AFRL-OSR-VA-TR-2013-0112

MURI-07) DYNAMIC DECISION MAKING IN COMPLEX TASK ENVIRONMENTS: PRINCIPLES AND NEURAL MECHANISMS

James McClelland, Jochen Ditterich, William Newsome and Philip Holmes

Stanford University

March 2013
Final Report

DISTRIBUTION A: Approved for public release.

AIR FORCE RESEARCH LABORATORY
AF OFFICE OF SCIENTIFIC RESEARCH (AFOSR)
ARLINGTON, VIRGINIA 22203
AIR FORCE MATERIEL COMMAND

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
06-02-2013

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
01-07-2007 - 30-06-2012

4. TITLE AND SUBTITLE
(MURI-07) DYNAMIC DECISION MAKING IN COMPLEX TASK ENVIRONMENTS: PRINCIPLES AND NEURAL MECHANISMS

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-07-1-0537

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
James L. McClelland (PI)
Jochen Ditterich, William Newsome, and Philip Holmes, Co-Investigators

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Stanford University, Stanford, CA
University of California, Davis, Davis, CA
Princeton University, Princeton, NJ

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 N. Randolph Street
Arlington, VA 22203
Dr. Jay Myung

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-OSR-VA-TR-2013-0112

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution A: Approved for Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT

We investigated the nature of human decision making in the leaky competing accumulator (LCA) model, a model that links behavior to underlying neural activity through abstract dynamical models in which noisy evidence is accumulated toward a decision, subject to leak and inhibition. The model predicts that evidence coming at different times can receive differential weight, and decisions outcomes should exhibit a mixture of discreteness and continuity. These predictions were confirmed across several studies using more complex task settings than previous studies. Additional studies investigated the optimization of decision making within a simplified reduction of the LCA, while others uncovered evidence establishing how differential payoff information affects decision making. Still other studies investigated the brain basis of decision making, using neuronal recording methods and noninvasive human brain activity measurement methods. Findings from these studies revealed that reward information can change the starting point of evidence integration and that evidence selection and integration are interwoven in the process of reaching a decision.

15. SUBJECT TERMS
Human decision making. Dynamical models of cognition. Mathematical models of mental processes. Human performance optimization.

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Jay Myung
U	U	U	U		19b. TELEPHONE NUMBER *(Include area code)* 703-696-8487

Reset

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Final Project Report for

DYNAMIC DECISION MAKING IN COMPLEX TASK ENVIRONMENTS: PRINCIPLES AND NEURAL MECHANISMS

MURI-07

Grant Award FA9550-07-1-0537

Jay Myung, AFOSR, Program Manager

James L. McClelland, PI
Department of Psychology
Center for Mind, Brain, and Computation
Stanford University

mcclelland@stanford.edu

Date of Submission: February 14, 2013

0. Preamble & Introduction

The ability to optimize behavior in the face of uncertainty and competing goals is of crucial importance to national defense. Theoretical and experimental investigation of the dynamic processes underlying human decisions should increase understanding of human decision making abilities, how these abilities can be optimized, and what the limits are of these abilities. In our MURI project, we have continued to develop a neurodynamic theory of decision making, using a combination of computational and experimental approaches, to address these issues. We pursued a three-pronged approach, (1) extending existing models of dynamic decision making to address the integration of outcome value, reward rate, perceptual uncertainty, and other factors in the decision making process, and assessing these models through behavioral investigations; (2) employing single and multi-single unit recording techniques to investigate the roles of neurons in several brain areas in the representation of decision relevant information and its use in the dynamical process leading to leading to overt decision in varying task situations; and (3) using fMRI, EEG, and MEG to monitor the real-time dynamics of the distributed neural processes underlying decision making in the brain. We pursued this three-pronged approach initially through investigations (Task 1) of fairly standard decision making task situations while also (Task 2) exploring more complex task situations and (Task 3) developing tasks motivated explicitly to address real-world decision-making situations facing aviators, in collaboration with scientists at the NASA Ames research laboratory at Moffett Field. While the specific tasks initially envisioned under Tasks 2 and 3 morphed into others in the course of the research, we feel we were able to address the goal in the call for the MURI 15 competition, of contributing to the integration of theory and experimental investigation across a broad range of levels of analysis, from single neurons to brain areas to the dynamic processes that unfold in real time through human behavior under time pressure.

A great deal of progress has been made developing and extending models of decision making and testing them against other models and against detailed aspects of experimental data, including data from human behavior, primate behavior and neurophysiology, and human brain activation studies. A total of 34 research articles have been published with support from this MURI award, and another 9 articles are still being completed, describing findings obtained during the final year of funding under this project. In addition, several book chapters and review articles have appeared providing synthetic overviews related to the themes of our proposal.

A central element of the progress made under this MURI Grant was the development of the Leaky Competing Accumulator (LCA) model that provided the theoretical foundation around which the proposal was originally developed. As detailed in **Progress Made and Results Obtained** below, findings from several studies conducted under support from this grant support the view that decision making arises as a result of a competition among alternatives, rather than from a race to a decision bound; that decision states both in the model and in participants' responses exhibit a hybrid blend of elements of discreteness and continuity; and that decision making does not stop after an initial decision state is reached but is subject to reversal, should later evidence strongly and persistently support a competing alternative, at least for some participants. The implications of this theory for future research investigating individual differences in the parameters of this complex, interactive and dynamic process and for research on how parameters of this process may be tuned by experience or task constraints to avoid some of its potential pitfalls are discussed in the **Significance of Results and Impact on Science** section.

1. Scientific Objectives of Research

Research on decision making has a long history in the field of human cognitive psychology. The theoretical foundations of this research can be traced back to signal detection theory (Tanner, & Swets, 1954) and the random walk model, providing the basis of the sequential probability ratio test (Wald and Wolfowitz, 1948). These two landmark theoretical innovations became interwoven in the drift diffusion model of Ratcliff (1978), which unified them in a theory of the time course of the integration of continuous noisy perceptual information toward a binary decision for or against a two-alternative forced-choice decision.

