Carnegie Mellon University

Annual Progress Report: 2006 Formula Grant

Reporting Period

July 1, 2009 – June 30, 2010

Formula Grant Overview

The Carnegie Mellon University received $649,424 in formula funds for the grant award period January 1, 2007 through December 31, 2010. Accomplishments for the reporting period are described below.

Research Project 1: Project Title and Purpose

Predicting 3D Protein Structures - Automated structural prediction for tertiary and quaternary protein folding, and finding corresponding coding motifs in the DNA sequence are crucially important in computational molecular biology. This research will develop algorithms and software to identify the multiple peptide sequences in multiplexed tandem mass spectroscopy spectra through database searching and will develop approaches for comparing collections spectra representing different states of complex proteomes, including disease states. Our intent is to make sufficient progress that would permit us to seek larger funding from NIH, NSF, and perhaps DOE or the pharmaceutical industry.

Duration of Project

7/1/2007 – 11/30/2008

Summary of Research Completed

This project ended during a prior state fiscal year. For additional information, please refer to the Commonwealth Universal Research Enhancement C.U.R.E. Annual Reports on the Department's Tobacco Settlement/Act 77 web page at http://www.health.state.pa.us/cure.

Research Project 2: Project Title and Purpose

Computational and Neural Mechanisms of Object Vision - The purpose of this project is to test the idea that the visual cortex of the primate brain adjusts in response to experience so as to represent with maximal efficiency the things that have been seen. We will test this idea by using two kinds of images: natural scenes and faces. In each case, we will carry out parallel computational and electrophysiological studies. The computational studies will assess what features of natural scenes and faces carry the most information. The electrophysiological studies will test the idea that neurons in the visual cortex selectively represent these informational features.

Duration of Project

1/1/2007 – 12/31/2009

Project Overview

The general aim of this project is to understand how representations of scenes and faces can be learned from visual experience. The basic hypothesis is that features associated with the scenes and objects tend to be highly correlated, or, in other words, that they vary coherently in our visual experience. For example, the color or texture at different parts of an object's surface tends to be correlated. The movements of different body parts of a person during motion are also correlated. We conjecture that neurons in the brain are able to learn these correlations, and form connections and circuits to encode these correlations as statistical priors useful for perceptual inference. Neurons in the primary visual cortex are known to encode small features such as bars or spots. The statistical regularities in natural scenes can drive the formation of connections that link the small features together. These connections in turn will predispose the neurons to generate correlated activity when they again encounter higher order structures that are consistent with such correlations. This correlated activity can signal higher order structure such as larger fragments of objects or surfaces, and form a powerful and effective drive that leads to the development of tuning properties for specific higher order structures at the individual neurons downstream. This process can be applied successively in visual areas that form a processing chain extending from V1 to V2 to V4 and finally to IT, to create a hierarchy of representations linking features, parts, shapes, and objects together. While similar proposals have been advanced before, there have been no attempts to test them in concrete computationally inspired empirical studies. Nor has neurophysiological evidence been put forward to support the notion that such processes occur in the brain. This project aims to remedy these deficiencies. In these studies, we will (a) analyze the statistics of visual scenes and faces and (b) record from neurons in the visual cortex. Carrying out parallel experiments based on the two approaches will allow us to test the general idea that visual cortical neurons are sensitive to those features of scenes and faces that, in a strict statistical sense, carry the most information about them.

Principal Investigator

Carl R. Olson, PhD
Center for the Neural Basis of Cognition
Mellon Institute 115
4400 Fifth Avenue
Pittsburgh, PA 15213

Other Participating Researchers

Michael S. Lewicki, PhD, Tai Sing Lee, PhD - employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

These experiments will cast light on how neurons in the primate visual system represent natural

scenes and faces. The information will be directly relevant to understanding several debilitating disorders of vision from which patients with congenital or traumatic brain pathology suffer, including topographic amnesia (an inability to recognize formerly familiar places) and prosopagnosia (an inability to recognize formerly familiar faces).

