density on reading times, this analysis does not provide guidance as to which decision variable participants are using for their eye movement decision. Figure 6.2 shows a more detailed view of the model in this chapter, analogous to Figure 3.1 in Chapter 3.

A decoupled noise schema As in the previous variant of the LLDT model, the decisions to plan a saccade, shift attention, and make a motor response are realized as Sequential Probability Ratio Tests. With the change in decision variable, the test used is formally the multihypothesis variant of the MSPRT (Baum & Veeravalli, 1994; Dragalin, Tartakovsky & Veeravalli, 2000). At each timestep, the model performs a Bayes update based on a noisy sample drawn from perception or memory, with the posterior at each timestep becoming the prior for the next timestep. The choice of word representation is unchanged from the previous model and follows Norris (2006) in representing a letter as a unit-basis vector encoding and a word as a concatenation of such vectors.

Samples are generated by adding mean-zero Gaussian noise to these vectors. With the introduction of different samplers, the model also introduces a set of decoupled noise components, as follows: to generate a perceptual sample, mean-zero Gaussian *perception noise* with standard deviation (SD) σ_p is added to each component of the word representation vector (this is also notated as the shorthand `pNoise` in some

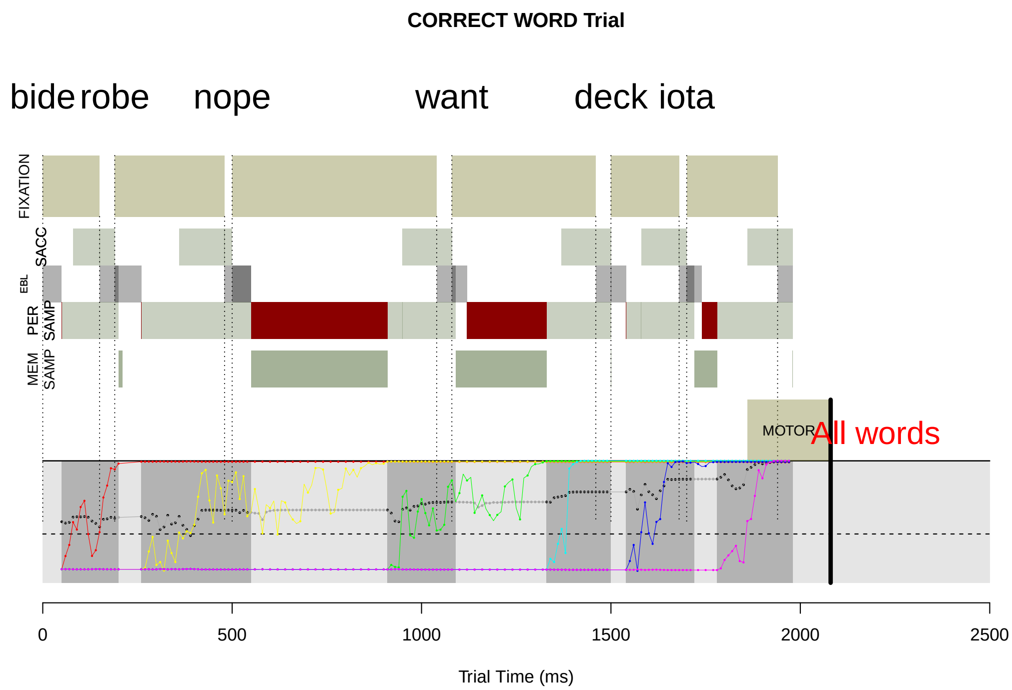

Figure 6.2: A typical all-words trial in the memory-sampling model. The six strings are plotted at the top of the figure, folowed by the fixations, and the oculomotor and lexical processing stages. Perceptual and memory sampling bars are plotted separately, with the pereptual sampling in dark red when perceptual samples are ignored in favor of memory. At the bottom is the max a posteriori belief for each position as a function of time in color and the trial-level belief in gray, with the background color indicating perceptual sampling (dark grey) or memory (lighter gray). The dashed horizontal line is the saccade threshold and the solid line is the attention threshold.

plots). Each perceptual sample is also stored in a memory buffer, and memory samples are generated by uniformly drawing a stored sample from memory (with replacement), and adding an additional mean-zero Gaussian *memory noise* with SD σ_m to each position (notated om plots as `mNoise`). Before each Bayesian update, whether using a sample from perception or memory, mean-zero Gaussian *update noise* with SD σ_u is added to each component of the word representation vector (notated on plots as `uNoise`). In the remainder of the chapter, all references to a noise value refer to one of these standard deviations (so ‘perception noise is 0.5’ means ‘the standard deviation of the mean-zero Gaussian pereption noise is 0.5’).

As before, the Bayes update uses the correct noise in its likelihood computation, i.e. $\sqrt{\sigma_p^2 + \sigma_u^2}$ as the noise in its likelihood computation for a perceptual update, and $\sqrt{\sigma_p^2 + \sigma_u^2 + \sigma_m^2}$ in the likelihood computation for a memory update. All noises are drawn independently. The three SD's, σ_p, σ_m and σ_u, are free parameters in the model, explored in some detail below and then fit to human data in the next chapter.

6.3 A computational rationality analysis

Unlike the model in chapter 4, this model has three distinct noise parameters that can no longer be interpreted as simply scaling the controlled sampling portion of the fixation, because now different permutations of the parameters will yield the same mean SFD in optimal policies. For example, one might expect memory sampling to be dispreferred when memory noise is high. In a setting where memory sampling is minimal, any combination of σ_p and σ_u that sums to the same amount will yield exactly the same behavioral outcome (because perceptual samples always suffer from both noises). An alternative pursued in this chapter is exploring the range of the free noise parameter values, showing the range of predictions with which the model is compatible, and explaining how aspects of the architecture (whether a priori determined or 'free') constrain the class of explanations possible.

Two analyses are undertaken. First, it is shown that top-performing policies (with respect to the same payoff given to human participants) yield frequency spillover consistent with human data, and poor-performing policies do not, as long as there is a relatively high noise in the update process. Second, the model's policy space is extended to allow it to prioritize perception over memory samples when both are available, eliminating spillover in those policies. In the resultant model, the spillover portions of the policy space perform better than non-spillover ones under any imposed speed-accuracy tradeoff in some noise settings, and never perform worse.

On the question of scaling. In setting aside the use of noise as a scaling factor it is useful to reevaluate what appropriate comparisons to human data look like. For example, high-noise models might show numerically large frequency effects that are nontheless far too small against the backdrop of 1500ms fixations. One solution is to use the sample rate, fixed at 10ms in the previous model variant, as another free parameter fit to fixation durations at each noise permutation. This would be at the cost of computational tractability, already a scarce resource in grid-searches of a six-parameter model. Backing off from grid searches to a faster optimizer would eliminate the ability to explore the noise and policy space and understand how behavior changes throughout the parameter space. A third option chosen here is to look at frequency and spillover effects scaled by mean fixation duration. This addresses the scaling issue and allows a comparison to human effect sizes, at the expense of being able to compare raw fixation durations themselves.

6.3.1 Spillover as computationally rational behavior

The first evaluation of the model asks the question of whether spillover effects emerge in approximately optimal policies. Since only balanced payoff data is available for all three masking conditions (unmasked and the two types of masking), the model was evaluated on this payoff. The policy space was discretized as follows: the saccade threshold θ_s ranged between 0.199 and 0.999 in steps of 0.05; the attention shift threshold θ_a between 0.19999 and 0.99999 in steps of 0.05, and also including $\theta_a = 0$ which prevents memory sampling; and the response threshold θ_r between 0.599 and 0.999 in steps of 0.1. All 1530 permutations were explored.

Figure 6.3 shows the distribution of spillover effect sizes in the top 5% of policies (evaluated by task payoff, not fit to human data), for a range of noise parameter settings. At higher noise settings, even the best policies are close to chance performance. The figure shows that top-performing policies show little to no spillover when update noise is low, positive but small spillover effects when update noise is moderate (especially when the other two noises are relatively high), and sizable positive spillover effects when update noise is relatively high and the other noises (especially perception) are relatively low. These results are consistent with spillover as a rational adaptation to belief update noise. Figure 6.4 shows the same measure on the same scale, in the human data. The effect size in the humans is noticeably higher, but otherwise the plots look similar. This includes the fact that the whiskers on the boxplots cross 0, i.e. that neither all participants nor all top policies show spillover effects.

6.3.2 Spillover-capable policies dominate

Even though the previous section established that spillover frequency effects can emerge as a computationally rational adaptation to a particular payoff in the LLDT, a potential concern is that the model excludes some better-performing variant that does not show spillover. The strong version of this criticism is that the policy space excludes the optimal policy: the model must either take no advantage of memory sampling by setting the threshold to 0 (thereby introducing slack time during the saccade and eye-brain lag), or set it higher and risk delaying better perceptual samples in favor of noisier ones from memory.

To address this concern, this section introduces a model that is capable of always giving priority to processing perceptual samples over memory samples. Such a model can take advantage of memory sampling to make up for slack time during the saccade without the risk of prioritizing memory over perception. This model adds a single binary policy parameter, the *perceptual priority* bit. If this bit is set, then when the model has the choice between memory sampling from word N-1 and perceptual sampling from word N, it always chooses the latter. Such an option is not available in the previous model—there is no setting of the saccade and attention shift thresholds that will always use memory samples when only they are available, but also never choose to use memory samples when perceptual samples can be used. With the perceptual priority bit set, the model is capable of exploiting the least noisy samples available to it, but is incapable of exhibiting spillover effects.

Figure 6.5 shows speed-accuracy tradeoffs for the model, with the perceptual-priority bit unset (spillover-capable) and set (spillover-incapable), in three representative noise settings. Individual points are policies and the lines mark the best accuracy available at a particular reaction time for the two classes of policies; i.e. these lines represent the best speed-accuracy tradeoff possible for both spillover-capable and spillover-incapable policies.

In the left plot of the figure, noise is low enough overall such that responses are very fast and spillover-capable policies do no worse and no better than spillover-incapable policies. In the middle plot, update noise is higher, and the optimal speed-accuracy tradeoff is better for the model that can yield spillover, consistent with the exploitation of memory sampling to mitigate update noise. In the right plot, perception and memory noise are high enough that it is no longer useful to sample from memory at the expense of perception.

All the noise settings explored (see Figure 6.3 for the range) yield one of these three patterns, or the uninteresting case of near-chance performance. In no setting does the spillover-capable model perform worse than the spillover-incapable one. The noise settings cover a range from implausibly-high accuracy to chance performance, leading to the conclusion that spillover-capable policies dominate, in that they do no worse, and occasionally do better, than those constrained to give priority to perception over memory.

6.4 Why spillover arises from sequenced thresholded samplers

The model yields frequency spillover through a composed sequence of perception and memory sampling, which presents the following puzzle: if there is a fixed threshold at which memory sampling is initiated, then it should be surprising for the duration of memory sampling to depend on the prior of the perception sampler. Given that memory sampling is what leads to delay and spillover, it should be therefore surprising that frequency-sensitive spillover arises at all in the model.

The crucial insight is that it is not always the case that the true word hypothesis reaches the threshold first; i.e., the decision to initiate saccade planning may be based on (partial) recognition of a different word than the true word. In such cases, at the start of memory sampling, the hypothesis for the true word is farther from the attention shift threshold θ_a than if the true word had been (partially) recognized. Incorrect decisions are more likely for low frequency words, so in expectation the memory-driven attention shift mechanism will start farther from its threshold for low-frequency words, and therefore take longer to reach threshold, delaying the following word more.

To illustrate this phenomenon, Figure 6.6 shows data summaries from a minimal two-sampler model, without any architectural constraints. The top panel of Figure 6.6 illustrates the dynamics of such a *recovery* trial. In this panel, the threshold is crossed for the incorrect hypothesis (green line) in the first sampler, triggering the start of the second sampler. The second sampler recovers from the mistake, allowing the correct (red) hypothesis to cross the threshold, but at the cost of additional time. The bottom left panel shows that incorrect (and thus eligible for recovery) trials are more frequent for low priors. The bottom right panel shows that the finishing time of the second sampler is proportional to the prior probability of the correct hypothesis for the first sampler. It is also inversely proportional to accuracy (bottom left plot), consistent with inaccurate trials driving the relationship between the first sampler prior and second sampler finishing times.

6.5 Rational parafoveal preview in the LLDT

The first portion of the chapter showed how spillover effects might emerge in the absence of parafoveal preview, such as in the masked variant of the LLDT. But it does not provide a full explanation of what happens in the unmasked LLDT, because in that case parafoveal preview is available (though likely not very useful due to the wide spacing). This section introduces parafoveal preview into the model and explores how it interacts with the memory delay mechanism. In particular, it shows the following: first, that a model capable of preview performs better than a preview-incapable model across a wide range of speed-accuracy tradeoffs, and performs only slightly worse elsewhere. Second, that the preview-capable model shows spillover frequency effects. And third, that the spillover effects it shows are still a result of memory-driven delay, but that this effect appears in more of the model's parameter space due to an interaction between parafoveal preview and the model's memory mechanism. Some speculation is provided on how such a model might explain why denying parafoveal preview sometimes but not always eliminates spillover effects (as reviewed in S5.1.1).

6.5.1 Adding parafoveal preview to the model

In the memory delay model presented earlier in this chapter, saccade planning proceeds until a *saccade threshold* is hit, followed by 'free' perceptual sampling while saccade planning is ongoing, followed by memory sampling if the *attention shift threshold* has not been hit during the saccade planning time. Memory sampling continues until the attention shift threshold was hit if the *perceptual priority bit* was unset, or until perceptual samples from the next word became available otherwise.

The proposed change to incorporate preview is to replace the aforementioned 'free' sampling portion with an adaptively controlled stage: if the *attention shift threshold* is crossed during saccade planning time, then the next word is previewed. If it is not crossed before the eye moves on, memory sampling will ensue, with the same interactions with the attention shift threshold and perceptual priority bit as in the no-preview model.

As with the other models in the thesis, the model should have the correct generative model of samples, i.e. generation noise of foveal and parafoveal samples is known so that the correct likelihood function can be computed depending on the memory sample drawn. But a model that labels memory samples according to their noise might as well use the objectively less noisy foveal samples and avoid the noisier parafoveal ones, and so the preview model only buffers foveal samples and only uses them for the memory update.

Parafoveal sampling is implemented by scaling the perceptual noise very similarly to the scheme provided by Bicknell & Levy (2010a). Bicknell & Levy multiplied the standard deviation of the noise generation function at each position by $\frac{1}{\lambda(\epsilon)}$, where ϵ is the eccentricity of the current character in characters [2], and:

$$\lambda(\epsilon) = \int_{\epsilon-0.5}^{\epsilon+0.5} \frac{1}{Z} \exp\left(-\frac{x^2}{2\sigma^2}\right) dx, \sigma = \begin{cases} \sigma_L, x < 0 \\ \sigma_R, x \geq 0 \end{cases} \tag{6.1}$$

This is the asymmetric processing rate function from the SWIFT model of eye movement control (Engbert et al., 2005), with Z as a normalization constant. For the dissertation model, Z is set to the perception noise parameter and the integral (which is one character wide and centered at the eccentricity of the letter) approximated by evaluating the function at the eccentricity of the current letter. σ_L is set to 2.41 and σ_R to 3.74, following Engbert et al. (2005), and close to the value used by Laubrock, Kliegl & Engbert (2006). This is smaller than the value of 4.55 used by Richter et al. (2006), which would allow more parafoveal preview, but the latter paper also set σ_L to 0.05, which allows implausibly little perceptual processing even one letter to the left of fixation. This choice fixes what would otherwise be another free noise parameter (the parafoveal noise).

For purposes of computing the eccentricity of a string, the eye is assumed to always be fixated on the second character of the fixated word (this is slightly left-of-center in a four-letter word). The result is that the noise generation function for the second letter of the fixated word uses exactly the standard deviation specified by the perceptual noise parameter, and grows according to $\lambda(\epsilon)$ in the rest of the word. This means that in this model, foveal noise is slightly higher overall than in the memory-only noise above, for the same setting of the perceptual noise parameter. For purposes of computing noise of the parafoveal word, the model correctly takes into account that it is separated by six spaces, as in the human task.

At one limit of the resultant policy space, where both thresholds are equal, the model previews the next word during the entirety of the saccade planning time. At the other limit where the difference between thresholds is large, the model will not use preview and instead look identical to the memory model discussed earlier in the chapter, except with slightly higher noise because of the eccentricity noise adjustment. Intermediate values of the thresholds will yield intermediate behaviors. Because the duration of saccade planning is random, intermediate values of the thresholds will

[2] The value reported in the paper is $\frac{1}{\lambda(\epsilon)^2}$ but the correct value is not squared as per Klinton Bicknell, p.c.

yield policies that include some preview and some memory-based review, depending on the particular saccade planning time drawn, as long as the perceptual bit is unset. If the perceptual priority bit is set, there will be mostly preview, with memory review only if the current word has a very short saccade execution time and eye-brain lag, and the next word has an unusually long eye-brain lag.

Figure 6.7 shows a sequence diagram of the preview-capable model for three illustrative cases: in the top diagram, attention shift happens during saccade planning, resulting in preview. In the middle, attention shift happens after the eye has moved, resulting in memory sampling. In the bottom, attention shifts exactly when saccade is planned, resulting in maximum preview. The middle figure is illustrated with the perceptual priority bit unset; with it set, the attention shift decision would be truncated at the end of EBL and there would be no memory-driven delay.

