Information and Motivation In Organizations

Thesis by

Sera Linardi

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2010

(Defended May 19, 2010)

To my father, Dr. Juwira Linardi,

who gives me purpose,

and

my husband, Will Halim,

who makes pursuing it possible.

Acknowledgements

I will begin by recognizing my advisor, John Ledyard, as the patron saint of graduate students. I'm not kidding. John is patient, knowledgeable, and generous, all while relentlessly challenging his students to continuously improve. I want to express my sincere thankfulness to Colin Camerer for both his mentorship and his endless creative ideas. Our family will never forget how Colin and his family, Sharon and Conor, have welcomed and included us. I am especially grateful to Tom Palfrey, Jean-Laurent Rosenthal, and Leeat Yariv: who have each spent long hours helping me form ideas, going through drafts of my papers, and preparing me for conference presentations. I would also like to thank Robert Sherman, Thayer Scudder, Ken Pickar, Matias Iaryczower, Matt Shum, Federico Echenique, Antonio Rangel, Charlie Plott, and Jasmina Arifovic for their support.

Many fellow students have transformed my experience at Caltech into a marvelous one. Those most responsible are as follows: my classmate Min Jeong (without whose catch-up sessions, I would never have gotten past my first year), my coauthor Maggie, and my friends Debrah, Michael, Ines, Mohamed, Guilherme, Stephanie, Shankar, Muruvvet, Theresa, Sang Mok, Stuart, Andrea, and the rest of the board members of Engineers for a Sustainable World. Likewise, I owe a debt of gratitude to Julian, Gabriel, and Aniol: for doing the impossible by making the job market experience actually fun. My special thanks go to the staff of HSS, namely, Laurel Auchampaugh, Tiffany Kim, Sabrina Boschetti, and Walter Yuan, for safeguarding me from total disorganization.

Finally, I would like to acknowledge my parents, Juwira Linardi and Rudiyanti Lemena; my late father-in-law, Iskandar Halim; my mom-in-law Sri Melani Intan; and my sister, Ellen Linardi, and her family, Ian and Kara. At a fundamental level, this would have been an enormously lonesome endeavor without my family; my preternaturally Zen husband, Will, and my children, Andreas and Nina, make coming home every day a wondrous event.

Abstract

My research focuses on incentive/information design for environments where contract enforcement is difficult and the information required for decision-making is dispersed. These environments are particularly challenging when the number of participants are small enough such that small perturbations have persistent influences. In these three chapters, I use theory, computation, and experiment to investigate the robustness of several basic economic mechanisms to stochastic noise.

The first chapter analyzes the basic unit of information aggregation – the Geanakoplos and Polemarchakis (1982) posterior revision process. I find that if stochastic noise is present, then 1) the posterior revision process does not reliably give public statistics that approach the full information posterior, and 2) methods exist to rank information structures based upon the likelihood that they produce good public statistics through the posterior revision process.

The last two chapters address the impact of stochastic noise on labor markets. The chapter coauthored with Margaret McConnell uncovers the image motivation behind prosociality by enforcing privately known stochastic stopping time in volunteering sessions. A unique cascade of quitting behavior suggests that volunteers are partially driven by stigma avoidance. The third chapter, coauthored with Colin Camerer, analyzes the robustness of contracting relationships to exogenous disruptions caused by stochastic drops in demand. We find that stochastic noise slows the formation of relational contracts, but high-quality contracts remain unaffected.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In Chapter 2, Accounting for Noise in the Microfoundations of Information Aggregation, I show that the basic unit of information aggregation described by the Geanakoplos and Polemarchakis (1982) posterior revision process does not always produce public statistics that are closer to the full information posterior than the common prior. I study this process of back and forth communication between two individuals with private signals by introducing white noise into payoff computations, defining the evolution of common knowledge, and providing conjectures on the resulting public statistics. I then develop a computational method to ex-ante rank information structures on likelihood of beating the prior. Through a laboratory experiment I find that subjects behave in ways consistent with the noise model: even though initial revisions bring reports closer to each other and to the full information posterior, subsequent revisions have little effect. A comparison of report errors across four information structures provides evidence that the ex-ante ranking matters; more than half of subjects' reports are further from the full information posterior than the common prior in the two lowest ranked structures.

Chapter 3, No Excuses For Good Behavior: Volunteering and Social Environment, analyzes the effect of stochastic noise in the social environment on the quantity and quality of prosocial behavior. We look specifically at volunteering, the most common example of prosociality. By extending Benabou and Tirole's (2006) image signaling framework, we derive theoretical predictions on time volunteered given (1) the availability of excuses to stop volunteering and (2) the presence of an authority figure. We test these predictions in an experiment where laboratory subjects are directly involved in a local nonprofit operation. We find that in the absence of excuses to stop volunteering, subjects volunteer longer without working less productively. This increase is partially driven by subjects reluctance to be the first to stop volunteering. The presence of an authority figure has little impact, but the presence of peers has a positive and significant impact.

Chapter 4, Can Relational Contracts Survive Stochastic Interruptions?, studies the robustness of contracting relationships to exogeneous disruptions caused by stochastic drops in demand. This paper investigates the robustness of the "two-tiered labor market" experimental results of Brown, Falk and Fehr (2004) by subjecting relationships to stochastic interruptions. Firms in Brown et al prepay a wage, request a level of effort, and then workers who accept the offer choose effort. Effort is costly to workers and valuable to firms, but cannot be explicitly contracted for. They find that when firms can make private offers to particular workers (identified by ID numbers), long term relationship with high wages and high effort level account for a large fraction of contracts in the market. We replicate and extend their paradigm to allow for exogenous random "downturns" in which firms cannot hire workers for three periods. Our hypothesis is that (1) job rents are lower in downturns (2) this will lower the wages and effort that can be sustained in equilibrium. We do find that job rents are lower, but surprisingly, the downturns do not harm aggregate market efficiency. On the contrary, wages in the second half are higher in labor markets with downturn than in the baseline markets. Wages and efforts in the bottom tier (public) market are significantly higher in the downturn, pushing the top tier (private) market to continue raising wages to achieve adequate separation. We conjecture that this is because firms' charge a 'relationship premium' for private offers in downturn market early on. In contrast with baseline market where private offers are always better in price and surplus sharing than public offers, private offers in new relationships (those who trade together twice or less) are worse than public offers. This delays the formation of long term relationships and necessitates the use of public offers, which increase the competitiveness of the short term market. However, good workers and fair firms are largely unaffected because a strong loyalty norm allows them to reconnect after the downturn

Chapter 2

Accounting for Noise in the Microfoundations of Information Aggregation

2.1 Introduction

One of the most common human interactions consists of individuals inferring information from the behavior of others and reacting to it. The assumption that people do this correctly serves as a theoretical basis for a large variety of economic mechanisms under uncertainty. Under rational expectations equilibrium (REE), prices are supposed to aggregate all the information about the states of the world. Those who are 'outsiders' to the information in the system become 'insiders' simply by observing these statistics (Plott, 2000). This has led to a large literature on the design and use of economic mechanisms for information aggregation. However, increasing experimental evidence shows that (1) prices in markets and other forms of public statistics, e.g., prices in market, odds in parimutuel, and numbers in polls,[1] only approximately converge to the full information posterior, and (2) increased complexity is accompanied by increased noise (Sunder, 1995).

Noisy and approximate aggregation implies that an outside observer may not always be better off abandoning the common prior in favor of the public statistics produced by the interactions of insiders. Various institutional components such as aggregate uncertainty (Plott and Sunder, 1988; Lundholm, 1991; Ackert and Church, 1998), number of assets

[1] Some examples are Plott and Sunder (1988) for double auctions, Plott and Roust (2006) for parimutuel, McKelvey and Page (1995) for iterative polling, and Ledyard, Hanson, Ishikida (2009) for poll with Market Scoring Rule. See Healy, Ledyard, Linardi, and Lowery (2009) for experimental comparisons of the aggregation properties of these four mechanisms.

(O'Brien and Srivastava, 1991), and signal informativeness (Goeree, et al., 2008) seem to contribute to informational efficiency. In this paper, I capture the interaction of these three components by analyzing information structures. Focusing on a basic aggregation process, I provide a method of ranking information structures on the relative likelihood that public statistics improve predictions.

I study the Geanakoplos and Polemarchakis (1982) posterior revision process which describes how common knowledge of the full information posterior is reached instantaneously when two individuals with private signals communicate. This setup is fundamentally different from the more commonly studied information cascades setting: individuals now have repeated opportunities to revise their previous actions after receiving feedback from others' actions. The posterior revision process forms the microfoundation of theoretical information aggregation; extensions by McKelvey and Page (1995) and Nielsen, et al., (1995) describe how dispersed information is incorporated into market prices.[2] Further insights stemming from this basic process continue to be influential in recent work.[3]

I perturb this basic unit of aggregation by introducing white noise into payoff computations. Both the information content of reports and the speed at which beliefs converge become dependent on the information structure and noise level. Each step of revising and reporting introduces new noise into the system, thus limiting the marginal impact of additional revisions. In some information structures, a low level of perturbation is sufficient to result in public statistics that are worse approximations of the full information posterior than the common prior. When the expected distance of reports from the full information posterior is monotonically increasing in the noise level, there is a unique solution to the probability of improving predictions in an information structure. An outsider with no information about the noise level can use this result to ex-ante rank a set of information structures on the likelihood of improving prediction using the posterior revision output.

I ex-ante rank four information structures on the likelihood of prediction improvement. I then provide the first in-depth experimental investigation of the dynamics of posterior revision by testing the process in these four structures.[4] I find that subjects behave in

[2]See Feigenbaum, Fortnow, Pennock, and Sami (2005); Chen, Mullen, and Chu (2006); and Ostrovsky (2009) for connection between the posterior revision process and prediction markets.

[3]See Hanson (2002); Menager (2005); Acemoglu, Chernozhukov, and Yildiz (2005); Liu (2008); and Zimper (2009)

[4]Weber's "dirty face game" (2001) is related to the common knowledge inference process but does not directly depict the posterior revision process.

ways consistent with the noise model: even though initial revisions bring reports closer to each other and to the full information posterior, subsequent revisions have little effect. Report errors remain significant throughout the revision process. Comparison of these errors across the four information structures provides evidence that this ex-ante ranking matters. I show that fragility to mistakes is not simply a function of states; indeed, experimental subjects produce more frequent and larger prediction improvement in a three-state information structure than in a two-state structure.

The rest of the paper proceeds as follows. Sections 2.2.1 and 2.2.2 define information structures, the posterior revision process and the prediction improvement metric. Section 2.2.3 analyzes the changes in common knowledge evolution from the posterior revision process when stochastic noise is introduced. In Section 2.3.1, I assume a particular behavioral model with white-noise-payoff disturbances and derive testable conjectures. I then describe the noise tolerance ranking, a computational method to rank information structures on the likelihood of prediction improvement. In Section 2.3.2, I first describe the hypothesis that will be tested in my experiment, including the four information structures that will be used as experimental parameters. I then provide numerical simulations of the posterior revision process within each structure and compare the outcome across the four structures. Section 2.4.1 presents the experimental design. Results are organized as follows: Section 2.4.2 utilizes individual level data to estimate noise levels and find evidence of noisy Bayesian updating, and Section 2.4.3 compares raw report error in the experiment with the simulation. The last two subsections ask whether the noise tolerance ranking is able to predict prediction improvement. Graphs of simulation and experimental results can be found in Section 2.4.4; regressions, in Section 2.4.5.

2.2 Theoretical Model

2.2.1 Posterior Revision as the Microfoundation of Aggregation

An information aggregation setup can be described as follows: Θ is a finite set of states with an associated probability distribution $p \in \Delta_\Theta$, where Δ_Θ is the set of all possible probability distributions over Θ. State $\theta \in \Theta$ occurs with probability $p(\theta)$. *Insiders* $i \in \{A, B, .., I\}$ privately observe signal $s_i \in S_i$ which is drawn with probability $q_i(s_i|\theta)$. The distribution of signals across all Insiders s is drawn from the finite set $S = S_A \times .. \times S_I$ with probability $q(s|\theta) = \prod_{i=1}^{I} q_i(s_i|\theta)$. As is standard in the literature, I consider only states that are distinguishable by signals, $q(s|\theta) \neq q(s|\theta')$ if $\theta \neq \theta'$.

Previous literature have separately investigated the challenge posed to information aggregation by various components of this setup. O'Brien and Srivastava (1991) experimentally demonstrate that aggregation degenerates as the number of assets $|\Theta|$ increases. Plott, Witt, and Yang (2003) show that aggregation improves as the private signal becomes more informative. However, the trade-off between these components remains unknown. For example, is aggregation better when there are three states with more informative signals or two states with less informative signals? How do the prior probabilities of the states affect aggregation? This paper departs from the previous literature by analyzing an information structure (number of states, priors, and marginal probability) as a unit in determining difficulty of aggregating information.

Definition 1. *Let $(\Omega, 2^{\Omega}, \psi)$ be a probability space where $\Omega = \Theta \times S$ and $\psi : 2^{\Omega} \rightarrow [0, 1]$ is the probability measure that represents the common prior. $\psi(\theta, s)$ is defined as $q(s|\theta)p(\theta)$, $\psi(\theta|s) = \sum_{s} \psi(\theta, s)$, and $\psi(s) = \sum_{\theta} q(s|\theta)$. I will refer to ψ as **information structure** .*

Whether an information aggregation mechanism takes the form of a market, a parimutuel, or a repeated polling, the central mechanism is as follows: individuals holding private information choose actions, which influence the formation of a public signal. The announcement of this public signal allows individuals to make inferences about other individuals' private information; these inferences, in turn, induce the individuals to revise their own actions. These new actions produce the next public signal, one that, in theory, contains more information than the previous public signal.[5] Communication between two Insiders comprises the most basic unit of the process: an individual, A, chooses his action; another, B, observes A's action, updates her information, and chooses her own action; then, A responds by revising his action. This is the Geanakoplos and Polemarchakis (1985) posterior revision process. When the basic unit works well and instantaneously, this process generates the full information posterior of A's and B's combined private signals. Observers of the resultant statistic incorporate A's and B's private signals into their own information before taking further actions; thus, the aggregation process advances (McKelvey and Page, 1995; Nielsen,

[5]Contrast this with information cascades, in which an individual only takes a single action that he will not be able to revise.

et al., 1995). If the posterior revision process works slowly or noisily, though, the generated statistic may not correspond to the full information posterior, as errors may accumulate throughout the aggregation process.

Geanakoplos and Polemarchakis (1985) assume that individuals are communicating truthfully. McKelvey and Page (1990) augment the process with the quadratic scoring rule (QPSR) and show that myopic truthful reporting at every step, by every subject, is a Bayes-Nash equilibrium.[6] The posterior revision process studied here follows their precedence in utilizing QPSR.[7] In this setup, individuals are directly assumed to be myopic; hence, they choose a_i^t to maximize period t's payoff.

Definition 2. *Let $i \in \{A, B\}$. Define the incentivized T-step* ***posterior revision process*** *as $PR^T = (\Delta_\Theta \times \emptyset, Y_i, \gamma)$. $\Delta_\Theta \times \emptyset$ is the finite set of actions that i can choose from. The action of individual i is $a_i = (a_i^1, ...a_i^T)$ where $a_i^t \in \Delta_\Theta$ when t is odd and $i = A$ or when t is even and $i = B$; $a_i^t = \emptyset$ otherwise. A scoring rule assigns a score to a probability report a_i^t based upon the realization of the state of the world $\tilde{\theta}$ after $t = T$. $Y_i(.) = \sum\limits_{t=1}^{T} Q_{\tilde{\theta}}(a_i^t)$ is the state contingent payoff function for player i. $Q(.)$ is the quadratic scoring rule defined as:*

$$Q_{\tilde{\theta}}(a_i^t) = 1 - \sum_{\theta \in \Theta} (a_i^t(\theta) - \eta(\theta|\tilde{\theta}))^2$$

where $\eta(\theta|\tilde{\theta}) = 1$ when $\theta = \tilde{\theta}$, or 0 otherwise. $Q_{\tilde{\theta}}(a_i^t) = 0$ when $a_i^t = \emptyset$. The outcome function γ produces a probability distribution over Θ from players' actions: $\gamma(a_A, a_B) = a_{I(T)}^T$ where $I(T) = A$ when T is odd and $I(T) = B$ when T is even.

Lemma 2.2.1 shows that myopic payoff optimizers facing QPSR report their current period beliefs; hence, by Geanakoplos and Polemarchakis (1982), the second step of the incentivized T-step posterior revision process produces the full information posterior.

Lemma 2.2.1. *Let $p_i^t \in \Delta_\Theta$ denote i's belief at time t. A myopic payoff optimizer i facing a payoff function of $\sum\limits_{t=1}^{T} Q_{\tilde{\theta}}(a_i^t)$ maximizes his payoff by reporting $a_i^t = p_i^t$.*

All proofs are in the Appendix.

[6] See McKelvey's and Page's extension of the posterior revision process into a three-person iterative poll.
[7] For an explanation of favorable properties of QPSR in terms of handling risk aversion, accuracy, and calibration: see Sonnemans and Offerman, 2001; and Palfrey and Wang, 2009.

Indicate the private signal realizations as $\tilde{s} = \tilde{s}_A\tilde{s}_B$ and the full information posterior as $p_{\tilde{s}}$. The information structure ψ and behavioral model is common knowledge[8].

Theorem 2.2.2. *Geanakoplos and Polemarchakis (1982). Let $i \in \{A, B\}$ be two myopic payoff maximizers participating in PR^T. Given private signal realization $\tilde{s}$, $\gamma(.) = p_{\tilde{s}}$ for all PR^T where $T \geq 2$.*

The intution is provided below, see the original paper for the full proof. Under Bayesian updating, $p^1_A(\theta) = \psi(\theta|\tilde{s}_A)$ and $p^1_B(\theta) = \psi(\theta|\tilde{s}_B)$. By Lemma 2.2.1 at step 1, $a^1_A = p^1_A$. Under common knowledge of Bayesian updating, $Pr(s_A|a^1_A = p^1_A) = 1$ for $s_A = \tilde{s}_A$ and 0 otherwise. Hence $p^2_B(\theta) = \psi(\theta|\tilde{s}_A\tilde{s}_B) = p_{\tilde{s}}$ and by Lemma 2.2.1 at step 2 $a^2_B = p_{\tilde{s}}$. Under common knowledge of Bayesian updating, $Pr(s_B|a^2_A = p^1_A) = 1$ for $s_B = \tilde{s}_B$ and 0 otherwise. Hence $p^3_A(\theta) = \psi(\theta|\tilde{s}_A\tilde{s}_B) = p_{\tilde{s}} = p^2_B(\theta)$, indicating that beliefs have converged. As before $a^3_A = p^3_A(\theta) = p_{\tilde{s}}$ and from this step onwards beliefs and reports remain at the full information posterior.

2.2.2 Measuring Aggregation Success

I now define a metric to measure aggregation success. Let $\omega^n = (\theta^n s^n) \in \Omega$ indicate the nth independent draw from $(\Omega, 2^\Omega, \psi)$. Let $p_{s^n} \in \Delta_\Theta$ indicate the full information posterior over Θ computed from s^n where $p_{s^n}(\theta) = \psi(\theta|s^n)$. Let $\alpha^n \in \Delta_\Theta$ be the observed output of PR^T given ω_n. Following convention,[9]. the loss of approximating p_{s^n} with α^n is measured with the Euclidean distance. Define *report error* of α_{s^n} as follows:

$$d(p_{s^n}, \alpha_{s^n}) = (\sum_\theta (p_{s^n}(\theta) - \alpha_{s^n}(\theta))^2)^{\frac{1}{2}} \tag{2.1}$$

Define an *Outsider* as an individual $O \notin \{A, B\}$ who only knows of the information structure ψ and does not receive any private signals s_i. The Outsider does not participate in PR^T, but he does choose whether to use the output α_{s^n} in place his common prior p to make predictions about Θ. Denote the set comprising of both full information posteriors and observed output distribution from N realizations of ω^n in ψ as $D^N_\psi = \{(p_{s^1}, \alpha_{s^1}), ...,, (p_{s^n}, \alpha_{s^n}), ..(p_{s^N}, \alpha_{s^N})\}$. Let *baseline error* be the average error over

[8]Common knowledge indicates a shared set of information, comprised of the following: that which A knows that B knows, B knows that A knows, A knows that B knows that A knows, and so on, ad infinitum.

[9]Other measures of the degree of aggregation that has been used are L1 (Koessler, Noussair, and Ziegelmeyer, 2006), Wurtz distance (Plott, Witt and Yang, 2003), and the Kullback Leibler divergence.

D_ψ^N from following the prior in ψ: this quantity indicates the loss suffered by the Outsider when no attempts to aggregate information are made.

$$BaseErr(D_\psi^N) = \frac{1}{N}\sum_{n=1}^{N} d(p_{s^n}, p) \tag{2.2}$$

By following PR^T's nth output in information structure ψ, the Outsider achieves an *error reduction* of $d(p_{s^n}, p) - d(p_{s^n}, \alpha_{s^n})$. An error reduction of 0.4 in an information structure where the baseline error is 0.8 indicates a prediction improvement. This would not be true for an information structure where the baseline error is 0.3. In the metric defined below, the size of error reduction will be evaluated relative to the baseline error.

Definition 3. ***Prediction Improvement*** *of the nth output in information structure ψ is the error reduction of the nth output normalized by the baseline error of ψ.*

$$PI(n, D_\psi^N) = \frac{d(p_{s^n}, p) - d(p_{s^n}, \alpha_{s^n})}{BaseErr(D_\psi^N)}$$

Write the average prediction improvement as $\bar{P}I(D_\psi^N)$.

This metric provides a standard with which aggregation across different information structures can be compared. $\bar{P}I(D_\psi^N) = 1$ indicates that $\alpha_{s^n} = p_{s^n}$ for all $n \in \{1, .., N\}$. $\bar{P}I(D_\psi^N) \leq 0$ indicates that the error of the aggregation process in D_ψ^N is on average larger than the baseline error. In this case, the Outsider who naively abandons his prior would have been better off without the aggregation process.

With myopic payoff optimizers, an Outsider to the system can always improve upon the common prior by following the output of PR^T. As described in the next section, this will no longer be the case when stochastic noise is present.

2.2.3 Common Knowledge Evolution with Stochastic Reports

Experimental evidence has shown that, instead of the exact and instantaneous aggregation predicted by theory, all information aggregation mechanisms proposed to date produce statistics that only approximately and noisily converge to the full information posterior. Three types of explanations have received the most attention: systematic deviations from

Bayesian updating,[10] the effect of risk attitudes,[11] and the existence of stochastic noise. The type of systematic deviations from Bayes rule that best fit the information aggregation setting is still an open question. This paper will focus instead on the third explanation, the existence of stochastic noise, and build a noise tolerance model that assumes risk neutrality. The experimental section will test whether this noise tolerance model is able to predict behavior.

Stochastic noise can occur in the perception of a private signal, in the belief updating process, or in payoff perception. In an experimental setting, private signals are often simple: for example, subjects might draw a colored ball from a covered urn. It seems unlikely that subjects will become confused about this type of basic information during the experiment; thus, the existence of stochastic noise in their actions is more likely a result of noisy belief updating or payoff computation. Modeling stochastic noise in payoff perception presents two advantages. First, it allows the failures of aggregation to be examined while retaining the assumption of Bayesian updating. Secondly, there is strong evidence of stochastic action in simple risky-choice experiments where Bayesian updating does not play a role (Camerer 1989, Starmer and Sugden 1989, Hey and Orme 1994, Ballinger and Wilcox 1997)[12]. This latter point suggests that noise in payoff perception exists independently of noise in belief updating.

Below, I will describe PR^T as played by agents who choose their actions stochastically due to disturbances in the payoff perception.[13] In the next subsection, I will assume that payoff disturbances can be represented with white noise ($\epsilon \sim \mathcal{N}(0, \sigma)$). I then propose several conjectures behavior of PR^T under this assumption to serve as the building blocks of a method to rank information structures over likelihood of prediction improvement. Support for the conjectures and the ranking is provided through numerical simulations (Section 2.3.2) and experiments (Section 2.4.1).

[10] These deviations include overweighting one's prior relative to private information (conservatism), neglecting one's prior when considering private information (base rate fallacy) (Kahneman and Tversky 1973), and overweighting one's private information relative to information from others (Hung and Plott 2001, Goeree et al. 2008). See Camerer (1995).

[11] Risk aversion has an effect of 'flattening' reported probabilities towards equal likelihood (Harrison and Rustrom, 2008).

[12] These studies found that more than 25% of subjects reversed their choices when presented a second time with the same binary-choice problem.

[13] Stochastic noise in belief updating and signal perception may result in actions that appear to be stochastic in the beginning. It seems that noise in belief updating and signal perception will eventually result in the correct action, while noise in payoff perception persist throughout – even when beliefs are correct, actions will continue to be stochastic.

Let $f(\sigma)$ be the distribution from which payoff disturbances are drawn where $\sigma = 0$ indicates no noise and $\sigma = \infty$ indicates pure noise. The noise level σ is a random variable drawn from an unknown distribution $g(.)$ each time $\omega = (\theta, s_A s_B)$ is realized.[14] Let F and G be the respective cumulative density functions for f and g. A and B, as Insiders, observe both s_i and the background noise level σ; O, as the Outsider, does not. As an example, imagine A and B as retailers in a busy mall. They can directly observe customers' reactions to a product, but the pair are distracted by activities in the mall while discussing this information. In comparison, O would be the CEO, far removed from the mall, who must try to infer customers' attitudes from the communications between A and B, all while unaware of the context in which these communications occured.

Let $Pr_\sigma(a_j^t|p_j^t)$ be the standard quantal response function of Luce (1957) denoting the probability of choosing an action a_j^t given beliefs p_j^t and noise level σ. As before, myopic players optimize only their current-period payoffs; however, this computation is now affected by stochastic disturbances.[15] With stochastic action in effect, any given report could have been induced by any private signal. Consequently, complete aggregation at the second step of the PR (Theorem 2.2.2) cannot occur in this absence of one-to-one mapping of report to private signal. The state of common knowledge among A and B must therefore be enlarged to contain each person's hypothetical beliefs about Ω given each realization of the private signal.

Definition 4. *Denote* ***common knowledge*** *among A and B at step t of the posterior revision process as* $K^t = (K_A^t, K_B^t)$. K_A^t *is the set* $\{\psi_{s_A}^t\}_{s_A \in S_A}$ *where* $\psi_{s_A}^t$ *is A's posterior over* Ω *at step t if A's private signal is* s_A. *Similarly,* K_B^t *is the common knowledge at time t about B's possible beliefs. The initial state of common knowledge is* $K^1 = (K_A^1, K_B^1) = (\{\psi_{s_A}^1\}_{s_A \in S_A}, \{\psi_{s_B}^1\}_{s_B \in S_B})$ *where* $\psi_{s_i}^1(\theta, s_j) = \psi(\theta, s_j|s_i)$.

A general behavioral model of common knowledge evolution is given by Proposition 2.2.3. I will discuss a specific behavioral model that utilizes the assumption of white noise payoff disturbances in Section 2.3.1. When actions are stochastic, an individual gradually learns about his partner's signals by interpreting the history of past actions within the context of the noise level. After observing A's report a_A^t, B can update his probability

[14] Assume that the communication process is short enough that noise level is constant from start to end.

[15] Since the players are myopic, this is not an equilibrium model such as McKelvey and Palfrey (1995, 1998) quantal response equilibrium (QRE).

distribution over A's private signals by comparing the probabilities of a_A^t given A's various hypothetical beliefs. Notice that A is aware of B's updating process, but not A of the private signal that initialized B's beliefs. The resulting evolution of common knowledge is path dependent. Player A perceives that the inaccuracies of his earlier reports affected B's beliefs; so, A takes this knowledge into account as he attempt to infer B's signal from B's subsequent reports.

Definition 5. *Let an* ***update rule*** *be a function $L(K^t, a_i^t, \sigma) = K^{t+1}$ with the following inputs: the current state of common knowledge K^t, noise level σ, and the most recent report a_i^t. The update rule returns the next period state of common knowledge K^{t+1}.*

Proposition 2.2.3. *Let σ be the noise level present when A and B are communicating. The update rule of PR^T at step t is*

$$L(K^t, a^t, \sigma) = \begin{cases} K_A^t, \{\psi_{s_B}^{t+1}\}_{s_B \in S_B} & \text{if } t \text{ is odd} \\ \{\psi_{s_A}^{t+1}\}_{s_A \in S_A}, K_B^t & \text{otherwise} \end{cases}$$

$$\text{where } \psi_{s_i}^{t+1}(\theta, s_j) = \frac{\psi_{s_i}^t(\theta, s_j) Pr_\sigma(a_j^t | p_{s_j}^t)}{\sum\limits_{s_j'} \psi_{s_i}^t(s_j') Pr_\sigma(a_j^t | p_{s_j'}^t)}$$

where $p_{s_j}^t \equiv \sum_{s_i} \psi_{s_j}^t(\theta, s_i)$ is j's posterior over Θ at step t ***if*** *j's private signal is s_j.*

The probability of observing a certain report as the output of this process depends on the history of actions that has preceded it; however, an identical state of common knowledge could result from a variety of different action histories. This consequence stems from the update rule itself, which is driven solely by the likelihoods of observing report a_j^t given j's hypothetical draws s_j. For private signal realization $\tilde{s}$ and noise level is σ, denote the set of all possible states of common knowledge at step t as $\mathcal{K}_{\tilde{s},\sigma}^t$. The *expected error* of the output of PR^T is:[16]

$$\mathbb{E}Err_{PR^T}(\tilde{s}, \sigma) = \mathbb{E}_{K^T \in \mathcal{K}_{\tilde{s},\sigma}^T} \sum_{a \in \Delta_\Theta} Pr_\sigma(a | p_{\tilde{s}_{I(T)}}^T) d(p_{\tilde{s}}, a)$$

Using Proposition 2.2.3, $\mathbb{E}Err_{PR^T}(s, \sigma)$ as a function of σ can be numerically computed for any behavioral model that provides a definition of $Pr_\sigma(a|p_i^t)$. This error can then be compared to the baseline error $d(p_s, p)$; however, since the random variable s is unobserved

[16] As an example, suppose T is odd. Then $I(T) = A$, and last report is chosen with $Pr_\sigma(a|p_{\tilde{s}_A}^T)$ where $p_{\tilde{s}_A}^T = \sum_{s_B} \psi_{\tilde{s}_A}^T(\theta, s_B)$ and where $\psi_{\tilde{s}_A}^T \in K_A^T \in \mathcal{K}_{\tilde{s},\sigma}^t$.

by the Outsider, this ex-ante computation must be done by taking expectation over all possible signals. The metric Expected Improvement (EI) below takes this lack on information on the Outsider's part into account.

Definition 6. *The* ***Expected Improvement*** *from using the output of PR^T in place of the common prior in information structure ψ is as follows:*

$$EI_{PR^T}(\psi|\sigma) = \mathbb{E}_s \mathbb{E}_{K^T \in \mathcal{K}^T_{s,\sigma}} \sum_{a \in \Delta_\Theta} Pr_\sigma(a|p^t_{s_{I(T)}}) \frac{[d(p_s, p) - d(p_s, a)]}{\mathbb{E}_s BaseErr(\psi)}$$

where $\mathbb{E}_s BaseErr(\psi) = \sum_{s \in S} \psi(s) d(p_s, p)$.

In the worst-case scenario, reports become pure noise: they are completely uninformed by either private signals or previous reports. Below, I will conjecture that, within any given information structure, following pure noise over discretized action space Δ_Θ should result in larger losses than remaining with prior probability p.

Conjecture 2.2.4. *Let $(\Omega = \Theta \times S, 2^\Omega, \psi)$ be a probability space and Δ_Θ be the discretized set of probabilities over Θ.*

$$\sum_s \psi(s) \sum_{a \in \Delta_\Theta} \frac{1}{|\Delta_\Theta|} d(p_s, a) > \mathbb{E}_s BaseErr(\psi)$$

If Conjecture 2.2.4 is true, then for all information structures there exists noise levels beyond which the Outsider would be better off remaining with his prior than following the output of the posterior revision process. The likelihood that an Outsider will be better off discarding his prior and following the output of PR^T in ψ can be approximated by the probability of drawing noise levels that satisfy $EI_{PR^T}(\psi|\sigma) > 0$. This is the intuition for the noise tolerance ranking described in the next section.

2.3 Posterior Revision with White Noise Disturbances

2.3.1 Behavioral Model and Noise Tolerance Ranking

I will begin by describing a simple behavior model, before suggesting several conjectures that will lead to the noise tolerance ranking of information structures. Let payoff disturbances be represented by white noise of level σ, i.e., let $\epsilon \in \mathbb{R}$ be a random variable distributed with $\mathcal{N}(0, \sigma)$. Individual i is aware of his true belief; he has arrived at these by updating his old belief using Bayes' rule. Assume that, when faced with QPSR, i can only choose to report either his true belief $p_i^t \in \Delta_\Theta$ or other probability distributions $a_i^t \in \Delta_\Theta$. The region of error that induces i to report a_i^t, when his true belief is p_i^t, may be expressed as follows:

$$R(a_i^t|p_i^t) = \left\{(\epsilon, \epsilon') \in \mathbb{R}^2 : \mathbb{E}_{p_i^t} Q(a_i^t) + \epsilon \geq \mathbb{E}_{p_i^t} Q(p_i^t) + \epsilon'\right\} \tag{2.3}$$

By Eq. 2.3, i chooses a_i^t when the difference in the payoff disturbances $\epsilon - \epsilon'$ exceeds the difference between i's expected payoffs from reporting a_i^t instead of p_i^t. Under QPSR payoffs, this difference is simply the Euclidean distance between p_i^t and a_i^t (Selten 1998). This allows the Outsider's loss function to be internalized within the Insiders' choice of actions.

Definition 7. *A* ***myopic*** σ ***optimizer*** *i faces payoff disturbances ϵ drawn from distribution F_σ, a centered normal distribution with variance σ. The probability an individual with belief p_i^t reports a_i^t is:*

$$Pr_\sigma(a_i^t|p_i^t) = \frac{\int\limits_{R(a_i^t|p_i^t)} dF_\sigma(\epsilon)}{\sum\limits_{a \in \Delta_\Theta} \int\limits_{R(a|p_i^t)} dF_\sigma(\epsilon)} \tag{2.4}$$

where $R(a_i^t|p_i^t)$ is the set of payoff disturbances such that the current-period payoff of reporting a_i^t exceeds the payoff of reporting the true belief p_i^t.

Lemma 2.3.1. *Let $F'_\sigma(d(p_i^t, a_i^t))$ be the probability of drawing errors larger than $d(p_i^t, a_i^t)$ from $N(0, 2\sigma)$.*

(i) When rewarded with QPSR, a myopic σ optimizer with beliefs p_i^t reports a_i^t with probability:

$$Pr_\sigma(a_i^t|p_i^t) = \frac{F'_\sigma(d(p_i^t, a_i^t))}{\sum\limits_{a \in \Delta_\Theta} F'_\sigma(d(p_i^t, a))}$$

(ii) For any $0 < \sigma < \infty$, $Pr_\sigma(a_i^t|p_i^t)$ is decreasing in $d(p_i^t, a_i^t)$

How does this probability distribution change as a function of the noise level? Since the noise level is the variance of a centered normal distribution, I will make use of the following property.

Proposition 2.3.2.

$$\forall \delta, k \in \mathbb{R}_+,\ \frac{F'_\sigma(\delta)}{F'_{\sigma+k}(\delta)} < 1$$

This ratio is decreasing in δ.

Proposition 2.3.2 provides Lemma 2.3.3 with the following property: when a decrease in noise level σ induces i to report a_i^t more frequently and, correspondingly, $\beta \in \Delta_\Theta$ less frequently, a_i^t must be closer to i's true belief p_i^t than β. As a consequence, the average distance of reports to true beliefs is an inverse function of noise level.

Lemma 2.3.3. *Let $\sigma < \sigma'$.*

$$(i) Pr_\sigma(a_i^t|p_i^t) > Pr_{\sigma'}(a_i^t|p_i^t) \text{ and } Pr_\sigma(\beta|p_i^t) < Pr_{\sigma'}(\beta|p_i^t) \Rightarrow d(p_i^t, a_i^t) < d(p_i^t, \beta)$$

$$(ii) \sum_{a \in \Delta_\Theta} Pr_\sigma(a|p_i^t) d(p_i^t, a) < \sum_{a \in \Delta_\Theta} Pr_{\sigma'}(a|p_i^t) d(p_i^t, a)$$

I now propose two conjectures within the context of this behavioral model; in the next section, numerical simulations will provide support for both of these, using each of the four information structures tested in the experiment.

The first conjecture states that for all private signal realizations, the report errors of myopic σ optimizers at any step of PR^T are monotonically increasing in noise level. Without loss of generality, suppose t is odd. By Lemma 2.3.3(ii), a reduction in noise level results in a higher frequency of reports a_A^t that are closer, as measured in Euclidean distance, to beliefs $p_{\tilde{s}_A}^t$. More accurate reporting on A's part would increase the weight that B places on $\tilde{s}_A$.[17] This, in turn, should lead to a reduction in the distance between $p_{\tilde{s}_B}^{t+1}$ and $p_{\tilde{s}}$. Implementing Lemma 2.3.3(ii) again at step $t+1$, a lower noise level would again increases the frequency of reports a_B^{t+1} that are closer to $p_{\tilde{s}_B}^{t+1}$. Since $p_{\tilde{s}_B}^{t+1}$ is on average closer to $p_{\tilde{s}}$ under lower noise, a_B^{t+1} should also be (on average) closer to $p_{\tilde{s}}$.

Conjecture 2.3.4. *Let $(\Omega = \Theta \times S, 2^\Omega, \psi)$ be a probability space and Δ_Θ be a discretized set of probabilities over Θ. For private signal realization $\tilde{s} \in S$ and for any step $T \geq 2$,*

$$\mathbb{E}Err_{PR^T}(\tilde{s}, \sigma) \text{ is monotonically increasing in } \sigma \in \mathbb{R}+ \text{ in } [0, \sum_{a \in \Delta_\Theta} \frac{1}{|\Delta_\Theta|} d(p_s, a)]$$

[17] As σ is common knowledge, a reduction in noise would also lead B to update more confidently on a_A^t.

