

ELICITING AND DETECTING AFFECT IN COVERT AND ETHICALLY SENSITIVE SITUATIONS

by

Philip Charles Davis


B.S., Physics
Brown University, 2000

B.A., Mathematics
Brown University, 2000

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Media Arts and Sciences
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2005



© 2005 Massachusetts Institute of Technology. All Rights reserved

Author

Program in Media Arts & Sciences
May 6, 2005

Certified by

Dr. Rosalind W. Picard
Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences
Thesis Supervisor

Accepted by

Dr. Andrew P. Lippman
Chair, Departmental Committee on Graduate Students
Program in Media Arts & Sciences

ROTCH

ELICITING AND DETECTING AFFECT IN COVERT AND ETHICALLY SENSITIVE SITUATIONS

by

Philip Charles Davis

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
on May 6, 2005, in partial fulfillment of the requirements
for the degree of Master of Science

Abstract

There is growing interest in creating systems that can sense the affective state of a user for a variety of applications. As a result, a large number of studies have been conducted with the goals of eliciting specific affective states, measuring sensor data associated with those states, and building algorithms to predict the affective state of the user based on that sensor data. These studies have usually focused on recognizing relatively unambiguous emotions, such as anger, sadness, or happiness. These studies are also typically conducted with the subject's awareness that the sensors are recording data related to affect. This thesis describes two experiments that are designed to explore these two areas that these studies have not examined.

The first experiment in this thesis focuses on eliciting and detecting the feeling of guilt, an emotion that is much more subtle than the emotions that most studies have examined. In the context of a specific experiment, we show that guilt level during a self-reflection period is correlated with overall electrodermal activity above a baseline value, but not well correlated with the frequency of skin conductivity responses or two features of an electrocardiogram. Our results also imply that collecting physiological data during a self-reflection period may contain useful information for lie detection above and beyond physiological data collected during the actual lie.

The second experiment focuses on eliciting and detecting affect covertly with an ordinary computer mouse. The widespread use of the computer mouse makes this a particularly interesting area to explore. We were unable to find a relationship between mouse trajectory features and valence. We did find that both velocity and acceleration off of a baseline level were correlated with arousal level.

Thesis Supervisor: Rosalind W. Picard
Title: Associate Professor of Media Arts and Sciences

Eliciting and Detecting Affect in Covert and Ethically Sensitive Situations

by
Philip Charles Davis

Thesis Committee:

Advisor:

Rosalind Picard

Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences
Massachusetts Institute of Technology

Reader:

Dan Ariely

Luis Alvarez Renta Professor of Management Science
Sloan School of Management and Program in Media Arts and Sciences
Massachusetts Institute of Technology

Reader:

John Stern

Professor Emeritus of Psychology
Department of Psychology
Washington University in St. Louis

Acknowledgements

The past two years have been incredible and I owe thanks to many people. I've grown on so many different levels by interacting with all the interesting, passionate people at the Media Lab. Thank you to everyone for making the Lab so dynamic and exciting!

Roz, you have been everything I could have asked for in an advisor. You always encouraged me to explore my interests no matter what they were. Your guidance on everything from pattern recognition algorithms to human subjects problems was invaluable. Thank you for everything.

Thanks to my thesis readers, Dan Ariely and John Stern. Dan, thank you for inspiring me to conduct the first experiment in this thesis and for all of your interesting perspectives. John, thank you for sharing your expertise in lie detection and EKG and EDA data collection.

Thanks to everyone in the Affective Computing group! You have all been great friends. Carson, I've collaborated with you the most and have always appreciated your helpful advice on everything from python scripts to subject recruiting. We certainly have many funny memories of our experiences running subjects! Shani, thank you for being a great officemate and a great person to bounce ideas off. Ashish, thanks for sharing your machine learning expertise. Win, thanks for helping me examine my work from a broader perspective. Yuan and Hyungil, thanks for all your insights about IBI analysis.

Thanks to all my friends and colleagues who have indirectly influenced this work and my other projects through our discussions. Amon, Andrea, Anmol, Cory, Gauri, Hugo, Jay, Joanie, Karen, Mike, Nathan, Nina, Patrick, Philippe, Push, Ron, Sunil, and Vadim – you have all been of great help and I hope we continue to keep in touch after leaving the Lab.

Most importantly, thank you to my friends and family. Lisa, you have always been there for me, through good times and bad, and I couldn't have gotten to where I am today without you. Phil, thanks for forcing me to think about technology from a broader government/policy standpoint. Mom and dad, thank you for always being there for emotional support even when I am a pain. Steve, you're the smartest guy I know and a great brother – thanks for your support and all of your interesting insights about technology, policy, and business. Thanks to all of my other friends from high school, Brown, MIT, Lincoln Lab, and elsewhere – our fun times take my mind off of work and keep me sane!

Table of Contents

Abstract.....	3
Acknowledgements	7
1 Introduction to this Document	11
2 Eliciting and Detecting Feelings of Guilt	13
2.1 Introduction	13
2.2 Background and Relevant Work	14
2.2.1 Electrocardiogram and inter-beat intervals	14
2.2.2 Electrodermal Activity	15
2.3 Experimental Overview	15
2.4 Understanding Subject Cheating Behavior	18
2.5 Estimating Feelings of Guilt.....	29
2.5.1 Estimating the intensity of the feelings of guilt felt	30
2.5.2 Expected Intensity of Guilt.....	33
2.6 The Experimental Setup	34
2.7 Sensor Data Analysis.....	39
2.7.1 Data Integrity.....	39
2.7.2 Analysis	46
2.8 Conclusions	53
2.9 Future Work.....	54
3 Covert Mouse Behavior Analysis.....	57
3.1 Introduction	57
3.2 Relevant Work.....	58
3.3 The Experimental Setup	58
3.4 Manipulation Check	62
3.5 Analysis: Detecting Affective State	64
3.6 Conclusions	68
3.7 Future Work.....	69
4 References.....	70
5 Appendices.....	72
Appendix A: Recruiting Poster	73
Appendix B: Screen shots from registration form.....	74
Appendix C: Experimenter’s Instructions for Running Experiment.....	80
Appendix D: Screen Shots from Experiment.	82
Appendix E: Complications with Experiment.....	97
Appendix F: Matlab Code used to implement EM algorithm	98
Appendix G: Matlab Code used to extract skin conductivity responses (SCRs)	100
Appendix H: Inter-beat Interval (IBI) extraction code.....	102
Appendix I: Spectral analysis code	106
Appendix J: Consent Form.....	114
Appendix L: Screen shots from craigslist surveys	118
Appendix M: Quizzes.....	124

Appendix N: Baseline EDA and SCR Frequency 127

1 Introduction to this Document

There is growing interest in creating systems that can sense the affective state of a user for a variety of applications. This thesis is an attempt to push the boundaries of affect sensing by exploring two areas that have not been explored in depth to our knowledge. In one experiment, we attempt to create an experiment to elicit different levels of a very subtle emotion, the feeling of guilt, and relate physiological sensor data to these levels of guilt. In another experiment, we attempt to relate the affective state of a user to his/her interaction with a simple computer mouse. Both of these experiments also bring up significant ethical issues.

The first experiment began as a joint project with Carson Reynolds, another student at the MIT Media Lab who is interested in the ethics of affect sensing in adversarial situations. The project was originally an open-ended exploration of how physiological sensor data varies in these situations. Through several revisions, the project evolved into an analysis of whether reflecting on the ethics motivating a lie, which is related to feelings of guilt, would impact physiology. This work has applications in lie detection, but these applications were not a primary motivator of the original work.

The second experiment, in which we attempted to relate a user's computer mouse behavior to his/her affective state, was particularly interesting to us because it involves a sensor that is already attached to virtually every desktop computer in the world. We would not expect a user's computer mouse use to be highly indicative of his/her emotional state, but the ubiquity of the computer mouse makes this an interesting area to explore. However, the widespread use of the computer mouse also implies that sensing of a user's affective state with a mouse could be done covertly or without the user's knowledge. Few affect sensing studies have been done where a subject is not aware that affect sensing may be taking place, which raises questions about whether it is even possible to sense affect covertly. More importantly, it also raises ethical questions about whether this type of covert affect sensing should be done at all.

The high level structure of this thesis reflects the fact that it comprises two very different experiments that are largely independent of each other. Chapter 2 describes in its entirety the first such experiment, which is designed to elicit feelings of guilt and detect these feelings with physiological sensors. Similarly, Chapter 3 describes in its entirety the second experiment, which is designed to elicit several affective states and sense them covertly with a simple computer mouse.

2 Eliciting and Detecting Feelings of Guilt

This chapter describes an experiment designed to elicit feelings of guilt and detect these feelings using physiological sensors. After an initial introduction motivating the experiment and a description of some relevant work, the third section describes an overview of the experiment. The fourth section describes how we estimated “cheating” behavior during a game. The fifth section presents an analysis of how much guilt we believe we elicited in this experiment for each condition relative to the others. The sixth section describes the experimental setup in detail. The seventh section presents the sensor data analysis. The eighth section describes the conclusions that we were able to draw from this work. Finally, the ninth section suggests possible avenues for future work in this area.

2.1 Introduction

There has been growing interest in recognizing a wide variety of affective states for a large number of possible applications. The goal of the majority of relevant experiments has been to recognize common, relatively unambiguous states, such as happiness, sadness, and anger. No work that we are aware of has explicitly aimed to elicit and detect a state as subtle as the feeling of guilt. This work represents an initial attempt to do just that.

This work began out of an exploration of how physiology changes in adversarial situations and, in particular, when lying or cheating with different ethical incentives. We wondered if we could create conditions in which subjects would be dishonest, but when asked to reflect on their dishonesty, would feel very different about the ethics behind their actions. If so, could we build a tool to distinguish between these different feelings based on the physiology during the reflection stage? This tool would, in a sense, be a “guilt detector” if we could do it properly. Investigating the feasibility of building such a “guilt detector” became a major motivator for the experiment outlined in this chapter.

2.2 Background and Relevant Work

In this section, we discuss work that has previously been done in related areas. First, we introduce relevant work on the electrocardiogram and then we introduce relevant work on electrodermal activity.

2.2.1 Electrocardiogram and inter-beat intervals

An electrocardiogram (EKG) is simply a representation of the electrical activity of the heart. Inter-beat intervals are the lengths of time between successive “beats” of the heart and can be extracted from an EKG waveform. Researchers have shown that dividing the power spectrum obtained from the inter-beat intervals, extracted by thresholding the R-waves from a raw EKG waveform, into three frequency spectra is useful in determining the amount of sympathetic and parasympathetic activity. More specifically, using the notation and results of (Akselrod et. al., 1981), the low frequency (LF) band from 0.01 to 0.08 Hz is primarily considered a measure of sympathetic activity with a minor parasympathetic component, the high frequency (HF) band from 0.15 to 0.5 Hz is predominantly a measure of parasympathetic activity, and the mid-frequency (MF) band between these two bands is a measure of both sympathetic and parasympathetic behavior, but predominantly the latter. Computing the integral of the power spectrum within these bands and comparing the relative magnitudes of these bands has proven to be useful for a wide range of purposes, including distinguishing between emotional states. A study by (Sloan et. al., 1994) found that mental stress increases LF activity relative to both MF and HF activity. Researchers from the Heartmath Institute showed a similar result for anger (McCraty et. al., 1995).

Karen Liu created a mobile system to collect and annotate this type of heart rate variability data during normal activities (Liu, 2004). Hyungil Ahn is currently analyzing some of Liu’s collected data; a preliminary analysis of this data has shown no strong correlation between subjects’ self-reported stress and heart rate variability, although additional analysis may show otherwise.

2.2.2 Electrodermal Activity

Electrodermal activity (EDA) is a measure of the alteration of electrical resistance of the skin associated with sympathetic nerve discharge. A number of studies have shown EDA to be highly correlated with arousal (Ekman et. al., 1983). EDA is one of the physiological signals recorded in a polygraph test and has also been shown to be independently correlated with lying behavior (Kircher et. al, 1997). It has been used widely in a variety of research projects in the Affective Computing group and elsewhere at the MIT Media Lab.

2.3 Experimental Overview

At a high level, our goal in designing this experiment was to create several conditions in which subjects will experience different levels of guilt across conditions, on average, so that we can correlate physiological sensor data obtained in the experiment with the average levels of guilt that each group experiences. In order to do this, we need to understand the nature of guilt, propose a method to elicit feelings of guilt, and have a method for estimating whether those feelings were induced properly.

The working definition of guilt that we will use is: “remorseful awareness of having done something wrong” (<http://www.answers.com/topic/guilt>). By this definition, in order to have feelings of guilt, one must first make a value judgment about what is ethical and then do something that goes against what is ethical. Thus, to induce guilt, we need to create situations in which subjects will do something (an “action”) that is what they consider unethical.

Written below is the basic scenario we have created to induce feelings of guilt:

The situation: Subjects play a game in which they receive a score.

The action: Over-reporting his/her score (i.e. “cheating”).

The unethical ramifications: “Cheating” would be perceived by many to be unethical in itself. Moreover, higher scores decrease the amount of money that goes to a

charity of their choice. Hence, by cheating, they are effectively stealing from a charity that they have specified is good.

However, there are several difficulties with creating any experiment that induces feelings of guilt. Three major difficulties are listed below:

- 1) If a subject judges the ramifications of an action to be unethical, s/he is likely not to do it without a strong incentive to do so based on his/her own sense of morality.
- 2) If a subject believes an experimenter is observing his/her behavior, s/he may be even less likely to perform this action, independent of his/her own ethical values, simply to prevent the experimenter from observing this behavior.
- 3) It is difficult to estimate the amount of guilt a particular subject experienced since cognitive dissonance in the subject may prevent honest feedback in the form of a traditional survey given after a task.

To overcome the previously mentioned three difficulties with creating an experiment that induces feelings of guilt, we do the following in our experiment:

- 1) We motivate the subject to over-estimate his/her score by providing a financial reward that goes up as his/her score increases.
- 2) We create a situation in which there is no way for the experimenter to observe whether or not an individual subject cheated. However, by comparing the performance of subjects in this condition with those in a control condition in which subjects cannot cheat, we can estimate the percentage of subjects that cheated.
- 3) We collect survey data from an independent pool of subjects giving them the hypothetical scenario and asking how much guilt they would feel if they cheated. Assuming that people's expectations for the level of guilt they would feel is

reasonably accurate, we can combine this information with the approximate percentage of subjects who cheated (as mentioned above) to determine the average level of guilt among all subjects in a condition.

So far, we have described an experiment in which there is a condition designed to produce feelings of guilt and a control condition in which subjects cannot cheat and hence should not experience feelings of guilt. These are our two most important conditions and with these two conditions alone, we could conduct the experiment and attempt to detect guilt by analyzing collected sensor data during times when we expect subjects to feel guilt. From now on, we will label these two conditions as *charity_loses* and *control*, respectively. Note that subjects in the *control* condition are compensated the same way as those in the *charity_loses* condition, but there is no mention of a charity in the *control* condition.

However, we have also introduced two additional conditions that have different ethical ramifications of cheating from the *charity_loses* condition described above. The purpose of introducing these conditions is both due to uncertainty in the behavior of subjects in the above-mentioned condition (i.e. we didn't know a priori if they would cheat) and uncertainty in how much guilt the subjects would experience. The two other conditions have the following differences from the *charity_loses* condition:

First additional condition: A charity does not receive any compensation at all, which we expect would eliminate guilt associated with taking money from the charity by cheating. In this condition, there was no mention of a charity at all as in the *control* condition. From now on, we will label this condition as *no_charity*.

Second additional condition: A charity of the subject's choice receives extra compensation when the subject scores well, which we expect would actually make some subjects believe it is ethical to cheat since it would help this charity. From now on, we will label this condition as *charity_gains*.

For all conditions, we also need a point in the experiment in which subjects are expected to reflect on the ethics of what they have done so that we can determine if there is a physiological response associated with their relative feelings of guilt. For this reason, after all scores have been reported, we ask subjects to think about the ethics of what they have done and how it has affected both themselves and the charity.

It is worth noting at this point that we are not attempting to build yet another lie detector, although this work has applications in lie detection. We would like to attempt to elicit the feelings of guilt associated with lying/cheating, which should vary depending on the ethics of the underlying payoff scheme. That is, even if there were similar cheating/lying behavior in two groups, the level of guilt felt during the period in which subjects are asked to reflect on the ethics of what they have done may be different depending on the condition. This work is relevant to lie detection applications, however, in that physiological measures obtained during this period of reflection could potentially be used in addition to physiological measures recorded during the actual lie to obtain better detection accuracy overall. This additional data may be particularly useful in situations where someone does not have any strong feelings during a lie, but feels a strong sense of guilt upon reflecting about the lie.

The issues discussed in this section are the key points of the protocol for this experiment. In the next section, we estimate the average levels of guilt experienced in these conditions based on data we collected before going into detail on the experimental setup. The reason we estimate the levels of guilt in these conditions prior to describing the experiment in detail is to justify the rationale behind running the experiment at all since at this point the reader may not be convinced that the levels of guilt experienced in these groups will be significantly different.

2.4 Understanding Subject Cheating Behavior

This experiment is designed so that subjects can cheat without any way for the experimenter to verify whether or not they cheated in all but the *control* condition. It was necessary to design the experiment this way or else subjects would be afraid of getting caught and would

not cheat (based on discussions with subjects during pilot experiments). However, in order to estimate how guilty subjects are likely to feel in each condition, we needed to estimate how many subjects were likely to have cheated in each condition. This section describes how we estimated the number of cheaters in each condition.

Figure 1 shows the score distribution for all four conditions. The means for the distributions in the order in which they appear in the figure are 11.1, 12.1, 13.7, and 15.1. Their standard deviations in the order in which they appear in the figure are 5.4, 4.4, 5.4, and 4.5.

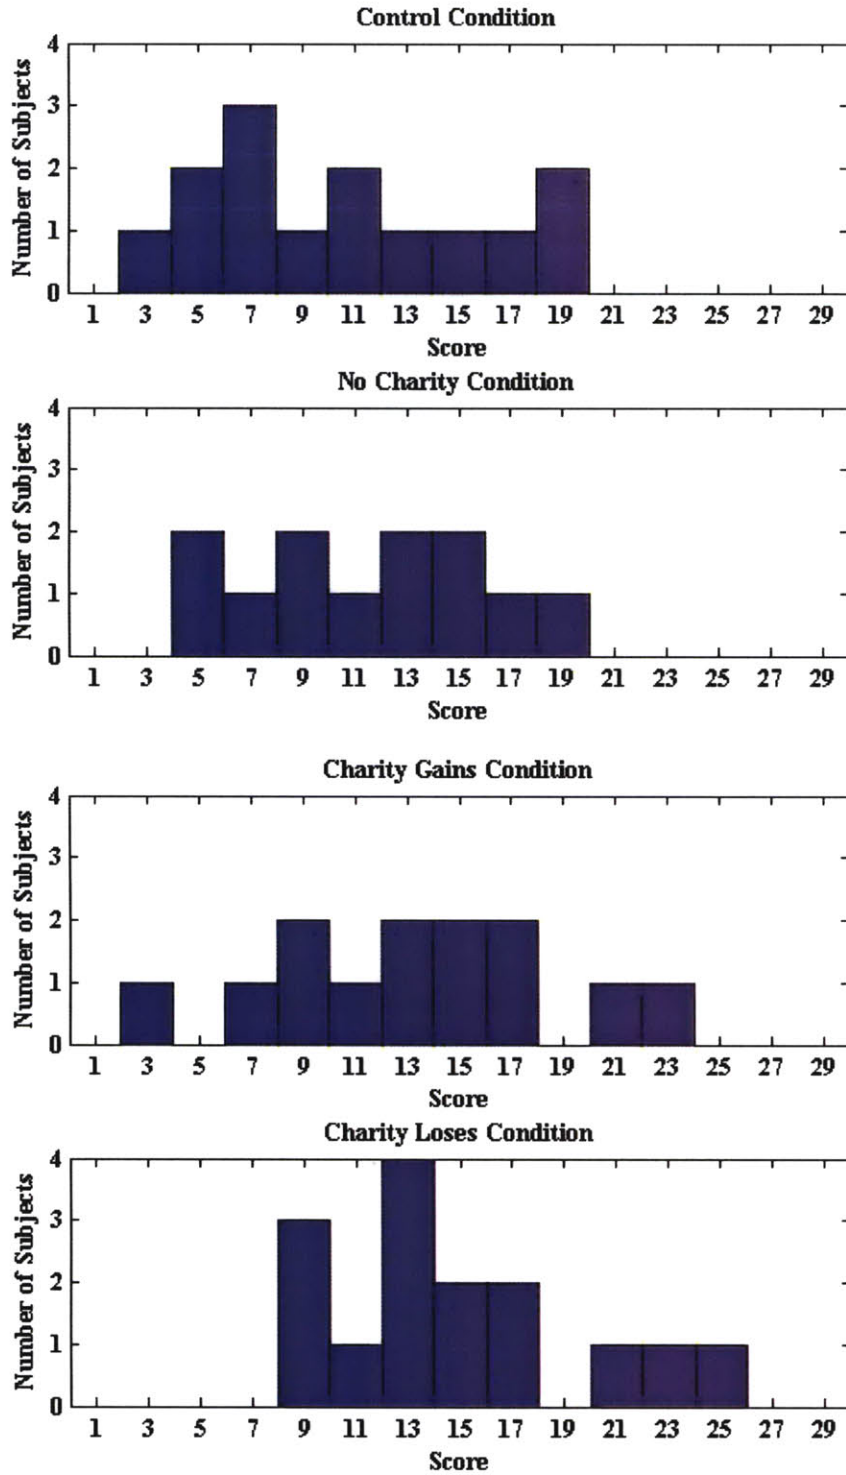


Figure 1: Score distributions for all four conditions. There were a total of 14 subjects in the *control* condition, 12 subjects in the *no_charity* condition, 13 subjects in the *charity_gains* condition and 17 subjects in the *charity_loses* condition.