As in the memory-only model, the parameter space was explored to illustrate the range of predictions the model is consistent with. Given that noise values above 2.5 were effectively at chance, higher noise values were not explored. In addition, the attention shift threshold range was from 0.499 to 0.999 in steps of 0.05. The perceptual priority bit was also allowed to take on both values to investigate whether memory review is still rational in the presence of parafoveal preview. The removal of the very low thresholds was done because preliminary experiments and the memory model above both showed that lower attention shift threshold values perform poorly. The move from five to three decimals allows exploring the maximum-preview policies discussed above, i.e. ones with saccade and attention shift thresholds equal.

6.5.2 Parafoveal preview is rational in most speed-accuracy tradeoffs in the LLDT

Figure 6.8 shows best speed-accuracy tradeoff curves, for the preview model introduced in this section and the no-preview memory model discussed in the previous section, in the same noise settings as Figure 6.5. These noise settings serve as representative cases for this comparison as well the previous one: the preview model achieves substantially better accuracies at higher RTs, but slightly worse ones at lower RTs. This is most likely because at low RTs, both models avoid memory sampling by keeping attention shift thresholds low, but the preview model replaces the 'free' foveal samples of the no-preview model with noisier preview samples. On the other hand, at high RTs, both models move the thresholds higher and the preview model replaces less informative memory samples (less informative because they only average over update noise) with more informative parafoveal samples (averaging over perception and update noise).

When it comes to performance on the Balanced payoff given to human participants, the preview model performs better in nearly all noise settings. There are two exceptions: first, when perceptual and update noise are very low, both models do equally well, likely at ceiling; second, when perceptual noise gets slightly higher but update noise is still low, the parafoveal preview model performs more poorly. A key confound in this result is that as mentioned above, the noise for foveal words is slightly higher in the preview-capable model than in the no-preview model because eccentricity is correctly computed for foveal words. Therefore the preview model is expected to perform slightly worse at equivalent noises, which might explain the slightly worse payoffs at lower RTs.

This result is largely consistent with parafoveal preview being adaptively useful to a rational reader, with the strong caveat that most of the differences in the model appear at response times above 2000ms, higher than the human RTs even in the Accuracy-emphasizing payoff.

6.5.3 Interaction of parafoveal preview and memory review

In the case of the memory-only model, policies with the perceptual priority bit unset dominated policies with this bit set. But it need not still be the case when parafoveal preview comes into the picture. Preview samples are enable complete averaging over sample generation noise and the recovery of the true word, whereas memory samples only permit averaging over memory and update noise and recovering the mean of the noisy perceptual samples gathered. A model that can choose between the two may well always choose the parafovea (i.e. set the attention shift threshold close to the saccade threshold) and not show any advantage to the memory-driven delay.

Figure 6.9 shows the same illustrative noise set as Figures 6.5 and 6.8. As in the no-preview model, delay policies do well when update noise is high (middle panel), likely for the same reason (because it is worth delaying to average over update noise). The picture in the rest of the space favors the delay-incapable model, with preview policies that are not at risk of delaying the foveal processing of the next word performing better. However, as in the case of comparing preview-capable to -incapable models, much of the difference exists at unrealistically high RTs.

6.5.4 Spillover effects in the preview model

The preview model is capable of recovering spillover effects in much of its policy space, but not in its entirety: by setting the saccade threshold low, the attention shift threshold high, and setting the perceptual priority bit, this model can choose to neither preview nor delay the next word, eliminating spillover. Therefore it is important to determine whether the model shows spillover in near-optimal policies.

Figure 6.10 shows distributions of spillover effects in the top 5% of policies (with respect to the Balanced payoff), across noise settings, separated by whether they have the perceptual priority bit set. The y-axis limits are set to be the same as the limits on figures 6.4 and 6.3 for ease of comparison, though this results in a truncation on the upper end of the range. It is apparent from this figure that spillover effects in the preview-capable model are still due to memory-driven delay rather than parafoveal preview. This is consistent with the claim that spillover effects are not driven by preview even in the unmasked LLDT. Note also that this model is even more reliant on the update noise for its explanation, as delay (and therefore spillover) in the preview model is only rational at high update noise settings.

This result presents another puzzle. If parafoveal sampling has minimal impact on spillover, then why is it the case that the effects in Figure 6.10 (the preview-capable model) are obviously larger than those in Figure 6.3 (the memory-only model), as well as present in more noise settings? A possibility has to do with the choice made above to only use foveal samples in the memory buffer. Under most policy settings in the preview-capable model, there will be some parafoveal and some memory sampling in a given trial (depending on architecture draws), though a single word will only see one type of sampling. In the preview-incapable memory-only model, there is only foveal and memory sampling. Therefore, under equivalent policies and noise settings, the memory-only model would have more foveal samples available, be better able to rationally average over update noise, and therefore have better speeds and accuracies overall.

This benefit is likely what drives better performance in the preview model when update noise is not high, but will also make memory sampling drive bigger delays into the following word. This is unlikely to be an effect of the deadline on the memory sampler in the preview model: the deadline was extended to 100 samples (1000ms) in the preview model to mitigate against this possibility, and spillover effects seem to appear even in noise settings where fixation durations are far less than a second long.

It is not clear whether this explanation should be treated as an interesting theoretical finding, or a curious modeling artifact. The memory buffer used is of no deep theoretical commitment, so relying on it for the apparent recovery of spillover is no great theoretical strength. Future work might investigate whether the same pattern holds with buffered parafoveal samples, or in a more sophisticated memory model (such as the one sketched in §7.4.3).

6.6 Discussion and Conclusion

This chapter built on the suggestion of a delay-based explanation for spillover effects derived from the E-Z Reader modeling efforts in Chapter 5 by introducing an explicit mechanism that could extend processing into the next word. This mechanism is a memory-driven attention shift: perceptual samples are buffered and continue to be processed from memory until a second adaptive threshold is hit. Such a mechanism is shown to be rational across a range of speed-accuracy tradeoffs, assuming independent noise in perception and the update process. The reason for this is that it may be rational to reprocess past perceptual samples with independent draws of noise in the update process, to rationally average over this noise. This mechanism is also shown to produce spillover effects in settings where update noise is high.

A puzzle is introduced for why the second sampling period duration would be sensitive to the prior of the first under a threshold policy, as the threshold is expected to make the starting probability of the second sampling period independent of prior. It is shown that while the threshold of the max a posteriori hypothesis (i.e. string) is equal to the threshold, the probability of that string being the true string is a function of prior. The second sampling period can recover from an incorrect crossing by the first, at the expense of additional time (because the correct string probability is more likely to be below threshold as the prior gets lower). This appears to be a fundamental property of combined samplers: information about the prior persists in the posterior in spite of the threshold nature of the sequential test.

Next, parafoveal preview is introduced in the model in order to understand how a model that can adaptively select between foveal, parafoveal and memory sampling navigates that adaptive space, and by what mechanism it may yield the spillover effect of interest. It is shown that the speed-accuracy tradeoff results are more mixed in this case: the use of parafoveal preview performs better than an equivalent no-preview model at higher RTs, but close or even slightly worse at lower RTs. Spillover effects in the preview capable model, assuming the large word spacing used in the LLDT, are entirely driven by the delay-based rather than preview-based explanation. This supports the assertion that the LLDT minimizes parafoveal processing, and that spillover effects in this task are not parafoveally driven even when the parafovea is not masked. With closer spacing (and therefore less noisy parafoveal samples), the tradeoff between parafovea and memory will become more apparent. In particular, for some values of memory noise and parafoveal noise, it will be advantageous to always choose the parafovea over memory, and parafoveal preview will drive spillover. Possibly even with both noises equal, parafoveal sampling may be preferred because it draws 'fresh' un-updated samples with completely independent noise draws rather than buffered samples with only an independent memory noise draw. As parafoveal noise increases (due to spacing or individual variation, or gaze-contingent manipulation), memory delay will increasingly drive spillover.

An interesting future direction in light of this work is to revisit the question of spillover results as a function of incorrect rather than obviously masked preview. As noted in §6.1, some prior results show the disappearance of spillover effects in the absence of parafoveal preview, while some show only a reduction. These results could be understood as a mixture of two different kinds of participants: those who notice a display change and show large spillover effects, and those who do not notice the change and show a reversal, i.e. slowdown for words that follow high-frequency words. Slightly different masking speeds and proportions of participants noticing the display change might yield the gradation in results seen in prior work.

To investigate this, the preview model above might be modified to reduce word spacing from six characters to one (increasing the fidelity of preview), but provide incorrect preview. Instead of drawing parafoveal samples from the correct preview, the model would draw them from a random set of four letters, which means that sequential sampling from the parafovea would systematically bias the model's belief away from the correct string. If higher frequency foveal words allow the model more parafoveal sampling, then this parafoveal sampling will bias the following word away from the correct decision.

This model would be structurally capable of producing the spillover effect reversal seen in the White et al. (2005) display-change-aware group, specifically by setting the attention shift threshold low and close to the saccade threshold. The model would still be structurally capable of producing the spillover speedup effect as well, under the memory mechanism discussed earlier in the chapter (by setting the attention shift threshold high).

A key challenge in this effort is to model the display awareness of participants. One possibility is to treat the no-preview model advanced in §6.2 as one extreme, where participants know that parafoeval preview provides no useful information and can completely ignore the parafovea. As shown above, memory-driven spillover effects may still emerge in this setting under certain noise conditions. The other extreme on this spectrum is the preview model in §6.5.1, where participants might find the rational adaptation to incorrect preview without being able to shut the parafovea down completely.

However, the actual empirical setting in which the spillover reversal has been found is neither of these two: parafoveal preview in those studies is disrupted in only a small proportion of words, which is part of the reason participants rarely notice the display change manipulation. In a typical design with a single critical word per sentence, long sentences typical in psycholinguistic experiments, and many fillers, only as few 5-10% of words have an incorrect preview. A reader who thinks display changes are uncommon may therefore choose to utilize parafoveal preview in these experiments, whereas a reader who thinks changes are common would ignore preview in favor of memory-based review. Such a policy would yield rational behavior (and presumably, parafoveal preview) in precisely the 90-95% of the experiment that is *not* being measured.

Understanding rational behavior in this setting, therefore, might require optimizing the model in a setting where only some words are given incorrect preview, and investigating what proportion of correct and incorrect previews drives which type of spillover behavior. Such a manipulation could of course also be undertaken in human participants.

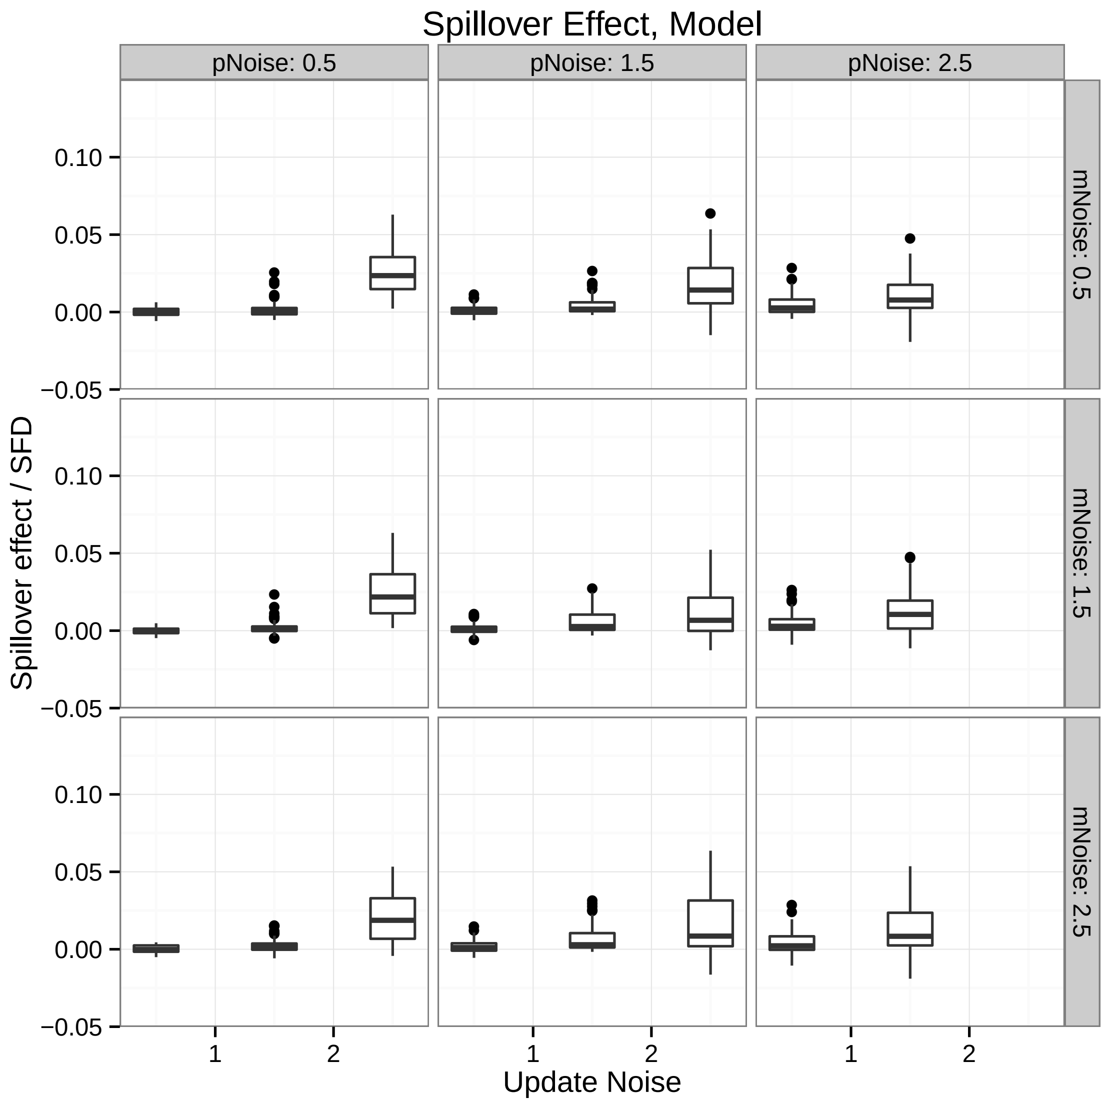

Figure 6.3: Spillover effects generated by the top 5% of policies across different settings of memory, perception, and update noise. Plotted is the ratio of the spillover effect to single fixation durations, which normalizes out the long fixations at higher noise settings. On each distinct machine defined by a combination of noise settings, policies (settings of $\theta_s, \theta_a, \theta_r$) were evaluated by the same task payoff given to human participants in the experiments described in Chapters 4 and 5. Boxplots show spillover effects of the top-performing 5% of policies. Spillover effects are the difference in mean single fixation durations on word N when word N-1 is low frequency and when word N-1 is high frequency (low/high determined by median split). The highest noise settings in the bottom row are not shown because performance was near-chance even for the best policies.

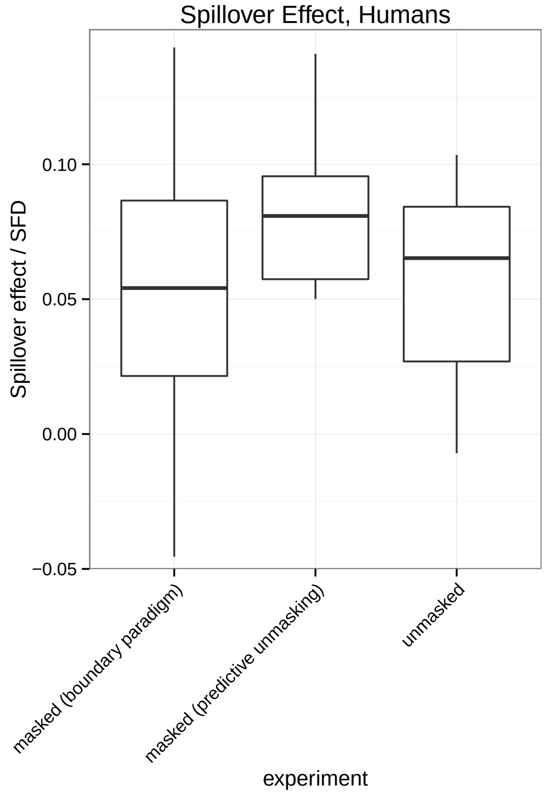

Figure 6.4: Spillover effects in the human participants. Effects plotted using the same normalization as the model (i.e. divided by SFD) and on the same scale. The observations that go into the boxplot calculation are effect sizes for each participant, computed in the same way as for the model.