Advances in theories of visual representation will provide new methods of coding and processing visual information. In addition, progress in these areas will lead to the design of better algorithms for computer vision which may one day allow computers and robots to interact with the natural sensory environment the way we do. Finally, by better understanding the properties of visual representations and the computational principles of the visual system, we will gain deeper insight into the neural basis of visual perception.

Summary of Research Completed

The inferotemporal cortex (ITC), a region of the temporal lobe in which neurons respond selectively to natural images, plays a critical role in visual object recognition in primates including humans and monkeys. Neuronal visual responses in ITC are modifiable by visual experience. In particular, if an image has been rendered familiar by repeated viewing, neurons respond to it with reduced overall strength but with enhanced differential strength leading to improved selectivity. These effects concern responses to images presented in isolation. However, under natural viewing conditions the visual system is bombarded by a continuous stream of ever-changing images.

In work conducted during the preceding year, the Olson lab has shown that image familiarization affects not only the strength of the response to an isolated image but also the dynamics of responses to rapidly changing displays. When images alternate at a rate of around 8 Hz, ITC neurons respond to each new item with a strong burst if the images are familiar but not if they are novel. In contrast to previously documented effects, which involve the modification of response strength and thus could arise from alteration of synaptic gain in bottom-up pathways, this sharpening of response dynamics must depend on a mechanism affecting response adaptation at the level either of the individual neuron or of the intrinsic circuit within which it is lodged. We have obtained this striking new finding in one monkey, have replicated it in a second monkey and are now writing up the results for publication.

We began by giving two monkeys repeated exposure to 16 digitized images, using a unique set of images for each animal. During a training session, all 16 images were presented on interleaved trials. On each trial, while the monkey looked at the center of the screen, an image appeared at fixation for a second. At the end of the training period, each monkey had viewed each image 90 times. Following training, we began to collect data from single ITC neurons in microelectrode recording sessions. To test each neuron, we arbitrarily selected two images from the training set and two images never before seen by the monkey. On each trial, a pair of images (either both novel or both familiar) appeared in back-to-back alternation with a switch every 120 ms until each had appeared three times (Fig. 1).

The neuronal responses to streams of novel and familiar images were strikingly different as demonstrated by the plot of population activity in Fig. 1. The population of neurons responded

with a strong burst to the onset of each successive familiar item (red curve). In contrast, they gave a phasic response only to the first item in a novel string blue curve).