The disparity in average scores across conditions could be due to cheating behavior or actual differences in the performance of subjects in these conditions. For example, it is possible that subjects in the *charity_loses* condition simply tried harder than subjects in the other conditions, resulting in a higher average score in that condition. Future work could determine if this is the case by introducing additional conditions identical to the *charity_gains* and *charity_loses* conditions except that subjects are unable to cheat. However, we believe it is highly unlikely that subjects would simply perform better in those conditions and will assume for the rest of this thesis that the disparity in scores between the different conditions is due to different amounts of cheating.

Since there is a lot of overlap in the score distributions among the groups and there are a relatively small number of data points within a group, there is no clear method that will unambiguously determine the correct number of cheaters within each group. A large amount of overlap is consistent with a large number of subjects cheating by a small amount, a small number of subjects cheating by a large amount, or a small number of subjects cheating by a small amount. Thus, every such method will rely on assumptions that cannot be verified. We used two methods to estimate the number of cheaters in each condition, which we believe correspond to conservative and high estimates for the true number of cheaters.

The first such method, which we believe produces a conservative estimate for the true number of cheaters, is the simplest: we fit scores in each condition to a normal distribution (normalized for scores between 0 and 30) and then estimate the extent to which the non-*control* distributions are skewed to the right of the *control* distribution by integrating between the two curves where the probability density function for the non-*control* distribution is greater than for the *control* distribution as showed in Figure 2. If cheaters do not over-report their scores by a large amount, which we believe to be the case, then this method will significantly underestimate the number of cheaters (i.e. if all subjects over-report their scores by 1 then this method will estimate that there are a small number of cheaters).

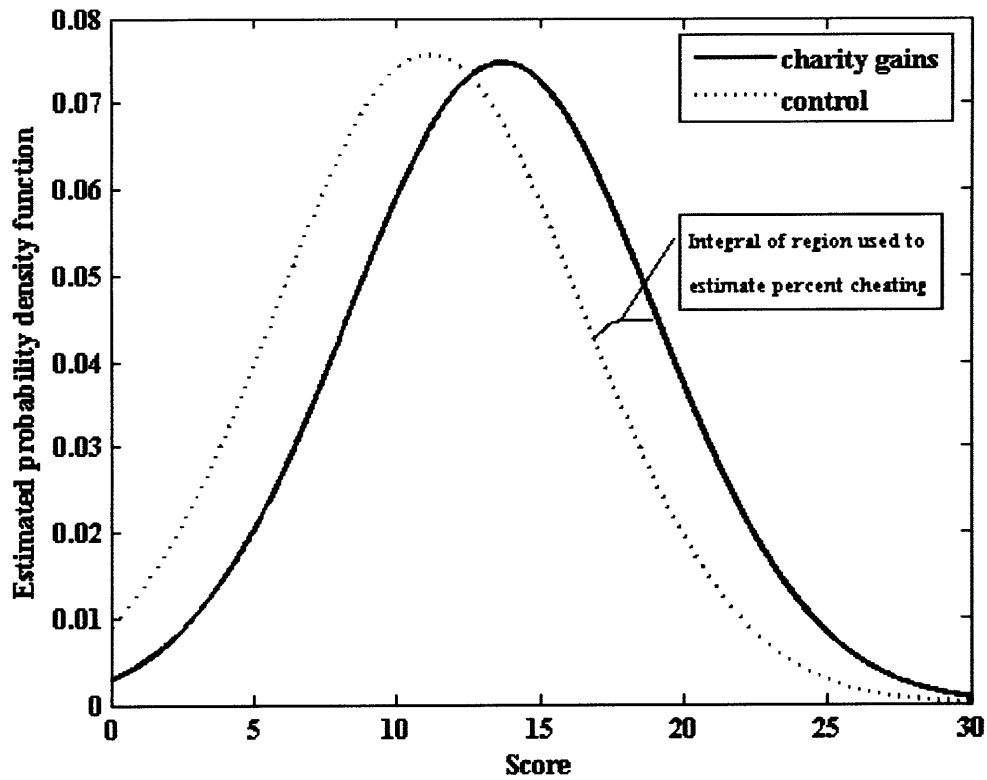


Figure 2: Example of how the number of cheaters are estimated in our first method for the *charity_gains* condition. This method is likely to underestimate the number of cheaters.

Figure 3 shows the percentage of subjects cheating as estimated using this method.




Figure 3: Estimates of the percentage of subjects cheating when estimated using method 1. We will treat these numbers as conservative estimates of the true percentage of subjects cheating within each group.

In our second method, which we believe produces high estimates for the true number of cheaters, we explicitly modeled the cheating behavior and then used the expectation maximization (EM) algorithm to estimate the model parameters. The EM algorithm has been used widely in the machine learning literature and is used to estimate parameters of an underlying probability distribution, which is assumed to produce one or more observable signals.

In our implementation, we wanted to make as few assumptions as possible about the actual nature of the cheating behavior since we did not have very many data points. The simplest reasonable model that we could come up with was one in which the distribution of scores for

cheaters was the same across conditions and the distribution of scores for non-cheaters was also the same across conditions. In this model, the only difference between the conditions was the percentage of subjects that cheated, which we included in our model with a variable, α . Our assumptions in this model were as follows:

- 1) If a subject does not cheat, his/her score is independent of the condition to which he or she is assigned.
- 2) If a subject cheats, his/her score is independent of the condition to which he or she is assigned (although, of course, cheating is not possible in the control condition).
- 3) The distribution for scores of non-cheaters is a normal distribution (i.i.d.). Similarly, the distribution of scores for cheaters is a normal distribution (i.i.d.).
- 4) The mean score for cheaters is at least as high as the mean score for non-cheaters.
- 5) The probability that a subject will cheat is dependent on the condition that the subject is in.

Stated differently, we used the EM algorithm to find the set of parameters in the equations written below that is most likely to have produced the observed score distributions in a maximum likelihood sense. The reader should interpret variables with the “NC” subscript as parameters for the “no cheat” distribution and variables with the “C” subscript as parameters for the “cheat” distribution.

$$f_{control}(x): \frac{1}{\sigma_{NC} \sqrt{2\pi}} e^{-\frac{(x-\mu_{NC})^2}{2\sigma_{NC}^2}}$$

$$f_{no_charity}(x): (1 - \alpha_{no_charity}) \frac{1}{\sigma_{NC} \sqrt{2\pi}} e^{-\frac{(x-\mu_{NC})^2}{2\sigma_{NC}^2}} + \alpha_{no_charity} \frac{1}{\sigma_C \sqrt{2\pi}} e^{-\frac{(x-\mu_C)^2}{2\sigma_C^2}}$$

$$f_{charity_gains}(x): (1 - \alpha_{charity_gains}) \frac{1}{\sigma_{NC} \sqrt{2\pi}} e^{-\frac{(x-\mu_{NC})^2}{2\sigma_{NC}^2}} + \alpha_{charity_gains} \frac{1}{\sigma_C \sqrt{2\pi}} e^{-\frac{(x-\mu_C)^2}{2\sigma_C^2}}$$

$$f_{charity_loses}(x): \left(1 - \alpha_{charity_loses}\right) \frac{1}{\sigma_{NC} \sqrt{2\pi}} e^{-\frac{(x-\mu_{NC})^2}{2\sigma_{NC}^2}} + \alpha_{charity_loses} \frac{1}{\sigma_C \sqrt{2\pi}} e^{-\frac{(x-\mu_C)^2}{2\sigma_C^2}}$$

where $\mu_C \geq \mu_{NC}; 0 \leq \{\alpha_{no_charity}, \alpha_{charity_gains}, \alpha_{charity_loses}\} \leq 1$

Equation 2.1: Probability distributions assumed in our model for each condition

With this model, our goal becomes to compute the maximum likelihood estimate of the means and standard deviations of these two distributions as well as the probability of cheating, α , in each condition. Figure 4 shows the probability of cheating by condition obtained by this method. The numbers reported in this figure should be treated as high estimates, rather than true estimates, of the actual cheating behavior since it is fairly clear that “overfitting” most likely occurred due to our small sample size. That is, the algorithm effectively fit the highest cheating case, *charity_loses*, to the pure “cheat” distribution. This is not a property of the model, but rather is the result of having too few data points.

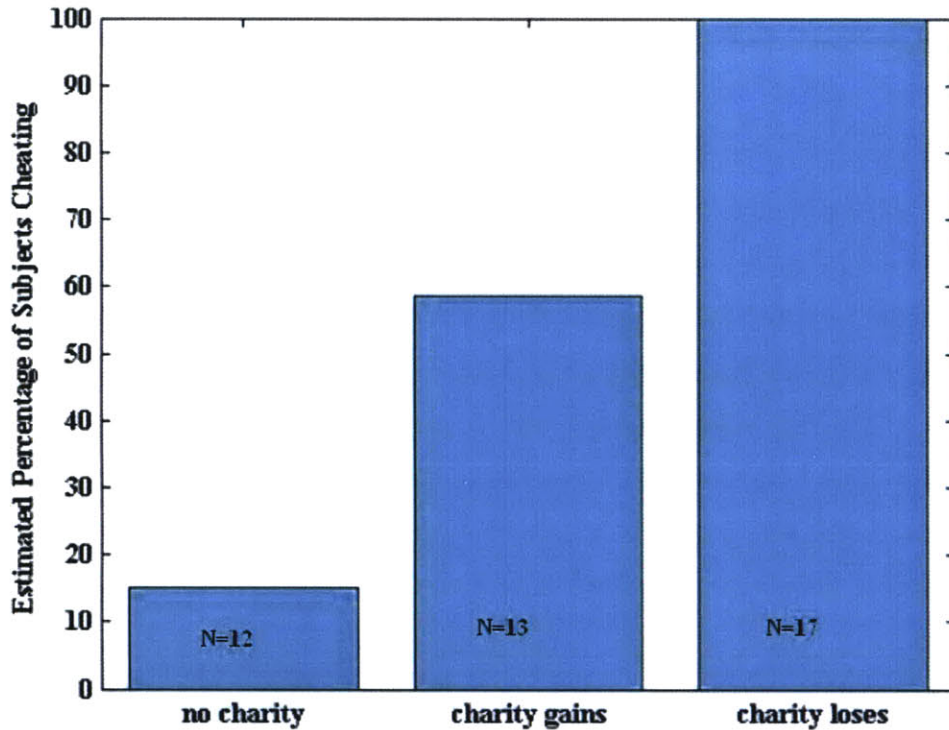


Figure 4: Estimates of the percentage of subjects cheating when estimated using method 2. We will treat these numbers as high estimates of the true percentage of subjects cheating within each group since this model is very likely to have over-estimated the percentage of subjects cheating, particularly in the *charity_loses* condition.

Prior to conducting this experiment, we expected that the *charity_loses* condition would have the least cheating, but instead it has the most cheating (by far) as estimated by both of our models. Although it was not an explicit goal of this thesis to understand the behavior (i.e. cheating patterns) of the subjects across conditions, we did investigate these patterns a bit due to the strange nature of these results.

First, we hypothesized that subjects in the *charity_loses* condition may simply not understand that they are hurting the charity by cheating. However, when we asked four subjects in this

condition after the experiment about whether they understood how the charity was compensated, all of them understood that they would be hurting the charity by cheating.

Next, we posted a survey on the Web (on craigslist, several different cities) to determine how people believe that they would behave in these conditions. We advertised under the “volunteer” section of the Web site and emphasized that the survey would only take a minute to complete since it was so short. When subjects came to the survey Web site, they were randomly assigned to one of the three non-*control* conditions. They were then given a brief description of the situations that subjects in their condition had experienced and were asked what the likelihood of them cheating would be under those circumstances. For the specific scenarios and questions that were shown to the subjects as well as an example of the message that we used to solicit subjects, see Appendix L.

When we plotted the average reported likelihood of cheating (on a scale of 1-8), we got a much more intuitive result, as shown in Figure 5.




Figure 5: Cheating survey results (with 95% confidence intervals) for 100 subjects in non-control conditions confirm what the author expected: that the *charity_loses* condition would have the lowest amount of cheating. However, these results are inconsistent with the empirically observed results that show cheating behavior in the *charity_loses* condition to be higher than in all of the other conditions.

We have not been able to solve the discrepancy between observed cheating behavior and the cheating behavior predicted by this survey. It is clear, however, that in this case, subjects' expected behavior differs from their true behavior.

We conclude this section with an estimate of the percentage of subjects cheating in each condition. Using our two methods, we have computed what we believe are conservative and high estimates for the true percentage of subjects cheating in each condition. We estimate the true percentage of subjects cheating by taking the midpoint between the conservative and high

estimates. When we do this calculation, our estimated percentage of cheaters by condition are: 13.7%, 39.1%, and 66.6% for the *no_charity*, *charity_gains*, and *charity_loses* conditions, respectively. These are the percentages that we will use as estimates for the remainder of this chapter.

2.5 Estimating Feelings of Guilt

In this section, we attempt to estimate the expected amount of guilt experienced by an individual in each of these groups relative to the other groups. The method we use is straightforward. First, we estimate the percentage of subjects who cheated in a given condition or the probability that an individual cheated in that condition as we have already done in the previous section. Second, we estimate the amount of guilt (on an arbitrary, linear scale) that an individual is expected to experience if s/he cheated and also if s/he did not cheat. We estimate these numbers using a survey posted on the Web. Last, we compute the expected amount of guilt experienced by an individual by substituting the above numbers into the following equation:

$$\begin{aligned}
 E(G_{no_charity}) &= \\
 &P_{no_charity}(cheat) \times G_{no_charity}(cheat) + P_{no_charity}(no_cheat) \times G_{no_charity}(no_cheat) \\
 E(G_{charity_gains}) &= \\
 &P_{charity_gains}(cheat) \times G_{charity_gains}(cheat) + P_{charity_gains}(no_cheat) \times G_{charity_gains}(no_cheat) \\
 E(G_{charity_loses}) &= \\
 &P_{charity_loses}(cheat) \times G_{charity_loses}(cheat) + P_{charity_loses}(no_cheat) \times G_{charity_loses}(no_cheat)
 \end{aligned}$$

Equation 2.2: Expected levels of guilt felt by condition

We have already estimated the percentage of subjects cheating per group, but we need to estimate the guilt felt while cheating and not cheating within each group. We do this in the next section.

2.5.1 Estimating the intensity of the feelings of guilt felt

We introduced a survey on craigslist to estimate the intensity of guilt felt (on a scale of 1-8) by a cheater and a non-cheater in each of the conditions. When subjects clicked on the link to the survey, our system assigned him/her to one of the three non-*control* groups and told the user of the associated scenario in that group. We used two surveys: one which asked the subject to assume that s/he did cheat and the other told the subject to assume that s/he did not cheat (see Appendix L).

The first survey instructed the subject to assume that s/he choose to cheat and asked how much guilt, on a scale of 1-8, the subject would experience from that. The results of this survey are shown in Figure 6.

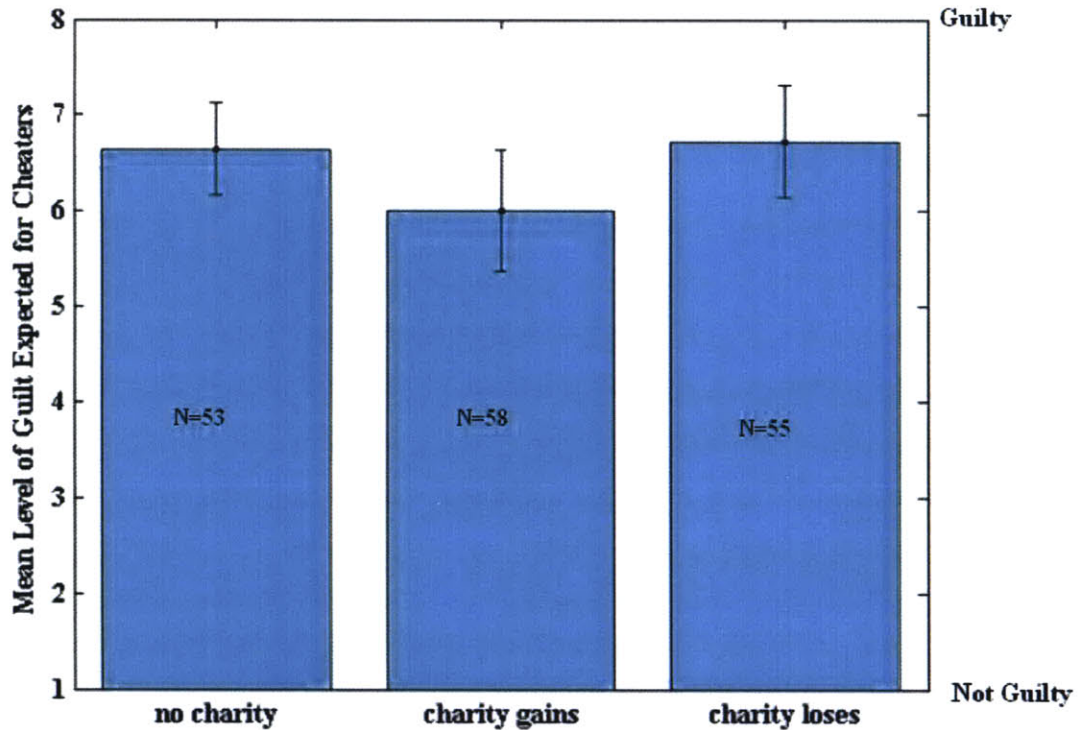


Figure 6: Guilt intensity survey results for cheating in non-control conditions. Although the relative ordering of the means is what the author expected, the error bars (95% confidence interval) are too large to draw any strong conclusions about the true relative levels of guilt felt when a subject cheats. Result implies that the probability of cheating in each condition is the dominant factor in determining the average level of guilt felt in a group.

Despite collecting survey data from 166 subjects, the 95% confidence intervals are still large enough as to prevent us from drawing any strong conclusions about the relative magnitudes of the expected guilt felt when a subject cheats. However, subjects are likely to report guilt levels more consistent with intuition (*charity_loses* condition highest and *charity_gains* condition lowest) if they are asked to rate guilt levels for all three conditions as opposed to only one condition as was the case in the above analysis. This assertion was also confirmed

by a preliminary analysis of data obtained using a survey with all three conditions displayed simultaneously. However, using the guilt levels reported above, the most important factor in determining the expected guilt felt in a condition is the probability of that subject cheating since all three guilt levels are similar.

We also estimated the guilt intensity felt when a subject does not cheat for each condition using the same type of survey. The results are shown in Figure 7.

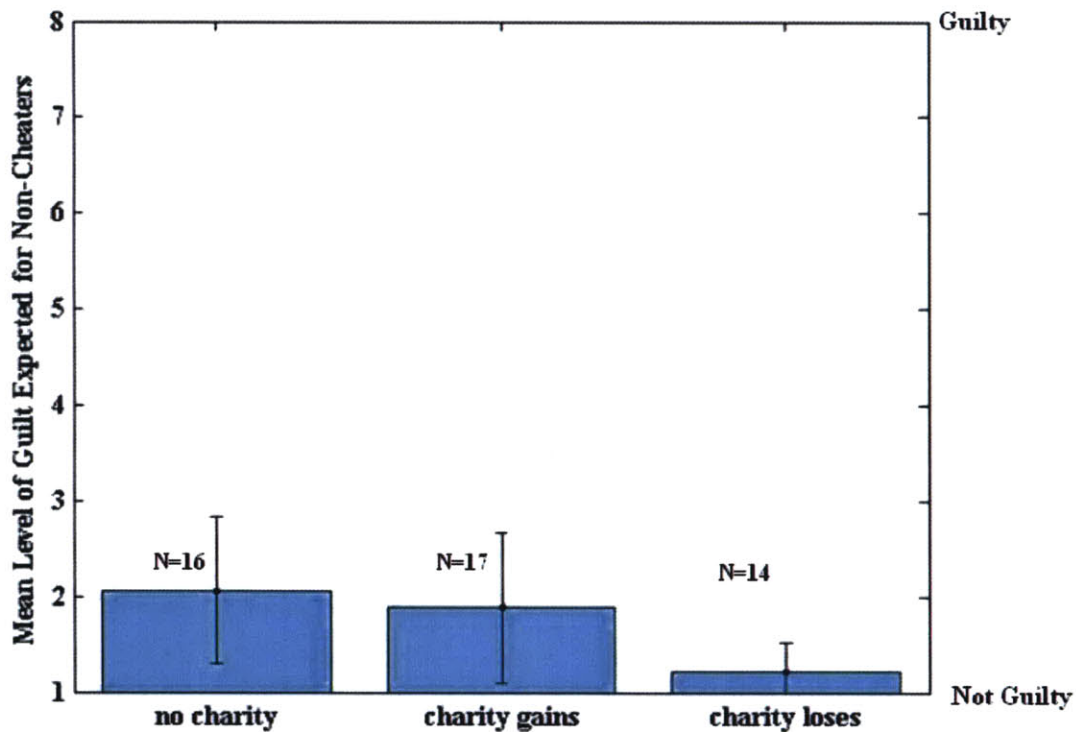


Figure 7: Guilt intensity survey results for no cheating in non-control conditions. We see that all conditions have low guilt when subjects are honest as we would expect.

Now, we will compute our final values: the expected intensity of guilt felt in a subject per condition.