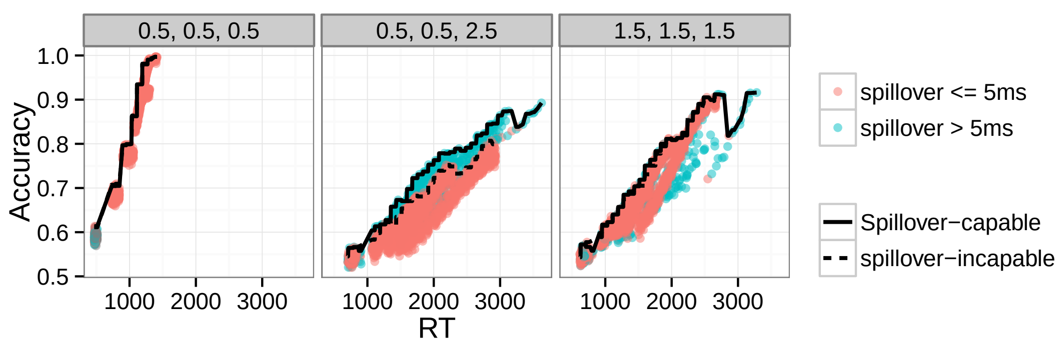

Figure 6.5: Speed-accuracy tradeoff curves for some representative noise settings. Points on the line correspond individual policies (i.e. setting of the three decision thresholds). Plotted are mean trial RT and accuracy (computed from 5000 simulated trials), color-coded by whether the policies yielded spillover frequency effects. Lines mark the best speed-accuracy tradeoff available to spillover-capable and incapable policies. Each plot is labeled at the top with the noise setting *(perceptual, memory, update)*.

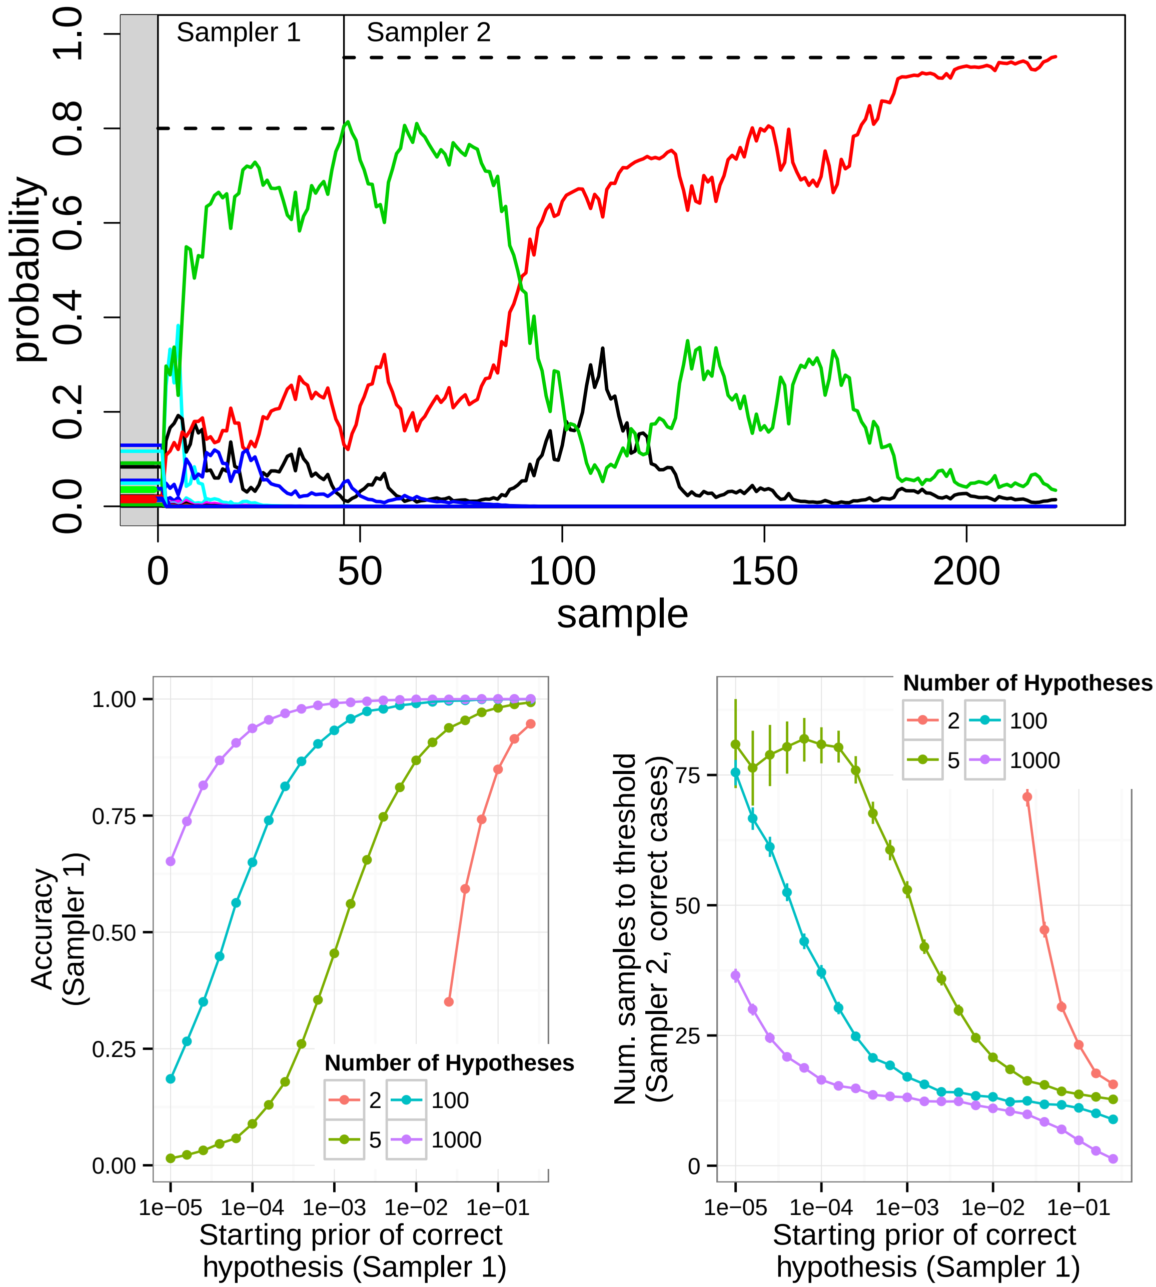

Figure 6.6: A simple example illustrating how the prior for a thresholded sampler affects its final posterior, and therefore the prior for a subsequent coupled sampler, despite the fixed threshold. *Top*: An example 'recovery' trial for 500 hypotheses (words). *Bottom Left*: Accuracy for the first sampler as a function of the prior of the true hypothesis. *Bottom Right*: Second sampler finishing times as a function of to the true-hypothesis prior in the first sampler.

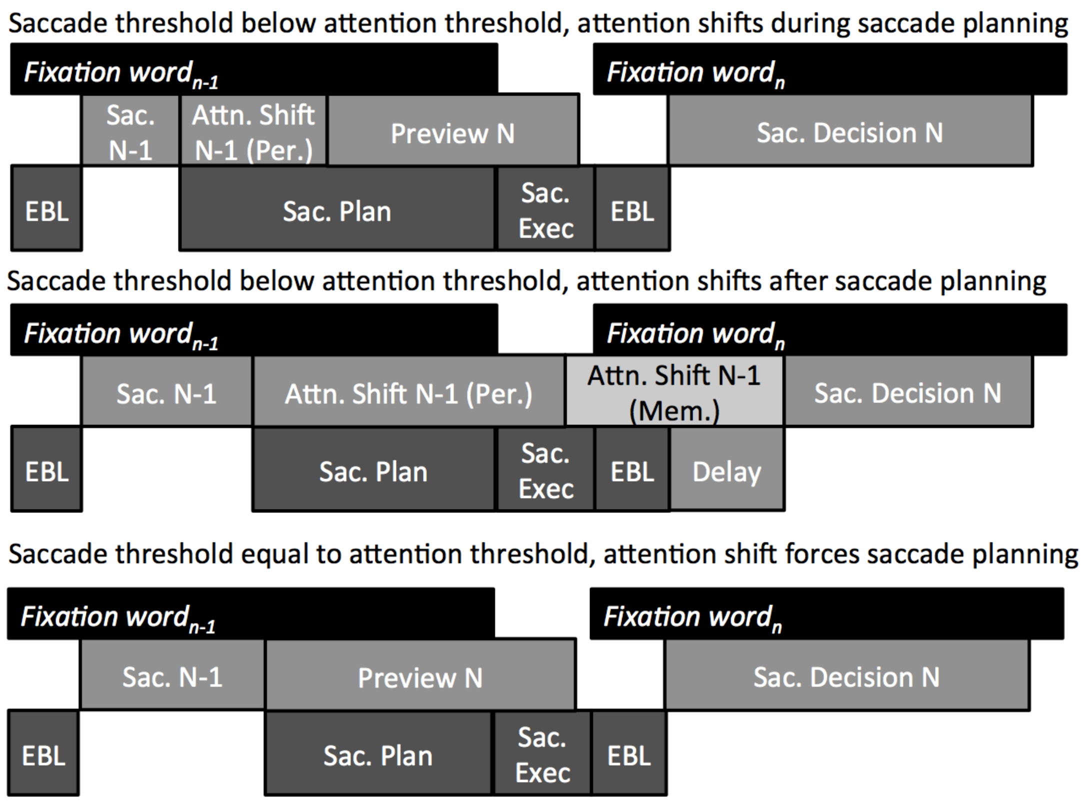

Figure 6.7: *Parafoveal preview capable model sequence diagram.* See § 6.5.1 for description.

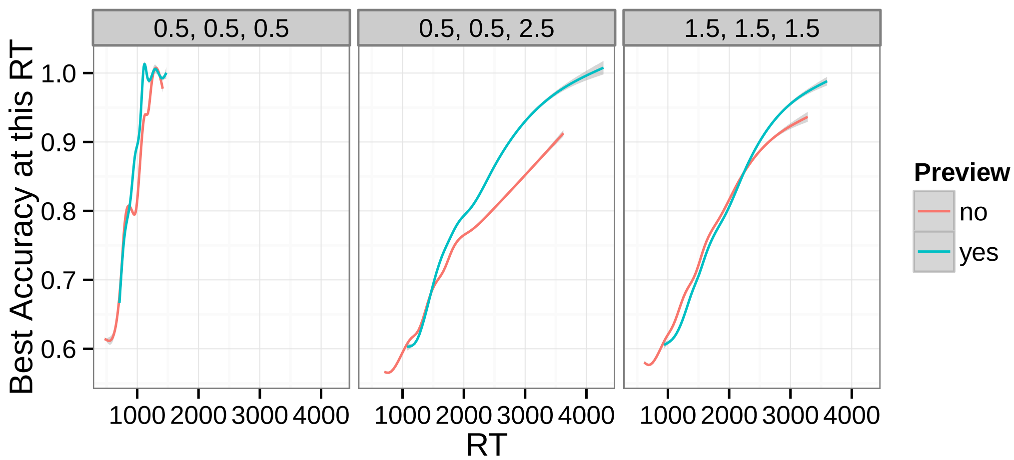

Figure 6.8: *Best Speed-accuracy curve for preview and non-preview model, at the same noises shown for the memory-only model.* Shown is a smoothed line connecting best accuracies available at each RT. At low noise, both preview and no-preview policies are effectively at ceiling. With high update noise, the preview policy performs substantially better at high RTs but slightly lower at at low RTs. The effect is more pronounced at more moderate noise settings, with the preview policy performing substantially better at slower RTs and slightly worse at faster RTs. The remainder of the settings look like one of these three examples, with the rightmost one most common. Vertical lines mark points of best performance under the Balanced payoff for the two model variants: the preview model can afford to wait longer because the boost in accuracy is worthwhile.

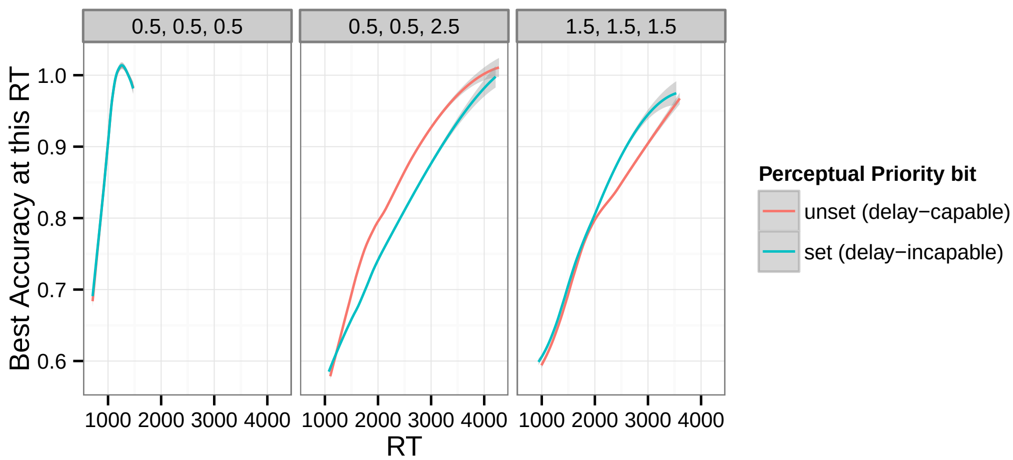

Figure 6.9: *Best Speed-accuracy curve for preview model with perceptual priority bit set and unset, selected noises.* Shown is a smoothed line connecting best accuracies available at each RT. At low noise, both parts of the policy space do equally well. With high update noise, the delay-capable policy performs better, consistent with results from the no-preview model. At other noise values (of which the right plot is one instance), delay-incapable policies do better.

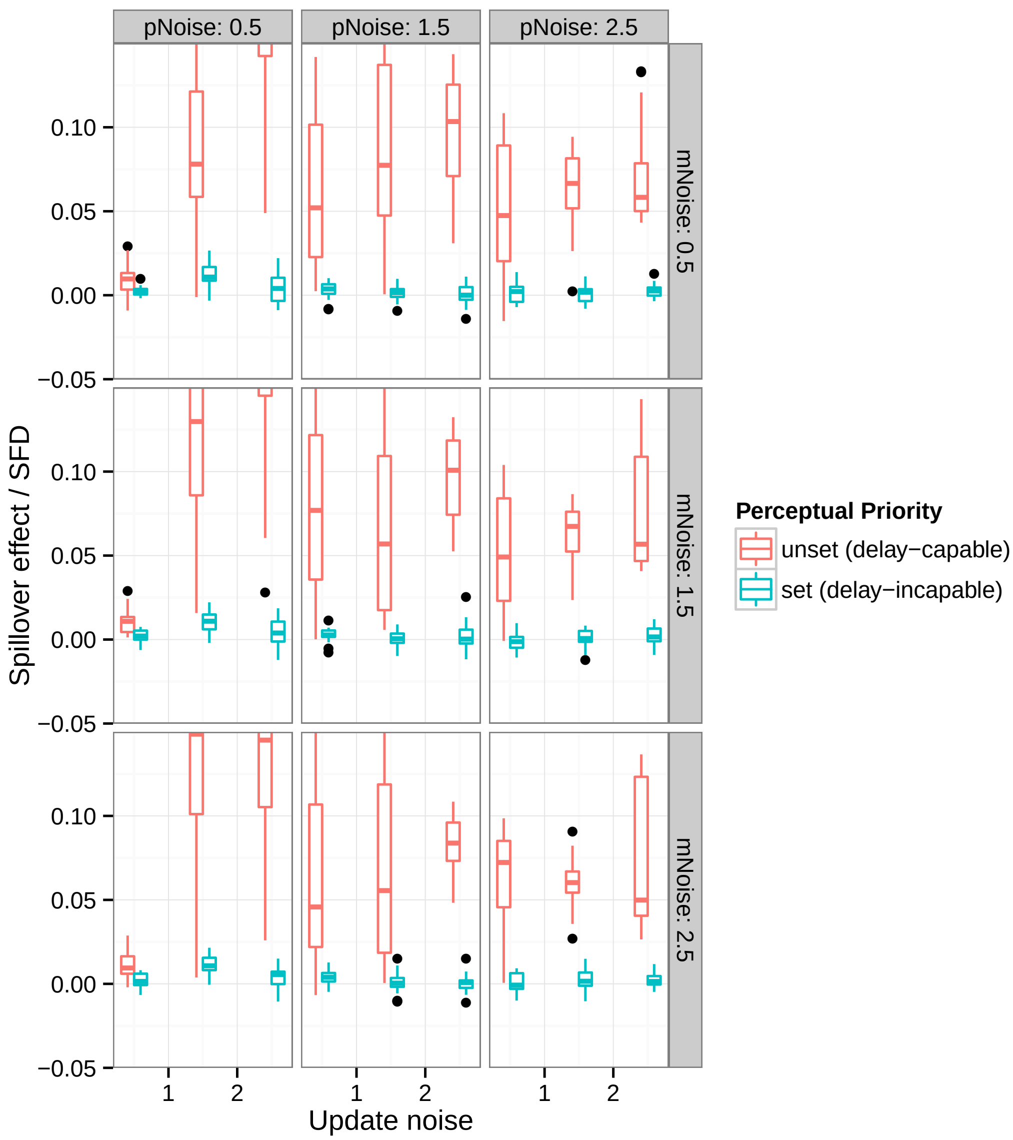

Figure 6.10: Spillover effect sizes scaled against SFD in the preview model, with the perceptual priority bit set and unset. Spillover effects are only present when the bit is unset, consistent with memory-based delay driving spillover effects even in the preview-capable model when word spacing is wide.

