

**Properties of Human Motor Control Under Risk  
and Risk Aware Control**

by

**Amber Lynn Dunning**

B.S., Arizona State University, 2011

M.S., University of Southern California, 2015

A dissertation submitted to the  
Faculty of the Graduate School of the  
University of Southern California in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
Department of Biomedical Engineering  
May 2017

This thesis entitled:  
Properties of Human Motor Control Under Risk and Risk Aware Control  
written by Amber Lynn Dunning  
has been approved for the Department of Biomedical Engineering

---

Dr. Terence Sanger

---

Dr. Jean-Michel Maarek

---

Dr. Maryam Shanechi

---

Dr. Francisco Valero-Cuevas

Date \_\_\_\_\_

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Dunning, Amber Lynn (Ph.D., Biomedical Engineering)

## Properties of Human Motor Control Under Risk and Risk Aware Control

Thesis directed by Dr. Terence Sanger

Recently, a lot of attention has been given to exploring the type of control algorithm humans implement in movement. A comprehensive theory of motor control is important for many reasons. It would allow us to compare symptoms of motor diseases to symptoms resulting from different interruptions and damages in the model of motor control to gain a better understanding of the pathophysiology and construct a focus for treatments. A complete understanding of motor control will also influence design of prosthetics and biomimetic robots. It could also have many implications in learning and may even transform the way we teach motor actions.

There are several proposed models, which predominantly focus on achieving a goal through a reference trajectory (Todorov, 2002; Todorov, 2004). However, motor control is not just about reaching a goal, but also avoiding predictable failure in the process. Risk is inherent in all activity, and avoidance of risk is fundamental to human survival, so response to risk must be an integral part of human movement as well. A new theory, Risk-Aware Control, emphasizes selecting motor actions that minimize risk (Sanger, 2014). Risk-Aware Control is distinctive from classical theories in that it is an entirely new way of approaching the relationship between cost and motor actions. It does not attempt to formulate a reference trajectory to the goal, but instead predicts that movement develops from maintaining a probability distribution of state, a detailed understanding of the cost function, and knowledge of the relationship between action and change in state. The result is a control theory that accounts for incorporation of uncertainty and cost in motor planning and execution, plans for unexpected error prior to perturbation, and does not require assumptions of system linearity.

The theory of risk-aware control suggests a reduction in computational burden com-

pared to current full models of human motor control because it allows for parallel computing. In an existential proof, we implement risk-aware control using a spiking neuron model of cortex to control a robotic arm in real time. Utilizing the framework of Stochastic Dynamic Operators, we were able to offload the majority of computation to a graphics processing unit to maintain a high operating rate. We explore the effects of gain and damping parameters on the control and demonstrate response behavior to perturbations.

Since evasion of risk is fundamental to survival we believe it must be a fundamental component in human motor control as well. In order to characterize and emphasize the influence of risk in the environment on human behavior, we have designed a series of experiments. In these studies, we are describing risk as the expected cost of behavior defined by the combination of cost of failure and probability of failure. The rest of this report details these experiments.

The role of motor uncertainty in discrete or static space tasks, such as pointing tasks, has been investigated in many experiments (Trommershauser et al. 2003a; Trommershauser et al. 2003b). These studies have already shown that humans hold a highly accurate internal representation of their intrinsic motor uncertainty and compensate accordingly for this variability. Furthermore, experiments imposing additional extrinsic motor and sensory variability have shown that subjects still respond near optimally, even as risk increases (Trommershauser et al. 2005). While static conditions provide an important foundation to understanding the relationship between risk and movement, they rarely appear in natural situations. The aim of our first study was to investigate how humans respond to uncertainties in a dynamic environment despite indeterminate knowledge of the outcomes of specific actions. Our hypothesis was that subjects would tune their statistical behavior to uncertainty based on cost in a dynamic, feedback-driven task.

In the first experiment, subjects maintained one-dimensional “steering” control of a vehicle in an iPad driving simulation. The speed of the car was determined solely by position on a two-lane road. While on the road, driving in a lane yielded the maximum possible

velocity, driving on the dashed line between lanes caused the vehicle to slow down, and hitting the grass along the side of the road brought the car to a complete stop. The road contained random curves so that subjects were forced to use sensory feedback to complete the task and could not rely only on motor planning. The points earned were inversely proportional to the time taken to complete each trial. Risk was manipulated by using horizontal perturbations to create the illusion of driving on a bumpy road, thereby imposing motor uncertainty. The baseline task contained five levels of uncertainty, including no additional variability. A subsequent task introduced high risk into the scenario by replacing grass on one or both sides of the road with water, which if hit would incur a very high penalty.

As expected, results depict position as a bimodal probability density function at low uncertainty, implying that subjects tended to keep towards the center of a single lane. As uncertainty increased, the peaks of the bimodal distribution tended toward one another. At high uncertainty, most subjects' position distribution exhibited a well-fit Gaussian function, indicating that they spent most of the time in the center of the road. This phenomenon was augmented when cost of error increased. Interestingly, this shift in behavior occurred even in the absence of errors, i.e. even if a subject never hit the side of the road at a particular uncertainty level (regardless of the cost), behavior was still significantly different when the cost increased. This is significant since the common model for learning in motor control is error-driven learning (Wei and Kording, 2008), and this observation suggests that human performance is often not driven by errors. The results demonstrate that subjects made predictions of both the likelihood and cost of failure, even if failure had never occurred, and are consistent with the existence of internal estimates of probability of failure and cost of failure.

The first experiment only investigated the role of motor uncertainty on behavior. We wanted to expand the paradigm of the first study to compare the effect of motor uncertainty (uncertainty in the control variable) with sensory uncertainty (uncertainty in the state variable). This has the same effect on control in motor control theories, however it is conceivable

that humans may perceive and interpret these uncertainties differently. The same iPad driving simulation will be used, but instead of physically slowing the car on the road, the cost function will be directly implemented by a point penalty. Subjects performed two tasks: one with the same imposed motor uncertainty used in the first experiment and the other with imposed sensory uncertainty. In order to implement sensory uncertainty, the contrast between the road and boundary was varied, and then the image was converted to 2-bit. In order to achieve this, the grayscale value of each pixel corresponded to a probability of being white versus black. Instead of contrast being calculated with Michelson’s formula, we transform this into a new equation that uses probability in place of luminance. Simple luminance contrast would be ineffective because the human eye can detect the difference between any two 8-bit colors, so the edge of the road would always be apparent. We characterized and validated this method of imposing sensory uncertainty in a previous experiment.

The first set of trials was a calibration phase, in which we matched the standard deviation of position when attempting to stay on a path under each type of uncertainty to specific levels. Subjects then completed the task when the cost function was the same as the first study (a bimodal cost function) and we compared subjects’ behavior under imposed sensory uncertainty to behavior under imposed motor uncertainty with equivalent statistics. Results showed that sensitivity to risk was significantly higher in response to visual uncertainty compared to motor uncertainty. This also allowed us to test whether human control of movement obeys the certainty equivalence property. A system in which the optimal solution is the same as the optimal solution for that system in the absence of uncertainty would be certainty equivalent. Certainty equivalence is the result of discrete-time centralized systems with only additive uncertainty and is a common assumption in motor control theories since it decreases computation complexity. However, the results of these studies demonstrate that this is not an appropriate assumption in models of human motor control.

From the first set of studies, it is evident that humans tune their statistical behavior

based on cost, taking into account entire probability distributions of possible outcomes with long tails in response to environmental uncertainty. In addition to modifying the control of movement to reflect the risk of the environment, we predict that humans will prepare for error in response to risk as well. Recent studies have demonstrated that humans have the ability to modulate the long latency stretch reflex based on the goal of the task, but typically utilize a simple go-don't-go paradigm or the goal is perturbation-dependent (Ludvig et al. 2007; Pruszynski et al. 2008). It is our hypothesis that awareness to risk is so fundamental that humans also maintain reflexes tuned specifically to the cost function of the environment, even when the goal of the task does not depend on the perturbation.

In the second set of experiments, the role of risk in tuning reflexes was examined. The first study extended the paradigm of the first study to include random visual displacements of the car. The accelerometer responses to perturbations increasing risk were compared to responses decreasing risk. A significant difference was found in the amplitude of response depending on risk. These perturbations were visual perturbations and therefore we were interested in seeing if this behavior extended to the stretch reflex as well.

The second two experiments in this series investigated the human stretch reflex response between risk conditions. In the first, the FDI reflex was studied. Healthy, adult subjects were positioned in front of a monitor with their right index finger attached to the arm of a robot that controlled a cursor on the screen. Surface electromyography from the first dorsal interosseous (FDI) was recorded (sample rate 1000 Hz, bandpass filter 25-250 Hz). The monitor displayed three rectangles: two cost regions on either side of a center reward region, which moved horizontally (remaining equidistant) in a randomized sinusoidal motion. Subjects were instructed to maximize points by keeping the cursor within the center reward region while avoiding the cost regions that would result in a loss of points. Nine cost environments were evaluated— all combinations of no penalty, low penalty, and high penalty— in order to evaluate the effect of both symmetric and asymmetric risk. Thus the goal of the task was always to remain in the center target, but the cost of hitting the penalty regions

was varied. The robot generated a constant 1 N force with randomized 4 N perturbations in both directions (randomized) at a mean rate of 3 seconds.

Only trials in the direction that provoked the FDI stretch reflex were analyzed. Reflex response was categorized into standard epochs for baseline, short latency, long latency, and voluntary response. The filtered EMG within each epoch was averaged to a single value for analysis. Short latency epoch was not significantly different between cost environments. A significant difference was found in the long latency epoch between cost functions. There was not a significant difference between cost functions that pushed toward higher cost versus away from higher cost. Therefore, results suggest that humans do plan for error by tuning reflexes to the risk of the environment, and that they do this independent of the goal of the task. However, subjects did not demonstrate the ability to set separate reflex responses for different directions when the direction of perturbation was unplanned. However, the results of this experiment were not as strong as expected, and we postulated that this might be attributed to the FDI. Therefore, we repeated the experiment using the bicep muscle.

The final experiment was very similar to the previous paradigm except a manipulandum was used to perturb the arm by applying torque to the elbow joint. The results from this study demonstrated more coherency. For the most part, the conclusions from the FDI reflex were confirmed. There was a significant increase in the long latency stretch reflex in response to increased risk. In fact, the average amplitude of the long latency stretch reflex between the symmetric high-cost was almost double that of the symmetric no-cost. Unlike the study on the FDI reflex, there was also a significant difference between the long latency reflex response to perturbations pushing toward higher risk compared to lower risk in asymmetric cost conditions when the cumulative risk was the same. This suggests that muscle stiffness was not the only method of modulating the stretch reflex since co-contraction is inherently symmetric.

The unifying aspect of these results is that they represent basic characteristics of human movement that are lacking or absent from current implementations of classical motor control

theories. Any complete emulation of human movement must reproduce these behaviors as well. The goal of these studies is not simply to demonstrate human behavior, but to persuade the reader to consider an alternative perspective on motor control that moves away from the traditional trajectory-based viewpoint and instead proposes that movement results from maintaining probability distributions of the probability of failure and cost of failure.



## **Dedication**

To my parents,

For their unwavering love and support.

And to Adam,

Without whom this book would have been written much sooner.



## Acknowledgements

First and foremost, I would like to thank my advisor, Terry Sanger, for teaching me more than I could have ever hoped to learn. There are no words that express how grateful I am for the past five years.

I would like to thank my lab including Cassie Borish, Shanie Liyanagamage, Enrique Arguelles, Won Joon Eric Sohn, Adam Feinman, Sirish Nandyala, John Rocamora, Sam Huynh, Ati Ghoreyshi, Matteo Bertucco, Scott Young, Minos Niu, Nasir Bhanpuri, Maryam Beygi, Shinichi Amano, Arash Maskooki, Diana Ferman, and Aprille Tongol for their help, support, and most of all friendship. There is no other group I would have wanted to go through this journey with.

I would like to thank my friends, for always being my willing subjects, no matter how boring the task. They helped me keep my sanity and perspective. Specifically I would like to thank Raisa Ahmad for over a decade of friendship and support, I would never have been an engineer without her.

I would like to thank my committee members, Francisco Valero-Cuevas, Jean-Michel Maarek, and Maryam Shanechi for their guidance and for taking time out of their busy schedules to participate in this process.

I would like to thank the Provost Fellowship for financial support, without which I would be living in a cardboard box. I would also like to thank the financial institutions that support our research.

I would like to thank my boyfriend for putting up with the long nights of work and studying, for picking up my slack when I need it most, and for being my partner in crime in everything we do.

And I would like to thank my parents, my biggest cheerleaders, who taught me I could do anything I put my mind to.



## Contents

### Chapter

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Background . . . . .	1
1.2	Specific Aims . . . . .	2
1.3	Theory . . . . .	3
1.3.1	Optimal Control Theory and Optimal Feedback Control . . . . .	3
1.3.2	Equilibrium Point Hypothesis . . . . .	4
1.3.3	Risk Aware Control . . . . .	5
1.3.4	Comparisons and Conclusions . . . . .	7
<b>2</b>	<b>Implementation</b>	<b>9</b>
2.1	Spiking Neuron Model . . . . .	9
2.1.1	Introduction . . . . .	9
2.1.2	Methods . . . . .	11
2.1.3	Results . . . . .	15
2.1.4	Discussion . . . . .	19
2.2	MATLAB Implementation . . . . .	22
2.2.1	Methods . . . . .	22
2.2.2	Results . . . . .	22
2.3	Spiking Neuron Implementation . . . . .	25

2.3.1	Methods . . . . .	25
2.3.2	Experiments . . . . .	27
2.3.3	Results . . . . .	29
2.3.4	Discussion . . . . .	33
<b>3</b>	<b>Human Motor Response to Risk</b>	<b>35</b>
3.1	Experiment 1: The Tuning of Human Motor Response to Risk in a Dynamic Environment . . . . .	35
3.1.1	Introduction . . . . .	35
3.1.2	Materials and Methods . . . . .	38
3.1.3	Results . . . . .	43
3.1.4	Discussion . . . . .	48
3.2	Experiment 2: Certainty Equivalence Assumption . . . . .	51
3.2.1	Introduction . . . . .	51
3.2.2	Methods . . . . .	51
3.2.3	Results . . . . .	55
3.2.4	Discussion . . . . .	57
<b>4</b>	<b>The Tuning of Reflexes to Risk</b>	<b>59</b>
4.1	Introduction . . . . .	59
4.2	Experiment 1: Response to Visual Perturbations . . . . .	61
4.2.1	Materials and Methods . . . . .	61
4.2.2	Results . . . . .	64
4.2.3	Discussion . . . . .	67
4.3	Experiment 2: Response to Mechanical Perturbation . . . . .	69
4.3.1	Materials and Methods . . . . .	69
4.3.2	Results . . . . .	73
4.3.3	Discussion . . . . .	76

4.4	Experiment 3: Role of Cocontraction in Tuning Reflexes . . . . .	79
4.4.1	Introduction . . . . .	79
4.4.2	Materials and Methods . . . . .	79
4.4.3	Results . . . . .	82
4.4.4	Discussion . . . . .	85
4.5	Chapter Conclusions . . . . .	87
<b>5</b>	<b>Concluding Remarks</b>	<b>89</b>
5.1	Conclusion . . . . .	89
5.2	Applications and Impact . . . . .	90
	<b>Bibliography</b>	<b>93</b>



## Figures

### Figure

2.1	Static Stochastic Sinusoidal Grating . . . . .	13
2.2	Psychometric Functions . . . . .	17
2.3	Mean Contrast Sensitivity Functions . . . . .	18
2.4	Probability of State for Left and Right Initial Conditions . . . . .	23
2.5	Probability of State for Larger Standard Deviations . . . . .	24
2.6	Probability of State for Cost to Move . . . . .	24
2.7	Control Diagram of the Robot and Risk-Aware Control . . . . .	27
2.8	Pictures of the robot set-up . . . . .	28
2.9	Tracking data of trajectories. . . . .	30
2.10	Stability of Stationary Tracking and Speed of Sinusoidal Tracking. . . . .	31
2.11	Response to perturbation. . . . .	33
3.1	Theoretical minimization of cost under uncertainty . . . . .	37
3.2	iPad Application Screen View. . . . .	39
3.3	Raw Population position data. . . . .	44
3.4	Continuous position probability distribution as a function of uncertainty. . .	45
3.5	Regression on the distance from the center of the road subjects maintained vs. level of motor noise. . . . .	46
3.6	Proportional hazard model of successful trials. . . . .	47

3.7	Application Images . . . . .	52
3.8	Results of Exemplary Subject For Calibration Session . . . . .	54
3.9	Histograms of Raw Position Data and Fits to Bimodal Distribution . . . . .	56
3.10	Distance from the Center of the Road vs. Uncertainty Level . . . . .	57
4.1	Types of visual displacements . . . . .	62
4.2	Specific Hypotheses . . . . .	63
4.3	Rapid Responses of Individual Subjects and Average of All Subjects . . . . .	64
4.4	Amplitude of Average Accelerometer Response to Push to Edge Perturbation	65
4.5	Amplitude of Average Accelerometer Response of the Push to Edge and Push to Center Perturbations . . . . .	66
4.6	Position of Car Post Lane Switch Perturbation . . . . .	67
4.7	Monitor Display and Set Up . . . . .	70
4.8	Change in EMG Between Conditions for Individual Subjects . . . . .	73
4.9	Symmetric Cost EMG . . . . .	74
4.10	Asymmetric Cost EMG . . . . .	75
4.11	EMG Normalized by Maximum Voluntary Contraction and by Preactivation	76
4.12	Manipulandum Set-Up . . . . .	80
4.13	Individual Subject EMG . . . . .	82
4.14	Reflex Response to Symmetric Risk . . . . .	83
4.15	Reflex Response to Asymmetric Risk . . . . .	84
4.16	Average EMG: Normalization . . . . .	85
4.17	Hypotheses Visualized . . . . .	86

# Chapter 1

## Introduction

### 1.1 Background

Human movement is performed using a musculoskeletal system that is redundant, non-linear, and constantly changing. It is controlled using feedback from a sensory system that is variably delayed and often imprecise. And it is executed in an environment that is full of potential harm and novel challenges. Nonetheless, human motor control reliably demonstrates success at generating complex, coordinated movements to achieve very difficult goals. Even using tremendous computational power, robotics is still unable to match the robustness, compliance, and flexibility of human movement in real time. How it is that humans accomplish this has been a widely debated and pursued topic of research (Diedrichsen, 2010; Haith and Krakauer, 2013).

There are several proposed models of motor control, which predominantly focus on achieving a goal through a reference trajectory (Todorov 2002; Todorov 2004). However, motor control is not just about reaching a goal or accomplishing a task, but also avoiding predictable failure in the process. Risk is inherent in all activity, and avoidance of risk is fundamental to human survival, so response to risk must be an integral part of human movement as well. In this report, we describe risk specifically as a combination of two factors: the probability of failure and the cost of failure. A high cost of failure but low probability is not generally considered risky (standing several meters away from the edge of a cliff). Likewise, a high probability of failure but low cost is not regarded as risky (standing on the

edge of a step). It is only where high likelihood of failure converges with high cost, when we stand on the edge of the cliff, that we venture into high risk.

A new theory, Risk-Aware Control, emphasizes selecting motor actions that minimize risk (Sanger, 2014). Risk-Aware Control is distinctive from classical theories in that it is an entirely new way of approaching the relationship between cost and motor actions. It does not attempt to formulate a reference trajectory to the goal, but instead predicts that movement develops from maintaining a probability distribution of state, a detailed understanding of the cost function, and knowledge of the relationship between action and change in state. The result is a control theory that accounts for uncertainty and cost in motor planning and execution, plans for unexpected error prior to perturbation, and does not require assumptions of system linearity. This report goes through the risk-aware control theory and demonstrates an implementation used to control reaching in a desktop robot. It then details a series of experiments designed to characterize and emphasize the influence of risk in the environment on human behavior.

## **1.2 Specific Aims**

The aim of this report is to articulate a new theory in motor control, termed risk-aware control (Sanger, 2014), to simulate the theory, and to present the results from a series of experiments that demonstrate the existence of analogous behavior in humans. The stages of implementation will be described, ultimately developing into a spiking neuron model of cortex to control a robotic arm in real time.

Additionally, a series of experiments were designed and executed to demonstrate fundamental characteristics of human behavior that is exhibited by risk-aware control but is lacking or absent from the classic motor control theories. These experiments were divided into two categories of study. The first series looked at the effect of probability of failure and cost of failure on behavior in a continuous environment task and the effect of errors, or lack of errors, on behavior. The second set of studies investigated the role of risk in tuning

reflexes and provided physiological evidence that humans plan for error.

The goal of this paper is to persuade the reader to consider an alternative perspective on motor control that moves away from traditional trajectory-based viewpoint and instead proposes that movement results from maintaining probability distributions of the probability of failure and cost of failure.

### **1.3 Theory**

In this section we will provide a very brief overview of two leading motor control theories, optimal feedback control and equilibrium point hypothesis. We will then describe a newer theory, risk aware control (Sanger, 2014), and highlight the differences between all these theories. Risk-aware control is a novel formulation of a classic feedback controller that replaces the state variables with probability densities and the control trajectory with a dynamic cost function defining both penalty regions and rewards. The result is a control theory that allows for uncertainty in state and control variables and is not based on assumptions of system linearity. The theory of risk-aware control suggests a reduction in computational burden compared to current full models of human motor control because it allows for parallel computing.

#### **1.3.1 Optimal Control Theory and Optimal Feedback Control**

Optimal Control Theory is perhaps the most widely accepted prediction of free human movement. In general, models of optimal control explain a large class of behavior. Existing variations of optimal control usually optimize a variety of specific cost functions. The form of the cost will depend on the goal of the task as well as a regularization term that constrains undesired features of movement (such as jerk or integrated torque change) (Diedrichsen, 2010). Optimal control can be categorized into two classes of models: open-loop and closed-loop.

Traditional optimal control uses feed-forward motor commands to execute a precal-

culated trajectory to minimize the specific cost function. Open-loop optimal control was extended to incorporate sensorimotor feedback in a closed-loop model, optimal feedback control (Todorov and Jordan, 2002). The feedback driven model transforms the current state estimate resulting from feed-forward (efferent copy of motor commands) and feedback (afferent sensory signals) into a new motor command. In optimal feedback control, the trajectory minimizing the cost function is recalculated at every time point so that a new best path is recalculated after any deviation from the previously calculated path.

An advantage of this theory is that it also accounts for observations of variability between trials and the uncontrolled manifold phenomenon. It also inherently solves the problem of mechanical redundancy and trajectory redundancy (Todorov, 2004).

### **1.3.2 Equilibrium Point Hypothesis**

Studies have shown that muscles possess the same properties as nonlinear springs with adjustable stiffness or adjustable threshold length (Shadmehr and Arbib, 1992; Shadmehr, 2010). Two springs acting antagonist to each other will naturally reach some point of equilibrium. This equilibrium point can be manipulated by controlling the stiffness of the springs (muscles). This is the basis of equilibrium point hypothesis (referent configuration hypothesis), which postulates that voluntary and involuntary movement arises from a trajectory of these equilibrium points (Latash 2010a).

Consequently, this theory asserts that descending motor commands do control force directly, but can only influence the set point of a local feedback circuit. The result is a nervous system that effectively uses the stretch reflex to control movement. Thus, equilibrium point hypothesis advocates that movement is an emergent property of the motor system and cannot be prescribed by any neural controller (Glansdorff and Prigogine, 1971; Latash, 2010a). The system itself is controlled by setting specific parameters and the output is a result of the interaction of those parameters with the system itself and outside dynamics.

### 1.3.3 Risk Aware Control

Risk-aware control (Sanger, 2014) governs movement based on estimates of risk, consequence of state and outcome uncertainty, as well as expectations of the cost of errors. Fundamentally the term risk is used to describe the expected cost of action. High risk occurs when high cost of failure converges with high probability of failure, as described in the introduction. Probability of error is regulated by unpredictability or uncontrollability of effect of actions or current state and cost of error is defined by predictions of the environmental cost and movement objective. The goal is to maximize reward or decrease cost. The novelty of risk-aware control is that it does not attempt to formulate or follow a trajectory at any point. Instead movement is the result of maintaining estimates as entire probability distributions.

The following theory is described in greater detail in Risk-Aware Control (Sanger, 2014). In risk-aware control, the state estimate,  $x$ , is replaced by a probability density changing in time describing the belief of state,  $p(x,t)$ . State is updated continuously according to the equation

$$\frac{\partial p(x,t)}{\partial t} = Lp(x,t) \quad (1.1)$$

where  $L$  is a linear operator describing a change in state. Therefore,  $L$  will be dependent on the specific choice of action,  $u$ . A common instance of such a linear operator is the Fokker-Plank equation describing a drift and diffusion process. An equivalent model for physical systems takes the form of the Ito stochastic differential equation, which often has nonlinear components.

An important characteristic of  $L$  is that since it is a linear differential operator, the effects of multiple operators can be summed, even if the underlying stochastic dynamics are nonlinear. This is central to the risk-aware control theory as it means that a superposition of different dynamics can be constructed out of the operators,

$$\dot{p} = \left( \sum_{i=1} u_i L_i \right) p \quad (1.2)$$

where  $u_i$  is a set of nonnegative weighting coefficients. A classical feedback controller can be constructed by making the control variable,  $u$ , dependent on the state variable,  $x$ . Instead of a trajectory, risk aware control implements a cost function,  $v(x)$ , that defines both the rewards and potential dangers of being in each state,  $x$ . Expected value can thus be calculated  $E[v] = \int v(x)p(x)dx$ . In order to maximize total expected value, the rate of change in expected value can be computed as

$$\frac{\partial E[v]}{\partial t} = \int v(x)\dot{p}(x)dt \quad (1.3)$$

which simplifies to  $vLp$ . In the most basic implementation, we select the action,  $u$ , by maximizing the expected change in reward (or minimizing cost) at any time point.

$$u(t) = \operatorname{argmax}_u v_t L(u) p_t \quad (1.4)$$

Furthermore, instead of selecting a single action, the superposition dynamics can be maximized by setting the weighting factor for each  $L$  operator proportional to the (positive) change in expected value corresponding to that operator. This is described by the equation

$$u_i = v^T L_i p \quad (1.5)$$

Since we are modeling a physical system, it can be assumed that state is continuous and can make no instantaneous jumps. Consequently,  $\dot{p}(x, t)$  will be nonzero only where  $p(x, t)$  is nonzero and the  $L$  operator will be near diagonal. Therefore, if these assumptions hold, equation 3 can be rewritten as

$$u_i(t) = \sum_k u_{ik}(t) = \sum_k v_k (l_{k,k-1}^i p_{k-1}(t) + l_{k,k}^i p_k(t) + l_{k,k+1}^i p_{k+1}(t)) \quad (1.6)$$

With the exception of the sum-reduction (which itself is largely local because nonzero elements will be clustered) all these operations become local operations. This is important because it means that neurons will only need information from and to interact with other neighboring neurons to contribute according to the overall goal of the system. A more detailed description of the risk aware control theory and the stochastic dynamic operators has been described in papers (Sanger 2011; Sanger, 2014). One of the most notable strengths of risk aware control is its computational efficiency. Unlike Optimal Feedback Control, it can control complex movement in real time using computers with ordinary processing power.