Subsequent investigations sought to build bridges between these abstract models of decision making and underlying neurophysiological mechanisms. In the mid 1990's, research in the Newsome lab led to the hypothesis that neurons in the lateral intra-parietal cortex in macaque monkeys functioned essentially as integrators like those described in the drift diffusion model (Shadlen & Newsome, 2001). Research in McClelland's lab, in collaboration with Marius Usher, and parallel work in the laboratory of XJ Wong led to neurally inspired models of decision making that incorporated known properties of neurophysiology to predict novel features of decision dynamics not captured by the earlier more abstract models. Research in the laboratories of Philip Holmes and Jonathan Cohen at Princeton (Bogacz et al, 2006) explicitly explored the links between all of these approaches and proposed a lattice of models extending from detailed physiological models such as that of Wong (2002) through the Leaky Competing Accumulator model of Usher and McClelland (2001) to the more abstract models of Ratcliff (1978) and of Busemeyer and colleagues (Busemeyer & Townsend, 1993; Roe et al, 2001).

This research set the stage for the MURI Topic #15 announced for competition in fall, 2006 for funding in 2007. This Topic called for a project designed to lead toward "a complete and thorough understanding of basic human decision making processes ranging from neuroscience through cognition to behavior". This was to be done "building a lattice of theoretical models with bridges that span across … neural recording and brain imaging in elementary decision to human … decision making with complex dynamic tasks."

The PI approached Usher, Newsome, Holmes, and Cohen with the goal of responding to this challenge, to build effectively on the theoretical foundations discussed above. The Newsome lab also brought neurophysiology while my laboratory and Usher's brought human behavioral investigations and Cohen's provided expertise in human functional brain imaging. We also solicited the collaboration of Dr. Nathan Urban at Carnegie Mellon in an effort to exploit his interest in brain dynamics and his access to the MEG facility at the University of Pittsburgh, as well as the collaboration of Drs. James Johnston and Joel Lachter of the NASA Ames Research Laboratories to confront our efforts with some of the complexities that face human decision makers (aviators) in real-life decision making situations, where decisions much be made in an constantly changing task situation against a backdrop of competing demands on attention.

In line with the MURI Topic Announcement, our goal was to develop and extend existing models of decision making to address issues that were only beginning to be considered by researchers investigating the process of decision making. Task 1 of our research sought to address the integration of outcome value, reward rate, perceptual uncertainty, and other factors including time pressure into theories and models of the decision making process, constraining the development of these models through experiments employing behavioral investigations in humans, single- and multi-single neuron recording studies in primates, and EEG, fMRI, and MEG studies in humans. A central focus under this task was the investigation of the role of prior reward bias in shaping the time course and outcome of the decision making process. To this end we built upon a study already underway in the Newsome laboratory, proposing to collect additional behavioral and physiological data; to test alternative models as possible accounts of the behavioral and physiological data; to vary the behavioral task in further experiments with human participants to more directly assess the time course of information integration, and to investigate the brain basis of this process in humans.

As originally proposed, Task II involved the extension of the study of decision making to continuous time and space, as an example of extending our investigations into more complex task settings. As we will detail in the sections that follow, our investigations did consider a number of more complex task settings than those described under Task I. Among the more complex task settings considered were: Decision making tasks with three or more alternatives; decision making tasks in which the basis for the decision can change from trial to trial; and tasks in which there are multiple display elements and the (human or primate) participant must find the target element as well as make a decision about its identity.

Under Task III, in collaboration with Johnston and Lachter at NASA Ames research laboratories, we initially proposed to extend our effort to investigate real-world situations faced by aviators – in particular, situations posing the need to continually re-assess the state of a decision in real time, such as the extent to which a plane is on course of make a smooth landing. As we proceeded toward the design of specific studies to investigate this matter, the collaborative team became convinced that more work was needed addressing a basic question whose answer could help inform how decision makers allocate their resources when several aspects of a situation are in contention for their attention. We therefore focused this part of the effort on addressing whether human decision makers are able to monitor their own decision states to the extent of being able to indicate, in real time, the state of their certainty about a noisy and ambiguous perceptual variable.

2. Technical Approach

Since our effort focused around the development of dynamical models of decision making, we focus here primarily on the model that served as the central theoretical organizing idea for our project: The leaky, competing accumulator model of Usher & McClelland (2001). In particular, we briefly describe the model and then discuss several questions that were completely open at the outset of our research on which we have now been able to make a great deal of progress. This progress is distributed across the research within the Tasks described above, and employed behavioral research in primates and human, neurophysiology in non-human primates, and non-invasive brain activation studies in humans using EEG and fMRI. After describing the modeling framework we will then describe the behavioral, neurophysiological, and non-invasive brain activation methods.

The leaky competing accumulator model. The leaky competing accumulator model serves as a bridge between detailed biological models on the one hand and completely abstract 'information processing' models on the other. The LCA posits that decision making involves the accumulation of noisy information by an ensemble of accumulators, one for each alternative in a decision-making situation. Each accumulator is thought to correspond to a large population of neurons likely to be distributed across multiple brain areas, all working in concert with each other and in competition with the neurons in other populations. The pattern of activation across the ensemble of accumulators in this model corresponds to the decision maker's *decision state.* We summarize the state of each accumulator with a single activation value, and describe the dynamics of accumulator activation, ultimately serving as a basis for decision, through coupled differential equations capturing the forces operating on each decision variable. Specifically, each accumulator's time evolution is described by the differential equation:

$$dx_i = \left(I_i - kx_i - \beta \sum_{j \neq i} x_j\right) dt + \sigma dW$$

The term I_i represents the external input to the accumulator which may include a common drive B shared by all accumulators plus an additional drive that depends on the input arising from a (possibly time-varying) external stimulus. The term $-kx_i$ captures the tendency of the state of the accumulator 'leak' or decay back toward 0, with $k > 0$ representing the strength of this tendency, and the final term represents the competition from other accumulators, with $\beta > 0$ corresponding to the magnitude of the competition. In the formulation of the model that we currently favor for reasons explained below, the values of the accumulator values are subject to a 'reflecting bound' at 0; that is, if the equation above would result in a negative value for x_i it is instead simply set to 0. The term dt represents the time step while the term σdW represents zero-mean Gaussian noise with standard deviation σ. The reflecting bound on activations at 0 makes the model non-linear and therefore challenging to understand analytically while also introducing very interesting features that have found support from the data in our studies.