CHAPTER 7

Fitting the model to human data

7.1 The advantage and importance of model fits

This chapter details the attempt to find a good fit to human data in the LLDT. While the larger-scale exploration in the previous chapter provides some insight into the conditions under which spillover might arise in the model, one of the benefits of a theory that makes millisecond-level predictions is exactly the ability to test those predictions against human data. As in Chapter 4, this section pursues fit against single fixation duration first. As noted in Chapter 6, there may well be multiple best fits on single fixation duration that behave differently with respect to other measures. Nonetheless, the recovery of frequency and spillover effects at reasonable SFD values should be a minimal requirement of any fit against human data.

7.2 Comparing single fixation durations in humans to the memory-delay model

Figure 7.1 shows mean fixation durations for high and low previous frequency bins for the top 5% of policies for the spillover-capable memory (no-preview) model discussed in the previous chapter. Dashed lines mark human performance in the Balanced payoff in the LLDT with predictive unmasking: this is the empirical target. It is easy to see that the range of noises where the spillover effect appears (indicated by a separation between the two solid lines), yields single fixation durations far higher than those in the humans. On the other hand, lower noises that yield single fixation durations in the correct range for the humans show no spillover.

One possibility is that the search using the noises above was using too-high noises overall. To investigate this, a model was tested with perception noise set to 0.1 and update noise set to 0.5, across a small range of memory noises. The same ratio in the higher noise data (0.5 perception and 2.5 memory) produced a sizable spillover effect in the top 5% of policies, so one might have expected to see the same with all the noises scaled down. However, this is not the case: Figure 7.2 shows single fixations only slightly shorter than those in the humans, but no sign of spillover effect.

A second possibility is that the sample rate is too low to see fine-grained effects of the prior, especially at such low noises. At 10ms per sample and 175ms non-controlled sampling (125ms saccade planning and 50ms eye-brain lag), there are only about 7-8 samples that are actually under adaptive control without any memory-based delay. With memory-based delay in play, this number shrinks further and the differences in controlled sampling duration may be lost. To investigate this, sample duration was reduced from 10ms to 2ms. Figure 7.3 explores fixation durations across previous-frequency bins at a 2ms sample rate. While the fixation durations in this plot are far closer to those in the humans, the spillover effect sizes are about an order of magnitude too small.

A third possibility is that the lack of saccade cancellation removes an adaptive locus available to human participants from the model. Evidence from saccade countermanding tasks (e.g. Becker & Jürgens, 1979) is consistent with a majority of the duration of saccade planning (as much as 80-90%) being *labile*, i.e. amenable to inhibition or cancellation. This was investigated by setting the saccade planning time to 25ms instead of 125ms. This can be thought of as an extreme variant of a labile-saccade model where saccades can be inhibited, canceled and resumed so precisely that the effective maximum delay for a saccade plan is the short non-labile component. Figure 7.4 shows a summary from model runs with mean saccade planning duration set to 25ms instead of 125ms. Fixations are faster overall as compared to Figure 7.1 and spillover effects do appear at slightly lower noise values (and therefore slightly faster mean fixation durations) than those in the long-SPT model.

This last modification is the most promising, and one next step is to combine some of the above ideas – i.e. a short saccade planning duration with lower noise and a faster sample rate, which might finally yield effects in the right magnitude at the right mean SFDs. But such a model would at best be equal parts computationally rational analysis and model-fitting exercise. Even with proper train-validate methods (i.e. training on one payoff and validating on another), a parsimonious explanation should not require a precise setting of five parameters (three noises, sample rate, and saccade planning time) to recover a single effect of interest.

This is particularly notable compared to effects like the frequency effect on single fixation duration, which is consistent in its appearance and size in the top 5% of policies across the noise spectrum (Figure 7.5). Likewise this is unappealing in contrast to the model in Chapter 4, which provided effects in the correct range over many measures using a single fit parameter.

A better option might be to pursue the saccade planning direction in a more sophisticated manner than turning it into a free parameter, by attempting to understand how a decision to *inhibit or cancel* a saccade as well as to initiate one might be triggered based on the ongoing state of lexical processing modeled as a sequential sampling machine. This would be a departure from the E-Z Reader inspired architecture, which only uses saccade cancellation to handle higher level post-lexical processing (Reichle et al., 2009).

The challenge in doing this would be to articulate how the saccade cancellation decision is conditioned over the model's hypothesis space: random walk models of the saccade countermanding paradigm (e.g. Boucher, Palmeri, Logan & Schall, 2007) or saccade inhibition in reading (e.g. Engbert et al., 2005) that one might draw inspiration from are expressed in terms of competition or inhibition of activation. A correct treatment of inhibition in a model like the one in the thesis would need to explain how inhibition of activation is reflected in the model's probabilistic belief space and how a cancellation action is conditioned on it.

7.3 Comparing single fixation durations in humans to the parafoveal preview model

In light of the problems above in the memory-only model, one might expect a fit using the preview model to be a futile exercise, with the preview model spillover effects still driven by memory delays. But this may not be the case: recall the curious result where parafoveal sampling limited the amount of information available to the memory system, and made memory-driven spillover effects present in more of the model's free parameter space.

Figure 7.6 shows fixation durations split by previous-word frequency bin, similarly to the figures above. The higher perceptual-noise setting is excluded because fixation durations in that setting start above the y-limit of the rest of the plot. This plot demonstrates that spillover effects do exist in the model's top 5% of policies, in a range of single fixation durations close to those in the humans (though the spillover effects are slightly small). This is a promising result, though caveatted by its apparent reliance on the modeling choice to not buffer preview samples reducing the number of memory samples available and reducing memory performance overall in the preview model.

7.4 On the problems with quantiatively fitting human data

This chapter shows that the values of free noise parameters that yield spillover in near-optimal policies also produce single fixation durations far higher than those in the human data. This is a problem for the theory as stated, though in some sense it is a good problem to have: a more abstract model without the ability to make millisecond-level predictions would not be able to exclude this particular explanation relying on a rational memory-buffer-driven delay. Neither would a non-rational model, which would find policies even at low noise settings showing spillover (ones with low saccade thresholds and high attention shift thresholds).

The model including parafoveal preview suffers from this issue to a lesser extent, only underestimating spillover effects in the correct range of fixation durations rather than missing them outright. However, parafoveal preview in this model has minimal impact on spillover except by changing the success of memory sampling via a quirk of the memory sampler implementation. Future work might introduce the memory buffering of parafoveal samples into the model, in order to understand how strongly this quirk of implementation is driving the effect.

Before considering alternate explanations, it is worth remarking that with the number of parameters and coarseness of discretization in the fitting exercise undertaken, the parameter space may not have been adequately characterized, and qualitatively and quantiatively better fits may have been missed. Unlike the single-sampler model in Chapter 4, the memory model is particularly vulnerable to this problem because multiple noise parameters jointly contribute to the scaling of the sampler and trade against each other. For example, the best fit could be when either update or perception noise is zero (with the other noise scaling the sampler alone) or there could be multiple equally good fits with different settings of the trading parameters.

That said, based on the data collected it certainly seems like the model specified in the previous chapter cannot predict spillover effects at human-magnitude SFDs. It is important to remark that what is rejected is not the full class of delay models, or of memory-delay models. Rather, it is this particular buffered memory model and noise scheme, and the rationality of using memory to average out update noise at human-magnitude SFDs.

This present section discusses a few other possibilities of how spillover might be explained in light of these results. Some of the variants sketched below retain the ideas of memory sampling, delay-based spillover, or both, while changing other aspects of the model. The intent is not to introduce free parameters or adjust the model architecture arbitrarily to make the effect appear rational. Rather, it is to understand how the simplifications taken in the model (in terms of memory architecture, oculomotor architecture, or the decision process itself) might have eliminated spillover effects at human-like SFD magnitudes in the model space.

The first revision considered is the possibility of viewing spillover as a departure from rationality. It is rejected because of how it would weaken the remainder of the rational results in the thesis. More plausible is the possibility considered next of treating spillover as an invariant consequence of architecture. Finally, some other architecture variants are sketched within which post-perceptual processing might be rational for other reasons.

7.4.1 Spillover as a departure from rationality

Within the framework of computationally rational analysis, behavior is viewed as approximately rational by assumption. The advantage of this is that the space of possible theories is constrained, but the disadvantage is that there is always the risk that participants are not actually rational in the same way that the model proposes. In light of evidence that participants do attempt to adapt their behavior to payoff (such as that provided in Chapter 4), a complete rejection of adaptation is not warranted. However, spillover could be a signature of some local minimum in the adaptive payoff surface. It could also be rational with respect to some additional constraints on sentence-reading but not rational in the LLDT, and participants may still use their strategies practiced and trained to normal reading. Finally, the non-spillover policy could be fragile or hard to optimize, again driving towards the selection of spillover policies for the LLDT (this is a variant of the *satisficing* argument from bounded rationality, Simon, 1955).

To accept any of these possibilities would exclude spillover from the class of phenomena explained as a rational adaptation. But the initial failure to explain spillover under the bounded optimality assumption should not be sufficient cause to abandon the computationally rational approach with respect to the phenomenon. Rather, the proper choice in the computationally rational approach is to iterate on the theory to understand the link between task constraints, agent bounds, and the behavioral outcome of interest. This is not to say that all phenomena must be understood as rational: a future theory of biological, cultural, and within-the task individual adaptation may provide a taxonomy of bounded-rational and satisficed behaviors.

7.4.2 Spillover as a consequence of architecture

The memory delay architecture makes the assumption that the existence of any post-perceptual processing is under adaptive control, and that the only goal of this processing is to facilitate perceptual recognition under update noise. Both of these assumptions can be plausibly relaxed: it is almost certain that post-perceptual processing is required for additional reasons. In the case of reading post-perceptual processing might be required to make contact with the conceptual or semantic system, and participants may leverage this system in the LLDT as well. This system may well be sensitive to frequency and capable of delaying perceptual processing like the simplistic memory sampler proposed in the dissertation, producing spillover.

Other architectural explanations are provided in all of the non-adaptive models of reading that provide explanations for spillover – these explanations are architectural by definition, because no adaptive component is present. The non-preview explanations are variations on the theme of hysteresis: in E-Z Reader, for example, the portion of processing after saccade planning begins is a fixed proportion to the time before saccade planning, and potentially extends into the next word. In SWIFT, decisions to delay a saccade can spill forward into the next saccade, also yielding spillover as a consequence of architecture.

Another variant of the architectural explanation would be to tie the noise parameters to something that is possible to estimate a priori outside of the model (e.g. neural firing rates in relevant areas), and show how those a priori estimates match those in the spillover model. Such a model would still need to be modified to yield short SFDs, but would do so with one degree of freedom fewer.

Building spillover into the architecture would provide a somewhat unsatisfying answer to the question of why spillover effects exist: they exist because the architecture is constrained to produce them. It would not be less true for being unsatisfying, and as with the rest of the architecture would require a priori empirical or theoretical motivation.

7.4.3 Spillover as rational under other constraints or architecture variants

This is the correct choie under both rational and computationally rational analysis: iterating on the model specification until the phenomenon of interest arises in rational behavior. A minimal model change might retain the idea of a post-perceptual sampling delay and investigate other reasons that such post-perceptual sampling might be adaptive, i.e. ones not relying on a particular set of noise settings.

One possibility is to investigate the quantitative aspects of the oculomotor machine. The parametric machine architecture presented in the dissertation made strong commitments to architecture parameters in state-of-the-art models, in particular E-Z Reader and SWIFT. Such a commitment was shown to be useful in Chapter 4, which showed how deviations from those architecture values worsen model fits. However, §7.2 showed that deviations from those consensus architecture parameters may also improve fits. Rather than turning the architecture parameters into free parameters as discussed in that section, one might reconsider the architecture in light of the a priori empirical work on architecture discussed in §3.2.3.

A second possibility is to replace the simple memory buffer with a more theoretically interesting memory retrieval process – even one still modeled as a sequential sampler[1] In such a sampler, the prior might be set by recency or some other function of memory, and the sequential updates are done based on noisy samples from a cue. When a threshold is reached, the cue is identified well enough to retrieve a memory trace it corresponds to. The noisy cue samples are drawn based on the current state of the perceptual sampler. More perceptual sampling would yield more reliable cues (i.e. ones better at averaging over perceptual noise; in the limit, the mean of the cue sampling distribution is the true percept out in the world, a noiseless representation of the word). If the memory hypothesis space is over words, then the closer match between the hypothesis space and sampling distribution of the cue would result in faster and more accurate memory retrievals. But a memory sampler prioritizing accuracy might be able to mitigate against a lower-threshold perceptual sampler. In this model, a memory delay may be adaptive not as a consequence of high update noise, but as consequence of noisy perception. Such a model might also predict that introducing noise in the visual display would magnify spillover effects by increasing the emphasis on memory. A non-sampler memory could be introduced as well (e.g. Lewis, Vasishth & Van Dyke, 2006; Lewis & Vasishth, 2005).

[1] Thanks are due to Rick Lewis for ideas regarding this memory mechanism.

Yet another possibility is to make a dynamic treatment of the prior. It may be the case that an informative prior is not instantly available, and instead comes on-line dynamically. Early on, such a model looks like a uniform-prior model, and later it looks like a highly-informative-prior model. The advantage of a memory-driven delay would be in waiting for more prior bias to become available in difficult decisions. The dynamic prior could itself be a rational adaptation to something else: for example, Hanks, Mazurek, Kiani, Hopp & Shadlen (2011) argue for a model that integrates a time-varying bias signal in the drift diffusion model of decision-making, which is analogous to a time-varying prior in the MSPRT. They make this argument on two grounds: first, to account for the fact that low-probability stimuli are identified both more slowly and less accurately than high-probability stimuli; and second, as a mechanism for the estimation of stimulus noise. The first of these is unexpected, since both the MSPRT and related DDM both show the pattern of lower speed and lower accuracy for lower-prior items (MSPRT: Figure 6.6; DDM: e.g. Wagenmakers et al., 2008). But the second is intriguing: the argument that Hanks et al. make is that time-on-decision can be used as a proxy for the reliability of incoming stimuli. If the decision happens quickly, then presumably the incoming perceptual information was reliable and should be weighted higher. On the other hand, if a long time passes without a decision then perhaps the incoming information was unreliable, and the prior should be trusted more.

7.5 Discussion and Conclusion

Based on the results from the previous section, one might be tempted to conclude on a negative note. After all, while spillover effects do appear to be computationally rational in some portion of the no-preview model's free parameter space, they do so far from the range consistent with human data, and fairly drastic parameter manipulations cannot seem to change this. Put differently, the model does not posses a degree of freedom that corresponds to the scaling method undertaken in the initial computationally rational explanations. The fact that the parameter-fitting exercise largely fails is indication that the model is not overly flexible, but this might seem like small consolation if it is constrained to *not* produce the result of interest. The preview-capable model does provide more promising fits on SFDs, but as discussed in §6.5.4, may be relying on a quirk in the memory sampler implementation.

Nonetheless, a number of useful conclusions can still be made from this work. First, a minor point: a non rational-analysis approach would have failed to reject the memory buffer explanation for spillover. Figure 7.7 shows a spillover effect generated from the model that allows policy parameters to fit to human data rather than optimizing to payoff performance.

Second – and more substantively – although a second sampler is an elegant way of capturing continued processing of perceptual inputs, spillover may arise from another process with access to the posterior of the perceptual sampling as long as such a process can recover from perceptually misidentified words and competes for processing resources with perception. Such a process could be a lexical completion stage as in E-Z Reader, probabilistic memory retrieval, or something else. Chapter 6 provides the first illustration of why such a second process will have access to the priors of the first even if the first is a fixed-threshold perceptual recognition machine.

In addition, while the present model motivates post-perceptual memory-based processing by mitigating noise in the update process itself, it is almost certainly the case that post-perceptual processing is required in the course of reading for independent reasons. Such processing could also yield spillover frequency effects in a way that the memory sampling process does, but without the limitation of only being rational at unreasonable noise settings. A challenge for such an alternate process is that spillover effects persist in the LLDT in the absence of required higher level syntactic or semantic processing, but a more realistic model of memory need not require those higher level properties.

Even the memory resampling process as yielding the delay is not categorically excluded: §7.4.2 noted how a priori estimates of noise might yield spillover in the memory-resampling model, but require other changes in scale to yield human-plausible SFDs.

The present model and results suggest several avenues for future work. One is a search for another delay-type mechanism that might yield spillover without being rational only at unusual noise settings. A promising avenue on this front is a time-varying effect of prior (Hanks et al., 2011).