### 1.3.4 Comparisons and Conclusions

In terms of output, risk aware control will share many characteristics with optimal feedback control. However, optimal control theory calculates an optimal trajectory and then follows that trajectory with standard feedback control. This means that a perturbation from the trajectory will result in movement back toward the trajectory regardless of the risk resulting from the perturbation. This is important because setting reflexes is unnecessary to follow a reference trajectory. However, when controlling the system through dynamics, such as in risk aware control, tuning reflexes is an inherent result of the system. This may be observed as planning for unexpected errors, which optimal feedback control does not inherently exhibit. However, it may be possible for tone to be adjusted appropriately in OFC by layering multiple feedback loops operating simultaneously on top of each other and controlling separate gain variables (Todorov, 2004).

Optimal feedback control must also make many underlying assumptions to simplify the computational burden of reoptimizing the trajectory at every time point. Currently, no implementations include state uncertainty or action outcome variability or complicated loss functions (Sanger, 2014).

A stated advantage of equilibrium point hypothesis is its dependence on physics and physiology instead of an evolution of control theory and robotics (Latash, 2010b). However,

there is also evidence that appears to refute equilibrium point hypothesis. One shortcoming of equilibrium point hypothesis is that it cannot explain the ability of people with proprioceptive loss to make voluntary movements (Shadmehr, 2010). Risk aware control updates the probability of state using both a predictive term and sensory feedback term, which accounts for how this population can still make accurate rapid arm movements toward a goal. They may create an accurate internal model and utilize other forms of sensory feedback.

These are only a couple of the theories of motor control that have been proposed. Certainly it seems that every motor control theory has evidence of being both highly supported and highly contradicted. It is very probable that specific counterexamples will be able to be found for any proposed theory of motor control, present and future, as a result of the flexibility and adaptability of humans. However, it cannot be denied that there are very fundamental characteristics of human movement that are still unaccounted for by these (and other existing) motor control theories. Ultimately, we propose that this may be the result of approaching the problem of motor control from the wrong perspective. Non-trajectory based motor control, such as risk aware control, still exhibits all the desirable characteristics of current motor control theories (perhaps with the exception of complex path finding that OFC solves and the current absence of physiological descriptors that EP possesses) while additionally describing an assortment of characteristics lacking from current models. The rest of this report will consist of simulations of risk-aware control and a series human behavioral studies that highlight the appropriateness of risk aware control as a model of free human movement.

## Chapter 2

### Implementation

The ultimate goal of the theoretical work was to implement risk-aware control on a biologically-realistic distributed network of spiking neurons. Therefore, the first step was to determine the representation of information. The first section explores the sufficiency of rate coding as a method of neuronal communication. The second and third sections detail an implementation of risk-aware control.

#### 2.1 Spiking Neuron Model

*A version of this section was prepared to be submitted Psychophysics, Attention, and Perception.*

##### 2.1.1 Introduction

It has been widely established that information can be coded in the average rate of spikes in a neural signal. However, it is less clear if a train of spikes carries additional information in the precise temporal pattern of firing or in the relative timing of spikes in different cells. Therefore, it was first important to determine if it was necessary to incorporate spike timing or if rate coding is sufficient to transmit information and biologically realistic. In this study, we investigate the fraction of information transmitted by spike rate alone in the visual system by examining the detection curves of images with pixels that flicker according to a Poisson distribution.

Almost all information in the mammalian central nervous system is transmitted as neural spike events, yet controversy remains as to whether the average rate of spikes is sufficient to capture the meaningful information or whether the detailed pattern and timing of individual spikes might carry additional information (Adrian, 1926; Bhumbra and Dyball, 2005; Knight, 1972; Meister and Berry, 1999; Rullen and Thorpe, 2001; Shadlen and Newsome, 1994; Softky, 1995; Stein et al., 2005). While the answer may be different for different parts of the brain, we investigate here the extent to which firing rate is sufficient to transmit information about contrast in visual images.

A sequence of spikes that is Poisson-distributed has the property that the number of spikes in any interval is independent of the number of spikes in any other interval, and thus information is carried only in the average spike rate. We generate images in which each pixel flickers with a Poisson distribution whose average rate is given by the desired contrast. The observed image contrast is always maximum, because pixels are either fully on or fully off. But the information carried in the rate can be more nuanced, representing varying degrees of contrast for more complex images. Since the original image contrast is only encoded by the rate, the contrast sensitivity will reflect the brain’s ability to extract information from rate-coded representation.

Rate coding presents a computational difficulty however. The brain must have a mechanism for decoding, or extracting the rate from a sequence of spikes. It is often assumed this is done by linear filtering, both in time and in space, so that sequential spikes are combined and spikes representing neighboring regions of the image are averaged (Kilikowski and King-Smith, 1974; Thibos, 1989). This corresponds to both spatial and temporal low-pass filtering. Linear filtering predicts two phenomena that we will test: (1) Spatial low-pass filtering will reduce contrast sensitivity at high spatial frequencies, (2) Temporal low-pass filtering will increase contrast sensitivity when the flickering rate-coded image can be viewed for a longer time. Our results will reject both of these predictions, suggesting that the brain’s mechanism of perception behaves as a nonlinear filter so that sharpness and rapid temporal

responses are not lost in the decoding process.

We test the psychometric properties of spatial-frequency grating images in which contrast is represented by the Poisson rate. If the psychometric curves for such images parallel those for normal grayscale images, then this supports the claim that contrast can be represented by rate coding. If the images retain sharpness at high frequencies, this contradicts spatial low-pass filtering as a decoding method. If the contrast sensitivity does not improve with longer presentations of the flickering images, this contradicts temporal low-pass filtering as a decoding method.

### 2.1.2 Methods

#### *Participants*

The study consisted of 9 adults (4 males, 5 females), ages 23 to 28, with normal or corrected-to-normal vision. Sample size was designed to detect approximately a 1 standard deviation difference with 80 % power. Subjects provided consent as approved by the University of Southern California Internal Review Board and received compensation for their time.

#### *Apparatus and Stimuli*

In addition to the typical achromatic analog sinusoidal gratings (Campbell and Robson, 1968), this study implemented rate-coded sinusoidal gratings. In these images, each pixel contained only one bit of information per frame, i.e. each pixel was either white or black (100% contrast). The information of the image was coded in probabilistic contrast instead of luminance contrast. Probabilistic contrast, the probability of a pixel being either white or black, was proportional to luminance contrast. This is a type of dithering, random dithering (Russ, 2016), generates a one bit-per-pixel image in place of the original analog image. When a sequence of such dithered images is generated from the same original analog image by choosing pixel probabilities independently for each frame, the time-averaged luminance

of the image sequence is the same as the luminance of the original analog image. Classical experiments use Michelson’s formula (Michelson, 1927) to calculate contrast of an image:

$$C = \frac{L_{max}(x) - L_{min}(x)}{L_{max}(x) + L_{min}(x)} \quad (2.1)$$

where  $L_{max(x)}$  is the luminance at the peak of the sinusoid,  $L_{min(x)}$  is the luminance at the minimum and  $C$  is the contrast of the image. The probabilistic contrast values referenced in this paper are derived from the same basic equation:

$$C = \frac{P_{max}(x) - P_{min}(x)}{P_{max}(x) + P_{min}(x)} \quad (2.2)$$

In this case,  $P_{max(x)}$  and  $P_{min(x)}$  are the probabilities that a pixel will be on at the maximum and minimum of the sinusoid respectively. We will refer to these images, such as figure 2.1, as stochastic images. Since pixels are chosen randomly, each stochastic image will be unique while still derived from the same original image. We refer to a sequence of stochastic images derived from the same original analog image as a “dynamic stochastic image”. We will measure the contrast sensitivity for dynamic stochastic images comprised of a sequence of 5 stochastic images presented in a cyclic series at 30 frames per second (fps).

Stimuli were presented over a black background on 24-inch, 1080p resolution light emitting diode screen (DELL 1280x1024 maximum resolution). This screen was limited to 32-bit color. At this color resolution, the edge between two regions of adjacent gray levels (differing only in the least significant bit) is visible with normal vision. Pilot studies indicated that subjects could detect these edges, resulting in an artificially high spatial grating sensitivity at low contrast. In order to simulate a true analog contrast grating, the analog images were generated using a “noisy-bit” method characterized by Allard and Faubert (Allard and Faubert, 2008). This method involves adding uncorrelated noise to the contrast of each color channel of a pixel in order to soften the visible edges between color blocks. Allard and Faubert verified that the noisy-bit method is perceptually equivalent to

a continuous display and does not significantly impact the contrast sensitivity function. It should be noted that while this method is a type of dithering, the technique and outcome are very different from the dithering implemented in our stochastic images. The noisy-bit method adds the smallest possible increment of white noise to each analog color channel. Therefore it affects only the least significant bit for each channel, leaving the remaining bits unchanged, whereas in our stochastic images each pixel is fully on or off and information is carried only in the probability. (A noisy-bit image clearly carries more information per frame than our stochastic images, but it cannot be used to assess for rate coding in our experiments, because temporal coding, including details of the spike sequence or timing, could potentially be used by the brain to encode the contrast of the non-dithered bits of the image.)

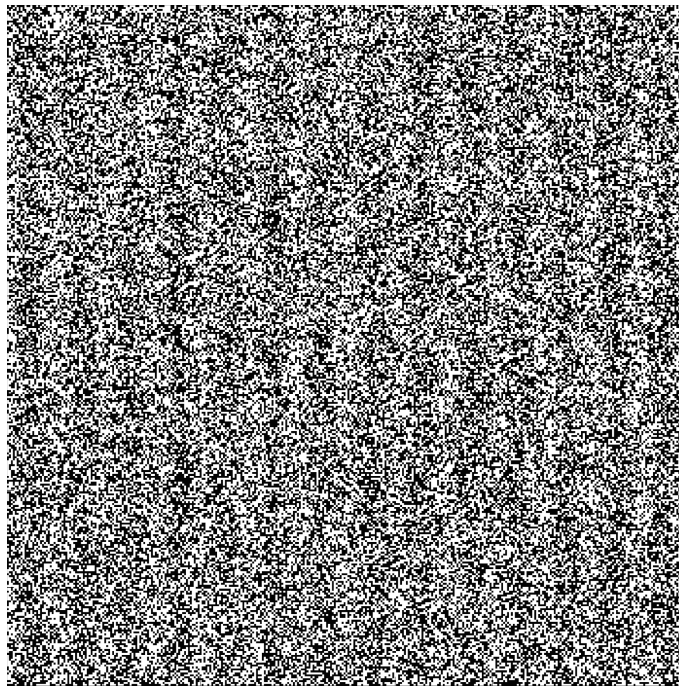


Figure 2.1: **Static Stochastic Sinusoidal Grating.** Subjects were presented with images such as the one above. The above figure is a quarter panel of the experimental images (350x350 pixels, 3.82x3.82 inches). This is an example of .125 probabilistic contrast and 4 cycles per degree.

In order to ensure more than one sinusoid cycle was visible for all spatial frequencies,

each original analog image was generated at 700x700 pixels (7.64x7.64 in). The sinusoidal grating varied in frequency, probabilistic or luminance contrast, and orientation. The frequencies implemented were .0920, .1840, .3680, .7361, 1.4721, 2.9443, and 5.8885 cycles per degree; the probability contrasts were .25, .125, .0625, .0313, .0156, and .0078 for the stochastic images and the deterministic contrast was 0.0625, 0.0313, 0.0156, 0.0078, 0.0039, and 0.0020 for the analog images. The direction of the grating was either vertical or horizontal. The frequencies were chosen from pilot studies to visualize the entire contrast sensitivity function and the contrast values were selected such that the stimulus could never be detected at the minimum and could very clearly be detected at the maximum. Analog (noisy-bit) images, static stochastic images, and dynamic stochastic image sequences were compared. Each of these conditions was repeated 4 times for a total of 1008 trials (2 orientations x 4 repeats x 6 contrasts x 7 frequencies x 3 types of images).

### *Procedure*

A two-alternative forced choice (2AFC) method was used (Blackwell, 1952) with automated experiment software written in MATLAB (version 7.13.0.564. Natick, Massachusetts: The MathWorks Inc., 2011). The screen was the only significant light source in an otherwise empty room. Subjects were positioned eye-level and centered on the image at a viewing distance of 22 in. Subjects were not restrained with a chin rest, but were asked to maintain the same position throughout the entire experiment to the best of their ability. Lack of significant movement was confirmed by direct observation. The stimulus resolution was 37 pixels/degree.

The experiment was divided into two 30-minute sessions in an attempt to facilitate a more constant attention level. Subjects viewed each image for 2 seconds, followed by 2 seconds of black screen during which subjects were prompted to identify the orientation of grating (vertical or horizontal). Subjects were told that no response would automatically be considered incorrect. These instructions were printed on an information sheet given to

subjects and reiterated by the experimenter prior to testing. Subjects were not provided with feedback on correct or incorrect responses.

### *Data Analysis*

All subjects neglected to answer some number of trials, so for analysis purposes, if the subject did not respond in the allotted time, the trial was counted as half correct. The underlying psychometric function of sensory perception cannot be observed directly, but is inferred from the empirical data. Therefore it is necessary to utilize certain analysis techniques in order to estimate the true psychometric functions and resulting contrast sensitivity function.

To construct psychometric curves, the raw data were fitted using a maximum likelihood estimation to the normal cumulative distribution function with an adjusted baseline.

$$F(x) = \frac{1}{4} \left( 1 + \operatorname{erf} \frac{x - \mu}{\sigma \sqrt{2}} \right) + \frac{1}{2} \quad (2.3)$$

In equation 2.3,  $x$  is the contrast,  $\mu$  is the mean of the normal distribution (shift of the psychometric function),  $\sigma$  is the standard deviation of the normal distribution (steepness of the slope of the psychometric function),  $\operatorname{erf}$  is the Gaussian error function, and  $F(x)$  is the probability of correct response. Threshold was determined using the nonparametric Spearman-Kärber method (Karber, 1931; Miller and Ulrich, 2001; Miller and Ulrich, 2004; Spearman, 1908) at each spatial frequency for each subject. All data and statistical analyses were done in MATLAB and R (A Language and Environment for Statistical Computing, version 3.0.1. Vienna, Austria: R Development Core Team, 2013).

### **2.1.3 Results**

While the psychometric function is arguably the most fundamental and widely used tool in visual psychophysics, it is conceivable that the normal psychometric function may

not be a well-suited analysis tool for stochastic images. Therefore, to assess the shape of the stochastic detection curves, the fits of the psychometric function were compared between image types. The averaged means of the R-squared values for the analog, static stochastic, and dynamic stochastic images respectively were 0.7369, 0.8213, and 0.8284. This suggests that the stochastic images are at least as well fit, if not better fit, to the cumulative normal distribution as the analog images within the constraints of this study.

The Spearman-Kärber method, which makes no assumptions about the underlying distribution of the data and provides more accurate estimates of location and dispersion parameters, was selected for its generally superior performance regardless of the underlying distribution (Miller and Ulrich, 2001; Miller and Ulrich, 2004). Within this study, the threshold values of individual subjects demonstrated more consistency implementing this analysis technique. However, the Spearman-Kärber method only quantifies the moments of the distribution, it does not provide a continuous estimate of the psychometric function.

A two-way repeated measures ANOVA comparing sensitivity (inverse of threshold) to all three types of images and spatial frequencies showed that there was a significant effect of image type ( $p < 0.001$ ,  $F(2,189)=40.13$ ) and spatial frequency ( $p < 0.001$ ,  $F(6,189) = 16.55$ ) as well as a significant effect of the interaction ( $p < 0.001$ ,  $F(8,189) = 8.46$ ). A two-way repeated measures ANOVA comparing between the two types of stochastic images revealed no significant difference in type ( $p=0.582, F(1,126)=8.2$ ) and no significant difference in spatial frequency ( $p = 0.176, F(6,126)=1.858$ ). In short, the sensitivity to stochastic images was significantly different from analog images, while sensitivity between the static stochastic and dynamic stochastic images was not, as shown in figure 2.3. Post hoc pairwise t-tests were used to compare the sensitivity between images at each spatial frequency. None of the static stochastic-dynamic stochastic pairs were significantly different ( $p < 0.05$ ). There was a significant difference ( $p < 0.05$ ) between sensitivity to the stochastic and analog images at the four highest spatial frequencies.

## Psychometric Functions

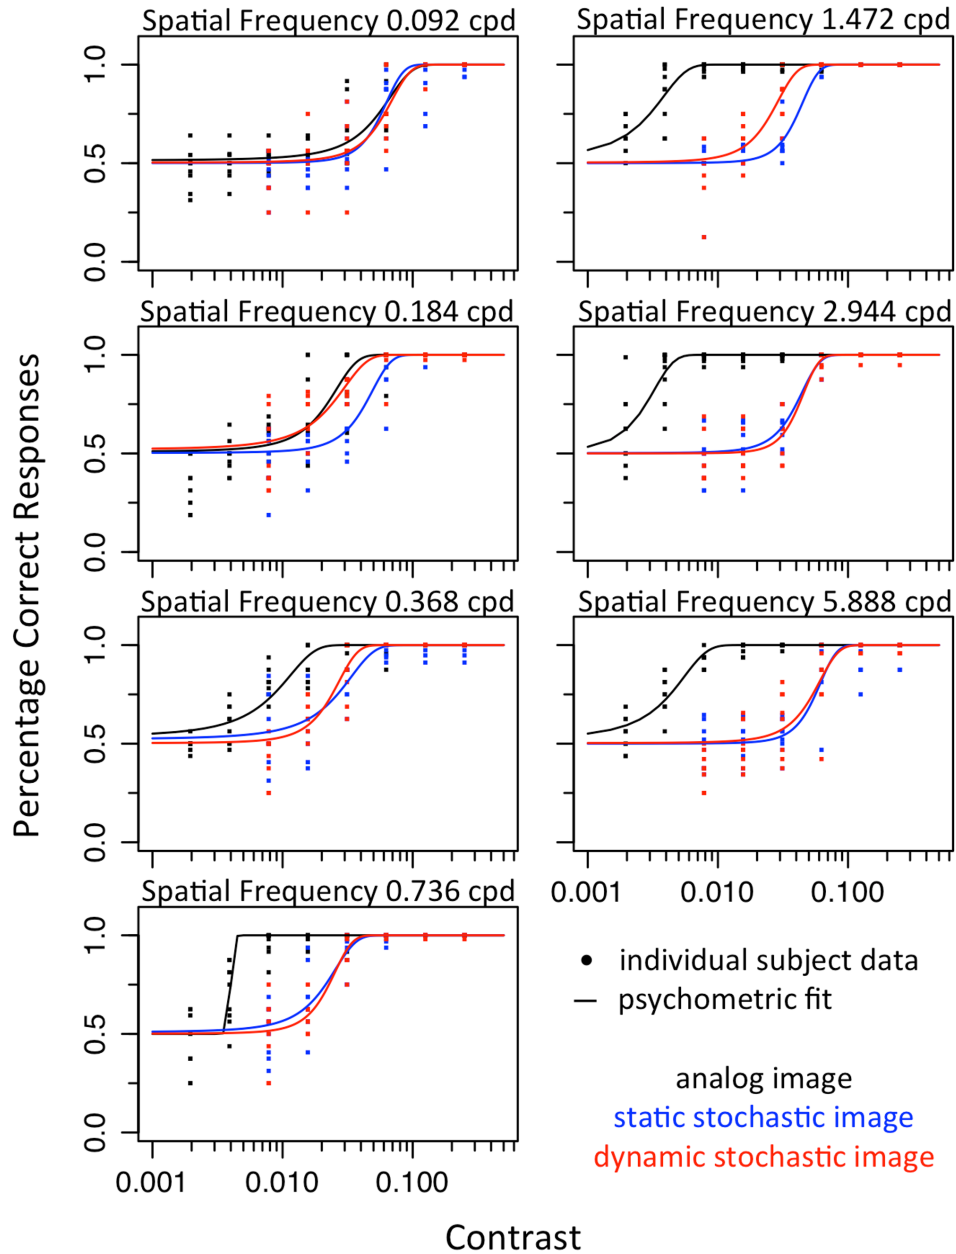


Figure 2.2: **Psychometric Functions.** These figures contain the psychometric functions, fit with equation 2.3, to the pooled raw subject data for each image type. The dots indicate the percentage of correct responses for each subject under each condition. The black, red, and blue lines indicate the fit resulting from the analog, static stochastic, and dynamic stochastic images respectively. The y-axis represents percentage of correct responses and the x-axis is contrast (calculated by equation 2.1 or 2.2) on a logarithmic scale. Above each panel is the spatial frequency value of the visual stimulus.

Nevertheless, while the sensitivity was significantly different between stochastic and analog images, the shapes of the detection curves were remarkably similar. All three contrast sensitivity functions peak at nearly the same frequency. This point is illustrated by multiplying the entire analog contrast sensitivity curve by 0.4. A two-way repeated measures ANOVA between the sensitivity to the static and dynamic stochastic images and sensitivity ( $\times 0.4$ ) to the analog images is no longer significantly different ( $p = 0.875$ ,  $F(2,189) = 0.133$ ). Moreover, in post hoc tests, not a single pairwise t-test between image types within spatial frequency was significantly different ( $p < 0.1$ ).

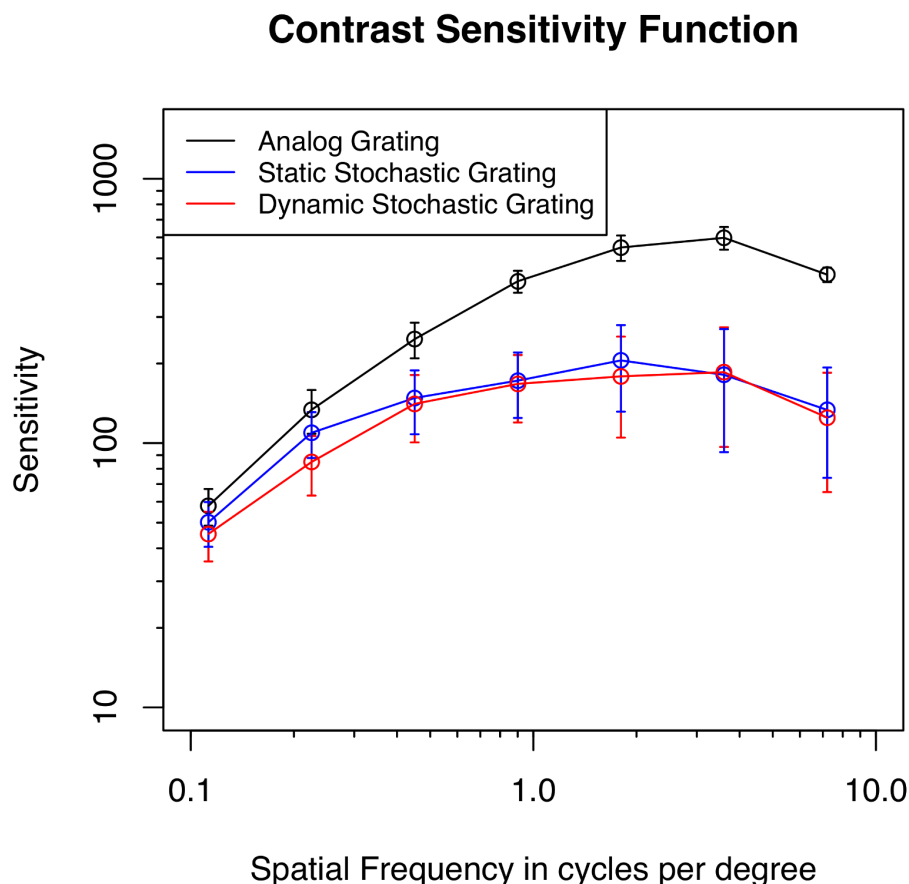


Figure 2.3: **Mean Contrast Sensitivity Functions  $\pm$  SE** Figure illustrates the contrast sensitivity function for each image type. The black, blue, and red circles indicate the sensitivity to the analog, static stochastic, and dynamic stochastic images respectively. The bars designate standard error at each point. Spatial frequency is represented by the x-axis and sensitivity on the y-axis. Both scales are logarithmic.

#### 2.1.4 Discussion

Most classical studies have reported peaks in the contrast sensitivity function between 2 and 7 cycles per degree (cpd) for normal human vision (Campbell and Robson, 1968; Owsley, 2003; Van Nes et al., 1967), and our measured peak of the analog contrast sensitivity function in this study is consistent, with a value between 1.8 and 7.2 cpd. The peak of the contrast sensitivity function for each of the three types of image occurred at approximately the same spatial frequency. The psychometric functions resulting from the stochastic images has a similar shape but with lower sensitivity compared to the original image. Part of the reason for the difference in sensitivity is that the stochastic images introduce high-frequency spatial and temporal noise due to the quantization and the highly visible pixel boundaries throughout the image. In addition, any single stochastic image contains less information (only one bit per pixel) than the original (8 bits per pixel). The fact that the psychometric curves have similar shapes suggests that the brain is capable of processing rate-coded data.

The most common decoding method for rate-coded data is to count the spikes over a period of time or average the spikes from neighboring pixels in order to obtain an estimate of the mean spike rate. However, averaging nearby pixels is a spatial low-pass filter that would be expected to have a specific deleterious effect on contrast sensitivity at high frequencies. No such effect was seen in our data, and the psychometric curves for the dithered images parallel those for the original analog images. Similarly, averaging the value of a single pixel over a period of time will eventually recover the exact analog value of that pixel, with better estimates as averaging occurs over a longer period of time. No change in the sensitivity with longer periods of time was seen in our data. Therefore our results are inconsistent with the use of either a spatial or temporal contrast sensitivity.

Diffraction, optical imperfections, and retinal issues cause a decrease in spatial frequency contrast sensitivity similar to a low-pass spatial filter (Campbell and Green, 1965; Thibos, 1989). However, since subjects can easily identify individual pixels and pixel edges

in our stochastic images, any inherent low-pass spatial filtering at the retinal level has cutoff frequencies higher than the spatial resolution of a single pixel in our images. We conjecture that the reason the contrast sensitivity function for the stochastic images parallels that for analog images is that both types of images are rate-coded in the same way by the retina. Each pixel in the analog image will cause a group of retinal ganglion cells to fire, and if this firing is rate-coded then the two types of images have no functional difference. This is analogous to the way in which a full-spectrum color image and an image coded only using red, green, and blue pixels both produce the same color percept, because both types of images yield the same retinal output.

The stimulus resolution was 37 pixels per degree or roughly 1,500 pixels per square degree. The fovea contains approximately 17,500 cones per square degree (Kolb, 2005). This means that at the fovea, there are just over ten cones per stimulus pixel. Under the rate-coding hypothesis, a single pixel from the analog image will stimulate ten cones each of which will fire independently at a rate proportional to the contrast. In the stochastic image, all ten respective cones will be on or off together. This discrepancy in resolution could account for up to approximately a magnitude of difference between sensitivity to the stochastic images and the analog images.

The difference between the analog and stochastic images is clearly visible. This is partly caused by the flickering itself being visible, because the 30hz monitor update is below the critical flicker frequency for much of the image. In addition, the dithering process introduces high spatial-frequency noise even in regions that were constant in the original analog image. Therefore although the stochastic images are conjectured to generate retinal ganglion output similar to the analog images, there is additional retinal ganglion output due to unavoidable introduction of high spatial frequencies by the dithering process. Even if this is not a cause of degraded contrast sensitivity, it will produce a perceptual difference between the two types of image.