Not all of the modeling work carried out in our studies used the LCA but a great deal of it was strongly influences by its tenets; and furthermore, several studies, to be described below, specifically addressed assumptions of this model. Another important body of work under our grant examines how performance can be optimized under the LCA, relying on simplified versions of the full model.

Behavioral methods. All of our studies employed behavioral decision making paradigms, sometimes in conjunction with primate neurophysiological investigations or human brain activation measurements. A typical study involved the collection of an extensive data set from each of a moderate number of non-human primate or adult human research participants. This approach differs from that of many investigators in the human decision making literature who collect at most a few score of trials from each participant in each experimental condition. Our approach makes it possible to provide detailed assessments of the goodness of fit of particular models to each individual participant, avoiding the need to pool data over participants, a process that necessarily obscures individual differences and makes model assessment far more difficult. Many studies are conducted using the free response paradigm, in which participants determine the timing of their responses, while other studies employ some variant of a time-controlled paradigm, in which the state of evidence accumulation at one or (better) many points after stimulus onset is used. One variant of this approach we have found particularly useful is the *interrogation procedure,* in which an imperative signal to respond within a very brief interval (250-300) msec is presented at different times post stimulus onset. In this way we have been able to trace separately the dynamics of the effect of reward and stimulus information on decision making. Complimentary work using the free response paradigm allows assessment of the optimization of decision criteria, a subject of many of the studies supported under this grant.

Neurophysiological and human brain activation methods. In conjunction with behavioral testing studies in the Newsome lab employ single or multi-single unit recording methods in non-human primates. These methods provide evidence about the ways in which individual neurons or ensembles of neurons encode information relevant to a decision. The single unit method involves isolating one neuron then recording while the monkey carries out hundreds of trials in a behavioral study. The electrode is removed from the brain overnight, and the process is repeated until a substantial set of individual neurons has been recorded from, usually over the course of years. Newer multi-single unit approaches rely on the implantation of an electrode array that remains implanted for an extended period, allowing many neurons to be recorded simultaneously. However, with this method is it not easy to establish that one is recording from the same individual neurons in different days.

The human brain activation methods we intended to employ in our studies included Electroencephalography (EEG), Magnetoencephalography (MEG), and functional magnetic resonance imaging (fMRI), based on the blood oxygenation level dependent (BOLD) magnetic susceptibility of the hemoglobin molecule. Both EEG and MEG can track brain activity at high temporal frequency in real time while the BOLD response is sluggish and delayed with respect to the underlying brain activity. Initially, one of the subprojects in our grant was dedicated to the use of MEG to track the temporal dynamics of decision making, but this part of the project proved infeasible. A lack of progress with this subproject was identified at the mid-term review of our MURI project, and the project was then wound down. A portion of the funding for that subproject was used to fund a new subproject added to our MURI grant in year four to fund neurophysiological investigations in the laboratory of Jochen Ditterich at UC Davis, and a portion has been returned to the Air Force.

3. Progress Made & Results Obtained

We begin by describing the progress made on exploring and testing the implications of the specific assumptions of the LCA that differentiate it from other models using human behavioral data. These assumptions include the presence of leakage and inhibition and the presence of a floor or reflecting bound on activation at 0. An emergent consequence of this constellation of assumptions is new characterization of the concept of 'decision state' and of what it means to 'make a decision'. Following this, we describe a substantial body of work exploring the optimization of decision processes and the neural mechanisms underlying this optimization, as explored in work on human participants. A third section describes neurophysiogical investigations of factors affecting decision making, with cross-reference to relevant human neuroscience studies. Studies relevant to both Tasks 1 and 2 described under scientific goals are integrated into all three sections of this narrative. The studies under Task 3 (Lachter et al, in preparation) are integrated into the first section.

Implications of the LCA. We begin by considering the roles of leakage and inhibition somewhat separately from the remaining issues. A mathematical analysis of the two-alternative version of the model in Usher and McClelland (2001), building on earlier work by Busemeyer & Townsend (1993) established that, whenever there is an imbalance between leak and inhibition in the model (i.e., whenever $k \neq \beta$ or equivalently whenever $\lambda = (k - \beta) \neq 0$), performance in a decision making task levels off below the level of 100% correct responding, at a level reflecting an interplay between the degree of imbalance and the strength of sensory evidence favoring the correct alternative. In particular the growth of accuracy as a function of processing time follows an exponential approach to asymptote, where the rate of approach to asymptote is determined by the degree of imbalance ($|\lambda|$)and the asymptotic accuracy measured by the signal detection variable d' is proportional to the signal strength of stimulus support divided by $|\lambda|$. Interestingly, however, leveling off occurs for qualitatively different reasons when λ is positive (i. e., leak is stronger than inhibition so that the process is *leak dominant*) or negative (inhibition stronger than leak, so the process is *inhibition dominant*). In the leak dominant case, early information entering the accumulators effectively decays away, favoring evidence coming in the recent period before a decision is required; while in the inhibition dominant case, evidence coming into the process early in a trial can give one alternative the upper hand, and the resulting inhibition can thereafter suppress the other alternative, even if on balance the evidence favors it. To investigate this issue, it is necessary to manipulate when during an evidence accumulation period critical stimulus information is presented. While Usher & McClelland (2001) made a preliminary investigation of this issue, studies supported by our MURI grant have considerably extended the investigation of this issue. Tsotsos, Gao, McClelland and Usher (2012) used random dot motion stimuli of the sort used in many primate neurophysiology studies. Evidence favoring a left or right response, in the form of different degrees of stimulus coherence, was presented either throughout the viewing period on a given trial (the constant condition), or either during only the first or only the second half of the trial (early and late conditions). Participants were required to respond within 300 msec of a go cue that occurred immediately after the end of the stimulus, which varied in duration from 300 to 1500 or 2000 msec. All participants showed a primacy effect (greater accuracy in the early vs the late condition) consistent with inhibition dominance. An alternative to inhibition dominance is the idea that evidence integration stops when a decision bound is reached (Kiani et al, 2008). Several aspects of the findings in Tsetsos et al (2012) and others of our studies (Gao, Totell & McClelland, 2010; Tsetsos et al 2011; Gao and McCelland, submitted) favor the inhibition dominance interpretation, including both qualitative and quantitative signatures of goodness of fit in Gao et al. and in Gao and McClelland.