The second conjecture states that in a given information structure, prediction improvement is more likely when the private signal draws are identical for A and B than when they are different. Let $s = (s_A s_B)$ where $s_A = s_B$ serve to illustrate the former; and $s' = (s'_A s'_B)$ where $s'_A \neq s'_B$, the latter. Two observations suggest support for this conjecture. First, since s moves the full information posterior further than s', baseline errors would be larger for s than s'. Second, when the private signal realization is s no aggregation is required to beat the baseline error since an individual's initial beliefs is already closer to p_s than prior p. The opposite is true when s' occurs: since initial beliefs are farther away from $p_{s'}$ than p, beating the prior depends crucially on each players learning that the other player has received a different signal, which may be challenging when the noise level is very high.

Conjecture 2.3.5. *Let* ${D_\psi}^N = \{(p_s, \alpha_{s^1}), ..., (p_s, \alpha_{s^N})\}$ *and* ${D'_\psi}^M = \{(p_{s'}, \alpha_{s'^1}), ..., (p_{s'}, \alpha_{s'^M})\}$ *where* $s = (s_A = s_B)$ *and* $s' = (s'_A \neq s'_B)$*.*

$$Pr(PI({D_\psi}^N) > 0) > Pr(PI({D'_\psi}^M) > 0) \ \text{as } N \text{ and } M \rightarrow \infty$$

Since EI is merely a normalization of report errors, Conjecture 2.3.4 is a sufficient condition for EI to be monotonically decreasing in noise level. Conjecture 2.3.4, together with the earlier Conjecture 2.2.4 ensures that for every ψ there will be a corresponding noise level σ^* such that $EI_{PR^T}(\psi|\sigma^*) = 0$; for any noise levels above σ^*, priors will on average generate smaller losses than the output of PR^T.

Proposition 2.3.6. ***Noise tolerance ranking:***

Suppose Conjecture 2.2.4 and 2.3.4 are true.

(i) For all ψ $\exists \sigma^* \in \mathbb{R}_+$ *such that* $\forall \sigma > \sigma^*$, $EI(\psi|\sigma) < 0$*.*

(ii) Let $(\Omega_1, 2^{\Omega_1}, \psi_1)$ *and* $(\Omega_2, 2^{\Omega_2}, \psi_2)$ *be two probability spaces. Also let* σ_1^* *and* σ_2^* *denote the respective solutions to* $EI_{PR^T}(\psi_1|\sigma) = 0$ *and* $EI_{PR^T}(\psi_2|\sigma) = 0$*.*

$$\sigma_1^* > \sigma_2^* \Rightarrow Pr(PI(D^N_{\psi_1}) > 0) > Pr(PI(D^M_{\psi_2},) > 0) \ \text{as } N \text{ and } M \rightarrow \infty$$

(iii) Suppose $EI_{PR^T}(\psi_1|\sigma)$ *and* $EI_{PR^T}(\psi_2|\sigma)$ *are separated by first-order stochastic dominance. Then*

$$\sigma_1^* > \sigma_2^* \Rightarrow \bar{P}I(D^N_{\psi_1}) > \bar{P}I(D^M_{\psi_2}) \ \text{as } N \text{ and } M \rightarrow \infty$$

Info structure	$Pr(X)$	$Pr(Y)$	$Pr(Z)$	$Pr(Black\|X)$	$Pr(Black\|Y)$	$Pr(Black\|Z)$
2E	0.40	0.60		0.79	0.13	
2H	0.33	0.67		0.30	0.60	
3E	0.40	0.50	0.10	0.15	0.85	0.30
3H	0.25	0.50	0.25	0.40	0.60	0.80

Table 2.1: Marginal and conditional distributions of experimental treatments

As σ^* decreases, the Insiders, A and B, must exhibit ever-greater accuracy to produce public statistics that approximate the full information posterior better than the prior. This concept of noise tolerance unifies disparate elements of an information structure by providing a single number that indicates the extent to which an Outsider's benefit from information aggregation depends on the Insiders' rationality. In the next section, I will simulate this process; then, I will present a series of testable hypotheses which form the basis of my experimental investigation.

2.3.2 Numerical Simulation

Table 2.1 describes four information structures: two with two possible states (2E and 2H) and two with three possible states (3E and 3H).[18] Each state can be thought of as a box, namely, X ,Y, or Z, containing some known numbers of Black and White balls. The game setup goes as follows. First, a box is randomly chosen with probabilities in Column 1-3; then, a ball is drawn with replacement from the chosen box for player A (according to probabilities stated in Column 4-6). Another ball is drawn similarly for player B. Finally, the two players engage in PR^T: guessing the identity of the box by taking turns reporting probability distribution over the set of possible boxes.

This section will describe simulations of the first five steps of PR^T between two myopic σ optimizers in information structure 2E. Simulations of 3E, 2H, and 3H are qualitatively similar and can be found in the Appendix 2.6. The left panel of Figure 2.1 illustrates the simulated average distance between player A's beliefs and player B's; the central panel describes the average distance between beliefs and full information posterior; and the right panel plots the average accuracy of reports.[19]

[18] 2E is a shorthand for 2 state Easy, 2H for 2 state Hard, 3E for 3 state Easy, and 3H for 3 state Hard. The Easy and Hard assignments correspond to the result of the noise tolerance ranking derived later in this section.

[19] The MATLAB code can be found in the Appendix 2.8.

The left panel shows that the distance between beliefs increases from step 1 to step 2 before decreasing consistently in the subsequent steps. This initial increase is due to the fact that when private draws are similar, A and B's initial beliefs are identical. B's knowledge about A's signal at step 2 initially moves B's belief away from A's belief, but at step 3 A catches up and the beliefs starts converging. The central panel shows an oscillating pattern to belief error (distance between beliefs and full information posterior) as player A and B takes turns. This is caused by more tentative updating on A's part: as the first mover A has update his belief not only taking into account B's potential mistakes but also the effect of his early mistake (from his first report) on B's posterior.

Looking only at A's reports ($t = 1, 3, 5$) shows that as reports are revised belief errors are on average decreasing; however, the first revision has a much larger impact than the second revision. Player B's initial belief formation is identical to A's, so taking the sequence $t = 1, 2, 4$ as B's average belief errors further confirm the decrease in belief errors and the decreasing marginal impact of additional revision.

Report errors, as plotted in the right panel, exhibit similar patterns as belief errors; however, their changes are dampened by the additional layer of noise resulting from stochastic action. Even when belief errors approach zero, report errors persist since the reports themselves remain stochastic. Note also that report error at each step increases as σ goes from 0.1 to 0.3 to 0.5, providing support for Conjecture 2.3.4.

Stronger evidence that report error monotonically increase in σ for each realization of private signal draws is shown in the left panel of Figure 2.2 for 2E. The solid black lines illustrate report errors when private signal draws are *Black Black* or *White White* and the solid gray lines plot report errors under draws of *Black White* or *White Black*.[20] The corresponding baseline error for each draw is plotted as dotted lines. Note that the solid lines intersect the dotted lines at much lower σ in the case of different draws (gray) than for similar draws (black). This indicates that successful aggregation in the former demands Insiders' action to exhibit much less stochastic noise than the latter, lending support to Conjecture 2.3.5.

Suppose that an Outsider observes only the step t report of PR^T; from this, alone, he must then decide whether to follow the output or to remain with the common prior. Among the four information structures under examination, which one would provide for

[20] Analogous figures for 3E, 2H, and 3H are included in Figure 2.6 Appendix 2.6.

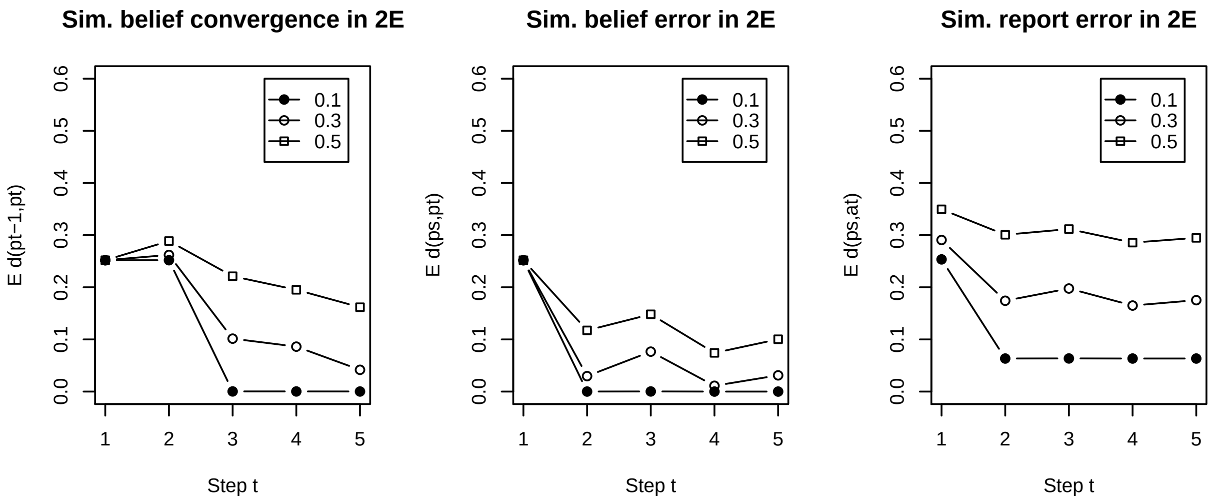

Figure 2.1: Simulation of PR^5 in 2E: revisions decrease distance between beliefs (L), average belief error (C), and report error (R).

the highest likelihood that he will be better off following the output? In other words, in which information structure is likelihood of achieving prediction improvement the highest?

Each of the four information structures has a distinct configuration of number of states, priors, and marginal probabilities. Previously no method exists to compare an Outsider's benefit in attempting to aggregate Insiders' information across these disparate information structures. One common approach is to rate three-state structures as harder than two-state structures. The literature is replete with examples of experimental design which purport to increase aggregation difficulty by increasing the number of possible states; however, they often lack theoretical basis. The other approach is to fall back on the assumption of perfect rationality: in such a world, full aggregation would be achieved by the second step, equalizing the difficulty rating across all four structures.

The noise tolerance model disagrees with both of the ranking described above. The right panel of Figure 2.2 plots EI (Definition 6) of the four information structures as a function of σ. EI is 1 for all information structures when $\sigma = 0$; in such a case, with no noise whatsoever, the full aggregation will indeed be achieved by the second step for all structures. However as σ increases, the graphs for the four structures separate with stochastic dominance.[21] In fact, the line for a three-state structure, namely, 3E, nestles

[21] As noise overwhelms actions, EI becomes negative: this supports Conjecture 2.2.4, which predicts that

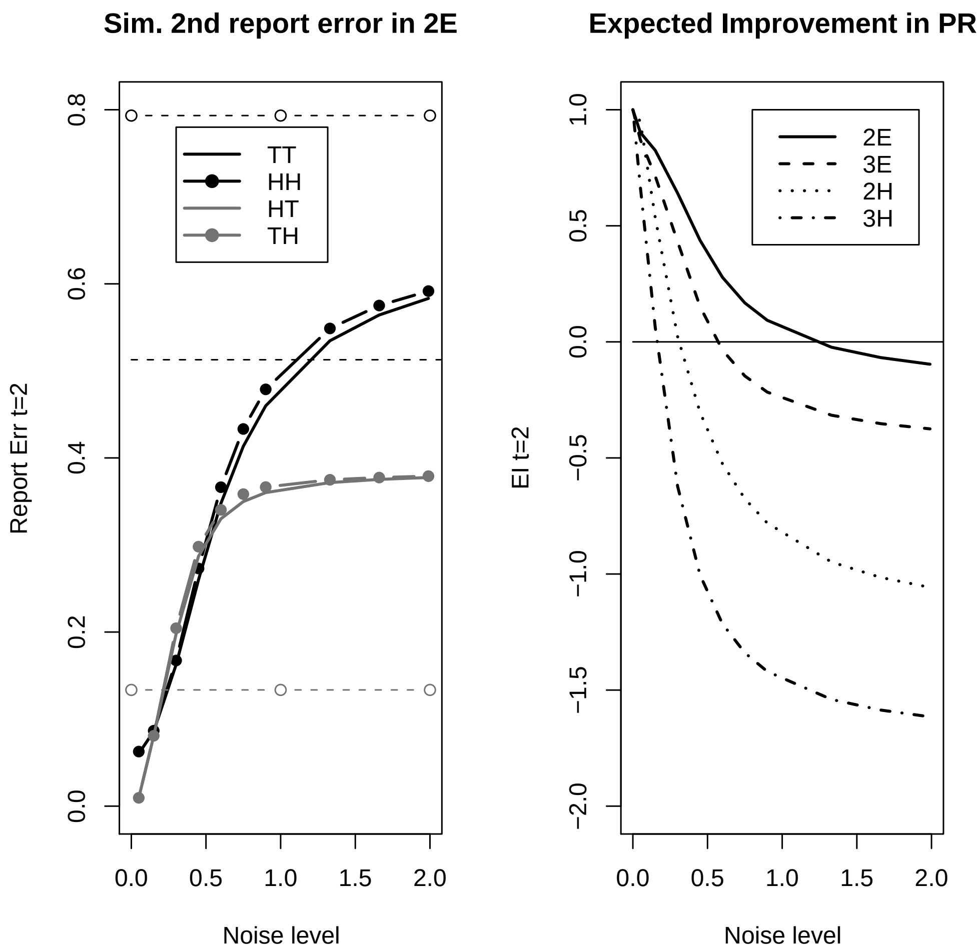

Figure 2.2: Simulated report error in 2E (L); and EI in all (R) at $t = 2$ as a function of σ

between the lines of a pair of two-state structures, 2E and 2H; thus, aggregation difficulty is not solely determined by the number of possible states.

It comes as no surprise that EI is monotonically increasing in σ, given that the simulated report errors increase monotonically in σ for all four structures.[22] As a result, there exists a unique σ^* such that $EI(\psi, \sigma) < 0$ for all $\sigma > \sigma^*$ in each information structure. Solving for $2E, 2H, 3E, 3H$ resulted in respective σ^* of 1.2, .3, .55, .15, which by Proposition 2.3.6(ii) implies the following hierarchy:

$$Noise\ Tolerance\ Ranking:\ 2E > 3E > 2H > 3H$$

priors yield lower losses than random noise over Δ_Θ.

[22]See the left panel of Figure 2.2 for 2E's graph; the remainder are contained in the Appendix.

This implies that PR^T is most likely to result in a prediction improvement in 2E and least likely in 3H. Referencing Proposition 2.3.6(iii), the stochastic dominance also implies that the size of prediction improvement will be largest for 2E and smallest for 3H. Simulation of EI for $t = 3, 4, 5$ is qualitatively similar and maintains the same ranking.

The numerical simulations presented lent support to the earlier conjectures and led to the following hypotheses:

- **H1:** Conjecture 2.3.5: Prediction improvement occurs more frequently when private signal draws are similar.

- **H2**: Proposition 2.3.6(ii): Prediction improvement occurs more frequently in those information structures ranked as more noise tolerant.

$$Pr(PI(D_{2E}) > 0) > Pr(PI(D_{3E}) > 0) > Pr(PI(D_{2H}) > 0) > Pr(PI(D_{3H}) > 0)$$

- **H3**: Proposition 2.3.6(iii): Average prediction improvement is greater in those information structures ranked as more noise tolerant.

$$\bar{PI}(D_{2E}) > \bar{PI}(D_{3E}) > \bar{PI}(D_{2H}) > \bar{PI}(D_{3H})$$

The next section will describe the design and implementation of an experiment to test these three hypotheses. This experiment will also test the following assumption of the behavioral model:

- **H0**: Actions are stochastic: the same individual will report different probability distributions, even when faced with the same private signal realization within the same information structure.

2.4 Experiment

2.4.1 Experimental Design and Implementation

The experiment consists of twenty-two rounds of a five-step PR^T. Subjects are randomly paired and assigned roles; each pair consists of a first-mover role, player A, and a second-mover role, player B. All subjects participate in all four treatments listed in Table 2.1, that is, 2E, 3E, 2H, and 3H; within these, the states of the world are represented as boxes and private signals as Black and White balls.[23]

After both A and B have privately viewed their ball draw, each player is asked to report a percentage chance corresponding to his perceived likelihood that the observed ball came from a given box.[24] Player A's first report a_A^1 is shown to B, but B's first report a_B^1 is not shown to A.[25] After observing a_A^1, B gets to revise his report; then, he announces his updated report a_B^2 to A. Next, player A revises and reports a_A^3. Yet another sequence of revisions follows: first from B (a_B^4), and then from A (a_A^5). Finally, the identity of the chosen box is revealed. To emphasize the number of revisions that each player undergoes, I refer to a_A^1 and a_B^1 as *initial reports*, a_B^2 and a_A^3 as *first revisions*, and a_B^4 and a_A^5 as *second revisions*.

The experimental sessions were conducted at California Institute of Technology during the summer of 2009. Using a within-subject design, thirty subjects, in random pairings, participated in the five-step PR^T for all information structures in Table 2.1. Data collection was conducted over the course of three sessions, each with ten subjects. After subjects arrived at the laboratory, the experimenter distributed the instructions then read them aloud.[26]

For PR^T, subjects interacted through zTree software written for this experiment.[27] A software screenshot is included in the Appendix. Subjects were given thirty seconds to submit initial reports; twenty seconds were allowed for each opportunity to revise their reports. Once the time limit had been reached, the experimenter announced, "Please enter your guess." This announcement was repeated, in fifteen-second intervals, until all subjects had entered their reports.

Each report was rewarded according to the Quadratic Proper Scoring Rule (QPSR): twenty-five cents was given for a correct answer. When a round concluded, subjects first observed their payoff for that round and then proceeded to the next round. At that point,

[23]To minimize the possibility that the visual display of the probabilities cues subjects to the difficulty level, information structures with equal likelihoods are excluded and two digit precision are included in every information structure.

[24]A subject in a two-state case only needs to enter the percentage chance for box X, and in a three-state case, the percentage chances for boxes X and Y: the software computes the percentage chance for the remaining box in each case.

[25]In fact, B's first report has no effect on the game: it is only elicited to learn about his initial belief.

[26]A full set of instructions is included in the Appendix.

[27]Pilot tests of the software were completed in the autumn of 2008. This data is available upon request.

they were randomly rematched. Each new pair was shown a new information structure with corresponding private signals. As each session comprised twenty-two rounds of posterior revisions, each subject made a total of sixty-six reports and earned between $19 and $22. The average session length was 75 minutes.

For the parameter set, a random-number generator produced 11 realizations of boxes and ball draws for the four information structures; hence, there were 44, not necessarily unique, aggregation problems. Another random-number generator ordered the 44 problems: grouping the first 22 as one parameter set and the last 22 as another. Experimental sessions alternated between the two parameter sets.

The presentation of results will be arranged as follows. Section 2.4.2 will focus on H0, providing evidence of stochastic behavior in initial reports and MLE estimates of noise level for each information structure. It will also disclose evidence that players are updating their own beliefs using their partners report. In Section 2.4.3 the noise levels estimated from initial reports will be used to simulate average report error at every step of PR^T. A comparison of simulated and actual report error follows. Section 2.4.4 evaluates report error in the context of baseline error through the prediction improvement metric. Evidence supportive of H1 will illustrate that the probability of beating baseline error is higher, indeed, for similar private draws than for different ones. I will then test both H2, the prediction of the noise tolerance ranking on relative frequency, and H3, the size of prediction improvement, over all four information structures. Out of the six possible pairings, strong evidence will present for four of them, with mixed results for the remaining two ($2E < 3E$ and $2H < 3H$). Section 2.4.5 will provide a joint test of the hypotheses through logistic regressions on the probability of positive prediction improvement and ordinary least-squares regressions on the size of the improvement.

2.4.2 Individual and Group Behavior

Over the course of 22 rounds of posterior revisions, each subject faces a total of 8 unique initial scenarios.[28] Subjects are therefore making decisions within the same scenario 2-4 times. When payoff disturbances are present in subjects' utility computation, subjects will give different responses to the same scenario. Indeed the data shows that a subject who

[28] A single bit of private information (Black or White) is given to a player in each of the four information structures.

faces the same scenario reports a different probability distribution 81% of the time.

However subjects are not just randomly choosing numbers to report in PR^T. Initial distance between reports are 0.348 (se=0.0155, N=660). 54.5% of first revisions and 37.4% of second revisions bring reports closer. The average convergence in reports are $d(a_A^1, a_B^1) - d(a_A^3, a_B^2) = 0.1425$ (se=0.014) for first revisions and $d(a_A^3, a_B^2) - d(a_A^5, a_B^4) = 0.0308$ (se=0.01) for second revisions. There is therefore strong evidence for belief convergence, which is consistent with the model of noisy bayesian updating.

I now provide a maximum likelihood estimate of the noise level. Assuming that subjects use Bayes rule to update the common prior to p_i^1 and that payoff disturbances are distributed $N(0, \sigma)$, the distribution of distance between report a_i^1 and the initial belief p_i^1 follows a stochastic process that depends on σ. As before, denote the set of full information posteriors and observed initial report distribution from N realizations of ω^n in ψ as $D^N = \{(p_{s^1}, \alpha_{s^1}), ..., , (p_{s^n}, \alpha_{s^n}), ..(p_{s^N}, \alpha_{s^N})\}$. Since B's initial report is elicited, we also have $\{\alpha_B^1, ..., , \alpha_B^n, ..\alpha_B^N\}$, doubling N.

Given realization $s^n = s_i^n s_j^n$ in ψ, subject i's initial belief over Θ is:

$$p^1_{s_i^n}(\theta) = \frac{\psi(\theta|s_i^n)}{\psi(s_i^n)}$$

Assuming independence across reports, the likelihood of observing initial reports of D^N in our experiment is:

$$l(D^N|\sigma) = \prod_{n=1}^{N} Pr_\sigma(\alpha_{s^n}|p^1_{s_A^n}) Pr_\sigma(\alpha_B^n|p^1_{s_B^n})$$

where

The estimation result of this model is given in Table 2.2. For comparison, we include the estimation result of the standard model rational behavior where $a_i^1 = p_i^1$. The likelihood is much lower, allowing us to reject the hypothesis that initial reports are chosen with $\sigma = 0$.

Estimating each information structure separately, we find that the MLE estimate for the noise parameter σ is 0.48 for 2E, 0.30 for 3E, and 0.33 for 2H and 0.21 H. All estimates of σ indicate that noise is present in subjects' actions. The estimates are highest for two state structures than three state structures, and within the same number of states, higher for the easier information structure. Interestingly, the same pattern appears in the time it

	σ	logL	2N
Stochastic	0.33	-3578.77	660
Deterministic		-4835.65	660
2E	0.48	-667.71	170
3E	0.30	-1086.63	160
2H	0.33	-608.42	160
3H	0.21	-1154.32	170

Table 2.2: Model estimates of variance of stochastic disturbances.

takes to make a decision. To make initial guesses, subjects spent 20.76 seconds (se=0.98) for 2E, 20.73 seconds (se=0.97) for 2H, 28.53 seconds (se=1.00) for 3E and 31.19 seconds (se=1.08) for 3H.[29]

Could risk preferences have been responsible for the deviation between initial report and initial beliefs? If subjects are updating using Bayes rule, no initial reports will indicate equal likelihood or degenerate probabilities. Hence reports of equal likelihood can be taken as evidence of risk aversion and degenerate probabilities could be taken as evidence of risk loving behavior. However, out of 660 initial reports, there were only 21 reports of 50-50 and 11 reports where each of the three states receives at least 30% chance. There were even fewer degenerate reports: less than 2% of initial reports place 95% or more weight on a single state.

Result 0: *Initial reports indicate stochastic behavior.*

2.4.3 Report Error

The black lines in Figure 2.3 plot the average error of experimental subjects' reports in PR^T. Initial report errors (a_A^1) are 0.343 (se=0.028) for 2E, 0.368 (se=0.023) for 3E, 0.256 (se=0.023) for 2H, and 0.245 (se=0.017) for 3H. In all the information structures, PR^T resulted in positive and significant error reduction.[30] However, the final report errors are positive and significant, which is consistent with the assumption that noise level does not

[29] Less time is needed to make a report for two-state structures than for three-state structures. Within the same number of states, subjects makes faster decisions in the easier information structure. First revisions took 14.09 seconds (se=0.74) in 2E, 17.25 seconds (se=0.93) in 2H, 22.89 seconds (se=0.86) in 3E and 24.51 seconds (se=0.91) in 2H. Second revisions took 10.15 seconds (se=0.62) in 2E, 11.16 seconds (se=0.71) seconds in 2H, 16.94 seconds (se=0.75) seconds in 3E and 18.6 seconds (se=0.80) seconds in 2H.

[30] Average reduction as measured in Euclidean distance are 0.068 (se=0.028) for 2E, 0.103 (se=0.025) for 3E, 0.059 (se=0.025) for 2H, and 0.037 (se=0.018) for 3H.

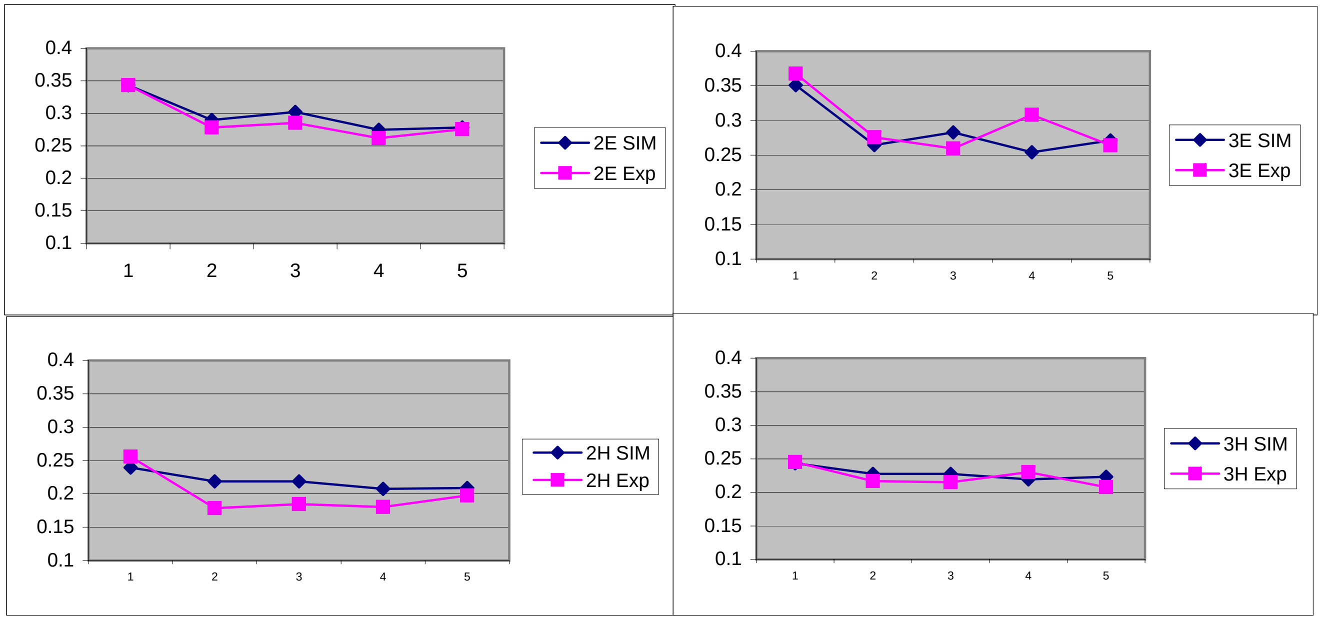

Figure 2.3: Report error for simulation vs experiment

decrease during a brief communication process.[31]

Does the average behavior of experimental subjects in PR^T follow the simulated behavior of myopic σ optimizers? The bright magenta lines in Figure 2.3 illustrates the simulated average report error in step $t = 1 - 5$ using estimated σ from Section 2.4.2 (0.48 for 2E, 0.30 for 3E, and 0.33 for 2H and 0.21 for 3H). Note that the noise levels σ were estimated using only the initial reports ($t = 1$) of A and B for each information structure.[32]

In every information structure, there is a close match between the simulated averages and the experimental data. As predicted by the simulations, the largest decrease in report errors happens at $t = 2$. Revisions from all following periods (t=3,4,5) result in little additional error reduction. The hypothesis that revisions do not bring reports closer to the full information posterior can be rejected at the $p < .001$ level for first revisions but not for second revisions (p=0.545).[33]

[31]Average final report error is 0.237. This is equivalent to predicting $\frac{1}{3}, \frac{2}{3}$ when the true probability is the equal likelihood $\frac{1}{2}, \frac{1}{2}$ or $\frac{1}{2}, 0.15, 0.35$ when the true probability is the equal likelihood $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$.

[32]For the experimental data, the average report error for the first step includes only player A's reports. Player B's initial reports are never shown to A. They are, therefore, not part of the posterior reporting process.

[33]When pooled across all information structures 55.5% of first revisions and only 37.4% of second revisions reduce the report error.

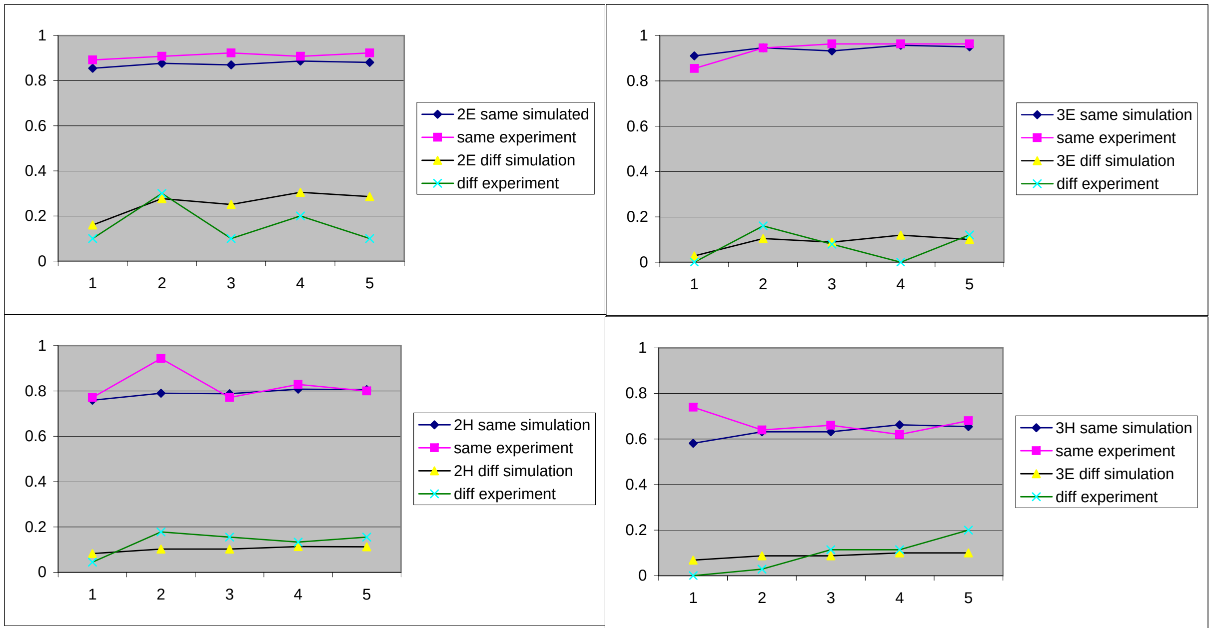

Figure 2.4: Frequency of prediction improvement for similar private signal draws and different private signal draws

2.4.4 Prediction Improvement

How do report errors compare to baseline errors? As previously discussed in the lead up to Conjecture 2.3.5, when A and B each draw a *Black* or *White* ball, their initial (pre-communication) beliefs are already closer to the full information posterior than the common priors are. The opposite is true when A and B receive opposing private signal, such as when A draws *Black* and B draws *White*.

The empirical frequencies of PI for each of the information structures are plotted with square and cross markers in Figure 2.4.[34] Note that when private signals are similar, at least 60% of reports at any step constitute prediction improvement. However, this frequency drops to 20% or less when private signals are different. Further revisions only improve the frequency of PI for different private signal draws but not much for similar draws. Binomial tests show that the difference in PI frequency given draws are significant for all t at $p < 0.001$ in each of the 4 structures, supporting Conjecture 2.3.5.

Result 1: *Prediction Improvement is more likely for similar private signal draws than for different draws.*

[34]Note that B's initial report a_B^1 is not used in the analysis of this section since this report is never observed by A.

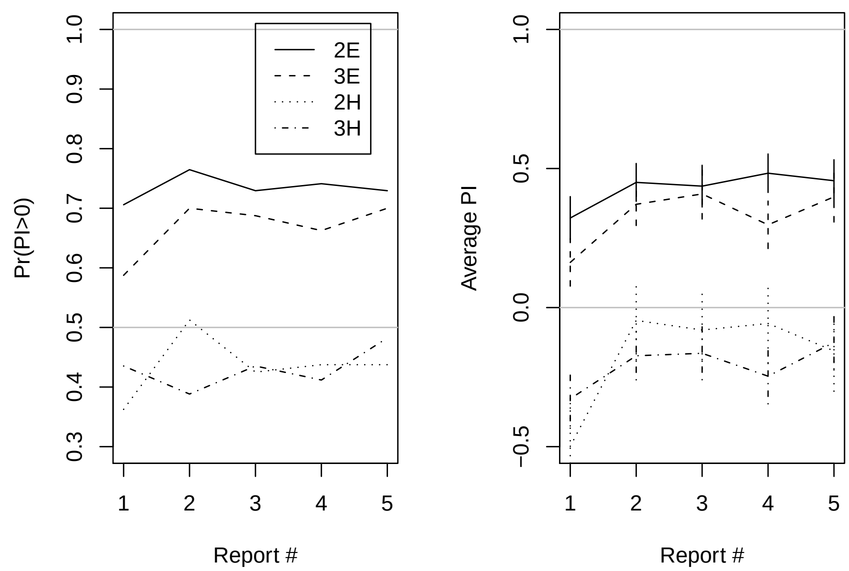

Figure 2.5: Frequency (L) and average size (R) of prediction improvement in experiment.

Does the model of myopic σ optimizer predict actual behavior? The diamond and triangle markers in Figure 2.4 plot numerically simulated frequency of PI using noise level estimated in Table 2.2. As in the graph of raw report errors in Figure 2.3 the frequency of PI from subjects' reports follows the simulation closely throughout all five steps of posterior revision.

The next question is whether the noise tolerance ranking of $2E > 3E > 2H > 3H$ (based on Proposition 2.3.6) accurately predicts the relative frequency and size of PI across the four information structures. There are several reasons why it may not. First, the ranking was based on the assumption that noise levels in different information structures were drawn from the same unknown distribution $G(.)$, while the MLE estimates show that noise levels were somewhat different. Secondly, while the ranking is computed using theoretical distributions implied by the information structure, the draws in the experiment were pre-generated randomly and are therefore different.[35]

Figure 2.5 illustrates empirical frequencies and average size of PI (with standard error bars) in each information structure at each step. The left panel shows that that PI is most

[35] Details of the differences can be seen in Table 2.4 in the Appendix.

often achieved by 2E, then 3E, followed by both 2H and 3H. Binomial tests confirm the following rankings of PI frequencies in all five steps at $p < 0.01$: $2E > 2H$, $2E > 3H$, $3E > 2H$, and $3E > 3H$. However, there is little support in the data for the ranking $2E > 3H$ and $2H > 3H$. The right panel show similar pattern on the size of prediction improvement. The fraction of possible improvement achieved in 2E and 3E are higher than in 2H and 3H ($p < 0.01$ for t-test), but the difference between 2E and 3E, and 2H and 3H, are not statistically significant.[36]

Result 2: There is strong evidence that frequency of prediction improvement follows the ex-ante ranking of $2E \geq 3E > 2H > 3H$. for $t = 2$. The evidence for $2 < t \leq 5$ is mixed for $2E > 3H$ and $2H > 3H$, implying a ranking of $2E \geq 3E > 3H \geq 2H$.

Result 3: At t=2, relative size of prediction improvement in the experiment is $2E \geq 3E > 2H \geq 3H$, which is consistent with the difficulty ranking. Again the evidence is mixed for $t = 5$ since the ranking implied by the experimental data is $2E \geq 3E > 3H \geq 2H$.

How does the difficulty ranking from this model of stochastic action compare to alternative models? In a model where subjects are perfectly rational, PI for all $t \geq 2$ is 1 in all information structures. If subjects ignore their private information, reports are centered around the prior and PI will remain around 0. Both imply a size ranking of $2H = 3E = 2H = 3H$, which is rejected by Result 2 and 3. A model where aggregation success depends on the number of possible states implies a ranking of $2E = 2H < 3E = 3H$, which is also rejected by Result 2 and 3. If subjects are following only their private information, beliefs will remain unrevised and the fraction of reports closer to p_s than the prior depends on the fraction of identical private signal draw (both A and B observe Black or White). This model results in the same difficulty ranking as the model of stochastic action ($2E > 3E > 2H > 3H$); this is expected since the slow updating of posterior is an implication of the stochastic action model. However Result 1,2, and 3 indicate that subjects are updating their posterior in response to their partners, thus rejecting the private information model.

2.4.5 Joint Test

As a joint test of Result 1, 2, and 3, Table 2.3 provides logistic regression on the probability of prediction improvement (Model 1-3) and OLS models (Model 4-6) on the size of improve-

[36]See Table 2.5 in the Appendix for more details of the statistical tests.