Very high standard deviations in a local region of an image are rarely, if ever, found

in the natural world (Attneave, 1954). There is no reason that humans should be good at interpreting this type of image. However, if visual information is carried in the rate of spikes, then the stochastic images are simply translating the visual image into an understandable neural code and bypassing part of the retinal encoding process. The results establish that rate coded spike data is largely sufficient for a full contrast sensitivity function. Furthermore, if Poisson-distributed spike trains are decoded with linear filters, then information present in higher frequencies will be lost. While there was a significant decrease in sensitivity between the stochastic stimulus and analog stimulus, this reduction was relatively proportional across the span of spatial frequencies. Our results are thus inconsistent with spatial or temporal linear filters as a decoding mechanism, and support the possibility that the brain uses a nonlinear filter to extract information.

We have built an iOS application that utilizes the camera feed and in real time dithers the image following the methodologies of this paper. (The application, “BabyCatnip”, can be downloaded from the Apple Inc. App Store.) The content of the original images is clearly visible, and object recognition, reading, and motion perception are possible. We conjecture that because of the high contrast of the dithered images, they are particularly engaging for infants and could potentially be useful for improving visual perception in people with decreased contrast sensitivity. Testing the perceptual thresholds for object recognition and motion and testing the potential utility of such images for patients with retinal disease will be topics of future research.

## 2.2 MATLAB Implementation

In order to demonstrate and quantify the characteristics of risk-aware control, we wanted to implement the theory in a real-time biologically-realistic system. Prior to a full implementation, we performed a basic proof of concept in MATLAB. This demonstration is the most basic form of the control theory and was not performed in real-time.

### 2.2.1 Methods

The distributions were organized by spatial representation, i.e. each index represented a specific location in space. In this implementation, all values and calculations were performed as floating point numbers. The probability distribution of state,  $p(x)$ , had an assumed constant standard deviation. The update equation, eq 1.2, was used to update the probability of state given a particular action. The possible actions in this case were move left, move right, or no movement, and were determined by equation 1.4. The gain was not proportional to the change in expected cost, each action was all-or-nothing, so it is important to note that this implementation will not have reflexes tuned to the risk.

The cost function was described by an image similar to a road, with a large cost outside the boundaries of the road and a smaller cost between lanes, such as 3.1. In the figures below, cost was represented by darkness; black was high cost and white was reward or no cost. While the cost function is represented by a two-dimensional image, the cost function was incorporated as a 1-dimensional spatial cost function changing in time. An analogy to this is driving at a constant speed, at any point in time the car can move left or right on the road, perpendicular to the automatic forward motion of the car.

### 2.2.2 Results

No experiments were performed with this implementation, only the demonstrations show below. The figures illustrate the probability distribution of state as the cost function

changes with time. The first set of figures demonstrate the effect of initial conditions on the movement. The figure on the left, 2.4(a), starts at an initial maximum likelihood of 40 pixels and the right figure, 2.4b, starts at 170 pixels. In each case, the maximum likelihood moves toward the center of the closest reward channel and remains there. In this example, the standard deviation is relatively low (10 pixels), so the maximum likelihood (the dark red line) remains nearly in the middle of the edge boundary (high cost) and centerline boundary (lower cost).

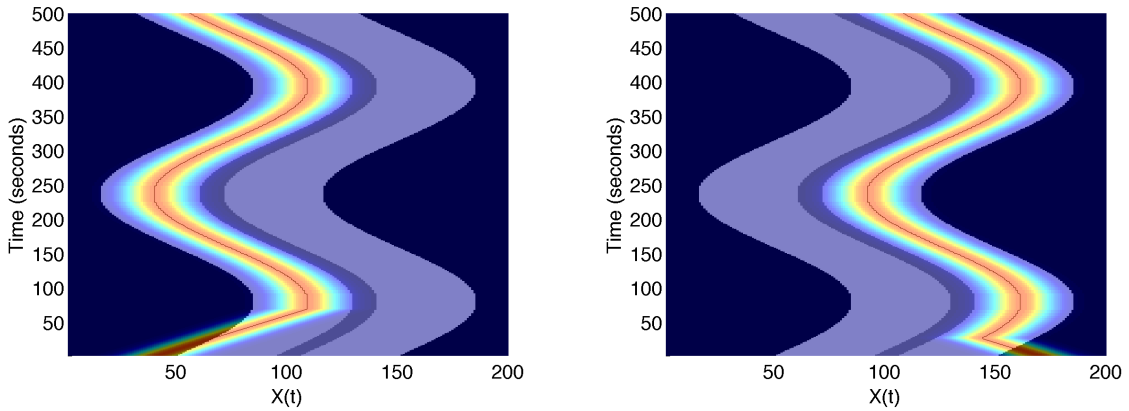


Figure 2.4: **Probability of State for Left and Right Initial Conditions.** These figures illustrate the probability distribution of state as the cost function changes with time. This can be imagined as driving, the movements are only 1-dimensional (in the horizontal plane), but the cost function of 1-dimensional state is changing with time. The figures demonstrate the effect of different initial conditions. In the figures, the colors represent the probability distribution of state (high probability in red and low probability in blue). The dark line indicates the maximum likelihood. The simulation demonstrates that the probability distribution of state moves to the nearest reward channel and stays there, since the channels are equivalent.

The second set of figures, figure 2.5 portrays the effect of the standard deviation of the probability distribution or certainty of state. As the uncertainty increases, the probability of state moves towards the center of the road, with the maximum likelihood remaining on top of the center low cost region when the uncertainty is high enough. The state shifts toward the lower cost to appropriately avoid the higher cost. The final figure, 2.6 demonstrates a

higher cost of movement. A small increase in expected cost is accepted to save energy. The result is a probability distribution that only shifts when most necessary to avoid cost.

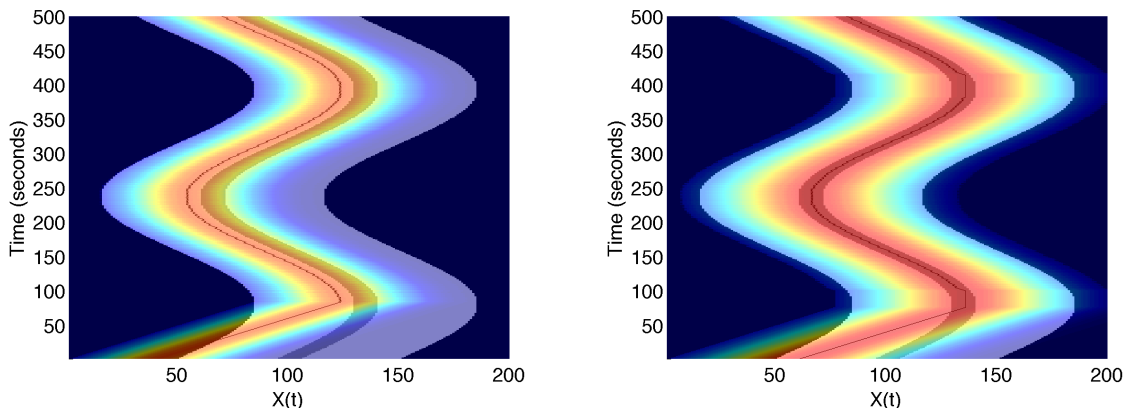


Figure 2.5: **Probability of State for Larger Standard Deviations.**

These figures are the same as in figure 2.4, but the standard deviations of the probability distribution of state,  $p(x)$ , are increased. When the Gaussian is widened to 20 pixels standard deviation, the maximum likelihood line hugs the center low cost region instead of being between both boundaries as in figures 2.4. When the Gaussian is further stretched to 40 pixels standard deviation, the maximum likelihood sits directly on the low cost centerline.

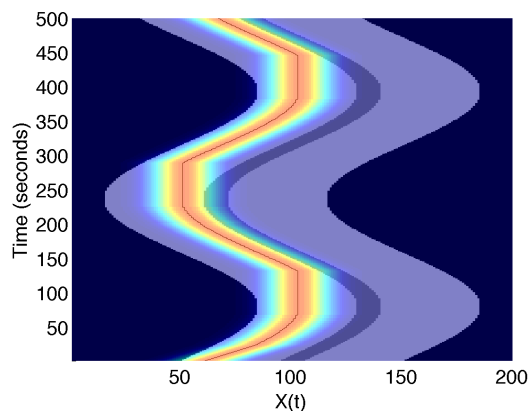


Figure 2.6: **Probability of State for Cost to Move.**

This figure is the same specifications as figures 2.4, with the cost to move or energy term increased. The maximum likelihood still remains between the high cost and low cost regions, however it only moves to the left of right when necessary to avoid contact with a penalty region.

## 2.3 Spiking Neuron Implementation

*A version of this section, along with section 1.3, was submitted to Journal of Neural Engineering.*

### 2.3.1 Methods

The first simulation demonstrates that risk aware control appropriately controls position according to a specified cost function. However, the question remains whether it is possible to implement Risk-Aware Control on a biologically-realistic distributed network of spiking neurons. In an existential proof, we were able to control a desktop robot (Sensable PHANTOM Omni Haptic Device) in real time using Risk-Aware Control to navigate a visually specified cost function. Real-time control was implemented utilizing the GPU (NVidia GTX 970) on a laptop computer (Alienware 17) programmed using a CUDA library (Accelerate) within Python (Anaconda2, Python 2.7). The cost function was derived from the built-in camera at 30Hz (OpenCV).

We simulated risk-aware control using a spiking neuron model of cortex. The model was comprised of layers of 640x480 neurons represented by Poisson-distributed binary spikes. The simulation demonstrated 2-dimensional movement, but the dimensions were controlled independently of each other.

The probability of state,  $p(x)$ , was encoded in a 640x480 layer of spatially tuned neurons. In the simulation, the probability of state was assumed to be Gaussian around the sensory estimate of the robotic arm. While this implementation only considered information from the robot position sensor, this estimate could be obtained from multisensory integration combined with an internal model estimate of state for a more biologically-realistic state estimate.

The cost of state,  $v$ , was encoded in two spatially analogous layers, a layer for reward regions and a layer for penalty regions (positive and negative costs). Cost of state was

derived from the camera feed at 30Hz. A neuron in the positive cost layer was on only if the associated pixel fell within the blue color range. Similarly, a neuron in the negative cost layer was turned on if the associated pixel fell within the red color range. This created a visually guided cost function with regions of both reward and potential harm.

Likewise, four layers were dedicated to kernel positive and negative representations of state for each spatial direction (vertical and horizontal). These layers represented neurons with input from multiple neighboring presynaptic cells that characterize a smoothed positive and negative probability of state. Neuronal layers were separated into positive and negative representations since spikes cannot inherently take on negative values, but negative changes in density must be accounted for.

Lastly, there was a layer for each action encoding the change in expected cost. The possible actions in this implementation were positive force and negative force in the each direction. These final layers represented the profit from each action for each individual neuron that were ultimately summed to determine the weighting factor,  $u_i$  in equation 1.2. This resulted in a total of 11 layers of approximately 300,000 neurons for a total of nearly 3.5 million neurons in the model.

As a result of this distributed representation, described in equation 1.5, the mathematical operations largely become local and/or operations. These types of operations lend themselves well to GPU computing. Utilizing the GPU for computational power, the implementation operated at approximately 30Hz, the upper limit for the frame rate using OpenCV to obtain the camera feed. Due to limitations of the robot drivers and computer hardware, the robot was run using a separate PC and communicated with the laptop running the Risk-Aware Control code via UDP connection. The control diagram is outlined in figure 2.7. Code for the Risk-Aware control is included for reference in the supplemental information.

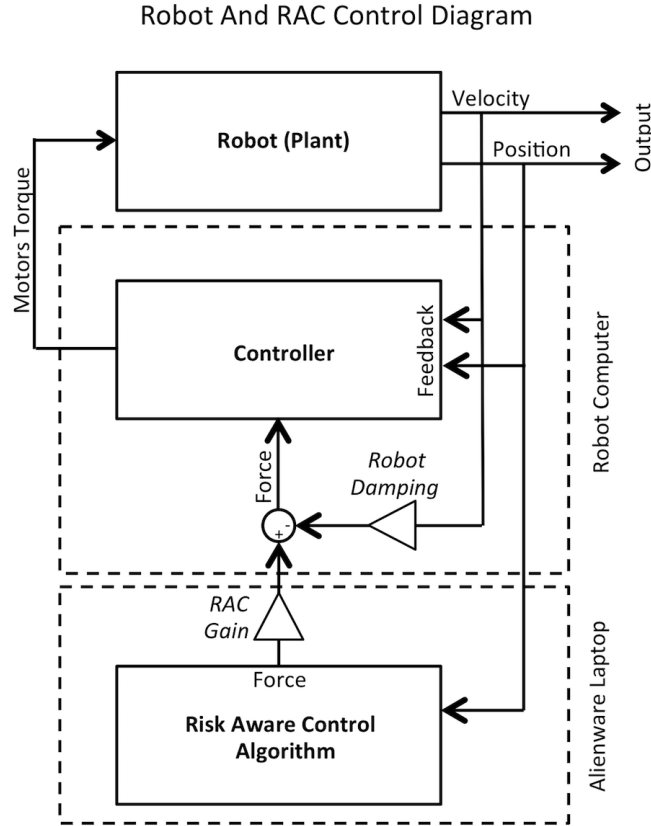


Figure 2.7: **Control Diagram of the Robot and Risk-Aware Control.**

The diagram illustrates the specific locations where the gain amplifier and damping amplifier enter the control loop. The dashed lines indicate the functions that are performed by each computer.

### 2.3.2 Experiments

The first test of the implementation was a simple sinusoidal tracking task. The robot was presented with a blue circle (reward region) moving in a horizontally sinusoidal motion with sweeping frequency. The lowest frequency, .03 Hz, was chosen to be a speed that the robot could easily and accurately track and the highest frequency, .6 Hz, was large enough so that the robot was unable to keep up with the speed of the cost function. The moving image was created in MATLAB and presented on an external display positioned in front of the camera, as seen in figure 2.8. The cost function display was positioned 27 inches away from the Alienware camera. The set up is depicted in figure 2.8.



Figure 2.8: **Pictures of the Robot Set-up.**

Figure A shows the view of the robot and computers and figure B displays the set up of the cost function monitor. The camera from the Alienware laptop observes the video cost function on the external monitor in real time. The cost function and spiking representation of the probability distribution can be seen on the laptop screen.

The stability of control will depend on the gain, damping, and control loop delay. The control loop was already operating at the maximum frame rate supported by OpenCV, therefore only gain and stiffness could be manipulated to optimize performance. If the gain is too high, the trajectory will demonstrate instability, specifically oscillations, due to overshoot of force. If the gain is too low, the robot will not have sufficient force to keep up with a rapidly moving cost function. The damping parameter will have a similar, but

opposite, outcome. The gain we describe is a gain within the risk aware control algorithm that can be thought of as an amplifier on the twitch strength, which we will refer to as RAC gain. Conversely, the damping parameter was a negative gain on velocity operating at 1000Hz added to the control loop on the robot. The location these parameters enter the control loop can be found in figure 2.7. The second experiment explored the effect of the damping and gain parameters on the control. The tracking paradigm of the first experiment was repeated for various damping and gain values and the average standard deviation of movement at rest and the average amplitude of movement across frequencies were compared between conditions. All combinations of three damping values, from 0.8 to 1.2, and ten gain values, from 1 to 10, were evaluated.

The final experiment investigated the implementation’s response to perturbation. In risk-aware control, the effect of a perturbation on the state is mathematically identical to a perturbation on the cost function. For ease of implementation, the perturbation was performed to the cost function. The shape of the cost function was an ellipse with the diameter of the major axis equal to twice the diameter of the minor axis. The orientation of the ellipse was varied so that the perturbation force occurred along either the major axis (horizontal orientation) or minor axis (vertical orientation). In addition to orientation of the cost function, the perturbation size and standard deviation of probability of state were also varied.

### 2.3.3 Results

The implementation was first assessed using a simple tracking task. The tracking results from a single experiment can be seen in figure 2.9A. The first segment is the results of the sinusoid tracking and the second segment tracks a stationary cost function. The gray shaded region indicates the reward region and the thick black line denotes the actual position of the robot arm. The robot follows the desired trajectory well within a range of frequencies. At very low desired frequencies (no motion), the robot produces small oscillations. At very

high desired frequencies, the robot cannot keep up. This tradeoff between stability and speed was explored in more detail in the second experiment.

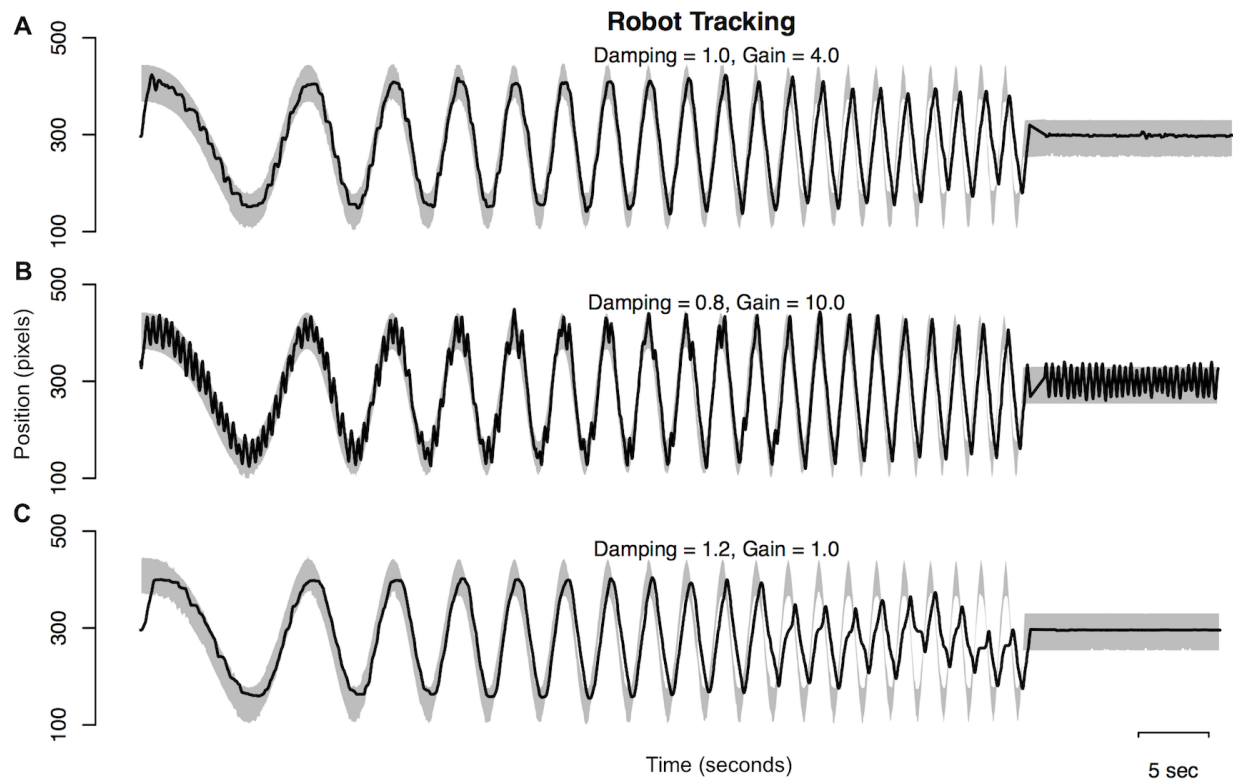


Figure 2.9: **Tracking Data of Trajectories.**

Figure A demonstrates the tracking data from the first experiment. The gray shaded region indicates the region of reward from the cost function (desired trajectory). The solid black line denotes the actual position of the robotic arm. Figures B and C display the tracking data from second experiment. Figure B shows the effect of decreased damping and increased gain. The robot keeps up with the cost function much better, however there are very prominent oscillations at the lower frequencies. Figure C shows the converse, increased damping and decreased gain coefficients. In this case, the robot arm is very stable, but has difficulty tracking the cost at high frequencies.

In the second part of the study, the first experimental setup was repeated under 30 different conditions to evaluate the best damping and gain parameter values. The raw tracking data from the combinations of low damping/high gain and high damping/low gain can be seen in figure 2.9B and 2.9C respectively for comparison. Two outcome measures were used for assessment: the stability of the robot, and the speed of the robot. Stability was measured by the standard deviation of the robot arm while tracking a stationary target di-

rectly following the sinusoidal tracking. The results can be seen in figure 2.10A. As expected, there is a general trend of increasing stability with decreasing RAC gain. Interestingly, there is a steep slope of change in stability at an RAC gain of around 5. The average standard deviation for all gains was 10.5 for a damping of 0.8, 8.0 for 1.0, and 6.0 for 1.2. Conversely, results demonstrated a general trend of increased speed with increased gain, shown in figure 2.10B. Speed was evaluated by calculating the average amplitude the robot reached across all frequencies. If the robot were able to keep up with the cost function, the average amplitude would be equal to the amplitude of the sinusoid. The average amplitude was 408.2 pixels at a damping of 0.8, 396.5 at 1.0, and 388.0 at 1.2. The relationship between the average amplitude and RAC gain demonstrated a more linear trend than the stability parameter.

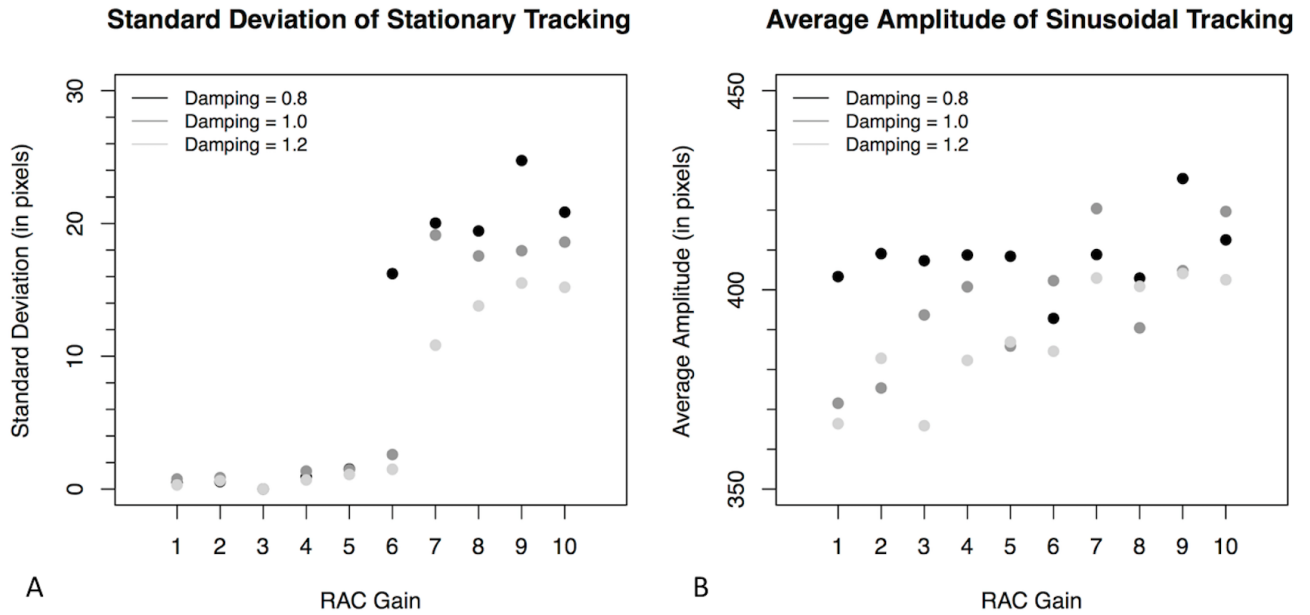


Figure 2.10: **Stability of Stationary Tracking and Speed of Sinusoidal Tracking.**

Figure A illustrates the effect of RAC gain and damping on the stability of tracking. Stability is measured as standard deviation of movement tracking a stationary target. Figure B depicts the effect of RAC gain and damping on the ability of the robot to keep up with a moving cost function. This was measured as the average amplitude the robot reached while tracking the sweeping frequency cost function. A larger amplitude corresponds to an increased tracking speed.

Different parameters may be desirable for different tasks. This is seen in human movement as well, where increased damping or stiffness is often observed for increased risk in order to decrease instability (Burdet et al., 2001; Perreault et al., 2002). It would be possible in the future to incorporate these state variables to be dependent upon risk in a similar manner.

The final test of the implementation was to evaluate the response to perturbations. In risk-aware control, a reflex response will only be initiated if the perturbation pushes the state towards risk or away from the goal of a task (Sanger, 2014). This property has been observed for the long latency stretch reflex in humans (Crago et al., 1976; Hammond, 1956; Ludvig et al., 2007; Rothwell, 1980). Moreover, the control theory dictates that the reflex response will be dependent on the risk. Therefore, perturbations that drive the state toward higher risk will have a proportionally higher reflex response. This phenomenon was observed in the results, shown in figure 2.11. In these trials, the perturbation either pushed the cost function along the major or minor axis of the ellipse, but the trajectory of the center of the cost function was identical between orientations. The dotted line indicates the perturbation displacement. At all distances, the vertical cost function elicits a larger response, in terms of both slope and change in position. In fact, overshoot is observed from the increased force of response for the vertical cost, but not for the horizontal cost. This is because the risk of a perturbation along the minor axis is greater than the risk of a perturbation along the major axis due to the width of the cost region. This behavior is exacerbated for smaller standard deviations of probability of state. However, when the standard deviation of the state probability density is near half the width of the cost, oscillations begin to occur. This could be mitigated by including a second order term to the state in the risk-aware control implementation.

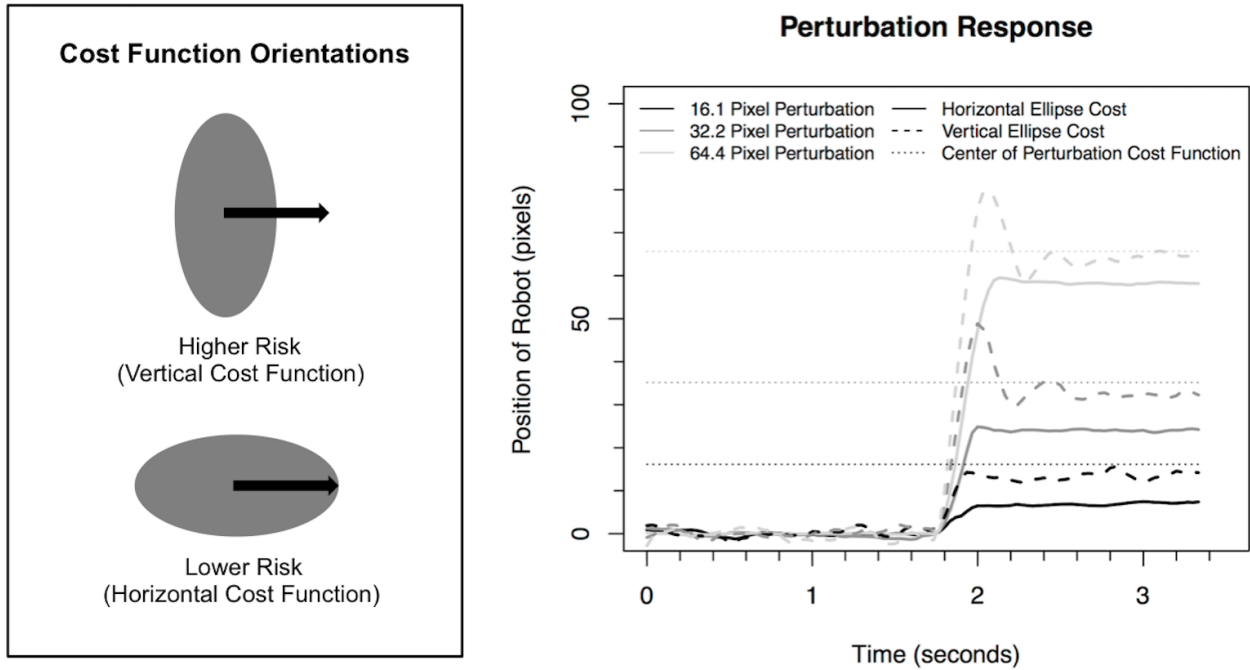


Figure 2.11: **Response to Perturbation.**

The left box illustrates the cost function implemented in the perturbation experiment. The plot depicts the response to different displacement perturbations (indicated by color) for the vertically (dashed) and horizontally (solid) oriented cost functions. The dotted lines indicate the actual size of perturbation. The standard deviation of the probability distribution was 80 pixels, the same as the width of the minor ellipse axis.