A further subtlety not predicted by bounded integration arises from the presence of the reflecting bound at 0 in the inhibition-dominant LCA. This is the fact that this reflecting bound makes possible the reversal of decision states when evidence changes, albeit subject to a delay in in the reversal process due to inhibition. Consider a situation in which evidence switches half way through a trial from favoring one alternative to favoring the other. For some levels of stimulus strength and a moderate degree of inhibition dominance, the LCA will favor the alternative that received greater support during the first half of the short trial, but favors the alternative that received greater support during the second half of the long trial. This feature of the model thus predicts an interaction, such that early is favored over late in short trials and late is favored over early in long trials. Although this pattern was seen in only a subset of participants in Tsetsos et al, it cannot be explained by a bound on information integration. Further evidence consistent with the reversibility of decision states and inconsistent with bounded integration models is also provided in Tsetsos, McClelland & Usher, 2011.

With this support for basic features of the LCA in place, we now turn to a consideration of the study of Gao, Tortell & McClelland (2012), which explored the role of reward and stimulus information using the interrogation protocol. This study built on the neurophysiological study the effects of reward on decision making of Rorie, Gao, Newsome, and McClelland (2010) and the model of the behavioral data from that study presented in Feng, Rorie, Newsome & Holmes (2008). Taken together, these studies provided support for the view that, at least under the task conditions used by Rorie et al, payoff information presented prior to the onset of stimulus information affected the starting activation of putative accumulator neurons so that their activation favored the higher reward alternative even before stimulus input began to influence these neurons' activation. The study by Gao et al. compared the predictions of this hypothesis with the predictions of two other plausible alternative accounts for the effect of prior reward information, and, further, assessed the overall adequacy of an account based on the inhibition-dominant LCA (with a reflecting bound on activation at 0) to capture the details of the pattern of choice responses exhibited in the extensive data sets from each of the five participants in the experiment. The LCA model provide an excellent fit to the data, accounting for the qualitative and quantitative form of the pattern of the effect of reward bias on response choices made to go cues presented at different times following stimulus onset, as well as for the invariant shape of time-accuracy curves for different levels of stimulus difficulty, a feature not predicted by bounded integration models.

With several studies pointing toward the inhibition-dominant LCA with a reflecting bound at 0 as a model capturing several features of data that other existing approaches could not explain, we began to consider the nature of decision states in the model and how these are translated into action – and we began also to consider the implications of this version of the LCA for what it means to 'make a decision'. In many models, to make a decision is to accumulate evidence over time until a decision bound is reached. In the inhibition dominant LCA, however, there is an alternative to this idea, namely that 'to make a decision' is to resolve a competition in favor of one or the other choice alternative, so that one accumulator remains active while the other(s) are suppressed to an activation of zero. Under this conception of what it means to 'make a decision', a decision state has a mixture of discrete and continuous properties. In particular, one accumulator wins and the other looses – a feature of a discrete decision – while the winner's activation remains a continuous random variable, whose mean value is a function of the support it receives from the sensory stimulus. The mean activation level of the winning accumulator is no longer dependent on inhibition from other accumulators once it has won the competition and the other's activations have been suppressed to 0. In this situation, however, it state is still dependent on leak, so activation levels off at a point reflecting the balance between stimulus support for the alternative on the one hand and leak on the other. The stronger the stimulus support, the stronger the activation of the winning accumulator. This then is the element of continuity remaining in the decision state.

Three further studies not yet published provide two very different kinds of support for the predictions of this version of the LCA. In one of these studies (Gao & McClelland, submitted) we further examined the data from the Gao et al (2010) study, looking at the time taken to respond to after the presentation of the go cue at different times after stimulus onset and with different stimulus conditions. Assuming, in accordance with the above discussion, that the decision state remains continuous until the go cue occurs, we considered how the exact level of accumulator activation would influence the time to respond following the go cue. Based on qualitative features of our own data as well as other recent studies, we proposed that the activation of each of the accumulators at the time the go cue occurs determines the strength of input to a response activation accumulator; this in turn determines the rate of activation of the response accumulator, so that stronger activation of the evidence accumulator will result in faster responding (modeled as a race between the response accumulators such that the response is determined by the first one to reach the bound). This then led to several specific implications for response times based on the model of the evidence accumulation process previously laid out in Gao et al. Three points, in particular, are of most interest. (a) the assumption that reward affected the starting point of the evidence accumulators meant that when the go cue occurred earlier, faster responses were associated with choices of the high reward alternative. (b) The assumption that reward affected the starting point of evidence accumulation further implied that reward information would influence the likelihood that the alternative associated with the higher reward would win the competition, and thus determine the response, even at long delays – but that the reward would not affect the degree of activation of the winning accumulator at long delays, since it would no longer be providing input to the accumulator. (c) Instead, at long delays, only the strength of stimulus support for the alternative chosen would affect the time take to respond. All three of these predictions were confirmed: That is, responses consistent with the reward bias were faster at short delays; response probability, but not response speed, was affected by reward bias at long delays; and stronger stimulus support was associated with faster responding at long delays. In addition to capturing these and other qualitative features of the data, the model also accounted for the relative sizes of the reward effect at short lags and the stimulus support effect at long lags. Other models appear to predict that reward should affect response speed to the extent that it affects response probability, and thus are inconsistent with an important an counter-intuitive feature of the data. Second, a recent EEG study conducted in the Holmes-Cohen labs develops a chain of LCA models for accumulation, threshold and response areas (van Vugt et al., submitted), showing that bistable neural activity can implement decision thresholds, and that lateralized readiness potentials (LRPs) reflect its dynamics.