2.5.2 Expected Intensity of Guilt

Now that we have estimated both the likelihood of a subject cheating in each condition and the average level of guilt felt when a subject cheats, we can compute the expected guilt felt per condition simply by multiplying these two quantities. To do this, we use Equation 2.2. When we substitute for the probabilities and guilt intensities that we estimated in our analysis, we get the result in Figure 8, where the three conditions amount to three distinct levels of guilt.

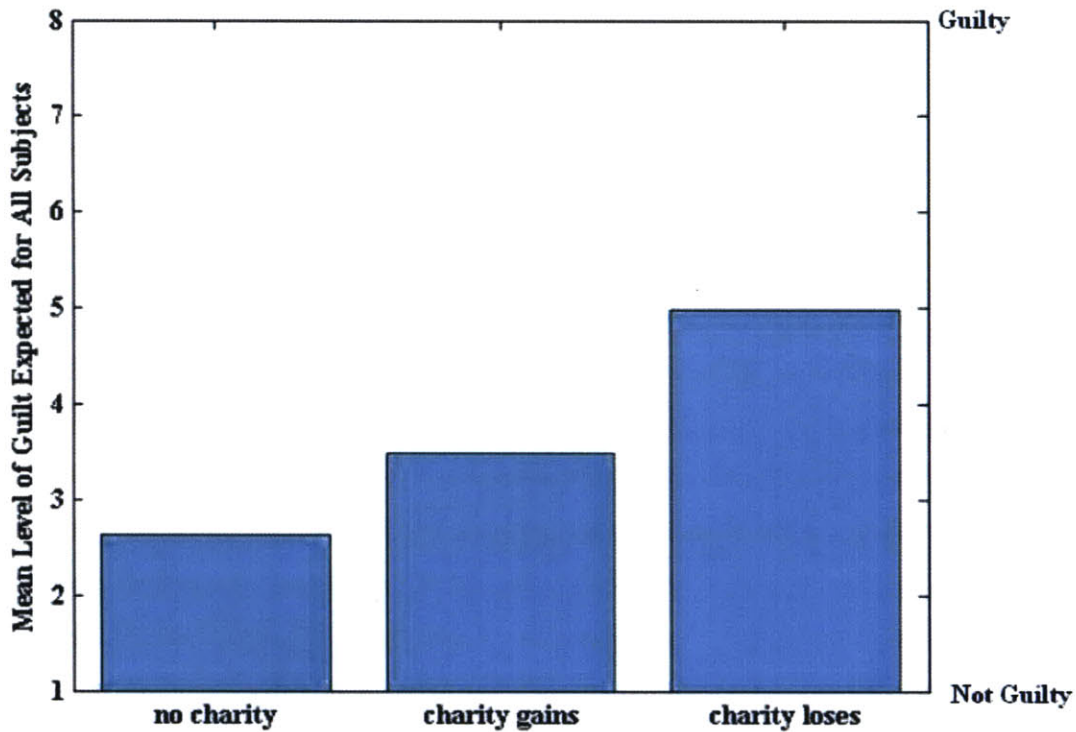


Figure 8: Our final estimate for the expected level of guilt experienced per subject by condition. We see three distinct levels of guilt experienced for the three different conditions.

Now that we have estimated the level of guilt experienced by condition, justifying the experimental setup as being meaningful for inducing differentiated levels of guilt, we will now explain the setup in more detail.

2.6 The Experimental Setup

There were 56 (30 female and 26 male) that participated in this study. Their ages ranged from 18 to 57 with a mean age of 27. The only screening criteria for participants was that they were at least 18 years of age and had a modest amount of experience using computers. Subjects were recruited through posters (see Appendix A) placed on bulletin boards at MIT and the surrounding areas. Of the 56 participants, 32 were MIT undergraduate or graduate students.

The participants were randomly assigned to one of four experimental conditions. As we described in the previous section, the four experimental conditions will be referred to as *control*, *no_charity*, *charity_gains*, and *charity_loses*.

Subjects began by signing a consent form, confirming their biographic information, and answering a few basic questions depending on the conditions to which they were assigned. In the case a subject was assigned to either the *charity_gains* or *charity_loses* condition, the subject was told that a charity could receive compensation depending on his/her performance and was asked to indicate a favorite charity s/he would like to see a donation made toward. Subjects could then choose from a list of 21 charities or specify another charity in an “other” box. Screen shots of all the screens that the subjects saw are shown in Appendix D.

They then were asked to place EKG electrodes on their bodies according to the diagram in Figure 11. The experimenter also placed electrodes on one palm (the one that did not use the mouse) and the corresponding forearm to sense electrodermal activity with the HandWave device designed by Marc Strauss (see Figure 14 for positioning of electrodes). It is worth noting that the electrode on the palm had a tendency of falling off due to perspiration and so the experimenters tried to secure the electrode better with additional medical tape, although it may still have come undone in some cases. Once the sensors were set up to record data, the

experimenter left the room. Subjects were told to do the rest of the experiment on their own while they were given instructions both orally through pre-recorded messages and written on their computer screens. They were also told that if they needed clarification about any of the instructions, they could send instant messages through an instant messaging (IM) client to the experimenter who would be in a nearby office.

Once the experimenter left the room, the automated system informed the subject that s/he would be compensated based on his/her score on a quiz and were given instructions on how the quiz would be scored. Subjects are told they will be given a piece of paper with 30 grids, each having 12 numbers. Their task was to find pairs of numbers contained within a single grid that add up to exactly 10. Their score is the number of such pairs that they find. Figure 9 shows two example grids along with their solutions.

0.23	1.86	7.34	9.96	8.26	4.79
3.40	6.60	1.71	2.81	3.38	9.58
1.07	1.97	5.23	7.72	6.55	0.04
6.01	9.13	0.55	8.84	1.50	9.33

Figure 9: Example grids shown to subjects along with their solutions

Subjects were then told that they would do a practice quiz that would not count toward their compensation before participating in the real quiz. When they were ready to take the quiz, they were told that the practice quiz was on a shelf in front of them. They were then timed for five minutes while they took the quiz. Once the quiz was over, subjects in the *control* condition were told to inform the experimenter that they were through so that he could collect their work. Subjects in the *no_charity*, *charity_gains*, and *charity_loses* conditions were told to shred their work in a paper shredder. The idea behind allowing subjects in the last three conditions to shred their work was so that they could report a score that was different from their true score (i.e. cheating) without the experimenter knowing. All subjects then reported their scores on the practice quiz on the computer.

Subjects were then told that the real quiz would be identical to the practice quiz, including what they do with the quiz when they are finished. Specifically, subjects in the *control* condition were told that they would need to turn in their work, while subjects in the *no_charity*, *charity_gains*, and *charity_loses* conditions were told that they would shred their work without the experimenter seeing their work.

Subjects were then told how they would be compensated. The compensation scheme as reported to the subjects was as follows:

- 1) They would receive a \$5 participation payment no matter what the outcome of the real quiz came out to be.
- 2) They would also get a 1 in 3 chance of winning a pot whose size would be \$1.50 per pair that they found. The reason we introduced the payoff as a gamble, rather than a guaranteed amount of money, was to increase the possible earnings on the same budget.
- 3) For the *charity_gains* and *charity_loses* conditions, a charity that they chose at the beginning of the game would also have a 1 in 3 chance of winning a pot of money. The size of this pot would be determined as follows:

- a. For subjects in the *charity_gains* condition, the charity's pot would be the same size as their own pot. So, by scoring well, the subjects in this group help both themselves and the charity.
- b. For subjects in the *charity_loses* condition, the charity's pot would be equal to \$45 minus the size of their pot (so that they both add up to \$45 combined). So, by scoring well, subjects in this condition help themselves, but hurt the charity.

For the *control* and *no_charity* conditions, there was no mention whatsoever of a charity at any point in the experiment.

Subjects were then told the location of the real quiz and were timed again for 5 minutes. When the 5 minutes was over, subjects in the *control* condition were told to turn in their work to the experimenter while subjects in the *no_charity*, *charity_gains*, and *charity_loses*

conditions were told to shred their work. They then reported their scores on the computer screen.

After filling out a brief survey on how they felt about the ethics of the affective sensors (as part of Carson Reynolds' work), they were then told to think about the following questions/issues (explained both orally and written):

- 1) Are you happy with your performance and the size of your pot?
- 2) Were you *honest* in reporting your score?
- 3) If you have overestimated your score, you have been *dishonest* and have unfairly influenced the size of your pot [and the charity's pot if in *charity_gains* or *charity_loses* conditions].

The goal in asking the subject to think about these questions was to get the subject to reflect on the ramifications of his actions and thus induce feelings of guilt in some subjects. By recording the subject's physiology during this portion of the experiment, we could compare how the physiology differed across conditions and relate these results back to the differences in expected guilt intensity for the different conditions. For the analysis portion of this chapter, the data collected during this time period will be referred to as the "region of interest."

Subjects were then told what their final compensation would be and were then given their final compensations. Subjects were also debriefed on the true nature of the experiment.

The entire experimental procedure is summarized in Figure 10 with an approximate timeline.

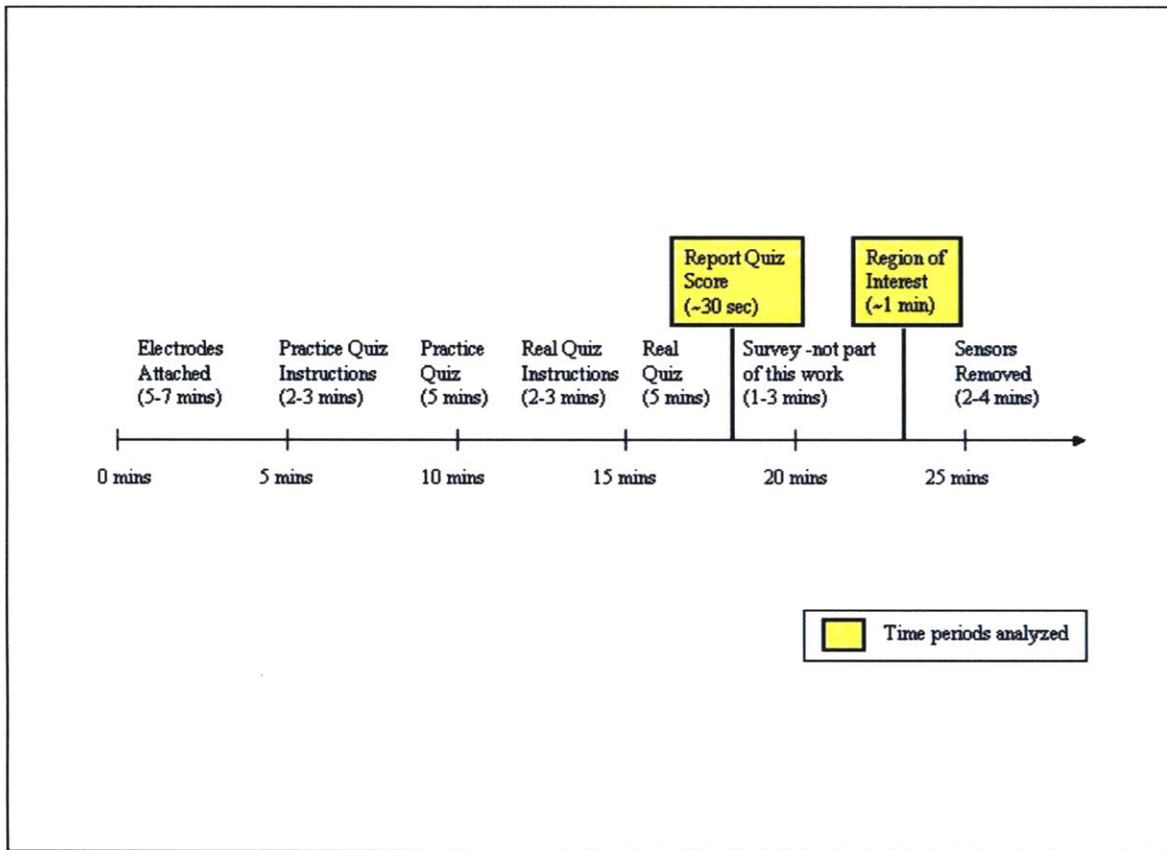


Figure 10: An overview of the experimental procedure with an approximate timeline.

It is also important to note that there were a number of extremely subtle aspects of the experiment that had to be controlled in order to ensure that people felt that they can cheat without getting caught. The experiment was inspired by an experiment by Dan Ariely who had been able to obtain cheating behavior in a group setting without sensors with a condition similar to our *no_charity* condition (there were no conditions similar to our *charity_gains* and *charity_loses* conditions). Thus, we had initially not anticipated great difficulty in getting subjects to cheat in non-*control* conditions. However, once Ariely tried the experiment with individuals outside of a group setting, he was and is still unable to achieve a significant amount of cheating behavior. This finding is consistent with the difficulties we encountered in our initial efforts. Appendix E lists some of the many changes we had to make to our initial protocol in order to get people to cheat at all. We highly recommend that anyone aiming to replicate the results of this experiment read this list first.

2.7 Sensor Data Analysis

We recorded time stamped electrocardiogram (EKG) and electrodermal activity (EDA) for all subjects. In this section, we first discuss the sensors that we used for the reader who is not familiar with these sensors. We then discuss the integrity of the data sets. Finally, we complete our data analysis.

2.7.1 Data Integrity

This section describes the limitations of the data sets that we collected.

2.7.1.1 Electrocardiogram

We used the ProComp+ system from Thought Technology Ltd. with disposable heart trace wet gel foam snap electrodes (cat # 8050). The system recorded the EKG signal at a rate that averaged 182 Hz and ranged between 164 Hz and 206 Hz for the data set used. Subjects were told to place the electrodes on their bodies according to the diagram in Figure 11. Subjects were told to ask the experimenter to leave the room for privacy if they wished.

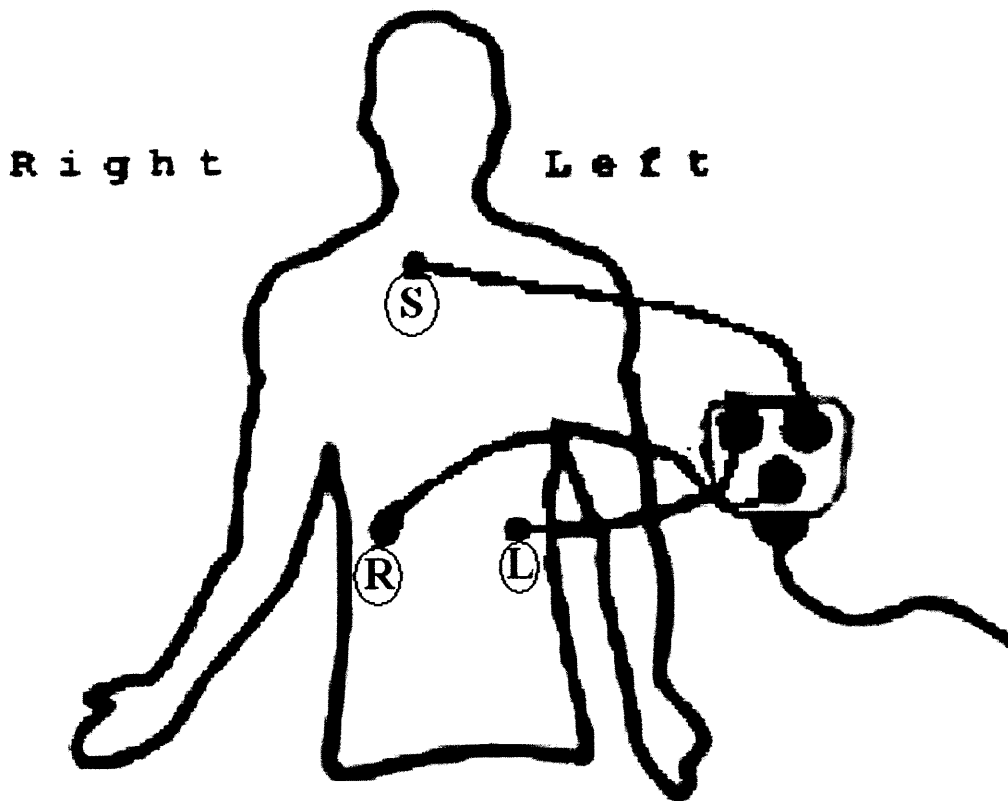


Figure 11: Subjects were told to place the electrodes on their bodies according to the above diagram. They were given disposable electrodes and leads with the labels “S,” “L”, and “R.”

We visually examined the EKG data coming out of the system from ourselves before beginning the formal experiment and the data looked like it was of good quality. Since the ProComp+ system is a commercial system, we had not performed extensive data quality analysis prior to the experiment. However, after collecting the data for all 56 subjects and extracting inter-beat intervals, we found that the majority of these data sets would not be valid for analysis. While this is unfortunate, the author and Carson Reynolds are currently collecting additional data, which could be used to verify the findings of this thesis.

We determined whether a particular subject’s EKG data would be used based on a visual inspection of the data set overall and an inspection specifically in the region of interest (i.e. the time period when the subject is asked to reflect on the ethics of the situation). We

identified portions of the data set as “clean” or “too noisy to analyze” and required that at least 80% of the overall EKG signal was “clean” and at least 95% of the EKG signal in the region of interest was “clean.” Since this process was done through visual inspection, the 80% and 95% requirements were approximate. Figures 12 and 13 show examples of regions that we identified as “clean” and “too noisy to analyze.”

When we examined the data sets overall, while nearly all of the data sets had portions of clean and noisy data, we found that only 19 of the data sets met the 80% clean EKG signal requirement. If we had a more robust method for dealing with noisy EKG data, we may still have been able to work with this data, but unfortunately, with the tools we had, this was impossible. Moreover, when we examined the signal specifically in the region of interest, we found that only 11 of the 19 data sets meeting the previous noise requirement met the 95% “clean” requirement in this region.




Figure 12: An example of an EKG signal that was identified as clean.




Figure 13: An example of an EKG signal that was identified as too noisy to analyze.

In sum, 11 of the EKG data sets obtained were determined to be valid for analysis. We are currently working to determine whether there is a way to reduce the noisiness of this signal or whether it will be necessary to use another machine to obtain clean enough data for analysis.

2.7.1.2 Electrodermal Activity

We used an EDA measuring device developed by Marc Strauss called the HandWave device with BioPac systems 1” solid gel electrodes (Cat. # EL504). The system recorded the EDA signal at 40 Hz. We placed the electrodes on the subjects’ non-dominant hand/arm (based on which hand they used the mouse with) according to the unipolar placement shown in Figure 14.




Figure 14: Electrodes were placed according to the unipolar placement in this diagram.

As we did for the EKG data sets, we visually examined EDA data coming out of the system from ourselves before beginning the experiment and the data looked good. Our initial tests included one complete experimental run and, when we examined the data set, it looked clean. Indeed, upon visual inspection of the data collected, 45 of the 56 data were clean enough to analyze. The 11 data sets that were identified as not being clean enough to analyze were of varying levels of quality from some that did not have any valid regions to others that had just a few noisy regions with the rest appearing clean. We suspect that the cause of the majority of the noisy data was due to a loose electrode attached to the palm, which tends to build up sweat. Figures 15 and 16 show examples of clean and noisy data sets, respectively.




Figure 15: An example of 10 minutes of clean EDA data.




Figure 16: An example of 10 minutes of EDA data that is too noisy to use.

While 45 data sets appeared clean enough to use in our analysis, it was highly unfortunate that we discovered that 29 of these data sets simply stopped prior to the end of the region of interest. We have since discovered that the cause of this behavior was likely the result of a slightly loose battery connection, which could open when the battery was moved rapidly, and/or a loose connection from the HandWave device to the electrodes. We have since taken steps to eliminate this behavior and are currently collecting additional data. However, this means that the data analyzed in this thesis will consist only of the remaining 16 sets. We hope to be able to confirm our preliminary findings from this thesis with the new data we are currently collecting.

In sum, 16 of the EDA data sets obtained were determined to be valid for analysis. We believe we have determined the cause for the frequent cut-offs in our data collection and have taken steps to reduce or eliminate these cut-offs. We are also exploring whether different electrodes reduce the noisiness of the EDA signal. We are currently repeating this experiment

with these variations and we expect that the majority of the new data sets obtained will be valid for analysis.

2.7.2 Analysis

It was our intention to build machine learning algorithms to distinguish between the different groups in this experiment in an attempt to build a “guilt detector,” in a sense. However, since the amount of valid data that we collected was small, our results from any classifier would be highly suspect. We have therefore refrained from using machine learning techniques and, instead, have simply reported aggregate data trends. Moreover, since we found only three subjects with both valid EKG data and valid EDA data, we have limited our analysis to each signal separately and have not done any joint signal analysis.

Before we could do this analysis, we needed to define a baseline time period to which we could compare to the region of interest in terms of the sensor data collected. We chose the period of time when the subject was given the instructions for the practice quiz and took the practice quiz as this baseline since the experimental setup was identical for all non-control subjects during this time period.

The first analysis that we did was to examine how the ratio of the average EDA level during the region of interest varied relative to baseline across the conditions. Subjects had similar baseline levels of EDA across conditions (see Appendix N). For each subject, we divided the average EDA output during the region of interest by the baseline value for each individual. We then averaged these ratios within a condition. Figure 17 shows the result of this analysis.




Figure 17: Normalized EDA in region of interest vs. baseline with 50% error bars.