### 2.3.4 Discussion

Risk-aware control is a simple theory with simple implementation, but powerful results. Current leading theories of motor control cannot be fully implemented in real time, even with immense processing power. The results from this study demonstrate successful control of visually guided reaching movement using a Poisson spiking neuron model. We have demonstrated that even a basic implementation of risk-aware control exhibits the ability to appropriately navigate a risky environment in real time on an ordinary laptop computer with a graphics processing unit. We have explored the parameters to optimize performance for a task and characterized the response to perturbations as well as the limitations of stability and speed in this particular implementation.

Humans are able to accomplish very complex behaviors in novel environments. Existing computers are certainly able to match the computational power in both magnitude and time as the human brain. However, the field of robotics has still not been able to reliably simulate the compliance, complexity, and robustness of human movement. Therefore, it may be speculated that the limitation lies primarily in inferiority of the control algorithms, plasticity, and organization, and therefore that alternative theories of movement control should be explored. Risk-aware control is ultimately an instance of a standard feedback controller, but has the unique approach of representing state and control variables in the probability domain allowing for uncertainty in state and control. In the future, this implementation will be extended to include more complex, learned dynamics and a predictive internal model for increased stability. It is the goal that the outcome will be an adaptive, compliant controller that can be executed on modern computers in real time.

## Chapter 3

### Human Motor Response to Risk

#### 3.1 Experiment 1: The Tuning of Human Motor Response to Risk in a Dynamic Environment

*A version of this was published in PLoS ONE.*

##### 3.1.1 Introduction

Previous studies (Landy et al., 2012; Trommershauser et al., 2003a; Trommershauser et al., 2003b; Trommershauser et al., 2005) have investigated the effect of risk on motor planning. Trommershauser and colleagues (Trommershauser, 2003a) have demonstrated that humans are able to maximize expected gain by using internal representations of the magnitude of outcome uncertainty. When outcome uncertainty was artificially enhanced by randomly perturbing trajectory end-points, subjects still demonstrated the ability to maximize reward based on end-point variability by shifting their mean trajectory endpoints in response to changes in penalties and location of the penalty region relative to the target region. Furthermore, it has been shown that subjects respond to changes in uncertainty when it is artificially increased or decreased without cue during an experiment. However, these experiments all investigate behavior in discrete-tasks, such as rapid pointing to a target. While there is a general lack of consensus on the degree of online error correction during motor program execution involved in these rapid movements, their duration is certainly too short to take full advantage of feedback control loops (Galen and John, 1995; Keele and

Posner 1968) and therefore they rely primarily on motor planning (Maloney et al., 2007). In such experiments, we can investigate the effect of uncertainty on motor planning but not the effect on ongoing control of continuous movements.

While there is good evidence that humans plan movements taking risk into account, it is not clear how this occurs. For example, people might avoid actions that have previously led to poor outcomes as predicted by error-driven learning (Wei and Kording, 2008). We consider the hypothesis that humans actively and continuously estimate both the probability of failure and the cost of failure, and that they make ongoing corrections to movement based on these estimates. In general, the probability and cost of failure may vary throughout the workspace, so to do this requires maintaining estimates of these values for all states that could possibly result from movement errors. This ability is a foundation of risk-aware control, a theory of motor control in humans that links ideas in optimal control with existing literature on risk behavior in humans (Sanger, 2014). If humans have this ability, then it is also possible to estimate risk without experiencing failure. Therefore we hypothesize that humans will respond to perceived risk even in situations where failure has not been experienced. At the most extreme, this means that humans will select movements that reduce risk even when the probability of failure is negligible.

To test this hypothesis, we designed a driving simulation experiment with a cost function similar to that of figure 3.1 and to the cost function used in the simulations implemented in chapter 1. Each lane became a reward region and driving off the road or between lanes resulted in a point penalty. If humans maintain estimates of both probability of failure and cost of failure, then where in a lane the subject drives should depend on the specific form of the cost function. We further predict that these changes in behavior do not require subjects to experience failure (driving off the road).

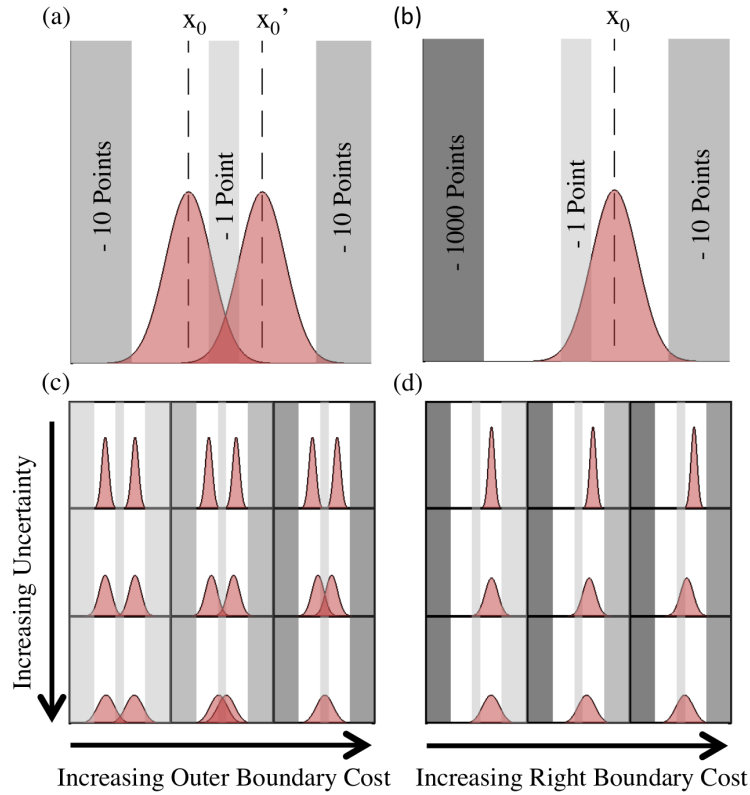


Figure 3.1: **Theoretical minimization of cost under uncertainty.**

In figures a-d, the shaded red distributions represent an uncertainty or variability in position. Grey bars signify penalty regions, the darker the grey, the higher the cost. The peaks of the curves illustrate the optimal position to minimize cost based on the standard deviation of uncertainty and the cost function. In (a) the loss function is symmetrical. The result is that there are two optimal positions that will minimize cost. Figure (c) demonstrates the effect of increasing the cost of the outer boundary, dark grey regions, from left to right (1, 10, 100). The result is a shift in peaks toward the lower cost region in the center. Similarly, as the standard deviation of uncertainty increases from top to bottom (.35, .75, 1) the optimal position again shifts toward the center lower cost region. At high standard deviation of uncertainty and high outer boundary cost, the optimal position becomes directly in the center of the middle region. Figures (b) and (d) illustrate the same phenomenon for an asymmetrical loss function. Here the left boundary penalty remains very high (1000 points) while the right boundary in (d) increases from left to right (1, 10, 100). In this case there are no longer two optimal positions, only one in the segment that is farther away from the high cost.

### 3.1.2 Materials and Methods

#### *Subjects*

Twelve nave subjects, ages 22 to 35, 9 males and 5 females, participated in the experiment. The University of Southern California Institutional Review Board approved the study protocol. All subjects gave informed written consent for participation and received compensation in proportion to their final score plus a base sum (Study IRB# UP\_09\_00263). Authorization for analysis, storage, and publication of protected health information was obtained according to the Health Information Portability and Accountability Act (HIPAA).

#### *Apparatus*

The experiment was performed on an iPad2 (iOS 6.0, resolution of 1024x740 pixels) in landscape orientation. A custom application was created using CoronaSDK (Version 2012.11.15. Palo Alto, California: Corona Labs Inc., 2012). The update rate of the screen and rate of data acquisition was 30 fps.

#### *Stimuli and Procedure*

The experiment took approximately one hour to complete, with small breaks as necessary, and was completed in a dimly lit room to avoid screen glare. For biomechanical uniformity, subjects were instructed to sit in a chair, maintaining their shoulders against the backrest, and to keep their elbows at approximately 90 degrees with their biceps inline with their torso during the entire experiment. Subjects grasped the sides of the screen with both hands at all times. Instructions for the experiment were verbally specified by the experiment administrator and presented again on the iPad screen for the subjects to read once they entered the application.

In the experiment, subjects maintained one-dimensional “steering” control of a vehicle in a driving simulation. The goal of the game was to complete each trial as quickly as

possible, where the speed of the car was determined solely by position on a two-lane road. While on the road, driving within a lane yielded acceleration to the maximum velocity (1100 pixels/sec), driving on the dashed line between the two lanes caused the vehicle to decelerate to 550 pixels/sec, and hitting the grass along the side of the road slowed the car to 2 pixels/sec (which will be referred to as “stopped” as the car could hardly be detected as moving). Figure 3.2 contains a screenshot of the application. Subjects were able to control the position of the car by tilting the iPad in the left/right directions. Points awarded were inversely proportional to the time taken to complete each trial. Subjects could earn a maximum of 100 points per trial if they maintained the maximum velocity along the entire length of the road and could not earn less than 0 points due to speed penalties. Implementing the cost function in this manner effectively reinforced the cost, since more successful trials were linked not only to increased points and therefore increased monetary reward, but also decreased experiment time.

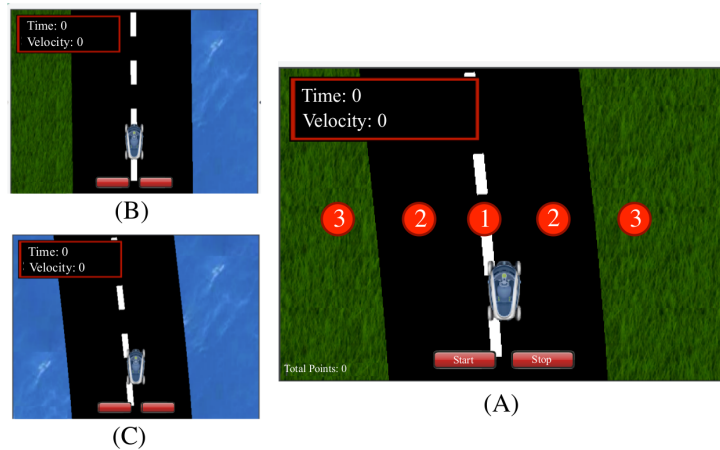


Figure 3.2: iPad Application Screen View.

The subjects pressed the red start button to begin each trial (and were asked to not press the stop button during any trial). The time and velocity of the car was provided in the upper left hand corner of the screen. The three regions of speed are labeled in the figure with circles. Region 1 produced acceleration to maximum speed of 1100 pixels/sec; region 2 decelerated the car to 550 pixels/sec; region 3 immediately stopped the car to 2 pixels/sec. Cost functions: (A) symmetric low-cost, (B) asymmetric, (C) symmetric high-cost.

In addition to inherent motor variability, uncertainty was artificially enhanced by cor-

rupting the responses of the subject with random, Gaussian-distributed horizontal perturbations at a frequency of 30Hz (the same frequency as the screen updates). Within the context of this study we will define this imposed variability as motor noise. The effect was similar to the sensation of driving on a bumpy road; the subject was able to determine the present car position, but was uncertain exactly where they may be in the next instant. Thus the effect of the motor noise was to increase uncertainty of future position and alter the probability of failure. It is important to note, however, that this is not identical to driving on a bumpy road, where noise is dependent on position on the road. The noise was generated at a constant time interval so that slowing down would not make the task significantly easier. There were five levels of imposed motor noise: 0 (no additional noise), 4, 8, 12, and 16 pixels standard deviation (psd). Each trial was 30,000 pixels in length and took approximately 30 to 60 seconds to complete. The first 10,000 pixels of each trial were practice, giving the subject enough time to get up to speed and adjust to the noise level, during which no points were accumulated or lost. The car always started a trial where it ended in the previous, unless it was the first trial of a block, then the car started in the middle of the road. The road was 500 pixels wide, the center dashed line was 15 pixels wide, and the car width was 40 pixels. The curves of the road were generated using Bezier curves (Farin, 1997) with random anchor points derived from a uniform distribution.

During the experiment, subjects' responses were tested to three cost functions, blocked into two sets of trials: block A) the symmetric low-cost and block B) the asymmetric and symmetric high-cost. In block A, subjects completed a random sequence of the 5 uncertainty levels 3 times, for a total of 15 trials using the cost function as described above (grass on both sides). During block B, water replaced the grass on one or both sides of the road respectively. Running into the water caused an immediate stop and replaced the car to the center of the road (the timer was stopped so that this was equivalent with respect to time to running into the grass), but with an additional 500-point penalty. In environments with water there were only 4 degrees of additive noise (0, 4, 8, and 12 psd) as during pilot testing the highest noise

level caused subjects to generally earn very negative points and discouraged subjects from heeding the point system. Therefore, in order to maintain a high sensitivity to risk, we did not include a noise level of 15 psd in environments with water. Each noise level was repeated twice with water on both sides and twice with water on one side (counter-balanced) for a total of 16 trials in a pseudo-random sequence.

Each subject first learned control of the car in the low-cost environment during 15 practice trials (same as block A). Each subject was informed that driving on the black part of the road would yield maximum velocity, while touching the white center lane would cause the car to slow down and hitting the side of the road would bring the car to a stop. Subjects were also told that they could earn a maximum of 100 points per trial and were encouraged to explore the road during the practice block during which points earned or lost would not count towards their monetary reward. After finishing the practice block and brief rest, subjects then completed block A and block B in random order with rest in between.

### *Data Analysis*

During the experiment, we recorded position of the car and the time it took to complete each trial. Points were recorded, but not used in analysis as they were rounded, and therefore less accurate, and only piecewise proportional (subjects could not earn less than 0 points from speed penalties). Trial time did not reflect the effects of falling into the water, but only one subject incurred this penalty.

All analysis was done in MATLAB (version 7.13.0.564. Natick, Massachusetts: The MathWorks Inc., 2011) and R: A Language and Environment for Statistical Computing (version 3.0.1. Vienna, Austria: R Development Core Team, 2013). Position data for subjects was pooled to represent average behavior of the sample population and fit to equation 1 using maximum likelihood estimation. In these functions, zero is center of the road and units are in pixels.

$$y = (p)f_1(x|\mu_1, \sigma_1) + (1 - p)f_2(x|\mu_2, \sigma_2) \quad (3.1)$$

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}}e^{\frac{-(x-\mu)}{2\sigma^2}} \quad (3.2)$$

In equation 1,  $x$  is position,  $\mu_1$  and  $\mu_2$  are the means of each Gaussian,  $\sigma_1$  and  $\sigma_2$  are the standard deviation of the respective Gaussians,  $p$  is the weighting factor between the Gaussians, and  $y$  is the resulting probability of position. The resulting parameters,  $\mu_1$ ,  $\mu_2$ ,  $\sigma_1$ ,  $\sigma_2$ , and  $p$ , were interpolated (using cubic spline interpolation) across motor noise to generate estimated continuous position probability distributions for each cost function.

Additionally, the position data for each subject, for each cost function at each noise level were fit to equation 1 using maximum likelihood estimation. In order to quantify the average distance from the center of the road that subjects attempted to maintain for each condition, the absolute value of  $\mu_1$  and  $\mu_2$  were weighted by the area under each Gaussian,  $p$  and  $p-1$ , and summed. Linear regressions were fit to the means across subjects of the four lowest levels of motor noise of each task condition.

A two-way ANOVA was performed using the `aov` function of the R statistical computing environment. The R model `aov(Position ~ Uncertainty * Task)` was used to test the differences in distance from the center of the road between the level of motor noise and task type. In this model, uncertainty is the quantifiable level of simulated motor noise imposed, and task was the type of environmental cost function. Post hoc pairwise comparisons were made between each of the three cost functions within each uncertainty level using paired-t tests with an alpha value of 0.05.

A two-way repeated measures ANOVA was also performed to test the differences in trial time between the level of motor noise and task type. The test was performed in R using the model `aov(Time ~ Uncertainty * Task + Error(Subject))`. Again, paired-t tests were used to make post hoc pairwise comparisons of the difference between each of the three cost functions within each uncertainty level.

We were also interested in the role that errors played in forming behavior. Therefore, the percentage of failed trials, trials in which the subject went outside the road, was calculated for each level of motor noise of the low-cost task and asymmetric task. (It is not presented for the symmetric high-cost task, because only one such failure occurred amongst all subjects in this environment.)

### 3.1.3 Results

In order to quantify any learning effect within the course of the experiment, distance from the center of the road for each condition from subjects who completed block A first were compared with those of subjects who completed block B first. They were not significantly different ( $p < 0.05$ ), therefore it was concluded that after the initial practice trials, there was no observable learning effect. As expected, position data resemble bimodal Gaussian distributions as shown in figure 3.3. As motor noise increases, the two peaks of the distribution tend toward each other, merging into a single normal distribution at high motor uncertainty. Essentially, subjects reacted accordingly to motor noise; they stayed within a lane at low levels of uncertainty and moved toward the center of the road at high levels of uncertainty, illustrated in figure 3.4. This reflects a tradeoff in which they accept the higher cost of driving on the median in order to avoid the risk of driving off the road. In the asymmetric risk environment, position data appropriately reflects the asymmetric cost with highly disproportionate peaks, so that subjects have a strong tendency to drive on the side of the road that is farthest from the water. However, as the noise increases, subjects drive closer to the middle of the road and thus closer to the water in order to balance the risk of falling to either side. Table 1 contains the percentage of time that subjects spent in the lane near the grass (away from the water) during this task.

Standard Deviation of Motor Noise (in pixels)	0	4	8	12
Percentage of Time Spent in Lane Away From Water	96.12%	96.23%	91.53%	88.87%

Table 3.1: **Percentage of Time Spent in Lane Farthest From Water in the Asymmetric Cost Environment.**

Subjects spent significantly more time in the lane opposite of the water. The increase in percentage of time with increased motor noise can be attributed to subjects moving closer to the center of the road to avoid hitting the grass. As they moved toward the center of the road, they crossed the centerline into the lane adjacent to the water more often, albeit still briefly.

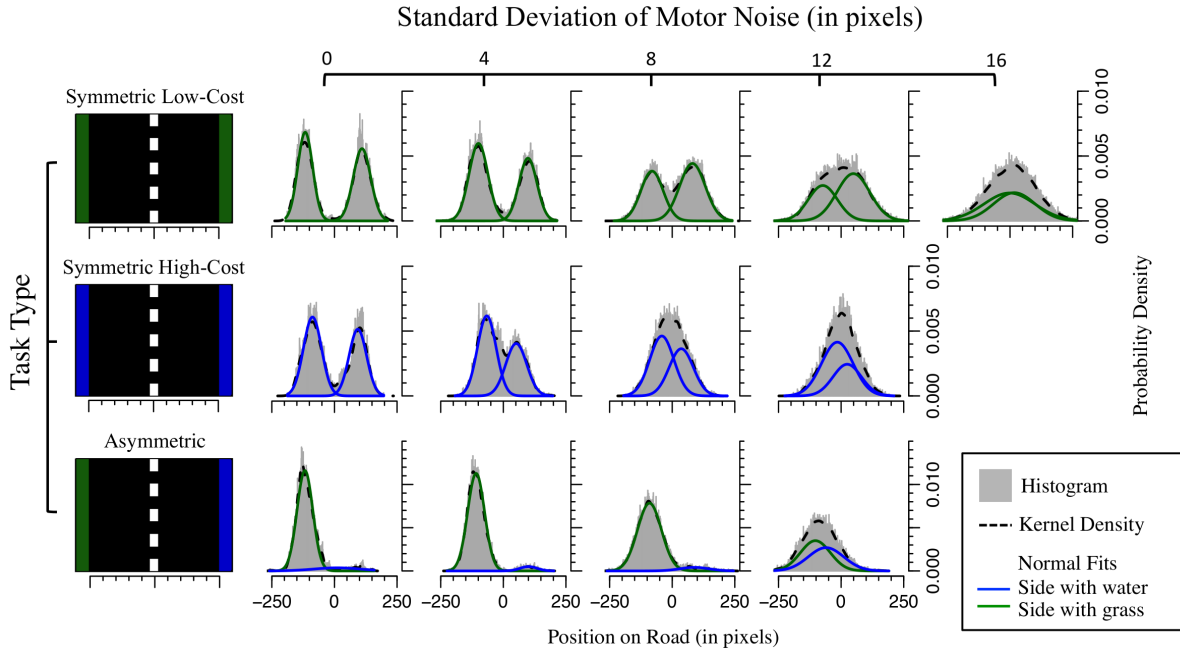


Figure 3.3: **Raw Population position data.**

Plots are the histograms of the pooled subject data for each task type (by row) and uncertainty level (by column). The dashed lines are the kernel densities of the data and the solid lines are the bimodal Gaussian fits (see methods). Green lines represent the position Gaussian near grass (low-cost) while blue represents a position peak near water (high cost). In the asymmetric task, the bottom row, it can be seen that subjects maintained a position far away from the side with water. The x-axis represents the position of the center of the car on the road in pixels. (The road is 500 pixels wide, and the car is 40 pixels, so the subject ran off the road at  $\pm 230$  pixels.) These images depict a trend similar to figure 3.1. As the outer boundary costs increased, the subjects moved toward the center of the road. Similarly, as the standard deviation of uncertainty increased, subjects also moved toward the center of the road.

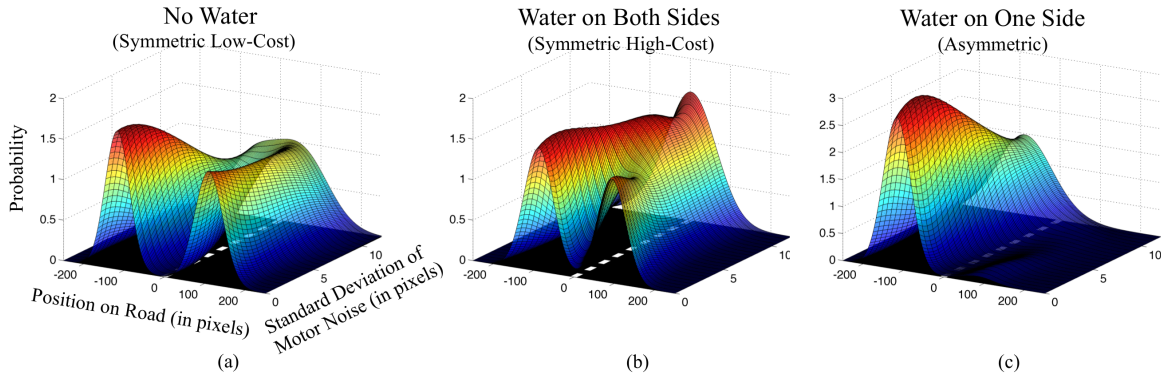


Figure 3.4: **Continuous Position Probability Distribution as a Function of Uncertainty.**

Variables of bimodal Gaussian fits ( $\mu_1$ ,  $\mu_2$ ,  $\sigma_1$ ,  $\sigma_2$ , and  $p$ ) from figure 3.3 were interpolated (using cubic spline interpolation) across noise levels. This demonstrates an estimate of the probability of where on the road a subject will be at any given instant as a function of motor noise.

Over the four lowest levels of uncertainty, subjects took an average of 26.69 seconds to complete a trial in the symmetric low-cost environment, 28.51 seconds in the asymmetric environment, and 30.30 seconds in the symmetric high-cost environment. A two-way repeated measures ANOVA showed that there was a significant effect of motor noise [ $F(4,408) = 102.469$ ,  $p < 0.001$ ], task type [ $F(2,408) = 43.892$ ,  $p < 0.001$ ], and interaction [ $F(6,408) = 6.417$ ,  $p < 0.001$ ] on the time it took to complete a trial. This is not especially informative since the implemented cost function directly affects trial completion time; increased motor noise will lead to larger trial times regardless of how the subject responds. The more interesting conclusions lie in the post hoc pairwise comparisons between task types at equal levels of uncertainty. Of the twelve comparisons, all pairs except one were significantly different ( $p < 0.05$ ), the symmetric low-cost task and the asymmetric task at 4 psd motor noise. In other words when risk was introduced into the environment, subjects sacrificed time and points to steer clear of the high cost regions, shown in figure 3.2. Comparing the symmetric tasks, there is a shift in the y-intercept of the regressions, but the slopes are almost identical. This indicates that there is a constant effect of the increased cost on subjects' responses independent of motor noise (at least within this range).

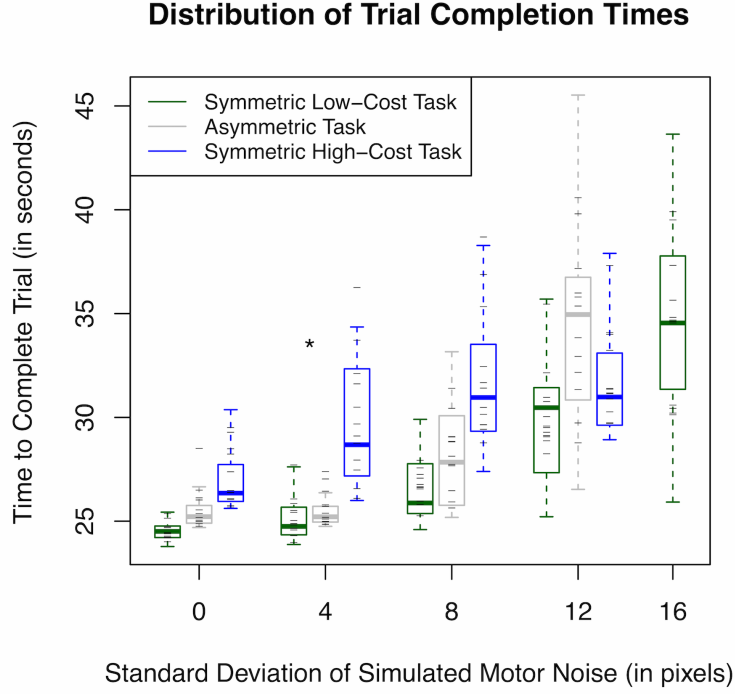


Figure 3.5: **Regression on the Distance from the Center of the Road Subjects Maintained vs. Level of Motor Noise.**

Points indicate the mean distance from the center of the road of all subjects derived from the peaks of the fitted probability density functions (see methods) for each task type and uncertainty level. As motor noise increased, subjects' position shifted proportionally toward the center of the road. Position is normalized to 250 pixels so 0 is the center of the road and 1 is the edge of the road. Errors bars indicate the standard error of subjects. Solid lines represent the linear regressions fit for each task type. Asterisks indicate the pairs of values with insignificant differences.

