ABSTRACT

Title of Dissertation: ESSAYS IN BEHAVIORAL AND EXPERIMENTAL ECONOMICS

Jonathan Eliot Alevy, Doctor of Philosophy, 2006

Dissertation directed by: Professor Robert G. Chambers
Department of Agricultural and Resource Economics
Professor John A. List
Department of Economics, University of Chicago

Behavioral economics aims to provide more realistic psychological foundations for economic models. Experimental methods can contribute to this effort by providing the ability to identify causal processes and motivations that can be confounded in field settings. The essays in this dissertation examine three critical issues in behavioral economics using lab and field experiments. The first two essays examine two core elements of economic rationality; expected utility theory and Bayesian updating. The essays consider, respectively, ambiguity, and information cascades, in environments in which limitations of the theories can be studied. The third essay examines a contracting game in which other-regarding preferences are explicitly considered.

Decision making under ambiguity has been of interest to economists since the 1920's (Knight (1921), Keynes (1921)). It has received renewed attention due to the work

of Ellsberg (1961). In the first essay I examine the stability of ambiguity attitudes using a within subject design across individual choice and market environments. The evidence favors stability, with attitudes elicited from individuals strongly correlated with trading decisions in asset markets. The comparative ignorance hypothesis of Fox and Tversky (1995) developed for individual choice is also supported in the market setting shedding light on the causes of ambiguity aversion.

Previous empirical studies of information cascades have used either naturally occurring data or laboratory experiments. In the second essay attractive elements of each line of research are combined by observing market professionals from the Chicago Board of Trade (CBOT) in a controlled environment. Analysis of over 1500 decisions suggests that CBOT professionals behave differently than a student control group. Professionals are better able to discern the quality of public signals and their decisions are not affected by the domain of earnings. These results have important implications for market efficiency.

The contracting game studies both one and two principal settings. With one principal, behavior is consistent with a reputational model in which principals are successful in structuring contracts to insure against defections by agents imitating inequity-averse behavior. The complexity of the two principal setting creates more difficulties, but there is evidence that reciprocity between principals partially mitigates the adverse payoff consequences.

ESSAYS IN BEHAVIORAL AND EXPERIMENTAL ECONOMICS

by

Jonathan Eliot Alevy

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2006

Advisory Committee:
Professor Robert G. Chambers, Co-Chair
Professor John A. List, Co-Chair
Professor Mark G. Duggan
Professor Bruce Gardner
Professor Tigran A. Melkonyan

Acknowledgements

The essays in this dissertation would not have seen the light of day without the help of a number of individuals and institutions. I thank the co-chairs of my committee, Robert Chambers and John List, for the insights, advice and opportunities that they provided. They have been examples of professionalism. I thank Tigran Melkonyan for the detailed comments on the first and third essays which led to substantial improvements. I am also appreciative of the input from committee members Bruce Gardner and Mark Duggan whose careful reading of the manuscript yielded many suggestions for clarifications and improvements.

I thank the International Foundation for Research in Experimental Economics for financial support. Their generosity allowed me to visit the Economic Science Laboratory at the University of Arizona where I became exposed to experimental methods in economics and where I conducted the research that resulted in the first essay. While in Arizona, conversations with Jim Cox, Vernon Smith, Steve Rassenti and Kevin McCabe proved especially helpful. Thanks to Rich Kiser for modifying the asset market software.

Thanks also to the taxpayers of Maryland who provided research funds for the remaining essays. For the work on informational cascades, I am grateful to my coauthors, John List and Michael Haigh, for their insights and efforts, as well as to seminar attendees who provided many useful comments. Lori Lynch was instrumental in the genesis of the final essay as was support from the Maryland Agricultural Experiment Station.

I have had many fruitful discussions about economic research and the dissertation process with fellow students. Especially important have been conversations with Dan

Gilligan, Tulika Narayan and Michael Price. Family members have also supported this work. I thank my brothers, Michael and Adam, for their encouragement and interest. My parents, Alice Learman and Daniel Alevy, have each in their way provided enduring support, and due to their professional interests in psychology, keen observations of issues relevant to behavioral economics.

Finally, I would like to express my deepest appreciation to Alejandra Palma, whose patience, humor, encouragement and love, along with her sharp economic insights, made this process more enjoyable.