The final type of support for the predictions of the LCA – and in particular, for the mixture of discrete and continuous features of decision making – is provided by a study by Lachter, Corrado, Johnston & McClelland (in preparation). In this study, we presented participants with two fields of dots, one containing one, three, five, seven or nine more dots than the other, and asked participants to indicate the judgment of the relative likelihood that there were more dots in the left field or in the right field. The relative likelihood scale had a zero point, such that responses to the side of the zero point associated with the field containing more dots could be scored as ‘correct’ while responses to the other side of the zero point could be scored as ‘incorrect’. On this basis, accuracy as measured by *d’* increased with the magnitude of the dot difference, as all extant theories would expect. Relevant to the mixture of discrete and continuous features in the decision states of the LCA, we observed just these features in the distribution of relative likelihood responses produced by many of the participants in this experiment. This is illustrated in the inset Figure, where the response distributions for each of the five stimulus difference levels (one, three, five, seven or nine dots) produced by one of the participants in the study are shown. Clearly, this participant’s responses show a degree of discreteness, in that there were no responses at the point on the scale corresponding to equal likelihood of each of the two alternatives. Yet they also show a degree of continuity, in that that they are farther from the indifference point when supported by stronger stimulus information.

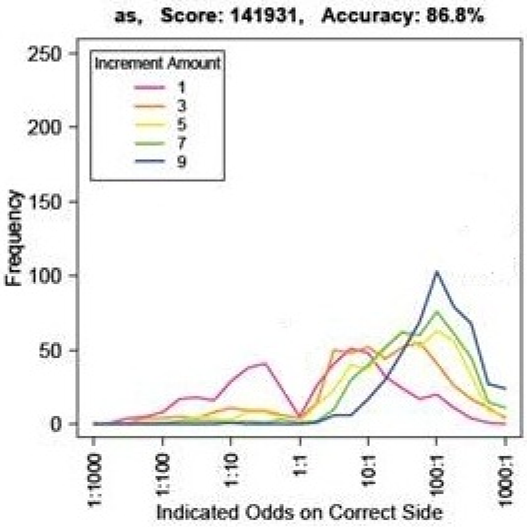

In general, in this and several others of our studies, not all participants produce a pattern of responding that requires the full complexity of the inhibition dominant LCA with a reflecting bound at zero. In this study and is some other earlier studies, some participants exhibit a pattern consistent with leak-dominance or equally consistent with inhibition dominance and bounded integration. Over all of these studies (those of Gao et al, 2010, Gao and McClelland, submitted; Lachter et al, in preparation; Tsetsos et al, 2011, 2012) only the LCA provides a complete account of the full range of patterns seen in the data, albeit requiring the full flexibility of the model to address patterns exhibited by different individuals. It should also be noted that there are features of some participants data in

Lachter et al. that require additional assumptions about the mapping from internal representations to responses to fully explain the pattern of data. Finally, a combined behavioral/EEG study of the role of reward bias in decision making revealed evidence that under some task conditions, models such as the LCA or other decision models must be supplemented by other assumptions, such as a fast-guess mechanism, to account for all aspects of participant's performance (Noorbaloochi et al, in preparation).

Optimization of decision making and its brain basis: theory, models, behavior, and human brain activation studies. A parallel strand of research growing out of the integrated analysis of the LCA and its simplified cousin the Drift Diffusion Model (DDM) offered by Bogacz et al (2006) explores the optimization of performance in free response settings, in which participants must establish a rule for terminating the evidence integration process. It is widely accepted that integration stops when a criterion is reached based on accumulated evidence, but whether this criterion is fixed or effectively changes during the course of a trial, or whether it is an absolute or relative criterion remains a subject of extensive ongoing research. A large body of theoretical and experimental work, based on the normative theory and optimal performance curve (OPC) developed in Bogacz et al. (2006), sheds important light on this issue (for review, see Holmes & Cohen, submitted).

One key finding is that participants often fail to achieve optimality, erring on the side of collecting evidence beyond the point where the improvement in accuracy is justified by the payoff contingencies of an experiment. Several alternative accounts for this effect have been considered through theoretical analyses and experiment (Bogacz et al., 2006, 2010; Zackenhouse et al., 2009; Simen et al., 2009; Balci et al, 2011). One possibility is that participants over-weight the importance of accuracy, perhaps implicitly assuming there are explicit costs (negative earnings) for errors. While this may be in play in many settings, practice tends to reduce this effect, resulting in enhanced reward rate. Another factor in computing optimal settings of response criteria may be the difficulty participants have in estimating the passage of time (Zackenhouse et al., 2009; Simen et al., 2011a, 2011b). Balci et al (2011) linked variability in participant's interval timing estimates to their ability to optimize reward rates. A third factor relevant to reward rate optimization is the cost or difficulty of exerting control over decision criteria (Todd et al., 2011). Such control may be experienced as effortful and/or can only be optimized by expending time and effort to track reward rate and make adjustments, and this may lead participants to adopt a single criterion across blocks of trials where adjustments of criteria would lead to greater overall reward (Balci et al., 2011).

A fourth factor has also emerged. Bogacz et al. (2006) had already shown that the *linearized* LCA, with large leak and inhibition, reduces to a 1-dimensional Ornstein-Uhlenbeck process, but that this is *only* an optimal DDM when leak and inhibition are balanced. Moreover, although biophysically-based spiking neuron models can be reduced to *nonlinear* accumulators (Eckhoff et al, 2009; 2011), these are not DDMs, and they exhibit more complex nonlinear dynamics than the LCAs of Usher-McClelland (2001). These theoretical studies suggest physiological constraints to optimality.

Other studies supported by this grant also considered how criteria may be adjusted on line to achieve, in some cases, a good approximation to optimization based on a fairly simple titration policy, or, alternatively, may actually reduce optimality by introducing variation in criteria that only serve to degrade performance (Yu et al., 2008). Extensions to LCAs and

DDMs that involve trial-to-trial threshold and starting point updates were created to account for sequential effects (Gao et al, 2009; Goldfarb et al, 2012), and diffusion models of interval timing were developed (Simen et al 2011a; Balci & Simen, submitted). Wong-Lin et al (2010) built an LCA model to predict optimal behavior in a countermanding task, and Zhou et al (2009) developed methods to distinguish between leak- and inhibition-dominated processes by injecting brief pulses of strong evidence, as also investigated in Tsetsos et al. (2012). Studies have also considered the possible neural basis of criterion adjustment (van Vugt et al., submitted; Simen et al., in preparation) and the relationship between these variables and disorders that affect decision making (Mulder et al, 2010).