The mean distance from the center of the road, calculated from the parameters fit to eq. 1 as explained in methods, for each condition can be seen in figure 3.6. Over the four lowest levels of uncertainty, the mean distance from the center of the road (normalized to the road width) was 0.3408 (SE  $\pm 0.0157$ ) for the symmetrical low-cost task, 0.3962 (SE  $\pm 0.0108$ ) for the asymmetrical task, and 0.2110 (SE  $\pm 0.0178$ ) for the symmetrical high-cost task. Linear regressions demonstrate a linear dependency of distance from the center of the road on uncertainty level [ $R^2$  (symmetrical low-cost task) = 0.982,  $R^2$  (symmetrical high-cost task) = 0.984,  $R^2$  (asymmetrical task) = 0.945]. A two-way ANOVA showed that there was a significant effect of motor noise [ $F(4,408) = 92.64$ ,  $p < 0.001$ ], task type [ $F(2,408)$

$= 112.61$ ,  $p < 0.001$ ], and interaction [ $F(6,408) = 4.157$ ,  $p < 0.001$ ] on the distance of the bimodal Gaussian position distribution peaks from the center of the road. In post hoc pairwise comparisons between task types within each uncertainty level, all were significant ( $p < 0.05$ ) except between the symmetric low-cost task and asymmetric task at 0 and 4 psd noise. Details can be found in figure 3.6.

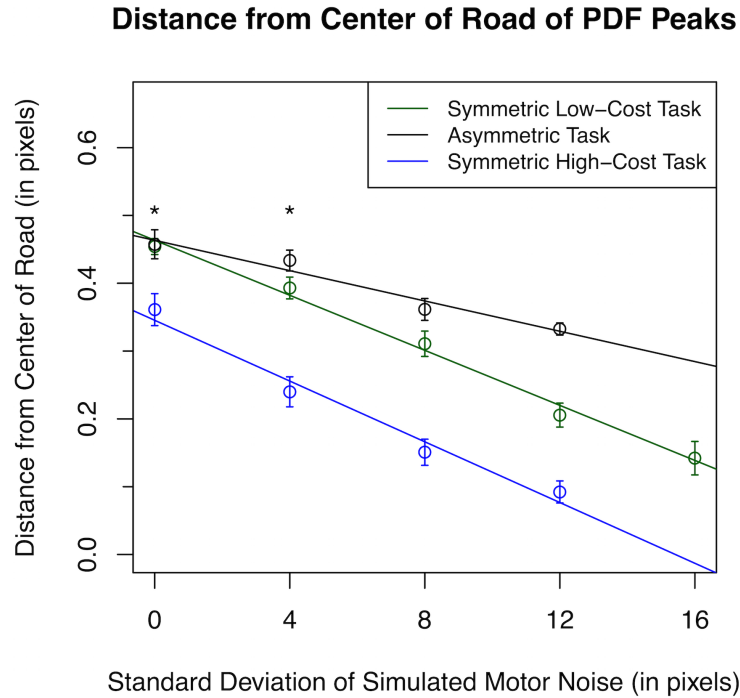


Figure 3.6: **Proportional Hazard Model of Successful Trials.**

Points indicate total percentage of successful trials for all subjects at each level of uncertainty, where a successful trial is defined as a trial during which the subject never ran off the road. (The green line represents the symmetric low-cost task, the black line is the asymmetric task, and the blue line indicates the symmetric high-cost task.) At all uncertainty levels, failed trials occurred more than twice as often in the asymmetric task than in the low-cost task. That is that subjects stayed so far away from the water that they hit the grass on the opposite side of the road much more frequently.

Subjects, on average, ran off the road in more than twice as many trials at every level of uncertainty in the asymmetric task than in the low-cost task, shown in Figure 3.6. These numbers do not include failures during practice. However, at the two lowest levels of motor noise in the low-cost task, even during the initial practice block, no subject ran

off the road. Additionally, only one subject ever fell off the road in the high-cost task. This demonstrates that subjects react to the probability of failure even when failure has not been experienced. This observation is inconsistent with an adaptive reduction in error, and instead must represent a mechanism that estimates and predicts failure that has not yet occurred.

### 3.1.4 Discussion

It has been previously suggested that humans act as Bayes optimal observers in motor planning tasks, such as rapid pointing, by modifying behavior to compensate for uncertainty (Faisal and Wolpert, 2009; Kording and Wolpert, 2004; Kording and Wolpert, 2006; Knill and Pouget, 2004; Maloney and Zhang, 2010; Tassinari et al., 2006; Wolpert and Landy 2012) In this study we were interested in investigating if this behavior extended to response to cost and uncertainty in a continuous task controlled with feedback, and if this behavior could be done without experiencing error in the task. At no additional motor noise, subjects on average maintained a bimodal Gaussian distribution near the center of either lane (approximately 10 pixels closer to the center line than the road boundary) in the low-cost environment. This shows that qualitatively optimal behavior could be performed with a bimodal cost function that is more complicated than the single target used in most prior studies. In the asymmetric environment, subjects stayed in the lane closer to the grass more than 95% of the time, and subsequently treated this lane almost identically to the low-cost task. In the symmetric high-cost environment, subjects moved more than an additional 20 pixels towards the center of the road compared to the symmetrical low-cost environment at equal uncertainty levels. Subjects shifted their behavior in the presence of risk even though no subject left the boundaries of the road in the symmetrical low-cost task at 0 and 4 psd motor noise. Based on the observation of error, there is in fact no reason to pull away from the side of the road when the cost to running off the road was increased. This behavior is either suboptimal or optimal with respect to an internally derived cost function that does not

match the empirical data. This suggests that predictions of failure not only carry very long tails, but predict possible error even when none has previously occurred. Additionally, only one subject ever fell into the water, but every subject still demonstrated a significant shift in behavior in the symmetric high-cost environment. This is significant since the common model for learning in motor control is error-driven learning, and this observation suggests that human performance is often not driven by errors. This demonstrates that subjects made predictions of both the likelihood and cost of failure, and our results are consistent with the existence of internal estimates of probability of failure and cost of failure.

As the uncertainty increased, subjects adjusted their position more towards the center of the road. The distance subjects moved away from the road boundary was dependent on the motor noise at least within the constraints of this task. While subjects behaved similarly in the asymmetric task to the symmetrical low cost-task at low levels of uncertainty, subjects adjusted their behavior differently at high levels of uncertainty. Subjects did not move towards the center of the road as much in the asymmetric task in order to avoid running over the centerline and into the high-cost region (water) on the opposite side. This occurred even though it caused subjects to hit the low-cost region (grass) much more often and meant taking significantly longer to complete the asymmetric task at the highest noise level than either of the other two tasks, see figure 3.5. However, since only one subject ever hit the water in any symmetrical high-cost environment, this again demonstrates that for most subjects this shift in behavior was not necessary and shows how sensitive humans are to high-risk regions.

It is important to recognize that the concept of risk we describe in this paper, “risk-awareness”, is derived from risk-aware control and a fundamentally different concept than the more ubiquitous “risk-sensitivity” originating from an economical decision-making perspective of motor control (Braun et al, 2011; Sanger, 2014). Risk-sensitivity is used to describe inter-individual differences in response to risk, where risk is defined in terms of higher moments of reward. In this experiment, an explicit cost function is provided so there should not

be much inter-individual difference. In the context of this paper we are defining awareness of risk as continuous estimates of both the cost of failure and probability of failure in a task. Unlike previous studies, we did not compare subjects' responses to the "optimal" response. It has already been demonstrated that in navigating 2-dimensional terrains humans' behavior is typically suboptimal (Zhang et al, 2010). It is certainly feasible to create a cost function sufficiently obscure or complicated to prevent humans from responding optimally. And there are many other considerations such as attention, fatigue, motivation, etc. that are impossible to quantify and implement in the estimation model, but that certainly affect the complete cost function a subject would theoretically minimize. Additionally, the results suggest that subjects are tuning their behavior to a probability function that is a result of both pre-existing assumptions about variability and measurements of the empiric variability of the task. Because we do not know the assumptions a subject makes of the underlying probability distribution, whether subjects are maximizing expected utility correctly and the appropriateness of the assumptions are not completely discernible with this study. It can be concluded, however, that subjects are responding to the increase of risk in the task. So it was not the focus of this study to determine how closely humans are able to reproduce the optimal response in continuous tasks, but whether they demonstrate on-going awareness of risk.

Humans are relatively fragile creatures. Only through constant vigilance and avoidance of risk do we remain safe from injury. We have shown that not only do we consider risk when initially planning a movement, but also that we are constantly evaluating the environmental cost function. Moreover, we are constantly making predictions of failure, even in cases where we have never experienced that failure. Our survival depends on knowing that falling off the cliff is going to be unpleasant without having to experience it first.

## **3.2 Experiment 2: Certainty Equivalence Assumption**

### **3.2.1 Introduction**

The previous study demonstrated that subjects continuously estimate their own uncertainty and appropriately tune their behavior to account for this uncertainty. The previous experiment imposed motor uncertainty, so that the imposed noise specifically affected the outcome of actions. However, uncertainty can generally exist in two forms: uncertainty in the state variables and uncertainty in the control variables. The former can be interpreted as sensory uncertainty (indeterminate knowledge of state) and the latter as motor uncertainty (indeterminate knowledge of the outcome of actions).

In the theory, state uncertainty and control uncertainty will have identical outcomes if the statistics of the uncertainties are equivalent. However, it is possible that the perception of these different types of uncertainty may affect the accuracy of their internal representation. The final segment of this report will revisit the role of uncertainty in behavior, and more specifically investigate the influence of both of these types of uncertainty on movement.

### **3.2.2 Methods**

In order to test the validity of the certainty equivalence assumption and compare responses between uncertainty types, we propose a third experiment expanding on the design of the first experiment. The same iPad driving simulation will be used as in the first chapter with a few minor modifications. Instead of physically slowing the car on the road, the cost function will be directly implemented by a point penalty (otherwise the subject may be able to infer the position of the car on the road from the velocity of the car). In addition to imposed motor uncertainty, identical to that from the previous experiment, subjects will complete the task with imposed sensory uncertainty as well, shown in 3.7(c-d).

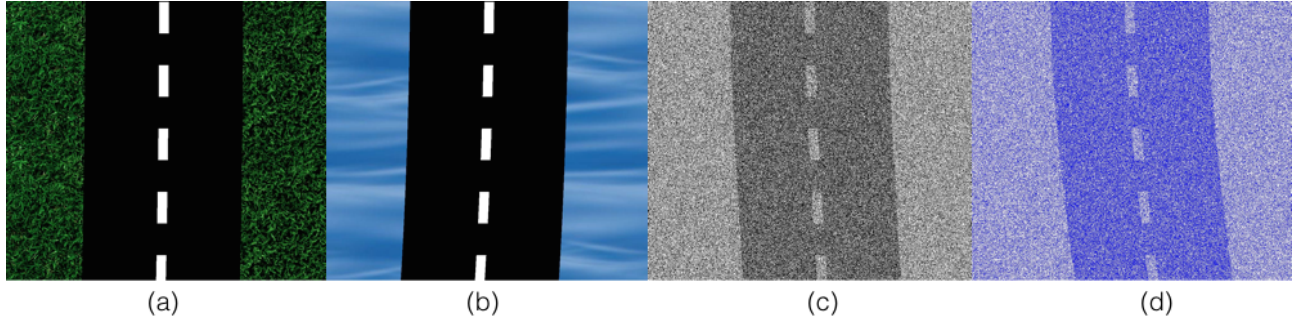


Figure 3.7: **Application Images.**

The above images depict the road for each of the conditions. Figure (a) shows the low-risk motor condition, (b) shows the high-risk motor condition, (c) shows the low-risk visual condition, and (d) shows the high-risk visual condition. The visual contrast of these images is 0.25.

In order to implement sensory uncertainty, the contrast between the road and boundary (both outside the road and the center dashed line) was varied and then the image was converted to 2-bit. In order to achieve this, the grayscale value of each pixel was proportional to the probability of that being black (versus white). Each pixel will be compared to an independent random variable and will be assigned a value (white or black) based on this comparison. For a single image, this is an implementation of random dithering that generates a one bit-per-pixel image from the original analog image. It may be important to recognize that this method does not increase the noise of an image so much as decrease the bandwidth or information of the image. We characterized and validated this method of imposing sensory uncertainty in section 2.1. The experiment was divided into two days, one day for calibrating the levels of uncertainty for each subject, and one day to compare the types of uncertainty.

### *Session 1: Calibration*

The first session was a calibration phase designed to match levels of motor uncertainty to statistically equivalent levels of sensory uncertainty. In this session, the application was similar to the previous experiment, but the cost function was altered. The subject was

rewarded points for driving on the dashed centerline, and there were no penalty regions. Subjects were not given any instantaneous feedback of performance, but were given a score from 0 to 100 at the end of each trial. The car velocity will not be altered in this experiment because this would indicate whether the car was on the road.

The contrast levels were chosen such that at the highest uncertainty, the road was essentially invisible to test subjects, and at the lowest level of uncertainty, the road was completely and easily visible. The levels were spaced logarithmically, so that there was finer differentiation at lower levels of contrast where visibility changes more rapidly. The motor uncertainty levels were evenly spaced from 0 pixel standard deviation to 14.5 pixels standard deviation (based on the results of the first experiment).

Subjects performed 3 blocks of 10 uncertainty levels under imposed motor uncertainty and the same number of trials under imposed sensory uncertainty. In one block the uncertainty went from high to low, in another the uncertainty went from low to high, and one block was randomized. The subjects performed these blocks in a random order. Prior to starting any trials, the subject completed one practice block identical to the increasing motor uncertainty block and one block of the increasing sensory uncertainty.

The position data for each trial was fit to a Gaussian distribution. The standard deviations of the Gaussian distributions for all trials from the motor uncertainty data was fit to a linear regression and the standard deviations resulting from the sensory uncertainty was fit to a decaying exponential. Four equidistant standard deviations of movement were selected for each subject and the corresponding contrast/noise levels for each standard deviation were calculated from the exponential/linear fits. Figure 3.2 illustrates the position data fit to Gaussian distributions and the standard deviations of the distributions fit to the regressions of one exemplary subject.

The four levels of sensory uncertainty and four levels of motor uncertainty from this calibration session were used in the second session. Any subject that had an  $R^2$  value for either fit that was less than 0.8 was not included in the second session.

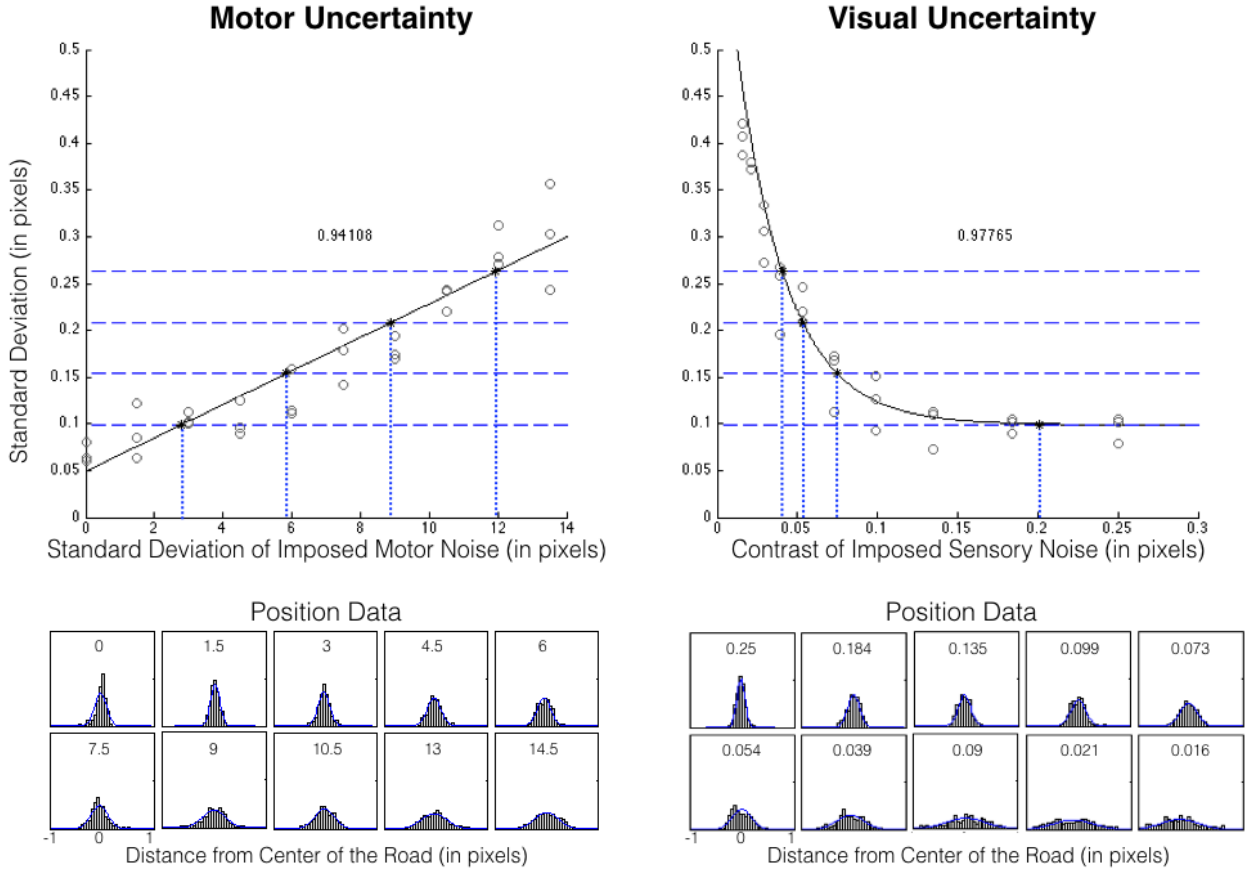


Figure 3.8: **Results of Exemplary Subject For Calibration Session.**

These figures are the results from the calibration session of one exemplary subject. The plots on the left are the results of the motor uncertainty and the plots of the right are from the visual uncertainty trials. Each figure on the bottom shows the histograms of the position data (distance from the center of the road) normalized by the road width. The numbers within each subplot indicate the level of noise of those trials. The blue lines represent the Gaussian fit for each histogram. The top plots show the standard deviation of the Gaussian fits and the levels of calibration. The circles represent the standard deviation from each trial. The black line indicates the linear/exponential fit to the standard deviations. The stars indicate the selected levels of standard deviation. The dotted and dashed lines show the conversion between standard deviation of uncertainty and level of motor/contrast noise. The numbers within in each plot indicate the  $R^2$  value of that regression.

### *Session 2: Complicated Cost Function*

In the second session, subjects completed the same game with the cost function from first study (a bimodal cost function). The levels of uncertainty were those established in

the calibration from the first session. The blocks were divided into high-risk and low-risk. In the low-risk task, subjects lost 5 points/second when the car was touching the centerline and 50 points/second at the boundary. In the high-risk task, subjects also lost an additional 300 points if they hit the outer boundary. Each block consisted of 12 trials, 3 for each level of uncertainty in a randomized order. At the beginning of the experiment, each subject completed a practice block identical to the low-risk block for both types of uncertainty. Additionally, the first third of each trial was practice for the subject to adjust to the uncertainty.

### *Subjects*

Fifteen subjects participated in the first session of the experiment. Of these, nine subjects did not meet the calibration criterion to participate in the second session.

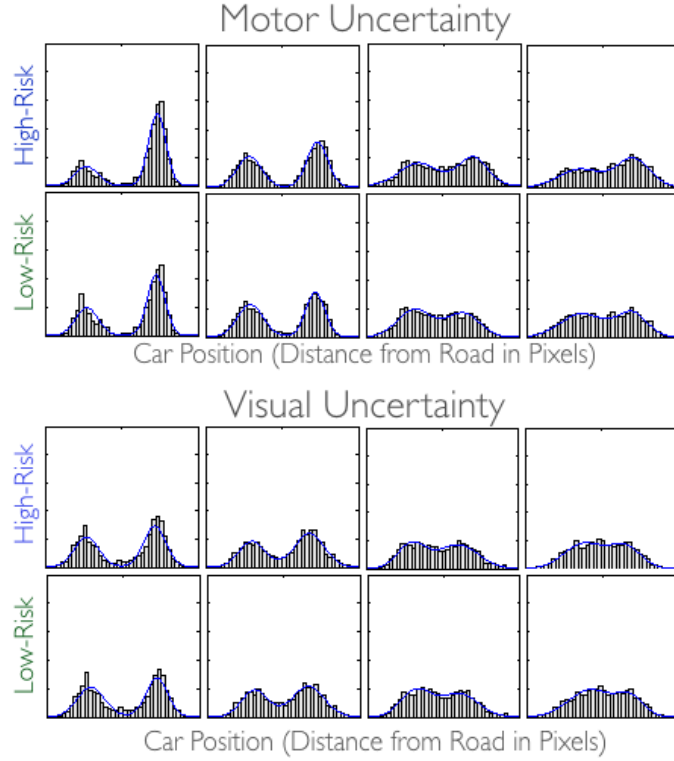
### *Data Analysis*

Data was analyzed in a similar manner as the previous experiment. The position data from the second session was fit to a bimodal Gaussian distribution, eq. 3.1. The distance from the center of the road was calculated as the absolute value of the peak of each distribution multiplied by the weighting factor of that distribution added together. This was computed for each condition for each subject individually. An ANOVA was performed comparing the effect of uncertainty type, uncertainty level, and risk on the distance from the center of the road.

## **3.2.3 Results**

All Gaussian fits from the calibration session data passed a chi-squared test. Results from the first session for a sample subject can be seen in figure 3.8. All the linear and exponential fits to the standard deviation for each subject were 0.8 or greater. The average  $R^2$  value resulting from the motor data regression was .8338 and the average of all subjects

for the exponential fit to the visual uncertainty data was .8854. The average motor noise levels were 1.63, 4.97, 8.31, and 11.66 pixels standard deviation and the contrast values were 0.2541, 0.0941, 0.0635, and 0.0457.



**Figure 3.9: Histograms of Raw Position Data and Fits to Bimodal Distribution.** The figure above display the histograms of the position data from all subjects combined. The position data is referenced to the center of the road and normalized by the size of the road. Therefore the left boundary of the plot signifies the left boundary of the right and the right boundary of the plots signify the right boundary of the road. The solid lines indicate the bimodal Gaussian fits to equation 3.1. The top set of figures are the results of the motor uncertainty and the bottom set are the results of the visual uncertainty. Each column is the results from each level of uncertainty, lowest uncertainty on the left and higher towards the right.

The bimodal Gaussian fits for each condition from all subject data can be seen in 3.9. Qualitatively, the distributions look similar. The average results indicating the distance from the center of the road for each condition can be seen in figure 3.10. The average distance for the visual uncertainty (across all levels and risks) is 0.3281 normalized to the road and the average distance for motor uncertainty is 0.40. The average distance (across all

uncertainty types and levels) for the symmetric low-risk condition is .3835 and the high-risk is .3346. There was a significant difference in uncertainty type [ $F(1,85) = 22.62$ ,  $p < 0.001$ ], uncertainty level [ $F(3,85) = 31.81$ ,  $p < 0.001$ ], and risk [ $F(1,85) = 15.13$ ,  $p < 0.001$ ].

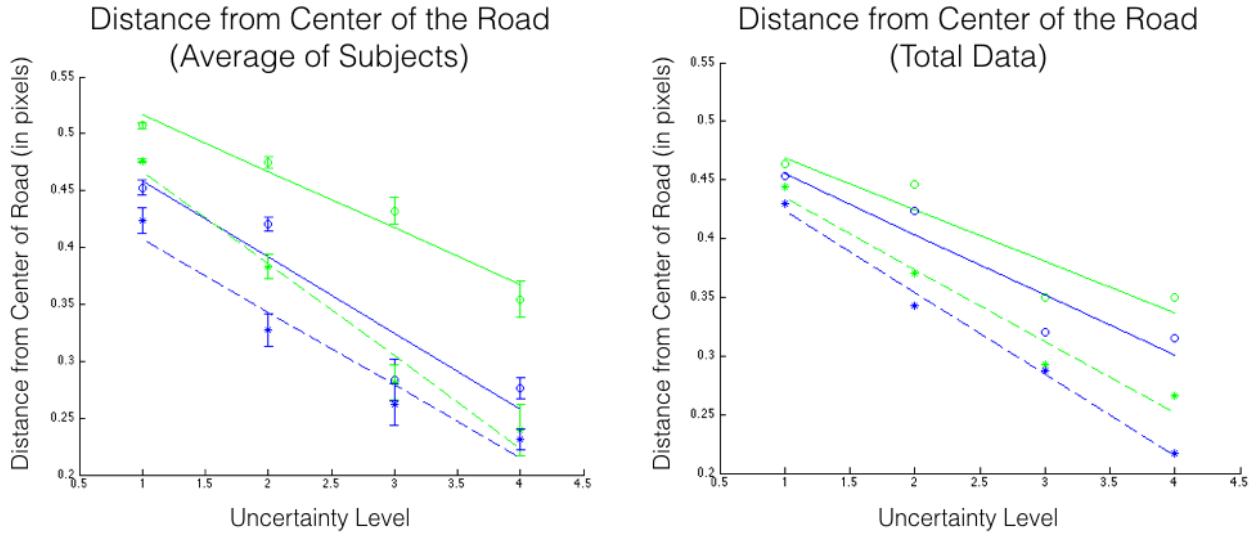


Figure 3.10: **Distance from the Center of the Road vs. Uncertainty Level.**

The plots above show the distance from the center of the road across uncertainty levels for the high risk and low risk conditions and the visual and motor uncertainty. The figure on the left are the results of the average distance from the center of the road calculating from the bimodal Gaussian fit for each individual subject. The figure on the right are the result of the bimodal gaussian fit from the total subject data, figure 3.9. The green line represents the low-risk condition and the blue line represents the high-risk. The stars indicate the visual uncertainty and the open circle indicates the motor uncertainty. The lines are the fit to a linear regression. The error bars indicate variance.

### 3.2.4 Discussion

Results demonstrated that there was a significant increase in sensitivity to sensory uncertainty over motor uncertainty. It cannot be determined if this is due to a difference in the manner in which each type of uncertainty is incorporated or if this is a result of an inaccurate perception of uncertainty.

Interestingly, the results from the motor uncertainty in this experiment appear differ-

ent than in the previous experiment, although the levels of motor noise and risk are similar in this study. The sensitivity to the risk, both in terms of response to uncertainty and cost, appears diminished. We hypothesize that this is a consequence of the cost not being reinforced by the duration of the experiment. It is possible that subjects consider a lengthier experiment a significant increase in cost as opposed to just monetary penalty. It is also possible that the velocity of the car provided instantaneous feedback to reinforce the sensory estimate of the car position. In this paradigm the subjects did not know when they hit the middle or the side of the road until the end of each trial.

### *Certainty Equivalence*

It is interesting to note that these results also have major implications about assumptions of certainty equivalence. In order to alleviate some amount of the computation burden, control theories often assume they are controlling a quadratic linear system. This means that the criterion function is quadratic and that the systems equations are linear. This simplification relies on many substantial assumptions including local linearity, additive Gaussian noise, and certainty equivalence (Sanger, 2014).

If a system is certainty equivalent, then the deterministic problem is equivalent to the stochastic problem (Kendrick, 2002). Therefore, a system in which the optimal solution taking into account uncertainty is the same as the optimal solution for that system in the absence of uncertainty is certainty equivalent (Thiel, 1957; Simon, 1956). This is a convenient assumption as it means the maximum likelihood can be substituted for the probability distribution. However, the results of this study demonstrate that this is not an appropriate assumption. In this study, the degree of uncertainty impacts behavior and therefore the deterministic system is not equivalent to the stochastic problem.