It appears that there is a trend toward increasing EDA relative to baseline with increasing levels of guilt (both go up as we move to the right on this graph based on the results calculated previously). Note that the small sample size makes it very difficult to draw strong conclusions from this data. We used a simple one-sided t-test and found that normalized EDA in the region of interest for the three non-*control* conditions is greater than the *control* condition at the 95.3%, 91.3%, and 99.7% confidence levels (in the order in which they appear in Figure 17) and thus the non-*control* conditions are all significantly different from the *control* condition based on this measure. The *charity_loses* condition was greater than the *charity_gains* condition at only a 75.2% confidence level based on this one-sided t-test.

It is worth noting that the above relationship does *not* hold at the time in which subjects actually report their scores. Figure 18 shows the normalized EDA while subjects report their scores with 50% error bars.




Figure 18: Normalized EDA while subjects report their scores with 50% error bars. There does not appear to be any statistically significant difference in these quantities.

The above result demonstrates that physiological information more closely associated with the cheating behavior and guilt level may be contained in the period of reflection (our region of interest) than in the actual time at which the subject cheats. Thus, to build a better lie detector, one may want to analyze physiological data recorded during a period of reflection about the consequences of the lie in addition to physiological data recorded during the lie itself. It would be straightforward to investigate whether this additional data would be useful in a common lie detection experiment where individuals are known to be lying or not, such as in (Tsiamyrtzis et. al., 2005).

The second analysis that we did was to determine the relative frequency of skin conductivity responses (SCRs) in the region of interest vs. baseline. SCRs, equivalent to electrodermal responses, are small “spikes” in EDA that we extracted with the Matlab code developed by Jennifer Healey that is in Appendix G. Subjects had similar baseline frequencies of SCRs across conditions (see Appendix N). We expected to see a similar trend to the normalized EDA trend in this data set, but the error bars were too large to draw any strong conclusions. Figure 19 shows our results with 50% error bars.




Figure 19: Relative frequency of SCRs in region of interest vs. baseline by condition with 50% error bars. We see no strong correlation between guilt and frequency of SCRs.

Last, we examined how EKG varied by condition. As mentioned in the background section of this chapter, a key finding by both (Sloan et. al., 1994) and (McCraty et. al., 1995) was that

LF activity increases relative to both HF and MF activity with increasing mental stress and anger. We examined whether this trend would also hold as feelings of guilt increased.

To test this hypothesis, we first needed to extract the frequency distribution for the region of interest and baseline for each subject. To do this, we used Yuan Qi's spectral estimation code (Appendix I), which estimates the frequency distribution at each time step. For the region of interest and baseline, we calculated the average frequency distribution (for the entire time interval) using this code. Figure 20 shows a sample spectrogram produced with this code.




Figure 20: An example of the spectrogram produced for one time step.

When we examined the LF activity relative to HF and MF activity for the region of interest vs. the baseline region, we did not obtain results similar to those obtained by (Sloan et. al., 1994) and (McCraty et. al., 1995) for anger and mental stress. Figures 21 and 22 show how the ratio of LF to HF and LF to MF changed relative to baseline for the region of interest as averaged over our subjects.

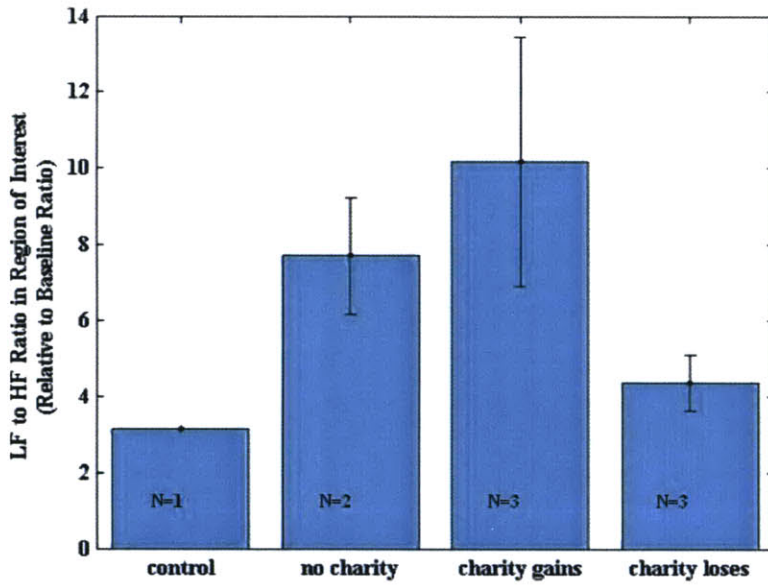


Figure 21: LF to HF ratio in region of interest relative to baseline LF to HF ratio.




Figure 22: LF to MF ratio in region of interest relative to baseline LF to MF ratio.

Last, we also examined the time the subjects took to report their scores across conditions. Prior literature, such as (Ekman et. al., 1991), has reported that answering a question too slowly is often an indicator of lying behavior. If this were true, we would expect to see the amount of time taken to report the quiz score increase with the amount of cheating behavior. However, when we performed this analysis, we found that the amount of time taken to report one's score was significantly higher in all non-*control* conditions than the *control* condition ($p < 0.00002$ for *no_charity* and *charity_gains*, $p < 0.004$ for *charity_loses* condition), but did not vary significantly among the non-*control* conditions. This result may imply that the increased time to report one's score is indicative of one's contemplating a lie, but not indicative of actual lying behavior. The average amount of time taken to report one's score by condition is shown in Figure 23.

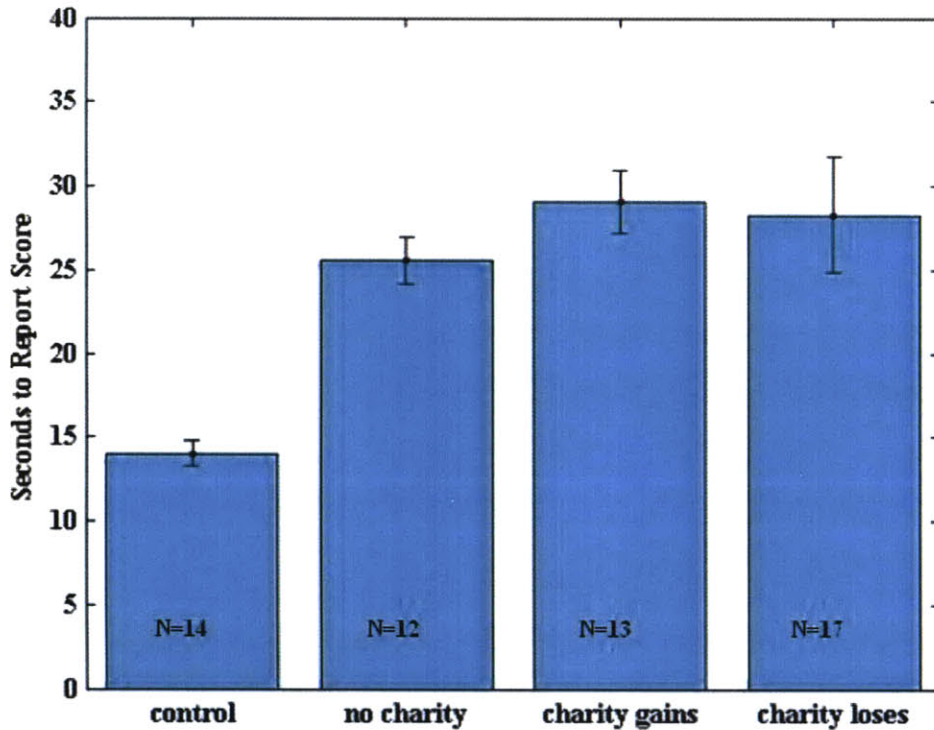


Figure 23: Amount of time taken to report score by condition with 50% error bars. Subjects in non-control conditions took significantly longer to report their scores than those in the control condition, but there was not significant variation among non-control conditions.

2.8 Conclusions

This work was an exploration of how physiology changes in response to cheating with different ethical incentives or different levels of the feeling of guilt. We created an experiment to produce varying levels of guilt and recorded EDA and EKG for subjects in each condition. We attempted to create a guilt detector using the physiological data recorded in this experiment.

The primary contribution of this work was a novel method of producing varying levels of guilt intensity for different groups. We encountered many difficulties along the way, most of which are outlined in Appendix E, but we believe we have created an experiment that will produce robust differences in average levels of guilt across conditions.

We showed that there was a relationship between guilt level and electrodermal activity (EDA) relative to baseline for a period of reflection that did not exist during the actual period of cheating. This finding is interesting because it may mean that one could build a more robust lie detector by not only examining physiology during the actual lie, but also during a self reflection period.

In our data set, we did not find a correlation between guilt level and frequency of skin conductivity responses (SCRs) relative to baseline during the period of reflection. This finding is inconsistent with some previous literature that has found SCRs to be a more robust indicator of emotional states than overall EDA level.

Our EKG analysis was inconclusive as we were not able to confirm our hypothesis that the LF/MF and LF/HF ratios should be higher (relative to baseline) during the reflection period in high guilt conditions than low guilt conditions. It is unclear whether this finding is accurate or whether it is a consequence of the small size of our data set.

Last, we found that the amount of time taken to report one's score was significantly higher when subjects were able to cheat than when they could not cheat (the *control* condition), but did not vary based on the amount of actual cheating behavior. This finding implies that the amount of time taken to report one's score is related to the act of contemplating cheating, but not necessarily indicative of actual cheating behavior.

2.9 Future Work

One follow-on to this work would be simply doing the same experiment over again with better integrity sensor data. We believe we have corrected the problems for EDA signal, although we are still not sure if we have corrected the problem for the EKG signal. We are

currently running another trial of this experiment and hope to obtain results that are more statistically significant. Another follow-on would be to do a similar experiment with additional sensors, such as respiration sensors or electromyogram sensors. There are a wide variety of sensors that could be used in this context.

Perhaps most importantly, since the small size of this data set forced us to merely correlate physiological signals with guilt levels, one could create a true “guilt detector” with larger amounts of data. It is significantly harder to create a classifier than simply examine aggregate data trends because of the variability in physiological response between subjects. However, by finding a rich set of features based on the sensor data and a large sampling of subjects, one could use a variety of pattern recognition tools, such as support vector machines, to distinguish between different levels of guilt based on these features.

Another area for future work is examining whether physiological data collected during a period of self-reflection following a lie would aid in producing a better lie detector. However, we hope that any researcher investigating this area will use this technology responsibly as lie detection is a technology that could be used irresponsibly in the wrong hands.

Lie detection researchers could also verify whether physiology recorded during a reflection period following a lie is useful in increasing overall lie detection accuracy. This type of analysis would be fairly straightforward in the context of a classic lie detection experiment where the experimenter knows whether a subject is lying or not.

Last, researchers could try to build a fully automated lie detection experiment similar in form to the experiment we created. Such a fully automated experiment would eliminate emotion contagion effects between the interviewer and the subject.

3 Covert Mouse Behavior Analysis

This chapter describes an experiment designed to elicit positive/negative affect and high/low arousal and detect these states based on simple computer mouse behavior. After an initial introduction motivating the work and explaining where the data set was obtained, the second section describes some relevant work in this general area. The third section describes the experimental setup in detail. The fourth section examines whether the states intended to be induced were elicited properly. The fifth section describes our data analysis. The sixth section summarizes the conclusions we were able to draw from our analysis. Finally, the seventh section suggests possible avenues for future work in this area.

3.1 Introduction

Detecting the affective state of a user based on ordinary mouse behavior has a wide range of applications. Although one is unlikely to find very rich affective information from ordinary mouse behavior, the ubiquity of the computer mouse makes applications based on this type of affect sensing technology potentially available to a very large population. The ubiquity of the computer mouse also opens up the possibility of detecting affect covertly and it remains an open question about whether this is possible (and also whether it is ethically acceptable, but this is not a subject of this thesis). In the experiment introduced in this section, we attempt to detect affect covertly from ordinary mouse behavior.

The author's contribution is the analysis of the mouse movement data in the context of this experiment. Patrick Gomez, Philippe Zimmermann, Sissel Guttormsen-Schar, and Brigitta Danuser of ETH Zurich conducted this experiment in its entirety (Gomez et. al., 2005).

Sensing in this experiment was covert since subjects were not told that the mouse behavior was being recorded until after the task was complete. The ETH team has published analysis of data collected from several physiological sensors during this experiment, but has not published any findings related to the mouse behavior exhibited.

3.2 Relevant Work

There has been very little work that we are aware of which investigates the role of affect in use of an unmodified computer mouse (i.e. without additional sensors attached) and its ordinary trajectories.

Andrea Lockerd, Florian Mueller, and Ted Selker from the MIT Media Lab studied mouse movements of 17 people as they browsed web pages. When asked to choose a CD online, the subjects that used the mouse to look through the list hesitated longest on the link corresponding to their second choice CD before choosing their first choice in 65% of the cases on one site and 75% of the cases on another site. This showed a correlation between user interest and mouse behavior. Lockerd also believes that a user's mood and personality may influence how the user interacts with a mouse based on her qualitative observations of users in her study, although she has not officially investigated this scientifically. (Mueller and Lockerd, 2001)

3.3 The Experimental Setup

There were 96 (46 female and 50 male) subjects that participated in this study. Their ages ranged from 17 to 38 with a mean age of 24. The only screening criteria for participants was that they were in good health and native German speakers.

The participants were assigned to one of five experimental conditions. The five experimental conditions were: positive valence / high arousal (20 subjects), positive valence / low arousal (18 subjects), negative valence / high arousal (19 subjects), negative valence / low arousal (19 subjects), and a control group of neutral valence / neutral arousal (20 subjects). For the rest of this experiment, these conditions will be abbreviated as *PH*, *PL*, *NH*, *NL*, and *control*, respectively.

Subjects were told that the purpose of the experiment was about physiological reactions to an active and passive task. They were not told that the experiment had anything to do with inducing or detecting affective states and were not told that their mouse behavior was being

recorded during the experiment. They were told about the physiological signals that were being measured in the experiment: respiration, EDA, heart rate, and corrugator activity (corrugator EMG or electromyography). Subjects were debriefed after the experiment about its true purpose.

After initial introductions to the experiment and several sensors being placed on the subjects' bodies (sensor data not analyzed in this thesis), subjects were asked to fill out a questionnaire with basic biographic information. They then completed a simple online shopping task as practice for the later parts of the experiment. Specifically, subjects were told to shop for office supplies on an e-commerce web site. This task was selected because it was expected to have little impact on mood and is a type of task that is familiar to most computer users. For each task, subjects were told to purchase 7 specific items and then type a final message to the shop operator. An example of the interface from the e-commerce site is shown in Figure 24.

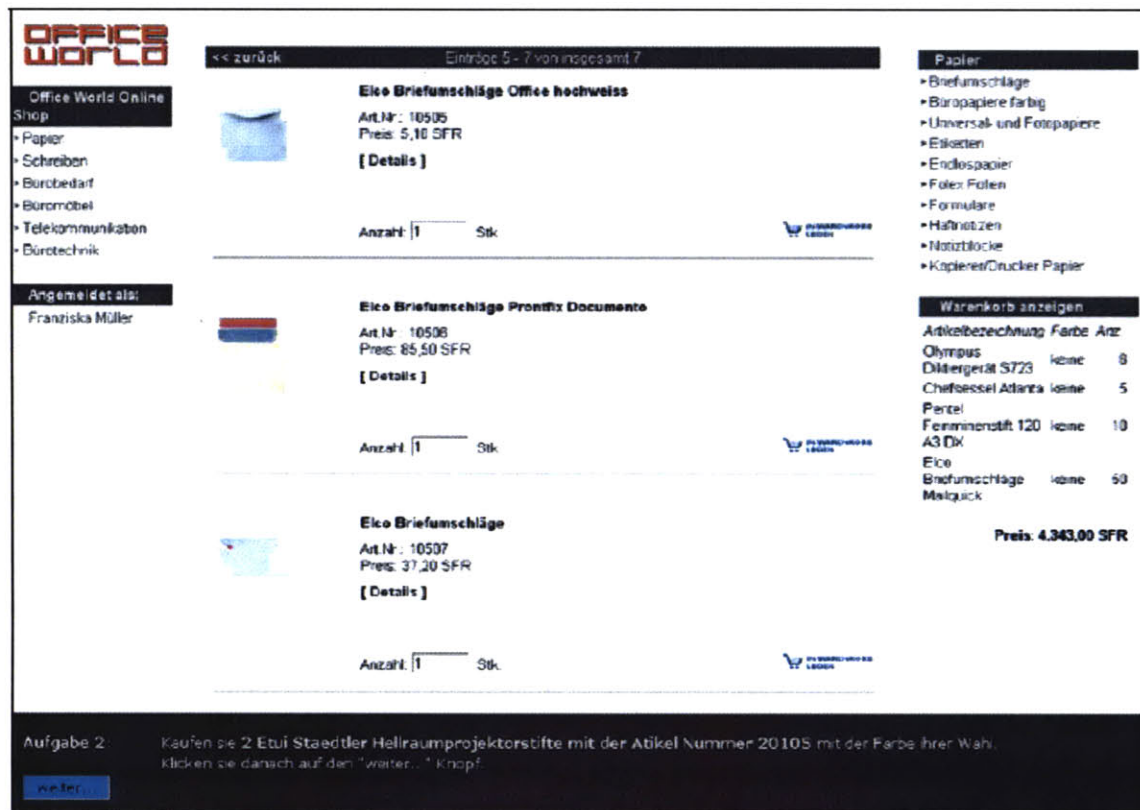


Figure 24: A view of the e-commerce site used for the subjects' tasks.

This initial practice task was followed by a mood assessment to determine a baseline mood. Participants rated their mood using the 9-point Self-Assessment Manikin or SAM (Lang, 1980) and a semantic differential with 6 bipolar adjective pairs in German, which was not analyzed as part of this work. SAM is a tool for measuring valence and arousal consisting of a graphic figure representing nine levels for both valence and arousal (see Figure 25).




Figure 25: Self-Assessment Manikin (SAM) used to allow subjects to self-report their moods.

Subjects then began two identical sequences with the following actions: a mood induction using a film clip designed to produce a specific affective state, a mood assessment with SAM, an online shopping task, and a second mood assessment. During the first mood induction stage, every participant saw the same film clip chosen to induce a neutral mood. The film clip in the second mood induction stage varied depending on the subject's condition. All film clips were between 7 and 11 minutes long and were pre-screened based on their ability to effectively and reliably induce the appropriate affective states while staying within ethical guidelines. The film clips used are summarized in Figure 26.

Mood	Content
<i>Control</i>	Educational movie about the characteristics of materials (i.e. wood, concrete)
<i>Neutral</i>	Documentary about the architect Louis Kahn and his work
<i>PH</i>	Clips of different sports (i.e. surfing, skiing, climbing) with rock/pop music
<i>PL</i>	Takes of landscapes and animals with classical music
<i>NH</i>	Extract from Deer Hunter (Cimino, 1978), depicting captives in Vietnam
<i>NL</i>	Documentary about the Borinage (Jean, 1999), a Belgian mining area and slum

Figure 26: Film clips used during mood induction phases. The film clip labeled “Neutral” was used as the first film clip for all subjects. The other film clips were used for mood induction in the second film clip for the corresponding conditions.

The entire experimental procedure is summarized in Figure 27.

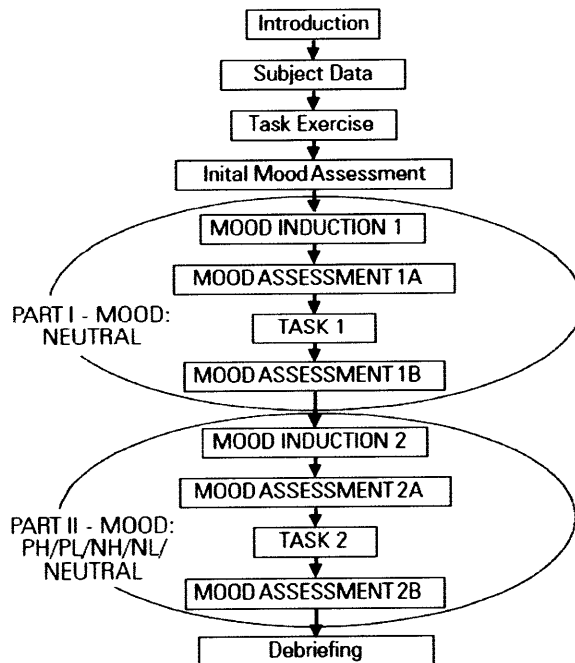


Figure 27: Experimental procedure summary.

3.4 Manipulation Check

The average self-reported valence and arousal changed considerably from their baseline levels (during the initial mood assessment) to their levels at mood induction 2 (after the second film clip). Valence and arousal levels changed predictably in all cases except for *PH*, where the valence level went unchanged off of baseline, and *NL* where the arousal level actually increased slightly off of baseline, although still within the 95% confidence level. Figure 28 shows how valence and arousal levels changed off of baseline for the self report after the second mood induction.