Neurophysiology of decision making in primates. As noted above, a study already in progress when our grant was submitted explored the behavioral consequences and neural basis of reward bias in decision making in a two-alternative forced-choice task in non-human primates. This study led to two important papers, one assessing the optimality of the primate's behavior (Feng et al., 2009) and one reporting the neurophysiological findings in relation to the data and offering a mechanistic computational model constrained, simultaneously, by both the brain and behavioral data (Rorie et al., 2010). Taken together these investigations clearly showed that under the constraints of the particular task used, both monkeys made near optimal use of reward bias information. The data supported the hypothesis that they did so by biasing the starting point of the evidence accumulation process; two other alternative accounts could explain the behavioral data alone but were ruled out by the combination of the physiological evidence. Specifically, Rorie et al reported that the activation of putative evidence accumulator neurons was offset by reward information presented, providing a starting point for evidence accumulation at the time of the presentation of the stimulus. Modeling work also reported in the same paper established that this offset was sufficient to explain the effect of reward bias on the neural activity data as well as the behavioral choice data.

Additional work from the Newsome lab has explored post-response choice tracking of eye-movement base decisions to allow subsequent outcome information to affect future choices. How does the brain track the identity of a stimulus and choice response during the period before a reward is received? How can it update the value of a given stimulus-response pairing when the corresponding sensory and motor representations are no longer active? Reppas & Newsome (submitted) describe a frontal-lobe choice-history signal that provides an enduring neural trace specific for the just-made eye movement during decision-making behavior. Neurons that carry this history signal are distinct from saccade-planning neurons, but exhibit preferential connectivity with those plan neurons with which they share a common choice preference. The history signal they describe may enable decisions to be faithfully linked to the outcomes they generate, even when those outcomes are deferred by temporal intervals of varying (and sometimes relatively long) duration. Two other studies from the Newsome lab supported by this MURI grant examine neural population activity that accounts for variance in saccadic latencies (Kalmar et al., submitted; Kiani et al., submitted).

A recent very exciting development in new work from the Newsome lab (Mante et al., submitted) applies advanced neural population modeling and analysis methods to reveal how the brain adaptively maps sensory information onto a response choice on a trial by trial basis. This work has the potential to link sophisticated ways of representing neural population activity to higher-level characterizations such as those provided by more abstract models such as the LCA and the DDM models that have been the backbone of brain-decision modeling up to the present. In particular, the authors study neural activity in prefrontal cortex in monkeys trained to flexibly select and integrate noisy sensory inputs towards a choice. They find that the observed complexity and functional roles of single neurons are readily understood in the framework of a dynamical process unfolding at the level of the population. The population dynamics can be reproduced by a trained recurrent neural network, which reveals a previously unknown mechanism for selection and integration of task-relevant inputs. This mechanism implies that selection and integration are two aspects of a single dynamical process unfolding within the same prefrontal circuits, and potentially provides a novel, general framework for understanding context-dependent computations.

Other neurophysiology studies in the laboratory of Jochen Ditterich explore in detail how neurons in parietal cortex compute net sensory evidence for one of several decision alternatives. Bollimunta & Ditterich (2011) trained monkeys on a perceptual decision task that allowed simultaneous experimental control over how much sensory evidence was provided for each of 3 possible alternative choices and recorded single unit activity and local field potentials (LFPs) from the lateral intraparietal area (LIP). While both the behavior and the spiking activity were largely determined by the difference between how much supporting sensory evidence was provided for a particular choice (pro evidence) and how much sensory evidence was provided for the other alternatives (anti evidence), the LFP reflected roughly the sum of these 2 components. Furthermore, the firing rates showed an earlier influence of the anti-evidence than the pro evidence. These observations indicate that LIP does not simply receive already pre-computed decision signals but that it plays an active role in computing the decision-relevant net sensory evidence and that this local computation is reflected in the LFP. A second study by Bollimunta et al. (2012) recorded simultaneously from multiple decision-related neurons in parietal cortex of monkeys performing a perceptual decision task and used these recordings to analyze the neural dynamics during single trials. Decision-related lateral intraparietal area neurons typically undergo gradual changes in firing rate during individual decisions, as predicted by mechanisms based on continuous integration of sensory evidence. Furthermore, we identify individual decisions that can be described as a change of mind: the decision circuitry was transiently in a state associated with a different choice before transitioning into a state associated with the final choice. These recent findings from the Ditterich lab support some predictions of the Leaky Competing Accumulator model but challenge others, and further modeling work assessing exactly how to relate neural activity to decision outcome, via population-level models such as the LCA, are underway.

4. Significance of Results & Impact on Science

The goal of developing a lattice of models spanning from neurons to behavior and taking full account of the physiological and psychological factors that influence decision outcomes is a very long-term goal. As progress is made toward this goal, we will have a greater and greater understanding of the limits of human decision making performance and of the bounds on its possible rationality.

The studies reported here reflect an ongoing transition in models of decision dynamics. While simpler perfect integration models can often provide fairly good accounts of the pattern of results in a decision making experiment, the work reported here suggests that a detailed consideration of data obtained in both relatively standard decision making settings (two-alternative forced-choice tasks with stationary evidence) and in somewhat more complex settings (e.g., with three or four alternatives, non-stationary evidence, or with the need to use a trial-specific cue to select the appropriate dimension of sensory evidence) has begun to suggest that the more complex dynamical characteristics of models like the non-linear Leaky Competing Accumulator model are reflected in subtler aspects of behavior and possibly also in brain activity. It is likely that future research will be increasingly influenced by such models, and by the effort to identify further behavioral and neural markers of the non-standard features of these models.