## Chapter 4

### The Tuning of Reflexes to Risk

#### 4.1 Introduction

In the previous chapter we describe how risk can be interpreted in terms of the probability of cost and the probability of failure. In general, we do not have control over the form of the cost function. The cost of failure is determined by the environment. However, we do have some influence over probability of failure. There are two ways to decrease the probability of failure. The first is that we may adjust our position to move away from the risk. From the previous study, we have demonstrated that humans do tune their statistical behavior based on the entire probability distribution of possible outcomes and the cost function of environment and task. The second is that we can prepare for the unexpected or plan for error. We may do this by tuning our rapid responses to the probability of failure and cost of failure. Therefore, in addition to modifying the control of movement to reflect the risk of the environment, we hypothesize that humans will also prepare for error in response to risk as well.

It has been well established that humans are able to modulate the long latency stretch reflex based on the goal of a task (Ludvig et al, 2007; Pruszynski et al, 2008; Hammond, 1956; Rothwell et al, 1980; Crago et al, 1976). Studies have implemented a variety of paradigms including verbal instructions (resist/let go) and target oriented tasks, single muscles and concurrent muscles. The commonality of all these experiments is that they look very specifically at goal-modulation. It is our hypothesis that awareness to risk is so fundamental, that

humans also maintain reflexes tuned specifically to the cost function of the environment, even when the goal of the task does not depend on the perturbation.

In order to appropriately respond to perturbations, the reflex response should take into account the risk of the environment as well. If we consider the previous example of driving, it would be very harmful if the reflex response to the mechanical perturbation from a bump in the road resulted in a reaction that pushed the car over a cliff. Alternatively, it would be very helpful if the same reflex response resulted in avoiding hitting another car and running into another empty lane instead. This chapter consists of three separate, but closely related, experiments. The first experiment implemented a paradigm similar to the experiment presented in the previous chapter to investigate the rapid response to visual perturbations. The second experiment further investigates this phenomenon by examining the stretch reflex response in the first dorsal interosseus in environments with different amounts of risk. The final experiment will repeat the second experiment with minor changes using the bicep and triceps muscle in order to investigate the role of co-contraction and tone in modulating the reflex response to risk.

This is an important set of experiments to distinguish characteristics indicative of risk aware control from optimal feedback control. In optimal feedback control, the reference trajectory is simply recalculated at each time point depending on the state. This accounts for redirecting movement not to the previous path, but to the best path to achieve a goal. This is important because setting reflexes is unnecessary to follow a reference trajectory. However, when controlling the system through dynamics, tuning reflexes is an inherent result of the system. This will present itself as planning for error before the error occurs.

## 4.2 Experiment 1: Response to Visual Perturbations

### 4.2.1 Materials and Methods

#### *Apparatus*

The experiment was performed on an iPad2 (iOS 6.0, resolution of 1024x740 pixels) in landscape orientation. A custom application was created using CoronaSDK (Version 2012.11.15. Palo Alto, California: Corona Labs Inc., 2012). The update rate of the screen and rate of data acquisition was 30 fps.

#### *Subjects*

Eight naive, healthy adult subjects participated in this second study. The University of Southern California Institutional Review Board approved the study protocol. All subjects gave informed written consent for participation and received compensation in proportion to their final score plus a base sum (Study IRB# UP\_09\_00263). Authorization for analysis, storage, and publication of protected health information was obtained according to the Health Information Portability and Accountability Act (HIPAA).

#### *Stimuli and Procedure*

This experiment utilized a very similar experimental set up as in the first study. The application itself was the same with a few modifications. The length of trial was increased from approximately 25-35 seconds to 50-60 seconds. Additionally, only 3 levels of motor noise were evaluated (0, 4, and 8 pixels standard deviation). At random time intervals throughout each trial, the car was displaced to a new position by visually moving the car from its current position to a new position in the next frame. There were three types of displacements: push to edge, push to center, and lane switch, as illustrated in figure 4.1.

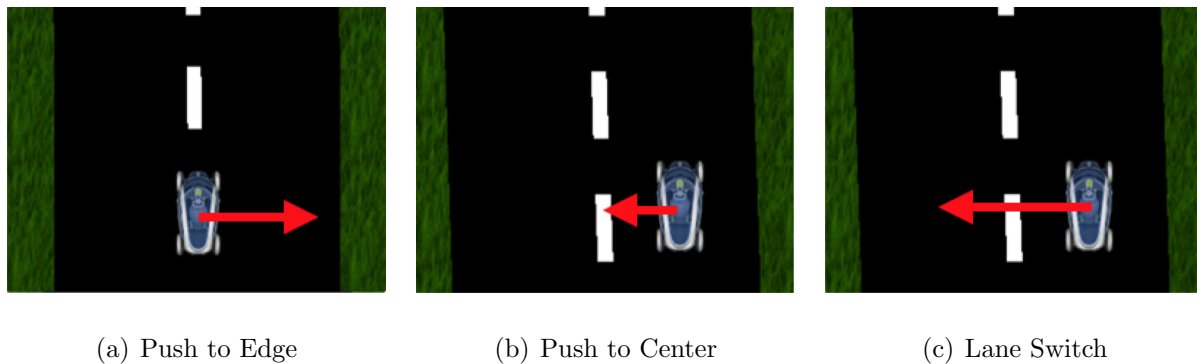


Figure 4.1: **Types of Visual Displacements.**

Three types of visual perturbations occurred during the experiment. Figure (a) illustrates the push to edge displacement. The car is moved from its current position to 50 pixels from the edge of the road in the current driving lane. Figure (b) depicts the push to center displacement where the car is moved to the very center of the road. Figure (c) shows a lane switch perturbation. Here the car is moved to the exact same position in the opposite lane.

The experiment was divided into 3 identical blocks with small breaks in between. Each block was subdivided into 3 sub-blocks for each cost function (grass-grass, water-grass/grass-water, or water-water). The order of sub-blocks was randomized. Within each sub-block the subject completed 3 trials, one of each noise level in a random order. During each trial the car was displaced 9 times, 3 per displacement type, with each displacement occurring at a random time in the trial but at least 5 seconds apart. This resulted in 9 displacements per noise level per cost function for a total of 81 visual perturbations throughout the experiment. Prior to the first block, the subject completed a practice block under the grass-grass cost function.

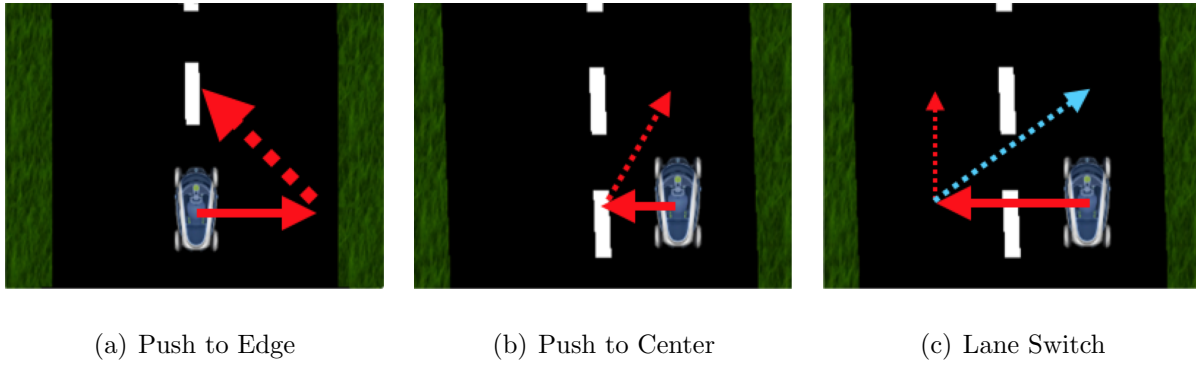


Figure 4.2: **Specific Hypotheses.**

If reflexes are tuned to the risk of the environment we expect to see differences in responses depending on the risk. In figure (a) we would expect to see a larger response to pull the car away from the high risk at the edge of the road. In figure (b) we anticipate that subjects' response will be less or slower to move away from the smaller risk of the center of the road. In figure (c) the response should depend on whether the cost function is symmetric or not. In a symmetric environment the risk is the same in the displaced position so we should see a much smaller response to the visual perturbation. In an asymmetric environment, we expect the subject to return to the previous lane.

### *Data Analysis*

Data from the accelerometer was primarily used for analysis since it is the most direct link to the subjects' physical reactions to the visual perturbations. Accelerometer responses were aligned to the direction of perturbation. Then the amplitude of response to each perturbation was calculated by subtracting the baseline accelerometer reading (averaged over .5 seconds prior to perturbation onset) from the maximum response. The maximum was defined as the first maximum accelerometer response (did not decrease for more than 3 consecutive frames) between perturbation onset and 100 frames post perturbation. The time of the amplitude was also computed as the frame this maximum response occurred.

The percentage of times that the subject returned to the previous lane after the lane switch perturbation was also analyzed. A trial was considered a return trial if at 100 frames post perturbation the car position was at least 50 pixels into the opposite lane. For symmetry, the trials in which the car position was in the same lane and at least 50 pixels from the centerline 100 frames after the perturbation was considered a same lane trial. All other trials

were discounted as ambiguous. A trial was also automatically discounted if the subject fell into the water within those 100 frames. The percentage was calculated as the number of return trials divided by the sum of the return and same lane trials.

#### 4.2.2 Results

Average accelerometer responses for each subject and the average of all subjects for the push to edge and push to center perturbations can be seen in figure 4.3. The previous study already established the generalized response to this task without the perturbations; therefore this investigation was only interested in specifically analyzing the rapid response to perturbations. In a three-way repeated-measures ANOVA, perturbation type, noise level, and cost function all had a significant effect on the response amplitude ( $p < 0.0001$ ).

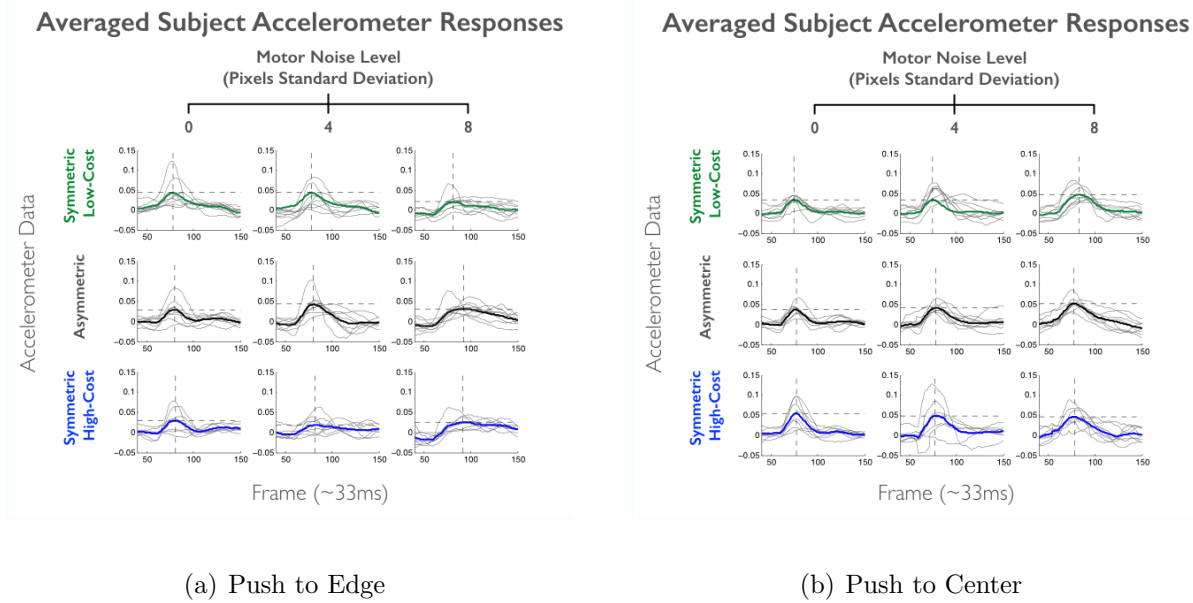
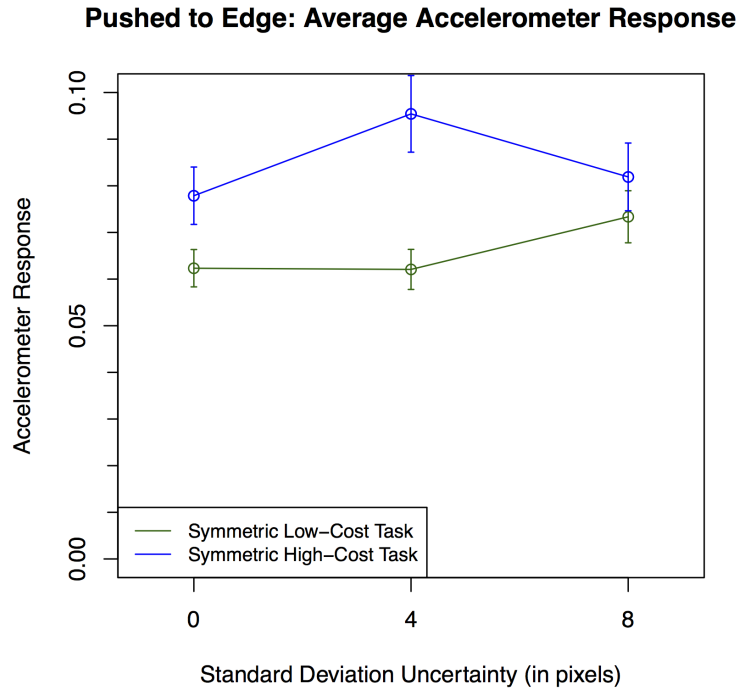


Figure 4.3: **Rapid Responses of Individual Subjects and Average of All Subjects.** In each grid the columns signify the level of motor noise (0, 4, 8 psd from left to right) and the rows characterize the cost function (symmetric low-cost, asymmetric, symmetric high-cost from top to bottom). Each grid is labeled with the type of visual perturbation (a: push to edge, b: lane switch, c: push to center). The x-axis is in seconds and the y-axis is accelerometer data. The grey lines represent individual subject responses (averaged over the 9 trials). The thick colored lines indicate the average of all subjects' responses. The perturbation occurred at frame 50.

Post hoc comparisons were performed with t-tests to more specifically interpret the effect of risk on response amplitude. In the push to edge displacements, subjects were either pushed towards the grass or towards the water depending on the environment. The response amplitude was significantly greater ( $p = 0.0028$ ) when the subject was being pushed toward the water or higher risk than when pushed toward the grass or lower risk, shown in figure 4.4. Considering now all the environments, the subjects also demonstrated significantly larger ( $p < 0.0001$ ) response amplitudes when pushed toward the edge of the road (higher risk) than towards the center of the road (lower risk), shown in figure 4.5.



**Figure 4.4: Amplitude of Average Accelerometer Response to Push to Edge Perturbation.**

Points represent the average accelerometer response to the push to edge perturbation of all subjects. Green indicates the symmetric low-cost task and blue indicates the symmetric high-cost task. The x-axis is the standard deviation of motor uncertainty in pixels and the y-axis is the accelerometer response. The bars represent standard error. The stars indicate significantly different pairs ( \* =  $< 0.05$ , \*\* =  $< 0.005$ , \*\*\* =  $< 0.0005$ ).

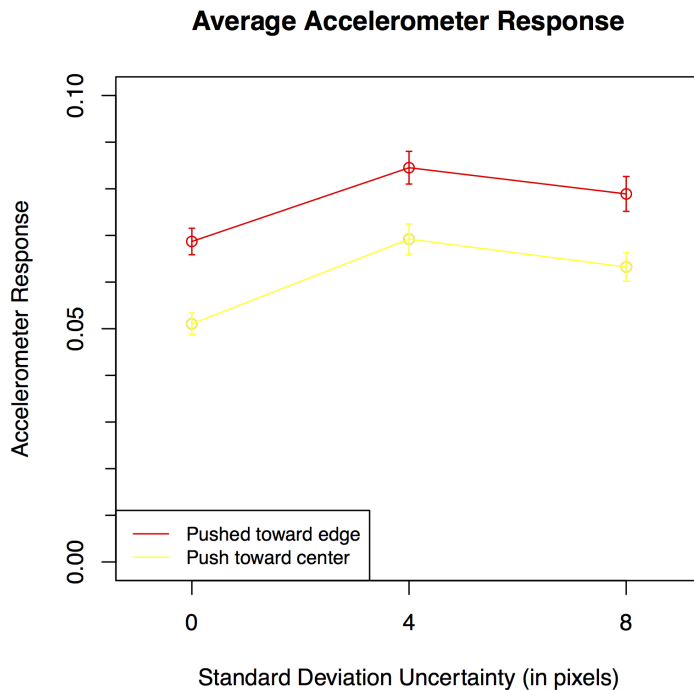


Figure 4.5: **Amplitude of Average Accelerometer Response of the Push to Edge and Push to Center Perturbations.**

Points represent the average accelerometer response of all subjects in all cost environments combined. Red indicates the response to the push to edge perturbation and yellow indicates the push toward center perturbation. The x-axis is the standard deviation of motor uncertainty in pixels and the y-axis is the accelerometer response. The bars represent standard error. The stars indicate significantly different pairs ( \* =  $< 0.05$ , \*\* =  $< 0.005$ , \*\*\* =  $< 0.0005$ ).

In the switch-lane perturbation, the position data is more informative than the accelerometer data, figure 4.6. Optimal feedback control would predict that, in the symmetric environment, subjects always remain in the new lane post switch lane perturbation because they recalculate the optimal trajectory. When subjects were pushed to the opposite driving lane, they generally stayed in that lane. It should be noted that this was not always the case though, sometimes subjects returned to the previous lane even in the symmetric environment. In the asymmetric environment, subjects usually returned to the lane away from the higher risk. Percentages of the returns can be seen for each condition in figure 4.6.

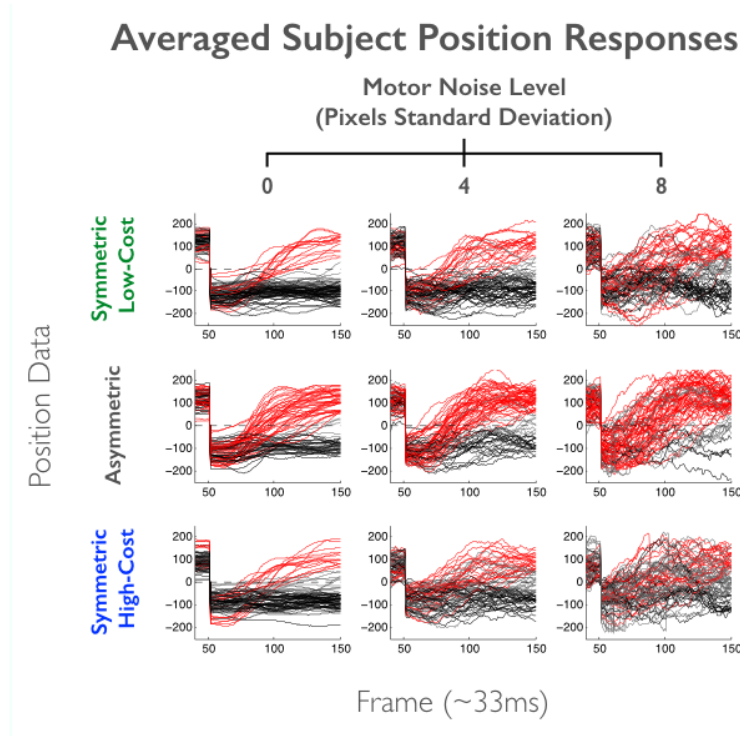


Figure 4.6: **Position of Car Post Lane Switch Perturbation.**

In each grid the columns signify the level of motor noise (0, 4, 8 psd from left to right) and the rows characterize the cost function (symmetric low-cost, asymmetric, symmetric high-cost from top to bottom). The x-axis is in frames and the y-axis is center of the car position data. The red lines indicate trials that the subject returned to the lane the car was perturbed from. The black lines indicate trials that the subject did not return to the previous lane. The grey lines designate trials that were deemed ambiguous (see methods). The numbers above each subplot are the percentage of trials that the subject returned to the previous lane.

### 4.2.3 Discussion

Results demonstrate that subjects' rapid response is modulated to the location of risk in the environment. Amplitude of response was higher for perturbations that pushed the subject toward higher risk, in this case toward the edge of the road instead of the center or toward water instead of toward grass.

This supports our hypothesis that reflex responses are modulated by risk, however, because the perturbations implemented are visual and not mechanical, these reaction times exist only in the voluntary region (Thorpe et al., 1996) and are not true reflex responses.

The next experiments will utilize mechanical perturbations to test the modulation of the stretch reflex response to environmental risk.

### 4.3 Experiment 2: Response to Mechanical Perturbation

#### 4.3.1 Materials and Methods

##### *Subjects*

Ten nave, healthy adult subjects participated in this study. The University of Southern California Institutional Review Board approved the study protocol. All subjects gave informed written consent for participation and received compensation in proportion to their final score plus a base sum (Study IRB# UP\_10\_00447). Authorization for analysis, storage, and publication of protected health information was obtained according to the Health Information Portability and Accountability Act (HIPAA).

##### *Stimuli and Procedure*

Subjects were positioned in front of a monitor with their entire forearm resting on a table and their right index finger braced in a splint. The split was attached to the arm of a robot magnetically, as a safety precaution. The physical arm of the robot was constrained to move in only one dimension, horizontal to the subject, by a plastic board with a rectangular section removed. The three distal fingers were taped together and all fingers, besides the index finger, were attached to the board with Velcro. Electrodes were affixed to belly of the first dorsal interosseus and abductor digiti minimi muscles and a ground electrode was placed on the opposite hand. Figure 4.7c shows this set-up.

The position of the robot arm corresponded to the position of the cursor on the screen. The monitor displayed three rectangles: two cost regions on either side of a center reward region, which moved horizontally (remaining equidistant) in a randomized sinusoidal motion. Subjects were verbally instructed to maximize points by keeping the cursor within the center reward region while avoiding the cost regions that would result in a loss of points. The center reward rectangle was twice as wide as the cursor dot so there was not much leeway in the position of the cursor that would earn points. This ensured that the cursor remained

approximately equal distance from either cost region. The robot generated a constant 1 N baseline force with randomized 4 N perturbations in both directions (randomized) at a mean rate of 3 seconds.

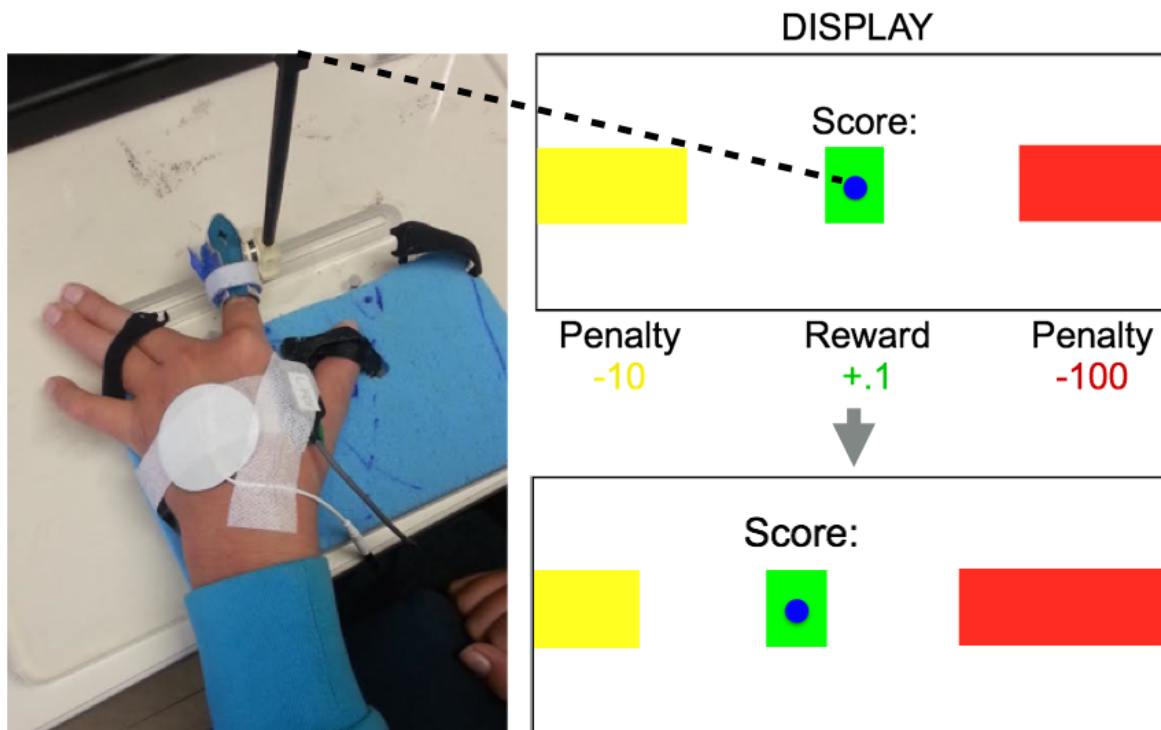


Figure 4.7: **Monitor Display and Set Up.**

The above images illustrate the display the subjects viewed during the experiment. There was always the center green reward region while the penalty regions changed throughout the experiment. The rectangles all moved together in a randomized sinusoidal motion. The subject's finger was attached to the robot that controlled the blue cursor on the screen. The subject's hand was strapped in using Velcro to avoid using the other fingers or adjusting the hand position mid-experiment. Electrodes were placed on the FDI and ADM and a ground electrode was placed on the back of the hand.

The two cost rectangles were colored to indicate penalty. Nine cost environments were evaluated: all combinations of no penalty, low penalty, and high penalty. No penalty meant the subject would not lose points when the cursor was inside the (white) cost rectangle, hitting a yellow rectangle resulted in a loss of 10 points (low penalty) while hitting a red rectangle resulted in a loss of 100 points (high penalty). Under this paradigm, responses to

both symmetric and asymmetric cost functions could be investigated.

The experiment was divided into 5 blocks of 90 perturbations. The cost function (cued by color) was changed every 10 perturbations in randomized order. A mandatory 5-minute rest period separated each block of trials. Prior to the first block, subjects were instructed to abduct their index finger against the hand of the experimenter as hard as possible in order to collect maximum voluntary control.

Ultimately, the goal of the task was always to remain in the center target, but the cost of hitting a penalty region was varied. If a subject always returned to the center target as quickly as possible, they would always maximize their points. However, we expected to still see a difference in the long latency stretch reflex between cost functions as a result of subjects tuning their reflexes to the risk of environment.

Prior to the beginning of the experiment, subjects were asked to push with as much force as possible against a stationary object in order to record maximum voluntary control.

### *Data Analysis*

All analysis was done in MATLAB (version 7.13.0.564. Natick, Massachusetts: The MathWorks Inc., 2011) and R: A Language and Environment for Statistical Computing (version 3.0.1. Vienna, Austria: R Development Core Team, 2013). Electromyography was recorded from the first dorsal interosseus and abductor digiti minimi muscles at 1000 Hz. The abductor digiti minimi was not used for analysis, but observed for a general sense of the stiffness of the hand. The EMG data was low pass filtered [500 Hz], rectified, then bandpass filtered [25 – 250 Hz, butterworth]. Data was normalized based on maximum voluntary control (MVC).

Data was time-aligned to perturbation onset. Only trials in the direction that activated the FDI stretch reflex were analyzed. This resulted in approximately 25 trials per subject per cost function. Reflex response was divided into standard epochs for baseline [-50 – 0 ms], short latency [R1, 20 – 45 ms], medium latency [R2, 45 – 75 ms], long latency [R3, 75

– 105 ms], and voluntary [VOL, 110 – 150 ms] response post perturbation. (From Rapid Motor Responses Are Appropriately tuned) Analysis compared mean EMG within each epoch for each perturbation trial.