	Logistic Regression Y= (PI>0)			Ordinary Least Squares Y = PI		
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Intercept	-2.381 ***	-1.607 ***	-3.893 ***	-0.812 ***	-0.586 ***	-0.284 *
	0.358	0.352	0.909	0.068	0.071	0.159
3E	-0.150			-0.016		
	0.259			0.039		
2H	-0.354			-0.166 **		
	0.317			0.081		
3H	-1.383 ***			-0.428 ***		
	0.260			0.067		
Difficulty		-0.454 ***	0.372		-0.144 ***	-0.260 ***
		0.085	0.281		0.024	0.072
Same private signals	3.809 ***	3.706 ***	5.514 ***	1.242 ***	1.228 ***	1.318 ***
	0.178	0.166	0.497	0.066	0.070	0.076
First Revision	0.578 ***	0.572 ***	1.818 ***	0.219 ***	0.219 ***	0.160 **
	0.184	0.183	0.689	0.032	0.032	0.074
Second Revision	0.427 ***	0.422 ***	0.960 **	0.206 ***	0.206 ***	0.169 ***
	0.137	0.135	0.393	0.032	0.031	0.061
Round	0.038 ***	0.033 **	0.087	0.011 ***	0.010 ***	-0.018 *
	0.015	0.013	0.060	0.004	0.004	0.010
Same privsig *Difficulty			-0.668 ***			-0.035
			0.163			0.037
FirstRevision*Difficulty			-0.433 *			0.024
			0.220			0.029
SecondRevision*Difficulty			-0.188			0.015
			0.134			0.029
Round*Difficulty			-0.017			0.011 **
			0.018			0.004
N	1972	1974	1970	1972	1974	1970
Likelihood Ratio	1171.4	1159.3	1186.9			
R2	0.598	0.594	0.604	0.437	0.434	0.434

* p<.1, ** p<.05, *** p<.01

Table 2.3: Logit regression (OLS) on frequency (size) of prediction improvement.

ment. Model 1 and 4 include a dummy variable for each of the information structures. A higher noise-tolerance ranking indicates a lower difficulty in obtaining useful output from information aggregation; therefore, I will use the variable *Difficulty* below, where 1 indicates 2E and 4 indicates 3H, to provide an easily interpretable interaction term for the noise tolerance ranking. *Same private signal* is 1 when the pair receives the same private signal and 0 otherwise. *First Revision* is a dummy variable used twice, which indicates A's and B's first opportunities to revise their reports (a_B^2 and a_A^3). Second Revision equals 1 for a_B^4 and a_A^5. The intercept can therefore be interpreted as initial reports ($t = 1$) when private signals are different. The robust standard errors are clustered on individuals.

The increasingly negative coefficient on harder-ranked information structures and the negative coefficient on *Difficulty* both provide evidence that noise tolerance predicts aggregation success, in accord with H2 and H3. The coefficient on *Same private signal* is significant and positive across all models, supporting H1. The interaction of *Difficulty* and *Same private signal* in Model 3 reveals when private signals are the same, PR^T fails to beat the common prior more frequently in harder ranked information structures. When private signals are different, the process fails equally frequently in all information structures, however the loss compared to the common prior is larger in harder ranked information structures (Model 6).

The coefficient on *First Revision* is significant and positive across all models, indicating that first revisions unilaterally improve predictions. The interaction term in Model 3 shows that first revisions are less successful at increasing the frequency of prediction improvement in harder-ranked structures; however, the same does not hold true for first revisions' effects on the size of prediction improvement, shown in Model 6. The coefficient on *Second Revision* is less than *First Revision* across all six models, indicating diminishing marginal returns on additional revisions. There is some evidence of learning: Model 6 indicates that as subjects gain experience, they seem to improve their performance in harder information structures, even though the size of the effect is quite small.

2.5 Conclusion

Experimental evidence has long shown that mechanisms which are theoretically expected to fully aggregate dispersed private information, such as that existing in asset markets, polls, and other forecasting systems, in reality, generate probability distributions that only approximately converge toward the full information posterior. By explicitly modeling the stochastic component of individual behavior in the posterior revision process, I have shown the dependence of a basic unit of the group inference process on the underlying information structure. To measure the contribution of the aggregation process, I introduced a metric of prediction improvement which captures the fraction of the possible error reduction that can be achieved in a given information structure. Then, I provided a method of ranking for a set of information structures over expected prediction improvement.

I derived the noise-tolerance rankings to predict relative aggregation difficulty across four experimental environments. Next, I tested the rankings with a laboratory experiment wherein subjects participated in the posterior revision process with a random partner.

Several aspects of the behavior of experimental subjects were consistent with the behavior of myopic σ optimizers. First, subjects frequently revised their beliefs in the direction of their partner's reports. Secondly, first revisions decreased initial report error. Finally, second revisions had little impact on either belief convergence or the reduction of report error.

The experimental data showed that relative prediction improvement across the four information structures is consistent with the noise-tolerance ranking. This held true for both the frequency and size of prediction improvement. One result that did prove not to be statistically significant, though, was the difference between the information structure pair $2E - 3E$ and $2H - 3H$: $(p > 0.30)$.

These results suggest that analyzing the tolerance of mechanisms and environments for stochastic noise may be a useful line of research. When noise is introduced into behavior, details of the environment, which did not matter in the models of deterministic behavior, start to matter. This approach may yield a better understanding of the relationship between mechanisms and environments; thus, a way might conceivably be found to anticipate the contribution of a given mechanism implemented in a particular setting.

	2E	3E	2H	3H
Simulation				
$\mathbb{E}_s Pr(sA == sB)$	0.73	0.72	0.54	0.56
$\mathbb{E}_s BaseErr(\psi)$	0.48	0.42	0.21	0.18
Experiment				
$Pr(sA == sB \mid D^N, \psi)$	0.76	0.69	0.44	0.59
$BaseErr(D^N, \psi)$	0.56	0.54	0.17	0.18

Table 2.4: Theoretical and empirical realizations of private signals.

p-value	$2E > 3E$	$3E > 2H$	$2H > 3H$
fraction of $PI > 0$			
1. $t = 2$	0.11	0.00	0.01
2. $t = 5$	0.32	0.00	0.76
size of PI			
3. $t = 2$	0.221	0.006	0.229
4. $t = 5$	0.311	0.001	0.569

Table 2.5: One-sided binomial test (t-test) of fraction (size) of prediction improvement.

2.6 Appendix: Additional Figures

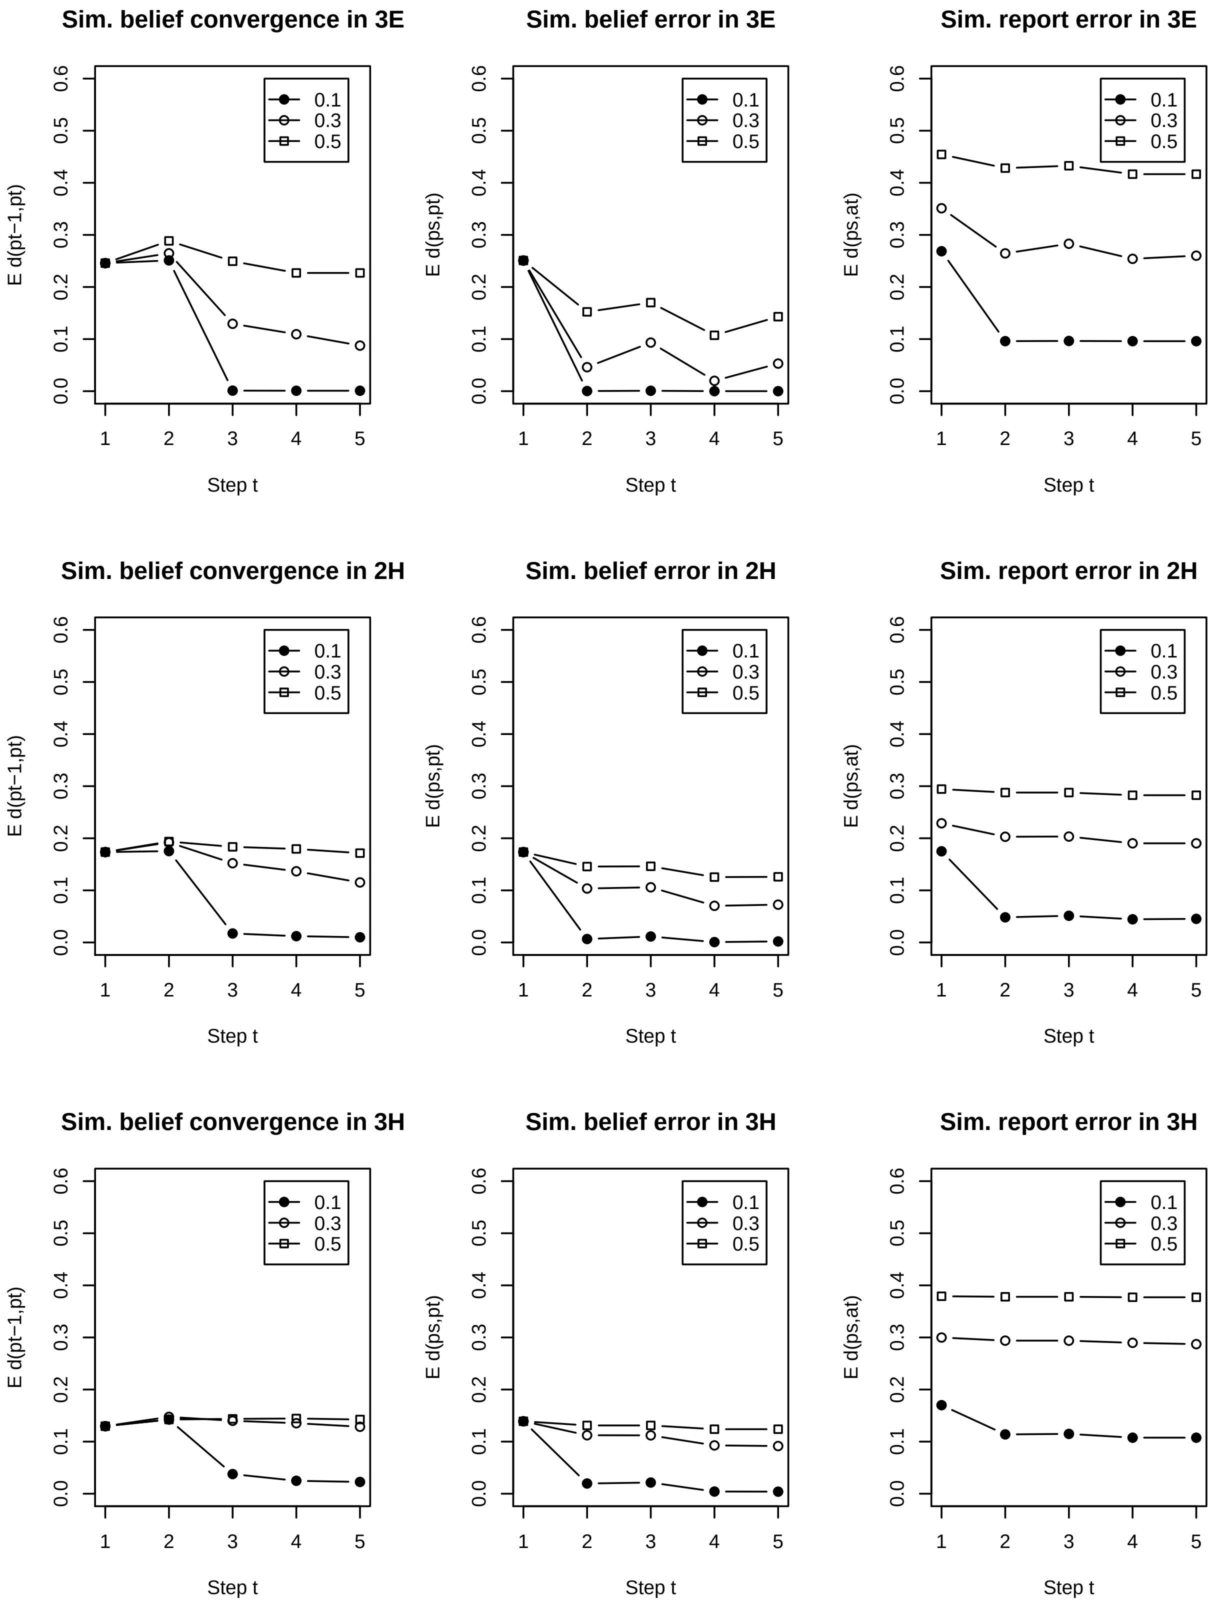

Figure 2.6: Simulated belief convergence, belief error, and report error in 3E, 2H, and 3H.

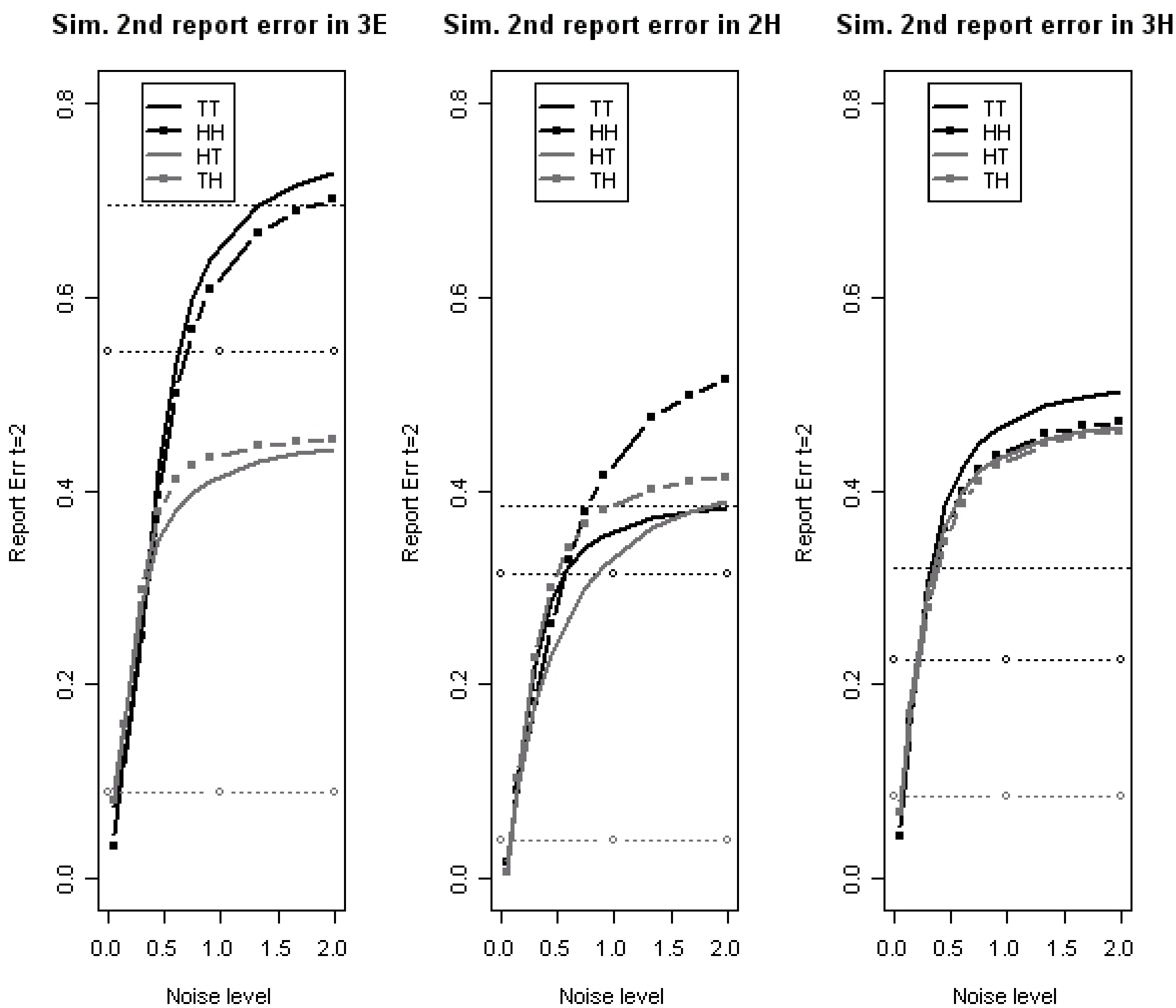

Figure 2.7: Simulated report error as a function of σ in 3E, 2H, and 3H.

2.7 Appendix: Proofs

Lemma 2.2.1

Proof. Let p_i^t be the vector of i's belief over Θ. Myopic payoff optimizer i chooses a_i such that a_i^t maximizes $argmax_{a_i^t} Q(a_i^t|\tilde{\theta})$. The expected payoff for reporting a_i^t when belief is p_i^t is $\mathbb{E}_{p_i^t} Q(a_i^t) = \sum_\theta p_i^t(\theta)[1-(a_i^t(\theta)-p_i^t(\theta))^2]$; this is maximized when $a_i^t = p_i^t$. Hence $a_i^* = (p_i^1, ..p_i^T)$. □

Proposition 2.2.3

Proof. Player j draws inference about s_i from a_i^t as follows:

$$Pr(s_i|a_i^t, K^t, s_j, \sigma) = \frac{Pr_\sigma(a_i^t|p_{s_i'}^t)\psi_{s_j}^1(s_i)}{\sum\limits_{s_i'} Pr_\sigma(a_i^t|p_{s_i'}^t)\psi_{s_j}^t(s_i')}$$

where each $p_{s_i'}^t$ is computed from $\psi_{s_i'}^t \in K_i^t$.
Updated beliefs about s_i drives the updating for the rest of the probability distribution.

$$\psi_{s_j}^{t+1}(\theta s_i) = \frac{\psi_{s_j}^t(\theta s_i)}{\sum\limits_{\theta'} \psi_{s_j}^t(\theta s_i')} Pr(s_i|a_i^t, K^t, s_j, \sigma)$$

Through simplification we arrive at

$$\psi_{s_j}^{t+1}(\theta, s_i) = \frac{\psi_{s_j}^t(\theta, s_i) Pr_\sigma(a_i^t|p_{s_i}^t)}{\sum\limits_{s_i'} \psi_{s_j}^t(s_i) Pr_\sigma(a_i^t|p_{s_i'}^t)}$$

Because player i do not update s_j from a_i^t, $\psi_{s_i}^{t+1}(\theta s_j) = \psi_{s_i}^{t+1}(\theta s_j)$. However, player i update his belief about j's possible beliefs (K_j^t). □

Proposition 2.3.6

Proof. (i) By Theorem 2.2.2, for all information structure ψ, $\mathbb{E}Err_{PRT}(s, \sigma = 0) = 0\, \forall s \in S$. Substituting into Def. 6 we arrive at $EI(\psi|\sigma = 0) = 1$ for all ψ.
When $\sigma = \infty$, actions becomes uniform random draws over Δ_Θ (pure noise) where

$$EI(\psi|\sigma = \infty) = 1 - \frac{\mathbb{E}_s \mathbb{E}_{a \in \Delta_\Theta} d(p_s, a)}{|\Delta_\Theta| \mathbb{E}_s BaseErr(\psi)} \tag{2.5}$$

By Conjecture 2.2.4, $EI(\psi|\sigma = \infty) < 0$. Since $EI(\psi|\sigma)$ is a normalization of $\mathbb{E}_s\mathbb{E}Err_{PR^T}(s,\sigma)$, Conjecture 2.3.4 implies that $EI(\psi|\sigma)$ is a monotonically decreasing function in $\sigma \in [0, \infty]$. Since $EI(\psi|\sigma)$ decreases monotonically from 1 to a negative number, a unique $\sigma^* \in \mathbb{R}+$ exists such that $\forall \sigma > \sigma^*$ $EI(\psi|\sigma) < 0$.

(ii) Let N_s be the number of times s occured in ${D_\psi}^N$. As $N \to \infty$, by Law of Large Numbers $\frac{N_s}{N} \to \psi(s)$ and hence $BaseErr({D_\psi}^N) \to \mathbb{E}_s BaseErr(\psi) = \sum\limits_{s \in S} \psi(s) d(p_s, p)$. Since σ is redrawn each time from distribution $g(.)$, as $N \to \infty$, $\frac{1}{N}\sum\limits_{n=1}^{N} ErrR(s^n) = \int \sum_s \psi(s)\mathbb{E}Err_{PR^T}(s,\sigma)dg(\sigma)$. Hence

$$\lim_{N \Rightarrow \infty} \bar{P}I_{PR^T}(D_\psi^N) = \int EI_{PR^T}(\psi|\sigma)dg(\sigma) \tag{2.6}$$

Let σ be distributed with cumulative density function $G(.)$. By Proposition 2.3.6(i),

$$Pr(\int EI_{PR^T}(\psi|\sigma)dg(\sigma) > 0) = Pr(\sigma < \sigma^*) = G(\sigma^*) \tag{2.7}$$

Since Eq. 2.6 and Eq.2.7 implies $\lim_{N \Rightarrow \infty} Pr(\bar{P}I({D_\psi}^N) > 0) = G(\sigma^*)$, $\sigma_1^* > \sigma_2^*$ implies $\lim_{N \Rightarrow \infty} Pr(\bar{P}I(D_{\psi_1}^N) > 0) > \lim_{M \Rightarrow \infty} Pr(\bar{P}I(D_{\psi_2}^M) > 0)$.

(iii) $\sigma_1^* > \sigma_2^*$ implies $EI_{PR^T}(\psi_1|\sigma_2^*) > EI_{PR^T}(\psi_1|\sigma_2^*)$. By first order stochastic dominance, this implies $\int EI_{PR^T}(\psi_1|\sigma)dg(\sigma) > \int EI_{PR^T}(\psi_2|\sigma)dg(\sigma)$. Therefore by Eq. 2.6, $\lim_{N \Rightarrow \infty} Pr(\bar{P}I(D_{\psi_1}^N) > 0) > \lim_{M \Rightarrow \infty} Pr(\bar{P}I(D_{\psi_2}^M) > 0)$. □

Lemma 2.3.1

Proof. By Selten, 1998, $\mathbb{E}_{p_i^t} Q(p_i^t) - \mathbb{E}_{p_i^t} Q(\alpha) = d(p_i^t, \alpha)$. This simplifies Eq.2.3 into:

$$R(\alpha|p_i^t) = \left\{(\epsilon_\alpha, \epsilon_p) \in \mathbb{R}^2 : \epsilon_\alpha - \epsilon_p \geq d(p_i^t, \alpha)\right\}$$

Let $\epsilon = \epsilon_\alpha - \epsilon_p$. Since ϵ_α and ϵ_p are iid distributed with mean 0 and variance σ, ϵ is distributed with mean 0 and variance $\sigma + \sigma = 2\sigma$.[37] Substituting this we arrive at $R(\alpha|p_i^t) = \left\{\epsilon \in \mathbb{R} : \epsilon \geq d(p_i^t, \alpha)\right\}$. Therefore

$$\int\limits_{R(a|p_i^t)} dF_\sigma(\epsilon) = F'_\sigma(d(p_i^t, \alpha))$$

[37] See Thm 4.2.12 Casella and Berger

Substituting this into Eq 2.4, we arrive at $Pr_\sigma(\alpha|p_i^t)$ as defined above.

(ii) Let $\delta, \delta' \in \mathbb{R}_+$ where $\delta > \delta'$. When $0 < \sigma < \infty$, $F'_\sigma(\delta) > F'_\sigma(\delta')$. Since the numerator of $Pr_\sigma(\alpha|p_i^t)$ decreasing in δ while the denominator is constant, $Pr_\sigma(\alpha|p_i^t)$ must be decreasing in $d(p_i^t, \alpha)$. $\square$

Proposition 2.3.2

Proof. When ϵ is distributed with $\mathcal{N}(0,\sigma)$, $F'(\delta)$ is $\frac{1}{\pi}\int_{\frac{\delta}{\sqrt{2}\sigma}}^{\infty} e^{-t^2}dt$ and its derivative with respect to δ is $-\frac{e^{-(\frac{\delta}{\sqrt{2}\sigma})^2}}{\sqrt{2\pi}\sigma}$. Hence,

$$\frac{F'_\sigma(\delta)}{F'_{\sigma+k}(\delta)} = \frac{\frac{1}{\pi}\int_{\frac{\delta}{\sqrt{2}\sigma}}^{\infty} e^{-t^2}dt}{\frac{1}{\pi}\int_{\frac{\delta}{\sqrt{2}(\sigma+k)}}^{\infty} e^{-t^2}dt} \tag{2.8}$$

Since $\delta, k \in \mathbb{R}_+$ implies $\frac{\delta}{\sqrt{2}\sigma} > \frac{\delta}{\sqrt{2}(\sigma+k)}$, Eq. 2.8 must be less than 1.

Taking the derivate of Eq. 2.8 with respect to δ we get

$$\frac{-\frac{e^{-(\frac{\delta}{\sqrt{2}\sigma})^2}}{\sqrt{2\pi}\sigma}\int_{\frac{\delta}{\sqrt{2}(\sigma+k)}}^{\infty} e^{-t^2}dt + \frac{e^{-(\frac{\delta}{\sqrt{2}(\sigma+k)})^2}}{\sqrt{2\pi}(\sigma+k)}\int_{\frac{\delta}{\sqrt{2}\sigma}}^{\infty} e^{-t^2}dt}{(\int_{\frac{\delta}{\sqrt{2}(\sigma+k)}}^{\infty} e^{-t^2}dt)^2}$$

Since the denominator is positive, the sign of the numerator (Eq.2.9) determines the derivative.

$$\frac{1}{\sqrt{2\pi}}[\frac{e^{-(\frac{\delta}{\sqrt{2}(\sigma+k)})^2}}{\sigma+k}\int_{\frac{\delta}{\sqrt{2}\sigma}}^{\infty} e^{-t^2}dt - \frac{e^{-(\frac{\delta}{\sqrt{2}\sigma})^2}}{\sigma}\int_{\frac{\delta}{\sqrt{2}(\sigma+k)}}^{\infty} e^{-t^2}dt] \tag{2.9}$$

We show by contradiction that the numerator is negative. Suppose this term is positive. Then:

$$\frac{e^{-(\frac{\delta}{\sqrt{2}\sigma})^2}}{e^{-(\frac{\delta}{\sqrt{2}(\sigma+k)})^2}}\frac{\sigma+k}{\sigma} < \frac{\int_{\frac{\delta}{\sqrt{2}\sigma}}^{\infty} e^{-t^2}dt}{\int_{\frac{\delta}{\sqrt{2}(\sigma+k)}}^{\infty} e^{-t^2}dt}$$

Which is a contradiction since the right hand side of the equation is less than 1 and the left hand side is greater than 1 since $e^{-(\frac{\delta}{\sqrt{2}\sigma})^2} > e^{-(\frac{\delta}{\sqrt{2}(\sigma+k)})^2}$. Hence Eq 2.9 must be negative, and Eq 2.8 is therefore decreasing in δ. $\square$

Lemma 2.3.3

Proof. (i) As a shorthand we write $d_\alpha = d(p_i, \alpha) < d_\beta = d(p_i, \beta)$, and $d_a = d(p_i, a)$. $\sigma' = \sigma + k$ for some $k > 0$. By Lemma 2.3.1 we write:

$$\frac{F'_\sigma(d_a)}{\sum\limits_a F'_\sigma(d_a)} > \frac{F'_{\sigma+k}(d_a)}{\sum\limits_{a'} F'_{\sigma+k}(d_{a'})} \; and \; \frac{F'_\sigma(d_b)}{\sum\limits_{a'} F'_\sigma(d_{a'})} < \frac{F'_{\sigma+k}(d_b)}{\sum\limits_{a'} F'_{\sigma+k}(d_{a'})}$$

Putting together and simplifying:

$$\frac{F'_\sigma(d_a)}{F'_{\sigma+k}(d_a)} > \frac{\sum\limits_{a'} F'_\sigma(d_{a'})}{\sum\limits_{a'} F'_{\sigma+k}(d_{a'})} > \frac{F'_\sigma(d_b)}{F_{\sigma+k}(d_b)}$$

From Prop 2.3.2(ii) we know $\frac{F'_\sigma(\delta)}{F'_{\sigma+k}(\delta)}$ is decreasing in δ, which implies that $d_b > d_a$.

(ii) Let $c(\sigma, \sigma'|\alpha, p_i^t) = Pr_\sigma(\alpha|p_i^t) - Pr_{\sigma'}(\alpha|p_i^t)$. Partition Δ_Θ into $L+ = \left\{\alpha|c(\sigma, \sigma'|\alpha, p_i^t) \geq 0\right\}$ and $L- = \left\{\beta|c(\sigma, \sigma'|\alpha, p_i^t) < 0\right\}$. By (i), $d_L \equiv \max\limits_{\alpha\in L+} d(p_i^t, \alpha) < d_H \equiv \min\limits_{\beta\in L-} d(p_i^t, \beta)$. Since $\sum\limits_\alpha Pr_\sigma(\alpha|p_i^t) = 1$, $\sum\limits_\alpha c(\sigma, \sigma'|\alpha, p_i^t)$ must be 0, which implies $\sum\limits_{\alpha\in L+} c(\sigma, \sigma'|\alpha, p_i^t) = -\sum\limits_{\alpha\in L-} c(\sigma, \sigma'|\alpha, p_i^t) = C \in \mathbb{R}_+$. Therefore $\sum\limits_{\alpha\in\Delta_\Theta} c(\sigma, \sigma'|\alpha, p_i^t) d(p_i^t, \alpha)$ must be smaller than $d_L k + d_H(-k) < 0$, which implies (ii). □

2.8 Appendix: Experimental Instructions and MATLAB scripts

Thank you for coming. Please pull up your divider.

[General]
This is an experiment in the economics of decision-making. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a considerable amount of money. You will be paid in cash at the end of the experiment.

[Choosing Lotteries]
Your have ten decisions to make. Each decision is a choice between "Option A" and "Option B" described in below. Even though you make ten decisions, only one of them will be used in the end to determine your earnings. We will roll a ten-sided die twice, once to select one of the ten decisions to be used, and a second time to determine what your payoff is for the option you chose, A or B, for the particular decision selected.

Suppose the first dice roll is 9. We will then be playing Lottery 9. We roll the dice again to find the result of the lottery. Suppose it is a 10. Then if you had circled Option A for Lottery 9, you will win $1.60, since 10>9. If you had circled Option B, you will win $0.10.

For each of these ten lotteries, which one do you prefer?
Please indicate your choice by circling it.

	A	B
Lottery 1:	1/10 of $2.00, 9/10 of $1.60	1/10 of $3.85, 9/10 of $0.10
Lottery 2:	2/10 of $2.00, 8/10 of $1.60	2/10 of $3.85, 8/10 of $0.10
Lottery 3:	3/10 of $2.00, 7/10 of $1.60	3/10 of $3.85, 7/10 of $0.10
Lottery 4:	4/10 of $2.00, 6/10 of $1.60	4/10 of $3.85, 6/10 of $0.10
Lottery 5:	5/10 of $2.00, 5/10 of $1.60	5/10 of $3.85, 5/10 of $0.10
Lottery 6:	6/10 of $2.00, 4/10 of $1.60	6/10 of $3.85, 4/10 of $0.10
Lottery 7:	7/10 of $2.00, 3/10 of $1.60	7/10 of $3.85, 3/10 of $0.10
Lottery 8:	8/10 of $2.00, 2/10 of $1.60	8/10 of $3.85, 2/10 of $0.10
Lottery 9:	9/10 of $2.00, 1/10 of $1.60	9/10 of $3.85, 1/10 of $0.10
Lottery 10:	10/10 of $2.00, 0/10 of $1.60	10/10 of $3.85, 0/10 of $0.10

[Making Guesses]

There are a total of 22 rounds. A sample round will proceed as follows:
There are two boxes, X and Y, with different number of BLACK and WHITE balls.
One of the boxes is randomly chosen with equal likelihood. A ball is drawn from the chosen box, the color is recorded for you, and the ball is returned to the box. A person in the room is randomly matched to you and is shown another (independently drawn) ball from the box. This is your private information – do not share them with anyone.

The setup of the round will be displayed on your computer screen the entire time.
You and your partner will both make initial guesses. Then one of you is chosen to "go first". Suppose you initial guess is randomly chosen to be shown to your partner. He/she will make a 2^{nd} guess, which will be shown to you (this means you get to see the 2^{nd} guess but not the initial guess). You then make a 2^{nd} guess.

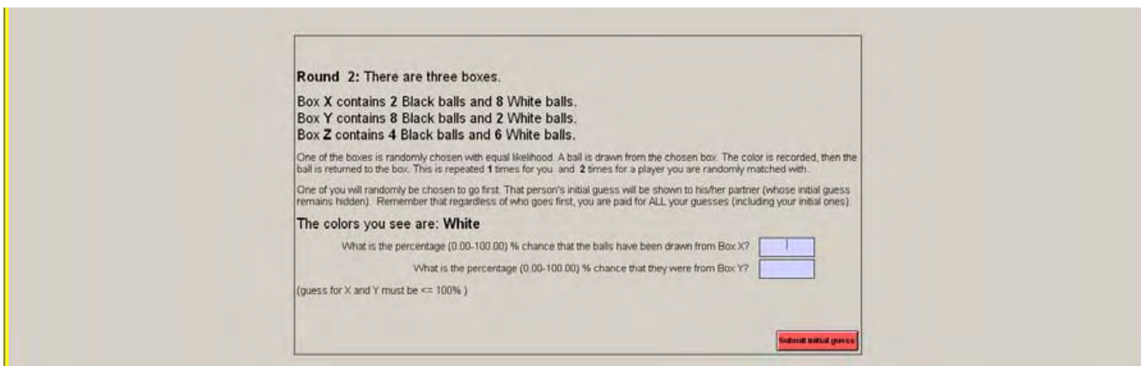

Your 2^{nd} guess is shown again to your partner who makes a 3^{rd} guess. You have the final chance to make a guess. Regardless of who goes first, your payment is determined by all three guesses that you have made.

Guesses are made in percentages, eg: here you enter "68.5" instead of "0.685"

The earnings formula is simple:
Given box S is chosen and r_S is your guess for that box,
you get $(1 – (1-r_S)^2)*$ 25 cents for each guess.

You will have a chance to make 3*22 = 66 guesses.

<u>When there are two boxes:</u>

Suppose the box is X and your guess for box X is r_X.

Your payment is $(1-(1-r_X)^2)*$ 25c

Suppose the box is Y. This implies your guess for box Y is $1-r_X$.

Your payment is $(1-r_X^2) *$ 25c

<u>When there are three possible boxes:</u>

The only difference is that now you are asked to enter r_X AND r_Y.

Suppose the box is Z and your guessed r_X and r_Y. This implies your guess for box Z is $1-r_X-r_Y$.

Your payment is (1-(r_X+ r_Y)^2)* 25c

WARNING:

If $r_X + r_Y > 100\%$, your entry is faulty and you will not earn anything for that guess.

After all a round is over, you will see how much you earn for each of your reports. A new round will start and a new set of boxes, private draws, and partners will be determined.