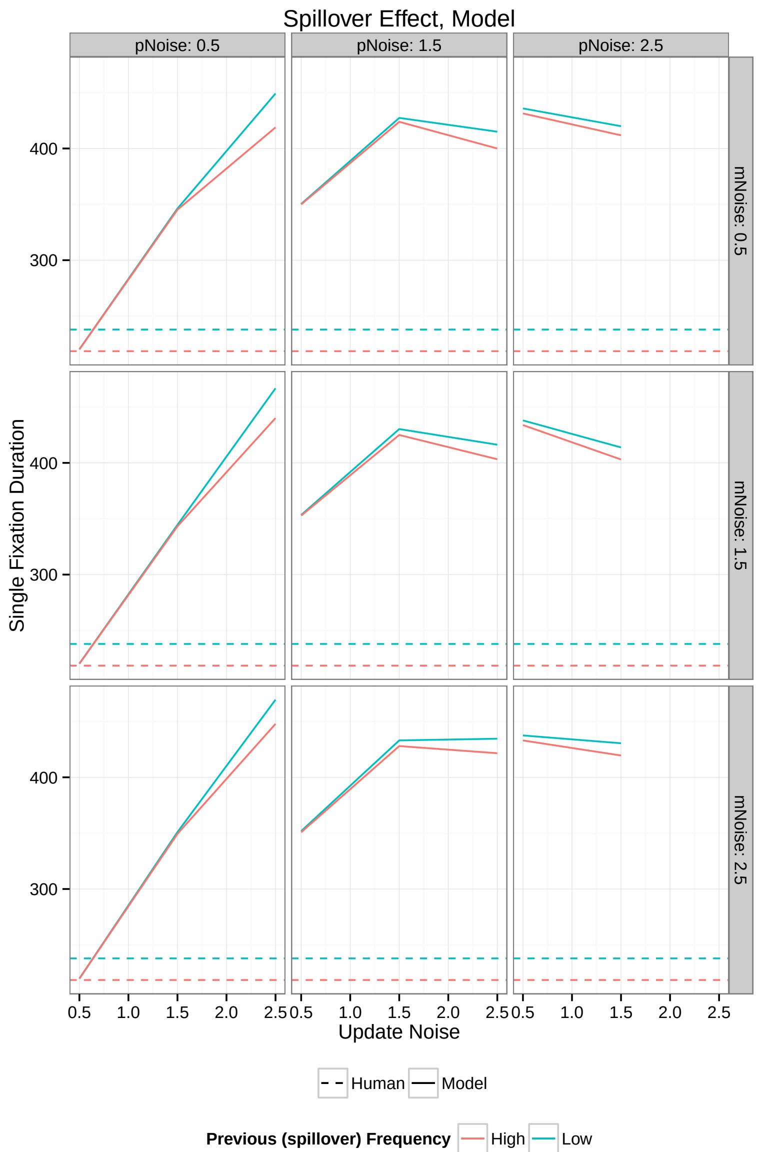

Figure 7.1: Model cannot simultaneously recover spillover effect and single fixation duration of the correct magnitude. Plotted are single fixation durations by previous-frequency bin for the top 5% of policies at each noise setting, with dashed horizontal lines indicating the values in the humans. Lattice plots marked with perception and memory noise, with update noise on the x-axis. It is clear that the model largely overestimates single fixation durations, and moreover that it fails to recover spillover effects in near-optimal policies at magnitudes of SFD that are consistent with those in the human LLDT data.

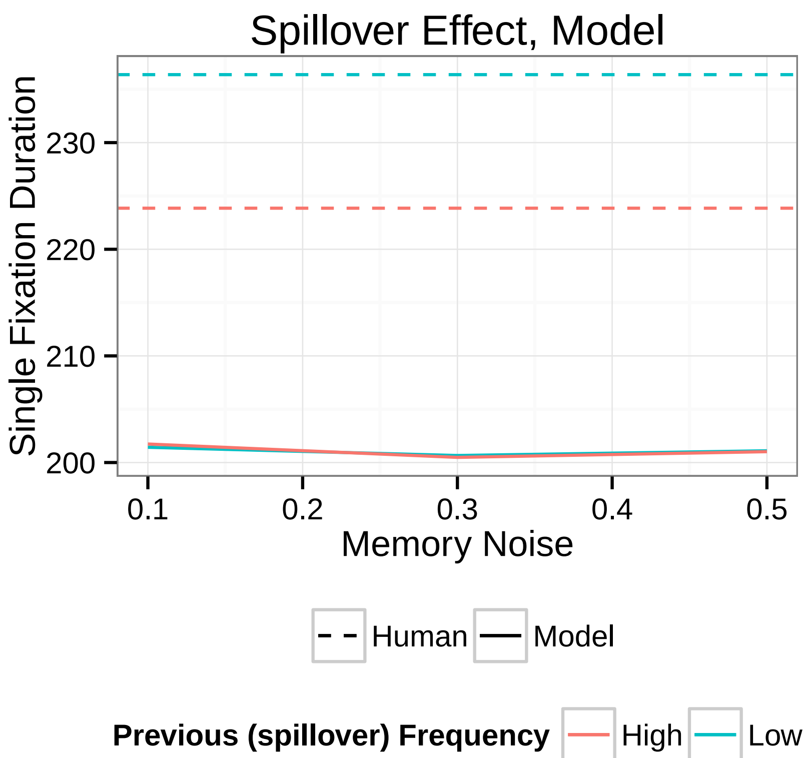

Figure 7.2: *Model at low noises fails to recover spillover, even with the same ratio of noises as the higher-noise model.* Points on the line correspond individual policies (i.e. setting of the three decision thresholds). Plotted are single fixation durations by previous-frequency bin for the top 5% of policies at each noise setting, with dashed horizontal lines indicating the values in the humans. While this model recovers fixation durations close to those in the human data, it fails to recover spillover effects entirely.

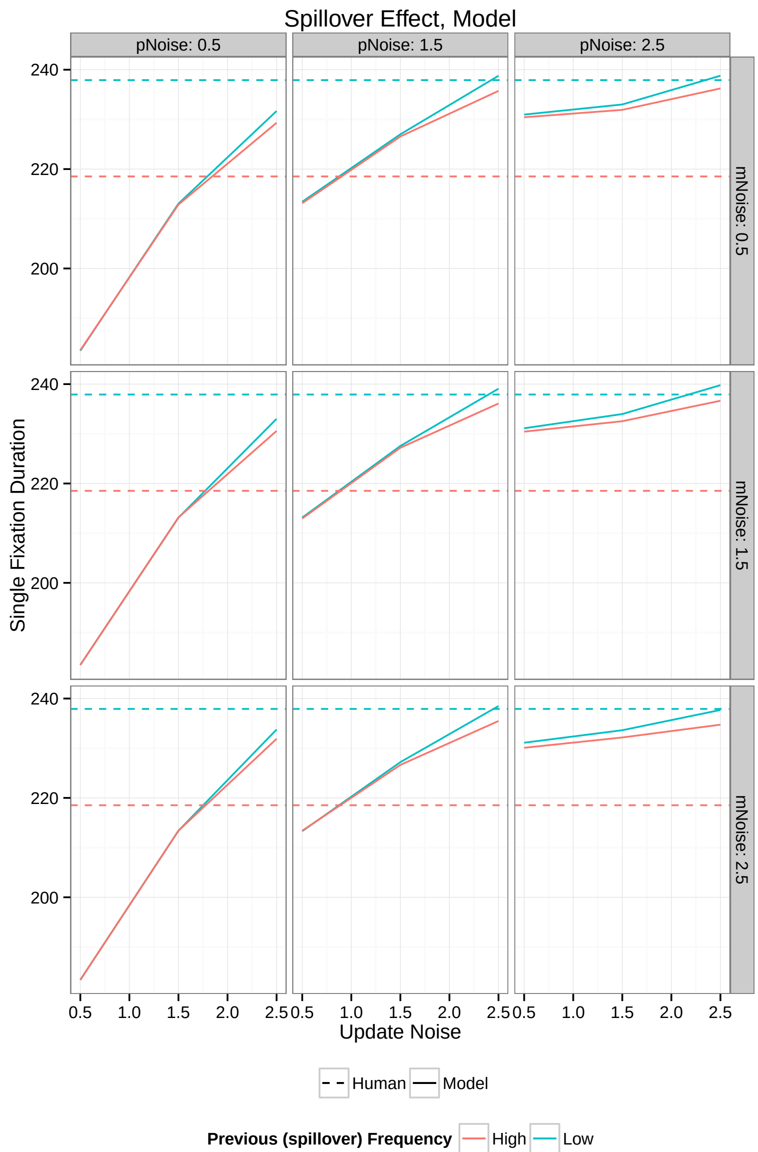

Figure 7.3: Model with fast sample rate yields fixation durations close to those in the humans, but the spillover effects are too small. Points on the line correspond individual policies (i.e. setting of the three decision thresholds). Plotted are single fixation durations by previous-frequency bin for the top 5% of policies at each noise setting, with dashed horizontal lines indicating the values in the humans. Lattice plots marked with perception and memory noise, with update noise on the x-axis. While this model recovers fixation durations close to those in the human data, it finds spillover effects about an order of magnitude too small.

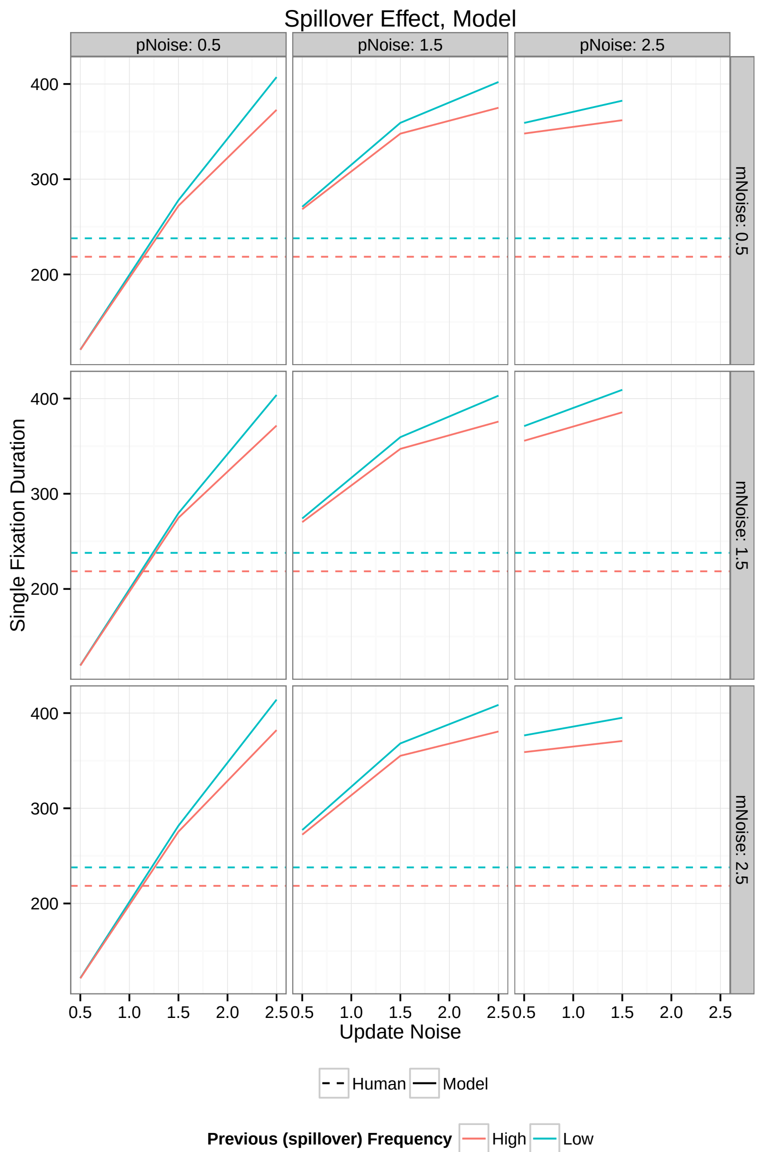

Figure 7.4: Model with fast saccade planning times yields spillover effects close to those in the humans, but overall fixation durations are too long. Points on the line correspond individual policies (i.e. setting of the three decision thresholds). Plotted are single fixation durations by previous-frequency bin for the top 5% of policies at each noise setting, with dashed horizontal lines indicating the values in the humans. Lattice plots marked with perception and memory noise, with update noise on the x-axis. This model overestimates fixation durations at which spillover is rational, but does so less than the original model. Moreover, it shows spillover effects in a larger proportion of the noises explored than the original model.

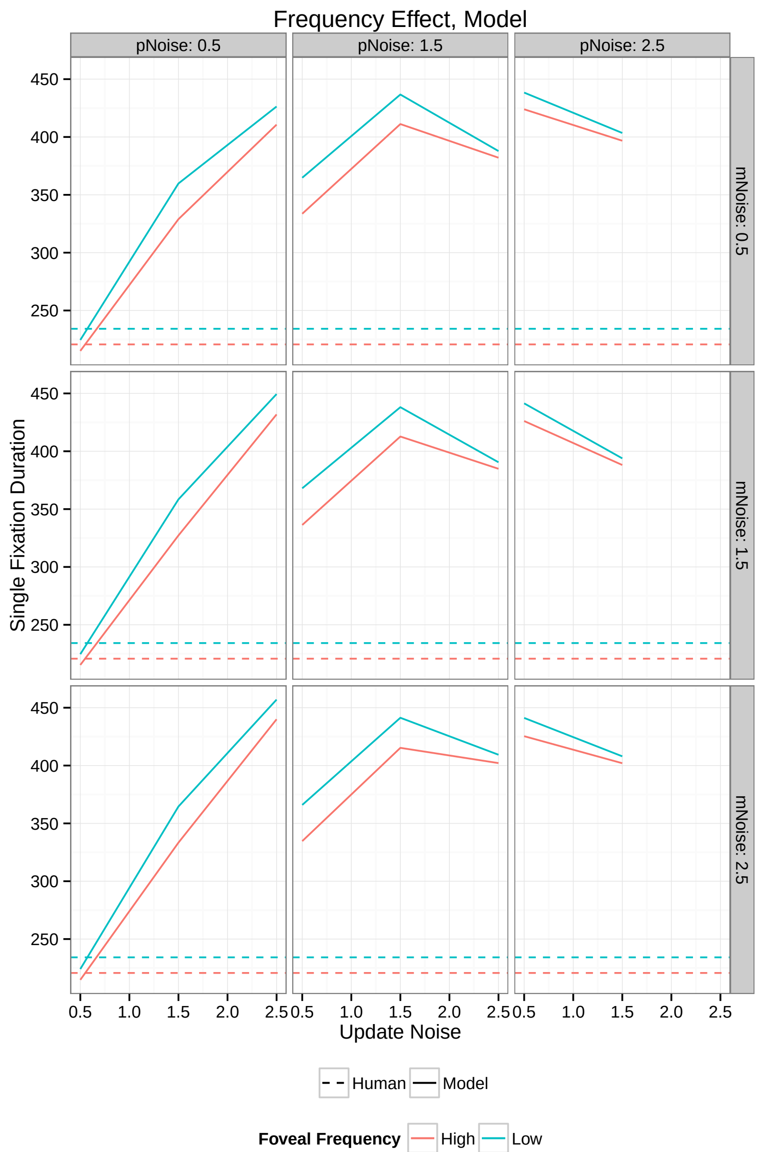

Figure 7.5: Frequency effects are present in the model across the range of free noise parameters. This is in contrast to spillover effects, which only appear at a particular setting of the noise parameters.

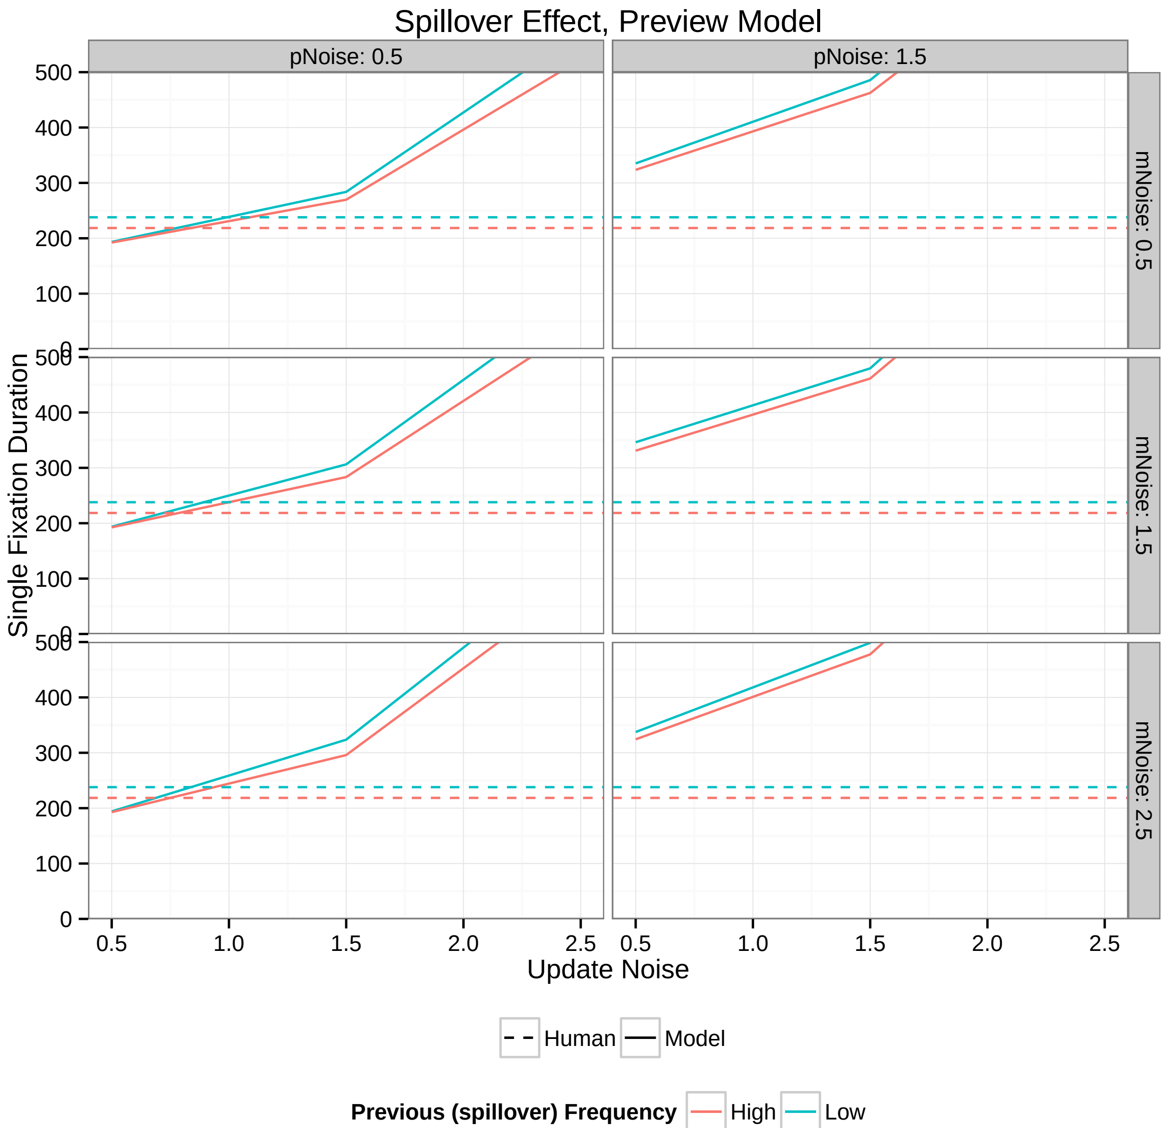

Figure 7.6: *SFD by previous-frequency bin, preview model.* Unlike the memory-only model, the preview model does succeed in recovering spillover effects close to those in the human data at similar fixation durations (indicated by the fact that the separation between the solid lines where they intersect the dashed lines). The effects are smaller than those in the humans (the separation in the solid lines is smaller than in the dashed lines).