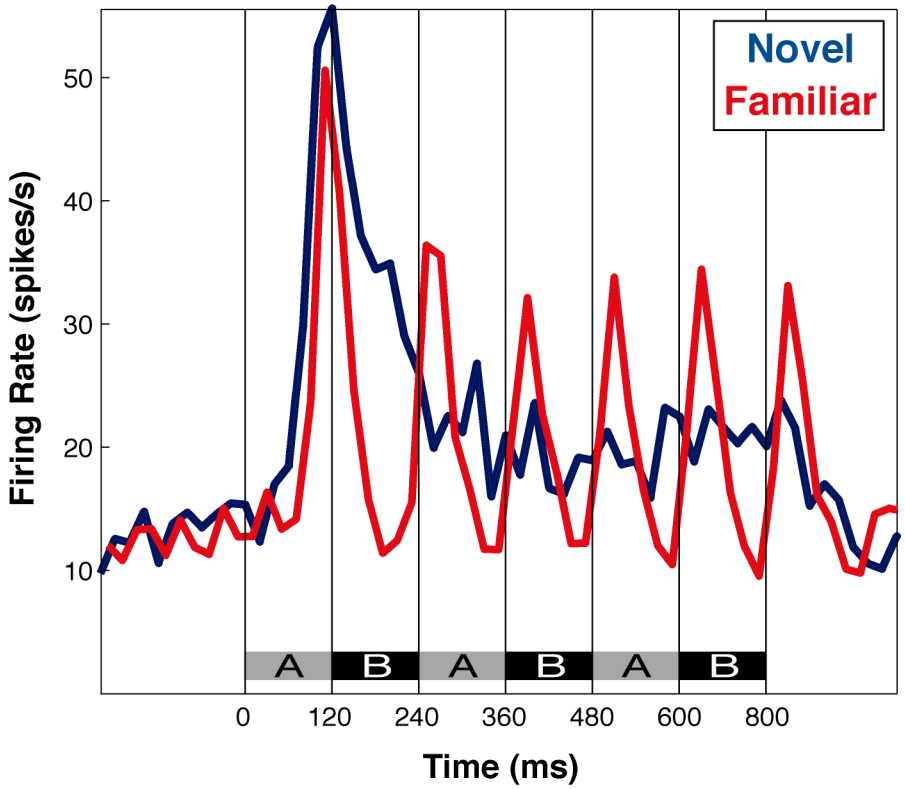

Figure 1. Population firing rate as a function of time during the trial for all neurons recorded in monkey 1. Two digitized images (labeled A and B) were presented in back-to-back alternation for 120 ms each. The neurons responded with a sharp burst of activity to each new onset if the images were familiar (red curve) but not if they were novel (blue curve). Thus, making images familiar radically alters response dynamics in ITC. Previous studies have demonstrated effects of learning only on the strength and selectivity of neuronal responses. The finding that response dynamics are altered by experience is entirely novel.

Research Infrastructure Project 3: Project Title and Purpose

Research Infrastructure Support for the Center for the Neural Basis of Cognition - This research infrastructure project is for renovation of 1,000 sq. ft. of space in the Mellon Institute (Room 134B) adjacent to the current quarters of the Center for the Neural Basis of Cognition (CNBC). This space will be converted to offices and laboratories that are badly needed to meet the needs of currently cramped CNBC faculty, postdoctoral researchers and students.

Anticipated Duration of Project

1/1/2007 – 12/31/2010

Project Overview

The broad objectives of the CNBC are to: 1) work jointly with allied departments at each

university to recruit and support the activities of outstanding researchers working on the neural basis of cognition, including development and disorders; 2) train outstanding graduate students and post-doctoral fellows so that they can carry on the research of the center in their future careers; and 3) provide mechanisms for fostering interactions and exchange of ideas among faculty, post-doctoral fellows, and graduate students.

This research infrastructure project is aimed primarily at supporting the latter two of CNBC's objectives. Successful completion will add 1000 square feet of space that will be used by students and post-docs of three of the center's investigators, fostering collaborations through the natural interactions that occur by co-location. The work those personnel will perform in the renovated space will be primarily computational data analysis, i.e., dry lab work, though a portion of the space will be configured to accommodate wet lab experimentation. The renovation plan has flexibility of space purpose built in. As such, should the opportunity and need arise, we can quickly re-purpose the space for use by the new Carnegie Mellon CNBC Co-Director, if he/she is hired (an active search is in progress) in a timeframe coincident with this project's schedule.

Rough plans and cost estimates for the renovation are in hand. Tasks to be performed in this project include letting the project out for bid, reviewing contractors' detailed proposals and price estimates, selecting a contractor and negotiating the contract, review and approval of final plans and drawings, removal of the existing ventilation system, construction of walls, installation of service lines, painting and flooring, certification of completion by the Carnegie Mellon Office of Campus Design and Facility Development, and purchase and installation of laboratory and office furniture.

Principal Investigator

Carl R. Olson, PhD
Center for the Neural Basis of Cognition
Mellon Institute 115
4400 Fifth Avenue
Pittsburgh, PA 15213

Other Participating Researchers

None

Expected Research Outcomes and Benefits

CBNC's facilities on the Carnegie Mellon campus are located on the first floor of the historic Mellon Institute building. Because of its success, CNBC has become cramped for space; laboratory space is especially at a premium. Upon completion of this project, the CNBC will have additional laboratory space to accommodate as many as 10-12 student and post-doctoral researchers. The planned subdivision of the 1000 square foot space, currently an outdated glass blowing shop, is for one 200 sq. ft. wet/dry lab and four 150-160 sq. ft. computational labs.