For total payment, please indicate in your receipt:

Show up fee: $5
Earnings: Sum of earnings from **[Making Guesses]**
Coupon: Result from lottery (**[Choosing Lotteries]**)

```
%Numerical simulation of the posterior revision process
%
%Author: Sera Linardi
%Last update: April 25 2010

%Experimental setting
Settings = [
   2, .4,.6, 0.79, 0.13, 1,1,0,0;   %2E
   3, .4,.5,.1, .15, .85, .6, 1, 1; % 3E
   2, 1/3, 2/3, 0.3, 0.6, 1,1,0,0;  % 2H
   3, 1/4,1/2,1/4, .4, .6, .8, 1, 1; %3H
];

numSet=size(Settings,1);
dFunc='Euclid';
errorF='normal';
baseActionSet = 0.0:0.1:1;

lambda=[0.1 0.3 0.5];
T=5;

reportErr = zeros(T,numlambda,numSet);
beliefDiv = zeros(T,numlambda,numSet);
EI = zeros(T,numlambda,numSet);
numlambda=length(lambda);

for i = [1 2 3 4]
    [numState, coin, outcome, nSig, actionSet, ps, gs, pFI, baseError] = setup(Settings(i,:), baseActionSet, dFunc);
    if(numState==3)
        actionSet(length(actionSet),:)=[1.00 0];
        actionSet = optimize3stateActionSet(actionSet,dFunc);
    end
    dist = generateDist(pFI, actionSet, dFunc); %distance of actionSet to actual posterior
    %idx refers to which column to look at given the player number
    [p1, idx1] = generatePartition(ps, sqrt(size(ps, 2)), 1);
    [p2, idx2] = generatePartition(ps, sqrt(size(ps, 2)), 2);

    for l=1:numlambda
         data = [length(actionSet), length(p1), sqrt(length(p1)), 1, T, lambda(l)];
         allOutput = ones(length(actionSet),length(p1), T+1);
         tic; allOutput = generateDistBetweenBeliefs(p1,idx1,p2,idx2, actionSet, data, dFunc, errorF, allOutput); toc;
         for t=1:T
            beliefDiv(t, l,i)=allOutput(t,:,T+1)*gs';
            reportErr(t, l,i)=gs*diag(dist(:,:)'*allOutput(:,:,t));
            EI(t,l,i)=(baseError-reportErr(t, l))/baseError;
         end
    end
end

%generate partition over all possible private signals given player number
%the common prior, and the number of signals.
%the indexing from this function controls the columns in the array of
%action frequency and beliefs that can be seen by a particular player.

function [post, idx] = generatePartition(commonP, nSignal, playerNum)
```

```
nDim = size(commonP);
post = commonP;
idx = ones(nSignal,nSignal);

for i=1:nSignal
   if (playerNum==1)
      idx(i,:)=(i-1)*nSignal+1:i*nSignal;
   elseif (playerNum==2)
      j=1:nSignal;
      idx(i,:)=(j-1)*nSignal+i;
   end
   temp = commonP(:,idx(i,:));
   total = sum(sum(temp));
   temp = temp/total;
   post(:,idx(i,:)) = temp;
end

% generateDistBetweenBeliefs generates all possible history of reports
   %
   % allOutput = generateDistBetweenBeliefs(p1,idx1,p2,idx2, actionSet,
   % data, dFunc, errorF, allOutput);
   % where data= [length(actionSet), length(p1), sqrt(length(p1)), t, T,
   % lambda(l)]
   %
   % pi,idxi is the belief of player i at time t
   % actionSet is the set of all possible reports
   % dFunc is the distance function
   % errorF is the distribution of report error
   % and allOutput is the data structure that is passed around for the
   % recursion
   %

function [allOut] = generateDistBetweenBeliefs(pT,idxT,pS,idxS, act, data, dFunc, errorF, allOut)

numaction = data(1); nsigTot=data(2); ndraw=data(3); cur=data(4); T=data(5); lambda=data(6);
periodsAhead = T-cur;
outcur=ones(numaction, nsigTot, periodsAhead)*0;
dist=zeros(numaction, nsigTot)*0;
pS_aold=pS*0;

priorS=generatePriorXY(pS,idxS, ndraw, 1) ;
priorT=generatePriorXY(pT,idxT,ndraw, 1);
allOut(cur,:,T+1)=distBeliefs(priorS,priorT);

afT= actionFrequency(act, dFunc, errorF, lambda, priorT(:,idxT(:,1)), ndraw);
allOut(:,:,cur)=copyGivenPrivSig(afT, [1 1], idxS);
generate = any((allOut(:,:,cur))'>0.01)';

if (cur<T)
   for a=1:numaction
         pS_a = updateCommonKnowledge(pS, idxS, afT(a,:), ndraw);
            if ((cur+1)==T)                  priorS_a=generatePriorXY(pS_a,idxS,ndraw,1);
               allOut(:,:,cur+1)= copyGivenPrivSig(actionFrequency(act, dFunc, errorF, lambda, priorS_a(:,idxS(:,1)),
ndraw), [1 1], idxT);
               allOut(cur+1,:,T+1)=distBeliefs(priorT,priorS_a); %%
            else
```

```
            allOut=generateDistBetweenBeliefs(pS_a, idxS, pT, idxT, act, [numaction nsigTot ndraw cur+1 T
lambda], dFunc, errorF, allOut);
          end

        for t=1:periodsAhead
          outcur(:,:,t)=outcur(:,:,t)+copyGivenPrivSig(allOut(:,:, cur+t), afT(a,:), idxT);
          dist(t,:)=dist(t,:)+copyGivenPrivSig(allOut(cur+t,:,T+1), afT(a,:),idxT);
        end
    end
    for i=1:periodsAhead
      allOut(:,:,cur+i)=outcur(:,:,i);
      allOut(cur+i,:,T+1)=dist(i,:);
    end
end
end

%distBeliefs generates Euclidean distance between beliefs over states
function[res]= distBeliefs(pT,pS)
    ndraw = length(pT);
    res = ones(ndraw,1);
    for i=1:ndraw
      res(i) = Euclid(pT(:,i),pS(:,i));
    end
end

% copyGivenPrivSig takes the distribution of action or beliefs that player j
% may have arrived at after updating based on player i's action and then
% multiplies it with the probability that player i plays that action.
% idx ensures that the copying takes into account of those private signal
% player i can distinguish from the ones he cannot.

function [result]=copyGivenPrivSig(original, multiplier, idx)
ndraw = size(idx,2);
result=ones(size(original,1), size(idx,1)*size(idx,2))*0;
for sT=1:ndraw % since S doesn't know if sT='T' or 'H', same action frequency there
     if(size(original,2)==ndraw)
        result(:,idx(sT,:))= original* multiplier(sT);
     else
        result(:,idx(sT,:))= original(:,idx(sT,:))* multiplier(sT);
     end
end
end


%generate priors over states from priors over private signals
function[prior]=generatePriorXY(p, idx, ndraw, copy)
    numState=size(p,1);
    if (copy)
      prior = ones(numState,length(p));
      for c=1:ndraw
        for r=1:numState
          prior(r,idx(c,:))=sum(p(r,idx(c,:)));
        end
      end
    else
      prior = ones(numState,ndraw);
```

```
      for c=1:ndraw
        for r=1:numState
          prior(r,c)=sum(p(r,idx(c,:)));
        end
      end
    end
end

% generate frequency distribution over all possible reports given errorF,
% the error distribution and player's prior.

function[actDist] = actionFrequency(actionSet, dFunc, errorF, lambda, prior, ndraw)
    actDist = ones(size(actionSet,1), ndraw)*0;
    dist = generateDist(prior, actionSet, dFunc);
    for j=1:ndraw
        actDist(:,j) = 1- cdf(errorF,dist(:,j), 0, lambda);
        actDist(:,j) = actDist(:,j)/sum(actDist(:,j));
    end
    actDist(find(actDist<0.00001))=0.00001;
end

% generate an array of distance according to the distance function given
% an array of target distribution and another array of approximate
% distribution
% hardcoded for euclidean distances
function [Dist] = generateDist(target, act, dFunc)
numA = size(act,1);
numSig =size(target,2);
Dist = zeros(numA, numSig);
if strcmp('Euclid',dFunc)
    for s=1:numSig
      for a=1:numA
        Dist(a,s)=Euclid([act(a,:), 1-sum(act(a,:))], target(:,s));
      end
    end
end

% Euclidean distance function
function [distance] = Euclid(pDist, qDist)
numElem=length(pDist);
distance=0;
for j=1:numElem
    distance=distance+(pDist(j)-qDist(j))^2;
end
distance=sqrt(distance);

%reads the array of experimental setting and setup the array of possible
%actions, baseline error, full information posterior (pFI), and distribution of
%private signal (gs)
function [numState, coin, outcome, nSig, actionSet, ps, gs, pFI, baseError] =setup(in, baseActionSet, dFunc)
    [numState, coin, outcome, nSignal] = getInfoStruc(in);
    actionSet = createActionSet(baseActionSet, length(coin));
    [ps, gs, pFI] = commonPrior(coin, outcome, nSignal, nSignal); %symmetric draws.
    nSig=nSignal+1;
    [baseError] = gs * calcBaseError(dFunc, coin, gs,  pFI); % matrix multiplication
end
```

```
 function [n,c,o,nsig1] = getInfoStruc(in)
    n=in(1);
    c=in(2:n+1);
    o=in(n+2:2*n+1);
    nsig1=in(2*n+2);
end

 % numA*numA vector of vector of [a(theta1), .., a(thetaNumState-1)..]
function a = createActionSet(baseActionSet, numState) % Subfunction
    if numState==2
         a= baseActionSet';
    elseif (numState==3)
        numA= length(baseActionSet);
        % for 3 states: the number of nested loops here has to be manually
        % added for more states
        a= zeros(numA*numA,numState-1);
        idx=1;
        for i=1:numA
            j=1;
            prob=baseActionSet(i)+ baseActionSet(j);
            while prob<1 && j<= numA
                prob=baseActionSet(i)+ baseActionSet(j);
                a(idx,:)=[baseActionSet(i), baseActionSet(j)];
                j=j+1;
                idx=idx+1;

            end
        end
        a=a(1:idx,:);
    else a=[];
    end

end

function [error] = calcBaseError(dFunc, coin, gs,  pFI)
    numSig = length(gs);
    numState=length(coin);
    error = ones(numSig,1);
    if strcmp(dFunc,'Euclid')
        for j=1:numSig
          error(j,1) = Euclid(pFI(1:numState,j), coin);
        end
    end
end

%optimize3stateActionSet rearranges array of possible reports
    %
    % sortedActionSet = optimize3stateActionSet(actionSet,
    % distanceFunction)
    %
    % group possible reports based on the confidence the report implies
    % about a particular state. This optimizes future loops for when
    % generateDistBetweenBeliefs is called. Hard coded for three states.

function actionSet = optimize3stateActionSet(actionSet, dFunc)
```

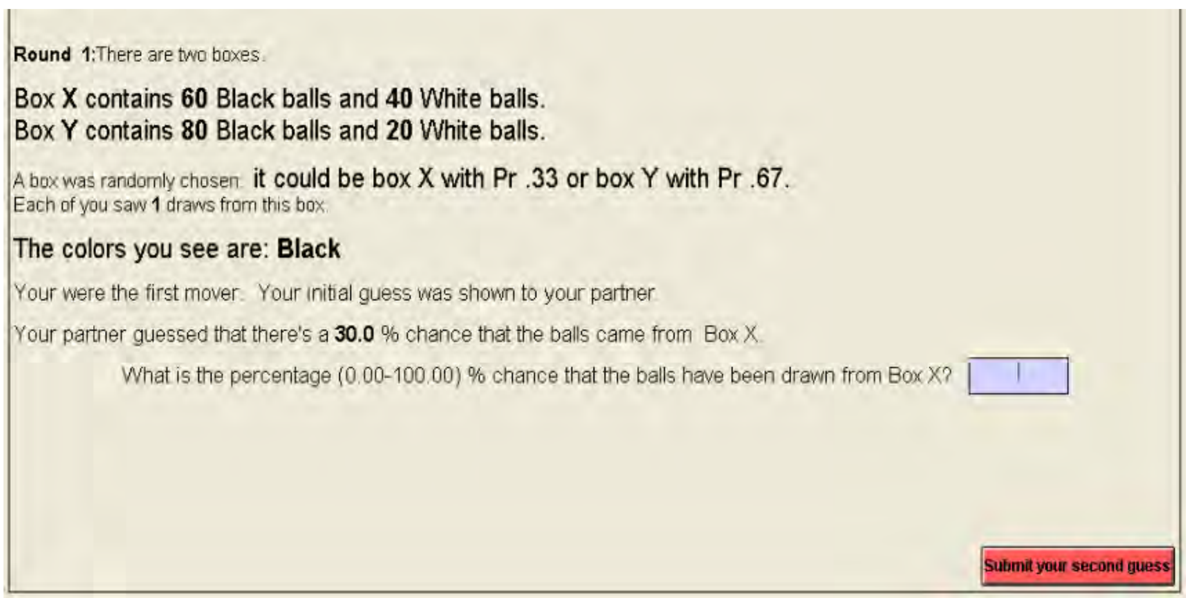

Figure 2.8: Screenshot of posterior revision software (ZTree)

Chapter 3

No Excuses For Good Behavior: Volunteering and Social Environment

with Margaret McConnell

3.1 Introduction

Economists have long been interested in the motivation behind prosocial behavior such as volunteering or donating money. Andreoni (1989) proposes a model in which individuals are intrinsically motivated by altruism to contribute to others' well-being. However, empirical evidence has shown that prosociality may be linked to public observation and can be crowded out by material rewards.[1] This evidence motivated several recent theoretical models where prosocial behavior is used as a signaling mechanism to gain social image benefits (Seabright 2004; Benabou and Tirole, 2003, 2006; Ellingsen and Johanssen 2008). In practice, prosocial behavior typically occurs in settings where multiple social mechanisms may be taking place simultaneously. Can the image signaling framework help us identify and manipulate components of a given social environment to encourage prosocial behavior?

We attempt to answer this question in the context of volunteering. Volunteering, an activity that involves 26.4% of the US population,[2] is crucial to the functioning of the non-profit sector. Studies consistently show that the value of individual volunteering is higher

[1] For example, Andreoni and Petrie (2004) and Rege and Telle (2004) find that removing the anonymity of gifts in a public goods game increases contributions. Frey and Jegen (2001) provide an early survey on crowding out. See Ariely, Bracha, and Meier (2007) and Carpenter and Myers (2009) for recent evidence.

[2] 61.8 million people, Bureau of Labor Statistics (2008), http://www.bls.gov/news.release/volun.nr0.htm.

than the value of household charitable giving.[3] We focus on two ubiquitous features of the volunteering environment. First, it is common knowledge that external circumstances can pose restrictions on some volunteers' ability to contribute time. These external restrictions are often difficult to verify, providing all volunteers with excuses for their own lack of contribution. Second, nonprofits often send representatives to informally supervise volunteers, under the assumption that these representatives' presence increases the pressure to contribute.[4]

We identify two social signaling mechanisms in the environment described above and theoretically derive predictions on contributed time using an extension of Benabou and Tirole's (2006) binary participation model. First, we predict that unverifiable excuses will dampen the stigma of not contributing. Removing excuses intensifies this stigma, and consequently, the image reward of working. Therefore the average time contributed will be higher in the absence of excuses. Second, the presence of a representative increases subjects' awareness of being observed, thus increasing the image reward of contributing. We predict that removing this 'monitor' will decrease the average contribution of volunteered time.

While little is known about labor contribution in the presence of excuses, a growing literature suggests that unconditional monetary transfers are less generous when others may not learn of a player's decision (Andreoni and Bernheim 2009; Tadelis 2007). On the other hand, existing literature provides conflicting predictions on the impact of a monitor's presence on labor. Dickinson and Villeval (2008), Falk and Kosfeld (2006), Frank and Schulze (2002), and Frey and Oberholzer-Gee (1997) argue that the presence of a monitor may be interpreted as distrust and decrease prosocial contributions. However, the demand effect literature posits that the desire to please authority figures drives laboratory subjects to be more altruistic when the experimenter is present.[5]

The impact of social environment manipulations on volunteering may be more complex than on monetary donations. First, unlike money, contributions of labor are multidimensional. Holmstrom and Milgrom (1991) have shown that incentives can increase the empha-

[3]For example, Independent Sector estimates that time volunteered in 2001 was valued at $240 billion (at $15.68 per hour) while household charitable giving was $153 billion: http://www.independentsector.org/programs/research/gv01main.html

[4]DellaVigna, List and Malmendier (2009) find that when individuals have the option to avoid being visited by a charity representative in person, their gifts are reduced by 30%.

[5]See Paulhus (1991), Levitt and List (2006), and Fleming, Townsend, Lowe and Ferguson (2007). Zizzo (2009) noted that vertical social pressure (experimenter) may be confounded with horizontal pressure (other subjects) in many studies.

sis on the rewarded dimension of a task to the detriment of unrewarded dimensions. Image rewards may encourage contributions of time, which are readily visible, but harm productivity. Second, monetary contributions are often studied in a static social environment, thus missing the dynamic changes that occur in a work environment over time. Third, given the higher degree of personal involvement inherent in labor contribution, manipulations that are effective in encouraging monetary contributions may not be effective in encouraging volunteering. In fact, Ellingsen and Johannesson (2009) find that fewer subjects in a bargaining game demand compensation for time investments compared to monetary investments.

We partner with the Los Angeles nonprofit School on Wheels (SOW)[6] to have lab subjects perform online internet search and data-entry to build SOW's database of educational resources, thus integrating the realism and context of volunteering into the controlled social environment of the laboratory. To test the effect of excuses, we utilize privately known random maximum stopping times that restrict the contribution of some subjects. The existence of this random mechanism provides excuses for subjects who do not face a stopping time (*unrestricted* subjects). To test our predictions about the role of a monitor, we use the experimenter as a representative of an authority figure.

We find that subjects volunteer less when external circumstances provide excuses for low contribution. Furthermore, we find evidence of differential departure patterns depending on the availability of excuses. Subjects are more likely to stop volunteering when others have stopped and are more likely to leave in clusters *only* in the absence of excuses. This behavior is consistent with stigma avoidance but not with framing, anchoring, conditional cooperation, or conformity.

We do not find, however, that removing the experimenter decreases volunteered time. Consistent with evidence from Frank (1998) that subjects are not sensitive to the payoff of the experimenter, our subjects do not appear to be affected by the presence of the experimenter. Subjects do, however, care about other subjects: the likelihood that an individual continues to volunteer increases with the number of subjects that are still volunteering. Average productivity, as measured by database entries per minute, remains unaffected throughout all the treatments, suggesting that the social environment can be manipulated to increase the average quantity of contribution without affecting average quality. Overall, our findings suggest that while image signaling mechanisms can increase prosocial behavior,

[6]School on Wheels provides tutoring for homeless children: http://www.schoolonwheels.org/

the effectiveness of these strategies depends on the details of the social environment.

The paper proceeds as follows. In Section 3.2 we describe the theoretical model and predictions for our experimental treatments. Section 3.3 describes our experimental design and the survey instrument. In Section 3.4 we present the results and Section 3.5 concludes. Proofs for Section 3.2 and experimental materials (instructions, software screen shots, and survey questions) can be found in the Appendix.

3.2 Theoretical Framework

A typical volunteering setup involves a representative from an organization and a group of potential contributors. Everyone knows that external circumstances may restrict some individuals' ability to contribute; these obstacles occur privately and are unverifiable. We present an extension of Benabou and Tirole's (2006) binary participation model[7] to illustrate the image signaling mechanisms that may be present in this environment. Formal details and proofs can be found in the Appendix.

Let v be an agent's intrinsic motivation to volunteer. We model v as a random variable with distribution function $g(v)$ and an associated density function $G(v)$. Let $x > 0$ be the visibility of volunteering, which represents an agent's awareness of being observed.

Following BT let the decision to volunteer be a binary choice $a = \{0, 1\}$. Let C be the cost of volunteering. An individual with type v who faces a choice to volunteer with visibility x has the following utility for volunteering:[8]

$$u(a = 1) = v - C + x(E(v|a = 1) - E(v|a = 0)) \tag{3.1}$$

Individuals participate if $v \geq C - x(E(v|a = 1) - E(v|a = 0)) \equiv v^*$ where the equilibrium threshold of altruism v^* is implicitly defined by the equation:

$$v^* - C + x(E(v|v \geq v^*) - E(v|v < v^*)) = 0 \tag{3.2}$$

BT show that when the distribution of altruism $g(v)$ is decreasing or constant in v,[9] there is a unique equilibrium threshold v^*. Without this assumption (e.g when $g(v)$ is increasing or

[7] Henceforth BT.

[8] Note that $u(a = 0) = 0$.

[9] There are fewer highly altruistic types in the population than less altruistic types.

unimodal in v), multiple equilibria exists for a large range of C and $g(v)$, making it difficult to derive theoretical predictions. We will therefore make the simplifying assumption that $g'(v) < 0$ for the rest of this paper.

We introduce excuses by considering some probability $\delta \in [0, 1]$ that individuals are prevented from volunteering by (unverifiable) external circumstances. When there are excuses for not participating, it is straightforward to infer the type of agents who participate, but more difficult to determine the type of agents who do not. This is because there are two reasons that an agent might not participate: with probability δ he has been prevented by circumstances, and with probability $1-\delta$ he is not altruistic enough to participate. In other words, unverifiable external circumstances provide *excuses* for all agents to not participate.

More formally, let $\Delta(v^*|x) = x(M^+(v^*) - M^-(v^*))$ be an agent's **image reward** from participating, where $M^+(v^*) \equiv E(v|v \geq v^*)$ is the *honor* for participating and $M^-(v^*) \equiv E(v|v < v^*)$ is the *stigma* of not participating. Credible excuses do not change the honor of participating but lessen the stigma of not participating:

$$M^-(v^*|\delta) \quad \equiv \quad \frac{\delta E(v) + (1-\delta)G(v^*)E(v|v < v^*)}{\delta + (1-\delta)G(v^*)} \tag{3.3}$$

We show that when excuses are available, participation can still be described by a unique equilibrium v^*. We then extend this binary participation framework to model an agent's contribution of time. This extended model identifies two image signaling mechanisms in the volunteering environment described earlier. First, the availability of excuses (δ) reduces the stigma of low contribution, thus reducing the image rewards from contributing time. Second, assuming that the presence of an authority figure increases an agent's awareness of being observed, the image reward of volunteering will increase when the experimenter is present. We formally derive two predictions on the impact of altering this social environment on average time volunteered. Since an agent's productivity has no image signaling value, it should remain unaffected by image treatments. All proofs are in the Appendix.

Lemma 3.2.1. *$\Delta(v^*|\delta, x)$ is increasing in v^*.*

Lemma 3.2.2. *Let $\bar{a}(\delta, x) \equiv N(1-G(v^*))$ denote the total participation among a population of N individuals.*

(i)Removing excuses increases participation.

$$0 = \delta < \delta' \Rightarrow \bar{a}(\delta, x) > \bar{a}(\delta', x)$$

(ii)Reduced monitoring decreases participation.

$$0 < x < x' \Rightarrow \bar{a}(\delta, x) < \bar{a}(\delta, x')$$

We now extend this binary participation model to our volunteering setup. Suppose there are t level of contributions from 1 minute up to a maximum of T minutes. Let $C(t)$ be the cost function for contribution level t where $C'(t) \geq 1$ (costs do not decrease over time). Let v_t^* be the threshold type for participation level t. Individuals contribute at level t if:

$$u(t) = vt - C(t) + \Delta(v_t^* | \delta, x) \geq 0$$

Treating each individual as facing t binary participation decision, let $v^* = (v_1^*, .., v_t^*, .., v_T^*)$ be the equilibrium threshold types induced by environment (δ, x). We show that higher levels of participation induce strictly higher thresholds than lower levels of participation; in other words individuals who do not choose to volunteer in level t will also not participate in level t' where $t' > t$. The monotonicity of v_t^* allows total time volunteered to be computed in intervals. This allows us to extend Lemma 3.2.2 to t levels of contribution.

Lemma 3.2.3. *Level t threshold type is strictly higher than level $t-1$ threshold type.*

$$v_t^* < v_{t-1}^*$$

Letting N be the total number of agents in the population, total time volunteered is:

$$0 = \delta < \delta' \Rightarrow \bar{a}_T(\delta, x) \equiv N \sum_{t=1}^{T-1} t(G(v_{t+1}^*) - G(v_t^*))$$

Proposition 3.2.4. *In a volunteering setup involving T levels of participation,*
(i) ***Excuses Prediction:*** *Removing excuses increases time volunteered.*

$$0 = \delta < \delta' \Rightarrow \bar{a}_T(\delta, x) > \bar{a}_T(\delta', x)$$

(ii) ***Monitoring Prediction:*** *Reduced monitoring decreases time volunteered.*

$$0 < x < x' \Rightarrow \bar{a}_T(\delta, x) > \bar{a}_T(\delta, x')$$

3.3 Experiment

To analyze how a particular social environment affects contributions of labor, the immediate impact of the environment needs to be isolated from other contributing factors. This is difficult to do in empirical studies, since volunteers may be motivated by long term concerns such as networking or resume building. The laboratory setting offers some advantages over the field in identifying short term image concerns: fewer opportunities for strategic reputation building, ease of constraining the audience, and precision in measuring the quantity and quality of contribution. To integrate realism and context into the controlled lab setting, we partner with the nonprofit School On Wheels (SOW) to design a real volunteering task.

3.3.1 Experimental Design

Subjects received an email publicizing an opportunity to participate in an experiment on decision making that did not mention volunteering.[10] The experiment consisted of two stages: training and volunteering.

The training session lasted 15-20 minutes. The experimenter started by introducing SOW and distributing SOW promotional materials.[11] After all subjects indicated they had adequate time to read the materials, the experimenter explained the volunteering task. SOW requested help in building a database of educational resources. This task consisted of doing internet searches and entering the information into a database; up to seven entries (subject, website address, grade level, etc) could be made per resource. Each subject received a task sheet listing the areas in the database assigned to them. Subjects were aware that they were all working on different portions of the database and that their work would not be redundant.[12] Subjects then practiced the task by performing one directed internet search

[10] Recruitment follows standard CASSEL (UCLA experimental lab) protocol.

[11] Promotional materials included SOW website, a People magazine article on SOW and a thank you letter from SOW's lead volunteer coordinator to the lab volunteers.

[12] The list contained several choices of grade levels and school subjects that has been randomly drawn, then adjusted to minimize overlap between subjects. We do this to increase the independence of the value of an individual's database entries from other subjects', thus decreasing concerns for free riding present in traditional public goods experiments.

and one data entry task. After everyone had completed the training session, we announced they had earned their show up fee ($20) and were free to go; if they chose to, they could stay in the lab and volunteer for SOW by performing the task they had just practiced. Subjects were informed that they were free to leave at any point and that the lab would be available for the next 90 minutes. We clearly stated that no additional monetary incentives would be forthcoming.

3.3.2 Treatments

All subjects in a session were assigned to one of the three treatments described below. See the Appendix for the script of instructions read to subjects.

Baseline=Excuses+Monitored

Excuses: A random mechanism embedded in the database software provided subjects with excuses to quit volunteering. Subjects clicked on a button on their screen to 'roll a die' after the training session. This die determined an individual's maximum time limit; a subject could stop at any point before the time limit but could not make any further database entries afterwards. This random mechanism introduced the probability δ of being prevented from working by external circumstances described in Section 2. Subjects were aware that each person could be limited by the randomly determined maximum time but were unaware of the true probability distribution of time limits. This approximates the natural occurrence of excuses where the true distribution of obstacles to prosocial behavior is unknown; all that is known is that $E(\delta) > 0$.

In our experiment, $\delta = 0$ with probability $\frac{2}{3}$, ensuring that a large share of the data was generated from subjects who had no time restriction and could be compared directly to subjects in the Remove Excuses treatment (see below). In order for it to be credible to subjects that there was a randomly generated stopping point, we set $\delta = 1$ with probability $\frac{1}{6}$, meaning that some subjects may leave the lab right away. The remaining $\frac{1}{6}$ of subjects received a time limit randomly chosen between 1 and 90. Neither the experimenter nor other subjects in the room know for certain if a subject had stopped by choice or because of the random mechanism.

Monitored: The experimenter stayed at the front of the room throughout the entire session and answered subjects' questions in person.[13]

Remove Excuses: No Excuses + Monitored

No Excuses: In this treatment, the random mechanism was disabled. After training, subjects were told that they could stay in the lab and volunteer for any amount of time they chose, up to 90 minutes.

Remove Monitor: Excuses + Unmonitored

Unmonitored: In this treatment, the experimenter left the room after training. In case questions about lab protocol or the volunteering task arose throughout the experiment, subjects could initiate contact with the experimenter through an anonymous chat software. Subjects randomly selected chat IDs out of a paper cup, thus fully assuring that their identities were protected from the experimenter.

3.3.3 Implementation

Pilot tests of the laboratory experiments took place at Claremont McKenna College in 2007 and the full set of experiments was run at UCLA[14] in Spring 2008 and Spring 2009.[15] The full set of experiments were run as 13 separate sessions with a total of 156 subjects. We ran 4 sessions of Remove Excuses, 5 sessions of Baseline, and 4 sessions of Remove Monitor; the average number of subjects per session in each treatment is 12.[16] We consider three outcome variables: the number of minutes worked by subjects, the number of entries completed and the number of entries completed per minute.

Over the course of running the experiments, subject volunteers completed a database of lesson plans before continuing on to educational activities.[17] The change of task was necessary to ensure that subjects' volunteered efforts continued to be useful for the organization. All data analysis controls for the task change.

[13] A lab technician was available to deal with computer problems if they arose.

[14] We attempted to replicate our experiment with actual SOW tutors, however logistical restrictions resulted in inadequate participation.

[15] The experiments ran at Claremont include only a subset of the treatments discussed in the paper. The pilot results support the findings of this paper and are available upon request.

[16] See Table 2 for session level statistics.

[17] The complete database of the results of subjects' volunteer work is available at http://www.hss.caltech.edu/~slinardi/data.xls

After the experiment, we collected data from subjects on demographic characteristics that have been found to be correlated with prosocial behavior.[18] To control for past volunteering experience, we ask subjects to report the length of time since their last volunteering experience and to rate that experience. We also asked them to rate the value of the lab volunteering task. In order to establish a measure of subjects' sensitivity to being paid, we asked them if they would prefer to work for an organization that pays volunteers for their time. Lastly, we asked the subjects to report the number of people in the room they knew by name to control for the relevance of social connections or peer pressures. The data collection was conducted by an online survey; subjects were automatically directed to that page when they clicked on a 'Finish Volunteering' button on the database software.

3.4 Results

Among our 156 subjects, 121 subjects were not affected by the random stopping time, receiving a time limit of 90 minutes. We classify these subjects as *unrestricted* and the remaining 35 subjects as *restricted*. Unless indicated otherwise, the data analysis focuses on comparing the behavior of *unrestricted* subjects across the treatments.[19]

3.4.1 Consistency of Lab Behavior with Natural Volunteering Behavior

To check whether the experimental setting induced behavior consistent with volunteering behavior in a natural setting, we perform several robustness checks. First we examined output to verify that subjects were actually working during the experiment. Figure 3.1 shows the relationship between the number of minutes worked and the entries completed. The strong positive trend between minutes worked and entries completed suggests that subjects were actually working instead of merely pretending to work.[20]

We then examine the relationship between the number of minutes worked and their valuation of the lab volunteering task. Consistent with evidence on the role of intrinsic

[18]See Mellström and Johannesson (2008), Schady (2001), and Freeman (1997) for gender and Brooks (2006) for religious activity.

[19]Excluding *restricted* subjects does not introduce selection effects since these subjects were randomly chosen by our mechanism. A duration model of the full sample with controls for time restrictions is included in Table 3.5.

[20]At the end of the experiment, we manually checked browser histories and found only 5 cases of internet usage unrelated to the volunteering task.

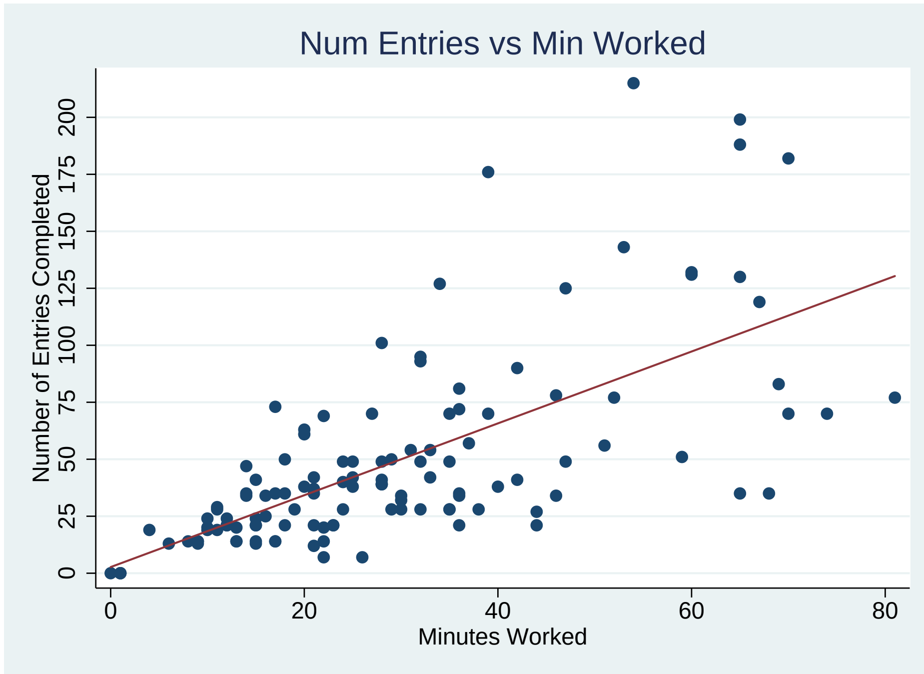

Figure 3.1: Amount of work completed is increasing in minutes volunteered

altruism,[21] the higher subjects rated the task, the longer they work (Figure 3.2).

3.4.2 Quantity of Contribution: Time Volunteered

Subjects exhibited a wide range of behavior in the experiment, with some subjects leaving right away while others remained to volunteer for nearly 90 minutes. Table 3.4.2 shows the average minutes volunteered in each of the three treatment groups. Figure 3.4.2 presents a comparison of the empirical distributions of minutes volunteered.

	Remove Excuses No Excuses Monitored	Baseline Excuses Monitored	Remove Monitor Excuses Unmonitored
Average	38.76	20.02	26.97
Standard Error	(3.06)	(1.78)	(2.19)
N	49	41	31

Unrestricted subjects only

Table 3.1: Average minutes volunteered per treatment

Consistent with the **Excuses Prediction**, removing excuses increased the average min-

[21]Finkelstein (2008) found that self reports of satisfaction predicted time spent by hospice volunteers.

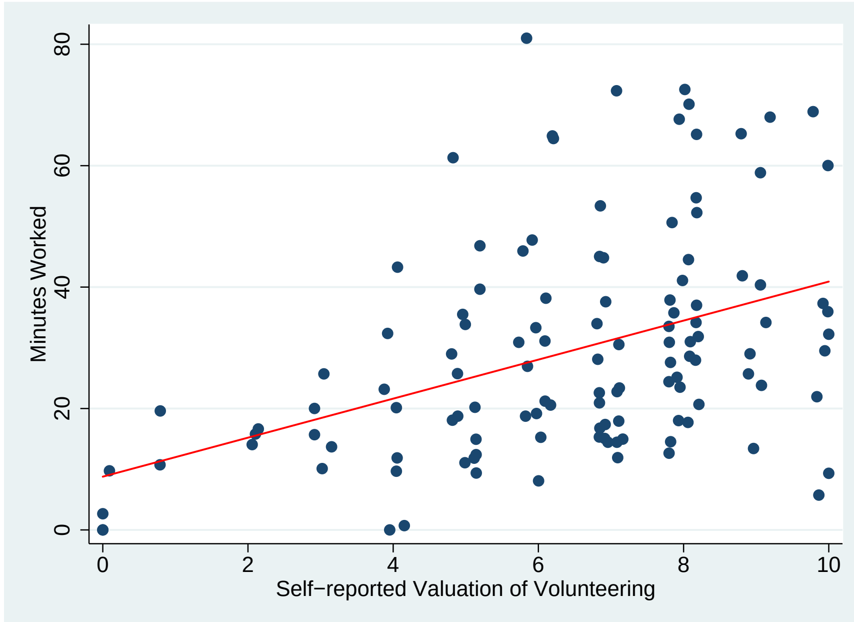

Figure 3.2: Self-reported value of work is increasing in minutes volunteered

utes volunteered. The difference between Remove Excuses and Baseline is positive and statistically significant at the 1% level using a non-parametric Wilcoxon (Mann-Whitney) test ($z = 4.26$). In contrast, the **Monitoring Prediction** was not supported by the data. The average minutes volunteered in Remove Monitor was significantly higher (at the 5% level) than minutes volunteered in Baseline (Mann-Whitney test statistic $z = 2.41$).

The cumulative density graph in Figure 3.4.2 tells the same story. The distribution of minutes worked in the Remove Excuses treatment and in the Remove Monitor treatment stochastically dominates the distribution of minutes worked in the Baseline.

Table 3.4.2 reports session level summary statistics for all 13 sessions. We see a consistent pattern of higher average minutes worked in Remove Excuses treatments. A Mann-Whitney test for the difference in average minutes worked at the session level across Excuses and Remove Excuses treatments yields z=1.39 (p-value of 0.08 for a one-sided test). The effect of removing excuses appears robust to localized social dynamics occurring at the level of the experimental session. On the other hand, we do not see a consistent pattern of higher average minutes worked in the Remove Monitor treatment compared to the Monitored treatments (Mann-Whitney $z = 0.77$).

Table 3.3 reports the regression results examining **Excuses Prediction** and **Monitor-**

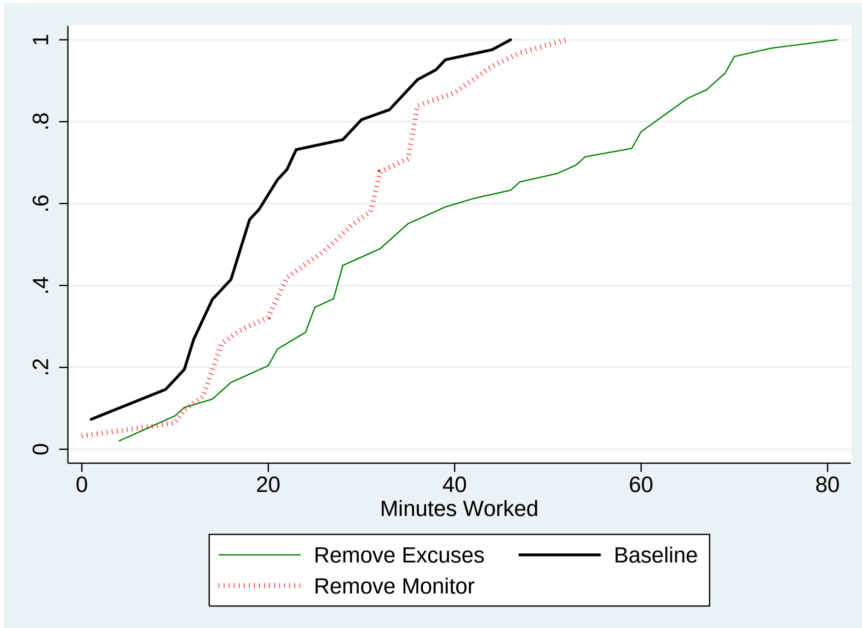

Figure 3.3: CDF of minutes volunteered

Treatment	Session Averages (Minutes worked)	Standard Deviation	Min	Max	Number of unrestricted Subjects	Total Subjects	Average of Session Averages
Excuses Monitored (Baseline)	18	13	1	39	6	12	
	18	9	9	35	10	13	
	31	9	18	44	6	11	
	22	14	1	46	10	13	
	15	7	1	23	9	12	
							21
Excuses Unmonitored (Remove Excuses)	26	21	0	52	5	10	
	28	11	11	42	7	12	
	29	13	13	47	10	13	
	25	7	15	36	9	11	
							27
No Excuses Monitored (Remove Monitor)	53	14	32	74	10	10	
	23	15	4	53	16	16	
	67	6	60	81	10	10	
	25	5	15	35	13	13	
							42

Unrestricted subjects only

Table 3.2: Session level summary statistics

ing Prediction. We included a dummy variable 'Task1' to control for the change in task from worksheet searches to educational activity searches. Model 1 is a least squares regression on time volunteered controlling for gender, religiosity, volunteer experience, and peer network.[22] Model 2 estimates a random effects model for experimental sessions to allow for the possibility of group specific norms, or other correlation in behavior within session. The estimated coefficient on 'Remove Excuses' suggests that removing excuses doubles the time volunteered when compared to Baseline. The treatment effect of Remove Monitor does not appear to be robust to controls for session level dynamics.

	Least Squares	Random Effects		
	Model 1	Model 2	Model 3	Model 4
	Minutes Worked		Entries	Entries/Min
(Intercept)	21.975***	21.907***	33.461***	1.393***
	(4.066)	(7.172)	(11.715)	(0.230)
Remove Excuses	19.959***	22.124**	36.465**	0.195
	(4.216)	(10.240)	(14.409)	(0.248)
Remove Monitor	7.724**	6.815	1.152	-0.309
	(3.218)	(10.280)	(14.656)	(0.256)
Task 1	-5.26	-4.231	18.655	0.805***
	(3.551)	(8.826)	(12.465)	(0.215)
Random Effects (by experiment)		YES	YES	YES
ρ		0.612***	0.190***	0.058
Breusch Pagan LM statistic		(185.97)	(15.46)	(0.45)
Covariates:				
Male	-2.622	-2.48	-12.232*	-0.105
	(3.065)	(2.190)	(6.652)	(0.158)
Religious	0.432	2.062	-1.479	-0.044
	(3.163)	(2.247)	(6.823)	(0.162)
Recent Volunteer	1.716	1.454	3.045	0.013
	(3.103)	(2.154)	(6.589)	(0.157)
Know other subjects	-2.783	-2.814	3.533	0.250
	(3.145)	(2.793)	(8.561)	(0.205)
N	121	121	121	121
Test Statistic	4.880	8.740	17.850	22.830
P-Value	0.000	0.272	0.013	0.002
Test	F-Test	Wald test	Wald test	Wald test

* significant at 10%; ** significant at 5%; *** significant at 1%
Robust standard errors in parenthesis
Unrestricted subjects only

Table 3.3: OLS and random effects model on minutes worked, entries completed, and productivity.