Subject data was analyzed both individually and as a group. A two-way repeated measures ANOVA using cost factor and epoch factor was performed to determine the effect of risk and epoch on the stretch reflex for each individual subject. The R model `aov(EMG Risk*Epoch)` was used to determine the effect of the risk of the environment on the EMG activity.

Five one-way repeated measures ANOVAs were also performed on the entire subject pool to determine the effect of risk on EMG activity within each epoch (baseline, R1, R2, R3, and voluntary). The R model was `aov(EMG Risk)` where the EMG factor was only the EMG activity in a single epoch. This was performed when the EMG was normalized by preactivation (100 ms prior to perturbation) instead of MVC as well.

Additionally, the asymmetric cost functions were analyzed in greater detail. The asymmetric cost functions were categorized by whether the perturbation was in the direction of overall higher risk or lower risk. Since only trials that activated the FDI reflex were analyzed, this meant that if the reflex response was greater in conditions with higher risk to the right than to the left, that the reflexes were tuned according to the location of risk in the environment. A two-way repeated measures ANOVA was performed on the EMG from the asymmetric cost functions. The R model was `aov(EMG Side*Epoch)` where the Side factor indicated the direction of higher cost. Post hoc pairwise comparisons were made of the effect of higher risk direction on the EMG activity within each epoch as well.

EMG from one subject did not exhibit a normal stretch reflex response to perturbation. It is unclear whether the data was accurately reflecting an abnormal response or the result of error in data collection, but this subject was removed from analysis.

### 4.3.2 Results

This experiment was designed to study the effects of the risk on the stretch reflex. Subjects did not demonstrate difficulty in performing the task and all subjects appeared to use the same strategy of staying within the center reward region and attempting to resist perturbation to avoid risk. We consistently observed a small peak in muscle activity at 50 ms post perturbation and a larger, distinct peak at 85 ms characteristic of the long latency reflex. Each subject was first analyzed individually. Seven of the nine subjects demonstrated a significant difference in muscle activity between cost functions (subjects 2, 4, 5, 6, 7, 11, and 13); six of those seven subjects tuned their muscle activity appropriate (generally increased reflex amplitude in higher risk conditions).

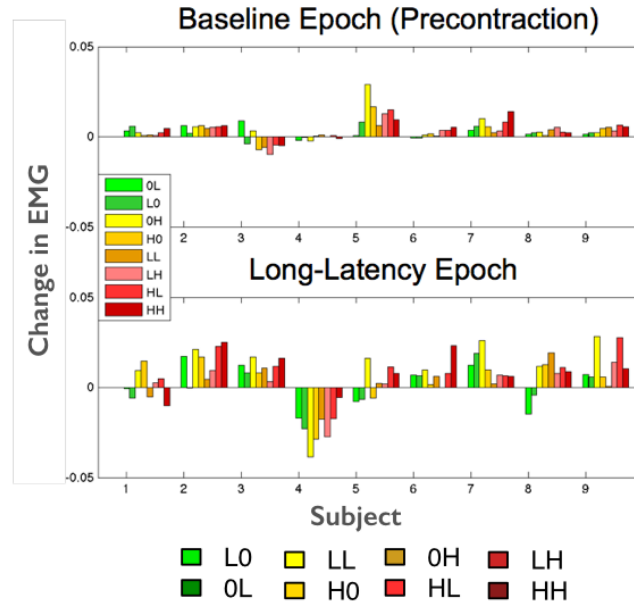


Figure 4.8: **Change in EMG Between Conditions for Individual Subjects.**

This figure contains the change in average baseline EMG (100ms prior to perturbation) between the symmetric no-cost/symmetric low-cost (yellow) and symmetric no-cost/symmetric high-cost (red) for each subject, and figure Aii contains the same for the long latency epoch. Subject number is located on the x-axis and EMG on the y-axis. Figure B contains the average EMG trace for all subjects for the symmetric no-cost (green), symmetric low-cost (yellow), and symmetric high-cost (red) tasks. The dashed black line designates the time of perturbation and the solid black lines indicate the long latency epoch.

EMG activity across all subjects was analyzed within each epoch as well. The baseline epoch showed a significant difference ( $F(8,1907) = 3.302$ ,  $p < 0.001$ ), there was no significant difference in the R1 epoch as expected ( $F(8,1907) = 1.407$ ,  $p = 0.189$ ). In the long latency epochs, R2 and R3, there was a significant different dependent on cost ( $F(8,1907) = 2.816$ ,  $p=0.00418$ ) and ( $F(8,1907) = 4.323$ ,  $p < 0.0001$ ). There was also a significant difference in the voluntary epoch ( $F(8,1907) = 4.739$ ,  $p < 0.0001$ ).

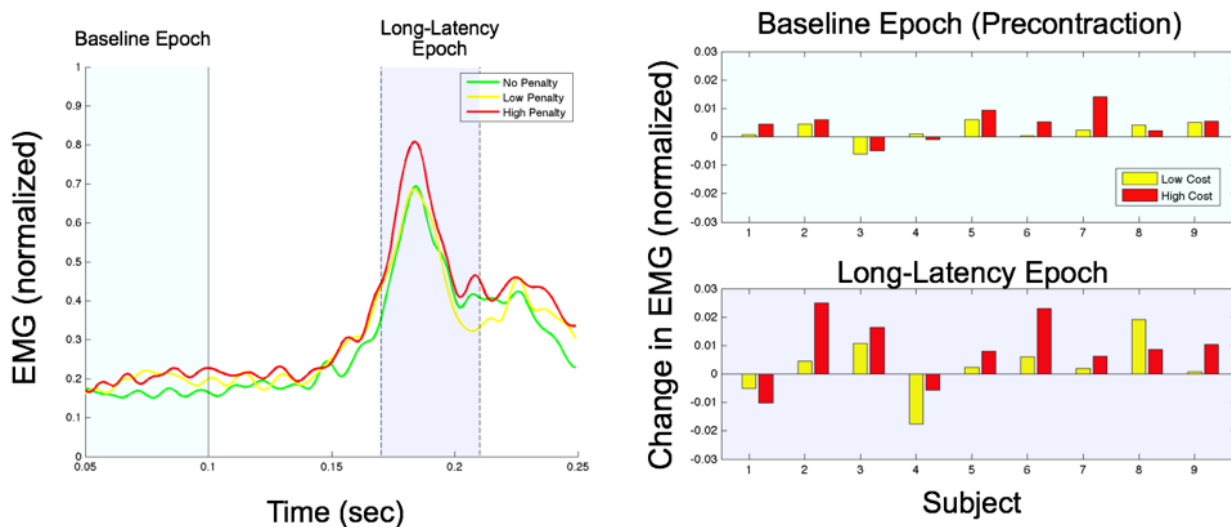


Figure 4.9: **Symmetric Cost EMG.**

Figure Ai contains the change in average baseline EMG (100ms prior to perturbation) between the symmetric no-cost/symmetric low-cost (yellow) and symmetric no-cost/symmetric high-cost (red) for each subject, and figure Aii contains the same for the long latency epoch. Subject number is located on the x-axis and EMG on the y-axis. Figure B contains the average EMG trace for all subjects for the symmetric no-cost (green), symmetric low-cost (yellow), and symmetric high-cost (red) tasks. The dashed black line designates the time of perturbation and the solid black lines indicate the long latency epoch.

The experimental design included both symmetric and asymmetric cost functions while the subject did not know the direction of perturbation. Therefore we were able to determine the degree of specificity of the environmental cost function that the reflex response would integrate. The asymmetric cost functions were categorized by whether the perturbation was

in the direction of overall higher risk or lower risk. Since only trials that activated the FDI reflex were analyzed, this meant that if the reflex response was greater in conditions with higher risk to the right than to the left, that the reflexes were tuned according to the location of risk in the environment. However, only five of the nine subjects actually demonstrated this tuning. As a whole, there was no significant effect of the location of the risk in the environment ( $F(1,6492) = 0.735$ ,  $p = 0.391$ ). A pairwise t-test comparing the asymmetric cost direction within the long latency epoch only did not exhibit a significant difference ( $p=.184$ ).

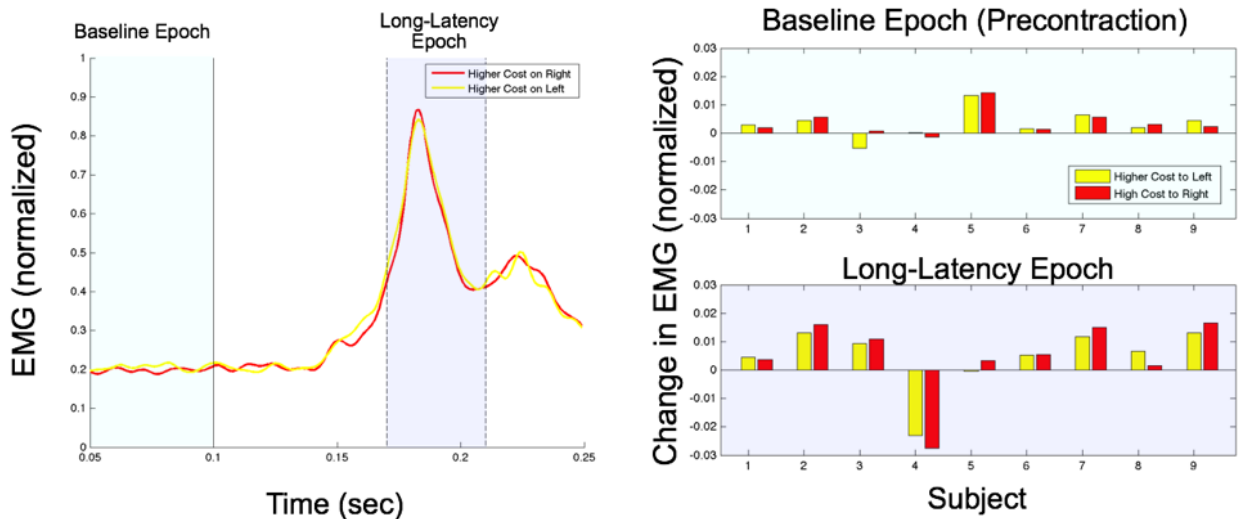


Figure 4.10: **Asymmetric Cost EMG.**

Figure Ai contains the change in average baseline EMG (100ms prior to perturbation) between the no-cost condition and asymmetric cost conditions; yellow indicates the conditions with perturbations in the direction of lower cost and red represents the perturbations in the direction of higher cost. Figure Aii contains the same for the long latency epoch. Subject number is located on the x-axis and EMG on the y-axis. Figure B contains the average EMG trace for all subjects for perturbations towards lower cost (yellow) and perturbations toward higher cost (red). The dashed black line designates the time of perturbation and the solid black lines indicate the long latency epoch.

In the previous analysis, the EMG was normalized by maximum voluntary control. If the EMG is instead normalized by muscle preactivation, there is no longer a clear signifi-

cant difference in EMG activation between risk environments. None of the epochs show a significant difference, including the voluntary epoch, R1: ( $F(8,1907) = 0.327$ ,  $p = 0.956$ ), R2: ( $F(8,1907) = 1.074$ ,  $p = 0.378$ ), R3: ( $F(8,1907) = 0.73$ ,  $p = 0.665$ ), Vol: ( $F(8,1907) = 0.974$ ,  $p = 0.454$ ).

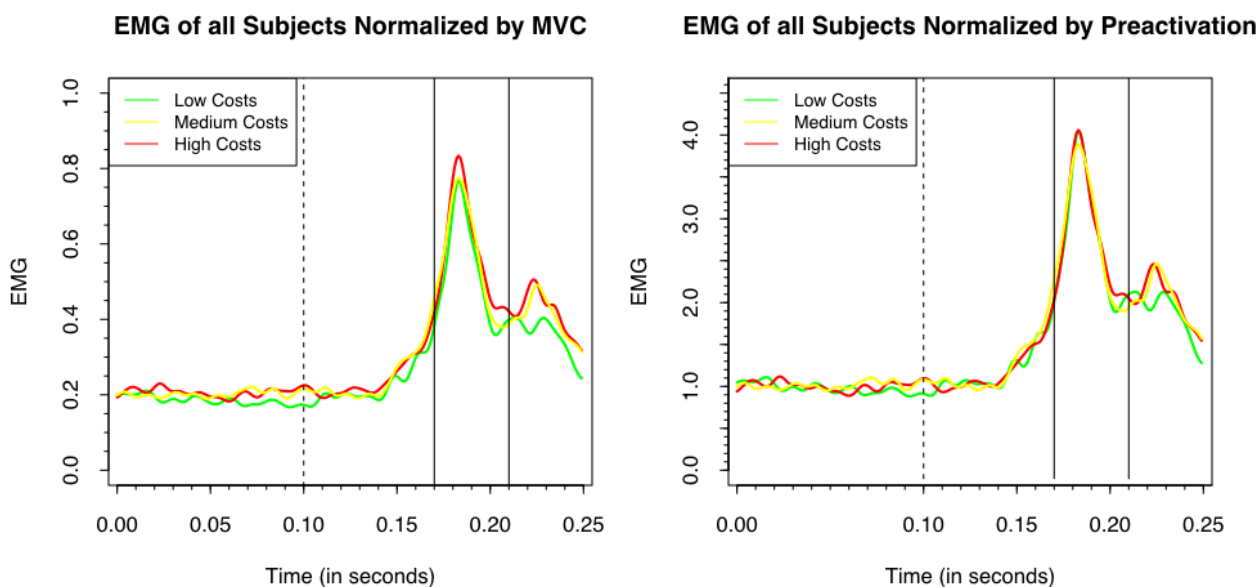


Figure 4.11: **EMG Normalized by Maximum Voluntary Contraction and by Pre-activation.**

Figure Ai contains the change in average baseline EMG (100ms prior to perturbation) between the no-cost condition and asymmetric cost conditions; yellow indicates the conditions with perturbations in the direction of lower cost and red represents the perturbations in the direction of higher cost. Figure Aii contains the same for the long latency epoch. Subject number is located on the x-axis and EMG on the y-axis. Figure B contains the average EMG trace for all subjects for perturbations towards lower cost (yellow) and perturbations toward higher cost (red). The dashed black line designates the time of perturbation and the solid black lines indicate the long latency epoch.

### 4.3.3 Discussion

This experiment was designed such that the goal of task remained the same through the experiment and did not depend on the perturbation. Therefore, if the subject had performed their best on every trial, they would have received as many points as it was possible for them

to get. However, most of the subjects still demonstrated appropriate modulation of the long latency reflex to the risk in the environment. Across all subjects combined, a significant difference was found in the long latency epoch between cost functions when the muscle activity was normalized based on maximum voluntary control. This suggests that humans to tune their long latency reflexes to the risk of the environment. However, there was not a significant difference between cost functions that pushed toward higher cost versus away from higher cost. This suggests that subject do not set separate reflex responses for different locations of risk simultaneously when the direction of perturbation was unplanned.

The results were also analyzed with the EMG normalized by preactivation instead of maximum voluntary control. Under this adjustment, there was no longer a significant difference in the muscle activity within any epoch. This suggests that muscle preactivation may be an essential component to modulating the reflex response to risk. Pruszynski et al. provide evidence that the long latency reflex may be composed of two functionally distinct components (Pruszynski et al, 2011). One that is modulated by the goal of the task independent of tone, and another that is sensitive to muscle preactivation. It is possible that the goal of the task predominantly modulates the first component of the long latency reflex and the risk of the environment primarily modulates this second preactivation dependent component.

The evidence that stiffness was affected by the risk of the environment indicates that subjects were planning for error prior to the perturbation. This is not represented in optimal feedback control. In optimal feedback control, the reference trajectory is recalculated at each time point depending on the current state. This accounts for redirecting movement not to the previous path, but to the best path to achieve a goal. In the case of this experiment, that would be toward the center target and away from the cost once the perturbation occurred. However, the results demonstrate not just that subjects redirected their goal appropriately post perturbation, but also planned for that error prior to its occurrence.

It is problematic to draw more specific conclusions about the role of co-contraction in

this study since the muscle antagonist to the FDI cannot be measured reliably with surface electrodes. Therefore, the final experiment in this chapter will repeat this experimental design using the bicep and triceps muscle. This will provide an opportunity to investigate more deeply the role of stiffness in tuning the reflex response to the risk of the environment.

## 4.4 Experiment 3: Role of Cocontraction in Tuning Reflexes

### 4.4.1 Introduction

The final study in this series of experiments was designed to validate the results of the previous experiment and specifically look more closely at the role of co-contraction in modulating the stretch reflex to risk. It is well understood that co-contraction will act as a gain on the reflex response (Akawaza et al., 1983; Lewis et al., 2010). It is possible that humans tune their stretch reflex to risk only by modulating their tone or muscle stiffness. In fact, stiffness likely has a large role to play in modulating reflexes of an unpracticed task. An example of this is walking on a balance beam. It is expected that the average person will experience no difficulty walking in a straight line on a rigid balance beam planted firmly on the floor. However, lift that balance beam a hundred feet into the air and most people will change the way they move across it, often in a disadvantageous manner. This is as the cliché "scared stiff" suggests, due at least in part to changes in muscle tone in response to danger.

However, modifying stiffness may not be the only manner in which humans tune reflexes to risk. It is possible that we tune reflexes independent of tone as well. The previous study established that reflexes are tuned to the risk, however the results were not particularly clear, especially in regards to stiffness. Therefore, the experiment was repeated in the bicep muscle. The paradigm was very similar to the previous experiment, but EMG from both the agonist and antagonist muscles were recorded.

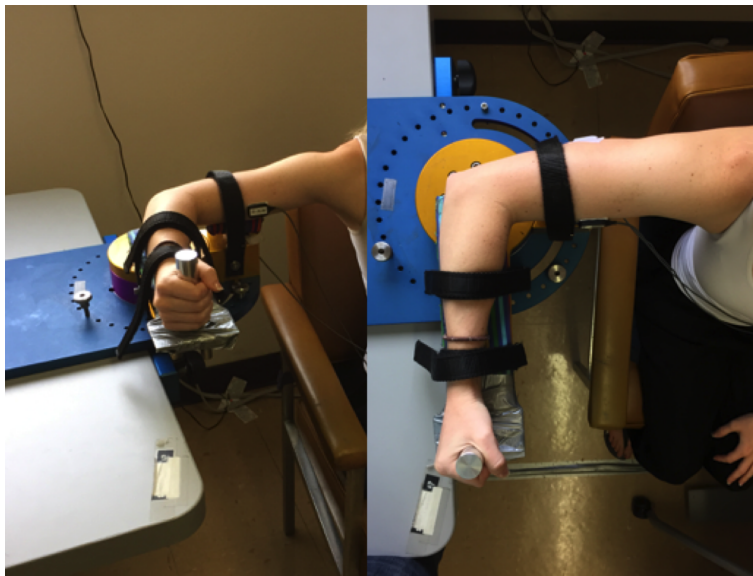
### 4.4.2 Materials and Methods

#### *Subjects*

Ten subjects, 7 males and 3 females, participated in this experiment. Other than standard inclusion criteria for a normal subject, each subject also had to be strong enough to complete the majority of the experiment.

### *Stimuli and Procedure*

In this experiment, subjects were positioned in front of a monitor with their right arm strapped to a manipulandum designed to apply torque at the elbow joint while maintaining all other arm joints immobile. The subject's hand gripped a rigid joystick attached to the arm of the manipulandum that controlled a cursor horizontally on the screen. The display provided to the subjects was identical to the previous experiment, 4.7a. The only change was that the rectangles were stationary in order to measure the co-contraction more accurately. Electrodes were placed on the bicep and triceps muscles and a ground electrode was placed on the surface of the opposite hand. The figures below portray the manipulandum and physical set-up of the experiment.



**Figure 4.12: Manipulandum Set-Up**

The figures above show the setup of the manipulandum and the position of the arm. The subject's arm was strapped to the robot that controlled the blue cursor on the screen. The elbow joint was positioned directly over the robot joint exerting the perturbation torque. Electrodes were placed on the bicep and triceps muscles.

The baseline force and the perturbation force were the same between subjects in order to compare reflexes between subjects. However, the strength of each subject influenced the ability of the subject to perform the task. To account for the difference in strength between

subjects, the scaling factor between the displacement of the robot and the displacement of the screen cursor was calibrated to each subject. Prior to the start of the described experiment, subjects completed a short calibration phase. They were presented with the same screen with no penalty regions and were asked to withstand the force of the perturbation as much as possible. There were ten perturbations, five in each direction, for the calibration. The maximum point the robot reached in response to each perturbation was recorded and averaged for all perturbations. The scaling factor was calculated so that the distance between the reward region and each penalty region was 90% of this value. Still, the strength of the subject affected their ability to perform the task. The stronger subjects never hit a single penalty region while some of the weaker subjects hit a penalty region on more than 25% of the trials.

The experiment was divided into two sessions to take place on separate days in order to minimize the effect of fatigue. Each session consisted of six blocks of 45 perturbations. The presented cost was randomly changed every 5 perturbations and the direction of the perturbation was random. Not all subjects were able to complete all the blocks due to fatigue, so subjects were given the option to stop the experiment once they felt they could no longer continue. Approximately half the subjects did not complete the full experiment, but all subjects included in the analysis completed more than half of the total trials.

### *Data Analysis*

Results were analyzed in a similar manner as the previous experiment. Only trials in the direction that activated the bicep stretch reflex were analyzed. Therefore, there were approximately 30 trials analyzed for each condition for each subject that completed the full experiment. EMG was divided into baseline  $[-50 - 0 \text{ ms}]$ , short latency  $[0 - 50 \text{ ms}]$ , long latency  $[50 - 105 \text{ ms}]$ , and voluntary  $[105 - 150 \text{ ms}]$  epochs. The EMG data was low pass filtered  $[500 \text{ Hz}]$ , full-wave rectified, then bandpass filtered  $[25 - 250 \text{ Hz}]$ , butterworth.

Statistics were done using the mean of the EMG within each epoch. In order to

account for fatigue, EMG was normalized by the average of the maximum EMG (from -50 – 150 ms) from all trials in the same direction within each block (i.e. each block had its own normalization factor). The analysis was repeated with the EMG normalized by baseline as well in order to determine the amount of risk modulation performed by adjusting muscle tone. Refer to section 4.3.1 for more details on data analysis details.

#### 4.4.3 Results

Five of the ten subjects significantly tuned their long latency reflex based on the risk of the environment ( $p < 0.05$ ). Four of these five subjects appropriately tuned their reflexes, so that their reflex response was generally higher for higher risk. Although not significant, several other subjects still demonstrated increased reflex response in environments with increased risk. Figure 4.13 shows the average change in EMG between risk conditions for each subject.

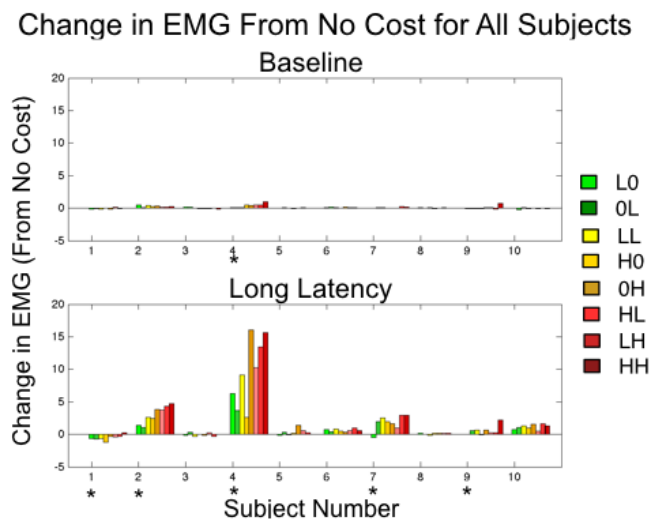


Figure 4.13: **Individual Subject EMG**

The figures above show the average change in EMG between the no cost condition and each other condition for each subject. The top panel shows the EMG from the baseline epoch and the bottom panel shows the long latency epoch. Stars indicate subjects whose long latency reflex was significantly different with respect to risk ( $p < 0.05$ ).

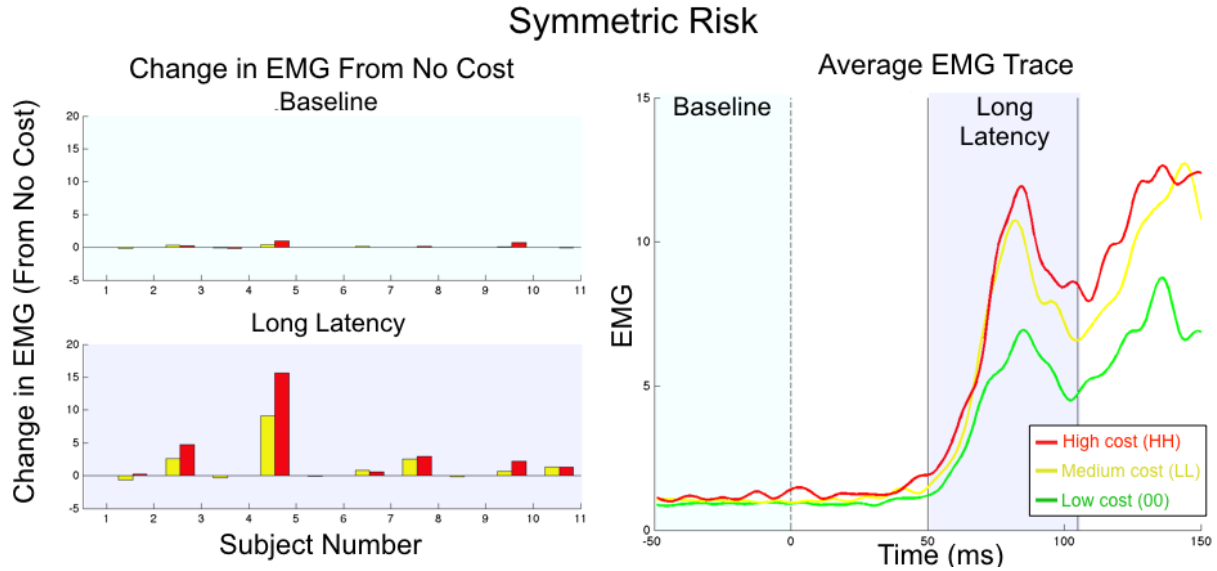


Figure 4.14: **Reflex Response to Symmetric Risk.**

The plot on the right shows the average EMG trace for the symmetric cost conditions. The red line indicates the EMG response to the highest symmetric risk (high cost-high cost), the yellow line indicates the response to the low symmetric risk (low cost-low cost), and the green line represents the response to the no cost condition (no cost-no cost). The bar plot indicates the average difference in EMG between the symmetric high risk/symmetric no cost (red) and the difference between the symmetric low risk/symmetric no cost (yellow). Subject number is indicated on the x-axis.

In the symmetric risk conditions, a repeated-measures ANOVA indicated there was a significant effect of risk on the long latency epoch ( $F(2,635) = 23.35$ ,  $p < 0.001$ ). There was also a significant difference within the baseline epoch ( $F(2,635) = 4.636$ ,  $p < 0.01$ ). The maximum of the peak in the long latency epoch was 11.93 for the symmetric high-risk, 10.75 for the symmetric low-risk, and 6.95 for the symmetric no-risk. This is a large difference in response since the amplitude of the highest risk was almost twice that of the lowest risk.