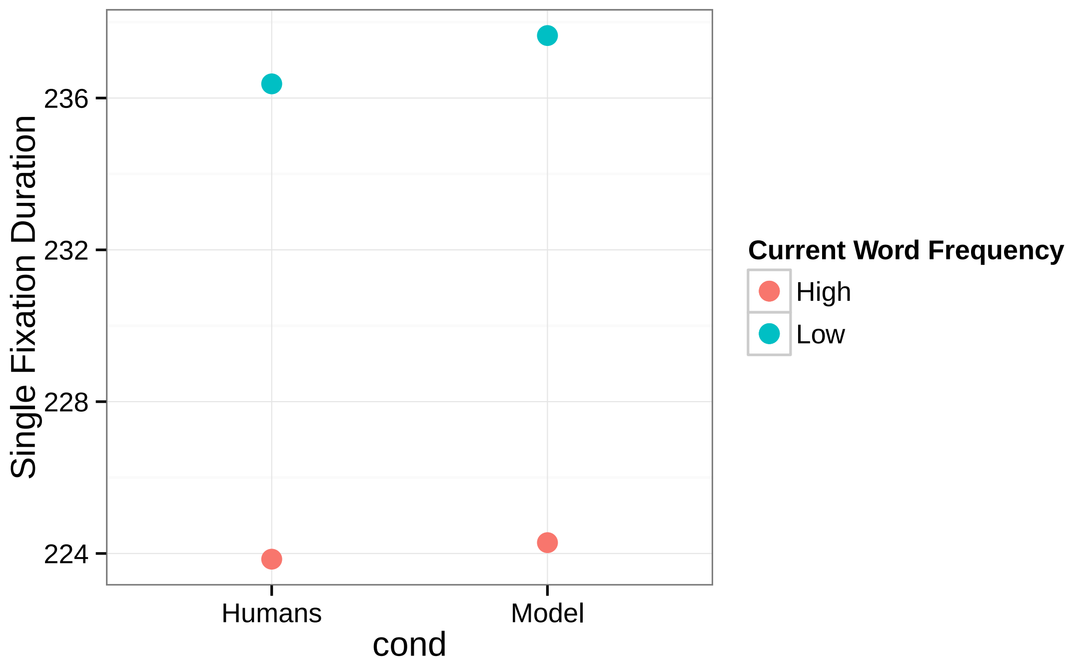

Figure 7.7: *The bounded optimality assumption excludes the spillover effects present in sub-optimal parts of the policy space.* Plotted are fixation durations by previous-frequency bin in the balanced condition in the human experiment, and in the model fit to this data. Parameter values $\sigma_p = 0.5, \sigma_m = 1.5, \sigma_u = 1.5, \theta_s = 0.449, \theta_a = 0.84999$. Balanced payoff is 2.83, as compared to 15.70 in the best performing model at this noise.

CHAPTER 8

Conclusion

This chapter discusses the key results in the thesis: the evidence for payoff adaptation in human participants and its recovery in the model, and the exploration of spillover effects derived from an adaptive tradeoff between processing foveal, parafoveal, and memory information. In addition, it outlines a few directions for future research and sketches how some of these directions might go forward.

8.1 An overview of results

The dissertation is concerned with exploring eye movement behavior from the perspective of computational rationality. It does so by introducing a simple empirical paradigm (the LLDT) that in some ways approximates full sentence-reading (e.g. proportions of single fixations, left-to-right reading, slight inter-word dependency). This simpler task allows the construction of a theory far closer to the data it is simulating than standard models in the field – i.e. the model 'does the task.'

Using this empirical paradigm and model, the dissertation demonstrates that humans adapt their reading strategies to subtle quantitative properties of the task payoff before them, adding to the body of work showing adaptation to qualitative task instructions (e.g. Grainger, 1990; Wagenmakers et al., 2008) and task structure manipulations (e.g. Wotschack, 2009). More importantly, it provides a theoretical model with the same apparent adaptive locus as the human participants (the ability to vary fixation durations) and shows that it recovers many of the same empirical phenomena seen in the humans, including subtle ones like variations in frequency effects and a serial position effect.

A surprising result that emerges from this initial empirical effort is that spillover lexical frequency effects, thought by a dominant theory to be driven by parafoveal preview, persist in the LLDT in spite of wide word spacing intended to keep preview to a minimum. The spillover effect is a longstanding benchmark phenomenon in the eye movement in reading literature in part because it is highly robust, and in part because the mapping from underlying word and sentence processing to the eye movement record is of substantial interest to related fields that use it to investigate underlying cognitive processing.

To verify that parafoveal preview is indeed not solely responsible for spillover in the LLDT, an LLDT experiment is performed with parafoveal preview gaze-contingently masked, in which spillover effects persist. In doing so, a new masking paradigm is introduced that creates masks and unmasks stimuli based on extrapolated saccade landing locations. This masking paradigm yields reading times far closer to unmasked data than the traditional boundary paradigm, suggesting that it provides less visual disruption and a more naturalistic reading experience.

With this empirical result in hand, the dissertation provides a new explanation for spillover frequency effects as arising from a delay driven by post-perceptual processing, and provides one explanation for why such delays might be rational in a particular memory architecture. It then introduces parafoveal preview into the model in an attempt to understand whether the adaptive interaction between memory-driven delays and parafoveally-driven preview might explain whether parafoveal preview is rational, and whether preview-driven spillover can arise in the LLDT. While preview does perform better across much of the speed-accuracy tradeoff curve, it does not seem to drive spillover effects directly in the LLDT, and moreover makes the memory-driven post-perceptual processing that drives preview suboptimal at many parameter settings.

Finally, quantitative millisecond-level fits against human data are attempted, with mixed results. The no-preview model does not seem to be able to simultaneously fit human-magnitude single fixation duration and spillover effects without substantial adjustments to machine parameters intended to be fixed, and even then does not provide strong quantitative fits. The preview model does provide better quantitative fits, but the improved fit seems to be driven by a choice in the memory architecture (i.e. that preview samples are not buffered, meaning fewer samples are available to memory overall) rather than by the availability of preview itself.

8.2 Future Directions

There are many interesting directions that may now be pursued. Some are already outlined in reflecting on the quantitative fit problems with fixation duration discussed in §7.4 and with accuracy discussed in §4.5. A few more ambitious ideas rely on the combination of cognitive architecture and rational analysis that computationally rational analysis provides.

8.2.1 Individual Variation

One especially promising and novel direction is building models that explain individual differences in reading strategies as bounded-optimal adaptations to individually-varying architectural constraints. This approach was demonstrated by Howes et al. (2009) in their individual differences models of the PRP (Psychological Refractory Period) task, where differences in dual-task costs were explained as signatures of near-optimal adaptations to low-level processing characteristics (stage durations and motor noise), not as differences in an underlying dual-tasking capacity. To undertake this type of analysis in the LLDT model, one might estimate the architectural parameters individually by participant, using the same kinds of tasks used to arrive at the consensus group averages used in the thesis. For example, saccade planning duration might be estimated using the countermanding paradigm, and eye-brain lag by finding the first EEG deflection to word or letter stimuli. The noise parameter would still be fit in the same way as in the thesis but individually by participant, and the model would make predictions of individual participant level performance characteristics, including saccade duration, response time, and accuracy distributions. The richer dataset of multiple architecture-performance pairings would provide both a stronger test of the theory, and more constraints for optimizing the model's free parameters. A model on this level of granularity may also reveal qualitative differences in how participants approach the task.

In addition to exploring architectural differences, it is also possible to explore effects of differences in internally modulated speed-accuracy tradeoffs—essentially differences in *intrinsic reward*. For example, Chapter 4 an increase in the size of lexical frequency effects under a payoff which emphasizes accuracy. This provides a potential theoretical connection between the findings that older adults tend to emphasize accuracy more so than young adults (Rabbitt, 1979; Smith & Brewer, 1995; Starns & Ratcliff, 2010) and also show higher lexical frequency effects on reading times (Laubrock et al., 2006). In addition, this framework may also provide a way to understand why undergraduate participants perform differently at different times of the semester (Grimm, Markman & Maddox, 2012; Nicholls, Loveless, Thomas, Loetscher & Churches, 2014), e.g. by varying the mapping between the objective reward in the task (the imposed payoff) and the intrinsic reward of the agent (the payoff that the model optimizes).

8.2.2 Sentence Reading

The LLDT as a way to understand adaptive control and eye movement behavior is interesting in its own right. As a way to understand reading and sentence processing, it is only interesting to the extent that it can be shown to be similar to more sophisticated reading tasks. Chapter 4 showed that the LLDT does show some empirical signatures of reading, such as proportions of single fixations and regressions. On the other hand, Chapter 6 suggests that the wide spacing in the LLDT largely eliminates parafoveal preview that could yield spillover – a convenient simplification, but also a departure from more closely-spaced reading behavior.

Therefore, one future direction must be scaling the model to more complex sentence-level tasks. The challenge is largely methodological and computational rather than theoretical: the Bayesian sequential sampling approach is agnostic as to both the likelihood function and the form of the belief distribution. Recall that the string-level belief update in the thesis computes two quantities: $P(S^k = s_i|e^k, T = t)$, i.e. the posterior probability that the string in position k is some string s_i conditioned on incoming evidence from position k and the multinomial T corresponding to trial type; and $P(T = t|e^k)$, i.e. the posterior probability that the trial is of a particular type (word or nonword) given incoming evidence. T could instead be a probability distribution over sentence structures (i.e. a grammar), and the theory becomes – in the abstract – a theory of sentence processing.

Practically speaking, the extension is less simple. On the computational front the challenges are in specifying a likelihood function that lets words of different lengths be confused for each other, maintaining and updating a belief distribution over potentially infinitely many sentence structures, and handling a broader set of actions than only saccading forward. On the methodological front, the challenge is in finding a sentence-reading task with a decision variable that can be computed from the distribution over sentence structures. These challenges are briefly considered below.

Handling words of different length The letter-level unit-basis vector representation of strings in the thesis model does not enable a likelihood function that confuses strings of different lengths. This is because the length of the vector implicitly provides noiseless information about the length of the string. An alternative is needed that can represent uncertainty over string length. One such alternative is a weighted finite state transducer (wFST) inspired by Levenshtein distance. Levenshtein distance is a string distance metric that counts the number of insertions, deletions, and substitutions needed to transform one string into another. A wFST probabilistically transforms its input 'tape' into an output 'tape'. One can define a wFST that takes in a sequence of letters and probabilistically inserts, deletes, or transposes letters to generate another sequence.

Such a transducer will probabilistically transform some true word being viewed into a sequence of letters. The likelihood of a particular word given this sequence of letters is the probability of generating the sequence from that particular word using that wFST. To compute it, one composes the deterministic finite state automaton representing the sample (letter sequence) with the reversed sample generation wFST. The resultant wFST represents a probability distribution over all the strings that could have generated the noisy sample (whether a word in the lexicon or not). As a side benefit, moving to this formulation of the likelihood function makes its parameters possible to estimate from data, for example from letter confusability in a speeded recognition task.

Maintaining and updating belief over many structures A model that can process sentences will need to maintain and update its belief distribution over sentence structures. Exhaustively listing all possible structures for all but trivial problems is difficult, and becomes impossible when any recursion is introduced and the grammar produces an infinite number of strings. The computational linguistics and natural language processing literature provides a number of solutions to this problem, however.

One option, and the one chosen by Bicknell & Levy (2010a) in their sequential sampling model used to understand regressions, is to represent the grammar as a weighted finite state automaton (wFSA). They induce their grammar by counting bigrams in a large corpus. The primary advantage of this approach is that it is fast, easy to estimate from data, and can have broad coverage. The chief disadvantage is that wFSAs make it difficult to engage with much of modern (psycho)linguistic theory, which primarily works with context-free and mildly-context-sensitive grammars.

This leads to a second option, which is to use the state of an incremental parser to maintain the current posterior probability over sentence structure. One such parser is the Earley-Stolcke parser (Earley, 1970; Stolcke, 1995), which parses probabilistic context-free grammars (PCFGs) and has previously been used to predict expectation-based sentence processing difficulty Hale (2001). In a footnote, Stolcke notes that the input-recognition *scanning* action of the parser could be extended to probabilistic input. Because the Stolcke parser is a dynamic programming parser (i.e. keeps track of partial derivations), it may be possible to iterate the scanning step (this would be the posterior update), normalizing correctly at each timestep. This option would be slower than intersecting wFSAs but would allow use of a hand-written PCFG to investigate psycholinguistically-interesting phenomena. The disadvantages of this approach are that it will be slower than using wFSAs, and that while a PCFG is a more expressive formalism than a wFSA, still cannot represent concepts like movement, and other state-of-the-art aspects of linguistic theory.

A third option relies on a result from Bar-Hillel, Perles & Shamier (1961), who showed a method for intersecting a context-free language with a regular language, and that such an intersection is again context-free. Nederhof & Satta (2003) extended the result to a probabilistic grammar and probabilistic finite state automaton, and Chen, Hunter, Yun & Hale (2014) provide an implemented tool that intersects wFSAs with a weighted multiple context-free grammar (WMCFG), a class of mildly context-sensitive grammars that can be derived from minimalist grammars (Stabler, 1997). By iteratively intersecting wFSAs representing samples with the grammar, this toolchain should allow sequential sampling with a probabilistic minimal grammar as the representation of the belief distribution over sentence structures. But the method is very slow: informal testing shows that it is about two orders of magnitude slower than using the Earley-Stolcke parser, itself slower than the wFSA intersection method or the method in the thesis.

Tasks and decision variables in sentence-reading With the question of the belief representation and update tackled using one of the methods above, what remains is the set of decision variables.

For the eye movement decision, if the action is to only move forward, then the decision can still be conditioned on the max a posteriori string probability (as in Chapter 6). But Bicknell & Levy (2010a) showed that agents that can also move their eyes back to previous words perform better at any speed-accuracy tradeoff than agents that cannot. They used threshold decisions for both the forward and backward eye movement actions. This larger policy space may work at a first approximation, but may also turn out to be suboptimal. The general problem in both the LLDT and sentence-reading is that of choosing which sensor to attend to when there is a time cost of switching sensors (i.e. the saccade itself), and neither the 'always go forward' nor the 'go back when confidence drops' policy space may contain the policy optimal in this setting. While any model with such a richer policy space could be compared to human data using the same position-based measures used in the thesis, the inclusion of movements to words other than the next one would also motivate the inclusion of novel scanpath-based analyses (von der Malsburg & Vasishth, 2011) and in particular explicitly computing similarity scores between the model's scanpaths and those of human participants.

For the trial-level decision, the model outlined above could support any decision variable that can be computed from a probability distribution over sentence structures. Judgments about structure like subject-verb relations ('was it the feisty ballerina who jumped?') or anaphor coindexing ('is *himself* referring to the tall milkman or the stout wrestler?') could be computed at each timestep and thersholded. A criticism of such a task may be that it does not tap into naturalistic reading, which likely involves discriminating between meanings rather than merely between structures. But the same criticism holds with respect to most psycholinguistic tasks, and in fact the verb relation question variant looks remarkably like a standard comprehension question in a psycholinguistic sentence-reading experiment). A notable exception to this characterization of psycholinguistic tasks is recent work on sentence processing during story comprehension (Brennan, Nir, Hasson, Malach, Heeger & Pylkkänen, 2012).

Moreover, in the computationally rational perspective the criticism about dissimilarity between the experimental task used and naturalistic reading is not necessarily valid because the notion of 'normal reading' is ill-formed. If reading strategies vary with task goals (as Chapter 4 and other work shows) then the goal of a theory of reading is not to explain some abstract platonic notion of reading, but to build a theory of the way reading strategy is adapted to the joint constraints of agent and task. A theory of a sentence structure query task is very much on the path towards that goal.