[22]We impute the demographic characteristics of one subject who failed to complete the survey.

In both models, demographic characteristics do not have predictive power in explaining time volunteered, although the signs of the coefficients follow field evidence to a certain extent.[23] Tests for the joint significance of all of the demographic controls yields an F-statistic of 0.53 for Model 1 and a χ^2-statistic of 3.67 for Model 2. While empirical studies suggest that demographic variables such as gender and religion are correlated with volunteering activity, they are not a central determinant of behavior in our experiments.

3.4.3 Quality of Contribution: Productivity

We now investigate whether our social environment manipulation affects the less visible dimensions of labor contribution. Model 3 is a random effects model with the number of database entries completed as the dependent variable. Consistent with the findings from Model 2 (minutes worked), we find that Remove Excuses doubled the number of entries completed while Remove Monitor has little effect.

Model 4 uses the number of entries per minute as a measure of productivity. The *Task1* dummy is positive and significant, suggesting that subjects searching for worksheets were working faster than subjects that were searching for educational activities.[24] The coefficient on Remove Excuses is close to zero and not significant. While not significant, the coefficient on Remove Monitor is negative, suggesting that while we see more time volunteered in the unmonitored sessions, the time volunteered may be slightly less productive. Unlike our estimation of treatment effect on contributed time, the Breusch Pagan test did not indicate statistically significant session level random effects (test statistic=0.45).

Overall, the results from Sections 3.4.1, 3.4.2 and 3.4.3 suggest that removing external obstacles that restrict a small fraction of volunteers has a powerful impact on the rest of the volunteers. The Remove Excuses treatment increases time volunteered without decreasing productivity. On the other hand, the impact of an authority figure's presence in the room is inconclusive. The coefficients for Remove Monitor weakly suggest that subjects work more productively for fewer minutes when the experimenter leaves the room. Overall, being observed seems to have little impact when there is little stigma associated with low

[23]For example, the negative coefficient of *Male* is consistent with empirical findings that women volunteer more than men. Stronger evidence for gender effects can be seen in the duration model in Table 4.

[24]Model 1 and 2 of Table 3 and the duration models in Table 4 and Appendix Table 1 suggest that subjects working on Task1 worked fewer minutes. Our conjecture is that worksheet searches may have been easier to conduct but less interesting than activity searches.

contributions.

3.4.4 Peer Effects

The literature on experimenter demand effects and leadership assume that people want to gain the esteem of an authority figure and will therefore behave more prosocially when such a person is present. However, in line with Frank's (1998) findings that the decisions of subjects in the lab are not sensitive to the payoffs of the experimenter, our results do not indicate that subjects are concerned with the experimenter. Who then, do the subjects care about?

In this section we investigate the possibility that the salient audience for subjects is their peers. Falk and Ichino (2006) find that individuals work more when working alongside others. This may be due to image signaling mechanisms (a peer group provides a larger audience), higher enjoyment or lower cost of working due to camaraderie (Rotemberg 1994), or conformism (a desire to do what everyone else is doing (Bernheim 1994)). Since the number of peers changes as people leave, volunteering creates an environment where the social factors affecting contribution change dynamically.

To address this we utilize a discrete time model, where we consider an individual's likelihood of continuing work in five minute intervals.[25] As before, we consider the subsample of 121 *unrestricted* individuals[26] and include a separate intercept to account for the change from worksheet searches to activity searches.[27] Model 1 estimates the baseline discrete time duration model without including any time varying social factors. In any time interval, subjects are 24% more likely to continue volunteering when excuses are removed and 10% more likely to continue working without the experimenter in the room. Willingness to work declines over time: with every additional five minutes, the likelihood that subjects continue to work decreases by 6%.

Model 2 of 3.4 investigates the influence of peers on an individual's decision to continue working. The variable '# subjects remaining in period' controls for the number of subjects present in the room at the beginning of each five minute interval. The presence of an

[25]The results are robust to smaller intervals of time. Since it takes less than five minutes to find an educational resource and enter the information into the database, intervals larger than five minutes are too large to capture the impact of changes in the social environment.

[26]We also include the full sample of 156 subjects in Table 3.5. All conclusions hold qualitatively.

[27]The *Task1* dummy is negative and statistically significant, consistent with our earlier conjecture that worksheet searches may have been easier to conduct but less interesting than activity searches.

	Model 1	Model 2	Model 3	Model 4
Probability of working	0.161	0.194	0.178	0.157
	dy/dx	dy/dx	dy/dx	dy/dx
Variable				
Remove Excuses	0.244***	0.178**	0.417***	0.337***
	(0.062)	(0.079)	(0.117)	(0.071)
Remove Monitor	0.108**	0.080	0.008	0.105**
	(0.045)	(0.086)	(0.063)	(0.049)
Period #	-0.057***	-0.027***	-0.058***	-0.058***
	(0.005)	(0.006)	(0.005)	(0.005)
Task1	-0.072**	-0.063	-0.065*	-0.072**
	(0.033)	(0.037)	(0.034)	(0.031)
Time varying social factors				
# subjects remaining in period		0.053***		
		(0.010)		
# subjects remaining x Remove Excuses		-0.012		
		(0.009)		
# subjects remaining x Remove Monitor		-0.005		
x Unmonitored		(0.009)		
Anyone left in prior periods			-0.121	
			(0.077)	
Anyone left in prior periods			-0.165**	
x Remove Excuses			(0.057)	
Anyone left in prior periods			0.117	
x Remove Monitor			(0.087)	
# subjects leaving in period				0.009
				(0.008)
# subjects leaving x Remove Excuses				-0.052***
				(0.014)
# subjects leaving x Remove Monitor				0.006
				(0.012)
Demographic controls				
Male	-0.051*	-0.086**	-0.069**	-0.048
	(0.030)	(0.037)	(0.032)	(0.029)
Religious	0.011	0.041	0.025	0.010
	(0.031)	(0.040)	(0.033)	(0.030)
Recent Volunteer	0.016	0.023	0.037	0.018
	(0.031)	(0.038)	(0.033)	(0.030)
Know other subjects	-0.013	-0.002	-0.023	-0.012
	(0.031)	(0.042)	(0.036)	(0.030)
AIC	0.584	0.538	0.539	0.515
N	2299	2299	2299	2299

Standard errors are clustered by individuals
Marginal effects after glm (Bernoulli distribution with complimentary log-log link function)
Periods are defined in minute intervals (0, 1-5, 6-10)
* significant at 10%; ** significant at 5%; *** significant at 1%
Unrestricted subjects only

Table 3.4: Duration model with time varying social factors – unrestricted only

	Model 1	Model 2	Model 3	Model 4
Probability of working	0.128	0.156	0.141	0.125
	dy/dx	dy/dx	dy/dx	dy/dx
Variable				
Remove Excuses	0.197***	0.139*	0.381**	0.273***
	(0.058)	(0.079)	(0.119)	(0.070)
Remove Monitor	0.079**	0.057	0.025	0.068*
	(0.022)	(0.071)	(0.045)	(0.037)
Period #	-0.036***	-0.010*	-0.036***	-0.038***
	(0.004)	(0.005)	(0.004)	(0.004)
Task1	-0.045*	-0.031	-0.039	-0.045*
	(0.025)	(0.029)	(0.027)	(0.024)
Remaining periods before	0.011***	0.015***	0.012***	0.011***
time limit	(0.002)	(0.003)	(0.003)	(0.002)
Time varying social factors				
# subjects remaining in period		0.039***		
		(0.008)		
# subjects remaining x Remove Excuses		-0.003		
		(0.007)		
# subjects remaining x Remove Monitor		-0.003		
		(0.007)		
Anyone left in prior periods			-0.027*	
			(0.042)	
Anyone left in prior periods			-0.127***	
x Remove Excuses			(0.035)	
Anyone left in prior periods			0.018	
x Remove Monitor			(0.047)	
# subjects leaving in period				-0.001
				(0.007)
# subjects leaving x Remove Excuses				-0.036**
				(0.011)
# subjects leaving x Remove Monitor				0.009
				(0.010)
Demographic controls				
Male	-0.032	-0.054*	-0.041*	-0.030
	(0.023)	(0.029)	(0.024)	(0.022)
Religious	0.009	0.032	0.018	0.009
	(0.023)	(0.030)	(0.024)	(0.022)
Recent Volunteer	0.003	0.004	0.015	0.004
	(0.023)	(0.029)	(0.024)	(0.022)
Know other subjects	-0.008	-0.031	-0.015	-0.008
	(0.024)	(0.029)	(0.028)	(0.024)
AIC	0.607	0.541	0.547	0.516
N	2964	2964	2964	2964

Standard errors are clustered by individuals
Marginal effects after glm (Bernoulli distribution with complimentary log-log link function)
Periods are defined in minute intervals (0, 1-5, 6-10)
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 3.5: Duration model with time varying social factors – all subjects

additional peer observer at the start of a five minute interval increases the probability of working through the end of the interval by 5%. The interaction with 'Remove Excuses' and 'Remove Monitor' are not significant, suggesting that peer effects remain consistent across treatment groups.

The marginal effect of 'Remove Excuses' is significant and positive. This suggests that unrestricted subjects in Remove Excuses are still 18% more likely to continue working than unrestricted subjects in the Excuses treatment even after controlling for the effect of restricted subjects' exits on group size. On the other hand, the coefficient 'Remove Monitor' is positive but not significant. This cannot be accounted for by our model, since removing the monitor should reduce visibility. However, this finding is consistent with image signaling models where image rewards depend on the identity of the audience (Ellingsen and Johanssen 2008) and related models of crowding out from monitoring (Dickinson and Villeval 2008).[28] The next section attempts to disentangle subject's image concern from other potential motivators.

3.4.5 Stigma Avoidance and Clusters

Could the higher contribution of time in Remove Excuses have been driven by reasons other than image concerns? We first consider whether the random time limit introduced framing and anchoring. Even though we attempted avoid framing effects by wording the instructions as similarly as possible across treatments,[29] we cannot entirely rule out the possibility that the treatments affected subjects' perception of the cost of time or the socially acceptable level of contribution. However we find that subjects' decisions in all treatments were highly sensitive to the immediate social environment (Table 3.4), which suggests that anchoring and framing from instructions read earlier were not major determinants of behavior.[30]

We now consider whether *unrestricted* individuals volunteered less in Excuses because they felt that it was unfair or unproductive for them to work when others were not also working. Subjects might have worked less out of retaliation if they perceived that they

[28] Monitoring causes crowding out if it communicates distrust without having a disciplining effect. Peer observation may be prefered to central monitoring in this setting since it does not communicate distrust. Another possibility is that subjects may be signaling altruism to their peers and signaling obedience to authority.

[29] See Appendix. All treatments state that subjects can stay and volunteer unpaid as long as they like up to 90 minutes. The random mechanism is explained as a method of ensuring the subject's privacy.

[30] Mann Whitney tests comparing survey responses indicated that the random time limit did not increase subjects preference for organizations that provide small compensation for volunteered time ($z = 0.11$).

were unfairly being asked to work more than others. However, higher time limits seem to induce higher volunteered time. We see a positive correlation between time limits and minutes worked when the restricted subjects are included in the duration model (Table 3.5). Conditional cooperation motives present in standard public good experiments are minimized by our task design since subjects work on independent tasks and the resulting database is unlikely to directly benefit student volunteers.

Could the increase of contributed time under Remove Excuses have been driven by conformity instead of image concerns? Subjects that are imitating each others' behavior would produce a similar pattern of departure regardless of the availability of excuses; the 'cascade' of departures would merely start earlier as the random mechanism induced restricted subjects to leave. However, image concerned subjects would be less affected by the departure of others when excuses are available. This is because image consideration has less impact on decision making when excuses reduce the stigma of low contribution, and consequently, the potential gains from signaling. [31]

Examining the raw data, we see some evidence that departure patterns depend on the availability of excuses. In Remove Excuses, subjects seem unwilling to be the first to quit volunteering, but once someone leaves, a large fraction of subjects follow suit. On the other hand, subjects leave earlier when excuses are available, but seem less affected by others' departures. For example, ten minutes after the first departure from the room, 49% of unrestricted subjects have left in Remove Excuses, compared to 16% when excuses are available.[32] This evidence suggests that stigma may not be linear in the amount of time volunteered. Below we provide a brief sketch of a possible 'bad apple' model, where an individual suffers disutility B from being the first person to stop working.

As before, let v be an agent's intrinsic motivation to volunteer. Let $\Delta_C(t) = C(t) - C(t-1)$ be the increase in cost from working an additional minute at time t. Denote the image rewards as $S(t|\delta)$, where as before δ is the probability of external obstacles. Let the

[31] The difference in sensitivity to others' departures will be largest on subjects with low altruism, who stand to gain the most in image rewards.

[32] The first unrestricted subject to leave the room volunteered an average of 27.7 minutes in Remove Excuses (se=12.19, n=4), 6 minutes in Baseline (se=3.38, n=5), and 9.75 minutes in Remove Monitor (se=3.35, n=4). Across the 9 excuses session, unrestricted subjects were the first to leave in 3 sessions. The departure times were minute 0, 1, and 15.

bad apple stigma be $B > 0$. Individual i's utility for volunteering an extra minute is:[33]

$$U(t) = v - \Delta_C(t) + S(t|\delta)$$
$$\text{where } \; S(t|\delta) = 1 + (1-\delta)B \text{ if no one has left, 1 otherwise}$$

Before anyone has left, individuals continue to volunteer either because they are altruistic ($v \geq \Delta_C(t) - 1$) or because they are avoiding the bad apple stigma ($\Delta_C(t) - (1 + (1-\delta)B) < v < \Delta_C(t) - 1$). Once someone leaves, this stigma is no longer a constraint, and those who only stayed to avoid B will depart immediately. The existence of unverifiable external circumstances (δ) lowers volunteering in two ways. First, it lessens the bad apple stigma to $(1-\delta)B$. Second, it may induce some early departures that completely eliminate the bad apple stigma. Since this means fewer people are staying due to stigma avoidance, subjects are less likely to leave in clusters when excuses are available. On the other hand, individuals who are simply following the behavior of others are equally likely to leave in clusters in both treatments.[34] The differential impact of others' departures on an individual's likelihood of continuing distinguishes stigma avoidance from conformity.

We investigate the implications of the bad apple stigma with the duration model. Model 3 estimates the probability that a person continues to volunteer given 'Anyone left,' a binary variable that is 1 if someone has left the room. By itself, 'Anyone left' is negative but not significant, however, it is negative and significant when interacted with 'Remove Excuses.' In Model 4 we estimate the probability that subjects continue working given the number of departures within that time interval. Again, the coefficient for '# of subjects leaving' is not significant by itself, but is negative and significant when interacted with Remove Excuses. We find that when excuses are not available, subjects are 16.5% more likely to leave when someone else has left and 5% more likely to leave for every subject that leaves within that time period. The marked increase in clustering behavior in the absence of excuses is not consistent with imitative behavior and is supportive of stigma avoiding behavior.

[33] Behavior in this model is not driven by expectations, so unlike the Benabou and Tirole's signaling model, no assumption about the distribution of altruism $g(v)$ or common knowledge of this distribution among the agents is necessary. An agent's strategy specifies the optimal minutes to work before and after someone else has left.

[34] Goeree and Yariv (2007), Bernheim (1994)

3.5 Conclusion

While a large body of literature addresses financial contributions, only a small literature exists on contributions of time and effort.[35] We focus on volunteering, the most common example of prosocial activity. In a typical volunteering environment, a representative from an organization orients and informally monitors a group of individuals, each of whom may be under external time restriction.

Motivated by recent theoretical and empirical studies showing that image concerns play a central role in prosocial behavior, we use an image signaling framework to investigate how each component of the social environment influences the contribution of time by volunteers. In particular, we hypothesize that the presence of a representative heightens agents' awareness of being observed, thus increasing time contributed, while the availability of excuses lowers time contributed by decreasing the stigma of low contributions.

We test these theoretical predictions with an experiment designed in partnership with School on Wheels, a nonprofit that tutors homeless children in Los Angeles. The nonprofit's own promotional material and volunteering task translate the core components of institutional volunteering into the laboratory. The laboratory setting allows time and effort to be precisely measured. Furthermore, the lab provides control over recruitment, task training, the presence of a monitor and external time restrictions.

Subjects contributed substantial time and effort in our experiment, producing several large databases of internet resources. The existence of a privately observed random time limit halved the average contribution of subjects who were unrestricted by the time limit. Subjects showed heightened sensitivity to others' departure in the absence of the random time limit: they were more likely to leave after someone else had left and were more likely to leave in clusters. These behavioral patterns are consistent with stigma avoidance and not with alternative mechanisms such as framing, anchoring, conditional cooperation or conformity.

We manipulate the presence of the experimenter to test whether being observed passively by an authority figure reduces shirking. We find no increase in volunteering when the experimenter is present. The data suggests that the salient audience for signaling in this

[35]Some examples of studies of financial contribution include Harbaugh (1998), Karlan and List (2006), Landry et al (2006) and Shang and Croson (2009). Studies in labor contribution include Gneezy and Rustichini (2000), Ariely, Bracha, and Meier (2009), Carpenter and Myers (2007).

experiment may actually be peers: subjects are 5% more likely to continue volunteering for every peer that still remains in the room.

Volunteers' productivity remains largely unaffected by our image treatments. This suggests that image treatments can influence the observable component of labor (time) without altering the unobservable dimensions (productivity).

Our results illustrate that the social environment is an important factor in determining volunteer behavior. Creating an environment where external circumstances cannot be used to justify low contributions may increase the quantity of contributions without impacting their quality. This sheds some light on the effectiveness of common nonprofit practices. Asking for contributions of time or money in public (Soetevent 2005, Martin and Randall 2008) prevents individuals from pretending that they were uninformed about the opportunity to contribute. Precommiting contributions (such as monetary pledges) makes it hard to claim prior commitments when the time to give comes. However the strategy of eliminating excuses is markedly less effective once a single bad apple openly stops contributing. While social image can be manipulated to increase prosocial behavior, the success of this approach is sensitive to the details of the social environment.

3.6 Appendix: Proofs

Lemma 3.2.1

Proof. Let $[v_L, v_H] \in \mathbb{R}_+$ indicate the interval from which v is drawn. By Proposition 6 (Benabou and Tirole, 2006), the assumption that $g(v)$ is decreasing implies that $\Delta(v^*|\delta, x)$ is increasing in v^* when $\delta = 0$. Since $\Delta(v^*|\delta, x)$ is composed of only M^+ and M^-, and M^+ is unaffected by δ, we only need to show that the slope of M^- when $\delta > 0$ lies beneath the slope of $E(v|v < v^*)$.

Let $f(v) \equiv E(v|v < v^*)$ and $f'(v)$ be its derivative. Let $f_H \equiv E(v|v \leq v_H) = E(v)$. Also define $e(v^*) \equiv \delta + (1-\delta)G(v^*)$ and $h(v^*) \equiv \frac{(1-\delta)G(v^*)}{e(v^*)}$. Rewrite $M^-(v^*|\delta) = \delta\ f_H e(v^*)^{-1} + h(v^*)f(v^*)$ and take its derivative:

$$\frac{\partial M^-(v^*|\delta)}{\partial v} = -\frac{\delta f_H}{e(v^*)^2} + h'(v^*)f(v^*) + h(v^*)f'(v^*) \tag{3.4}$$

Taking the derivate of $h(v^*)$ and substituting in $e(v^*)$ we get:

$$h'(v^*) = \frac{(1-\delta)G'(v^*)e(v^*) - (1-\delta)G(v^*)e'(v^*)}{e(v^*)^2} = \frac{(1-\delta)G'(v^*)\delta}{e(v^*)^2} \tag{3.5}$$

Substituting Eq.3.5 into Eq.3.4 and simplifying, we are left to show that:

$$\frac{\delta(1-\delta)G'(v^*)f(v^*) - \delta f_H}{e(v^*)^2} < f'(v^*)(1-h(v^*))$$

Since $0 < h(v^*) < 1$ and $f'(v^*) > 0$, $f'(v^*)(1-h(v^*)) > 0$. Since by assumption $g'(v^*) < 0$, $(1-\delta)G'(v^*)f(v^*) < f_H$, which implies that the slope of $M^-(v^*|\delta > 0)$ is smaller than $M^-(v^*|\delta = 0)$. Hence $\Delta(v^*|\delta > 0, x)$ must be increasing in v^*. □

Lemma 3.2.2

Proof. (i) Let v' the solution to $v + \Delta(v|\delta', x) - C = 0$. Honor remains unchanged by excuses while stigma is lowered, hence $\Delta(v|\delta', x) < \Delta(v|\delta, x)$. When excuses become unavailable $v' + \Delta(v'|\delta, x) - C > 0$, which implies v' will still participate. By Lemma 3.2.1 we know that $\Delta(v^*|\delta, x)$ increases in v^*, hence the new cutoff type v^* whom is now indifferent about volunteering must be a lower type. Since participation is decreasing in type, $v^* < v'$ implies higher total participation when $\delta = 0$.
(ii) Let v' the solution to $v + \Delta(v|\delta, x') - C = 0$. When visibility is decreased, $v' + \Delta(v'|\delta, x) - C < 0$ hence type v' will no longer participate. By Lemma 3.2.1 we know that $\Delta(v^*|\delta, x)$ increases in v^*, hence the new cutoff cannot be smaller than v'. Hence $v^* > v'$, and since participation is decreasing in type, this implies lower total participation. □

Lemma 3.2.3

Proof. The utility of the cutoff type at each level is zero:

$$v_t^* t - C(t) + \Delta(v_t^*|\delta, x) = v_{t-1}^*(t-1) - C(t-1) + \Delta(v_{t-1}^*|\delta, x) = 0$$

Note that $v_t^* = \frac{C(t) - \Delta(v_t^*|\delta, x)}{t}$. Subtracting the utilities we get:

$$(v_t^* - v_{t-1}^*)(t-1) + v_t^* - (C(t) - C(t-1)) + \Delta(v_t^*|\delta, x) - \Delta(v_{t-1}^*|\delta, x) = 0 \tag{3.6}$$

Substituting v_t^* into Eq 3.6 and simplifying we arrive at:

$$(v_t^* - v_{t-1}^*)(t-1) + \Delta(v_t^*|\delta, x) - \Delta(v_{t-1}^*|\delta, x) = \frac{\Delta(v_t^*|\delta, x)}{t} + C(t) - C(t-1) - \frac{C(t)}{t}$$

From the assumption that $C'(t) \geq 1$, $C(t) - C(t-1) - \frac{C(t)}{t} \geq 0$. Since $\frac{\Delta(v_t^*|\delta,x)}{t} > 0$ the entire right hand expression is positive. By Lemma 3.2.1 we know that $\Delta(v^*|\delta, x)$ increases in v^*, hence v_t^* can't be smaller than $v_{t-1}^* < 0$ since this implies $\Delta(v_t^*) - \Delta(v_{t-1}^*) < 0$ and that the left hand expression is negative. Hence $v_t^* > v_{t-1}^*$. □

Proposition 3.2.4

Proof. (i) As before let $0 = \delta < \delta'$. Let $v' = (v_1', .., v_t', .., v_T')$ denote the vector of cutoff types induced by environment (δ', x) while $v^* = (v_1^*, .., v_t^*, .., v_T^*)$ denotes the vector of cutoff types induced by environment (δ, x). Hence v_t' is the solution to $v_t t + \Delta(v_t|\delta', x) - C(t) = 0$ while v_t^* solves $v_t t + \Delta(v_t|\delta, x) - C(t) = 0$. Following the proof of the binary case Lemma 3.2.2(i) we arrive at $v_t^* < v_t'$. This implies that $\bar{a}_T(\delta, x) > \bar{a}_T(\delta', x)$. (ii) Using same steps and application of Lemma 3.2.2(ii) we show that $v_t^* > v_t'$, hence $\bar{a}_T(\delta, x) < \bar{a}_T(\delta, x')$ for $0 < x < x'$. □

3.7 Appendix: Experimental Instructions

http://corinth.hss.caltech.edu/slinardi/volInc2/databaseP.cgi

Educational Activity Resource Database

Help us build a database of targeted educational activities to help tutors engage their students. Please work carefully. If you cannot find the information from the webpage, please write "N/A". Click **Next** to proceed to the next entry. Click **Finish Volunteering** if you have completely finished working.

Your practice task today is to find *instructions for an art activity using recycled materials*.
Please open another tab (Ctrl T) to perform searches and use this screen to enter information. Do not close this screen.
Use this practice session as an opportunity to ask any questions you have.

1. Subject:

2. Grade level:

3. Description/topic area (algebra, history, painting, etc):

4. Website address:

5. Approximate duration of time needed to complete (please estimate):

6. Description of online resource or the activity itself (worksheet, field trip, experiment, etc):

7. (Optional) What is interesting about this resource? What advice do you have for the tutor who chooses to do this activity with his/her student? Does it require special preparation/skills?

1 of 2 6/1/2009 6:07 PM

Figure 3.4: Screenshot of volunteering software (CGI/Perl)

Master Subject Instructions

1. Thank you for coming. During this experiment, please do not talk, or use the web for any activities outside of the experiment. If you have any questions please raise your hand and an experimenter will come to you to answer it in private. This experiment is different from other experiments you may have participated in because we will be actually be working with a local nonprofit. Today's session will consist of a 15 minute training session, for which you will earn $10 and another $10 showup fee. After the training session, you may stay and volunteer **unpaid** as long as you like up to 90 minutes.

   After volunteering, you will complete a brief survey.

   > ***Excuses*****:** (***Unmonitored*** in parenthesis)
   > This experiment is completely anonymous, not only to other subjects but also to the experimenter (who will not be present during the experiment). Your decisions and answers to the survey will be tagged by only an ID number, allowing us to analyze the data without using any identifying personal information.
   >
   > ***Unmonitored***: Again, it is important that you do not communicate with each other. After the training session, the experimenter will have no further involvement with anyone in this experiment. However, you may ask questions to the experimenter throughout today's session using the AIMExpress. You have received a piece of paper with a username and password for the chat software. The experimenter will be on your buddy list when you sign in. If you have any problems signing in, raise your hand and a lab assistant will help you.

2. On your keyboard there is information about School On Wheels, the nonprofit that we will be working with today. Please read the article about the organization. Our job today is to compile a database of educational activities for School on Wheels tutors. These tutors often do not have a teaching background and may find it difficult to come up with age appropriate activity for kids that can be done with their limited resources. The list of activities you suggest today will help the tutors connect with homeless kids more effectively.

3. I will now pass along a sheet of paper on the type of activity that you are in charge of finding. We have staggered your task for minimum overlap with other students so that we get to cover as many areas as possible. Please take a look at your task and ask me any questions you have.

4. Before we start the actual work, we will do a five minute practice task. Click *Start Practice Task*. You are now in the database window. Please take extra care to not close this screen during the ENTIRE session.

5. ***Excuses*****:**

   > Notice that in the bottom of the screen there's a button that says *Roll Dice*.
   > You will click this LATER when you have completed your practice task. This mechanism protect the privacy of your choice of how long to volunteer. When you click *Roll Dice*, the computer will roll a dice and randomly pick the maximum number of minutes you will volunteer today. This number will be between 0-90 minutes and the computer will automatically stop you from starting a new entry once time is up. It will not interrupt you so do not worry about losing any work. Remember: you do not have to do the number of minutes the computer picked: how long you want to work is completely up to you. Again, your privacy is guaranteed: when you leave the room, nobody will know whether you chose to finish or were forced to by the time limit. After rolling the dice, if your time limit is zero,

click *Start Survey*. Any questions?

6. Please press Ctrl T to open a new tab, look online for an art project using recycled materials, and input the information you found into the database window. When you are finished with your practice task please wait for further instructions before you click on anything.

7. Thank you for completing the training. You are now free to go. Please keep the database window open and fill the survey before you go. The lab manager outside will process your show up fee. If you want to stay to volunteer, you will now look for the activity listed in your sheet. Remember that you choose how much you want to work – there will be a button that says *Finish Volunteering*. It is very important that you work carefully, since the information you produce today will be given out to tutors as a searchable database of educational activities. The lab is available for us for the next 90 minutes.

No Excuses: You can click *Start Volunteering* now.
Excuses: You can click *Roll Dice* now.

Unmonitored: I will now leave the room. To reach me at any time you can contact me through AimExpress. Please open the AimExpress, and send thx.experimenter a test message. If you have any software problems, raise your hand and a lab assistant will help you.

Survey

1. We will compute the average minutes volunteered by UCLA students today. Please write down your guess for this average. If your guess comes closest to the actual average, you will receive a $20 gift certificate to Amazon sent to the email address you indicate at the end of the survey.

   I think the average number of minutes volunteered by other students is _______ minutes.

2. If your guess (see no.1) comes closest to the average, which email address should we send the $20 gift certificate to?
3. Gender: Male __ Female __
4. Do you identify with any particular religious tradition, denomination, or church?
   a. NO
   b. YES
5. How many people do you know in this room by first name (including the experimenter)?

   I know __________ people in this room.

6. When was the last time you volunteered?
   a. I have never volunteered before
   b. Within the last week
   c. Within the last month
   d. Within the last year
7. If you have volunteered before, what organization did you volunteer for?

   ____

8. Refer to question 6. On a scale from 0-10, how valuable was the volunteering work you *did for this organization*?

   0 1 2 3 4 5 6 7 8 9 10

9. Refer to question 6. How did you hear about this organization? (check all that applies)
   a. I was referred to it by a friend or relative
   b. I saw it in the media (print, TV or radio)
   c. I myself or someone I know are personally affected by the cause this organization works on
   d. I found it when looking for volunteer opportunities
10. On a scale from 0-10, how valuable was the volunteering work you did *today*?)

0 1 2 3 4 5 6 7 8 9 10

11. Suppose there is an organization with two different regional offices that needs volunteers for doing data entry from their home. Which organization would you recommend that your friend work for?
    a. A regional office that pays volunteers 5 cents per minute worked.
    b. A regional office that doesn't pay volunteers.
12. Refer to your answer to question 10. Why did you recommend that organization?

    ____

13. How can we improve today's volunteering experience?

    ____

Chapter 4

Can Relational Contracts Survive Stochastic Interruptions?

with Colin Camerer

4.1 Introduction

It is well-known that typical employment contracts very loosely specify the duration and terms of employment.[1] Implicit contracts which reflect historical norms are often called "relational contracts". Because these contracts are implicit, it is often difficult to measure their terms directly, but these limited duration contracts often lead to repeated future contracts. However, circumstances outside of the workers' performance such as stochastic drops in demands for a firm's products or changes in personal circumstances of an employer often affect the probability of future contracts. This paper is an attempt explore how much these stochastic interruptions affect the labor market.

Because individual level data linking a worker's pay with productivity are scarce, experiments in which contract terms can be exogenously manipulated play an important role in understanding relational contracts. Our paper replicates and extends an experimental "gift exchange" paradigm used by Brown, Falk and Fehr (2004, 2008, henceforth BFF1 and 2) to study relational contracts. In this paradigm firms offer prepaid wages and workers who accept wages then exert costly effort. Since effort is not contractible, there is a classic moral hazard problem; one possible outcome is that workers exert little effort and firms, anticipating this, pay minimum wages. BFF1 extends the paradigm by attaching ID num-

[1] Most US firms uses "employment at will" clauses, which states that an employment contract of indefinite duration can be terminated by either the employer or the employee at any time for any reason.

bers to all firms and workers, hence allowing "private" offers to be extended to a particular worker. A two tier labor market emerges rapidly in which private contracts sustain high wages and effort while public contracts sustain low wages and effort.

Our paper makes two contributions. First, we directly replicate the BFF1 experiment and check its robustness to subject pool and time horizon. The central stylized facts replicate rather well, though there are some small differences in behavior. Second, and more importantly, we investigate the robustness of relational contracting and market efficiency to stochastic downturns. In every period of trading, each firm faces a probability p of being hit by a downturn, where no hiring is possible for three periods.[2]

Temporarily removing firms from the labor market creates uncertainty about future relationships that may neccesitate the emergence of *loyalty norms*. These norms determine the type of firm behavior that signals commitment to rehire a worker after the downturn and worker behavior that signals commitment to maintaining or increasing effort if rehired. These norms determine the effect of stochastic interruptions on relationships, and consequently a worker's future utility of not shirking (*job rents*). If loyalty norms are weak, workers will start relationships with new firms as soon as their firm is hit by a downturn and returning firms will dip into the labor market for a new worker. Shirking will be less costly for a worker since the probability of being rehired is high, and the market may unravel. However, we theoretically show that with strong loyalty norms, high quality contracts can be maintained despite stochastic interruptions .

Our experiment shows that while stochastic interruptions have some disruptive effects, the net effect is actually beneficial in an interesting way. Consistent with our theoretical predictions, we did find that job rents in downturns are lower than that in baseline. However, wages are actually higher in downturn, especially in the second half. Effort is *not* higher, and therefore average worker profit per trade is significantly higher. The reason for this is complex: firms, perhaps out of fear of loss earnings from nontrading periods, demand more out of workers in downturns before renewing relationships. This delays the separation of the two tiers of the market, improves the short term market to such an extent that it is difficult to get workers to resist public offers and stick with private contracting. In the end firms

[2]Note that our specification was not designed to resemble any particular macroeconomic type of downturn. (In the economy downturns are often correlated across firms, due to business cycles, and their length is stochastic rather than fixed.) When downturns are uncorrelated, strategic behavior from anticipating future downturns is minimized, thus allowing us to better isolate the resilience of relational contracts from stochastic interruptions.

are ironically forced to spend more for the same level of effort. However, the best pairings are largely unaffected by the interruptions: firms and workers who are sharing their surplus 50-50 before the downturn have strong loyalty norms and reconnect after the downturns. Like many previous experiments, this positive effect of downturns illustrates the fact that an expanding labor economics analysis including considerations of reciprocity and fairness can overturn conventional results.

4.2 Theoretical Framework

The celebrated folk theorem in game theory established that payoffs which are Pareto-superior to those achievable in the one-shot Nash equilibrium can arise given sufficient probable repetition of a one-shot game. However, these results are sensitive to disruptions and noise. For example, tacit collusion among firms to produce their share of monopoly output dissolves into Cournot behavior when demand fluctuations not directly observable by firms lower prices sufficiently (Green and Porter, 1982). Furthermore, in theory, if a worker-firm game is only repeated finitely, there should be shirking in the last period which unravels to shirking immediately. However, in experiments and naturally-occurring settings with known finite horizon, cooperation typically occurs until close to the end (e.g., Camerer, 2003). One explanation is gift exchange due to norms of reciprocity or moral obligation (e.g., Akerlof, 1982, 1984; Rabin, 1993; Fehr, Kirschsteiger and Reidl, 1993): firms pay above-market wages and workers reciprocate the "gift" with costly effort.

BFF1 shows that introducing fair types into a labor market with excess supply makes it possible to attain a high-efficiency equilibrium. In this equilibrium all firms offer fair wages and selfish firms offer rent at the last period, inducing both fair and selfish types to put in the requested effort until the end. A worker who shirks is immediately punished by nonrenewal of contract; shirking hence has a very high cost since reemployment is unlikely given the excess supply in this labor market. Selfish workers and fair workers act undistinguishably until the final period, when the selfish worker will finally shirk. Their followup paper, BFF2, shows that the existence of a high-efficiency equilibrium does not depend on the threat of unemployment. In a labor market with excess demand with inequity-averse types,[3] the high equilibrium can be attained when public offers feature lower wages than private repeated

[3] Note that BFF2 employ a different behavioral assumption of fair workers than BFF1

offers.

How would stochastic interruption affect behavior? Specifically, what is the impact of temporarily removing firms from the labor market on relational contracts? First consider the effect of the increased supply of labor. The increased uncertainty of being repaid for high efforts with a renewed contract may drive workers towards more shirking. On the other hand, the higher threat of unemployment may drive employed workers to try to protect their jobs by consistently delivering high effort.

A second possible effect is that the job rents to workers in the high wage-high effort equilibrium might be lower in the downturn market than in the baseline market because of the probability of future disruptions. This will generally lead to lower sustainable effort level in the downturn market compared to the baseline market.

A third possible effect comes from the prospect of interruption serving as a selection criterion. In repeated games where the folk theorem applies, there are always inefficient equilibria as well as efficient ones. The fact that there is a possibility of interruption, per se, could serve as some sort of public correlating device that implements the inefficient equilibrium.

In this paper we focus on the ability of loyalty norms to influence the net effect of stochastic interruption. Suppose that all firms form relational contracts, and wages and efforts are high. Then a firm is hit with a shock and exits the market for a known number of periods. When loyalty norms are strong, the firm's previous worker simply waits out the interruption period, and when the firm returns to the market, the firm makes a private offer to 3, thus continuing the relationship. This strong loyalty norm would produce aggregate behavior very similar to the perfectly efficient relational contracting with "missing data" (from the downturn periods). On the other extreme, suppose loyalty norms are non-existent. Then workers laid off in downturns will immediately accept any public or private offer, and firms returning from downturns will make new public offers and find new partners. Here shirking is less costly for the workers and relationship lengths will be shorter.