In the asymmetric condition, the overall cumulative risk was the same, but the risk is the direction of the perturbation was considered. The results of the change in the baseline and long latency epoch for each subject for each cumulative risk can be seen in figure 4.15(a). The trace depicting the average response towards higher risk in the asymmetric costs vs towards lower cost can be seen in 4.15b. There is a significant difference ( $F(1,1286) = 11.17$ ,

$p < 0.001$ ) within subjects for the long latency reflex responses pushing toward higher risk vs lower risk. There is no significant difference in the baseline ( $F(1,1286) = .954$ ,  $p = 0.329$ ), which we would expect since the baseline should reflect the co-contraction.

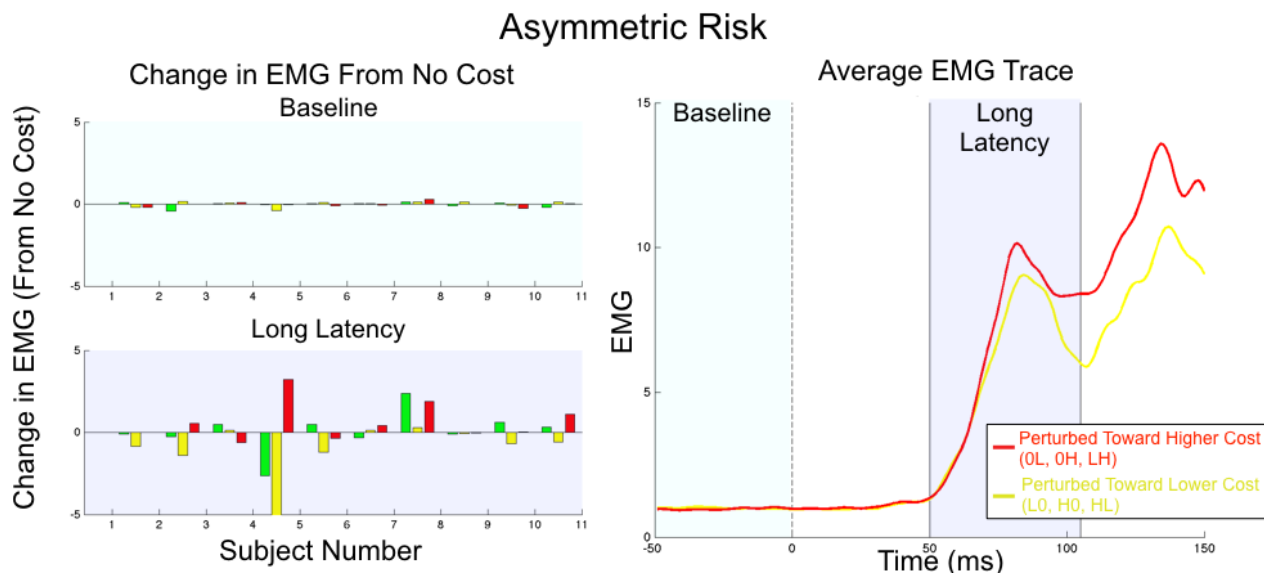


Figure 4.15: **Reflex Response to Asymmetric Risk.**

The plot on the right shows the average EMG trace for the asymmetric cost conditions. The red line indicates the EMG resulting from perturbations toward higher cost and the yellow line indicates toward lower cost. The cumulative cost of the conditions are the same, the only difference is the side of the risk. The bar plot indicates the difference in the average long latency EMG between the higher risk in the direction of the perturbation and the higher risk away from the perturbation for the same cumulative cost for each subject. Subject number is indicated on the x-axis.

The average traces for the cumulative risk was also calculated, figure 4.16(a). Considering all subjects, an ANOVA assuming risk was numerically equal to the cumulative risk of the environment showed a significant difference in long latency response between risks ( $F(1,3214) = 23.71$ ,  $p < .001$ ). Additionally, if the risk was assumed numerically equivalent to the asymmetric risk in the direction of perturbation, there was also a significant difference in long latency response to risk ( $F(1,3214) = 12.99$ ,  $p < .001$ ). This indicates that the modulating of the long latency stretch reflex to risk was most likely a combination of co-contraction and asymmetric tuning. The traces for the EMG response normalized by the

baseline of that trial can be seen in 4.16(b). There was a significant difference in the long latency response overall to risk as well ( $F(8.1926) = 2.086, p < .05$ ).

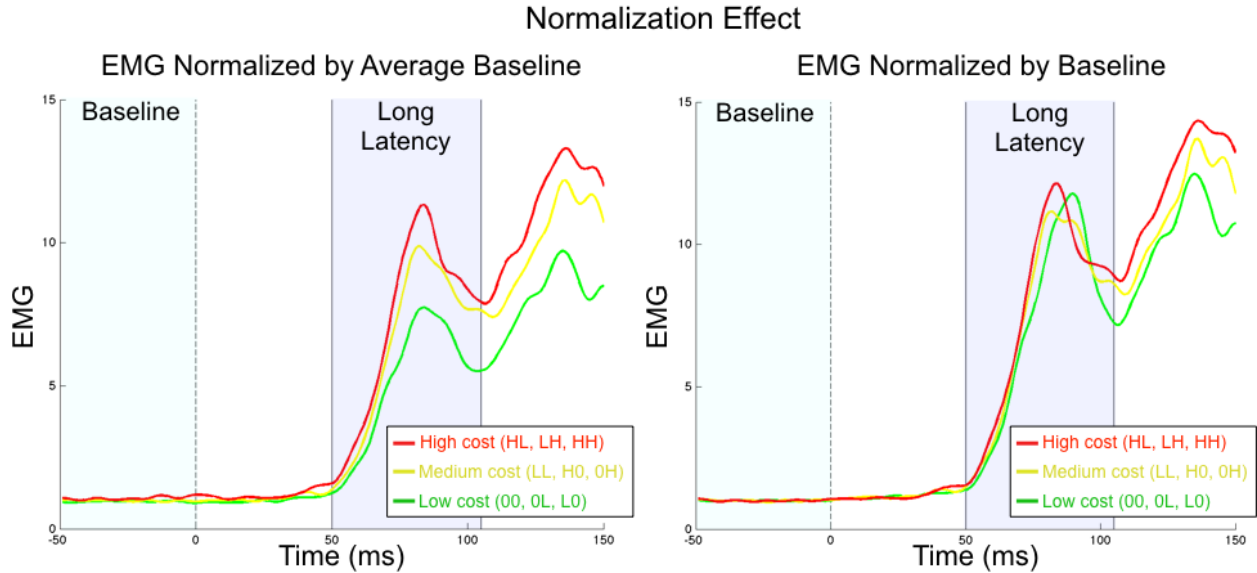


Figure 4.16: **Average EMG: Normalization.**

The lines indicate the average EMG trace for all subjects. The green line indicates the three lowest cumulative costs (no cost-no cost; low cost-no cost; no cost-low cost), the yellow line indicates the average of the three middle cumulative costs (low cost-low cost, high cost-no cost, no cost-high cost), and the red line indicates the average EMG response to the three highest cumulative costs (high cost-low cost; low cost-high cost; high cost-high cost). In the left plot, each EMG trace is normalized by the average baseline of that block. In the right plot, each EMG trace is normalized by the baseline of that EMG trace.

#### 4.4.4 Discussion

We focus on two methods that humans may utilize to modify their stretch reflex in response to risk. The first is the co-contraction/symmetric hypothesis, in which subjects tune their reflexes to risk by stiffening their joints. If this were the case then subjects would likely tune a single parameter representing the cumulative risk of the environment. The second is the asymmetric reflex, which hypothesizes that humans tune their reflexes specifically to wherever risk lies in the environment. This cannot be done with co-contraction since an asymmetry in contraction would produce movement. The null hypothesis is that humans

do not tune their reflexes to risk. These hypothesis are shown in figure 4.17 and can be compared to subject data in figure 4.13.

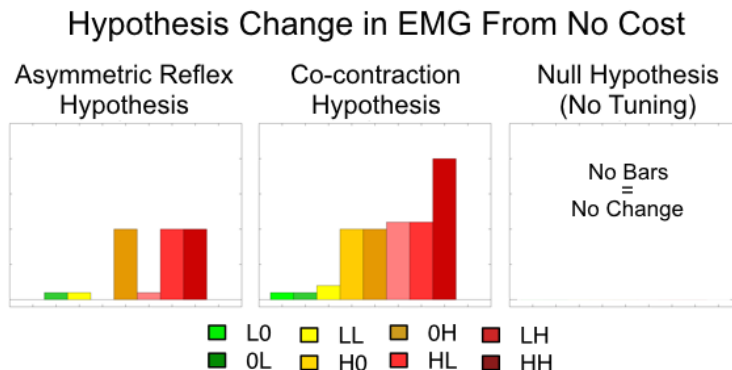


Figure 4.17: **Hypotheses Visualized.**

The figures above illustrate the null hypothesis as well as the asymmetric and co-contraction hypothesis. Each bar indicates the difference between the no cost condition and each other condition. This study was designed to determine which of the above hypotheses were valid. The co-contraction or symmetric hypothesis is that subjects tune their stretch reflex to the overall or cumulative risk of the environment. The asymmetric hypothesis is that humans tune their stretch reflex not only to the risk of environment, but more specifically to the risk in the direction of the perturbation as well. The null hypothesis is that subjects do not modulate their long latency reflex based on risk.

Response to risk was not uniform for all subjects, however overall subjects did tune their long latency reflex to risk. It appeared that subjects primarily used stiffness to compensate for risk, however subjects did tune their long latency reflex asymmetrically as well. Subjects long latency reflex was significantly increased in the direction of higher risk between trials when the cumulative risk was the same and the direction of perturbation could not be anticipated. This suggests that stiffness is not the only modulator of reflexes in response to risk since co-contraction is inherently symmetric.

This was not found in the case of the FDI. It is possible that this is a learned behavior and that some humans are more capable of asymmetric tuning than others. This component of the response was very variable between subjects. More likely this reflects the physiological differences between the FDI and bicep reflexes noted by other studies (Thilmann et al., 1991).

## 4.5 Chapter Conclusions

In order to keep us safe, risk must play a critical role in shaping movement. The first chapter demonstrated that both uncertainty and a detailed understanding of the environmental cost function modulate feedback-driven behavior in humans. However, if risk is as influential as we expect, modulation of behavior to risk should be prevalent at all levels of movement. This chapter explored the human reflex response and the role of risk in modulating this response.

The first step was to survey the very general characteristics of rapid response to perturbations to determine if this was a valid avenue of research. Response to visual perturbation confirmed an increase in response amplitude when the perturbation increased risk. The next step was to directly test the stretch reflex response. Many studies have demonstrated that humans have the ability to modify their long latency reflex depending on the goal of the task (Ludvig et al, 2007; Pruszynski et al, 2008; Hammond, 1956; Rothwell et al, 1980; Crago et al, 1976). We were very careful to design an experiment that did not simply repeat these findings, and instead we attempted to specifically address the role of risk and not the effect of the goal on the reflex response. While it is true that these two concepts are intimately linked, in the experiment implemented, the goal of the task remained constant throughout the entire experiment. If the subject performed their best on every trial, they would have received maximum points (with respect to that subject). However, results still show a significant difference in the long latency reflex response dependent on the level of risk in the task.

The bicep reflex in particular demonstrated the ability to tune the reflex response using a combination of both co-contraction and asymmetric tuning. It would be interesting in the future to see how complicated a cost function the asymmetric tuning could represent.

This set of studies is significant because in optimal feedback control, redirecting the movement toward the target post-perturbation arises from the recalculation of the optimal

trajectory at every time point. The consequence is that while the response to the perturbation will demonstrate modulation of the reflex response, it does not account for planning for perturbation based on risk (such as setting tone). However, when controlling the system through dynamics, tuning reflexes is an inherent result of the system.

## Chapter 5

### Concluding Remarks

#### 5.1 Conclusion

The first experiment determined that in a feedback-driven task, humans maintain a detailed understanding of uncertainty and the form of the cost function. Moreover, results suggested that humans make predictions of the likelihood of failure, without having experienced this failure first. This is a requisite ability if movement is governed by risk, as there are failures (such as falling off a cliff) that one cannot experience first. The importance of this study is that it established evidence that humans may maintain not just maximum likelihoods, but entire probability distributions of state and cost of state. The second set of experiments extended the idea of risk governing movement to the role of risk in modulating the most fundamental type of movement: the stretch reflex. We found that subjects did tune their long latency stretch reflex to the overall risk of the environment. Results suggested that humans might do this primarily through adjusting stiffness in response to risk. Another experiment will be performed to examine more closely the function and necessity of cocontraction for adapting the reflex response to the risk of the environment.

The last experiment will revisit the role of uncertainty in motor control. It will compare two fundamentally different types of uncertainty: sensory or current state uncertainty with motor or future state uncertainty. The purpose of investigating sensory uncertainty is to determine if certainty equivalence is a valid assumption to make in motor control. If the system, human movement, proves to behave differently under state uncertainty than when

the state is certain, then the certainty equivalence property does not hold. This is relevant because it is a very common, though rarely addressed, assumption.

The unifying aspect of these results is that they represent fundamental characteristics of human movement that are lacking or absent from current implementations of classical motor control theories. Any complete model of human movement must exhibit these behaviors as well. The goal of these studies is not simply to demonstrate human behavior, but to persuade the reader to consider an alternative perspective on motor control that moves away from the traditional trajectory-based viewpoint and instead proposes that movement results from maintaining probability distributions of the probability of failure and cost of failure.

## **5.2 Applications and Impact**

Building a true model of human movement has many potential benefits. These benefits mostly fall into one of two categories: exploiting the model to gain insight into motor diseases or implementing this control in artificial systems to replicate the advantages of human movement.

Human movement is consistently effective in achieving a goal while still avoiding risk in the environment, and at the same time is robust to perturbations and easily adaptable to new variations. The field of robotics is still unable to generate movement with these indispensable characteristics in real time in artificial systems. Additionally, we can use this knowledge to help replace function in individuals with motor impairment. Prosthetics can be designed with the specific form of control in mind to reproduce as natural of movement as possible.

The second class of advantages a deeper understanding of motor control will provide is greater insight into motor disease. We can compare symptoms of specific motor diseases with symptoms resulting from breaking the working model to help us understand the origin of dysfunction. Furthermore, it will help generate new ideas for solutions and treatments

once the source of impairment has been identified. Once we understand the system, it will have many implications in learning and may even transform the way we teach motor actions.



## Bibliography

- [1] Adrian E (1962) The Impulses Produced by Sensory Nerve Endings: Part I. *J Physiol* 61: 49-72.
- [2] Akazawa K, Milnder T, Stein R (1983) Modulation of Reflex EMG and Stiffness in Response to Stretch of Human Finger Muscle. *J Neurophysiol* 49(1): 16-27.
- [3] Allard R, Faubert J (2008) The Noisy-Bit Method for Digital Displays: Converting a 256 Luminance Resolution into a Continuous Resolution. *Behav Res Methods* 40: 735-743.
- [4] Attneave F (1954) Some Informational Aspects of Visual Perception. *Psychol Rev* 61: 183-193.
- [5] Bhumbra G, Dyball R (2005) Spike Coding from the Perspective of a Neurone. *Cogn Process* 6: 157-176.
- [6] Blackwell R (1952) Studies of Psychophysical Methods for Measuring Visual Thresholds. *J Opt Soc Am* 42: 606-616.
- [7] Braun D, Nagengast A, Wolpert D (2011) Risk-sensitivity in sensorimotor control. *Front Hum Neurosci* 5(1): 1-10.
- [8] Burdet E, Osu R, Franklin D, Milner T, Kawato M (2001) The Central Nervous System Stabilizes Unstable Dynamics by Learning Optimal Impedance. *Nature* 44: 446-449.
- [9] Campbell F, Green D (1965) Optical and Retinal Factors Affecting Visual Resolution. *J Physiol* 181: 576-593.
- [10] Campbell F (1968) The Human Eye as an Optical Filter. *P IEEE* 56: 1009-1014.
- [11] Campbell F, Robson J (1968) Application of Fourier analysis to the visibility of gratings. *J Physiol* 197: 551-566.
- [12] Crago P, Houk J, Hasan Z (1976) Regulatory Actions of Human Stretch Reflex. *J Neurophysiol* 39: 925-935.
- [13] Diedrichsen J, Shadmehr R, Ivry R (2010) The Coordination of Movement: Optimal Feedback Control and Beyond. *Trends Cogn Sci* 14(1): 31-39.

- [14] Dunning A, Ghoreyshi A, Bertuccio M, Sanger T (2015) The Tuning of Human Motor Response to Risk in a Dynamic Environment Task. *PLoS ONE* 10(4): e0125461. doi: 10.1371/journal.pone.0125461
- [15] Faisal A, Wolpert D (2009) Near Optimal Combination of Sensory and Motor Uncertainty in Time During a Naturalistic Perception-Action Task. *J Neurophysiol* 101: 1901-1912.
- [16] Farin, G (1997) *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*. San Diego, CA: Academic Press, Inc. 33-79.
- [17] Galen G, John W (1995) Fitts' Law as the Outcome of a Dynamic Noise Filtering Model of Motor Control. *Hum Mov Sci* 14: 539-571.
- [18] Glansdorf P, Prigogine I (1971) *Thermodynamic Theory of Structures, Stability and Fluctuations*. Wiley.
- [19] Guilford J (1954) *Psychometric Methods* (New York: McGraw-Hill).
- [20] Hammond P (1956) The Influence of Prior Instruction to the Subject on an Apparently Involuntary Neuro-muscular Response. *J Physiol* 132: 17P-18P.
- [21] Haith A, Krakauer J (2013) Theoretical Models of Motor control and Motor Learning n: Gollhofer A, Taube W, Bo Nielsen J, (Eds). *The Routledge Handbook of Motor Control and Motor Learning*. 7-28. Routledge, London.
- [22] Harvey L (1986) Efficient estimation of sensory thresholds. *Behav Res Meth Ins C*, 18: 623-632.
- [23] Inouye J, Valero-Cuevas F (2016) Muscle Synergies Heavily Influence the Neural Control of Arm Endpoint Stiffness and Energy Consumption. *PLoS Compute Biol* 12(2): 1-24.
- [24] Karber G (1931) Beitrag zur kollektiven Behandlung pharmakologischer Reihenversuche [A contribution to the collective treatment of a pharmacological experimental series]. *Archib fur experimentelle Pathologie und Pharmakologie* 162: 480-483.
- [25] Keele S, Posner M (1968) Processing of Visual Feedback in Rapid Movements. *J Exp Psychol* 77(1): 155-158.
- [26] Kendrick D (2002) *Stochastic Control for Economic Models* 2nd Edition. The University of Texas.
- [27] Kording K, Wolpert D (2004) Bayesian Integration in Sensorimotor Learning. *Nature* 427: 244-247.
- [28] Kording K, Wolpert D (2006) Bayesian decision theory in sensorimotor control. *Trends Cogn Sci* 10(7): 319-326.

- [29] Knight B (1972) Dynamics of encoding in a population of neurons. *J Gen Physiol* 59: 734
- [30] Knill D, Pouget A (2004) The Bayesian Brain: The Role of Uncertainty in Neural Coding and Computation. *Trends Neurosci* 27(12): 712-719.
- [31] Kolb, H. Facts and Figures Concerning the Human Retina. 2005 May 1 [Updated 2007 Jul 5]. In: Kolb H, Fernandez E, Nelson R, editors. *Webvision: The Organization of the Retina and Visual System*. Salt Lake City (UT): University of Utah Health Sciences Center; 1995-. Available from: <http://www.ncbi.nlm.nih.gov/books/NBK11556/>
- [32] Kulikowski J and King-Smith P (1973) Spatial arrangement of line, edge and grating detectors revealed by subthreshold summation. *Vision Res* 13: 1455-1478.
- [33] Landy M, Trommershauser J, Daw N (2012) Dynamic Estimation of Task-Relevant Variance in Movement under Risk. *J Neurosci* 32(37): 12702-12711.
- [34] Latash M (2010a) Motor Synergies and the Equilibrium-Point Hypothesis. *Motor Control* 14(3): 294-322.
- [35] Latash M (2010b) Two Archetypes of Motor Control Research. *Motor Control* 14(3): e41-e53.
- [36] Lewis G, MacKinnon C, Trumbower R, Perreault E (2010) Co-contraction Modifies the Stretch Reflex Elicited in Muscles Shortened by a Joint Perturbation. *Exp Brain Res* 207(1-2): 39-48.
- [37] Ludvig D, Cathers I, Kearney R (2007) Voluntary Modulation of Human Stretch Reflexes. *Exp Brain Res* 183:201-213.
- [38] Maloney L, Trommershauser J, Landy M (2007) Questions Without Words: A Comparison between decision making under risk and movement planning under risk. In: Gray, W. (Ed), *Integrated Models of Cognitive Systems*. New York, NY: Oxford University Press. 297-313.
- [39] Maloney L, Zhang H (2010) Decision-theoretic models of visual perception and action. *Vision Res* 50(23): 2362-2374.
- [40] Meister M, Berry M (1999) The neural code of the retina. *Neuron* 22: 435-450.
- [41] Michelson A (1927) *Studies in Optics*. (Chicago: University of Chicago Press).
- [42] Miller F, Ulrich R (2001) On the analysis of psychometric functions: The Spearman-Kärber method. *Percept Psychophys* 63: 1399-1420.
- [43] Miller F, Ulrich R (2004) Threshold estimation in two-alternative forced-choice (2AFC) tasks: The Spearman-Kärber method. *Percept Psychophys* 66: 517-533.

- [44] Morel P, Baraduc P (2010) Dissociating the Impact of Sensory and Motor Noise in Human Saccades. *Neurocomp* 10.
- [45] Mussa-Ivaldi FA, Human N, and Bizzi E (1985) Neural, Mechanical and Geometric Factors Subservicing Arm Posture in Humans. *J Neurosci* 5: 2732-2743.
- [46] Nachmias J (1981) On the psychometric function for contrast detection. *Vision Res* 21: 215-223.
- [47] Urban G, Wolpert D (2011) Representations of Uncertainty in Sensorimotor Control. *Curr Opin Neurobiol* 21:1-7.
- [48] Owsley C (2003) Contrast Sensitivity. *Ophthalmol Clin North Am* 16: 171-177.
- [49] Perreault E, Kirsch R, Crago P (2002) Voluntary Control of Static Endpoint Stiffness During Force Regulation Tasks. *J Neurophysiol* 87: 2808-2816.
- [50] Pruszynski A, Isaac K, Scott S (2011) The long-latency reflex is composed of two functionally independent processes. *J Neurophysiol* 106(1): 449-459.
- [51] Pruszynski A, Kurtzer I, Scott S (2008) Rapid Motor Responses Are Appropriately Tuned to the Metrics of a Visuospatial Task. *J Neurophysiol* 100:224-238.
- [52] Quick R (1974) A vector magnitude model of contrast detection. *Kybernetik* 16: 65-67.
- [53] Ravichandran V, Honeycutt C (2013) Instruction-dependent modulation of the long-latency stretch reflex is associated with indicators of startle. *Exp Brain Res* 230:59-69.
- [54] Rothwell J, Traub M, Marsden C (1980) Influence of voluntary intent on the human long-latency stretch reflex. *Nature* 286: 496-498.
- [55] Rullen R, Thorpe S (2001) Rate Coding Versus Temporal Order Coding: What the Retinal Ganglion Cells Tell the Visual Cortex. *Neural Comput* 13: 1255-1283.
- [56] Sanger T (2011) Distributed Control of Uncertain Systems Using Superpositions of Linear Operators. *Neural Comput* 23: 1911-1934.
- [57] Sanger T (2014) Risk-Aware Control. *Neural Comput* 26(12): 2669-91.
- [58] Shadmehr R, Arbib M (1992) A Mathematical Analysis of the Force-Stiffness Characteristics of Muscles in Control of a Single Joint System. *Biol Cybern* 66:463-477.
- [59] Shadmehr R, Mussa-Ivaldi FA, Bizzi E (1993) Postural Force Fields of the Human Arm and their Role in Generating Multi-joint Movements. *J Neurosci* 13: 45-62.
- [60] Shadmehr R (1993) Control of Equilibrium Position and Stiffness through Postural Modules. *J Mot Behav* 25: 228-241.
- [61] Shadmehr R (2010) From Equilibrium Point to Optimal Control. *Motor Control* 14(3): e25-e30.

- [62] Shadlen M, Newsome W (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4: 569-579.
- [63] Simon H (1956) Dynamic Programming Under Uncertainty with a Quadratic Criterion Function. *Econometrica* (24): 74-81.
- [64] Simpson T (1995) A Comparison of Six Methods to Estimate Thresholds from Psychometric Functions. *Behav Res Meth Ins C* 27: 459-469.
- [65] Softky W (1995) Simple Codes Versus Efficient Codes. *Curr Opin Neurobiol* 5: 239-247.
- [66] Spearman C (1908) The Method of 'Right and Wrong cases' ('Constant Stimuli') without Gauss's formulae. *Brit J Psychol* 2: 227-242.
- [67] Stein R, Roderich G, Jones K (2005) Neuronal Variability: Noise or Part of the Signal? *Nat Rev Neurosci* 6: 389-97.
- [68] Strasburger H (2001) Converting Between Measures of Slope of the Psychometric Function. *Percept Psychophys* 63: 1348-1355.
- [69] Tassinari H, Hudson T, Landy M (2006) Combining Priors and Noisy Visual Cues in a Rapid Pointing Task. *J Neurosci* 26(40): 10154-10163.
- [70] Thibos L (1989) Image Processing by the Human Eye. *P Soc Photo-Opt Ins* 1-7.
- [71] Thiel H (1957) A Note on Certainty Equivalence in Dynamic Planning. *Econometrica* 25(2): 346-349.
- [72] Thiamin A, Schwarz M, Topper R, Fellows S (1991) Different Mechanisms Underlie the Long-Latency Stretch Reflex Response of Active Human Muscle at Different Joints. *J Physiol* 444: 631-643.
- [73] Thorpe S, Fize D, Marlot C (1996) Speed of Processing in Human Visual System. *Nature* 381(6582): 520-522.
- [74] Trommershauser J, Gepshtein S, Maloney L, Landy M, Banks M (2005) Optimal Compensation for Changes in Task-Relevant Movement Variability. *J Neurosci* 25(31): 7169-7178.
- [75] Trommershauser J, Maloney L, Landy M (2003a) Statistical Decision Theory and Trade-offs in the Control of Motor Response. *Spat Vis* 16(3-4): 255-275.
- [76] Trommershauser J, Maloney L, Landy M (2003b) Statistical Decision Theory and the selection of rapid, goal-directed movements. *J Opt Soc Am* 20(7): 1419-1433.
- [77] Todorov E (2004) Optimality Principles in Sensorimotor Control. *Nat Rev Neurosci* 7(9): 907-915.
- [78] Todorov E, Jordan M (2002) Optimal Feedback Control as a Theory of Motor Coordination. *Nat Neurosci* 5(11): 1226-1235.

- [79] Todorov E (2005) Stochastic Optimal Control and Estimation Methods Adapted to the Noise Characteristics of the Sensorimotor System. *Neural Comput* 17(5): 1084-1108.
- [80] Van Nes F, Bouman M (1967) Spatial Modulation Transfer in the Human Eye. *J Opt Soc Am* 57: 401-406.
- [81] Wei K, Kording K (2008) Relevance of Error: What Drives Motor Adaptation? *J Neurophysiol* 101: 655-664.
- [82] Wolpert D, Landy M (2012) Motor control is decision-making. *Curr Opin Neurobiol* 22:1-8.
- [83] Zhang H, Maddula S, Maloney L (2010) Planning routes across economic terrains: maximizing utility, following heuristics. *Front Psychol* 1(214): 1-10.