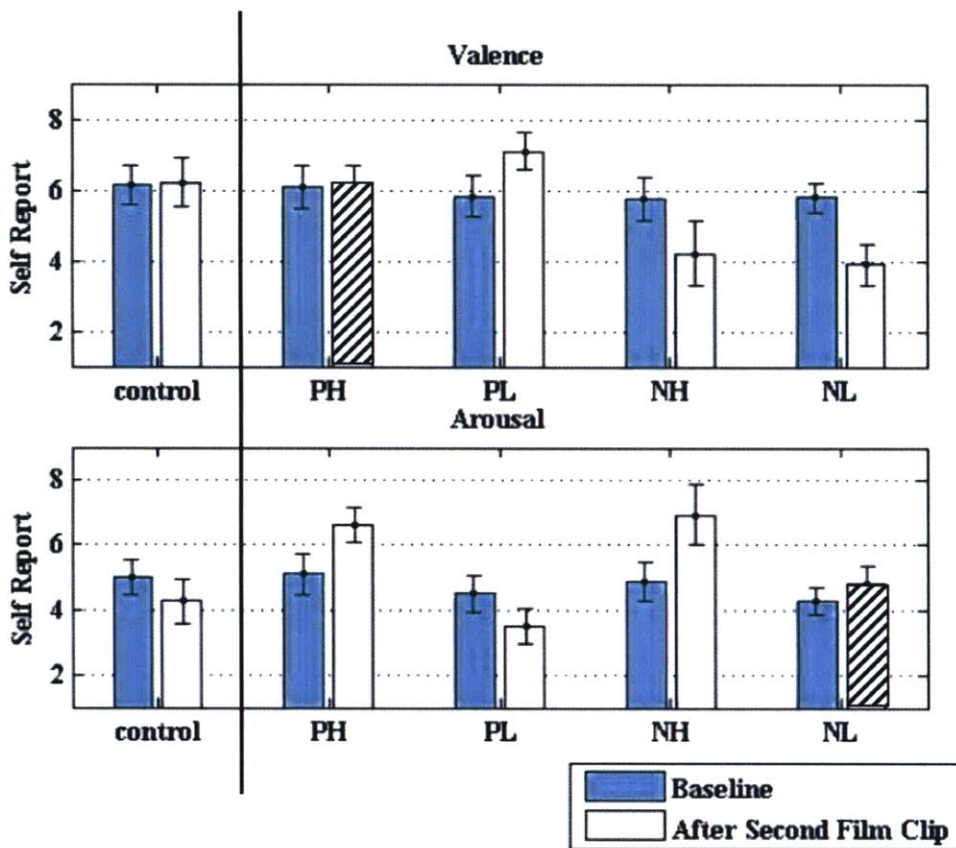


Figure 28: Comparison of self reported valence and arousal for baseline (initial self report) and after the second film clip. The error bars shown are for 95% confidence level. Slanted lines indicate that induction was not successful.

An alternative explanation of the results in Figure 28 is that all non-*control* conditions were elicited successfully, but the baseline and *control* states (neutral valence, neutral arousal)

were not elicited properly and, instead, elicited positive valence, low arousal states. This interpretation was the implicit view taken in (Gomez et. al., 2005) when the authors ignored the control condition and, after verifying that baselines did not differ significantly across conditions, analyzed only the absolute levels of valence and arousal. However, for the rest of this work, we will assume that the *control* and baseline states were elicited properly (neutral valence, neutral arousal), but the *PH* condition did not elicit valence properly and the *NL* condition did not elicit arousal properly.

While the mood induction was generally successful, the effect did not last through the self report at the end of the second task. Figure 29 shows that self report ratings were back to the baseline level by the end of the second task.




Figure 29: Comparison of self reported valence and arousal for baseline (initial self report) and after the second task. Error bars shown are for 95% confidence level. This figure shows that the mood induction wore off by the end of the second task.

These two charts show that the effect of the mood induction wore off somewhere during the task itself, but we do not know exactly when this occurred.

3.5 Analysis: Detecting Affective State

While there are a variety of features one could use in this analysis, we focused on two highly intuitive features: speed and acceleration. Just as we did in the previous chapter on guilt detection, the analysis focuses on the differences in these features during a time period of interest (the second task once the appropriate mood has been elicited) from a baseline (the first task when all subjects are placed into a neutral mood).

Figure 30 shows how the average velocity varied with condition during task 2 off of baseline. Note that the arousal elicitation was not successful in the *NL* condition based on the self report after the film clip. The average velocity in the *PL* condition was less than the average velocity in the other conditions with p-values less than 0.25, 0.08, 0.11, and 0.07 for the *control*, *PH*, *NH*, and *NL* conditions, respectively. Thus, the average velocity in the *PL* condition was less than in the other conditions with reasonable statistical significance for all conditions but the *control* condition. This finding implies that average velocity off of baseline may be a reasonable indicator for arousal level.




Figure 30: Average velocity during task 2 as a percentage off of the baseline (task 1) velocity by condition with 50% error bars.

Similarly, Figure 31 shows how the average acceleration varied with condition during task 2 off of baseline. It appears that average acceleration is also an indicator of arousal level, although perhaps not quite as strong an indicator as velocity. The *PL* condition is less than the *control*, *PH*, *NH*, and *NL* conditions with p-values less than 0.27, 0.09, 0.17, and 0.04, respectively.

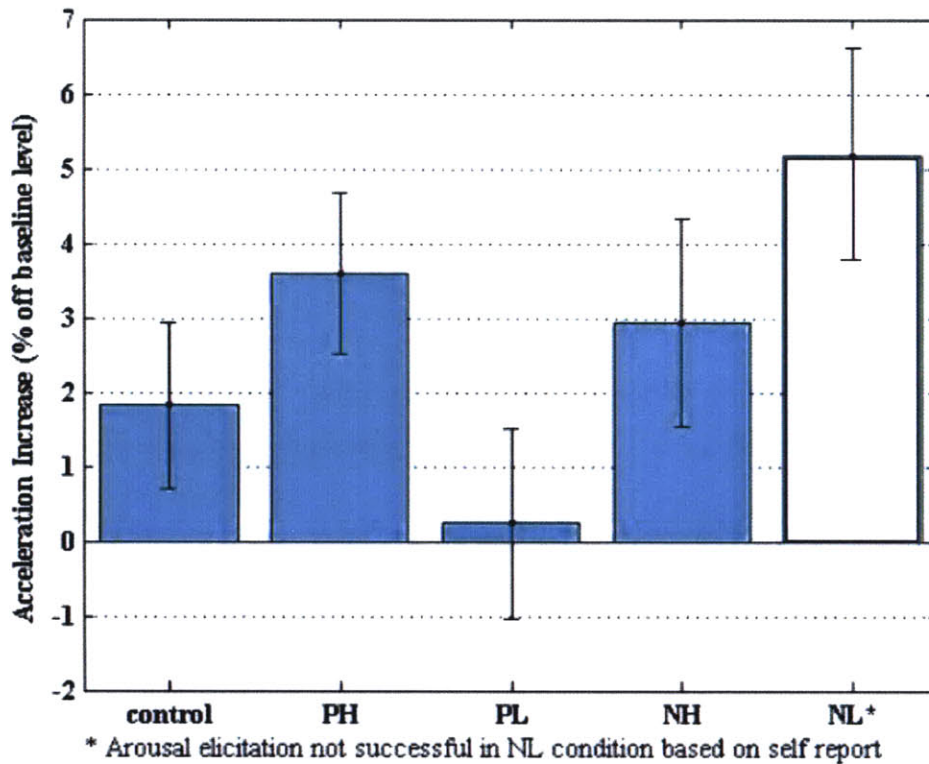


Figure 31: Average acceleration during task 2 as a percentage off of the baseline (task 1) velocity by condition with 50% error bars.

We also hypothesized that these relationships would hold better in the first 60 seconds of task 2. This hypothesis turned out to be correct for velocity, but not as significant for acceleration. Figures 32 and 33 show the results that we obtained when we looked only at the first 60 seconds of task 2. For the velocity during the first 60 seconds, the *PL* condition turned out to be less than the *control*, *PH*, *NH*, and *NL* conditions with p-values less than 0.10, 0.01, 0.03, and 0.03, respectively. For the acceleration during the first 60 seconds, the *PL* condition turned out to be less than the *control*, *PH*, *NH*, and *NL* conditions with p-values less than 0.22, 0.03, 0.07, and 0.06, respectively.




Figure 32: Average velocity during the first 60 seconds of task 2 as a percentage off of the baseline (task 1) velocity by condition with 50% error bars.

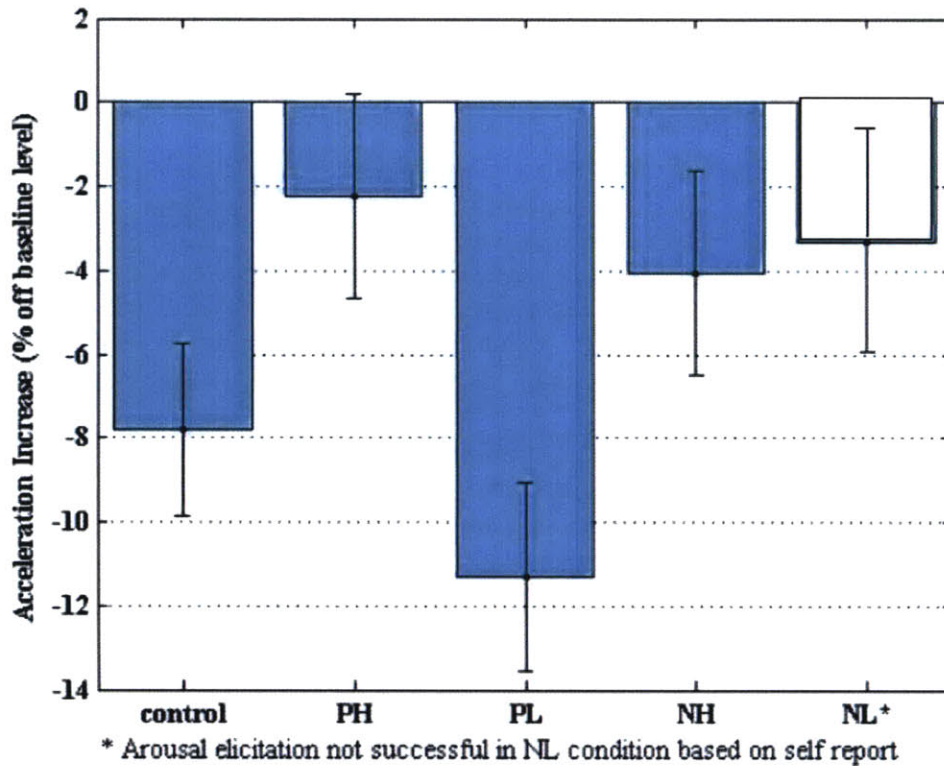


Figure 33: Average acceleration during the first 60 seconds of task 2 as a percentage off of the baseline (task 1) velocity by condition with 50% error bars.

We also examined how curvature varied with valence and arousal, but were not able to find any significant correlations. The feature related to curvature that we computed was the average tangential acceleration component during curves (trajectories that were not straight).

3.6 Conclusions

Our goal was to determine whether a user's behavior with an ordinary computer mouse during a typical online shopping task would be related to his/her affective state. We examined how mouse behavior varied with valence and arousal.

We found that average velocity and average acceleration of a user's computer mouse varies positively with arousal level during a shopping task based on results obtained from 96

subjects. These conclusions seem to hold more robustly for the first 60 seconds, when the level of arousal is likely to be closest to that which was elicited, than for the entire length of the shopping task. We were not able to draw any significant conclusions about any other potential features of the computer mouse trajectory or about how these velocity or acceleration may vary with valence.

3.7 Future Work

There are two primary ways in which this work could be extended.

The first way is by trying out different features, possibly some borrowed from the literature on on-line handwriting recognition, and seeing if any of these correlate well with valence and arousal. There are a wide variety of such features in the literature, such as those used in (Rumelhart, 1993).

The second way, likely to be done in combination with the first, is to create machine learning based classifiers to try to detect a user's level of arousal and valence from mouse movement. This was not attempted in this work due to the relatively small size of this data set. A simple classifier could be created by first extracting some features from the entire data set and then using support vector machines to distinguish between different levels of valence and arousal during the second task. A more sophisticated approach would model the temporal nature of the trajectory better with hidden Markov models or another technique. An example of this type of classifier used in an on-line handwriting recognition task is in (Neskovic et. al., 2000).

4 References

- Akselrod, S., D. Gordon, F. Ubel, D. Shannon, A. Barger, and R. Conen (1981), "Power Spectrum Analysis of Heart Rate Fluctuation: a Quantitative Probe of Beat-to-Beat Cardiovascular Control," *Science*, 213: 220-222, July 1981.
- Dennerlein, J., T. Becker, P. Johnson, C. Reynolds and R. W. Picard (2003). "Frustrating Computer Users Increases Exposure to Physical Factors," *Proceedings of the International Ergonomics Association*, Seoul, Korea, August 24-29, 2003.
- Ekman, P., R. Levenson and W. Friesen (1983). "Autonomic Nervous System Activity Distinguishes Among Emotions," *Science*, 221:1208-1210, September 1983.
- Ekman, P. and M. O'Sullivan (1991). "Who Can Catch a Liar?," *American Psychologist*, 46:913-920.
- Gomez, P., P. Zimmermann, S. Guttormsen-Schar, and B. Danuser (2005). "Respiratory Responses Associated with Affective Processing of Film Stimuli," *Biological Psychology*, 3:223-35, March 2005.
- Kircher, J., T. Packard, B. Bell and P. Bernhardt (1997). "Comparison of Skin Conductance and Skin Resistance Measures for the Detection of Deception," Report No. DoDPI02-R-0001. Department of Defense Polygraph Institute, Fort Jackson, SC 29207-5000, October 2001.
- Liu, K (2004). "A Personal, Mobile System for Understanding Stress and Interruptions," MIT SM Thesis, June 2004.
- McCraty, R., M. Atkinson, W. Tiller, G. Rein, and A. Watkins (1995). "The Effects of Emotions on Short-Term Power Spectrum Analysis of Heart Rate Variability," *The American Journal of Cardiology*, November 1995, pp. 1089-1093.
- Mueller, F. and A. Lockerd (2001). "Cheese: Tracking Mouse Movements on Websites, a Tool for User Modeling," Proceedings of CHI, Seattle, WA, April 2001.
- Neskovic, P., P. Davis, and L. Cooper (2000). "Interactive Parts Model: An Application to Recognition of On-Line Cursive Script," *Advances in Neural Information Processing Systems (NIPS) 13*, Denver, CO, 2000.
- Reynolds, C. (2001). "The Sensing and Measurement of Frustration with Computers," MIT SM Thesis, May 2001.

- Rumelhart, D.E. (1993) "Theory to practice: A case study – recognizing cursive handwriting," in E.B. Baum, editor, *Computational Learning and Cognition: Proceedings of the Third NEC Research Symposium*, SIAM, Philadelphia, 1993.
- Sloan, R., P. Shapiro, E. Bagiella, S. Boni, M. Paik, J. Bigger, R. Steinman and, J. Gorman (1994). "Effect of Mental Stress Throughout the Day on Cardiac Autonomic Control," *Biological Psychology*, March 1994, pp. 89-99.
- Tsiamyrtzis, P., J. Dowdall, D. Shastri, I. Pavlidis, M.G. Frank and P. Eckman (2005). "Lie Detection: Recovery of the Periorbital Signal through Tandem Tracking and Noise Suppression in Thermal Facial Video," *SPIE Conference*, 2005.

5 Appendices

Appendix B: Screen shots from registration form

Initial Screenshot at <http://arsenal.media.mit.edu/study/>

Welcome to Experiment Registration

The purpose of this experiment is to collect data to evaluate the acceptability of applications that recognize and communicate emotions.

- We are seeking subjects ages 18 and over.
- You will be asked to either participate in an online interview or perform a simple task at the Media Lab.
- Additionally, you will be asked to fill out questionnaires about your experience.
- Participants will receive at least \$5 and up to \$50 depending on the experiment you are chosen for and your performance during that experiment.

The experiment provides subjects with an opportunity to interact with cutting edge technologies and applications. The results of this research will help us better understand how to design systems that sense and communicate emotions.

[Start Registration](#)

Step 1 Screenshot

Demographics: Step 1 of 3

Please enter the following demographic information.

AGE

GENDER

NATIONALITY

EDUCATION

Next Page

Step 2 Screenshot

Time, Date, and Email: Step 2 of 3

Your session will take approximately 30 minutes.
You will receive between \$5 and \$50 for participating.
A modest amount of computer experience is required.

Note: If none of the following times work for you, please try back again later.

TIME SLOT

EMAIL

Step 3 Screenshot

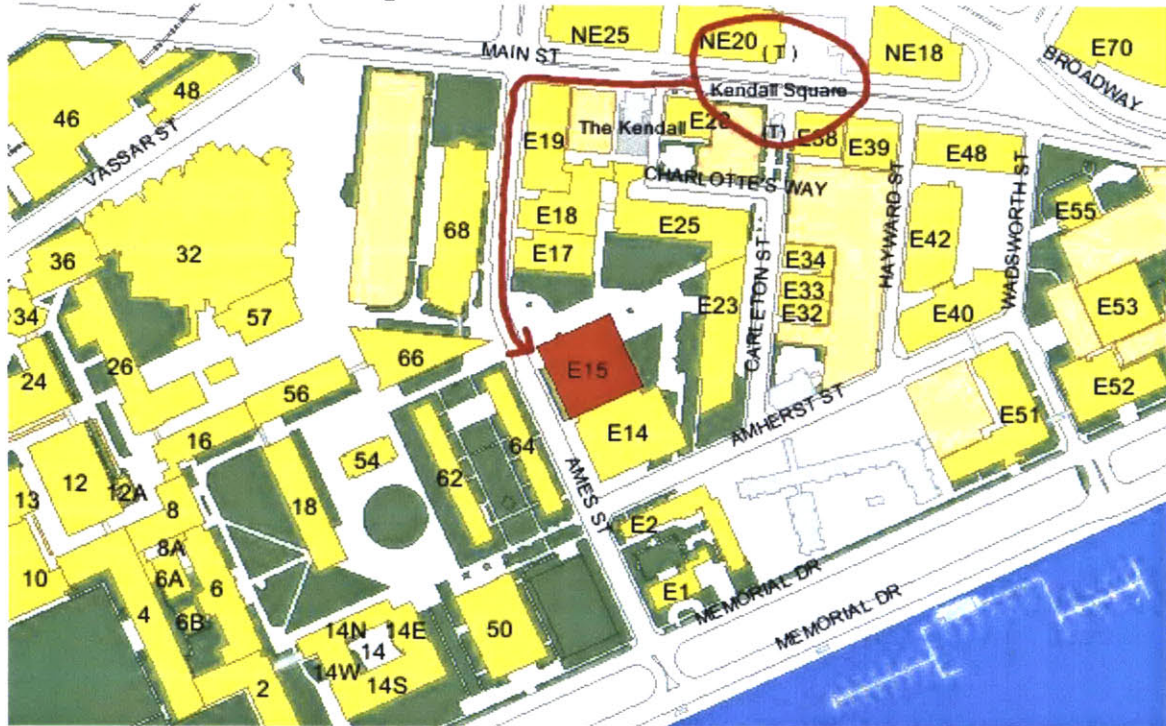
Take note: Step 3 of 3

Thank you for registering for our experiment. Your appointment is at the following time:

11:00-11:30 AM 5/10/2005

You will receive a confirmation email shortly. Take an opportunity to write this in your calendar.

Here are directions to the building and to the room:



The Media Lab is in Building E15 (a.k.a. The Weisner Building), at 20 Ames Street on the east side of campus. It's the building with the big white tiles, rounded corners and the big modern concrete arch connected to the side of it. Room 001 is located on the lower level (LL) of the lab directly across from the elevators.



If you have any problems, please email experiment@media.mit.edu.

First email confirmation sent immediately after registration. Other two are sent 24 hours in advance of appointment and 1 hour in advance of appointment, respectively, but are not shown in this Appendix.

Greetings,

You have signed up for an experiment at the MIT Media Lab at:

11:00-11:30 AM 5/10/2005

Here are directions to the building and to the room:

The Media Lab is in Building E15 (a.k.a. The Weisner Building), at 20 Ames Street on the east side of campus. It's the building with the big white tiles, rounded corners and the big modern concrete arch connected to the side of it. You can find a map of the area at this URL:

<http://arsenal.media.mit.edu/study/walk-e15.jpeg>

The Lab is easily accessible by T (see the Kendall Square T stop on the above mentioned map).

When you arrive at the Lab, please come to Room 001 on the lower level, directly across from the elevators. Have a seat on the couches; we will be with you shortly.

If you will be unable to make it to your appointment or know that you will be late, please contact experiment@media.mit.edu at least 24 hours in advance so that the experimenters can make alternative arrangements.

It is very important that you arrive on time to your appointment since the experiment may involve multiple participants. The whole group will be delayed if you are late so please be considerate of the other subjects and show up on time.

Thanks for your interest and participation.