The new space will facilitate the work of three CNBC investigators based at Carnegie Mellon; it should also support occasional use by two CNBC investigators based at the University of Pittsburgh. An important benefit of the planned layout is that it will co-locate researchers, thus promoting interactions among research teams that inevitably lead to new collaborative research projects.

Summary of Research Completed

Renovation of the former glassblower space is progressing slowly. Additional time was required for the demolition of the space and some asbestos removal. Intermittent delays have occurred due to construction of the Vivarium directly above this space.

Research Project 4: Project Title and Purpose

Risk and Addiction in Smoking - The purpose of this project is (1) to assess how perceptions of addiction moderate adolescents' reported willingness to try tobacco, and (2) to explore the relationship between their perceptions of the risk of smoking and risk of addiction. This project will draw on work showing that people over-estimate the risks of smoking, as they do with most other risky health behaviors, as well as work showing that people see their own risk as lower than other people's risk. We plan to illuminate this apparent contradiction by showing the role of *perceptions of addiction risk* in determining smoking behavior, through their effect on the perceived relevance of smoking risks. We believe that making addiction risk real to adolescents can enhance the effect of anti-smoking education. However, doing so requires conveying the phenomenology of an inherently unfamiliar state.

Duration of Project

7/1/2007 – 6/30/2010

Project Overview

In this project we will conduct two studies, grounded in the "mental models" approach to investigating and improving people's understanding of risk. The first study will involve formative interviews to gain insight into adolescents' understanding of the forces at work in determining smoking risk, covering their beliefs regarding how people become addicted to substances, how that addiction drives future behavior, how patterns of behavior develop and are driven by social and physiological factors, how treatment of nicotine addiction works, how other treatments aimed at quitting work (or fail), and, finally, the role of visceral drives and emotions in smoking behaviors. To better understand this last factor, we will recruit adolescents who report having never smoked and adolescents reporting being regular smokers, and will randomize the smokers, into two groups, one of which agrees to refrain from smoking for a period of time prior to the interview, while the remainder make no such promise. The contrast will provide exploratory insight into the impact of visceral deprivation and drive states on these perceptions.

The second study will use structured questions to follow up on insights generated from the

formative interviews to explore how common different perceptions are and how they are interrelated, in larger groups than is possible with labor-intensive one-on-one interviews. We will recruit participants from three groups: current smokers, former smokers and never smokers. They will be randomized into two groups: one designated as "drive state" and a control group. Those in the drive state will be asked to refrain from smoking for a period prior to their survey (for half the smokers in the drive state group), or to refrain from drinking water for a shorter period prior to their survey (for non-smokers, former smokers, and for the other half of smokers in the drive-state group). This will allow us to examine effects of the visceral drive states, either normal (water) or addictive (smoking), on risk and benefit judgments, with manipulation checks used to evaluate their intensity.

The content of the questions posed in each study will be defined by an "expert model," summarizing critical information from the scientific literature, supplemented by additional beliefs arising in the formative interviews. The structured questions will take advantage of the terms used by adolescents in the formative interviews, so as to minimize misunderstanding of teens' beliefs, due to unfamiliar formulations.

Principal Investigator

Baruch Fischhoff, PhD
Carnegie Mellon University
Porter Hall 219-E
5000 Forbes Avenue
Pittsburgh, PA 15213-3890

Other Participating Researchers

Julie S. Downs, PhD, Wändi Bruine de Bruin, PhD - employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

The goal of this project is to understand better a puzzle that has been much debated in recent years: what role risk perceptions play in decisions to smoke. We believe that apparently inconsistent results are partially due to the incompletely examined role of perceptions of addiction risk. Unless that risk is plausible, the often-taught (and well-learned) risks of smoking will seem irrelevant – because they pertain to other people, those who continue smoking. In the interviews, we anticipate finding that adolescents do not widely consider addiction when contemplating the decision to smoke. Rather, we expect that they will focus on social factors that make smoking attractive and which, they believe, are under their control, making it possible to quit whenever they really want to. Indeed, we anticipate finding that many adolescent smokers will have plans to quit smoking once they reach adulthood, plans that do not take physiological addiction into account.