The theoretical model here illustrates the intuition behind our experiment. For simplicity, we model the downturn as a single period interruption instead of the three period interruption in our experiment. Like the framework from BFF1, there is an excess supply of labor and selfish workers imitate fair workers at period $t < T$ for fear of unemployment. However, we use BFF2 behavioral assumptions: all firms are utility maximizers, p workers

$\tilde{e}_T$	1	2	3	4	5	6	7	8	9	10
1. $c(\tilde{e}_T)$	0	1	2	4	6	8	10	12	15	18
2. $\hat{w}(\tilde{e}_T)$	5	12	16	22	28	34	40	46	53	59
3. $\hat{w}(\tilde{e}_T) - c(\tilde{e}_T)$	5	11	14	18	22	26	30	34	38	41
4. $\hat{w}(\tilde{e}_T) - 5$	0	7	11	17	23	29	35	41	48	54
5. $\tilde{e}_{T-1}$	1	5	7	9	10	10	10	10	10	10

Table 4.1: Theoretical fair contracts at T and corresponding effort level at $T-1$

are *fair* and the rest are utility maximizers. In order to model employment choices while a firm goes into downturn, we assume two firms and $n > 2$ workers, which depart from BFF1's model of a single firm and many workers or BFF2's model of a single worker and many firms. We first discuss our prediction for the baseline labor market before analyzing the downturn market.

Without stochastic interruptions, a single period of the game proceeds as follows:

1. Firm simultaneously offers a contract that stipulates wages and desired effort $[w, \tilde{e}]$ to all workers through public offers or to a single worker through private offers.

2. Workers receive all contracts at once and decide which one (if more than one) to accept.

3. If the worker accepts the contract, he chooses an actual effort level $e \in [1, 10]$ to deliver. If a worker rejects a contract, the firm can offer another contract.

4. After e is chosen, firm earns $10e - w$ and worker $w - c(e)$ (see Table 4.2 Row 1 for cost of effort). Unmatched firms earn 0; unmatched workers earn 5.

Stochastic interruptions happens at the beginning of the period before the firm makes offers. With probability δ, a firm experience a temporary (observable) shock in demand and cannot hire for one period.

The utility of a fair worker follows BFF (2009) model: workers have a bad conscience if they do not fulfill a contract that offers an equal (or better) split of surplus. Letting $\hat{w}(e)$ denote fair wages, Table 4.2 Row 2 provide $\hat{w}(\tilde{e}) = 5e - c(e)/2$ for each effort level. The marginal disutility b of not fulfilling a fair contract is assumed to be high enough such that fair workers will always provide the requested effort if they accept a fair contract.

Definition 8. *A fair worker is an agent whose utility for an accepted a contract of* $[w, \tilde{e}]$ *is* $w - c(e)$ *if* $w < \hat{w}(\tilde{e})$ *and* $w - c(e) - bmax(\tilde{e}, e)$ *otherwise.*

Proposition 4.2.1. *Consider a game of* $T = 1$ *period with two firms and* $n > 2$ *workers where a proportion* $p \in (0, 1)$ *is fair as defined above. If* $p < .55$*, there exist no PBE where the firm offers more than* $(5, 1)$ *and worker perform* $e > 1$*. If* $.55 \leq p < .6$*, there exist a PBE where fair workers perform* $e = 2$ *and selfish workers perform* $e = 1$*. If* $.6 \leq p < .65$*, there exist a PBE where fair workers perform* $3 \leq e \leq 8$ *and selfish workers perform* $e = 1$*. If* $p \geq .65$*, there exist a PBE where fair workers perform* $e > 8$ *and selfish workers perform* $e = 1$*.*

We now consider a multi period model; we will index all the variable with t to denote the period. Since all offers will be fair in equilibrium and fair workers will always deliver the requested effort of a fair offer, we only need to consider the shirking behavior of selfish workers. Let V_{t+1} be the expected future utility of not shirking at period t and U_{t+1} be the expected future utility of shirking at t. The selfish type will shirk if future job rents, defined as $V_{t+1} - U_{t+1}$, is not larger than current cost of delivering the requested effort.

$$V_{t+1} - U_{t+1} > c(\tilde{e}_t) \tag{4.1}$$

Proposition 4.2.2. *Consider a game of* $T > 1$ *periods with two firms and* $n > 2$ *workers where* $p = .6$ *are fair as in Definition 8. The following strategies and beliefs constitute a perfect Bayesian equilibrium in which both worker types perform maximum effort in all non final period* $t < T$*:*

- *At* $t = 1$ *all firms make public offers with the identical payoff splitting contract* $[59, 10]$*.*
- *A firm offers to his previous worker a contract of* $[59, 10]$ *in all periods* $1 < t < T$ *and* $[46, 8]$ *in period* T *if the worker performed the demanded effort in all previous periods. If the worker ever shirked, the firm privately offers the contract to a never employed worker. If all workers have shirked, firm makes public offers of* $[5, 1]$ *in all future periods.*
- *In period 1, two workers accept the two public offers and perform the desired effort* $e_t = 10$ *if he is selfish or fair. At period* T *a selfish worker performs* $e_T = 1$ *while a fair*

worker performs $e_T = 8$. Pairing happens at the first period and remain throughout; $n-2$ workers remain unemployed for all t.

- *Out of equilibrium belief: a firm believes that if a worker ever shirks he is selfish.*

We now consider downturns. Let μ be the probability of receiving employment offers from someone other than the current firm. μ depend on contracting norms: for example, if all firms that went into downturn make new offers upon returning, then $\mu = \delta$. If firms only make private offers to their previous trading partner, then $\mu = 0$. It is also possible that some firms are unmatched in period $t > 1$ either because they are hit by a downturn before forming relationships, or that their offers are rejected, thus necessitating public offers. Here $0 < \mu < \delta$.[4]

Let V^{δ}_{t+1} be the expected future utility of not shirking at period t and U^{δ}_{t+1} be the expected future utility of shirking at t when firms go into downturn with probabilty δ. Suppose firms attempt to reconnect with workers that did not shirk after a downturn but some firms remain unmatched at period $t > 1$ ($\mu > 0$). Let j_t be the shorthand for $\hat{w}(\tilde{e}_t) - c(\tilde{e}_t)$.

$$V^{\delta}_t = (1-\delta)(j_t + V^{\delta}_{t+1}) + \delta[\mu(j_t + V^{\delta}_{t+1}) + (1-\mu)(5 + j_{t+1} + V^{\delta}_{t+2})] \tag{4.2}$$

With probability $1-\delta$ the firm does not go into downturn and the pair continues their relationship. With probability δ firm goes into a downturn. When this happens, with probability μ agent accepts a public offer from another firm, and with probability $1-\mu$ agent is unemployed at period t and returns to the incumbent firm when he returns from the downturn.

Let U^{δ}_{t+1} be the expected future utility of shirking at t when firms go into downturn with probabilty δ.

$$U^{\delta}_t = \mu(j_t + V^{\delta}_{t+1}) + (1-\mu)(5 + U^{\delta}_{t+1}) \tag{4.3}$$

With probability μ agent accepts a public offer from another firm, and with probability $1-\mu$ agent is unemployed at period t and faces period $t+1$ with the same uncertainty.

[4] μ may also depend on subjective beliefs (optimism).

Proposition 4.2.3. *For $\delta > 0$ and $\mu \geq 0$*

(i) job rents are positive $V_t^\delta - U_t^\delta > 0 \; \forall t \leq T$ for all $e_T > 1$

(ii) stochastic interruptions lower job rents

$$V_t - U_t > V_t^\delta - U_t^\delta \quad \forall t \leq T$$

Since job rents are lower in downturn, the incentive constraint in Eq.4.1 binds more frequently. For example, suppose $(1-\mu)(1-\delta) = 0.5$. The second to last period effort level that can be sustained given $\tilde{e}_T = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ has now decreased to $\tilde{e}_{T-1} = \{1, 3, 4, 6, 7, 8, 9, 10, 10, 10\}$ from $\{1, 5, 7, 9, 10, 10, 10, 10, 10, 10\}$ when there is no downturn. Notice that given the same marginal cost of effort (e.g $c'(e) = 2$), higher effort levels are affected less than the downturn (for $e_T = \{3, 4, 5, 6, 7, 8\}$, the reduction in e_{T-1} due to downturn is $\{3, 3, 3, 2, 1, 0\}$).

The next proposition shows that an equilibrium exists where workers and firms have strong loyalty norms.

Proposition 4.2.4. *Consider a game of $T > 1$ periods where a firm can be hit by a downturn at any period with probability $\delta < 0.5$.*[5] *There are two firms and $n > 2$ workers where $p = .6$ are fair as defined above. The following strategies and beliefs constitutes a perfect Bayesian equilibrium in which both worker types perform maximum effort in all non final period $t < T$:*

- *At $t = 1$ all firms that are not in downturn make public offer of* $[59, 10]$.
- *If a firm is in a downturn at $t = 1$, when returning it makes a private offer to never employed worker of* $[59, 10]$. *If there are no workers that are never employed, firm makes a public offer of* $[5, 1]$.
- *A firm offers to his previous worker a contract of* $[59, 10]$ *in all periods $1 < t < T$ and* $[46, 8]$ *in period T if the worker performed the demanded effort in all previous periods. If the worker ever shirked, the firm privately offers the contract to a previously unmatched worker. If all workers have shirked, firm makes public offers of* $[5, 1]$ *in all future periods.*

[5] Except when a firm has just returned from the downturn. A firm cannot be hit by two consecutive downturns.

- *When a firm goes into a downturn, it makes a private offer to the previous worker if he has not shirked.*

- *Workers accept all offers and performed the desired amount* $e_t = 10$ *at* $t < T$. *Worker chooses the incumbent firm when he receives identical competing offers.*

- *Out of equilibrium belief: a firm believes that if a worker ever shirks he is selfish.*

- *In period 1, two workers accept the two public offers and perform the desired effort* $e_t = 10$ *if he is selfish or fair. At period* T *a selfish worker performs* $e_T = 1$ *while a fair worker performs* $e_T = 8$. *Pairing happens at the first period and remain throughout;* $n - 2$ *workers remain unemployed for all* t.

Stochastic interruptions reduce job rents, however, if the equilibrium wage /effort provides high enough profit (j_T) and reconnection norms are adequately strong, relationships and quality of contracts will survive downturns. Lower equilibrium wage/effort will be more affected by the discounting caused by the stochastic interruptions, and lower reconnection norms (large μ) will further reduce the job rents ($V_t - U_t$), thus making the incentive constraint binds more frequently and reducing the quality of contracts that can be sustained in equilibrium.

Our experimental investigation will focus on the following questions: first, are job rents are lower in downturns? If so, will it lower wages and effort? What sort of loyalty norms exist between firms and workers and how do these norms affect contracts?

4.3 Experiment

Many experiments have been done in the basic "gift exchange" paradigm (see Camerer and Weber (2008), Fehr and Gchter (2000)). The general finding is that wages and efforts are persistently above the minimum level.[6] The general interpretation of these effects is that many workers feel a sense of positive reciprocity toward firms that pay generous wages. However, there is also typically a noticeable drop in effort in the last period, which is

[6]When the total surplus from higher effort is smaller, efforts are lower too. Some experiments have shown boundary conditions under which wages and efforts are not very much higher than the minimum (e.g., Rigdon (2002)). Note, however, that in all cases (including the latter) when efforts are regressed against prepaid wages there is a strong, significant correlation.

a reminder that some workers are purely self-interested and will shirk when reputational concerns are absent.

BFF1 ran their experiments with 7 firms and 10 workers. They found that when effort is non-contractible and firms can make private offers (using worker IDs and privately-known effort history), about half of the total trades take place in relationships that last at least half of the total number of trading periods. Effort becomes a function of wage and private offer and the length of the relationship.

In their experiments "two-tiered" labor markets emerge reliably. Long term relationships conform to the theoretical prediction above and are characterized by repeated private offers, high wages, and high (costly) effort. The longer the previous length of relationship and the more positive the effort surprise is,[7] the more likely it is that a private contract will be repeated. On the other hand short term relations also exist in which contracts are mostly initiated through public offers with low offered wages and delivered efforts are low. Some of these results are detailed below in Section 4.4.1 and compared to our baseline replications of their design.

We ran two types of experimental sessions: baseline sessions replicating the structure of BFF1 with 9 firms and downturn sessions in which there are commonly-known probability (.05 or .10) of a firm-specific downturn which lasts 3 periods. The experiments were run in Spring 2008. Table 4.2 shows the treatments and subject pools in all sessions.[8]

Our session features were modified step-by-step to pursue what we thought were the most interesting variations. We first explore a labor market where in every period firms face a downturn probability of .05. This did not produce many downturns so we increased the probability to .10. We incorporated the effect of subject pool heterogeneity by running sessions at both California Institute of Technology (Caltech) and University of California in Los Angeles (UCLA). The subject pool at Caltech is small, often unusually cooperative, and have substantial experience with laboratory experiments, while UCLA subjects resemble the typical large-university group more closely.

[7]Surprise is defined in BFF1 as the difference between delivered effort and expected effort, which is elicited from firms after the workers have accepted the contract but before effort is revealed. We will use the same definition here.

[8]Two pilot baseline session with 15 periods were ran in January 2008. These pilot sessions replicated the BFF1 results.

Location	Treatment	Firms	Workers	Periods	Number of sessions
Caltech	Baseline	9	10	30	2
	Downturn $p = 0.05$	8	9	30	1
	Downturn $p = 0.05$	9	10	30	1
	Downturn $p = 0.1$	9	10	30	1
UCLA	Baseline	9	10	30	2
	Downturn $p = 0.1$	9	10	30	2

Table 4.2: Experimental treatments (Caltech and UCLA)

4.4 Results

Our analysis is summarized in four results. Section 4.4.1 shows the replication of baseline markets in UCLA and Caltech. In Section 4.4.2 we present the impact of downturns on job rents and contracting: even though job rents are lower, market efficiency is unharmed. We investigate the reason for this robustness in the last two sections. Section 4.4.3 shows that temporary market (short term relationships) actually improves with downturn threats. Ironically, this is because firms demand more out of public workers before renewing relationships, thus delaying the separation of the two tiers of the market and shortening relationships. The improvement in wages and effort in the temporary market increases the reservation value of workers, which induces firms desiring relationships to raise wages in order to achieve adequate separation. However, Section 4.4.4 shows that firms and workers who are sharing their surplus 50-50 remain unaffected. These pairs have strong loyalty norms that allow them to reconnect after downturns. Reconnection is fairly sharply predicted by whether firms offer surplus shares around 50% .

4.4.1 Replication of Baseline Result in Two Subject Pools

Our experiments replicate the basic patterns of contracting, wages and effort in BFF1 quite closely using two different subject pools and extending number of contracting periods from 15 to 30 (see Figure 4.1). The efficiency of the UCLA market is lower than that of the Caltech sessions which is slightly higher than the BFF averages. Wage, effort, private offers, and relationship length are remarkably consistent.

Figure 4.1.a (Top Left) shows the pattern of wages for accepted offers. Average contracted wages are increasing as a function of time. Qualitative and quantitatively, wages are very similar in BFF1's sessions and our baseline replications. The average plots mask

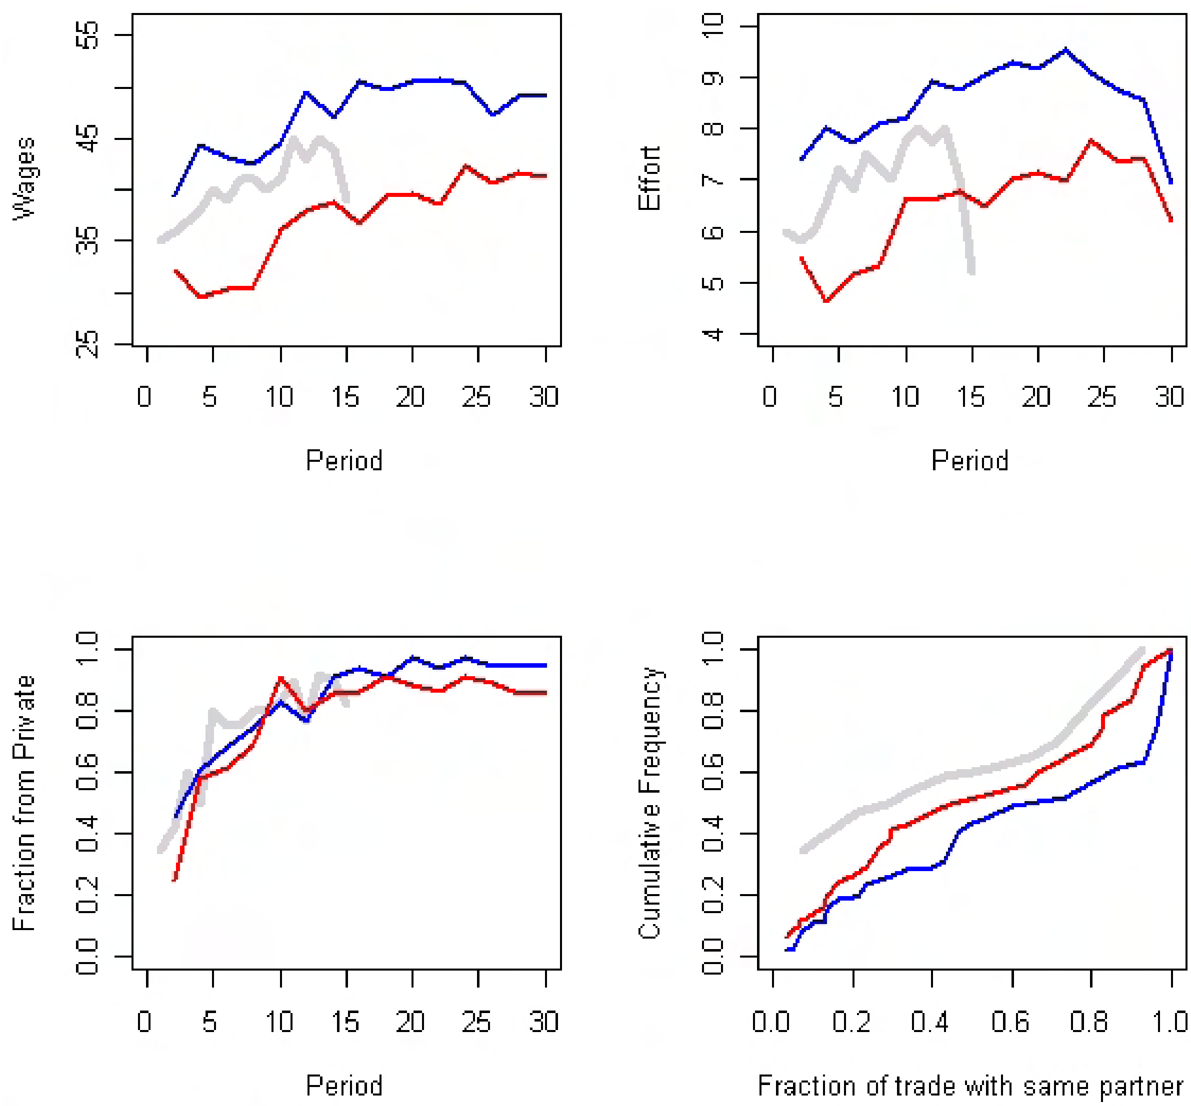

Figure 4.1: Patterns of relational contracting in BFF (gray) replicated at Caltech (dark blue) and UCLA (light red).

a lot of cross-firm variability however. Many private offers are made persistently around a wage of 60, while public offers are lower, around 20.

Figure 4.1.b (Top Right) shows the pattern of worker efforts. Average delivered effort increases in time with a last period effect. There is upward movement across the time with a last period effect; the effort levels are similar across BFF1's sessions and both of ours. Figure 4.1.c (Bottom Left) shows the percentage of offers that are private (i.e., which earmark an offer for a particular ID-numbered worker). That percentage starts off at about 1/3 and moves upward to around 80% before levelling off. Private relational contracting clearly comes to dominate in all subject pools.

One useful statistic is the average job tenure, summarized by the distribution of the fraction of total periods in which a firm hires the same worker. (Note that these do not have to be continuous runs of identical hiring). Figure 4.1.d shows the cumulative distribution of this statistic. There are more short term relationships in BFF1; about 1/3 of trades in BFF1 take place between partners that interacted 1/10th of the time (1 to 2 periods) while in LC (Linardi-Camerer) sessions only 1/6 of total trades fall within this category. In contrast, about 1/3 of the matches in our replications last through more than 9/10th of the session (27-30 periods); this number is much smaller in BFF1. Except for these modest differences in extreme short and long matching, however, the typical job tenures are not far apart. The differences are likely due at least in part to the fact that longer relationships develop over 30 periods (LC) than over 15 periods (BFF).

Surplus sharing and the relation between profits and relationship length are similar in BFF1's sessions and our sessions (see Figure 4.4 in the Appendix). Surplus sharing in very short relationship (temp market) is filled with shirking, which benefits workers. Then in the entry level of the high tier relationship, the workers pay their dues in higer effort to gain trust of the firm, sacrificing their surplus. As the relationship grows, the surplus sharing becomes more equal eventually reaching approximately 50-50 for the longest relationships. Earnings per trade for both firms and workers rise with relationship length in both sessions.

4.4.2 Impact of Downturn on Job Rents and Contracting

The key mechanism in the argument for sustaining a high equilibrium in finitely repeated contracting is the existence of positive rent for workers in the final period. This rent exists if there is difference in earnings between workers with a job and those without one. This

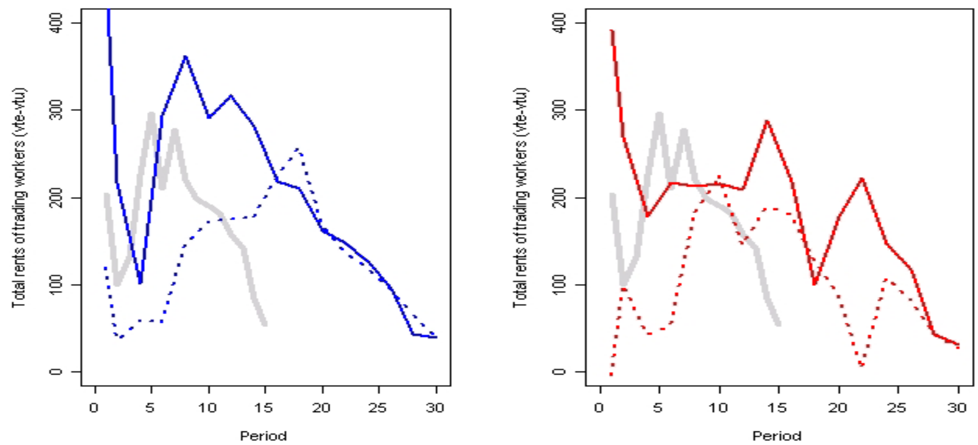

Figure 4.2: Job rents in Baseline (solid) is higher than Downturn (dotted) at both Caltech (L) and UCLA (R).

rent makes it optimal for workers to not shirk in earlier periods, which makes it profitable for firms to offer high wages (as in efficiency wage theory). We predict that when stochastic interruptions can happen, job rents are lower because workers that are employed may be unemployed in the future. We computed the average of all present and future incomes of all workers trading in period t. The average value of a job in period t (denoted Vte) is calculated from total current and future income of workers employed in period t. The average value of being without a job in period t (Vtu) is the average future income of a worker who is unemployed in period t.[9] Generally, Vte and Vtu are declining in t since there are fewer periods remaining. In a spot market where there are no relational contracts and workers freely move from firm to firm, Vte and Vtu will be similar. In a highly relational market with excess labor supply, in which all workers form long-term attachments, there will be a large separation between Vte and Vtu which represents the cumulative job rent from having a relational contract. BFF finds large and positive job rents in all 15 periods.

Figure 4.2 compares job rents in Baseline (Solid black)to Downturn (dotted black) and BFF (gray, 15 periods). As predicted, job rents (Vte-Vtu) are lower in downturn almost throughout the entire period. Rents are always large and positive in all markets (as comparison, we rescale BFF's job rents by two to take into account the difference in the number

[9]This includes both matched workers who are left unemployed when their firms go into downturn and those who have been unmatched all along. We will look at the matched workers more carefully in Section 4.4.4.

of periods in our experiment). The evolution of job rents across the three markets show remarkable similarites: the downturn job rents seem to be a lower, lagged version of the baseline job rents.

Consistent with the lower job rents, relationship lengths are also slightly shorter - trading parties in downturns have on average interacted two periods less than in the baseline (Table 4.3 Row 9). There is a 13% increase in trade taking place in relationship of length 4 or shorter (there's a 6% increase in 1 shot contracts alone). Row 12 shows that in the downturn 5% fewer contracts led to renewal compared to baseline. To take account of the differences in total trading periods introduced by the downturns, we define (firm) a.r.l (adjusted relationship length) as the number of trades between two partners divided by the total number of period the firm has been active. An a.r.l of 1 means that the firm has traded only with this partner. A.r.l in downturns are 6% shorter than in baseline (Row 10). These shorter relationships meant more pairings: 52% of all possible matches in the downturn markets actually occur, compared to 47% in baseline markets.[10]

Does this imply that the wages and effort levels that can be sustained in the downturn market are lower? Suprisingly, the answer is no. Table 4.3 shows that offered wages are actually significantly higher (Row 1) and so are contracted wages (Row 5). Effort level remains the same (Row 11) at around 7.5, resulting in a small shift of surplus sharing towards workers in the downturn treatment (Row 3). Figure 4.5 in the Appendix plots the patterns of contracting (wages, effort, private offers, and relationship lengths) against time in the baseline and downturn markets.

4.4.3 Slow Separation Leads to Higher Prices

The instability from the downturn market creates smaller job rents and shorter relationships. Since the size of job rents determines the severity of the threat of unemployment and relationship lengths are positively correlated with high wages, this makes the higher wage even more of a puzzle.

Plotting the average public and private wages offers separately against periods (Figure 4.3), we see an interesting pattern.The first thing to notice is that in the downturn market (dotted line), there is little separation between the two tiers of the market (public and

[10]Note that when there are 9 firms and 10 workers, 90 distinct firm-worker matches are possible in each experimental session.

	Caltech		UCLA	
	diffe	diffSE	diff	diffSE
Offers:				
1.Wage	3.70	0.80	2.04	0.98
2.Desired Effort	0.43	0.09	0.39	0.14
3.Worker/Total Surplus	0.01	0.01	0.04	0.01
4.Fraction Private	0.03	0.02	-0.04	0.02
Accepted Offers				
5.Wage	2.21	0.81	-1.70	1.28
6.Desired Effort	0.22	0.09	-0.52	0.16
7.Worker/Total Surplus	-0.01	0.02	0.05	0.01
8.Fraction Private	-0.02	0.02	-0.06	0.03
9.Relationship Length	-2.25	0.43	-2.54	0.43
10.A.r.l	-0.07	0.02	-0.08	0.02
Actual				
11.Delivered Effort	0.04	0.14	-0.30	0.20
12.Effort Surprise	0.13	0.09	0.00	0.12
13.Buyer Profit	-1.85	0.91	-1.35	1.26
14.Seller Profit	2.28	0.58	-0.77	0.96
15.Worker/Total Surplus	0.04	0.04	0.04	0.03
16.Renewed Immediately	-0.01	0.03	-0.14	0.03

Note:
Surprise = Expected Effort(not reported) – Delivered effort
A(djusted).r(elationship).l(ength) = Previous trades/Number of active trading period

Table 4.3: Summary statistics of the effect of downturn.

private) until the second half. In the baseline market (solid lines), the two tiers separated early. Private offers in the first half of the downturn are only 3.35 (1.26) higher than public offers. But private offers in the first half of baseline are 12.16 (1.29) higher than over public contracts.

There is a striking increase in the quality of the short term market in the downturn treatment. The average downturn wage offer is 29.75 (dotted line with white marker), compared to 18.59 (solid line with black marker) in baseline. The difference is significant at $p < 0.01$. The improvement in offers resulted in a much higher contracted wage/effort level for the second half as well. For the downturn market the average second half wage is 27.8, sustaining average effort level of 4.3. The baseline wage in the 2nd half of the 30 period contracting is a much lower 9.3, sustaining low average effort of 2. ($p < 0.01$)

We now focus on the determinants of wage, acceptance, effort, and renewal. In Table 4.4 first look only at the first period because each period affects the following period. Similar to that of BFF, the strongest drivers of wages are requested effort. This wage effect has been found, strongly, in every lab experimental study of this type (see Camerer and Weber,

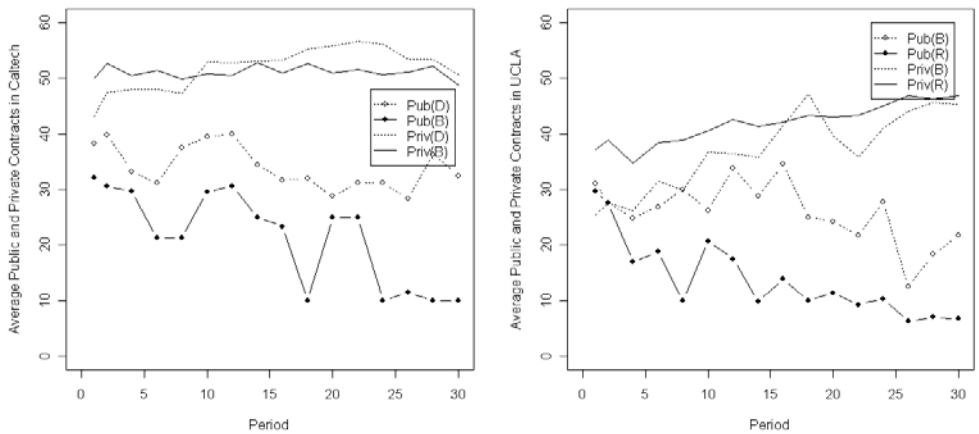

Figure 4.3: Stochastic noise delays separation of the two tiers of the labor market in Caltech (L) and UCLA (R).

2008). Private offers in baseline provide the workers with a bonus, but in downturns comes with a premium charge for the workers. Offers with higher surplus sharing are accepted more readily: higher wages and lower desired efforts. Effort is a function of increase in wage and requested effort. A contract renewal occurs when a firm makes a private offer to the employee hired in the previous period. In downturns firms are less likely to renew contracts. As we have seen before, the level of relational contracting in UCLA is slightly lower: equilibrium wages and effort are lower in UCLA and renewals are also lower.

Table 4.5 includes all 30 periods with controls for each period, relationship length, and last period effect. As play continues, relationship length becomes increasingly important: it increases wages offered, rate of acceptance, and contract renewals. The private offer bonus continues to be lacking in the downturn. Because stable relationships are relatively harder to come by for workers in downturn, the workers are willing to reward long term firms more: long term offers are even more likely to be accepted and rewarded with high effort in downturn. Firms' unwillingness to renew contracts gradually lessened as periods go on, the coefficient in the contract renewal for downturns continue to be negative, but is no longer significant. Other factors such as whether the delivered effort level meets expectation (surprises), relationship lengths, and whether the contract was private become more important.

Overall the result suggests that in a finitely repeated game, delay in forming relationships have long lasting repercussions. Firms in downturn correctly anticipated that with

	OLS: Wage Offer			Logit: Acceptance			OLS: Effort			Logit: Renewal		
	Coef		S.E	Coef		S.E	Coef		S.E	Coef		S.E
Intercept	-6.13		7.40	3.02	*	1.63	0.73		1.12	1.63		1.41
Wage				0.07	**	0.03	0.07	**	0.03			
Surprise										0.25		0.22
Desired Effort	4.84	***	0.87	-0.52	**	0.22	0.45	**	0.20			
Actual Effort										-0.12		0.18
Private	6.43	*	3.34	-2.66	***	0.65	0.62		0.74	0.86		0.91
UCLA	-5.38	**	2.29	-0.14		0.39	-1.19	**	0.48	-1.27	*	0.67
Downturn	11.36		10.21	-1.73		2.01	0.42		1.65	-6.53	***	2.46
Downturn*Wages				-0.02		0.03	0.00		0.04			
Downturn*Surprise										0.08		0.40
Downturn*DesEffort	-1.15		1.30	0.25		0.29	-0.08		0.27			
Downturn* Ac.Effort										0.68	**	0.32
Downturn*Private	-10.37	**	4.71	0.64		0.84	-0.76		0.99	2.00		1.56
N	186			186			71			71		
Adjusted R2	0.29			0.40			0.56			0.43		

***p=<1%, **p<5%, * p<10%

Table 4.4: Initial ($t = 1$) determinants of firm and worker behavior.
Note: Firms in the Downturn treatment start out with lower private wages (-10.37) and are less likely to renew private offers (-6.53). This suggests that firms are being choosy about their private contracting partner under the threat of stochastic interruptions.

	OLS: Wage Offered			**Logit: Acceptance**			**Ols: Effort**			**Logit: Renewal**		
	Coef		S.E	Coef		S.E	Coef		S.E	Coef		S.E
Intercept	-16.38	***	2.67	-1.90	***	0.48	-0.08		0.38	-3.97	***	0.46
Wage				0.05	***	0.01	0.10	***	0.02			
Surprise										0.39	***	0.09
Desired Effort	4.45	***	0.34	-0.06		0.06	0.28	***	0.08			
Actual Effort										0.13	***	0.05
A.r.l	18.42	***	2.17	6.58	***	0.71	0.33		0.36	3.44	***	0.42
Private	11.59	***	2.18	-2.57	***	0.36	0.94	***	0.26	2.01	***	0.37
Period	0.05		0.07	0.08	***	0.01	0.00		0.01	0.03	**	0.01
lastPeriod	-0.93		1.08	-0.98	***	0.19	-2.18	***	0.37	-14.20	***	0.69
UCLA	-3.32		1.61	0.73	***	0.27	-0.49	***	0.14	0.28		0.26
Downturn	7.97	**	3.92	1.49	**	0.60	0.44		0.55	-1.03		0.86
Downturn:Wage				0.00		0.01	-0.02		0.03			
Downturn:suprise										-0.11		0.19
Downturn:DesiredEffort	-0.09		0.52	-0.15	*	0.08	-0.03		0.15			
Downturn:ActualEffort										0.17		0.10
Downturn:A.r.l	-0.21		3.13	1.12	**	0.55	1.08	*	0.61	-0.74		0.68
Downturn:Private	-8.61	***	2.97	-0.05		0.29	0.18		0.35	0.13		0.58
Downturn:Period	0.13		0.11	0.01		0.01	0.01		0.01	-0.01		0.02
N	3435			3434			2091			2091		
adj R2	0.72			0.64			0.77			0.60		

***p=<1%, **p<5%, * p<10%

Robust standard error. For wage offered and contract renewal, standard error is clustered on firms. For acceptance and effort, standard error is clustered on workers.

Table 4.5: Determinants of firm and worker behavior, pooled across all periods. Note: Firms in the Downturn treatment continues to lower private wages (-8.61) but are no longer reluctant to renew private offers as the period progresses.

fewer firms active, workers experience higher competition for jobs. However, firms reacted by demanding more from a worker before extending long term relationships. First, they offer lower surplus sharing in private offers. For the earliest five periods, private offers in downturns are actually less fair than public offers (average private offer price is 30.5 with a surplus of 0.22 compared to public offers of 29.2 with surplus share of 0.29). This is the opposite of the situation in baseline, where average private offer prices of 34 (and surplus sharing of 0.23) serve as a bonus for the worker whose outside options are public offers of 24 (with surplus sharing of 0.18). Since fair surplus offers are accepted more frequently, private contracts constitute a significantly smaller fraction of total contracting in downturn market (42.5% in downturn) than in baseline (51.5% in baseline). This difference is significant at $p < .10$ with a one sided t-test).

The difference in surplus offered leads to further delay in relationship formation. Public offers are renewed less frequently than private offers. In addition, since public offers are more fair in downturn, workers' effort matches firms' expectations much more closely in downturns. Because public offers in downturn carry little penalty in terms of lower profits for both parties or suprises to firms,[11] there is little incentive for both partners to move away from public offers. Average rate of renewals in the first 5 periods of downturn is 41%, compared to 49% in baseline.

Even though these early period behaviors (such as penalty for private offers and firms' unwillingness to renew) decrease in time,[12] the effects lingered. In fact, aggregating across all periods, the difference in workers' earnings in a public offer (24.6) and a new private offer (26.9) is not statistically significant ($p = 0.12$). This is starkly different from baseline, where the average seller earning from public offer is 16.9 and new private offer is 23.85 ($p < 0.01$).[13] As a result, private contracts in downturn are conducted within a shorter relationship: average a.r.l of private contracts in downturn are 8% (0.01) shorter than in baseline while average a.r.l of public contracts in downturns are 1% (0.03) longer than in baseline.[14] Firms' unwillingness to commit in the early periods of downturns ironically

[11]Firms' average surprise from public offers is -1.36 (0.22) in the baseline and -0.31 (0.18). Baseline public contract suprises is significantly more negative ($p < 0.01$).

[12]Table 4.3 Row 3, Private contracts make up 81% of total contracts in baseline vs 78% in downturns

[13]For firms, the difference between public offer earnings and new private offer earnings is 19.8 vs 24.6 in downturn ($p < 0.05$), and 18.3 vs 28 in baseline ($p < 0.01$).

[14]The fact that the difference in a.r.l (adjusted relationship length) of private contracts is higher (-8%) than that of total contracting (-6%) suggests that as firms use private offers renewals less, repeated interactions are partially driven by workers who sought after public offers from their previous firms.

results in higher reservation prices for workers and loss of bargaining power due to shorter relationships, which requires the firm to pay much more in later periods for the same level of effort.

4.4.4 50-50 Surplus Sharing and Loyalty Norms

In the previous section we see that stochastic noise slows down the formation of long term private contracting. In this section we ask whether downturn affects established relationships by looking more closely at the periods surrounding downturns.

There were 79 instances of downturn. Among them, there were 70 instances where firms have the opportunity to reconnect with the workers they have previously traded with. At the period before the downturn, average wage contracted was 41.81 with average desired effort level of 7.9 (this implied offered surplus sharing of 0.39). Workers provide an average effort level of 7. (See Table 4.10 in the Appendix). When downturn prevents firms from hiring, these workers go to a lower quality (short term) market where contracts are lower in wage and surplus sharing (Table 4.10 Column 2). 2/3 of the total 118 interim trades came from private offers. Those workers that have received many private offers in the past or have longer relationship lengths are more likely to get private offers (Table 4.6 Model 1). Workers simply grabbed whatever offers they have (Table 4.6 Model 2) - the primary determinant of their activity is the number of private offers they receive. Workers treat these contracts with the same norm of fairness as they had with their previous employers, delivering effort that ensure themselves 0.68 of the total surplus.