Regards,
The Experimenters

Appendix C: Experimenter's Instructions for Running Experiment

Prior to subject showing up:

- 1) Place quizzes in appropriate locations (red on shelf, blue in drawer)
- 2) Place consent form by computer
- 3) Bring subject's demographic information up on screen
- 4) Start instant messaging client on subject's computer (Interviewee120g)
- 5) Start instant messaging client on your computer (Interviewer120g)
- 6) Make sure speakers are plugged in and turned onto 3/4 volume
- 7) Make sure there is sufficient scrap paper out

Beginning of experiment:

- 1) Have subject sign consent form
- 2) Have subject click through first couple of pages
 - a. If subject is in sensor condition, help subject place sensors on then turn on devices and start python scripts.
- 3) Reiterate that both written and oral instructions will be given over the computer
- 4) Show subject how to use IM in case s/he has never used it.
- 5) Explain that you must be contacted over IM if s/he has any questions so that sensor data will not be corrupted and so that no subjects will have an unfair advantage
- 6) Show subject where scrap paper is
- 7) Place one piece of scrap paper in shredder to show subject how it works (explanation: need to be sure next subject doesn't see work)
- 8) Wish subject 'good luck' and leave the room

During rest of experiment:

- 1) Make sure instant messaging client is constantly being monitored
 - a. Also, stay in office 120g (my office) for experiment if subject is in task 1
- 2) Answer any questions the subject asks
- 3) If subject is in task 1, s/he will tell you over IM when s/he is done with each quiz
 - a. Go into office and stop alarm with code **1209**
 - b. Collect quiz and make wait until s/he enters the number of pairs
 - i. If subject didn't circle pairs in practice, tell subject that it is essential that s/he do this in the real quiz
 - ii. Remind subject that you will verify number of pairs
 - iii. Remind subject that you must be notified within 30 seconds of the alarm going off or s/he will be disqualified
 - c. Leave office

After experiment:

- 1) Turn off sensors
- 2) Close Python scripts
- 3) Give subject compensation as explained on the screen
- 4) Have subject sign compensation form

- 5) Explain to subject what experiment was about and answer any questions s/he has
- 6) Thank subject and remind him/her not to tell anyone else about the experiment
- 7) Log compensation for both the subject and the charity in the accounting log
- 8) File both consent form and cash receipt
- 9) Record the subject ID on a list of subjects for that day

At end of day:

- 1) Replace batteries in sensors
- 2) Make sure everything is recorded properly in accounting and everything is filed
- 3) Move the subject folders for the subjects that came in that day to one of the following folders:
 - a. REAL_SUBJECTS : for successful experiments
 - b. Do Not Use : for unsuccessful experiments
 - c. Use If Necessary : for experiments where something went somewhat wrong, but the data could be used, if necessary
 - d. Get everything set up for first subject in morning

Appendix D: Screen Shots from Experiment.

Note: All screens after Step 1 were accompanied by recorded verbal instructions.

Step 1: As shown to control subjects. Experimenter remains in the room for this step.

Initial Questions

Please indicate with which hand you use the mouse:

- Left
- Right

Next Page

Step 1: As shown to non-control subjects. Experimenter remains in the room for this step.

Initial Questions

Please indicate with which hand you use the mouse:

- Left
- Right

Depending on your performance, a charity of your choice may receive compensation. Please indicate a favorite charity you'd like to see a donation made toward.

- Action Against Hunger
- American Red Cross
- American Jewish World Service
- AmeriCares
- Asia Foundation
- BAPS Care International
- CARE
- Direct Relief International
- GOAL
- Habitat for Humanity International
- International Federation of Red Cross and Red Crescent
- Islamic Relief Worldwide
- Karuna Trust
- Network for Good
- Oxfam International
- Quarters From Kids: Tsunami Relief and Rebuilding
- Sarvodaya
- Save the Children
- UNICEF
- World Food Programme (UN)
- World Vision
- Other Charity


Next Page

Step 2: As shown to all subjects. Experimenter leaves at the end of this step after sensors are set up to record data. Experimenter also leaves room for privacy if subject prefers.

Sensor Preparation

The experimenter will now ask you to attach sensors. You may go to the bathroom for privacy if you would like.

The sensors should be placed to match the following diagram:



The HandWave device will also be attached by the experimenter to your left hand.

Note: There will be no video, no audio, and no other sensors used in this experiment.

The experimenter should now show you how to use the paper shredder and then leave the room.

If you have any additional questions for the experimenter during the experimenter, please use the instant messaging client to contact him.

[Next Page](#)

Step 3: As shown to *control* subjects

Instructions: Practice Quiz

- You will soon participate in a game that will determine your compensation.
- However, before you begin this game, you will do a **practice quiz** that will not count toward your compensation.
- The quiz will be done on paper, not on the computer, so please locate a pen or pencil now.
- It will contain 30 grids of numbers that look like these 2 grids shown below.

0.23	1.86	7.34	9.96	8.26	4.79
3.40	6.60	1.71	2.81	3.38	9.58
1.07	1.97	5.23	7.72	6.55	0.04
6.01	9.13	0.55	8.84	1.50	9.33

- Each grid will contain exactly one pair of numbers that adds up to 10.
- To be sure you are clear on the instructions, the correct pairs of numbers for these two grids are shown in **bold** below.

0.23	1.86	7.34	9.96	8.26	4.79
3.40	6.60	1.71	2.81	3.38	9.58
1.07	1.97	5.23	7.72	6.55	0.04
6.01	9.13	0.55	8.84	1.50	9.33

- Your task is to find as many of these pairs as possible in 5 minutes.
- When your 5 minutes is up, you will:
 - a) tell the experimenter that you are done through the instant messaging client so he can take your work
 - b) enter the number of pairs that you found on the computer
- Note: The experimenter will verify that the number of pairs you report is accurate.
- Please make sure that you understand these instructions before going on to the next page.

Next Page

Step 3: As shown to *non-control* subjects

Instructions: Practice Quiz

- You will soon participate in a game that will determine your compensation.
- However, before you begin this game, you will do a **practice quiz** that will not count toward your compensation.
- The quiz will be done on paper, not on the computer, so please locate a pen or pencil now.
- It will contain 30 grids of numbers that look like these 2 grids shown below:

0.23	1.86	7.34	9.96	8.26	4.79
3.40	6.60	1.71	2.81	3.38	9.58
1.07	1.97	5.23	7.72	6.55	0.04
6.01	9.13	0.55	8.84	1.50	9.33

- Each grid will contain exactly one pair of numbers that adds up to 10.
- To be sure you are clear on the instructions, the correct pairs of numbers for these two grids are shown in **bold** below:

0.23	1.86	7.34	9.96	8.26	4.79
3.40	6.60	1.71	2.81	3.38	9.58
1.07	1.97	5.23	7.72	6.55	0.04
6.01	9.13	0.55	8.84	1.50	9.33

- Your task is to find as many of these pairs as possible in 5 minutes.
- When your 5 minutes is up, you will:
 - a) shred your paper in the paper shredder
 - b) enter the number of pairs that you found on the computer
- Please make sure that you understand these instructions before going on to the next page.

Next Page

Step 4: As shown to all subjects

Do the Practice Quiz!

- Please open the red folder marked **Practice Quiz** on the shelf in front of you and take out the practice quiz.
- You have exactly 5 minutes to complete this quiz.
- Good Luck!

Step 5: As shown to *control* subjects. Note that the text box at the bottom of the page is so that the experimenter can put in a code to go to the next page. This makes it so that the subject cannot go to the next page before the experimenter takes his/her quiz.

YOUR TIME IS UP!!

Stop writing and tell the experimenter that you are done through the instant messenger.

If you do not tell the experimenter that you are done within 30 seconds, you will be DISQUALIFIED.

Seconds Elapsed Since Alarm Went Off:

For Experimenter: Stop Alarm

Step 5: As shown to non-control subjects

YOUR TIME IS UP!!

Stop writing and click **Stop Alarm**:

Stop Alarm

Step 6: As shown to non-Control subjects (control subjects skip this step)

Please Shred Your Quiz

IMPORTANT: Please shred your quiz in the paper shredder located to your left before going on to the next page.

- To use the paper shredder, simply slip your quiz through the slit at the top.
- If you have any trouble with this, please notify the experimenter over the instant messaging client.
- **Note: Please remember the number of pairs that you found. You will be asked this number on the next page.**

Next Page: I have shredded my quiz

Step 7: As shown to control subjects

Practice Quiz Score

- Please enter the number of pairs found in the practice quiz below:

Number of pairs found in Practice Quiz:

- Note: The experimenter will verify that the number of pairs you report is accurate.

[Next Page](#)

Step 7: As shown to non-control subjects

Practice Quiz Score

- Please enter the number of pairs found in the practice quiz below:

Number of pairs found in Practice Quiz:

- **REMINDER: Please be sure that your quiz has been shredded before going on to the next page.**

[Next Page](#)

Step 8: As shown to control subjects.

Instructions: Real Quiz

Your Task:

- You have now finished your practice quiz and will complete the real quiz that will determine your compensation.
- This quiz will be identical in format to the practice quiz.
- As in the practice quiz, each grid will contain exactly one pair of numbers that add up to exactly 10.
- Your task is to find as many of these pairs as possible in 5 minutes.
- When your 5 minutes is up, you will
 - a) tell the experimenter that you are done through the instant messaging client so he can take your work
 - b) enter the number of pairs that you found on the computer
 - c) answer some final questions about the experiment before your compensation is determined.

- Note: The experimenter will verify that the number of pairs you report is accurate.

[Next Page](#)

Step 8: As shown to non-control subjects.

Instructions: Real Quiz

Your Task:

- You have now finished your practice quiz and will complete the real quiz that will determine your compensation.
- This quiz will be **identical** in format to the practice quiz.
- As in the practice quiz, each grid will contain exactly one pair of numbers that add up to exactly 10.
- Your task is to find as many of these pairs as possible in 5 minutes.
- When your 5 minutes is up, you will:
 - a) shred your paper as you did before
 - b) enter the number of pairs that you found on the computer
 - c) answer some final questions about the experiment before your compensation is determined.

Next Page

Step 9: As shown to *control* and *no_charity* subjects. Note: subjects were told through verbal recording that score would vary linearly with the score and that the 10, 20, and 30 pair levels were simply examples of pot sizes.

Compensation

Your Compensation:

- You will receive at least \$5 for your participation.
- You will also have a 1 in 3 chance at winning a larger pot of money.
- The size of this pot depends on your score on the quiz.
- If you find **10** pairs, the pot will be **\$15** in cash (\$20 in total, including participation payment)
- If you find **20** pairs, the pot will be **\$30** in cash (\$35 in total, including participation payment)
- If you find all **30** pairs, the pot will be **\$45** in cash (\$50 in total, including participation payment)
- In sum, the better your score, the larger the amount of money you can win will be.

- **Please try as hard as you can in this quiz to maximize the size of your pot of money!**

Additional Questions:

- If you have any questions about how you will be compensated, please ask the experimenter now.
- You should not go on to the next page until you are clear on how you will be compensated.

Next Page

Step 9: As shown to *charity_gains* subjects (example: for UNICEF charity choice). Note: subjects were told through verbal recording that score would vary linearly with the score and that the 10, 20, and 30 pair levels were simply examples of pot sizes.

Compensation

Your Compensation:

- You will receive at least \$5 for your participation.
- You will also have a 1 in 3 chance at winning a larger pot of money.
- The size of this pot depends on your score on the quiz.
- If you find 10 pairs, the pot will be \$15 in cash (\$20 in total, including participation payment)
- If you find 20 pairs, the pot will be \$30 in cash (\$35 in total, including participation payment)
- If you find all 30 pairs, the pot will be \$45 in cash (\$50 in total, including participation payment)
- In sum, the better you score, the larger the amount of money you can win will be.

Charity Compensation:

- In addition, UNICEF will have a 1 in 3 chance at winning another pot of money, which is equal in size to your own personal pot.
- So, by scoring well, not only do you help yourself, but you also help UNICEF by making the total amount it can win higher.

- **Please try as hard as you can in this quiz to maximize the size of your pot of money!**

Additional Questions:

- If you have any questions about how you will be compensated, please ask the experimenter now.
- Some subjects have asked if we will really give money to the charity. The answer is: **YES!** We will give real money to the charity you have chosen if we determine that it has won any money.
- You should not go on to the next page until you are clear on how you will be compensated.

[Next Page](#)

Step 9: As shown to *charity_loses* subjects (example: for UNICEF charity choice). Note: subjects were told through verbal recording that score would vary linearly with the score and that the 10, 20, and 30 pair levels were simply examples of pot sizes.

Compensation

Your Compensation:

- You will receive at least \$5 for your participation.
- You will also have a 1 in 3 chance at winning a larger pot of money.
- The size of this pot depends on your score on the quiz.
- If you find 10 pairs, the pot will be \$15 in cash (\$20 in total, including participation payment)
- If you find 20 pairs, the pot will be \$30 in cash (\$35 in total, including participation payment)
- If you find all 30 pairs, the pot will be \$45 in cash (\$50 in total, including participation payment)
- In sum, the better you score, the larger the amount of money you can win will be.

Charity Compensation:

- In addition, UNICEF will have a 1 in 3 chance at winning another pot of money.
- However, since our budget is limited, the larger your pot of money is, the smaller UNICEF's pot of money will be.
- If you do not find any pairs, UNICEF's pot will be \$45.
- If you find 10 pairs, UNICEF's pot will be \$30.
- If you find 20 pairs, UNICEF's pot will be \$15.
- If you find all 30 pairs, UNICEF's pot will be \$0.
- In sum, the better you score, the larger the amount of money you can win will be.
- So, by scoring well, you help yourself, but you hurt UNICEF by making the total amount it can win smaller.

• **Please try as hard as you can in this quiz to maximize the size of your pot of money!**

Additional Questions:

- If you have any questions about how you will be compensated, please ask the experimenter now.
- Some subjects have asked if we will really give money to the charity. The answer is: **YES!** We will give real money to the charity you have chosen if we determine that it has won any money.
- You should not go on to the next page until you are clear on how you will be compensated.

Next Page

Step 10: As shown to all subjects

Complete the Real Quiz!

- Please open the blue folder marked **Real Quiz** in the top drawer of the filing cabinet next to your left leg and take out the real quiz.
- You have exactly 5 minutes to complete this quiz.
- The average score so far on this quiz is: 21.7.
- Good Luck!

Step 11: As shown to *control* subjects. Note that the text box at the bottom of the page is so that the experimenter can put in a code to go to the next page. This makes it so that the subject cannot go to the next page before the experimenter takes his/her quiz.

YOUR TIME IS UP!!

Stop writing and tell the experimenter that you are done through the instant messenger.

If you do not tell the experimenter that you are done within 30 seconds, you will be DISQUALIFIED.

Seconds Elapsed Since Alarm Went Off:

Step 11: As shown to non-control subjects

YOUR TIME IS UP!!

Stop writing and click **Stop Alarm**:

Step 12: As shown to non-control subjects (control subjects skip this step)

Please Shred Your Quiz

IMPORTANT: Please shred your quiz before going on to the next page.

- To use the paper shredder, simply slip your quiz through the slit at the top.
- If you have any trouble with this, please notify the experimenter over the instant messaging client.
- **Note: Please remember the number of pairs that you found. You will be asked this number on the next page.**

Step 12: As shown to control subjects

Your Score

- ◆ Please report the number of pairs that you found below.
- ◆ Note: The experimenter will verify that the number of pairs you report is accurate.

Number of pairs found in this quiz:

- ◆ When you are finished, please click the **Next Page** button.

Step 12: As shown to non-control subjects

Your Score

- ◆ Please report the number of pairs that you found below.

Number of pairs found in this quiz:

- ◆ **REMINDER: Please shred the quiz so that later subjects will not see your answers.**
- ◆ When you are finished, please click the **Next Page** button.

Step 13: As shown to all subjects (for Carson Reynolds' analysis):

Questionnaire

Please fill out the following questionnaire before your compensation is determined.

Each of the following questions asks about the *situation* you encountered during the experiment. *Situation* refers to the methods of observation by the other individual or system with which you interacted.

Do you think the situation is:

Unethical Ethical

No Opinion

Do you think the situation is:

Invasive Respectful

No Opinion

Does the situation make you feel:

Comfortable Uncomfortable

No Opinion

When performing the task was the situation a:

Help Hindrance

No Opinion

Do you think the situation is:

Moral Immoral

No Opinion

Which of the following does the situation make you feel:

Suspicious Trustful

No Opinion

Do you feel the situation is:

Fair Unfair

No Opinion

Given the choice between two situations you'd prefer the situation:

Without Sensors that Collect Information About Emotion With Sensors that Collect Information About Emotion

No Opinion

Were the intentions with which you did the quiz:

Bad Good

No Opinion

[Next Step](#)

Step 14: As shown to *control* subjects. Pot size calculated based on score.

Determining Your Compensation

Thank you for participating!

Based on your score of 2, you now have a chance of winning a pot of money equal to **\$3.00**.

If you have any questions about this, please ask the experimenter over the instant messaging client now.

[Next Page](#)

Step 14: As shown to non-*control* subjects (example: for UNICEF charity choice). Pot sizes calculated based on score and condition.

Determining Your Compensation

Thank you for participating!

Based on your score of 2, you now have a chance of winning a pot of money equal to **\$3.00**.
Also, UNICEF has a chance of winning a pot of money equal to **\$3.00**.

If you have any questions about this, please ask the experimenter over the instant messaging client now.

[Next Page](#)

Step 15: As shown to *control* and *no_charity* subjects.

Compensation: Last Step

Please think about the following questions before clicking on the button below to determine your compensation:

- Are you happy with your performance and the size of your pot?
- Could you have done better?
- Were you **honest** in reporting your score?
- If you have overestimated your score, you have been **dishonest** and have unfairly influenced the size of your pot of money.

When you are ready, please click the button below to determine whether you won the pot of money.

[Determine Compensation](#)

Step 15: As shown to *charity_gains* and *charity_loses* subjects (example: for UNICEF charity choice).

Compensation: Last Step

Please think about the following questions before clicking on the button below to determine your compensation:

- Are you happy with your performance and the size of your pot and UNICEF's pot?
- Could you have done better?
- Were you **honest** in reporting your score?
- If you have overestimated your score, you have been **dishonest** and have unfairly influenced the size of your pot and UNICEF's pot.

When you are ready, please click the button below to determine whether you and UNICEF have won the pots of money.

[Determine Compensation](#)

Step 16: Final step as shown to *control* and *no_charity* conditions. Random number generator used to determine whether subject wins pot.

Compensation

◆ **Your Compensation:**

Unfortunately, you didn't win the pot of money, but you do get the \$5 participation payment.

Thank you for participating! The experiment is now over.

Please let the experimenter know that you are finished and he will give you your compensation.

IMPORTANT: Please do not tell anyone who may participate in this experiment later about the details of this experiment.

Step 16: Final step as shown to *charity_gains* and *charity_loses* conditions (example: for UNICEF charity choice). Random number generator used to determine whether subject wins pot.

Compensation

◆ **Your Compensation:**

Congratulations! You win **\$8.00** including the participation payment!!

◆ **Charity Compensation:**

Unfortunately, **UNICEF** did not win the pot of money.

Thank you for participating! The experiment is now over.

Please let the experimenter know that you are finished and he will give you your compensation.

IMPORTANT: Please do not tell anyone who may participate in this experiment later about the details of this experiment.

Appendix E: Complications with Experiment

It took many pilot subjects to get the experimental setup to a point where subjects would both understand the compensation scheme properly and also be confident that they could cheat and get away with it. Listed below are the representative pieces of feedback we encountered in the pilot studies and the steps we took to eliminate these complications.

- 1) “You can see through the window and determine whether we’re cheating or not.”
 - Have them work in an office without a glass wall.
- 2) “I thought there might be a video camera somewhere in the room to determine whether I’m cheating or not”
 - Explicitly state that there are no video cameras in the room and no other sensors other than the ones placed on the subject’s body used in the experiment
- 3) “The paper shredder doesn’t have cross shreddings so it’s possible to piece together my work.”
 - Get a shredder that does cross shredding
- 4) “I felt like you were monitoring me.”
 - Automated the whole system so that there was no need for the experimenter to come in at any point during the experiment except when subjects in the *control* condition turned in their work. Also introduced an instant messaging client so that subjects could contact the experimenter without the experimenter coming in the room.
- 5) “I didn’t want to hurt the results of your experiment.”
 - Tried not to engage in any small talk or even introduce myself to decrease the probability that the subject would be afraid to “hurt” the results
- 6) “I didn’t know I was supposed to to shred my work.”
 - Introduced a separate step for the sole purpose of shredding the subject’s work. The subject must click on a button that says “Next Step: I’ve shredded my work.”
- 7) “The whole paper shredding thing was weird. I felt like you were trying to get us to cheat, which made me wary of the whole situation.”
 - Came up with a reasonable explanation to tell them: that it was important to shred their quiz so that the next subject would not see their work.
- 8) “I suspected this might be about cheating because the registration page mentioned that this was a study about ethics.”
 - Took out any reference to ethics on the registration page.
- 9) “I suspected this might be about cheating because the URL we used for the study had ‘ethics’ in the name.”
 - Hid the URL on the web browser used in the experiment.