In the surveys, we anticipate finding that people underestimate the power of addiction, especially when they are in a satiated state. Thus, those who have been asked to abstain from smoking or drinking water are expected to rate higher perceptions of the power of addiction, and higher

benefits of engaging in the forbidden behavior. We also expect that people who have had less experience trying to quit addictive habits will have less insight into the power of addiction. Thus, length of time smoking and number of prior quit attempts (whether successful or not) should predict higher estimates of the power of addiction. The practical goal of this understanding is making the reality of addictiveness meaningful to adolescents, so that they can better manage these situations. Achieving it requires conveying the facts of addictiveness in a way that is scientifically sound (so that teens feel that they are being told the truth) and intuitively meaningful (so that it complements their existing mental models, and affords them active mastery). Our approach is designed to provide the basic scientific knowledge to accomplish both.

Summary of Research Completed

We conducted two studies in addition to Study 1 reported last year.

Study 2: Gut Feelings and Risk Perception

In this study we explored when and why people rely on intuitions or "gut feelings" in judgments and decisions versus relying on a more rational analysis, as a possible mechanism for the varying role of risk perceptions in decisions.

Post hoc evaluation of decisions: Participants responded to several trivia questions. The questions were designed such that participants would be unlikely to be able to answer with much confidence. For example:

In a recent survey, a representative sample of 5,000 Americans was asked, "Who do you think was the greatest American president?" Of the following, which president do you think was chosen most often?

A. Abraham Lincoln
B. Franklin D. Roosevelt
C. Ronald Reagan
D. John F. Kennedy
E. George Washington

Immediately after answering each question, the participant was randomly assigned to receive feedback indicating that his or her chosen answer was correct or incorrect. Next, after receiving feedback, participants indicated how they settled on their answer by answering the following questions:

How did you reach your answer on the question a moment ago?

A. I guessed
B. I had a gut feeling or intuition
C. I knew the answer
D. Other (please specify):

On the question a moment ago, to what extent did you have a gut feeling or intuition about the correct choice?

1 ----------2----------3----------4----------5
No feeling *Definite feeling*
or intuition *or intuition*

Did you answer according to your feeling/intuition?

Yes
No
Had no feeling or intuition

Confidence and post hoc evaluation of decisions: One possible origin of false confidence is the tendency to be dismissive of one's errors ("I was just guessing") but to take credit for one's correct judgments ("I knew it all along"). We collected data on participants' confidence in their responses to trivia questions like those above in order to test whether increased confidence is related to an increased propensity for outcome-contingent beliefs about decision processes. In other words, we expected that people who are more confident in their judgments about similar trivia questions would be more likely to say they knew or felt the answer after learning that they judged correctly, and more likely to say they were simply guessing after learning that they judged incorrectly. Confidence in answers were assessed with the question:

On the question you just answered, how confident are you that you chose correctly?

1 ----------2----------3----------4----------5
Not at all *Extremely*
Confident *Confident*

Availability and immediacy of outcome feedback: People may be more likely to rely on "gut feelings" in situations where unambiguous outcome feedback is unavailable – in other words, in situations where one's "gut feelings" are unlikely to be proven incorrect. This may be a second, related source of people's beliefs in the validity or diagnosticity of their gut feelings. People were first presented with a trivia question similar to those above, for example:

What state in the U.S. has the smallest population?

A. North Dakota
B. Alaska
C. Montana
D. Wyoming
E. Rhode Island

Before allowing participants to answer, we informed them that they either would (condition 1) or would not (condition 2) learn the correct answer to the question immediately after responding. Next, participants answered the questions below:

On the question you just saw, to what extent do you have a gut feeling or intuition about the correct choice?