Firms try to squeeze a little bit upon return (lowering surplus offers from 0.39 to 0.35) (Table 4.10 Column 3). From the 70 firms returning from downturn, 40 tried to reconnect with previous workers. Reconnection attempts depend mostly on previously delivered effort and relationship length (Table 4.6 Model 3). The 30 other firms attempted to hire other workers (14) or made public offers (16).[15] Six reconnection attempts were rebuffed by workers, and two public offers were grabbed by previous workers, thus continuing the relationship. In the end, the reconnection rate was 54%.

Table 4.6 Model 3 and 4 model the reconnection norms.[16] *Firm attempted* is 1 if the

[15]The public offers made by returning firms are often more generous in surplus sharing than average public offers (the mean returning firms' public offer is 0.37 as opposed to 0.29), which help keep the high surplus sharing norms in the downturn public offer.

[16]Model 1 is an OLS on the number of private offers received by workers left unemployed by downturns.

	1. OLS: # interim private offers			2. OLS: # of interim trades			3. Logit: firm attempt to reconnect			4. Logit: Pair reconnected		
	Coef		S.E	Coef		S.E	Coef		S.E	Coef		S.E
Intercept	1.34	**	0.53	1.01	***	0.36	-2.73	*	1.62	-3.92	**	1.96
Pre downturn variables:												
Wage	-0.02	**	0.01	0.00		0.01						
Delivered Effort							0.36	***	0.13	0.30	**	0.14
Offered Surplus	1.23	*	0.74	-0.25		0.42	-0.44		1.61	4.98	***	1.88
A.r.l (seller)	-0.21		0.63	0.56		0.38						
A.r.l (buyer)							2.70	***	1.03	2.45	**	1.09
Number of past private offers to seller	0.83	***	0.25	-0.19	*	0.10	0.21		1.27	0.39		1.40
# interim private offers				0.43	***	0.07						
# interim trades							-0.55	*	0.33	-0.96	***	0.38
N	70			70			70			70		
Adjusted R2	0.16			0.35			0.43			0.51		

***p=<1%, **p<5%, * p<10%

Robust standard error. For interim behavior, standard error are clustered on workers. For post downturn behavior, standard errors are clustered on firms.

Table 4.6: Determinants of firm and worker behavior surrounding downturn.

firm made a private offer to its previous employee and *Pair reconnected* is 1 if the firm made such an offer and it was accepted. More interim job searches lower the chance that firms attempt to reconnect and that workers accept reconnection offers, which supports our model of strong reconnection norms.[17]

There is a very striking effect of pre-downturn surplus sharing on successful reconnection: Higher previous surplus offers have a positive effect on successful reconnection and a negative effect on worker rejection of reconnection (i.e., workers tend to reconnect more often if the employer shared more surplus). In fact, Table 4.7 shows that the average predownturn surplus sharing of succesful reconnections is 0.51 - close to the 50-50 split. If the worker share was less than this, firms attempt to reconnect but will be refused by workers. If the worker's share was more than this, firms do not attempt to reconnect. The firms attempt and reciprocation from worker also depends on relationship length and wages offered. In conclusion, strong loyalty norm can protect relationships against stochastic interruptions:

Model 2 is an OLS on the number of trades these workers engage in. Model 3 is a logit regression on whether or not firms make a private offer to its previous worker after returning from the downturn. Model 4 is a logit regression on whether the firm-worker pair reconnected after the downturn.

[17]Note the sample sizes are low (N=79) because downturns are rare ($\delta <= 0.1$), so strong results are unlikely to emerge.

	Reconnection					
	No attempt by firm		Rejected by worker		**Accepted by worker**	
	Mean	S.E	Mean	S.E	Mean	S.E
Price	33.92	3.58	37.00	6.35	47.97	2.74
Offered Surplus	0.34	0.03	0.31	0.04	0.43	0.03
Actual Surplus	0.87	0.23	0.38	0.07	**0.51**	**0.03**
Relationship Length	2.80	0.54	4.50	1.18	7.89	1.02
	N=30		N=6		N=36	

Table 4.7: Characteristics of pre-downturn contracts and post-downturn relationships: 50-50 surplus sharing leads to reconnections.

good workers and fair firms reconnect after the downturn.

4.5 Conclusion

This paper replicates the important experimental paper on wage-effort gift exchange by Brown, Falk and Fehr (2004) and extends it by creating exogenous layoff periods ("downturns" in which firms cannot hire for three periods). The key feature of their design, compared to earlier experiments, is that firms can make private offers that can only be accepted by a specified worker. The result is that a "two-tiered labor market" emerges spontaneously (even there are no true skill differences among workers). Some firms make public offers which any worker can accept and others lock in to "relational contracts" in which they make repeated offers to the same worker period after period. Firms sometimes "fire" their worker-by not repeating their private offer– if the worker's effort is too low, compared to a nonbinding level of effort requested by the firm. It is crucial to note that these are not truly long-term contracts since they only last one period, and there is no communication at all that supports the relationship or clarifies what is expected, except for the time course of wages, effort requests, and actual efforts.

Our experiments replicate the basic patterns of contracting, wages and effort in BFF1 quite closely using two different subject pools and extending number of contracting periods from 15 to 30. The efficiency of the UCLA market is lower than that of the Caltech sessions which is slightly higher than the BFF averages. Wage, effort, private offers, and relationship length are remarkably consistent.

The novel contribution of our paper is exploring the effect of exogeneous temporary drops in labor demand (firm downturns). We explored the hypothesis that anticipating stochastic downturn would undermine relational contracting. In a rational framework the lower capitalized that long-run value of a relational contract will lead workers to exert low effort and firms to expect low effort and pay lower wages. On the other hand, if the market is driven by strong "loyalty" norms that firms rehire their old (pre-downturn) workers then relational contracting might be immune to disruptions.

Consistent with our prediction, job rents in downturns are lower. This leads to shorter relationships: the average relationship length in private contracts are 8% lower in downturn and there is a 13% increase in trade taking place in relationship where partners' total trading history is 4 periods or less. Suprisingly, average wages in downturn are higher.[18] Since average effort remains the same, this results in a shift of surplus sharing towards the workers.

The reason for this is the delayed separation of public and private contracts in labor market with downturns. In the baseline markets, firms provide a private offer bonus to workers in terms of higher wages and better surplus sharing compared to public offers. In the earlier periods of downturn, not only is the private offer bonus missing, but the surplus sharing is actually less fair than public offers. Renewal rates of comparable effort levels are also lower in downturns. One possibility is that firms in downturn correctly anticipated that with other firms unable to hire due to downturns, workers will compete harder for jobs. However, firms' initial strategy of demanding more from a worker before extending long term relationships actually backfired and resulted in higher reservation prices for workers (in the form of a much improved temporary market). The net effect is that firms' bargaining power is reduced and firms ended up paying much more for the same level of effort. Overall the result suggests that in this finitely repeated contracting game, delay in forming relationships have long lasting repercussions.

We observe the following loyalty norms. When downturn prevents firms from hiring,

[18]BFF2 studied relational contracting under full employment by inverting the labor market conditions of BFF1; their market is populated by 10 firms and 7 workers. Relationship lengths in BFF2 are shorter, confirming field studies (Bleakely et al., 1999) that workers are more likely to quit their jobs under full employment than when unemployment prevails. However, the lower frequency of long-term relationships does not affect aggregate performance across labor market conditions because it is compensated for by higher performance in short-term relationships. With no unemployment, high-performing workers receive higher wage offers from their current firm than from outside firms. This motivates workers to perform at a high level of effort, rather than to shirk and then switch firms.

workers left unemployed go to a lower quality (short term) market where contracts are lower in wage and surplus sharing. Workers simply grabbed whatever offers they could. When firms return from the downturn, there is a very striking effect of pre-downturn surplus sharing on successful reconnection. The average predownturn surplus sharing of succesful reconnections is 0.50 - the 50-50 split that has appeared frequently in various laboratory and field experiments. If the worker share was less than this, firms attempt to reconnect but will be refused by workers. If the worker's share was more than this, firms do not attempt to reconnect.

A wonderful property of the gift exchange paradigm is that one can think of many natural experiments to do next. The downturns here are independent across firms, but one could correlate them to study business cycle effects. Typically downturn lengths are stochastic rather than lasting a fixed number of periods, which is an easy design feature to change.

4.6 Appendix: Additional Tables and Figures

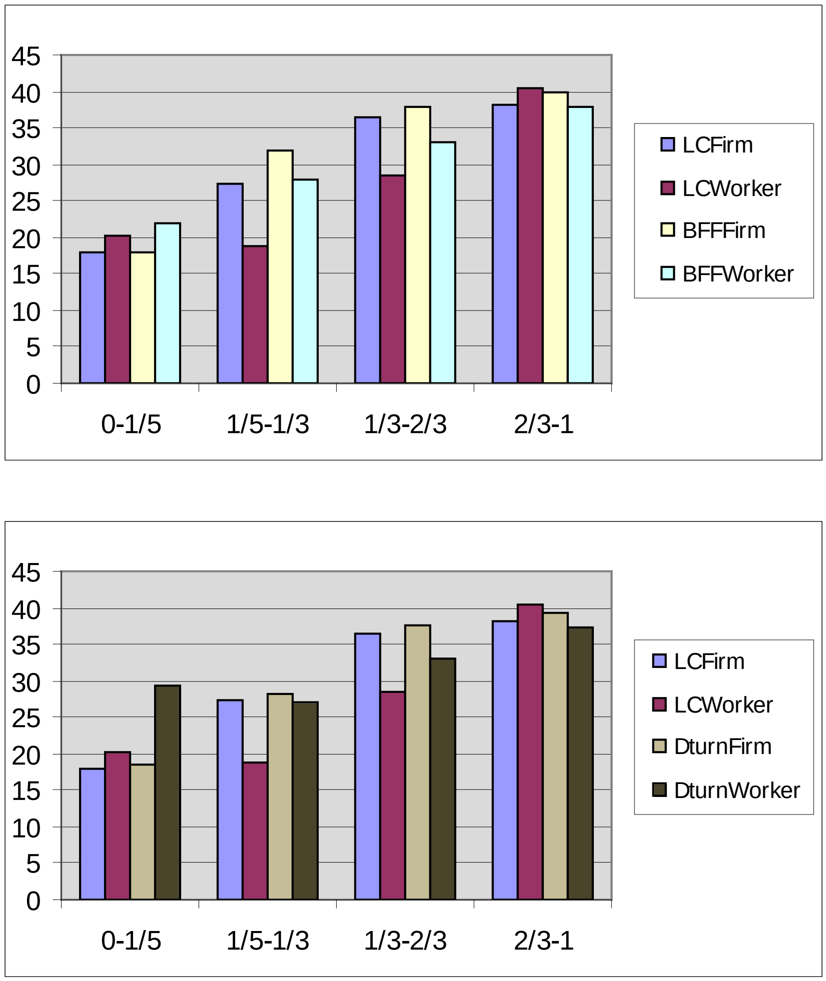

Figure 4.4: Surplus sharing (y) as a function of relationship length (x).
In the top panel, 30 period baseline LC sessions are compared to 15 period BFF sessions by binning relationship length into fractions of entire trading period. For example fraction 0-1/5 indicated relationship lengths up to 5 (out of 30) periods in fLC and up to 2 (out of 15) periods in BFF. Surplus sharing in shortest relationship (temp market) is filled with shirking, which hence benefits workers. The next shortest relationships are entry level of the high tier relationship where workers put in their dues to gain trust of the firm, sacrificing their share. As the relationship grows, the surplus sharing becomes more equal. The bottom panel shows the similarity of downturn sessions with the baseline session. Because wages are higher, workers in temp market gain a much larger share of the surplus when they shirk.

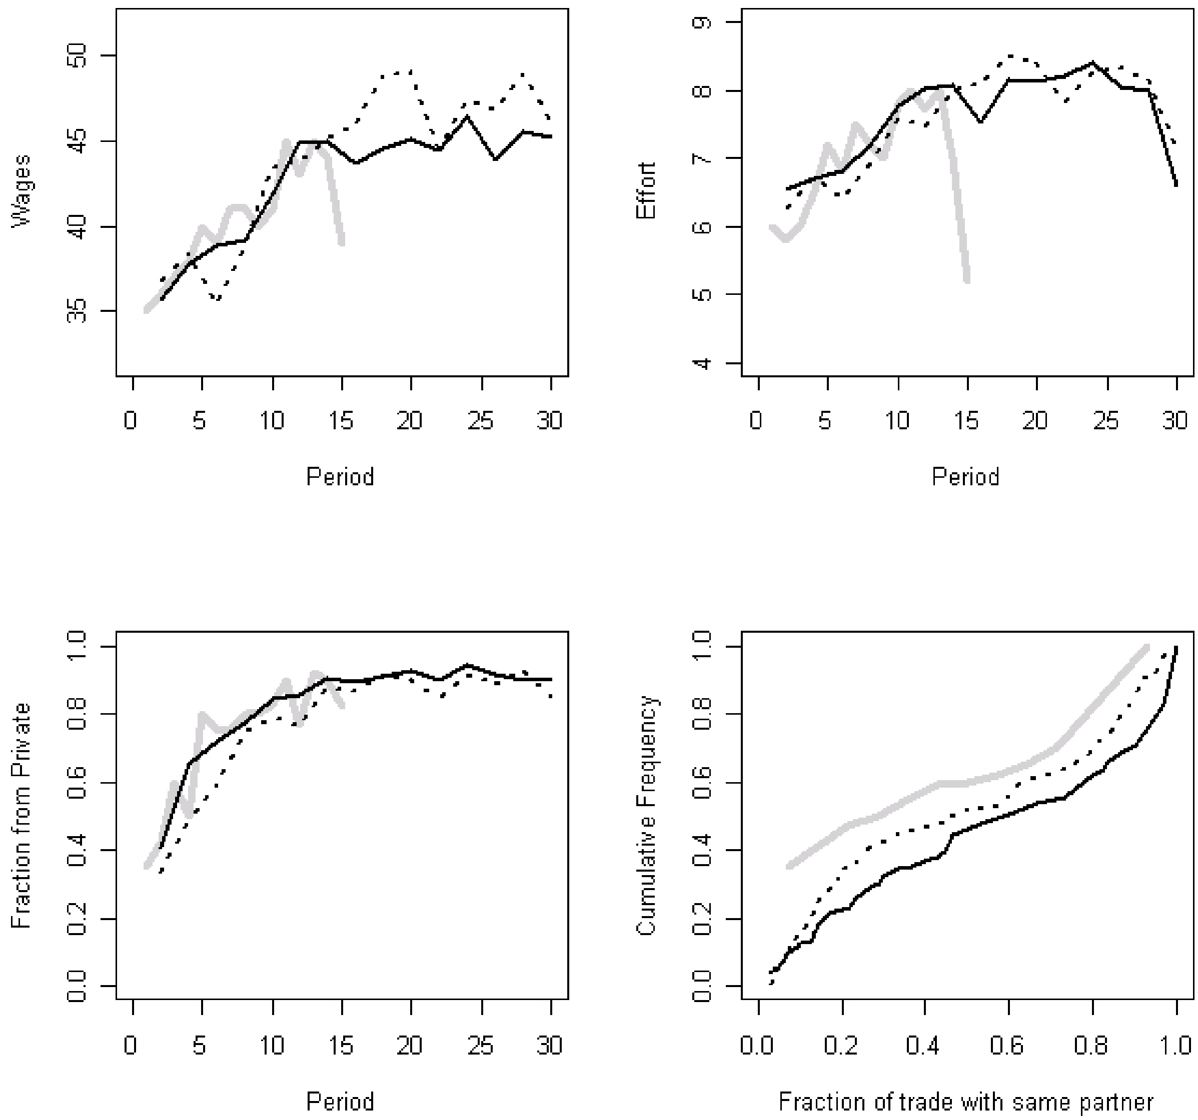

Figure 4.5: Wages, effort, private offers and relationship lengths in Baseline (solid) compared to Downturn (dotted).
The top left panel illustrates that average wages are suprisingly higher in the second half in Downturn (dotted lines) than it is in Baseline (solid lines black for LC, gray for BFF). The top right panel and bottom left panel shows no difference in average delivered effort or in fraction of contracting through private offers. The bottom right panel shows that relationships are shorter in downturns. Graph shows data pooled from Caltech and UCLA subjects.

	Model 1	Model 2	Model 3	Model 4
Downturn	2.99* (1.30)	2.54* (1.08)	2.42*** (0.46)	11.31** (3.61)
Desired effort	5.46*** (0.34)	5.47*** (0.30)	4.67*** (0.28)	4.80*** (0.32)
Relationship Length	0.68*** (0.07)			
A.r.l		16.92*** (2.49)	15.31*** (2.03)	17.19*** (2.89)
Private			7.74*** (1.65)	9.94*** (1.43)
Period			0.11** (0.04)	0.12* (0.05)
lastPeriod			-1.82* (0.9)	-1.81* (0.9)
UCLA			-2.87*** (0.77)	-3.09*** (0.83)
Dturn* Desired effort				-0.38 (0.51)
Dturn*A.r.l				-4.36 (3.64)
Dturn*Private				-4.08 (2.97)
Dturn*Period				0.00 (0.07)
Intercept	-11.12*** (3.28)	-15.05*** (3.03)	-13.78*** (2.9)	-17.62*** (3.16)
N	2105			
Adj R2	0.62	0.63	0.67	0.68

Table 4.8: Determinants of accepted wages.

	Baseline		Downturn	
	diff	**diffSE**	**diff**	**diffSE**
Offers:				
Wage	18.30	0.93	10.89	0.92
Desired Effort	0.78	0.16	1.25	0.16
Worker/Total Surplus	0.16	0.01	0.01	0.01
Contracts:				
Wage	25.28	1.10	16.27	1.01
Desired Effort	2.07	0.18	1.58	0.18
Worker/Total Surplus	0.18	0.02	0.05	0.02
Relationship Length	8.04	0.30	6.00	0.25
A.r.l	0.25	0.02	0.17	0.02
Actual:				
Delivered Effort	4.29	0.21	3.20	0.18
Suprise	1.02	0.19	0.36	0.15
Buyer Profit	17.65	1.31	15.72	1.39
Seller Profit	16.08	0.81	9.06	0.80
Worker/Total Surplus	-0.06	0.04	-0.20	0.06
Renewed Immediately	-0.61	0.03	-0.58	0.03

Table 4.9: Difference between private and public contracts are diminished in downturn.

	1.Firm: Pre-downturn		**2. Worker: Interim**		**3. Firm: Post-downturn**	
	Mean	S.E	Mean	S.E	Mean	S.E
Wage	41.81	2.20	33.84	1.51	41.00	1.99
Desired Effort	7.90	0.33	7.83	0.23	8.20	0.29
Offered Surplus	0.39	0.02	0.29	0.02	0.35	0.02
Fraction Private	0.80	0.05	0.67	0.04	0.78	0.05
Relationship Length	5.71	0.66	2.88	0.23	5.87	0.70
Delivered Effort	7.03	0.37	5.90	0.28	7.09	0.36
ExpectedEffort	7.16	0.35	6.23	0.24	7.41	0.31
Suprise	-0.13	0.17	-0.33	0.18	-0.32	0.16
Buyer Profit	28.47	2.25	25.14	1.84	29.87	2.29
Seller Profit	30.76	1.56	25.35	1.11	29.94	1.44
Actual Surplus	0.64	0.09	0.68	0.08	0.60	0.07
Renewed Immediately	0.61	0.06	0.46	0.05	0.57	0.06
	N=79		N=118		N=70	

Table 4.10: Contracts before, during, and after downturns.

4.7 Appendix: Proofs

Proposition 4.2.1

Proof. Selfish workers will always provide $e = 1$ in the one period game. Fair worker will shirk if a contract is unfair, hence firms are restricted to making fair offers $\hat{w}(e) = 5e - c(e)/2$. Firm's expected utility for making a fair offer for effort level e is $\pi(e) = p(10e - \hat{w}(e)) - (1 - p)(10 - \hat{w}(e)) = 10pe - \hat{w}(e) + 10 - 10p$. Since there is an excess supply of labor and because firms do not have any information about individual workers' fairness, there is competition among workers for firms and no competition between firms for workers.[19] This allows firms to maximize their profit. Taking a derivative of $\pi(e)$ over $\hat{w}(e)$ and substituting $c(e)$ for $\hat{w}(e)$, we find that firm's profit is maximized when $c'(e) = 20p - 10$. Since $c'(e) = 1$ for $e = 2$,$c'(e) = 2$ for $3 \leq e \leq 8$, $c'(e) = 3$ for $e > 8$, we can solve for minimum and p that support each effort level and arrive at the PBE above. □

Proposition 4.2.2

Proof. The steps of our proof are adapted from BFF2 for multiple firms and excess supply of labor. We prove this more generally for any $\tilde{e}_T$, $\tilde{e}_t$ for all $t < T$, $w = \hat{w}(\tilde{e})$ when $p = .6$, which includes $\tilde{e}_T = 8$,$\tilde{e}_t = 10$.

Step 1 (behavior of fair workers): All fair workers will always perform the desired effort since all wages will be fair.

Step 2 (behavior of selfish workers): In final period T, since there are no unmatched firms, value of unemployment $U_T = 5$. Since $V_T = \hat{w}(\tilde{e}_T) > U_T$, it is a uniquely best strategy of the selfish worker to accept the contract of the incumbent firm and then deliver $e = 1$. In any period $t < T$, since firms renew their private offers when workers deliver requested effort, expected future utility of not shirking is:

$$V_t = \hat{w}(\tilde{e}_t) - c(\tilde{e}_t) + V_{t+1} \tag{4.4}$$

and since firms who experienced shirking do not offer private contracts to workers who were previously matched and after $t = 1$ public offers are $[5, 1]$, expected future utility of not

[19] Suppose Firm A and B make an offer to the same person, who rejects A for B. Firm B and the worker are then matched, and Firm A can make another offer to the remaining $n > 1$ workers as a monopolist. Since time is not costly, there is no loss of utility here for Firm A for losing its initial worker to Firm B.

shirking is:

$$U_t = 5 + U_{t+1} \tag{4.5}$$

Since $V_{t+1} - U_{t+1} > c(\tilde{e}_t)$ for all $\tilde{e}_t$ above, it is therefore the best strategy for $e_t = \tilde{e}_t$.

Step 3 (firm behavior): Given the discussion in Prop 4.2.1 firms are not under competition for workers and can maximize profits at all periods. Since $p = .6$, firms' are indifferent between contracts of $c'(e) = 2$ at period T. At period T-1, firm offers $[\hat{w}(\tilde{e}_{T-1}), \tilde{e}_{T-1}]$ that corresponds to $\tilde{e}_T$ that was previously chosen (Table 4.2 Row 5). For $\tilde{e}_T > 4$, $\tilde{e}_{T-1} = 10$ can be supported, and hence $\tilde{e}_t = 10$ for all $t < T$. For $\tilde{e}_T \leq 4$, T-1 offers are $\tilde{e}_{T-1} = 9$ for $\tilde{e}_T = 4$ and $\tilde{e}_{T-1} = 7$ for $\tilde{e}_T = 3$. These two period T and T-1 can then support $\tilde{e}_t = 10$ for all $t < T - 1$. All firms are matched initially at $t = 1$ through identical public offers of $[59, 10]$, in equilibrium all following offers are private renewals to previous trading partner. Hence $n - 2$ workers remain unemployed for all t. □

Proposition 4.2.3

Proof. As before, let $j_t = \hat{w}(\tilde{e}_t) - c(\tilde{e}_t)$. Without downturn, job rents are equal to difference in earnings from current period and expected job rents from future periods.

$$V_t - U_t = j_t + V_{t+1} - 5 - U_{t+1} = j_t - 5 + (V_{t+1} - U_{t+1})$$

where $V_T - U_T = w_T$ for selfish types and $w_T - c(e_T)$ for fair types.
With downturn, job rents are

$$V_t^\delta - U_t^\delta = (1-\delta)(j_t + V_{t+1}^\delta) + \delta\mu(j_t + V_{t+1}^\delta) + \delta(1-\mu)(5 + j_{t+1} + V_{t+2}^\delta) - \mu(j_t + V_{t+1}^\delta) - (1-\mu)(5 + U_{t+1}^\delta)$$

Simplifying we arrive at

$$= (1-\mu)\left[(1-\delta)(j_t - 5) + (\delta - \mu)(j_{t+1} + V_{t+2}^\delta) - (1-\mu)(5 + U_{t+2}^\delta)\right]$$

$$= (1-\mu)\left[(1-\delta)(j_t - 5) - \mu(j_{t+1} - 5 + (V_{t+2}^\delta - U_{t+2}^\delta) + \delta j_{t+1} - 5 + \delta V_{t+2}^\delta - U_{t+2}^\delta\right]$$

$$= (1-\mu)\left[(1-\delta)(j_t - 5) - \mu(\delta V_{t+1} - U_{t+1})\right]$$

where $V_T^\delta - U_T^\delta = (1-\mu)(1-\delta)(V_T - U_T)$.

(i) Since $j_t > 5$ for all $e_T > 1$ and $V_{t+1} > U_{t+1}$, job rents are always positive.

(ii) Since current period job rent $j_t - 5$ is discounted by $(1-\mu)(1-\delta)$ and future period rents are also discounted, job rents are lower in the downturn. □

Proposition 4.2.4

Proof. The strategy profile describes a very strong reconnection norm ($\mu = 1$). We prove that this reconnection norm is a Perfect Bayesian Equilibrium with the following steps:

Step 1 (behavior of fair workers): All fair workers will always perform the desired effort since all wages will be fair.

Step 2 (behavior of selfish workers): As in the baseline case, it is a unique best strategy of the selfish worker to accept the contract of the incumbent firm at T and perform $e_T = 1$. It is also best response to not shirk, since no firms make private offers to previously employed workers, driving the outside option to 5.

Step 3 (firm not hit by downturn at $t = 1$): Proof follows the baseline case. We now show that if a worker shirks, the firm's best response is to extend a private offer to never employed workers or to make a public offer of $[5, 1]$. Given out of equilibrium belief that only selfish workers shirk, when no firms are hit by a downturn, all previously employed workers are selfish. When the other firm is in downturn, his previously employed worker enters the pool of unemployed workers for one period. If the firm can identify this worker, the best strategy is to privately offer him the single period contract of $[46, 8]$, since the worker will return to his employer after the downturn.[20] However, since firms only observe which workers are employed but do not observe pairings between workers and firms, it cannot perfectly identify this worker. Because the population has $n-1$ selfish workers and one worker with probability .6 of being fair, the firm cannot do better than to offer $[5, 1]$.

Step 4 (firm hit by downturn at $t = 1$): If the firm competes with an incumbent firm for his worker, he will earn 0 profit, which is lower than the profit of 5 from public offer of $[5, 1]$. Hence a firm unable to make offers at the first period make offers to never employed workers. Given out of equilibrium belief that only selfish workers shirk and the inability to perfectly identify workers left by downturn in Step 3, firms cannot do better than to make

[20] Trying to prevent this worker from returning to his employer will result in a profit of 0 for the firms, which is lower than the profit of 5 from public offer of $[5, 1]$.

a public offer of $[5, 1]$. $\square$

4.8 Appendix: Experimental Instructions

Instructions for Buyers

You are now taking part in an economic experiment. Please read the following instructions carefully. Everything that you need to know in order to participate in this experiment is explained below. Please note that communication between participants and usage of computer for other purposes is strictly prohibited during the experiment. Should you have any difficulties in understanding these instructions please notify us.

At the beginning of the experiment you will receive an initial endowment of **20 Francs**. During the course of the experiment you can earn money by gaining **points**. The amount of points that you gain depends on your decisions and the decisions of other participants. Points will be exchanged into Francs at the end of the experiment at the rate of **1 point = 0.10 Francs**

At the end of the experiment you will receive the money that you earned during the experiment in addition to your endowment of 20 Francs.

The experiment is divided into periods. In each period you have to make decisions, which you will enter on a computer screen. There are 30 periods in all.

Prior to the experiment the N participants were divided into 2 groups: B buyers and S sellers. **You are a buyer throughout the whole experiment.** All participants have received an identification number which they will keep for the entire experiment.

An Overview of the Experimental Procedures

In each period of the experiment every buyer can buy a product from a seller. The seller earns a profit by trading if he sells the product at a price, which exceeds his production costs. The buyer earns a profit by trading if the price he pays for the product is less than what the product is worth to him. The production costs and the product's value for the buyer depend on the quality of the product.

Out of the 30 periods, the first 4 periods last 220 seconds to give you an opportunity to familiarize yourself with the software, and the remaining 22 periods last 150 seconds. In each period the procedures are as follows:

1. Each period starts with a **trading phase** which lasts 90 second (or 120 second for the first 4 periods). The first thing that happens is the determination of economic conditions. With a probability of 0.05, each buyer can enter **a downturn** where there is no demand for their products for three periods. For example, if at period 4 a buyer gets hit by a downturn, she will not be able to submit offers (and hence will make 0) until period 7 where a new economic condition will be drawn for her.

2. Buyers not facing downturns can submit offers, which can be accepted by sellers. When submitting an offer a buyer has to specify **three things**:

• which price he offers to pay,
• which product quality he desires,
• and finally, which seller he wants to submit the offer to.

Buyers can submit two types of offers; private offers and public offers. **Private offers are submitted to one seller only** and can only be accepted by that seller. **Public offers are submitted to all sellers** and can be accepted by any seller.

As a buyer you can - in each period - submit as many offers as you like. Submitted offers can be accepted at any time during the trading phase. **Each buyer and each seller can at most conclude one trade in each period.** As there are S sellers and B buyers, several sellers will not trade in each period.

3. Following the trading phase each seller who has concluded a trade determines which product quality he will supply. **The seller is not obliged to supply the product quality desired by his buyer.** Once every seller has chosen a product quality each participant's earnings in the current period are determined. After this the next period starts.

The points gained from all 30 periods will be summed up at the end of the experiment, exchanged into Francs and paid together with your endowment in cash.

The Experimental Procedures in Detail

1. The Trading Phase

Each period starts with a trading phase. During the trading phase each buyer gets random draw of economic condition. With a probability .05 a buyer faces a downturn and where they cannot make offers for three periods. These buyers are indicated with a checkbox of the grid B1, B2, etc. **Buyers that do not face a downturn can submit as many offers as he wishes**. In each trading phase you will see the following screen:

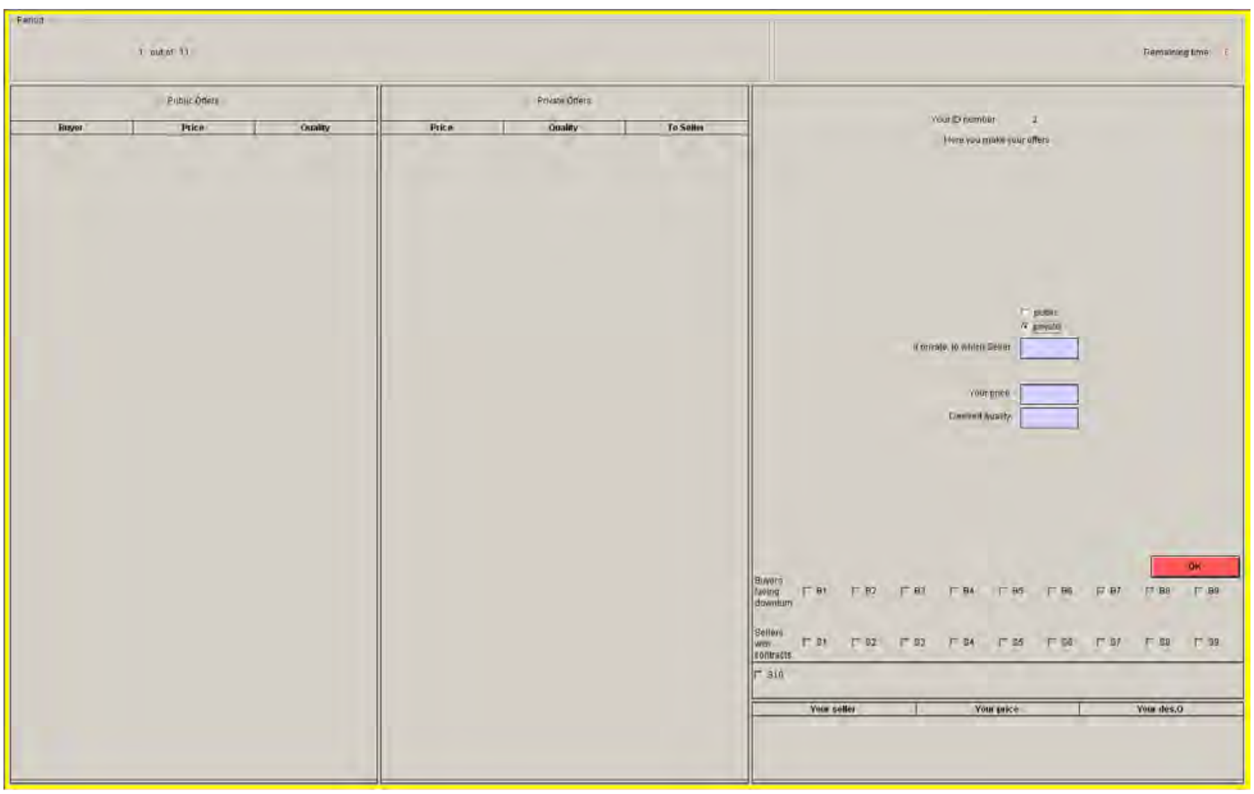

In the top left corner of the screen you see in which period of the experiment you are. In the top right corner of the screen you will see the time remaining in this trading phase, displayed in secs.

When this time is up the trading phase is over. Hereafter, no further offers can be submitted or accepted in this period.

- Once the above screen is displayed the trading phase starts. As a buyer you now have the opportunity to submit offers to the sellers. In order to do so you have to enter three things on the right hand side of the screen:

a) First you have to specify whether you want to submit a public or a private offer:

- **Public trade offers**
  Public offers will be communicated to all participants in the market. All sellers see all public offes on their screens. A public offer can therefore be accepted by **any seller**. As a buyer you will also see all public offers submitted by all buyers.
  If you want to submit a public offer, click on the field "public", using the mouse.

- **Private trade offers**
  A private offer is submitted to **one seller only**. Only this seller is informed about the offer and only this seller can accept the offer. No other seller or buyer will be informed about that offer.
  If you want to submit a private offer, click on the field „private“ using the mouse. After that you specify **which seller** you want to submit the offer to in the field below. Each of the 10 sellers has an identification number (seller 1, seller 2, ..., seller 10). Each seller keeps his identification number throughout the whole experiment. To submit an offer to a specific seller you enter the number of that seller (e.g., "4" for seller 4).

b) Once you have specified to whom you want to submit an offer, you must determine **which price you offer**. You enter this in the field "Your price". The price you offer is a number between 0 and 100:

0 ≤ price offered ≤ 100

c) Finally you have to specify which product quality you desire. You enter this in the field "Desired quality". Your **desired quality** is a number between 1 and 10:

1 ≤ desired quality ≤ 10

After you have completely specified your offer, you must click on the "OK" button to submit it. As long as you have not clicked "OK" you can change your offer. After you click "OK" the offer will be displayed to all sellers you have submitted it to.

- On the left side of your screen you see the header "public offers". All public offers in the current trading phase are displayed here. Your public offers as well as those of all other buyers will be displayed. You can see which buyer submitted the offer, which price he offered and which quality he desired. All buyers also have an identification number, which they keep throughout the whole experiment.

- In the middle of your screen under the header "Your private offers". You see all private offers, which you have submitted in the current trading phase. You see which price you offered, which quality you desired and which seller you submitted an offer to.

- **Each buyer can submit as many private and public offers as he wishes in each period.** Each offer that you submit can be accepted at any time during the trading phase.

- **In any given period each buyer can conclude at most one trade.** Once one of your offers has been accepted you will be notified which seller accepted which of your offers. In the bottom right corner of your screen the identification number of the seller will be displayed as well as your offered price and your desired quality. As you can conclude only one trade in each period all your other offers will be automatically cancelled. Also, you will not be able to submit any further offers.

- **In any given period each seller can conclude at most one trade.** You will be continuously informed which sellers have not yet accepted an offer. On the right bottom of the screen you see S fields, each field for one of the ten sellers. Once a seller has accepted an offer a "x" will appear in the field next to the seller’s identification number. You cannot submit private offers to a seller who has already concluded a trade.

- Once all buyers not facing downturns have concluded a trade or after time have elapsed, the trading phase is over.

- No buyer is obliged to submit offers, and no seller is obliged to accept an offer.

2. Determination of the Product Quality

- Follwoing the trading phase, all sellers who have concluded a trade determine which product quality they supply to their respective buyers. **The product quality, which you desired in your offer, is not binding for your seller.** Your seller can choose the exact quality you desired, but he can also choose a higher or a lower product quality. The product quality which your seller chooses has to be between 1 and 10:

1 ≤ product quality ≤ 10

- While your seller determines the actual product quality, we ask you to specify which quality you expect him to supply on a separate screen. In addition we ask you to state how sure you are about this expectation.

How are the Incomes Calculated?

Your income:

- If you do not conclude a trade during a trading phase you receive an income of 0 points in that period.

- If one of your offers is accepted, your income depends on the price you offered and on the product quality. Your income is determined as follows:

Your income = 10 * product quality - price

- As you can see from the above formula your income is higher, the higher the product quality actually supplied by your seller. At the same time your income is higher, the lower the price you paid for the product.

Income of your seller

- If a seller has not concluded a trade during a trading phase he gains an income of 5 points in that period.

- If a seller has accepted an offer his income equals the price he receives minus the production costs he incurs. The income of the seller is determined as follows:

Income of your seller = Price - production costs

- The production costs of a seller are higher, the higher the quality he chooses. The production costs for each product quality are displayed in the table below:

Product quality	1	2	3	4	5	6	7	8	9	10
Production costs	0	1	2	4	6	8	10	12	15	18

- The income of your seller is higher, the higher the price. Further, his income is higher, the lower the product quality he supplies.

The income of all buyers and sellers are determined in the same way. **Each buyer can therefore calculate the income of his seller and each seller can calculate the income of his buyer.** Further, each buyer and seller is informed about the identification number of his trading partner in each period.