As our conception of decision dynamics becomes richer, so, too, does our conceptual framework for understanding individual differences in decision dynamics. A highly parsimonious model may have only two or three free parameters (a sensitivity, a bias, and an integration threshold) allowing for relatively few possible ways for performance to be affected by underlying neural mechanisms. But when the parameter structure is richer (including leak, inhibition and a common drive or offset parameter that affects the engagement of the reflecting bound at 0) there may be more model freedom but there is also the possibility of understanding in a far richer way how individual differences may affect aspects of decision performance. This opens the way for an important body of new research exploring individual differences in decision making – both in terms of the behavioral phenotype and the neural mechanisms and factors that support this phenotypic variation.

The further prospect of a richer theory of decision dynamics is a greater opportunity to explore ways to foster its optimization. Much of the work discussed here explores how participants can control the placement of the decision threshold – but there are many other parameters, with identifiably different effects, that also may be tunable, and much remains to be learned about the controllability of these parameters and/or what factors influence their values.

The work described above also breaks new ground in linking neural population activity to behavioral outcomes. Most of the models to date have assumed that the members of the neural population that support a decision all have the same neural response function, perhaps differing in scalar parameters such as gain or variability but otherwise essentially identical. The exciting new work by Mante et al (submitted) breaks new ground in this regard, allowing us to begin to see how a population of neurons, each with a potentially quite different response profile, can each play a part in implementing a population-level computation that affords a richer array of computations (selection of the sensory dimension on which a response is based, as well as integration of the relevant evidence to a decision bound) than are naturally captured by populations made up of essentially identically distributed neurons.

5. Publications Resulting from Research

Balci, F., D. Freestone, P. Simen, L. deSouza, P. Holmes and J.D. Cohen (2011a). Optimal Temporal Risk Assessment. *Frontiers in Integrative Neuroscience 5*: 56.

Balci, F., P. Simen, R. Niyogi, A. Saxe, P. Holmes and J.D. Cohen (2011b). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception & Psychophysics 73 (2):*640-657.

Bogacz, R., P. Hu, P. Holmes and J.D. Cohen (2010). Do humans produce the speed-accuracy tradeoff that maximizes reward rate? *Quart. J. Exp. Psychol., 63* (5):863-891.

Bogacz,R., Usher, M., Zhang, J. & McClelland, J. L. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. Theme issue on Modeling Natural Action Selection. *Philosophical Transactions of the Royal Society: B. Biological Sciences. 362.* 1655-1670.

Bollimunta A, Ditterich J (2012) Local Computation of Decision-Relevant Net Sensory Evidence in Parietal Cortex. *Cerebral Cortex, 22*, 903-917

Bollimunta A, Totten D, Ditterich J (2012) Neural Dynamics of Choice: Single-Trial Analysis of Decision-Related Activity in Parietal Cortex. *Journal of Neuroscience, 32,* 12684-1270

Broderick, T., K.F. Wong-Lin and P. Holmes. (2010). Closed-form approximations of first passage distributions for a stochastic decision making model. *Applied Mathematics Research eXpress, Vol. 2009, No.2,* 123-141.

Eckhoff, P., K.F. Wong-Lin and P. Holmes (2009) Optimality and robustness of a biophysical decision-making model under nonepinephrine modulation. *J. Neurosci. 29,* 4301-4311.

Eckhoff, P., K.F. Wong-Lin and P. Holmes (2011). Dimension reduction and dynamics of a spiking neuron model for decision making under neuromodulation. *SIAM J. on Applied Dynamical Systems 10*, 148-188.

Feng, S., P. Holmes, A. Rorie and W.T. Newsome (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Computational Biology 5 (2),* e1000284.

Gao, J., Tortell, R., & McClelland, J. L. (2011). Dynamic integration of reward and stimulus information in perceptual decision making. *PLoS ONE 6(3):* e16749.

Gao, J., K.F. Wong-Lin, P. Holmes, P. Simen and J.D. Cohen. (2009). Sequential effects in two choice reaction time tasks: Decomposition and synthesis of mechanisms. *Neural Computation, 21(9),* 2407-2436.

Goldfarb, S., K.F. Wong-Lin, M. Schwemmer, N.E. Leonard and P. Holmes (2012). Can post-error dynamics explain sequential reaction time patterns? *Frontiers in Psychology 3,* 213.

Henderson, C. M. & McClelland, J. L. (2011). A PDP model of the simultaneous perception of multiple objects. *Connection Science, 23,* 161-172.

Kumaran, D. & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review, 119*, 573-616.

Lake, B. M., Vallabha, G. K. & McClelland, J. L. (2009). Modeling unsupervised perceptual category learning. *IEEE Transactions on Autonomous Mental Development, 1*, 35-43.

McMillen, T., P. Simen and S. Behseta (2011). Hebbian learning in linear-nonlinear networks with tuning curves leads to near-optimal, multi-alternative decision making. *Neural Networks, 24*, 417-426.

Mulder, M., D. Bos, J. Weusten, J van Belle, S. van Dijk, P. Simen, H. van Engeland and S. Durston (2010). Basic impairments in regulating the speed-accuracy tradeoff predict symptoms of ADHD. *Biological Psychiatry, 68,* 1114-1119.

Rorie, AE, J Gao, JL McClelland and WT Newsome. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex. *PLoS One, 5(2),* e9308, 2010.

Simen, P. (2011). Preventing combinatorial explosion in a localist neural network architecture using temporal synchrony. *Connection Science, 22,* 131-144.

Simen, P., F. Balci, L. deSouza, J.D. Cohen and P. Holmes (2011a). A model of interval timing by neural integration. *J. Neurosci. 31,* 9238-9253.

Simen, P., F. Balci, L. deSouza, J.D. Cohen and P. Holmes (2011b). Interval timing by long-range temporal integration. *Frontiers in Integrative Neuroscience 5 (1-2),* 28.

Simen, P. and J. D. Cohen (2009). Explicit melioration by a neural diffusion model. *Brain Research, 1299,* 95-117.

Simen, P., D. Contreras, C. Buck, P. Hu, P. Holmes and J.D. Cohen (2009). Reward rate optimization in two alternative decision making: Empirical tests of theoretical predictions. *J. Exp. Psychol.: Human Perception & Performance 35 (6)*, 1865-1897.