Appendix F: Matlab Code used to implement EM algorithm

This was used to determine percentage of cheaters (in a maximum likelihood sense) given the assumptions outlined in this thesis for each group.

```
clear all
load dataset
% data should be saved in this form:
% group is a cell array with as many elements as there are groups.
% group{1} should be the array of scores for the control group, which must
contain 0% cheaters.
% the rest of the groups should be saved in group{2} and so on.

% This program assumes that both the cheat and non-cheat distributions are
% Gaussian (different means and variances) and estimates the means and
% variances for each distribution as well as the probability of being in
% cheat vs. non-cheat for each group by running 200 iterations of the
% expectation-maximization (EM) algorithm.

lst = [];
for iter = 1:200
    %Initialization of paremeters:
    Mu_C = rand*25; Sigma_C = rand*25;
    Mu_N = rand*25; Sigma_N = rand*25;

    %Prob that they didnt cheat
    %alpha=[1; rand; rand; rand];
    alpha=[1; rand(length(group)-1,1)];

    delta_LLK = 100; LLK = -10000;
    while delta_LLK >= 10^(-5)
        %E-STEP
        old_LLK = LLK;
        LLK = 0;
        for i = 1:length(group)
            X = group{i};
            pr_n = normpdf(X, Mu_N, sqrt(Sigma_N))*alpha(i);
            pr_c = normpdf(X, Mu_C, sqrt(Sigma_C))*(1-alpha(i));

            pr_lambda_0{i} = pr_n ./ (pr_n + pr_c);
            LLK = LLK + sum(log(pr_n + pr_c));
        end
        delta_LLK = abs(LLK - old_LLK);
        %if LLK <= old_LLK
        %    fprintf(1,'Error!!! LLK decreased: %f\n', LLK);
        %end

        %M-STEP
        normalizer_N = 0; normalizer_C = 0;
        sum_Mu_N = 0; sum_Mu_C = 0;
        sum_Sigma_N = 0; sum_Sigma_C = 0;

        for i = 1:length(group)
```

```

        X = group{i};
        sum_Mu_N = sum_Mu_N + sum(pr_lambda_0{i}.*X);
        sum_Mu_C = sum_Mu_C + sum((1-pr_lambda_0{i}).*X);

        normalizer_N = normalizer_N + sum(pr_lambda_0{i});
        normalizer_C = normalizer_C + sum(1-pr_lambda_0{i});
        alpha(i) = sum(pr_lambda_0{i})/length(X);
    end
    Mu_N = sum_Mu_N/normalizer_N;
    Mu_C = sum_Mu_C/normalizer_C;

    for i = 1:length(group)
        X = group{i};
        sum_Sigma_N = sum_Sigma_N + sum(pr_lambda_0{i}.*((X -
Mu_N).^2));
        sum_Sigma_C = sum_Sigma_C + sum((1-pr_lambda_0{i}).*(X -
Mu_C).^2));
    end
    Sigma_N = sum_Sigma_N/normalizer_N;
    Sigma_C = sum_Sigma_C/normalizer_C;
end
fprintf(1, '\n');

lst = [lst; LLK Mu_N Sigma_N Mu_C Sigma_C alpha'];
end

[v, i] = max(lst(:,1));
lst(i,:);
Mu_N = lst(i,2);
Sigma_N = lst(i,3);
Mu_C = lst(i,4);
Sigma_C = lst(i,5);
alpha = lst(i,6:end)';

LLK = 0;
for i = 1:length(group)
    X = group{i};
    pr_n = normpdf(X, Mu_N, sqrt(Sigma_N))*alpha(i);
    pr_c = normpdf(X, Mu_C, sqrt(Sigma_C))*(1-alpha(i));

    pr_lambda_0{i} = pr_n ./ (pr_n + pr_c);
    prob_cheat_individual_subjects{i} = 1-pr_lambda_0{i};
    LLK = LLK + sum(log(pr_n + pr_c));
end
fprintf(1, '%f\n', LLK);

percent_cheat_by_group = (1-alpha)*100
honest_mean = Mu_N
honest_stdev = sqrt(Sigma_N)
cheat_mean = Mu_C
cheat_stdev = sqrt(Sigma_C)

```

Appendix G: Matlab Code used to extract skin conductivity responses (SCRs)

Adapted from Jennifer Healey's code.

```
function [ startles_beg, startles_end ] = msdetect( s, thresh, gapthresh,
t, Fs )
% MSDETECT
% @param s signal
% @param thresh optional threshold to define significant startle (default
to 0.003)
% @param gapthresh optional threshold to define significant gaps --
second-order diff (default to 31)
% @param t optional time
% @param Fs optional sampling frequency (2*pi/mean dt by default)
% @return vectors startles_beg and startles_end containing the indices of
the beginning and end of the startles

if ~exist('Fs')
    if exist('t')
        Fs = mean(diff(t));
    end
end

if ~exist('thresh')
    thresh = 0.003;
end

if ~exist('gapthresh')
    gapthresh = 31;
end

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%%Filter out the high frequency noise%%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
if exist('Fs')
    [b,a] = ellip(4, 0.1, 40, 4*2/Fs);
    sf = filter(b,a,s);
    sf_diff = diff(sf);
else
    sf_diff = diff(s);
end

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%%Find Significant Startles%%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
overthresh = find(sf_diff>thresh);
gaps = diff(overthresh);

% big_gaps contain the indices of beginning of startles
big_gaps = find(gaps>gapthresh);

iend = overthresh(big_gaps);
```

```

ibegin = iend+1;

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%%Fine Tuning%%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% find zero crossings
overzero = find(sf_diff>0);
zerogaps = diff(overzero);
z_gaps = find(zerogaps>1);

idown = overzero(z_gaps);
iup = idown+1;

% find the crossing closest to ibegin
new_begin = [];
for i=1:length(ibegin),
    temp=find(iup<ibegin(i));
    ltemp = length(temp);
    if ltemp>0
        choice=temp(ltemp);
        new_begin = [new_begin iup(choice)];
    end
end

n = length(new_begin);
% find the end of the startle as the maximum between
% startle beginnings
new_end = zeros(1, n);
for i=1:n,
    startit = new_begin(i);
    if i<n
        endit = new_begin(i+1);
    else
        endit = length(s);
    end
    [val,loc] = max(s(startit:endit));
    new_end(i) = startit+loc;
end

startles_beg = new_begin;
startles_end = new_end;

```

Appendix H: Inter-beat Interval (IBI) extraction code

This was written in Delphi by Vadim Gerasimov

```
// (C) MIT, Media Laboratory
// Vadim Gerasimov

unit Unit1;

interface

uses
  Windows, Messages, SysUtils, Classes, Graphics, Controls, Forms, Dialogs,
  StdCtrls, Menus, DateUtils;

const
  TimeStamp = 0;
  DataStamp = 1;
  SamplingRate = 150;
  PeakScale = 0.67; // at least 2/3 of the previous spike
  MaxPulse = 200; // maximum expected heart rate
  MinPulse = 30; // minimum expected heart rate
  DecreaseRatio = 1.5; // maximum permitted pulse rate decrease
  IncreaseRatio = 1.5; // maximum permitted pulse rate increase
  NSpikeSamples = SamplingRate div 30; // 8 samples, ~33 ms slope
  CR_LF = #13#10;

type
  TForm1 = class(TForm)
    TimeStamps: TListBox;
    OpenDataDialog: TOpenDialog;
    LoadFile: TButton;
    SaveDialog1: TSaveDialog;
    procedure ConvertFileClick(Sender: TObject);
  private
    { Private declarations }
    procedure ConvertData (const FileName: string);
  public
    { Public declarations }
  end;
  EDataError = class(Exception);
  TRWaveDetector = class
    function Spike (Threshold: Integer): Boolean;
    function GetSlopeToPeak: Integer;
    procedure Update (Sample: Integer);
    constructor Create;
  private
    Samples: array [0..NSpikeSamples-1] of Integer;
  end;

var
  Form1: TForm1;
```

implementation

```
{ $R *.DFM }
```

```
procedure TForm1.ConvertData (const FileName: string);
```

```
var
```

```
  IBIData: TStringStream;  
  IBIFile: TFileStream;  
  i: Integer;  
  Pulse, LastPulse: Real;  
  SampleTime, LastRTime: Cardinal;  
  LastR: Cardinal;  
  DetectR: TRWaveDetector;  
  TextFile: TStringList;  
  Sample: Integer;  
  IBI: Integer;
```

```
begin
```

```
  DetectR:=TRWaveDetector.Create;  
  IBIData:=TStringStream.Create('');  
  TextFile:=TStringList.Create;  
  TextFile.LoadFromFile(FileName);
```

```
  try
```

```
    SampleTime:=0;  
    LastRTime:=0;  
    for i:=0 to TextFile.Count-1 do  
      begin
```

```
        Sample:=Round(StrToFloat(TextFile.Strings[i]));  
        DetectR.Update(Sample); //shifts prior values in buffer to make room
```

```
      at end
```

```
        if (SampleTime-LastRTime>4*SamplingRate) then LastR:=0;  
        // Sets LastR to zero if no spikes detected for 4 seconds  
        if (SampleTime>LastRTime+NSpikeSamples) and  
        // if buffer holds new samples since last R-wave and  
        DetectR.Spike(Round(LastR*PeakScale)) then  
        // if the prior sample was a peak then
```

```
        begin
```

```
          IBI:=SampleTime-LastRTime; // Current IBI  
          IBIData.WriteString(IntToStr(IBI)+CR_LF);  
          LastPulse:=(60/(SampleTime-LastRTime))*SamplingRate;  
          // calculate LastPulse = pulse rate in beats per minute  
          LastRTime:=SampleTime; // replace with current R-wave time  
          LastR:=DetectR.GetSlopeToPeak; // save slope to current R-wave
```

```
        peak
```

```
          if (SampleTime > 3*SamplingRate) and  
          // if running for at least 3 seconds and  
          (LastPulse<MaxPulse) and (LastPulse>MinPulse) and  
          // LastPulse is within limits and  
          ((LastPulse<Pulse) and (Pulse<DecreaseRatio*LastPulse)) or  
          // LastPulse was slower than current pulse, but not by too much or  
          ((LastPulse>Pulse) and (LastPulse<IncreaseRatio*Pulse))  
          // Last Pulse was faster than current pulse, but not by too much  
          then Pulse:=LastPulse; // reset pulse if conditions satisfied
```

```
        end;
```

```
        Inc(SampleTime); // increment sample count
```

```
      end;
```

```
  finally
```

```

    TextFile.Free;
    DetectR.Free;
end;
if SaveDialog1.Execute then
begin
    IBIFile:=TFileStream.Create(SaveDialog1.FileName, fmCreate);
    try
        IBIData.Seek(0, soFromBeginning);
        IBIFile.CopyFrom(IBIData, IBIData.Size);
    finally
        IBIFile.Free;
    end;
end;
IBIData.Free;
end;

procedure TForm1.ConvertFileClick(Sender: TObject);
begin
    if not OpenDataDialog.Execute then Exit;
    ConvertData(OpenDataDialog.FileName);
end;

{ TSpikeDetector }

constructor TRWaveDetector.Create;
var
    i: Integer;
begin
    for i:=0 to NSpikeSamples-1 do Samples[i]:=0;
end;

function TRWaveDetector.GetSlopeToPeak: Integer;
begin
    Result:=Samples[NSpikeSamples-2]-Samples[0];
    // Samples[NSpikeSamples-2] holds peak; Samples[NSpikeSamples-1] = 1st
dip
end;

function TRWaveDetector.Spike (Threshold: Integer): Boolean;
var
    i: Integer;
begin
    Result:=False;
    if Samples[NSpikeSamples-2]=Samples[0] then Exit; // no increase over
base
    for i:=0 to NSpikeSamples-3 do
        if Samples[i]>Samples[NSpikeSamples-2] then Exit;
        // returns false if any prior sample before possible peak > than
possible peak
    if (Samples[NSpikeSamples-1]>=Samples[NSpikeSamples-2]) then Exit;
    // returns false if current sample not less the predecessor (possible
peak)
    if (Samples[NSpikeSamples-2]-Samples[0]) < Threshold then Exit;
    // returns false if slope of possible peak is less than
// threshold (67% of slope to last detected peak)
    Result:=True;
end;

```



```
procedure TRWaveDetector.Update(Sample: Integer);
var
  i: Integer;
begin
  for i:=0 to NSpikeSamples-2 do Samples[i]:=Samples[i+1];
  // shift buffer contents one position to make room at end of buffer
  Samples[NSpikeSamples-1]:=Sample;
  // put incoming sample at end of buffer
end;

end.
```

Appendix I: Spectral analysis code

This was written in Matlab by Yuan Qi

Example of how to run:

```
[amp, phs, vtar, xa] = pkalman(x,'Ns',128,'Fs',fs, 'ts', 700, 'smoothing',0,  
'LogScale',0,'disp',1);
```

```
function varargout = pkalman(x, varargin)
% PKALMAN Power Spectral Density (PSD) estimate via Kalman filtering and
smoothing.
% INPUT
% x: the input signal, which is a column vector
%
% The following optional input arguments can be specified in the form of
name/value pairs:
% [default value in brackets]
% Ns: the frequency resolution. If it is in the input argument list, Ns
must
%     show before Fs. [Ns = 64]
% Fs: Fs can be either a scalar or a row vector.
%     (scalar) the sampling frequency Fs in Hz, or twice of the frequency
range of a
%     given signal. [Fs=6 (Hz)]
%     (vector) the given frequencies bases for spectrum estimation
% t: the sampling time. [(1:n)/(Fs*2)]
% sigma: the process noise variance [1e1*eye(2*length(Fs)+1)]
% lambda: the observation noise variance [1e-2]
% m0: the mean of the prior distribution [ones(2*length(Fs)+1,1)]
% v0: the variance of the prior distribution [1e2*eye(2*length(Fs)+1,1)]
% disp: 0-no display; 1-displaying the result
% fig_sz: display size in inch [matlab default size]
% color: 0-gray display 1-color display [1];
% LogScale: 0-not plotting the image in log scale 1-plotting the image in
log scale.
% ts: the index of one time slice. If it is a positive integer, then
plot the
%     spectrum at that time slice. if it is an empty matrix, then plot
the
%     whole spectrogram. [ [] ]
% tp: the start point in time for displaying a spectrogram
% te: the end point in time for displaying a spectrogram
% smoothing: 0- no smoothing, 1 for smoothing
% axis: matlab display argument [[]]
% XTICK: matlab display argument [[]]
% YTICK: matlab display argument [[]]
% FontSize: matlab display argument [[]]
%
% OUTPUT
% varargout = [amp, phs, vtar]
% amp: estimated frequency amplitudes from time 1 to n
% phs: estimated frequency phases from time 1 to n
```

```

% vtar: the variances from last block's smoothing; it's an empty
%       matrix when we only do Kalman filtering
%
% EXAMPLE 1: Estimate the spectrum of a signal with fast decaying amplitude
and 10 percent missing data
% fs = 500;
% tt = 0:(1/fs):2;
% perc = 0.9;
% tr = randperm(length(tt));
% tc = sort(tr(1:floor(length(tt)*perc))); % truncated index of tt
% t = tt(tc);
% w = 50;
% aplmd = exp(-t*1);
% x = (sin(2*pi*t*w).*aplmd)';
% m = pkalman(x,'t',t,'Ns',50,'Fs',200,'smoothing',1,
'LogScale',0,'disp',1);
% EXAMPLE 2: Estimate the spectrum of a signal consisting of three
frequency components which are close to each other
% fs = 50;
% t = 0:(1/fs):3;
% y = sin(2*pi*21*t+20) + 1*sin(2*pi*20*t+20)+.5*sin(2*pi*19*t+20);
% nslevel = sqrt(1e-1);
% x=y' + nslevel*randn(size(y))';
% pkalman(x,'Ns',64,'Fs',fs,'ts',length(x),'axis',...
% [0 25 -55 10],'XTICK',[0 5 10 15 19 20 21 25], 'YTICK',[-40 -20
0],'FontSize',6)
%
% References
% [1] Bayesian Spectrum Estimation of Unevenly Sampled Nonstationary
Data",
% Yuan Qi, Thomas P. Minka, and Rosalind W. Picard, MIT Media Lab
% Technical Report Vismod-TR-556, 2002
% http://web.media.mit.edu/~yuanqi/papers/qi-spec-02.pdf
% [2] A shorter version in ICASSP 02, May 2002, Orlando, Florida
% http://web.media.mit.edu/~yuanqi/papers/qi-icassp02.pdf
%
% Author: Yuan Qi
% MIT Media lab
% Date: April 12, 2003

x = x(:);
n = length(x);

% parsing the parameters
[t, Fs, Ns, smoothing, n_block, n_overlap, sigma, lambda, m0,v0, disp_flg,
sparse] = parse_par(n,varargin{:});

% estimate the posterior mean by Kalman filtering/smoothing
[m, vtar] = kalman_est(x, t, Fs, Ns, smoothing, n_block, n_overlap, sigma,
lambda, m0,v0, sparse);

ds = length(Fs);
amp = [];
for j = 1:ds
    rowwd = [j j+ds];
    % amp: amplitude estimate
    amp(j,:) = sqrt(sum(m(rowwd, :).^2));

```

```

end
amp = ([abs(m(end,:)); amp]);
if nargout > 0
    varargout{1} = amp;
end
if nargout > 1
    phs = [];
    for j = 1:ds
        rowwd = [j j+ds];
        % phs: phase estimate
        phs(j,:) = atan(m(j, :)./m(j+ds,:));
    end
    phs = ([zeros(1,n); phs]);
    varargout{2} = phs;
end
if nargout > 2
    varargout{3} = vtar;
end

if nargout == 0 | disp_flg == 1
    disp_spec(amp, t, Fs, varargin{:});
end

function [t, Fs, Ns, smoothing, n_block, n_overlap, ...
    sigma, lambda, m0,v0, disp_flg, sparse] ...
    = parse_par(n, varargin)
% PARSE_PAR parses the parameters

% Set default values for the parameters: t, Fs, Ns
t = [];
Ns = 64;
Fs = 6;
spc = Fs/2 /Ns;
Fs = (0+spc):spc:(Fs/2-spc);
% Parse the list of the input arguments: t, Fs, Ns
args = varargin;
nargs = length(args);
for i=1:2:nargs
    switch args{i},
        case 't', t = args{i+1};
        case 'Ns', Ns = args{i+1};
        case 'Fs',
            if length(args{i+1})==1
                Ftp = args{i+1}/2;
                spc = Ftp/Ns;
                Fs = (0+spc):spc:(Ftp-spc);
            else
                Ftp = [];
                Fs = args{i+1}; % In this case, Ns is not used.
                Fs = Fs(:)';
            end
    end
end
% set value for t if t is not specified by input arguments
if isempty(t)
    t = (1:n)/(Ftp*2);
end

```

```

% Set default values for the parameters:sparse, smoothing, sigma, lambda,
m0,v0, and disp_flg
sparse = 0;
smoothing = 0;
n_block = 1000;
n_overlap = 100;
n_bs = 2*length(Fs)+1;
sigma = 1*1e1*eye(n_bs);
lambda = 1e-2;
m0 = ones(n_bs,1);
v0 = 1e2*eye(n_bs);
disp_flg = 0;
% Parse the list of the input arguments:smoothing, sigma, lambda, m0,v0
for i=1:2:nargs
    switch args{i},
        case 'sparse', sparse = args{i+1};
            case 'smoothing', smoothing = args{i+1};
            case 'n_block', n_block = args{i+1};
            case 'n_overlap', n_overlap = args{i+1};
            case 'var_pro_noise', sigma = args{i+1};
            case 'var_obs_noise', lambda = args{i+1};
            case 'mean_prior', m0 = args{i+1};
            case 'var_prior', V0 = args{i+1};
            case 'disp', disp_flg = args{i+1};
            case 't', ;
            case 'Ns',;
            case 'Fs',;
            case 'fig_sz',;
            case 'tp', ;
            case 'te', ;
            case 'ts', ;
            case 'color',;
            case 'axis', ;
            case 'XTICK',;
            case 'YTICK', ;
            case 'FontSize',;
            case 'LogScale', ;
            otherwise,
                error(['invalid argument name: ' args{i}]);
    end
end

function disp_spec(amp, t, Fs, varargin)

% default value
tp =1;
te = length(t);
ts = []; % the index of one time slice
fig_sz = []; % fig_sz = [2.1 2.1];
color = 1;
log_scale = 1;
ax = []; xtick = []; ytick = []; font_sz = [];

% Parse the list of the input arguments
args = varargin;
nargs = length(args);

```

```

for i=1:2:nargs
    switch args{i},
        case 'fig_sz', fig_sz = args{i+1};
        case 'tp', tp = args{i+1};
        case 'te', te = args{i+1};
        case 'ts', ts = args{i+1};
        case 'color', color = args{i+1};
        case 'axis', ax = args{i+1};
        case 'XTICK', xtick = args{i+1};
        case 'YTICK', ytick = args{i+1};
        case 'FontSize', font_sz = args{i+1};
        case 'LogScale', log_scale = args{i+1};
    end
end

xa=[0 Fs];
figure, %
if ~isempty(fig_sz)
    figure(figsize(fig_sz));
end
if isempty(ts)
    if log_scale == 1
        imagesc(t(tp:te),xa,(20*log10(amp(:,tp:te)))),axis xy;
    else
        imagesc(t(tp:te),xa, amp(:,tp:te) ),axis xy;
    end
    hy = ylabel('Frequency (Hz)');
    if ~isempty(font_sz), set(hy,'FontSize', font_sz), end
    hx = xlabel('Time (Second)');
    if ~isempty(font_sz), set(hx,'FontSize', font_sz), end
else
    if log_scale == 1
        plot(xa, 20*log10(abs(amp(:,ts))));
        hy = ylabel('Amplitude (dB)');
    else
        plot(xa, abs(amp(:,ts)));
        hy = ylabel('Amplitude');
    end
    if ~isempty(font_sz), set(hy,'FontSize', font_sz), end
    hx = xlabel('Frequency (Hz)');
    if ~isempty(font_sz), set(hx,'FontSize', font_sz), end
end
colormap(jet);
grid on
if color == 0
    load grayflp;
    colormap(grayflp);
end

if ~isempty(ax), axis(ax), end
if ~isempty(xtick), set(gca,'XTick', xtick), end
if ~isempty(ytick), set(gca,'YTick', ytick), end
if ~isempty(font_sz), set(gca,'FontSize', font_sz), end

function [m, vtar] = kalman_est(x, t, Fs, Ns, smoothing, n_block,
n_overlap, sigma, lambda, m0,v0, sparse)
% KALMAN_EST is called by PKALMAN