Please note that buyers and sellers can incur losses in each period. These losses have to be paid from your initial endowment or from earnings made in other periods.

You will be informed about your income and the income of your seller on an "**income screen**". On the screen (see below) the following will be displayed:

- which seller you traded with
- which price you offered
- your desired quality
- the product quality actually chosen by your seller
- the income of your seller in this period
- your income in this period.

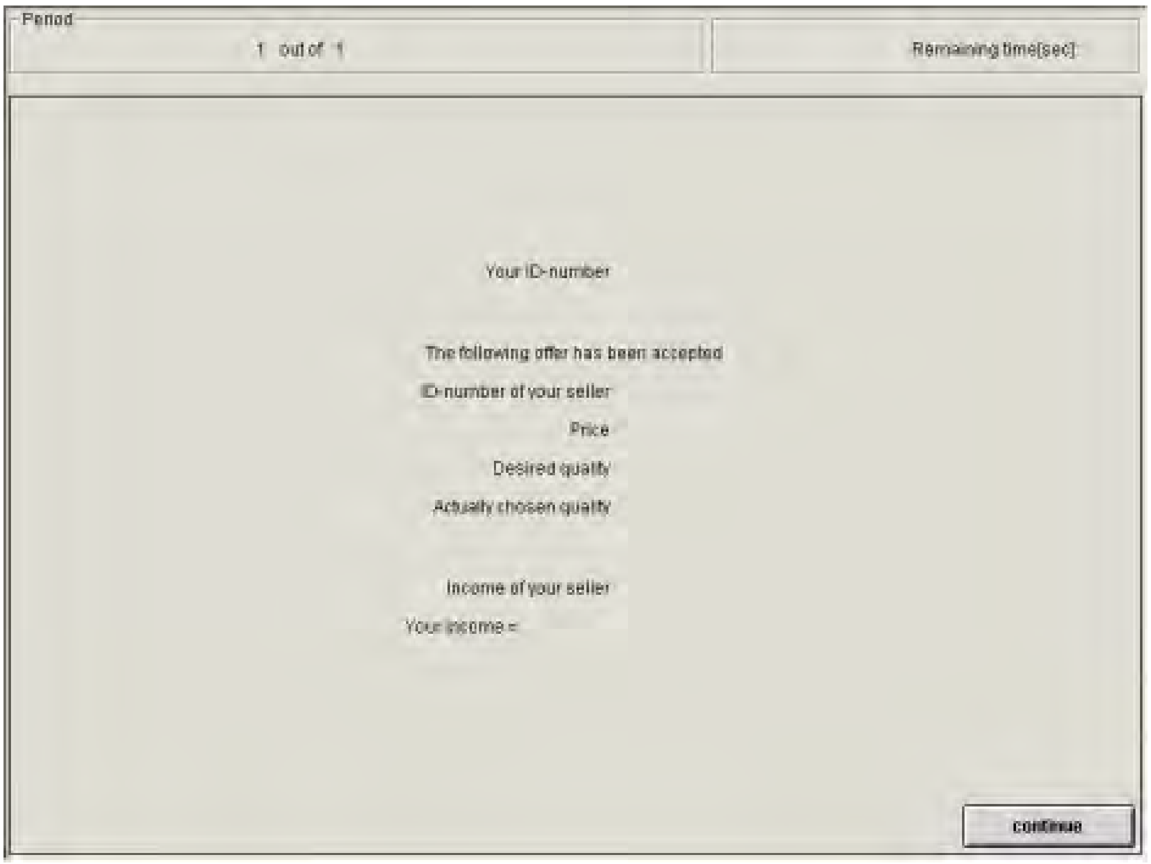

Please enter all the information in the documentation sheet supplied to you. After the income screen has been displayed, the period is over. Thereafter the trading phase of the following period starts. Once you have finished studying the income screen please click on the "continue" button.

The sellers also see an income screen, which displays the above information. They see the ID of their trading partner, the price, desired and actually supplied product quality as well as both incomes.

The experiment will not start until all participants are completely familiar with all procedures. In order to secure that this is the case we kindly ask you to solve the exercises below. There will not be a practice period, so please ask any questions now.

Control Questionnaire

Please solve the following exercises completely. If you have questions ask the experimenter.

Exercise 1
You did not make an offer during a trading phase. What is you income in this period?
Your income =

Exercise 2
You offered a price of 30 and indicated a desired quality of 9. A seller accepts your offer and actually chooses a quality of 8.
Your income =
Income of your seller =

Exercise 3
You offered a price of 60 and indicated a desired quality of 9. A seller accepts your offer and actually chooses a quality of 6.
Your income =
Income of your seller =

Exercise 4
You offered a price of 10 and indicated a desired quality of 2. A seller accepts your offer and actually chooses a quality of 5.
Your income =
Income of your seller =

Exercise 5
You offered a price of 10 and indicated a desired quality of 6. A seller accepts your offer and actually chooses a quality of 2.
Your income =
Income of your seller =

Exercise 6
A seller did not accept an offer during a trading phase. What is the income of this seller in this period?
Income of your seller =

Exercise 7
You made several offers during a trading phase. None of your offers has been accepted by a seller. What is your income in this period?
Your income =

If you have finished the exercises we recommend to look again at the exercises and the solutions provided. After this please think about the decisions you want to make during the experiment.

Documentation sheet: Buyer

This documentation sheet is meant for your orientation. Please complete the respective row in each period.

period	**ID of your seller**	**price**	**desired quality**	**actual quality**	**income of your seller**	**your income**
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						

Instructions for Sellers

You are now taking part in an economic experiment. Please read the following instructions carefully. Everything that you need to know in order to participate in this experiment is explained below. Please note that communication between participants and usage of computer for other purposes is strictly prohibited during the experiment. Should you have any difficulties in understanding these instructions please notify us.

At the beginning of the experiment you will receive an initial endowment of **20 Francs**. During the course of the experiment you can earn money by gaining **points**. The amount of points that you gain depends on your decisions and the decisions of other participants. Points will be exchanged into Francs at the end of the experiment at the rate of **1 point = 0.10 Francs**
At the end of the experiment you will receive the money that you earned during the experiment in addition to your endowment of 20 Francs.

The experiment is divided into periods. In each period you have to make decisions, which you will enter on a computer screen. There are 30 periods in all.

Prior to the experiment the N participants were divided into 2 groups: B buyers and S sellers. **You are a seller throughout the whole experiment.** All participants have received an identification number which they will keep for the entire experiment.

An Overview of the Experimental Procedures

In each period of the experiment every buyer can buy a product from a seller. The seller earns a profit by trading if he sells the product at a price, which exceeds his production costs. The buyer earns a profit by trading if the price he pays for the product is less than what the product is worth to him. The production costs and the product's value for the buyer depend on the quality of the product.

The first 4 periods last 220 seconds to give you an opportunity to learn the software, while the remaining 22 periods last 150 seconds. In each period the procedures are as follows:

1. Each period starts with a **trading phase** which lasts 90 second (or 120 second for the first 4 periods). The first thing that happens is the determination of economic conditions. With a probability of 0.05, each buyer can enter **a downturn** where there is no demand for their products for three periods. For example, if at period 4 a buyer gets hit by a downturn, she will not be able to submit offers (and hence will make 0 points) until period 7 where a new economic condition will be drawn for her.

2. Buyers not facing downturns can submit offers, which can be accepted by sellers. When submitting an offer a buyer has to specify **three things**:

• which price he offers to pay,
• which product quality he desires,
• and finally, which seller he wants to submit the offer to.

Buyers can submit two types of offers; private offers and public offers. **Private offers are submitted to one seller only** and can only be accepted by that seller. **Public offers are submitted to all sellers** and can be accepted by any seller. Buyers can submit as many offers as they like in each period. Submitted offers can be accepted at any time during the trading phase. **Each buyer and each seller can at most conclude one trade in each period.** As there are S sellers and B buyers, several sellers will not trade in each period.

3. Following the trading phase each seller who has concluded a trade determines which product quality he will supply. **The seller is not obliged to supply the product quality desired by his buyer.** Once

every seller has chosen a product quality each participant's earnings in the current period are determined. After this the next period starts.

The points gained from all 30 periods will be summed up at the end of the experiment, exchanged into Francs and paid together with your endowment in cash.

The Experimental Procedures in Detail

1. The Trading Phase

Each period starts with a trading phase. During the trading phase each buyer gets random draw of economic condition. With a probability .05 a buyer faces a downturn and where they cannot make offers for three periods. These buyers are indicated with a checkbox of the grid B1, B2, etc. **Buyers not facing a downturn can submit offers. As a seller you can accept one of the offers.** In each trading phase you will see the following screen:

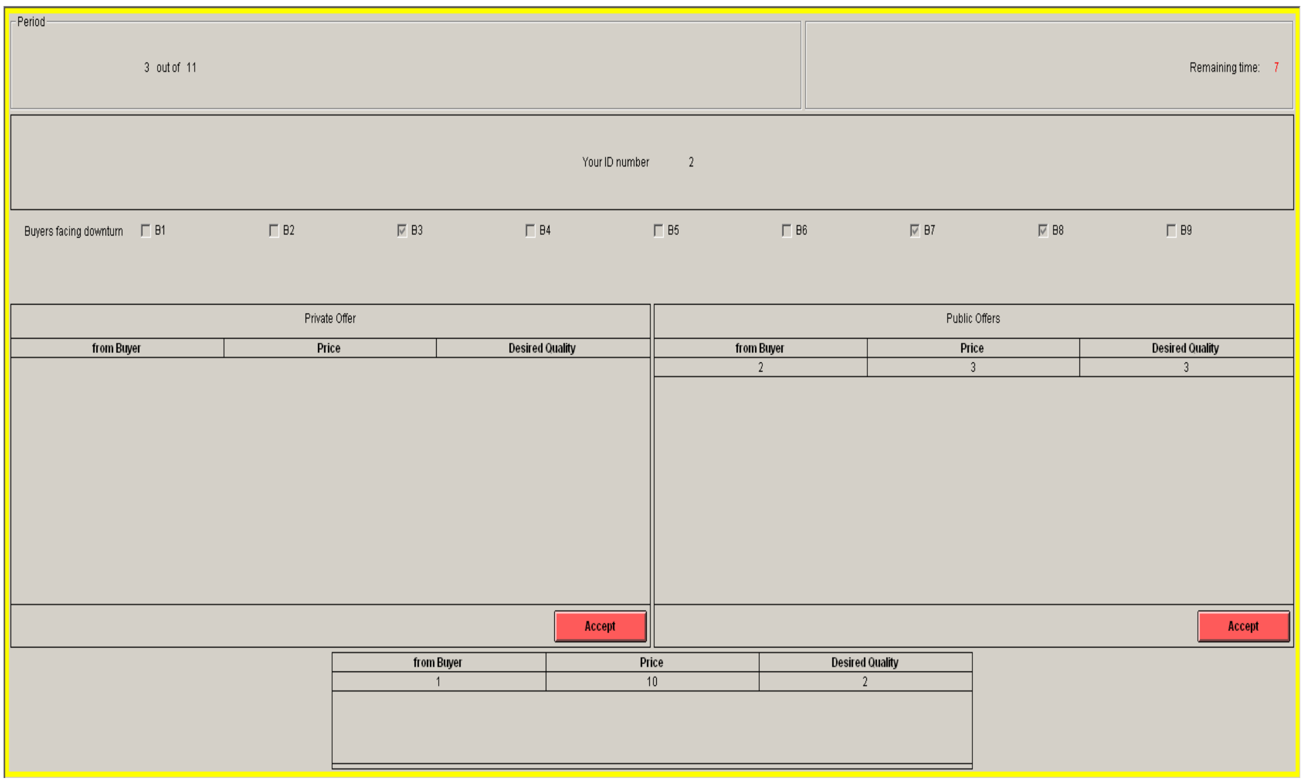

- In the top left corner of the screen you see in which period of the experiment you are. In the top right corner of the screen you will see the time remaining in this trading phase, displayed in seconds. When this time is up the trading phase is over. Hereafter, no further offers can be submitted or accepted for this period.

- Once the above screen is displayed the trading phase starts. As a seller you can now accept offers submitted by the buyers. There are two types of offers which you can accept:
  - **Private offers to you**
    Each buyer has the opportunity to submit private offers to you. **You alone will be informed about these offers and you alone can accept them**. No other seller or buyer is informed about these offers.

If you receive private offers, they will appear on the left side of your screen, below the header "Private offers to you". The offer of a buyer contains the following information: the identification number of the buyer who submitted the offer, the price which he offers for the product and which product quality he desires. If you want to accept a private offer, you click first on the respective row in which the offer is displayed. When you do this, the offer will be highlighted. If you are sure you want to accept the offer you then click on the button "accept" which you find at the bottom of the screen. As long as you do not click "accept" you can alter your choice.

- **Public offers**
  Each buyer also has the possibility to submit public offers. All sellers are informed about these offers and **any seller can accept them**. If a buyer submits a public offer it appears on the right side of your screen, below the header "Public offers". The offer of a buyer again contains the identification number of the buyer who submitted the offer, the price which he offers for the product and which product quality he desires. This information is also displayed to all other sellers and all buyers. If you want to accept a public offer you follow the same procedures as with private offers. You click first on the respective row in which the offer is displayed. When you are sure that you want to accept the offer you click on the button "accept" which you find at the bottom right corner of the screen. As long as you do not click "accept" you can alter your choice.

- As soon as you have pressed the „accept“ button you will see which offer you have accepted in the bottom row of your screen.

- **Each seller can conclude at most one trade in each period**. Once you have accepted one offer you cannot accept any further offers.

All buyers have to observe the following rules when submitting trade offers:

- The price offered by the buyer must be between 0 and 100:

$$\mathbf{0 \leq price \leq 100}$$

- The **desired quality** of the buyer must be between 1 and 10:

$$\mathbf{1 \leq desired\ quality \leq 10}$$

- **Each buyer can - in each period - submit as many private and public offers as he wishes.** Each offer submitted by a buyer can be accepted at any time during the trading phase.

- **Each buyer can conclude at most one trade in each period.** Once an offer of a buyer has been accepted he will be informed about which seller accepted the offer. As each buyer can conclude only one trade in each period all other offers of the buyer will automatically be cancelled. Also, he cannot submit any further offers.

- Once all 7 buyers have concluded a trade or after 3 minutes have elapsed, the trading phase is over.

- No buyer is obliged to submit offers, and no seller is obliged to accept an offer.

2. Determination of the Actual Product Quality

- Following the trading phase, all sellers who have concluded a trade determine which product quality they supply to their buyers. **The product quality desired by your buyer is not binding for you as a seller.** You can exactly choose the quality desired by your buyer, but also a higher or lower product quality. If you have concluded a trade during a trading phase, the following screen will appear. Here, you have to enter the product quality:

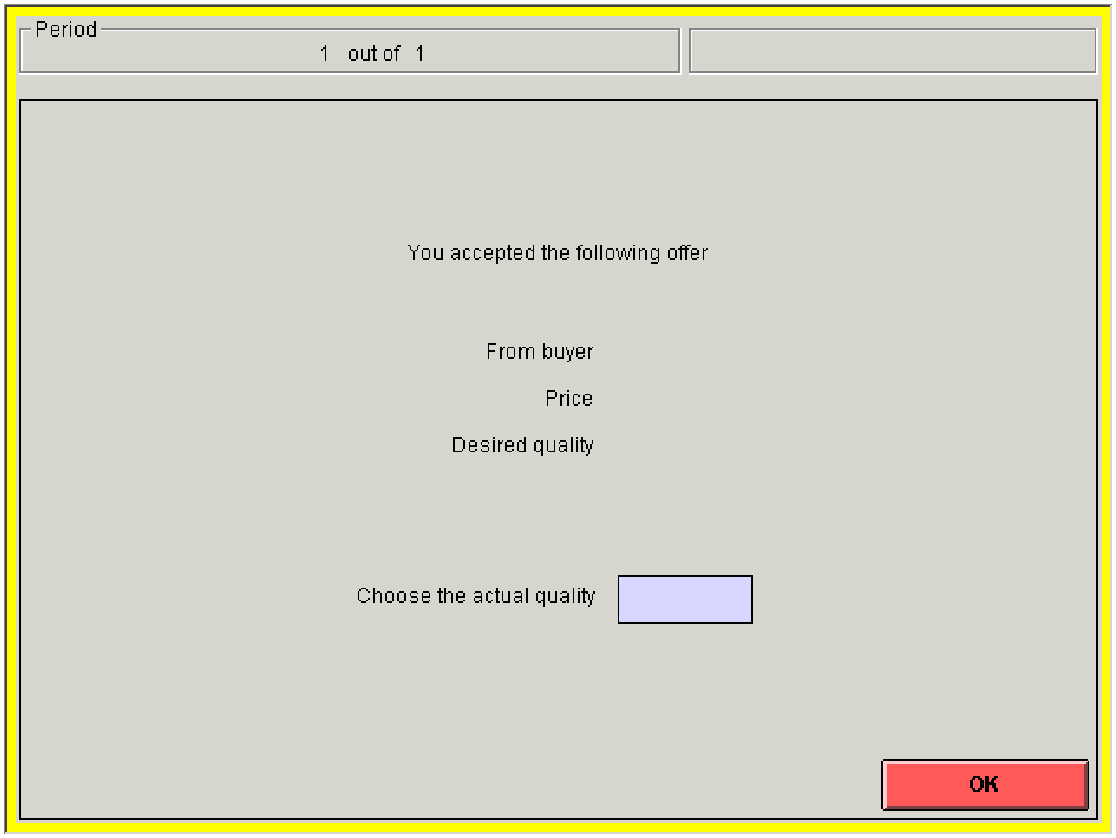

In order to choose the actual product quality, you enter the value for the quality in the field "Choose the actual quality" and press the "OK" button to confirm your choice. As long as you have not pressed "OK" you can alter your choice.

- The product quality you choose must be an integer between 1 and 10:

1 ≤ product quality ≤ 10

How are the Incomes calculated?

<u>Your income</u>

- If you have **not concluded a trade** during a trading phase you receive an income of **5 points** in that period.

- If you have accepted an offer your income depends on the price you accepted and the product quality you choose to deliver. Your income is calculated as follows:

Your income = Price - production costs

- Your production costs are higher, the higher the quality of the product you chose to deliver. The production costs for each product quality are displayed in the table below:

Product quality	1	2	3	4	5	6	7	8	9	10
Production costs	0	1	2	4	6	8	10	12	15	18

- Your income is therefore higher, the lower the product quality. Further, your income is higher, the higher the price.

<u>The income of your buyer:</u>

- If a buyer does not conclude a trade during a trading phase he receives an income of 0 points in that period.
- If one of his offers is accepted, his income depends on the price he offered and the product quality. The income of your buyer will be determined as follows:

Income of your buyer = 10 * product quality - price

- As you can see from the above formula the income of your buyer is higher, the higher the product quality actually supplied by you. At the same time his income is higher, the lower the price he paid for the product.

The income of all buyers and sellers are determined in the same way. **Each buyer can therefore calculate the income of his seller and each seller can calculate the income of his buyer.** Further, each buyer and each seller is informed about the identification number of his trading partner in each period.

Please note that buyers and sellers can incur losses in each period. These losses have to be paid from the initial endowment or from earnings made in other periods.

You will be informed about your income and the income of your buyer on an "**income screen**". On the screen (see below) the following will be displayed:
- Which buyer you traded with
- Which price he offered
- The desired quality of your buyer
- The product quality actually chosen by you
- The income of your buyer in this period
- Your income in this period.

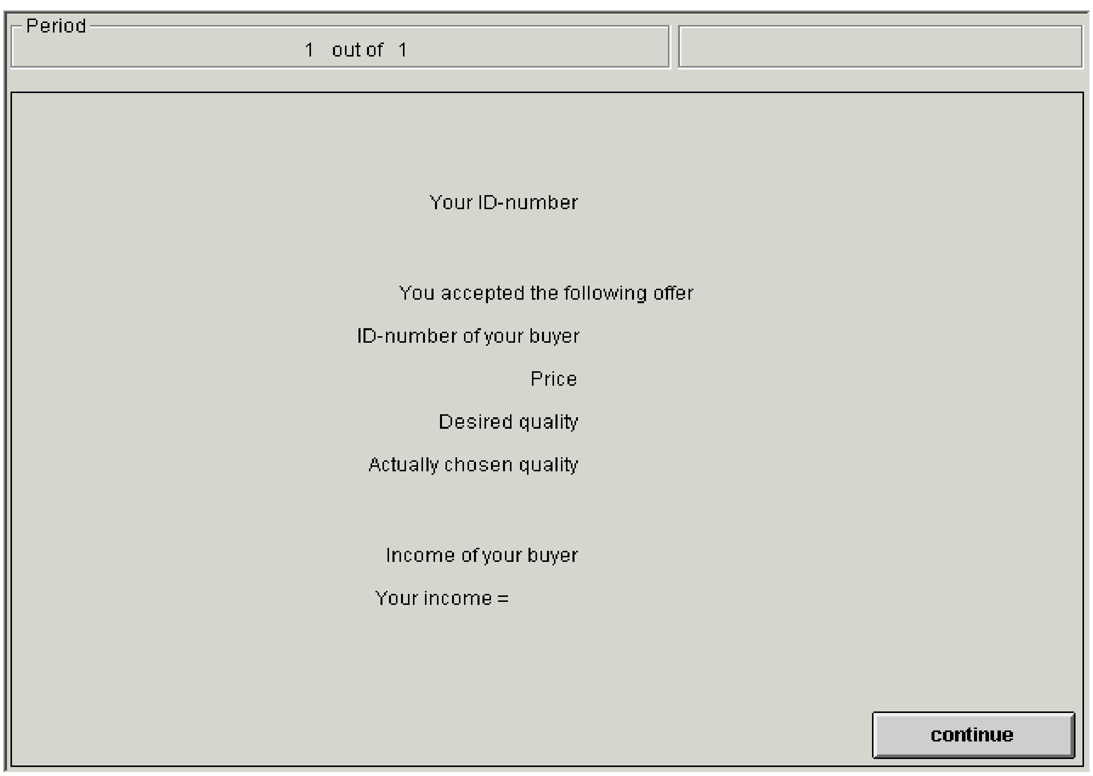

Please enter all the information in the documentation sheet supplied to you. After the income screen has been displayed, the period is over. Thereafter the trading phase of the following period starts. Once you have finished studying the income screen pleas click on the "continue" button.

The buyers also see an income screen, which displays the above information. They see the ID of their trading partner, the price, the desired and the supplied product quality as well as both incomes.

The experiment will not start until all participants are completely familiar with all procedures. In order to secure that this is the case we kindly ask you to solve the exercises below.

In addition we will conduct **2 trials of the trading phase**, so that you can get accustomed to the computer. During the trial phases no money can be earned. After the trial phases we will begin the experiment, which will last for 15 periods.

Control Questionnaire

Please solve the following exercises completely. If you have questions ask the experimenter.

Exercise 1
You did not accept an offer during a trading phase. What is you income in this period?
Your income =

Exercise 2
You accepted an offer with a price of 30 and a desired quality of 9. You supplied an actual quality of 8.
Your income =
Income of your buyer =

Exercise 3
You accepted an offer with a price of 60 and a desired quality of 9. You supplied an actual quality of 4.
Your income =
Income of your buyer =

Exercise 4
You accepted an offer with a price of 40 and a desired quality of 2. You supplied an actual quality of 5.
Your income =
Income of your buyer =

Exercise 5
You accepted an offer with a price of 30 and a desired quality of 6. You supplied an actual quality of 6.
Your income =
Income of your buyer =

Exercise 6
A buyer has made several offers during a trading phase. None of these offers has been accepted by a seller. What is the income of the buyer in this period?
Income of buyer =

If you have finished the exercises we recommend to look again at the exercises and the solutions provided. After this please think about the decisions you want to make during the experiment.

Documentation sheet: Seller

This documentation sheet is meant for your orientation. Please complete the respective row in each period.

period	ID of your buyer	price	desired quality	actual quality	income of your buyer	your income
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						

Bibliography

[1] Akerlof, G. A. (1982), "Labor Contracts as Partial Gift Exchange", *Quarterly Journal of Economics*, 97, 543-569.

[2] Andersen, S., Fountain, J., Harrison, G., annd E. Rustrom (2009), "Estimating Subjective Probabilities", Working Paper 09-01.

[3] Andreoni, J. (1989), "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence", *Journal of Political Economy*, 97 (6), 1447-1458.

[4] Andreoni, J. and D. Bernheim (2009), "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects", *Econometrica*, 77(5), 1607 - 1636.

[5] Andreoni, J. and R. Petrie (2004), "Public Goods Experiments Without Confidentiality: A Glimpse Into Fundraising", *Journal of Public Economics*, **88**, 1605 - 1623.

[6] Ariely, J., A. Bracha and S. Meier (2009), "Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially " *American Economic Review*, 99(1), 545 - 555.

[7] Aumann, R. (1976), "Agreeing to Disagree", *Annals of Statistics*, 4, 1236-39.

[8] Baker, G., M. Gibbs, and B. Holmstrom (1994a), Internal Economics of the Firm: Evidence from Personnel Data", *Quarterly Journal of Economics*, 109, 881-919.

[9] Ballinger, P. and N. Wilcox (1997), "Decisions, error and heterogeneity", *Economic Journal* 107, 1090-1105.

[10] Benabou, R. and Tirole, J. (2003), "Intrinsic and Extrinsic Motivation" *Review of Economic Studies*, 70, 489-520.

[11] Benabou, R. and Tirole, J. (2006), “Incentives and Prosocial Behavior”, *American Economic Review*, 96 (5), 1652-1678.

[12] Bernheim, D. (1994), “A Theory of Conformity”, *Journal of Political Economy*, 102(5), 841 - 877.

[13] Brier, G. (1950), “Verification of Forecasts Expressed In Terms of Probability”, *Monthly Weather Review*, 78(1):13.

[14] Brooks, A. (2006), “Who Really Cares: The Surprising Truth about Compassionate Conservatism” *Basic Books*.

[15] Brown, M., A. Falk and E. Fehr (2004), “Relational Contracts and the Nature of Market Interactions”, *Econometrica, Econometric Society*, vol. 72(3), pages 747-780, 05.

[16] Brown, M., A. Falk and E. Fehr (2009), “Competition and Relational Contracts: The Role of Unemployment as a Disciplinary Device”, *IZA Discussion Papers 3345*, Institute for the Study of Labor (IZA).

[17] Camerer, C. (1989), “An experimental test of several generalized utility theories”, *Journal of Risk and Uncertainty* 2, 61-104.

[18] Camerer, C. (1995), Individual decision making. In: Kagel, J.H. & Roth, A.E. (eds.): *Handbook of Experimental Economics*, 587-703. Princeton, NJ: Princeton University Press.

[19] Camerer, C. (1998), “Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting”. *Journal of Political Economy* 106:3, 457-482.

[20] Camerer, Colin F and Weber, R. Experimental organizational economics. *Handbook of Organizational Economics* (in prep).

[21] Carpenter, J. and Myers, C. (2007), “Why Volunteer? Evidence on the Role of Altruism, Reputation and Incentives” *CEPR Discussion Papers 3021*.

[22] Chen,Y., Mullen, T. and C.H. Chu (2006), “An in-depth analysis of information markets with aggregate uncertainty”, *Electronic Commerce Research*, 6(2), 201-221.

[23] Friedman, D. (1979), "An effective scoring rule for probability distributions", *Discussion Paper 164*, University of California, Los Angeles.

[24] Friedman, D. (1983), "Effective scoring rules for probabilistic forecasts", *Management Science*, 29(4), 447454.

[25] Diamond, D. W. and R. E. Verrecchia (1981), "Information aggregation in a noisy rational expectations economy", *Journal of Financial Economics* 9, 221-235.

[26] DellaVigna, S., J. List and U. Malmendier (2009), "Testing for Altruism and Social Pressure in Charitable Giving", *NBER Working Paper 15629.*

[27] Dickinson, David and Villeval, M. (2008), "Does Monitoring Decrease Work Effort?: The Complementarity Between Agency and Crowding-Out Theories", *Games and Economic Behavior*, 63(1), 56-76.

[28] Ellingsen, T. and M. Johannesson (2008), "Pride and Prejudice: The Human Side of Incentive Theory", *American Economic Review*, 98(3), 990 - 1008.

[29] Ellingsen, T. and M. Johannesson (2009),"Time is not money", *Journal of Economic Behavior & Organization*, 72(1), 96-102.

[30] Falk, A. and Ichino, A. (2006), "Clean Evidence on Peer Effects," *Journal of Labor Economics*, 24(1), 39-58.

[31] Falk, A. and Kosfeld, M. (2006), "The Hidden Costs of Control", *American Economic Review*, 96(5), 1611-1630.

[32] Fehr, E.and S. Gchter (2000), "Fairness and Retaliation: The Economics of Reciprocity", *Journal of Economic Perspectives*, 14, 159-181.

[33] Fehr, E., Kirchsteiger, G. and A. Riedl (1993). "Does Fairness Prevent Market Clearing? An Experimental Investigation", *Quarterly Journal of Economics*, 108, 437-459.

[34] Feigenbaum, J., Fortnow, L., Pennock,D. and R. Sami (2005), "Computation in a distributed information market", *Theoretical Computer Science*, 343(1-2):114-132.

[35] Forsythe,R., Nelson,F., Neumann, G., and J. Wright (1992), "Anatomy of an Experimental Political Stock Market", *American Economic Review*, 82(5):1142-1161.

[36] Finkelstein, M. (2008), "Correlates of Satisfaction in Older Volunteers: A Motivational Perspective", *The International Journal of Volunteer Administration*, 26(5), 6-12.

[37] Fleming, P., Townsend, E., Lowe, K. C., and Ferguson, E. (2007), "Social Desirability Effects on Biotechnology Across the Dimensions of Risk, Ethicality and Naturalness", *Journal of Risk Research*, 10, 989-1003.

[38] Frank, B. (1998), "Good News for Experimenters: Subjects Do Not Care About Your Welfare", *Economic Letters*, 61(2), 171 - 174.

[39] Frank, B. and Schulze, G. (2002), "Deterrence Versus Intrinsic Motivation: Experimental Evidence on the Determinants of Corruptibility", *Econometric Society World Congress 2000 Contributed Papers*.

[40] Freeman, R. B. (1997), "Working for Nothing: The Supply of Volunteer Labor", *Journal of Labor Economics*, 15(1), S140-166.

[41] Frey, B., and Oberholzer-Gee, F. (1997), "The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out", *The American Economic Review*, 87(4), 746-755.

[42] Frey, B. and Jegen, R. (2001), "Motivating Crowding Out Theory", *Journal of Economic Surveys*, 15(5), 589-611.

[43] Geanakoplos, J. and H. M. Polemarchakis (1982), "We Can't Disagree Forever", *Journal of Economic Theory*, 28(1): 192-200.

[44] Gneezy, U. and Rustichini,A. (2000), "Pay Enough-Or Don't Pay At All", *The Quarterly Journal of Economics*, 115(3), 791-810.

[45] Goeree, J.K., Palfrey, T.R., Rogers, B.W., and R.D McKelvey (2007), "Self-Correcting Information Cascades", *Review of Economic Studies*, 74(3), 733-762.

[46] Goeree, J. and Yariv, J. (2007), "Conformity in the Lab", *California Institute of Technology Working Paper*.

[47] Grus and Ledyard (2005), experimental data.

[48] Harbaugh, W. (1998), "The Prestige Motive for Making Charitable Transfers", *American Economic Review*, 88, 277-282.

[49] Healy, P, J. Ledyard, S.Linardi, and R.Lowery (2007), Aggregation Mechanisms and Environmental Complexity", Working Paper.

[50] Hey, J.D. and C. Orme (1994), "Investigating generalisations of expected utility theory using experimental data", *Econometrica* 62, 1291-1326.

[51] Hey, J.D. (2005), "Why We Should Not Be Silent About Noise", *Experimental Economics*, 8(4), 325-345.

[52] Holmstrom, B., and Milgrom, P (1991), "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design", *Journal of Law, Economics and Organization*, 7, Special issue, 24-52.

[53] Holt, C. and S. Laury (2002), "Risk Aversion and Incentive Effects in Lottery Choices", *American Economic Review*, 92(5), 1644-55.

[54] Hung, A.A and C. R. Plott (2001), "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions", *American Economic Review*, 91(5), 1508-1520.

[55] Hung, A.A and J. Dominitz (2009), "Empirical Models of Discrete Choice and Belief Updating in Observational Learning Experiments", *Journal of Economic Behavior and Organization*, 69(2),94-109.

[56] Karlan, D. and List, J. (2007), "Does Price Matter in Charitable Giving? Evidence From a Large-Scale Natural Field Experiment", *American Economic Review*, 97(5), 1774-1793.

[57] Landry, C., Lange, A., List, J., Price, M., and Rupp, N. (2006), "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment", *Quarterly Journal of Economics* 121(2), 747-782.

[58] Ledyard,J., Hanson, R. and T. Ishikida (2009), "An experimental test of combinatorial information markets", *Journal of Economic Behavior and Organization*, 69(2), 182-189.

[59] Levitt, S. and List, J. (2007), "What Do Laboratory Experiments Tell us About the Real World?", *Journal of Economic Perspectives*, forthcoming.

[60] Luce, R.D (1959), *Individual choice behavior: a theoretical analysis.* New York: Wiley.

[61] Lundholm, R. (1991), "What affects the efficiency of a market? Some answers from the laboratory", *The Accounting Review* 66, 486-515.

[62] Martin, R. and Randal, J. (2008), "How is Donation Behaviour Affected by the Donations of Others?", *Journal of Economic Behavior and Organization*, 67(1), 228-238.

[63] McKelvey, R. and T.R. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior*, 10, 6-35.

[64] McKelvey, R. D. and T. R. Palfrey (1998), "Quantal Response Equilibrium for Extensive Form Games", *Experimental Economics*, 1(1), 9-41.

[65] McKelvey, R and T.Page (1986), "Common Knowledge, Consensus and Aggregate Information", *Econometrica*, 54(1), 109-127.

[66] McKelvey, R. and T. Page (1990), "Public and Private Information: an Experimental Study of Information Pooling", *Econometrica*, 58(6), 1321-1339.

[67] Mellström, C. and Johannesson, M. (2008), "Crowding Out in Blood Donation: Was Titmuss Right?", *Journal of the European Economic Association*, 6(4), 845-863.

[68] Nielsen, L., A. Brandenberger, J. Geanakoplos, R. McKelvey and T. Page (1990), "Common Knowledge of an Aggregate of Expectations", *Econometrica*, 58, 1235-1239.

[69] Offerman, T., Sonnemans J., van de Kuilen G. and P.P. Wakker (2009) "A Truth-Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes", *Review of Economic Studies*, 76, 1461-1489.

[70] Palfrey, T., and S.W. Wang (2009), "On eliciting beliefs in strategic games", *Journal of Economic Behavior & Organization*, 71(2), 98-109.

[71] Paulhus, D. L. (1991), "Measurement and Control of Response Bias", In J. P. Robinson, P. Shaver, and L. S. Wrightsman (Eds.), *Measures of Personality and Social Psychological Attitudes*, San Diego: Academic Press, 17-59.

[72] Plott, C. and S. Sunder (1988), "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets", *Econometrica*, 56(5), 1085-1118.

[73] Plott, C.R. (2000), "Markets as information gathering tools", *Southern Economic Journal*, 67(1), 215.

[74] Plott, C.R. and K. Roust (2006), "Design and Testing of Information Aggregation Mechanisms: A Two-Stage Parimutuel IAM", Social Science Working Paper.

[75] Rabin, M.(1993), "Incooperating fairness into game theory and economics", *American Econonomics Review*, 83, 1281-1302.

[76] Rege, M. and K. Telle (2004), "The Impact of Social Approval and Framing on Cooperation in Public Good Situations", *Journal of Public Economics*, 88, 1625-1644.

[77] Rigdon, M (2002), "Efficiency wages in an experimental labor market", *Proceedings of the National Academy of Sciences*, 99(20), 13348-13351.

[78] Rotemberg, J. (1994), "Human Relations in the Workplace", *The Journal of Political Economy*, 102(4), 684-717.

[79] Schady, N. (2001), "Who Participates : the Supply of Volunteer Labor and the Distribution of Government Programs in Rural Peru", *The World Bank Policy Research Working Paper Series Number 2671.*

[80] Seabright, P. (2004), "Continuous Preferences Can Cause Discontinuous Choices: an Application to the Impact of Incentives on Altruism", *IDEI Working Paper Number 257.*

[81] Selten, S (1998), "Axiomatic Characterization of the Quadratic Scoring Rule", *Experimental Economics*, 1(1), 43-61.

[82] Shang, J. and R. Croson (2009), "The Impact of Social Comparisons on Nonprofit Fundraising", *Research in Experimental Economics*, forthcoming.

[83] Soetevent, A. (2005), "Anonymity in Giving in a Natural Context. A Field Experiment in 30 Churches", *Journal of Public Economics*, 89, 2301-2323.

[84] Starmer, C. and Sugden, R. (1989), "Probability and juxtaposition effects: An experimental investigation of the common ratio effect", *Journal of Risk and Uncertainty*, 2, 159-178.

[85] Sunder, S. (1995), Experimental Asset Markets: a Survey, in J. Kagel and A. E. Roth (eds) *Handbook of Experimental Economics*, Princeton, NJ. 13.

[86] Tadelis, S. (2007), "The Power of Shame and the Rationality of Trust", unpublished.

[87] Toda, M. (1963), "Measurement of Subjective Probability Distribution", Report 3, State College, Pennsylvania, Institute for Research, Division of Mathematical Psychology.

[88] Weber, R. A. (2001), "Behavior and learning in the dirty faces game", *Experimental Economics*, 4(3), 229-242.

[89] Winkler, R.L. (1969), "Scoring Rules and the Evaluation of Probability Assessors", *Journal of the American Statistical Association* 64, 10731078.

[90] Zimper, A. (2009), "Half empty, half full and why we can agree to disagree forever", *Journal of Economic Behavior and Organization*, 71(4), 283-299.

[91] Zizzo, D. (2009), "Experimenter Demand Effects in Economic Experiments", *Experimental Economics*, forthcoming.