1 ----------2----------3----------4----------5
No feeling *Definite feeling*
or intuition *or intuition*

On the question you just saw, how confident are you that you will choose correctly?

1 ----------2----------3----------4----------5
Not at all *Extremely*
Confident *Confident*

Finally, participants answered the trivia question and either receive feedback (condition 1) or not (condition 2). We hypothesized that participants would be less likely to claim to have an intuition (and will be less confident) when they knew that outcome feedback would be immediately provided.

We conducted pilot studies to develop the materials to test the above hypotheses, and to generate approximate effect sizes in order to power full studies. Unfortunately, many of our anticipated effects did not materialize, and so after thorough pilot testing we abandoned this line of inquiry. As described below, we turned to a more methodological approach to the question of risk perception, rather than the exploration of hypothesized effects that we had initially planned but that was not panning out empirically.

Study 3: A Contextualized Risk Scale for Measuring Perceived Lethality of Risks

The valid measurement of risk beliefs is central to effective risk communication. A reliable elicitation procedure is a prerequisite for judging the quality of people's risk beliefs and for evaluating the effectiveness of interventions aimed at changing those beliefs. At worst, a poor measure of risk beliefs can lead people to draw incorrect conclusions about issues that concern human lives or have large monetary consequences. The research proposed here involves the development, validation, and use of a new method for eliciting risk perceptions – specifically, perceptions of the lethality of various activities and conditions. The procedure, a contextualized lethality scale, is designed to elicit beliefs in a way similar to how people naturally think about risks, and it results in an estimate that can be evaluated objectively. We tested the validity of this scale alongside other commonly used measures of risk beliefs.

Development of the contextualized lethality scale

We developed a scale consisting of a list of hazards and health conditions, ranked according to their lethality, which is used to provide a framework within which respondents can express their beliefs about a target risk. Rather than seeking to elicit risk beliefs that encompass all dimensions of risk (e.g., controllability, voluntariness), we restrict each comparison to a single metric – the number of years of life one could expect to lose contingent on engaging in the activity or having the condition. A secondary format we explore is a series of paired comparisons in which respondents judge the target risk against each of the comparison risks individually. We believe that eliciting risk beliefs in this way (i.e., relative to other risks of varying lethality) has several advantages over existing measures.

By eliciting perceptions of the target risk via comparisons with other hazards and conditions, we seek a measure of risk lethality which will avoid problems with focusing effects, numeracy and unrealistic optimism. Further, if people indeed have a kind of "internal 'scale' of relative risk", then asking respondents to describe a target risk in relation to other risks should tap into the beliefs they already have articulated, thus producing a more accurate assessment and reducing the need for respondents to construct or infer responses on the spot.

Definition of lethality. Departing from much previous work examining perceptions of society level risks, we operationalize lethality as the expected years of life lost (EYLL) contingent on engaging in an activity or having a condition. In contrast, prior research has often conflated lethality with the base rates of risk related activities and conditions. Studies have asked, for example, about the number of people killed annually by influenza in the US without any reference to how many people contract influenza annually. Importantly, society-wide estimates that neglect the prevalence of risky behaviors will make extremely dangerous but low frequency activities (e.g., using heroin) appear less risky than high frequency activities with a much lower expected loss of life (e.g., driving a car). Moreover, the tendency for people to adopt risk prevention and risk mitigating behaviors for the most dangerous activities (e.g., using hard drugs) means that, on a societal level, highly dangerous activities also tend to be low frequency. By defining lethality as the number of years of life one would expect to lose contingent on performing an activity or having a condition, we seek a metric that equates rather than disambiguates the concepts of lethality and dangerousness.