```

```

% Yuan Qi
% April 2003

x = x(:);
n = length(x);
twopi = 2*pi;
m = zeros(length(m0), n);
C = [sin(twopi*Fs*t(1)) cos(twopi*Fs*t(1)) 1];

v0C = v0 * C';
k1 = v0C * inv( C* v0C + lambda);
vt = v0 - k1*v0C';
m(:,1) = m0 + k1*(x(1) - C*m0);

vtar = [];
if smoothing ==1
    % preparing for smoothing
    vtar = [];
    n_block = min(n,n_block);
    vtar = zeros([size(v0), n_block]);
    ptar = zeros([size(v0), n_block - 1]);
    vtar(:, :, 1) = vt;
end

mdi = 1;
i_old = 0;
for i = 2:n
    inttp = t(i)-t(i-1);
    % update the state noise variance
    sigma_t = sigma * (inttp);
    % if mod(i,150)==0, vt = v0; end
    pt = vt + sigma_t;

    C = [sin(twopi*Fs*t(i)) cos(twopi*Fs*t(i)) 1];
    ptC = pt * C';
    kt = ptC * (1./( C * ptC + lambda));
    vtp = pt-kt*(C*pt);
    mtp = m(:,i-1) + kt*(x(i)- C*m(:,i-1));

    if sparse == 1 & mod(i, 100) == 0 & i > 50 % used for obs4 % if mod(i,
40) == 0 & i > 50 % used for obssp
        % OBS pruning
        [m(:,i) vt] = obs4(mtp,vtp,v0);
        % vt = v0;
    else
        m(:,i) = mtp;
        vt = vtp;
    end

    if smoothing == 1
        mdi = mdi + 1;
        if mdi == 0
            nn = i;
            ptar(:, :, n_block-1) = pt;
            vtar(:, :, n_block) = vt;
        elseif mdi == 1
            vtar(:, :, 1) = vt;
        end
    end
end

```

```

elseif mdi >= 2
    vtar(:,:,mdi) = vt;
    ptar(:,:,mdi-1) = pt;
end
if (mdi == n_block) | (i == n)
    ind = i+1 - (n_block:-1:1);
    i_old = i;
    for j = mdi-1:-1:1
        jt = vtar(:,:,j)*inv( ptar(:,:,j)); % ss+le-
4*eye(size(ptar(:,:,j))) );
        m(:,ind(j)) = m(:,ind(j)) + jt * (m(:,ind(j)+1) - m(:,ind(j)));
        vtar(:,:,j) = vtar(:,:,j) + jt * (vtar(:,:,j+1) -
ptar(:,:,j))*jt';
    end

    vtar(:,:,1) = vtar(:,:,n_block - n_overlap + 1);
    for g = 2:n_overlap
        vtar(:,:,g) = vtar(:,:,n_block - n_overlap + g);
        ptar(:,:,g-1) = ptar(:,:,n_block - n_overlap + g - 1);
    end
    mdi = n_overlap;
end
end
end
end

```

```

function [w] = IBithresh(z)
%% This function returns a thresholded IBI in bpm where 50 < HR < 100

%% throw away any value less than 50 bpm
% m=find(z*60 <= 52); % z<=0.667
% m = find(z <= 55); %% Already in bmp
%z(m)=82;
%w=z;

%% -- hiahn
if ( z(1) < 600 ) z(1) = 601; % z in millisecond, 60000/600 = 100
end

if ( z(1) >= 1200) z(1) = 1199; % 60000/1200 = 50
end
%% --

n=1;
while (sum(n) ~= 0)
%if 1
    % m=find(z*60 >= 150); % z>=3.34
    n = find(z < 600);
    z(n)=z(n-1);
    w=z;
end

%% throw away any value greater than 100
m=1;
while (sum(m) ~= 0)
%if 1

```



```
% m=find(z*60 >= 150); % z>=3.34
m = find(z >=1200);
z(m)=z(m-1);
w=z;
end
```

Appendix J: Consent Form

CONSENT TO PARTICIPATE IN NON-BIOMEDICAL RESEARCH

You are asked to participate in a research study conducted by Rosalind Picard, Philip Davis, and Carson Reynolds from the Media Laboratory at the Massachusetts Institute of Technology (M.I.T.). The results of this experiment will be contributed to Philip Davis' and Carson Reynolds' Master's and Doctoral theses, respectively. You should read the information below, and ask questions about anything you do not understand, before deciding whether or not to participate.

• PARTICIPATION AND WITHDRAWAL

Your participation in this study is completely voluntary and you are free to choose whether to be in it or not. If you choose to be in this study, you may subsequently withdraw from it at any time without penalty or consequences of any kind. The investigator may withdraw you from this research if circumstances arise which warrant doing so.

• PURPOSE OF THE STUDY

The purpose of this study is to develop methods to improve software and Web site interfaces.

• PROCEDURES

If you volunteer to participate in this study, we would ask you to do the following things:

In a few minutes, you will be asked to participate in a task on a personal computer. Videotapes, audiotapes, mouse input, keyboard input, force pressure profiles, and biosignals may be collected of your participation and interview. Collection of heart rate through electrocardiograms (EKG), muscle movement through electromyograms (EMG), temperature, and mental workload through Electrodermal Response (EDR), will require the use of sensing electrodes. Sensing respiration rates, blood pressure levels, and temperature may involve sensors that need to be strapped on to your body. This data will be used for experimental purposes only, and after the data collection is over, they will be permanently stored in a private archive. In the future, portions of this record may be published or presented in scientific journals, scientific conference proceedings or book chapters. Further, no other information (such as the subject's name, address or other private information) will be included in these publications. Apart from this possible usage, such data will only be viewed/used for experimental purposes. Again, at any time during or after the experiment you may request that your records be destroyed.

Periodically you will be asked to fill out questionnaires about your experience. Afterwards you may be asked to participate in a brief interview. The entire process should take approximately one hour. If at any time you are uncomfortable, either physically, or with what you are being asked to do, you are free to ask that the experiment be suspended.

Any responses that are collected during the experiment will be completely anonymous. From this point forward, you will be referred to only as the ID number that appears on the upper right corner of this packet.

- **POTENTIAL RISKS AND DISCOMFORTS**

There are very few anticipated risks and discomforts associated with participation in this experiment. Subjects may experience mild psychological stress as they participate in the experiment. However, the level of stress should not exceed anything that they would encounter in everyday computing use. Since the EKG procedure only monitors the electrical impulses and does not emit electricity, there is no risk of shock. As to EMG, there are no risks associated with nerve conduction studies. Nothing is inserted into the skin, so there is no risk of infection. The voltage of electrical pulses is not high enough to cause an injury or to be felt.

- **POTENTIAL BENEFITS**

You will receive at least \$5 and up to \$50 for your participation and will have the opportunity to help improve state-of-the-art technology that we are developing.

- **PAYMENT FOR PARTICIPATION**

Up to \$50 in the form of cash or a voucher for payment will be given to you following your participation in this study. Payment will be prorated for early withdrawal.

- **CONFIDENTIALITY**

Any information that is obtained in connection with this study and that can be identified with you will remain confidential and will be disclosed only with your permission or as required by law.

Video and audio recordings will be captured, along with mouse and keyboard input. This data will be preserved and analyzed but will be anonymous and labeled with a subject number. Additionally, any data attained from the biosensors in this experiment and your use of the mouse will be recorded and labeled in the same manner. After the experiment is complete, you will be referred to by an ID number instead of your name. No one will have access to your personal information.

These records will be stored on password-protected computers at the MIT Media Laboratory. The data obtained from this experiment will be analyzed and published in scientific journals. No personally identifying information will be included in those publications. Records will be destroyed after data analysis is completed.

- **IDENTIFICATION OF INVESTIGATORS**

If you have any questions or concerns about the research, please feel free to contact:

Philip Davis (Co-Investigator): (617) 253-6341 or pcdavis@media.mit.edu
Carson Reynolds (Co-Investigator): (617) 253-8628 or carsonr@media.mit.edu
Rosalind Picard (Principal Investigator): (617) 253-0611 or picard@media.mit.edu

- **EMERGENCY CARE AND COMPENSATION FOR INJURY**

“In the unlikely event of physical injury resulting from participation in this research you may receive medical treatment from the M.I.T. Medical Department, including emergency treatment and follow-up care as needed. Your insurance carrier may be billed for the cost of such treatment. M.I.T. does not provide any other form of compensation for injury. Moreover, in either providing or making such medical care available it does not imply the injury is the fault of the investigator. Further information may be obtained by calling the MIT Insurance and Legal Affairs Office at 1-617-253 2822.”

- **RIGHTS OF RESEARCH SUBJECTS**

You are not waiving any legal claims, rights or remedies because of your participation in this research study. If you feel you have been treated unfairly, or you have questions regarding your rights as a research subject, you may contact the Chairman of the Committee on the Use of Humans as Experimental Subjects, M.I.T., Room E32-335, 77 Massachusetts Ave, Cambridge, MA 02139, phone 1-617-253 6787.

SIGNATURE OF RESEARCH SUBJECT OR LEGAL REPRESENTATIVE

I understand the procedures described above. My questions have been answered to my satisfaction, and I agree to participate in this study. I have been given a copy of this form.

Name of Subject

Name of Legal Representative (if applicable)

Signature of Subject or Legal Representative

Date

AGREEMENT TO INCLUDE DATA IN DE-IDENTIFIED CORPUS

In addition, I agree to allow the data collected in this experiment to be used in a larger corpus, which may be analyzed by other investigators in the future. There will be no identifying information in this corpus, such as name, social security number, or photographs/videotape of subjects' faces. Permission to use your data in this corpus is not a prerequisite for participating in this experiment.

Name of Subject

Signature of Subject

Date

SIGNATURE OF INVESTIGATOR

In my judgment the subject is voluntarily and knowingly giving informed consent and possesses the legal capacity to give informed consent to participate in this research study.

Signature of Investigator

Date

Appendix L: Screen shots from craigslist surveys

Volunteers filled out one of the following surveys, chosen randomly, through a posting on craigslist.com. We solicited volunteers from the craigslist communities in New York City, Chicago, Los Angeles, and San Francisco.

Written below is an example of a posting to one of these communities. Postings were slightly different from one another due to a requirement on craigslist that no two postings be identical.

Example:

Hi,

Do you have ONE MINUTE to spare? If so, please help a Master's student out by filling out a one question survey! You can find the survey here:

http://arsenal.media.mit.edu/cgi-bin/study/pcd_survey/data2.py

Thank you!

The surveys that were shown to subjects are shown on the following pages.

Survey to estimate guilt felt when subjects cheat in *no_charity* condition.

Instructions: Please report how you are most likely to feel in the situation described below:

Imagine that you are participating in a game in which you can win up to \$50 based on your performance.

Your earnings are high when you score well and they are low when you score poorly.

Now imagine that you recognize an opportunity to cheat without anyone knowing, which would increase your earnings.

You choose to cheat in the game.

Your compensation is thus higher than it would have been had you been honest.

In this situation, you would likely feel:

Not Guilty Guilty

Survey to estimate guilt felt when subjects cheat in *charity_gains* condition.

Instructions: Please report how you are most likely to feel in the situation described below.

Imagine that you are participating in a game in which you can win up to \$50 based on your performance.

Your earnings are high when you score well and they are low when you score poorly.

In addition, your favorite charity will receive compensation that is high when you score well and low when you score poorly.

So, by scoring well, you help both yourself and your favorite charity.

Now imagine that you recognize an opportunity to cheat without anyone knowing, which would increase both your earnings and the charity's earnings.

You choose to cheat in the experiment.

Both your compensation and the charity's compensation are thus higher than they would have been had you been honest.

In this situation, you would likely feel:

Not Guilty Guilty

Survey to estimate guilt felt when subjects cheat in *charity_loses* condition.

Instructions: Please report how you are most likely to feel in the situation described below:

Imagine that you are participating in a game in which you can win up to \$50 based on your performance.

Your earnings are high when you score well and they are low when you score poorly.

In addition, your favorite charity will receive compensation, but due to budget constraints, the charity's compensation decreases as you score higher.

So, by scoring well, you help yourself, but you hurt your favorite charity

Now imagine that you recognize an opportunity to cheat without anyone knowing, which would increase your earnings, but lower the charity's earnings.

You choose to cheat in the experiment.

Thus, your compensation is higher than it would have been had you been honest, but the charity's compensation is lower than it would have been had you been honest.

In this situation, you would likely feel:

Not Guilty Guilty

Note: To estimate expected guilt felt while being *honest*, subjects were also presented with one of three surveys that were identical to the above three surveys except that they state "You choose not to cheat in the experiment."

Survey to estimate expected cheating behavior for *no_charity* condition

Instructions: Please answer these two questions truthfully:

Imagine that you are participating in a game in which you can win up to \$50 based on your performance.

Your earnings are high when you score well and they are low when you score poorly.

Now imagine that you recognize an opportunity to cheat without anyone knowing, which would increase your earnings.

How likely are you to cheat in this situation?

Definitely not Definitely

Submit

Survey to estimate expected cheating behavior for *charity_gains* condition

Instructions: Please answer these two questions truthfully:

Imagine that you are participating in a game in which you can win up to \$50 based on your performance.

Your earnings are high when you score well and they are low when you score poorly.

In addition, your favorite charity will receive compensation that is high when you score well and low when you score poorly.

So, by scoring well, you help both yourself and your favorite charity.

Now imagine that you recognize an opportunity to cheat without anyone knowing, which would increase both your earnings and the charity's earnings.

How likely are you to cheat in this situation?

Definitely not Definitely

Submit

Survey to estimate expected cheating behavior for *charity_loses* condition

Instructions: Please answer these two questions truthfully:

Imagine that you are participating in a game in which you can win up to \$50 based on your performance.

Your earnings are high when you score well and they are low when you score poorly.

In addition, your favorite charity will receive compensation, but due to budget constraints, the charity's compensation decreases as you score higher.

So, by scoring well, you help yourself, but you hurt your favorite charity

Now imagine that you recognize an opportunity to cheat without anyone knowing, which would increase your earnings, but lower the charity's earnings.

How likely are you to cheat in this situation?

Definitely not Definitely

Appendix M: Quizzes

PRACTICE QUIZ

2.08	1.76	9.95
8.08	9.13	6.02
4.68	2.94	1.16
4.07	7.78	5.32

Got it!

9.32	7.85	7.90
9.76	7.53	9.07
0.24	6.33	9.66
9.57	8.23	2.04

Got it!

0.41	2.71	2.02
8.08	7.98	0.94
2.33	8.41	7.85
7.95	6.84	9.60

Got it!

0.52	5.07	9.48
7.98	4.33	7.99
5.37	1.93	7.27
7.14	0.44	3.38

Got it!

1.31	0.07	8.01
9.11	4.40	4.02
3.79	5.60	7.58
3.14	1.96	6.95

Got it!

4.28	8.53	7.82
9.43	0.57	3.55
2.59	5.36	3.79
8.01	1.54	8.71

Got it!

6.41	2.19	9.69
3.52	0.26	2.24
6.13	9.80	4.89
1.07	5.44	8.93

Got it!

0.93	3.77	6.54
8.94	7.51	9.29
9.07	6.04	1.71
6.07	7.84	9.31

Got it!

2.83	4.65	5.64
6.17	3.49	6.97
3.03	2.23	9.63
2.90	0.96	4.15

Got it!

2.70	6.59	0.40
6.58	6.81	8.83
0.45	2.04	6.69
1.17	8.02	8.39

Got it!

1.01	0.74	4.21
0.15	8.98	1.94
4.88	8.07	1.02
5.97	3.12	1.75

Got it!

5.18	4.82	2.91
7.91	4.75	2.82
3.10	7.00	4.73
0.29	9.04	7.49

Got it!

1.48	6.55	4.85
3.19	3.45	8.79
1.94	1.66	6.38
6.60	5.12	6.21

Got it!

0.45	1.84	8.65
6.59	8.25	2.21
9.27	7.42	4.42
3.04	6.96	2.04

Got it!

3.05	0.16	8.72
1.99	8.36	1.62
1.33	9.84	6.44
7.17	2.11	0.66

Got it!

2.50	7.61	2.10
5.49	3.46	0.01
7.48	7.06	2.52
2.62	1.83	9.85

Got it!

7.35	3.35	3.82
7.73	6.66	4.38
7.68	1.96	5.37
3.74	8.04	0.48

Got it!

1.48	2.68	8.33
3.73	4.92	9.10
2.54	4.09	3.24
7.69	2.24	1.67

Got it!

1.05	0.32	5.38
8.22	5.56	8.19
1.81	8.87	1.04
7.50	8.23	6.38

Got it!

4.30	5.46	5.23
7.15	3.18	4.02
5.42	1.62	7.10
8.56	4.58	7.58

Got it!

3.70	4.05	1.93
1.23	0.18	5.95
6.34	8.50	9.38
3.67	5.21	1.84

Got it!

6.35	2.08	7.57
6.52	6.25	6.49
3.65	7.53	8.36
0.41	0.30	1.04

Got it!

9.15	9.95	4.97
4.10	0.41	8.93
3.11	5.90	5.55
5.72	8.00	7.73

Got it!

6.79	0.53	2.05
7.04	2.89	7.84
7.95	4.06	5.46
7.05	4.42	8.84

Got it!

5.04	3.44	3.86
7.76	3.62	7.20
5.99	8.99	4.67
2.80	4.07	7.70

Got it!

6.04	3.54	1.57
2.78	5.15	7.93
6.74	6.86	3.44
4.85	3.72	0.27

Got it!

2.62	1.35	0.80
6.04	4.17	3.27
3.37	4.46	8.65
0.24	6.57	2.19

Got it!

0.87	6.53	6.93
6.35	4.74	1.86
6.95	6.37	7.13
2.87	8.76	9.09

Got it!

1.78	8.86	7.03
0.47	6.32	9.03
9.45	2.63	9.12
7.39	4.68	7.37

Got it!

4.07	2.76	6.20
3.47	4.93	6.53
6.44	3.97	8.82
1.83	0.79	8.93

Got it!

REAL QUIZ

6.57	9.18	3.43
2.58	7.73	1.66
9.08	5.50	0.73
0.49	9.70	6.40

Got it!

9.54	2.62	2.27
1.77	9.14	7.38
3.69	2.31	5.43
5.23	5.98	0.94

Got it!

3.76	6.92	4.73
9.99	8.88	2.46
6.36	1.77	9.31
2.80	0.00	1.12

Got it!

8.64	7.47	5.25
6.19	2.97	2.99
2.89	4.00	6.94
0.05	7.81	1.36

Got it!

7.77	8.25	6.20
1.49	1.69	0.13
1.35	2.51	1.75
1.71	4.94	8.84

Got it!

9.74	0.77	0.12
5.49	9.85	7.41
5.23	9.39	7.09
6.05	4.51	4.41

Got it!

5.65	1.99	1.56
7.38	7.98	9.04
4.48	2.39	8.34
3.47	5.52	6.69

Got it!

0.46	3.25	9.63
9.75	2.42	1.48
0.61	2.04	0.80
9.54	3.32	4.44

Got it!

0.36	8.22	1.78
6.02	9.70	7.43
0.65	7.55	5.32
6.27	8.27	9.58

Got it!

7.75	2.97	8.10
1.29	6.16	2.34
9.53	2.25	0.85
9.73	3.46	1.26

Got it!

3.26	9.66	1.72
2.82	5.56	4.59
3.89	8.02	8.28
9.59	1.76	4.34

Got it!

1.55	8.36	2.06
2.29	5.29	0.70
8.45	4.89	6.63
4.07	3.56	0.09

Got it!

3.15	2.27	7.91
2.09	1.01	0.04
6.98	8.88	3.12
8.19	6.05	6.54

Got it!

3.41	0.25	4.63
8.75	9.57	9.69
8.97	5.51	6.93
4.94	6.46	9.75

Got it!

3.64	7.16	6.99
6.92	9.73	6.81
3.01	5.64	3.53
7.24	4.86	3.74

Got it!

6.89	4.73	3.34
8.63	4.69	0.63
0.94	9.35	8.86
3.11	6.23	4.91

Got it!

9.26	8.55	6.70
2.60	7.87	7.56
3.87	8.62	8.48
5.55	3.78	1.38

Got it!

1.54	0.40	7.28
0.56	3.49	1.16
7.85	8.87	2.49
2.15	2.76	8.58

Got it!

9.40	7.45	1.77
3.38	0.57	2.83
0.90	8.35	7.54
9.43	8.20	4.29

Got it!

1.75	5.31	9.30
2.62	8.59	6.11
5.03	2.11	2.68
7.89	7.86	8.17

Got it!

4.13	7.92	1.62
7.30	5.91	6.81
5.04	4.96	2.52
6.51	0.35	0.20

Got it!

4.63	3.93	6.11
9.95	7.96	5.87
0.46	3.01	9.54
9.42	1.68	9.01

Got it!

6.69	2.94	6.76
3.83	8.94	0.98
8.81	6.68	8.06
3.50	1.94	7.33

Got it!

0.25	7.42	6.28
9.17	3.87	2.07
5.54	8.53	1.09
6.13	7.20	8.70

Got it!

1.08	4.97	6.90
6.64	5.38	5.80
2.84	9.79	9.38
0.78	1.25	3.36

Got it!

7.35	1.98	2.49
7.52	4.29	6.80
8.02	1.97	2.35
0.54	7.29	5.34

Got it!

0.08	6.90	2.14
7.60	1.63	4.48
0.74	4.11	2.35
1.33	5.89	0.14

Got it!

4.18	1.36	3.51
7.42	2.30	3.51
2.58	1.80	8.82
9.40	1.67	1.41

Got it!


9.57	7.17	2.66
3.76	8.37	7.14
1.63	6.04	8.61
7.37	6.21	7.64

Got it!


8.84	6.77	9.06
5.56	7.30	3.29
8.19	0.95	6.71
6.03	1.19	9.04

Got it!

Appendix N: Baseline EDA and SCR Frequency



The above figure shows the EDA level in the baseline region with 75% confidence intervals. One outlier in each of the *control* and *charity_loses* skewed the overall levels higher in those two conditions than the other two, but the differences are not statistically significant.



The above figure shows the frequency of SCRs in the baseline region with 75% confidence intervals. SCR frequency is not significantly different across conditions for the baseline region.