8.2.3 Adaptive control of attention

The model in the dissertation assumes that the attentional spotlight only applies to a single word at a time, in light of the wide spacing of the LLDT. In practice, a sequential sampling framework can rationally integrate information arriving from multiple words and straightforwardly encode multiple perceptual streams in this way (as done by Bicknell & Levy, 2010a, though see Reichle et al., 2009 for another opinion). However, these previous efforts on parallel information integration – both the

Bayesian work by Bicknell and colleagues, and the processing functions of models like Glenmore (Reilly & Radach, 2006) and SWIFT (Engbert et al., 2005) – all assume some fixed, asymmetric parallel processing function. This is consistent with the finding that the perceptual span is asymmetric in the direction of reading (McConkie & Rayner, 1976), but puzzling given that foveal acuity drop-off is symmetric.

An interesting direction of future research is an attempt to understand why such a processing function arises in spite of the symmetric acuity function of the fovea. Such a model would be a generalization of the model in Chapter 6 that can adaptively allocate a gradient attention beam over fovea and parafovea rather than only allocate it sequentially word by word. The goal would be to understand the asymmetry of the perceptual span, and integrate this understanding with a theory of adaptive saccade targeting of the kind Legge et al. (2002) and Bicknell & Levy (2010b) were interested in. Such a theory might also have some bearing on why n+2 preview effects and n+1 parafoveal-on-foveal effects appear to be fragile in the literature (Kennedy et al., 2002; Kennedy & Pynte, 2005; Kliegl et al., 2007; Angele & Rayner, 2011).

8.3 A final word

The thesis work is a set of first steps in a theory of eye movement control in reading that assumes a key role for the oculomotor machine humans possess, the task environments they find themselves in, and their apparent ability to adapt their behavior to the above. Already many challenges are seen in the effort, both in finding good quantitative fits and in having a deep understanding of which aspects of architecture and task drive particular aspects of behavior.

Future work building on these first steps will face these formidable computational and empirical challenges. But in light of the ubiquity of adaptive effects at all levels of human performance, these are challenges to any approach to understanding cognition. The combination of rational analysis and mechanistic architecture modeling has the strong advantage of putting the role of architecture and adaptation front-and-center, and facing these challenges head-on.

APPENDIX A

Mathematical details of the sequential sampler

Here is the MSPRT as applied to the LLDT decisions. First, consider the case of reading a single isolated word (analogous to Norris' Bayesian reader). Let $\mathcal{S}$ denote the set of strings in the lexicon. Let S denote the random variable that can take on values in $\mathcal{S}$, and let s_i denote a particular realization of S. Let each string in $\mathcal{S}$ be represented using unit-basis indicator coding (i.e. a vector of size 26 with a 1 in the position corresponding to the given letter and 0 elsewhere). Each string can be represented as a concatenation of such vectors – or more conveniently, as a $c \times a$ matrix where c is the length of the string (4 characters) and a is the length of the alphabet (26 characters). Let μ denote this matrix for some particular word $s \in \mathcal{S}$, such that $\mu_{pq} = 1$ if the p^{th} position in s has the q^{th} letter, otherwise $\mu_{pq} = 0$.

A noisy evidence sample e_t is created at time t by adding mean-zero Gaussian noise with some standard deviation σ to each cell in this matrix. Since the samples are independent and identically distributed (i.i.d.), the indexing by time is omitted going further. The setting of σ is discussed in the main text of the dissertation. Based on a stream of such noisy samples, the best that an ideal observer can do is integrate the noisy evidence using Bayes' rule:

$$P(S = s_i|e) = \frac{P(e|S = s_i)P(S = s_i)}{P(e)} \tag{A.1}$$

The model assumes that the reader has a good enough implicit model of their visual system to use the sample generation function to compute the likelihood term $P(e|S = s_i)$:

$$P(e|S = s_i) = \prod_{p,q} f(e_{pq}; \mu_{pq}, \sigma^2) \tag{A.2}$$

where $f(x; \mu, \sigma^2)$ is the probability density function of the Gaussian distribution with mean μ and standard deviation σ.

The prior $P(S = s_i)$ for the first update is set by normalizing the set of frequency counts in the lexicon (i.e. dividing each by the total sum to map from counts to probabilities). Then the posterior becomes the prior and the update repeats for the next sample. As in many practical problems, there is no way to estimate $P(e)$ directly but one can rely on the fact that the posterior probability over all possible values of S must sum to 1, compute the posterior for all all strings in $\mathcal{S}$, and divide by the sum to normalize.

Now, consider the list-reading case. For a list of length l there are l variables $S^k, k = 1, 2, \ldots, l$, each taking on some realization s_i from $\mathcal{S}$. One might expect to perform the updates above independently for each word as it is fixated – but this would ignore an interdependence introduced by the constraints on trial construction, namely that at most a single nonword can appear and that words are not repeated. The simplest assumption is that readers are able to be sensitive to both constraints, requiring that both be included in the update calculation. However, obeying the latter constraint in the update would substantially complicate the computation to no clear gain: with large lexicons, its practical effect would be minimal. The former constraint is somewhat more important: with only two possible outcomes (word and nonword), its effect might be more substantial. The correct update would ensure that as the probability of a nonword in some position k goes up, the probability mass over nonwords in other positions $j \neq k$ must go down.

One can address this on the level of the trial: let T denote the categorical distribution over trial types and t its realization. The possible types of t are an all-words trial $T = w$), and a nonword trial, subdivided into possible nonword positions ($T = n^k, k \in 1..l$). One can decompose the possibilities in T into whether the nonword is at the current position, $T = n^k$, or it is not, $T \neq n^k$. This allows the decomposing of the update into two components:

$$P(S^k = s_i|e^k, T = t) = P(S^k = s_i|e^k, T = n^k) + P(S^k = s_i|e^k, T \neq n^k) \tag{A.3}$$

Notice that in each of these only part of the support of S^k need be considered: conditioned on $T = n^k$, s_i can only be one of the nonwords and conditioned on $T \neq n^k$ it can only be one of the words. Put differently, one of the two right hand side values above is zero, for any s_i. Let n_m denote some nonword realization of S^k and w_o denote some nonword realization of S^k. Therefore, the equation above can be equivalently written as:

$$P(S^k = s_i|e^k, T = t) = P(S^k = n_m|e^k, T = n^k) + P(S^k = w_o|e^k, T \neq n^k) \tag{A.4}$$

Computing the Bayes update for both:

$$P(S^k = n_m|e^k, T = n^k) = \frac{P(e^k|S^k = n_m, T = n^k)P(S^k = n_m|T = n^k)}{P(e^k|T = n^k)} \tag{A.5}$$

$$P(S^k = w_o|e^k, T \neq n^k) = \frac{P(e^k|S^k = w_o, T \neq n^k)P(S^k = w_o|T \neq n^k)}{P(e^k|T \neq n^k)} \tag{A.6}$$

But $P(S^k = n_m, T = n^k) = P(S^k = n_m)$ by definition, since string S^k must be a nonword if the trial has a nonword in position k, and likewise $P(S^k = w_o, T \neq n^k) = P(S^k = w_o)$. In addition, recall that the set of possible words is the set of exhaustive and exclusive outcomes conditioned on the string being a word and likewise for nonwords, so one can again replace the denominator with the sum over possible numerators. Combining the two:

$$\begin{aligned}
P(S^k = n_m|e^k, T = n^k) &= \frac{P(e^k|S^k = n_m)P(S^k = n_m|T = n^k)}{P(e^k|T = n^k)} && \text{(A.7)}\\
&= \frac{P(e^k|S^k = n_m)P(S^k = n_m|T = n^k)}{\sum_m P(e^k|S^k = n_m)P(S^k = n_m|T = n^k)} && \text{(A.8)}\\
P(S^k = w_o|e^k, T \neq n^k) &= \frac{P(e^k|S^k = w_o)P(S^k = w_o|T \neq n^k)}{P(e^k|T \neq n^k)} && \text{(A.9)}\\
&= \frac{P(e^k|S^k = w_o)P(S^k = w_o|T \neq n^k)}{\sum_o P(e^k|S^k = w_o)P(S^k = w_o|T \neq n^k)} && \text{(A.10)}
\end{aligned}$$

Since each string s_i is either a nonword n_m or word w_o, the likelihood term is now equivalent to that in Equation A.1. The prior terms $P(S^k = n_m|T = n^k)$ and $P(S^k = w_o|T \neq n^k)$ imply keeping track of the word and nonword belief distributions separately, conditioned on the trial possibilities, but are otherwise also straightforwardly handled.

The model must also simultaneously update the trial-level belief (i.e. its belief over the outcome of T):

$$P(T = t|e^k) = \frac{P(e^k|T = t)P(T = t)}{P(e^k)} \tag{A.11}$$

As above, it is useful to decompose t into relevant possibilities $T = n^k$ and $T \neq n^k$. But $T \neq n^k$ actually corresponds to two very different possibilities with respect to the LLDT: the possibility that the trial is a word trial $T = w$ and the possibility that the trial is a nonword trial, which just happens to have the nonword in some other position $T = n^{j \neq k}$.

Therefore:

$$P(T = t|e^k) = \begin{cases} \frac{P(e^k|T=w)P(T=w)}{P(e^k)} & \text{if } t = w \\ \frac{P(e^k|T=n^k)P(T=n^k)}{P(e^k)} & \text{if } t = n^k \\ \frac{P(e^k|T=n^j)P(T=n^j)}{P(e^k)} & \text{if } t = n^j, j \neq k \end{cases} \tag{A.12}$$

The denominator term can be decomposed:

$$P(e^k) = P(e^k|T = n^k)P(T = n^k) + P(e^k|T \neq n^k)P(T \neq n^k) \tag{A.13}$$

Conveniently, the conditioned terms are the normalization terms from equation A.6, and the others are the prior for T.

Next, the likelihood terms are handled slightly differently for each possibility. For $T = n^k$, the likelihood is again the denominator of equation A.6, expanded in equation A.10. For the other two, the key insight is that words are selected identically in word and nonword trials, and therefore that $P(e^k|T \neq n^k) = P(e^k|T = w) = P(e^k|T = n^j), j \neq k$, and again the likelihood $P(e^k|T \neq n^k)$ is available in the denominator of equation A.6. The final form is therefore:

$$P(T=t|e^k) = \begin{cases} \frac{P(e^k|T=w)P(e^k|T\neq n^k)}{P(e^k)} & \text{if } t=w \\ \frac{P(e^k|T=n^k)P(T=n^k)}{P(e^k)} & \text{if } t=n^k \\ \frac{P(e^k|T\neq n^k)P(T=n^j)}{P(e^k)} & \text{if } t=n^j, j\neq k \end{cases} \tag{A.14}$$

BIBLIOGRAPHY

Anderson, J. R. (1990). *The adaptive character of thought.* Routledge.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, *14*, 471–517.

Anderson, J. R. & Matessa, M. (1990). A Rational Analysis of Categorization. In *Proceedings of the Seventh International Conference (1990) on Machine Learning*, (pp. 76–84)., San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Anderson, T. W. (1960). A Modification of the Sequential Probability Ratio Test to Reduce the Sample Size. *The Annals of Mathematical Statistics*, *31*(1), 165–197.

Andrews, S. (1997). The effect of orthographic similarity on lexical retrieval: Resolving neighborhood conflicts. *Psychonomic Bulletin & Review*, *4*(November 1995), 439–461.

Angele, B. & Rayner, K. (2011). Parafoveal processing of word n + 2 during reading: do the preceding words matter? *Journal of experimental psychology. Human perception and performance*, *37*(4), 1210–20.

Ballard, D. H. & Hayhoe, M. M. (2009). Modelling the role of task in the control of gaze. *Visual cognition*, *17*(6-7), 1185–1204.

Balota, D. A., Pollatsek, A., & Rayner, K. (1985). The interaction of contextual constraints and parafoveal visual information in reading. *Cognitive Psychology*, *17*(3), 364–390.

Bar-Hillel, Y., Perles, M., & Shamier, E. (1961). On formal properties of simple phrase structure grammars. *Zeitschrift für Phonetik, Sprachwissenschaft und Kommunikationsforschung*, *14*, 143–172.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting Linear Mixed-Effects Models using lme4. *Journal of Statistical Software ...*, 51.

Baum, C. & Veeravalli, V. V. (1994). A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, *40*(6), 1994–2007.

Becker, W. & Jürgens, R. (1979). An analysis of the saccadic system by means of double step stimuli. *Vision research*, *19*(9), 967–83.

Bicknell, K. & Levy, R. (2010a). A rational model of eye movement control in reading. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, (pp. 1168–1178). Association for Computational Linguistics.

Bicknell, K. & Levy, R. (2010b). Rational eye movements in reading combining uncertainty about previous words with contextual probability. *32nd Annual Conference of the Cognitive Science Society.*

Bicknell, K. & Levy, R. (2014). The mind leads the eyes: ungrammaticality detection from two words back in reading. In *26th annual CUNY conference on human sentence processing.*

Blanchard, H. E., Pollatsek, A., & Rayner, K. (1989). The acquisition of parafoveal word information in reading. *Perception & psychophysics*, *46*(1), 85–94.

Bogacz, R. & Gurney, K. N. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, *19*(2), 442–77.

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychological review*, *114*(2), 376–97.

Bouma, H. (1973). Visual interference in the parafoveal recognition of initial and final letters of words. *Vision Research*, *13*(4), 767–782.

Brennan, J., Nir, Y., Hasson, U., Malach, R., Heeger, D. J., & Pylkkänen, L. (2012). Syntactic structure building in the anterior temporal lobe during natural story listening. *Brain and language*, *120*(2), 163–73.

Chater, N. & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in cognitive sciences*, *3*(2), 57–65.

Chen, Z., Hunter, T., Yun, J., & Hale, J. (2014). Modeling sentence processing difficulty with a conditional probability calculator. In *Proceedings of the 36th annual cognitive science conference*, (pp. 1856–1857).

Clark, V. P., Fan, S., & Hillyard, S. A. (1994). Identification of early visual evoked potential generators by retinotopic and topographic analyses. *Human Brain Mapping*, *2*(3), 170–187.

Cook, E. P. & Maunsell, J. H. R. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature neuroscience*, *5*(10), 985–94.

Dehaene, S. (1993). Temporal oscillations in human perception. *Psychological Science*, *4* (4), 264–270.

Deutsch, A. & Rayner, K. (1999). Initial Fixation Location Effects in Reading Hebrew Words. *Language and Cognitive Processes*, *14* (4), 393–421.

Dragalin, V., Tartakovsky, A., & Veeravalli, V. (2000). Multihypothesis sequential probability ratio tests. II. Accurate asymptotic expansions for the expected sample size. *IEEE Transactions on Information Theory*, *46* (4), 1366–1383.

Draglia, V., Tartakovsky, A., & Veeravalli, V. (1999). Multihypothesis sequential probability ratio tests .I. Asymptotic optimality. *IEEE Transactions on Information Theory*, *45* (7), 2448–2461.

Drieghe, D., Rayner, K., & Pollatsek, A. (2008). Mislocated fixations can account for parafoveal-on-foveal effects in eye movements during reading. *Quarterly journal of experimental psychology (2006)*, *61* (8), 1239–49.

Earley, J. (1970). An efficient context-free parsing algorithm. *Communications of the ACM*, *13* (2), 94–102.

Edwards, W. (1965). Optimal Strategies for Seeking information : Models for Statistics , Choice Reaction Times , and Human Information Processing. *Journal of Mathematical Psychology*, *329*, 312–329.

Elliott, M. a. & Müller, H. J. (2000). Evidence for 40-Hz oscillatory short-term visual memory revealed by human reaction-time measurements. *Journal of experimental psychology. Learning, memory, and cognition*, *26* (3), 703–18.

Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: a dynamical model of saccade generation during reading. *Psychological review*, *112* (4), 777–813.

Erlhagen, W. & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109* (3), 545–572.

Feng, G. (2003). Throwing the baby out with the bathwater: Problems in modeling aggregated eye-movement data [Commentary on Reichle, Rayner & Pollatsek 2003]. *Behavioral and Brain Sciences*, *26* (4), 482–483.

Findlay, J. M. & Walker, R. (1999). A model of saccade generation based on parallel processing and competitive inhibition. *The Behavioral and brain sciences*, *22* (4), 661–74; discussion 674–721.

Fodor, J. A. (1983). *The Modularity of Mind: An Essay on Faculty Psychology.* MIT Press.

Forster, K. I. (1979). Levels of processing and the structure of the language processor. In *Sentence processing: Psycholinguistic studies presented to Merrill Garrett*, (pp. 27–85). Hillsdale, NJ.

Frazier, L. & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, *14*(2), 178–210.

Giraud, A.-L. & Poeppel, D. (2012). Cortical oscillations and speech processing: emerging computational principles and operations. *Nature neuroscience*, *15*(4), 511–7.

Gold, J. I. & Shadlen, M. N. (2007). The neural basis of decision making. *Annual review of neuroscience*, *30*, 535–74.

Grainger, J. (1990). Word frequency and neighborhood frequency effects in lexical decision and naming. *Journal of Memory and Language*, *29*(2), 228–244.