Selection of risks for the construction of the contextualized scale. In the past, researchers have compiled risk catalogs and lists primarily under the presumption that cross-hazard comparisons should inform societal debates and help prioritize our collective management of risk. In constructing such tables and comparisons, researchers have often sought out disparate hazards with similar rates of lethality, with the aim of influencing the perceived acceptability of one or more of the hazards.

Contrasting sharply with such approaches, in developing our contextualized risk scale, we sought out reasonably well known activities and conditions, and our goal was to construct a scale with understandable gradients of lethality on which respondents could express their beliefs about other hazards. Several constraints guided our selection of health conditions. The first prerequisite was the availability of data on expected years of life lost (EYLL), either from studies in the epidemiological literature or statistics from related organizations (e.g., Cystic Fibrosis Foundation). Second, we chose health risks and conditions that spanned a range from highly lethal to somewhat trivial, in a way analogous to previous research in which events have been selected to span a wide range of event probability. Selected conditions and behaviors include cystic fibrosis, acute lymphocytic leukemia, becoming addicted to heroin, type 2 diabetes, HIV with early diagnosis and treatment, smoking several packs of cigarettes a day, high alcohol consumption, smoking a pack of cigarettes a day, being obese, working as a coal miner, idiopathic epilepsy, being overweight, and having asthma. Table 2 displays the conditions (ranked from highest to lowest EYLL) and brief descriptions that are part of the scale.

Validation of the contextualized lethality scale
Our elicitation procedure was validated by examining (1) respondents' subjective impressions of the scale on dimensions such as ease of use and how well it captures their beliefs and (2) correlations between the accuracy of judgments and self-reported knowledge of hazards. Performance of the contextualized lethality scale was compared with a "pair wise comparison" version of the scale in which respondents judge the target risks compared to each hazard on the scale individually (rather than simply placing the target risk into the ranked list). Both of these formats were also compared with a set of commonly used elicitation procedures, including: a probability measure of the chance of dying from having the condition (e.g., "If you have cystic fibrosis, what is the probability you will die from cystic fibrosis?"), a frequency measure of lethality ("Out of 100 people with cystic fibrosis, how many of them will die from cystic fibrosis?"), a linear scale with tick-marks on which subjects can rate the probability of dying from engaging in the activity, and an absolute (fill-in-the-blank) measure of expected years of life lost (EYLL) contingent on having the condition or engaging in the behavior.

Response modes were varied between-subjects. The target risks that respondents evaluated were activities and conditions drawn from the contextualized risk scale itself. These target risks include moderate smoking, heavy smoking, cystic fibrosis, being overweight, and type 2 diabetes. For the group using the contextualized risk scale, these items were removed and respondents were asked to decide (sequentially) the appropriate placement of each of the target risks. Respondents in the other five groups will evaluate the lethality of these same five risks in isolation.

In all conditions, after making their lethality judgments, respondents then evaluated the response mode on its ease of use and the accuracy with which they believe it captures their beliefs about the lethality of the target risks. These responses were made on verbal scales. Next, respondents assessed their level of familiarity and personal experience with each target risk.
As our primary outcome measure, we measured particular knowledge about the key risk, focusing on aspects that would be unrelated to perceived magnitude of risk. For example, we included knowledge questions asking participants to identify complications of diabetes from a list that contained 50% actual complications and 50% distractors. Participants were told that 50% of the items in the list were actual complications, thus discounting effects of identifying larger numbers of distractors as a measure of knowledge that might reasonably correlate with perception of diabetes as somehow worse and leading to greater expected years of life lost. As secondary outcome measures, we asked participants to rate how easy the scale was to use (for whichever scale they had seen), and asked how confident they were that their answers reflected their true perceptions of the risk.

We recruited 183 participants into this study. The key hypothesis was confirmed, in that ratings on the contextualized lethality scale were significantly correlated with knowledge of the risk ($r=.30$, $p<.05$), whereas ratings in the pairwise comparison scale and the fill-in-the-blank scale were uncorrelated with knowledge (r's range from -.02 to +.03). The contextualized lethality scale was not, however, rated easier to use than the other scales.