Shlizerman, E., and P. Holmes (2012). Neural dynamics, bifurcations and firing rates in a quadratic integrate-and-fire model with a recovery variable. I: deterministic behavior. *Neural Computation 24,* 2078-2118.

Sternberg, D. A. & McClelland, J. L. (2012). Two mechanisms of human contingency learning. *Psychological Science, 23(1),* 59-68.

Tsetsos, K., Gao, G., McClelland, J. L., & Usher, M. (2012). Using time-verying evidence to test models of decision dynamics: Bounded diffusion vs. the leaky competing accumulator model. *Frontiers in Neuroscience, 6,* 79.

Tsetsos, K., Usher, M., & McClelland, J. L. (2011). Testing multi-alternative decision models with non-stationary evidence. *Frontiers in Neuroscience, 5,* 63.

van Vugt, M.K., P. Simen, L. Nystrom, P. Holmes and J.D. Cohen (2012). EEG oscillations reveal neural correlates of evidence accumulation. *Frontiers in Neurosci. 6,* 106.

Wong-Lin, K.F., P. Eckhoff, P. Holmes and J.D. Cohen. (2010). Optimal performance in a countermanding saccade task. *Brain Research, 1318,* 178-187.

Yu, A.J., P. Dayan and J.D. Cohen (2009). Dynamics of attentional selection under conflict: Toward a rational Bayesian account. *J. Exp. Psychol.: Human Perception & Performance, 35 (3),* 700-717.

Zacksenhouse, M., R. Bogacz and P. Holmes (2010) Robust versus optimal strategies for two-alternative forced choice tasks. *J. Math. Psychol. 54,* 230-246.

Zhang, J., R. Bogacz and P. Holmes (2009). A comparison of bounded diffusion models for choice in time controlled tasks. *Journal of Mathematical Psychology, 53 (4),* 231-241.

Zhou, X., K.F. Wong-Lin and P. Holmes (2009). Time-varying perturbations can distinguish among integrate-to-threshold models for perceptual decision making in reaction time tasks. *Neural Computation 21 (8),* 2336-2362.

Research articles in the pipeline (submitted or in preparation)

van Vugt, M.K., P. Simen, L.E. Nystrom, P. Holmes and J.D. Cohen (submitted). Lateralized readiness potentials reveal properties of a neural mechanism for implementing a decision threshold.

Niyogi, R.K. and K.F. Wong-Lin (submitted). Dynamic excitatory and inhibitory gain modulation can produce flexible, robust and optimal decision making.

Balci F. and P. Simen (submitted). A drift-diffusion account of response times in temporal bisection.

Lachter, J., Corrado, G., Johnston, J. C. & McClelland, J. L. (in preparation). Graded and discrete features of decision outcomes based on a continuous response measure.

Mante, V, D Sussillo, KV Shenoy and WT Newsome (submitted). Selective integration of sensory evidence by recurrent dynamics in prefrontal cortex.

Reppas, JB, and WT Newsome (in preparation). A choice-selective history signal in the frontal lobe of macaque monkeys.

Kiani, R, CJ Cueva, JB Reppas, D Peixoto and WT Newsome. (in preparation) Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex.

Gao, J. and J. L. McClelland. (submitted) Continuous decision states and their translation into action.

Gao, J. and J. L. McClelland. (in preparation). Dynamics of multi-sensory integration.

Noorbaloochi, S. D. Sharon, and J. L. McClelland. (in preparation). The role of payoff bias in a deadlined perceptual decision making task: Behavior, electrophysiology, and modeling.

Book chapters and review articles:

Bogacz, R., Usher, M., Zhang, J. & McClelland, J. L. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. Theme issue on Modeling Natural Action Selection. In A. K. Seth, T. J. Prescott and J. J. Bryson (Eds.), *Modelling Natural Action Selection*. pp 91-119. Cambridge, UK: Cambridge University Press.

McClelland, J. L. (2009). The place of modeling in cognitive science. Topics in Cognitive Science, 1(1), 11-38. [PDF] | [DOI]

McClelland, J. L. (2010). Emergence in cognitive science. Topics in Cognitive Science, 2, 751-770.

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T.T., Seidenberg, M. S., and Smith, L. B. (2010). Letting Structure Emerge: Connectionist and Dynamical Systems Approaches to Understanding Cognition. Trends in Cognitive Sciences, 14, 348-356.

McClelland, J. L. & Vallabha, G. (2009). Connectionist models of development: Mechanistic dynamical models with emergent dynamical properties. In J.P. Spencer, M. S. C. Thomas, & J. L. McClelland, (Eds). Toward a unified theory of development: Connectionism and dynamic systems theory re-considered. 3-24. New York: Oxford.

Simen, P., J.D. Cohen and P. Holmes (2009). On the neural implementation of optimal decisions. In ``The Oxford Handbook of Human Action,'' ed. E. Morsella, J.A. Bargh, and P.M. Gollwitzer, pp 533-548, Oxford University Press, Oxford, UK.

Thomas, M. S. C. & McClelland, J. L. (2008). Connectionist models of cognition. In R. Sun (Ed). Cambridge handbook of computational psychology. Cambridge University Press. 23-58.

Tytell, E.D., P. Holmes and A.H. Cohen (2011). Spikes alone do not behavior make: Why neuroscience needs biomechanics. Current Opinion in Neurobiology 21:816-822.

Usher, M., Elhalal, A. & McClelland, J. L. (2008). The neurodynamics of choice, value-based decisions, and preference reversal. In Chater, N. & Oaksford, M. The Probabilistic Mind. New York: Oxford University Press. 278-300.

Book chapters and review articles in preparation:

Holmes, P. and J.D. Cohen (submitted). Optimality and some of its discontents: Successes and shortcomings of existing models for binary decisions.

h

6. References Cited

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J.D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review*, *113*, 700-65.

Busemeyer, J. R. & Townsend, J. T. (1993). Decision Field Theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432-459.

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

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review, 108,* 370–392.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol, 86*(4), 1916-1936.

Tanner, W.P., Jr., & Swets, J.A. (1954). A decision-making theory of visual detection. *Psychological Review, 61,* 401-409.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108,* 550–592.

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

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron, 36*, 955-968.