Green, D. M. & Swets, J. A. (1966). Signal detection theory and psycho-physics.

Grimm, L. R., Markman, A. B., & Maddox, W. T. (2012). End-of-semester syndrome: How situational regulatory fit affects test performance over an academic semester. *Basic and applied social psychology*, *34*(4), 376–385.

Hale, J. T. (2001). A probabilistic Earley parser as a psycholinguistic model. *Proceedings of the second meeting of North American Chapter of the Association for Computational Linguistics*.

Hale, J. T. (2011). What a Rational Parser Would Do. *Cognitive Science*, *35*(3), 399–443.

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *31*(17), 6339–52.

Henderson, J. M. & Ferreira, F. (1990). Effects of foveal processing difficulty on the perceptual span in reading: Implications for attention and eye movement control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(3), 417–429.

Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. *Psychological review*, *116*(4), 717–51.

Howes, A., Vera, A., & Lewis, R. L. (2006). Bounding Rational Analysis: Constraints on Asymptotic Performance. In W. D. Gray (Ed.), *Integrated models of cognitive systems*, number January (pp. 403–413). New York, NY: Oxford University Press, USA.

Howes, D. H. & Solomon, R. L. (1951). Visual duration threshold as a function of word-probability. *Journal of Experimental Psychology*, *41*(6), 401–410.

Inhoff, A. W. & Rayner, K. (1986). Parafoveal word processing during eye fixations in reading: Effects of word frequency. *Perception & Psychophysics*, *40*(6), 431–439.

Inhoff, A. W., Starr, M., & Shindler, K. L. (2000). Is the processing of words during eye fixations in reading strictly serial? *Perception & psychophysics*, *62*(7), 1474–84.

Just, M. A. & Carpenter, P. A. (1980). A theory of reading: from eye fixations to comprehension. *Psychological review*, *87*(4), 329–54.

Kennedy, A. & Pynte, J. (2005). Parafoveal-on-foveal effects in normal reading. *Vision research*, *45*(2), 153–68.

Kennedy, A., Pynte, J., & Ducrot, S. (2002). Parafoveal-on-foveal interactions in word recognition. *The Quarterly journal of experimental psychology. A, Human experimental psychology*, *55*(4), 1307–37.

Kennison, S. M. & Clifton, C. (1995). Determinants of parafoveal preview benefit in high and low working memory capacity readers: implications for eye movement control. *Journal of experimental psychology. Learning, memory, and cognition*, *21*(1), 68–81.

Kieras, D. E., Wood, S. D., & Meyer, D. E. (1997). Predictive engineering models based on the EPIC architecture for a multimodal high-performance human-computer interaction task. *ACM Transactions on Computer-Human Interaction*, *4*(3), 230–275.

Kliegl, R., Risse, S., & Laubrock, J. (2007). Preview benefit and parafoveal-on-foveal effects from word n + 2. *Journal of experimental psychology. Human perception and performance*, *33*(5), 1250–5.

Kučera, H. & Francis, W. N. (1967). *Computational Analysis of Present Day American English*. Brown University Press.

Laubrock, J., Kliegl, R., & Engbert, R. (2006). SWIFT explorations of age differences in eye movements during reading. *Neuroscience and biobehavioral reviews*, *30*(6), 872–84.

Legge, G. E., Hooven, T. a., Klitz, T. S., Mansfield, J. S., & Tjan, B. S. (2002). Mr. Chips 2002: new insights from an ideal-observer model of reading. *Vision research*, *42*(18), 2219–234.

Legge, G. E., Klitz, T. S., & Tjan, B. S. (1997). Mr. Chips: an ideal-observer model of reading. *Psychological review*, *104*(3), 524–53.

Lehky, S. R. (2010). Decoding Poisson Spike Trains by Gaussian Filtering. *Neural Computation*, *22*(5), 1245–1271.

Levy, R., Bicknell, K., Slattery, T., & Rayner, K. (2009). Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(50), 21086–90.

Lewis, R. L., Howes, A., & Singh, S. (2013). Computational Rationality: Linking Mechanism and Behavior through Bounded Utility Maximization. *Topics in cognitive science*.

Lewis, R. L., Shvartsman, M., & Singh, S. (2013). The adaptive nature of eye movements in linguistic tasks: how payoff and architecture shape speed-accuracy tradeoffs. *Topics in cognitive science*, *5*(3), 581–610.

Lewis, R. L. & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, *29*, 375–419.

Lewis, R. L., Vasishth, S., & Van Dyke, J. (2006). Computational principles of working memory in sentence comprehension. *Trends in Cognitive Science*, *10*, 447–454.

Liu, Y. & Reichle, E. D. (2010). The emergence of adaptive eye movements in reading. In Ohlsson, S. & Catrambone, R. (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, number 2, (pp. 1136–1141).

Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin*, *81*(12), 899–917.

McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*(5), 375–407.

McConkie, G. W. & Rayner, K. (1975). The span of the effective stimulus during a fixation in reading. *Perception & Psychophysics*, *17*(6), 578–586.

McConkie, G. W. & Rayner, K. (1976). Asymmetry of the perceptual span in reading. *Bulletin of the Psychonomic Society*, *8*(5), 365–368.

McConkie, G. W., Rayner, K., & Wilson, S. J. (1973). Experimental manipulation of reading strategies. *Journal of Educational Psychology*, *65*(1), 1–8.

Medler, D. A. & Binder, J. R. (2005). MCWord: An On-Line Orthographic Database of the English Language.

Meyer, D. E. & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms. *Psychological Review*, *104*(1), 3–65.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1974). Functions of graphemic and phonemic codes in visual word-recognition. *Memory & cognition*, *2*(2), 309–21.

Morrison, R. E. (1984). Manipulation of stimulus onset delay in reading: evidence for parallel programming of saccades. *Journal of experimental psychology. Human perception and performance*, *10*(5), 667–82.

Morrison, R. E. & Rayner, K. (1981). Saccade size in reading depends upon character spaces and not visual angle. *Perception & psychophysics*, *30*(4), 395–396.

Mouchetant-Rostaing, Y., Giard, M. H., Bentin, S., Aguera, P. E., & Pernier, J. (2000). Neurophysiological correlates of face gender processing in humans. *The European journal of neuroscience*, *12*(1), 303–10.

Nederhof, M.-J. & Satta, G. (2003). Probabilistic Parsing as Intersection. In *8th International Workshop on Parsing Technologies*, volume 0, (pp. 137–148).

Newell, A. (1973). You can't play 20 questions with nature and win. In W. G. Chase (Ed.), *Visual Information Processing* (pp. 283–308). New York: Academic Press.

Nicholls, M. E. R., Loveless, K. M., Thomas, N. a., Loetscher, T., & Churches, O. (2014). Some participants may be better than others: Sustained attention and motivation are higher early in semester. *The Quarterly Journal of Experimental Psychology*, (September 2014), 1–9.

Norris, D. (2006). The Bayesian reader: explaining word recognition as an optimal Bayesian decision process. *Psychological review*, *113*(2), 327–357.

Norris, D. (2009). Putting it all together: a unified account of word recognition and reaction-time distributions. *Psychological review*, *116*(1), 207–19.

Oaksford, M. & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608–631.

O'Regan, J. K. & Lévy-Schoen, A. (1987). Eye-movement strategy and tactics in word recognition and reading. In M. Coltheart (Ed.), *Attention and Performance 12: The psychology of reading* (pp. 363–383). Hillsdale, NJ: Lawrence Erlbaum Associates.

Pinheiro, J. C. & Bates, D. M. (2000). *Mixed Effects Models in S and S-Plus*. Springer.

Pitt, M. a. & Myung, I. J. (2002). When a good fit can be bad. *Trends in cognitive sciences*, *6*(10), 421–425.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, *109*(3), 472–491.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Rabbitt, P. (1979). How old and young subjects monitor and control responses for accuracy and speed. *British Journal of Psychology*, *70*(2), 305–311.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychological review*, *111*(1), 159–82.

Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, *7*(1), 65–81.

Rayner, K. (1978). Eye movements in reading and information processing. *Psychological bulletin*, *85*(3), 618–60.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124*(3), 372–422.

Rayner, K. & Bertera, J. (1979). Reading without a fovea. *Science*, *206*(4417), 468–469.

Rayner, K. & Duffy, S. (1986). Lexical complexity and fixation times in reading: effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*, *14*, 191–201.

Rayner, K. & Fischer, M. H. (1996). Mindless reading revisited: eye movements during reading and scanning are different. *Perception & psychophysics*, *58*(5), 734–47.

Rayner, K. & Raney, G. E. (1996). Eye movement control in reading and visual search: Effects of word frequency. *Psychonomic bulletin & review*, *3*(2), 245–8.

Rayner, K., Slowiaczek, M. L., Clifton, C., & Bertera, J. H. (1983). Latency of sequential eye movements: implications for reading. *Journal of experimental psychology. Human perception and performance*, *9*(6), 912–22.

Reichle, E. D. & Laurent, P. a. (2006). Using reinforcement learning to understand the emergence of "intelligent" eye-movement behavior during reading. *Psychological review*, *113*(2), 390–408.

Reichle, E. D., Liu, Y., & Laurent, P. (2011). The emergence of adaptive eye movement control in reading: Theory and data. *Studies of Psychology and ...*, *2*(3).

Reichle, E. D., Liversedge, S. P., Pollatsek, A., & Rayner, K. (2009). Encoding multiple words simultaneously in reading is implausible. *Trends in cognitive sciences*, *13*(3), 115–9.

Reichle, E. D., Pollatsek, A., & Rayner, K. (2006). EZ Reader: A cognitive-control, serial-attention model of eye-movement behavior during reading. *Cognitive Systems Research*, *7*(1), 4–22.

Reichle, E. D., Pollatsek, A., & Rayner, K. (2012). Using E-Z Reader to simulate eye movements in nonreading tasks: a unified framework for understanding the eye-mind link. *Psychological review*, *119*(1), 155–85.

Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The E-Z Reader model of eye-movement control in reading: Comparisons to other models. *Behavioral and Brain Sciences*, *26*(04), 445–76; discussion 477–526.

Reichle, E. D., Reineberg, A. E., & Schooler, J. W. (2010). Eye movements during mindless reading. *Psychological science*, *21*(9), 1300–10.

Reichle, E. D., Warren, T., & McConnell, K. (2009). Using E-Z Reader to model the effects of higher level language processing on eye movements during reading. *Psychonomic bulletin & review*, *16*(1), 1–21.

Reilly, R. G. & O'Regan, J. K. (1998). Eye Movement Control During Reading: A Simulation of Some Word-targeting Strategies. *Vision Research*, *38*(2), 303–317.

Reilly, R. G. & Radach, R. (2006). Some empirical tests of an interactive activation model of eye movement control in reading. *Cognitive Systems Research*, *7*(1), 34–55.

Richter, E. M., Engbert, R., & Kliegl, R. (2006). Current advances in SWIFT. *Cognitive Systems Research*, *7*(1), 23–33.

Roberts, S. & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, *107*(2), 358–367.

Rothkopf, C. A., Ballard, D. H., & Hayhoe, M. M. (2007). Task and context determine where you look. *Journal of vision*, *7*(14), 16.1–20.

Russell, S. J. & Subramanian, D. (1995). Provably Bounded-Optimal Agents. *Journal of Artificial Intelligence Research*, *2*, 575–609.

Salverda, A. P., Brown, M., & Tanenhaus, M. K. (2011). A goal-based perspective on eye movements in visual world studies. *Acta psychologica*, *137*(2), 172–80.

Salvucci, D. D. (2001). An integrated model of eye movements and visual encoding. *Cognitive Systems Research*, *1*(4), 201–220.

Sanders, C. A. (2013). *Computational Models of Regressive Eye Movements in Reading*. Honors thesis, University of Michigan.

Schilling, H. E. H., Rayner, K., & Chumbley, J. I. (1998). Comparing naming, lexical decision, and eye fixation times: Word frequency effects and individual differences. *Memory & Cognition*, *26*(6), 1270–1281.

Schroyens, W., Vitu, F., Brysbaert, M., & D'Ydewalle, G. (1999). Eye movement control during reading: foveal load and parafoveal processing. *The Quarterly journal of experimental psychology*, *52*(4), 1021–46.

Schumacher, E. H., Lauber, E. J., Glass, J. M., Zurbriggen, E. L., Gmeindl, L., Kieras, D. E., & Meyer, D. E. (1999). Concurrent response-selection processes in dual-task performance: Evidence for adaptive executive control of task scheduling. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(3), 791–814.

Seeck, M., Michel, C. M., Mainwaring, N., Cosgrove, R., Blume, H., Ives, J., Landis, T., & Schomer, D. L. (1997). Evidence for rapid face recognition from human scalp and intracranial electrodes. *Neuroreport*, *8*(12), 2749–54.

Shvartsman, M., Lewis, R. L., Singh, S., Smith, M. R., & Bartek, B. (2011). Predicting Task Performance from Individual Variation in Eye-Movement Control Strategies. In *The CUNY conference of sentence processing.*

Simen, P., Cohen, J. D., & Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural networks : the official journal of the International Neural Network Society*, *19*(8), 1013–26.

Simon, H. (1955). A Behavioral Model of Rational Choice. *The quarterly journal of economics*, *69*(February 1955), 99–118.

Simon, H. (1991). Cognitive Architectures and Rational Analysis: Comment. *Architectures for Intelligence: The 22nd Carnegie . . .*, *0*.

Slattery, T. J., Angele, B., & Rayner, K. (2011). Eye movements and display change detection during reading. *Journal of experimental psychology. Human perception and performance*, *37*(6), 1924–38.

Smith, G. A. & Brewer, N. (1995). Slowness and age: Speed-accuracy mechanisms. *Psychology and Aging*, *10*(2), 238–247.

Stabler, E. (1997). Derivational minimalism. In C. Retoré (Ed.), *Logical aspects of computational linguistics* (pp. 68–95).

Starns, J. J. & Ratcliff, R. (2010). The effects of aging on the speed-accuracy compromise: Boundary optimality in the diffusion model. *Psychology and aging*, *25*(2), 377–90.

Starr, M. & Inhoff, A. W. (2004). Attention allocation to the right and left of a fixated word: Use of orthographic information from multiple words during reading. *European Journal of Cognitive Psychology*, *16*(1-2), 203–225.

Stolcke, A. (1995). An efficient probabilistic context-free parsing algorithm that computes prefix probabilities. *Computational Linguistics*, *21*(2), 165–201.

Sutton, R. S. & Barto, A. G. (1998). Reinforcement learning: an introduction. *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, *9*(5), 1054.

Tarkiainen, a., Helenius, P., Hansen, P. C., Cornelissen, P. L., & Salmelin, R. (1999). Dynamics of letter string perception in the human occipitotemporal cortex. *Brain : a journal of neurology*, *122 (Pt 1*, 2119–32.

Tartakovsky, A. G. (1988). Sequential Testing of Many Simple Hypotheses with Dependent Observations. *Probl. Inform. Transm*, *24*, 1988.

Tatler, B. W., Hayhoe, M. M., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision: reinterpreting salience. *Journal of vision*, *11*(5), 5.

VanRullen, R. & Thorpe, S. J. (2001). The Time Course of Visual Processing: From Early Perception to Decision-Making. *Journal of Cognitive Neuroscience*, *13*(4), 454–461.

Vitu, F., O'Regan, J. K., & Mittau, M. (1990). Optimal landing position in reading isolated words and continuous text. *Perception & psychophysics*, *47*(6), 583–600.

von der Malsburg, T. & Vasishth, S. (2011). What is the scanpath signature of syntactic reanalysis? *Journal of Memory and Language*, *65*(2), 109–127.

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A Diffusion Model Account of Criterion Shifts in the Lexical Decision Task. *Journal of memory and language*, *58*(1), 140–159.

Wald, A. (1945). Sequential Tests of Statistical Hypotheses. *The Annals of Mathematical Statistics*, *16*(2), 117–186.

Wald, A. (1947). *Sequential analysis*. Wiley mathematical statistics series. New York, London: J. Wiley & sons, inc.; Chapman & Hall, ltd.

Wald, A. & Wolfowitz, J. (1948). Optimum Character of the Sequential Probability Ratio Test. *The Annals of Mathematical Statistics*, *19*(3), 326–339.

White, S. J., Rayner, K., & Liversedge, S. P. (2005). Eye movements and the modulation of parafoveal processing by foveal processing difficulty: A reexamination. *Psychonomic bulletin & review*, *12*(5), 891–6.

Williams, C. C., Perea, M., Pollatsek, A., & Rayner, K. (2006). Previewing the neighborhood: the role of orthographic neighbors as parafoveal previews in reading. *Journal of experimental psychology. Human perception and performance*, *32*(4), 1072–82.

Wotschack, C. (2009). *Eye Movements in Reading Strategies*. PhD thesis, Potsdam University.

Yang, S. N. & McConkie, G. W. (2001). Eye movements during reading: a theory of saccade initiation times. *Vision research*, *41*(25-26), 3567–85.

Zacksenhouse, M., Bogacz, R., & Holmes, P. (2010). Robust versus optimal strategies for two-alternative forced choice tasks. *Journal of mathematical psychology*, *54*(2), 230–246.

Zhang, J. & Bogacz, R. (2010). Optimal decision making on the basis of evidence represented in spike trains. *Neural computation*, *22*(5), 1113–48.