Table of Contents

List of Tables

List of Figures

1 Ambiguity in Individual Choice and Market Environments

1.1 Introduction

Decision making under ambiguity has been of interest to economists at least since the 1920's, when Knight (1921) and Keynes (1921) raised issues similar to those that distinguish risk from ambiguity in the current study. Knight distinguished measurable uncertainty (risk) from unmeasurable uncertainty (ambiguity) while Keynes argued that equal probabilities can have different impacts on behavior as a result of the weight of the evidence through which they were derived.[1] The inadequacy of subjective expected utility theory (SEU) to account for vague probabilities, as well as its empirical relevance for economic decision making, was brought into focus by the work of Ellsberg (1961).[2]

In this paper we conduct experiments to examine the behavior of individuals under both sure (risk) and unsure (ambiguous) probabilities. We consider, within subjects, responses to risk and ambiguity in both individual choice and market environments and find that attitudes to ambiguity are conserved across these institutions. This result is interesting in its contrast to that for risk preferences which have proven malleable in laboratory settings, and in light of the view of Epstein (1999) and others that ambiguity is both a more prevalent and more fundamental aspect of economic environments than risk.[3]

---

[1] Throughout the paper we consider risk (ambiguity) to be associated with the ability (inability) to formulate subjective probabilities. Thus, both risk and ambiguity belong under the umbrella of uncertainty.

[2] Savage (1954) also took note of the issue, remarking that "there seem to be some probability relations about which we seem relatively 'sure' as compared with others…The notion of sure and unsure introduced here is vague, and my complaint is precisely that neither the theory of personal probability, as it is developed in this book, nor any device known to me renders the notion less vague."

[3] Mukerji and Tallon (2004) review a variety of economic applications that explicitly incorporate ambiguity.

The paper makes two additional contributions, first by extending the investigation of how ambiguity responses are triggered in markets. Sarin and Weber (1993) demonstrate that ambiguous probabilities have a significant effect on asset prices when trade of risky and ambiguous assets occurs simultaneously. The current design examines sequential exposure to risk and ambiguity for which Sarin and Weber's results were inconclusive. This aspect of the inquiry investigates how information conditions affect ambiguity attitudes generating further tests of Fox and Tversky's (1995) hypothesis of comparative ignorance, originally investigated in individual choice settings.

A final contribution of the paper is to offer a preliminary assessment of the empirical relevance of two alternative approaches to modeling ambiguity. These approaches differ in their conception of what constitutes an ambiguity-neutral act and thus in their definition of ambiguity aversion. Ghirardato and Marinacci (2002) assume that constant acts, in which the same outcome is achieved in all states are ambiguity neutral. Since constant acts also are an appropriate benchmark for risk neutrality (Yaari 1969) the authors acknowledge that their modeling choice is restrictive. Their strategy implies that probabilistically sophisticated preferences that are inconsistent with SEU, such as the probabilistically risk averse behavior associated with the Allais paradox, are not ambiguity neutral. Epstein (1999) develops an alternative approach in which probabilistically sophisticated preferences not just those consistent with SEU are ambiguity neutral.[4] Thus, the two approaches differ with respect to whether preferences

[4] Probabilistic sophistication implies that beliefs over relative likelihoods for subjective events are consistent with probability theory. Machina and Schmeidler (1992) demonstrate that probabilistically sophisticated beliefs are possible outside the framework of SEU, which has beeen shown to be inadequate due to violations of the "sure-thing principle" as seen for example in the Allais paradox.

that are probabilistically risk-averse fall within the ambit of ambiguity. Section 2 and Appendix 1.3 provide additional detail on these alternative approaches.

Our empirical methods differ according to our purposes. For the measurement of attitudes across the choice and market institutions, we use relatively simple descriptors. In the individual choice setting we use a nonparametric measure of ambiguity aversion that assumes that ambiguity affects utility directly and by comparing responses to lottery choices, within subject, under both risk and ambiguity (Smith 1969). Simple counts of asset accumulation by subject and across risky and ambiguous asset types along with bidding behavior and asset prices are the focus of the analysis in the market setting.

To investigate the empirical relevance of the alternative theoretical conceptions of ambiguity, maximum likelihood techniques are used in a two-stage procedure. First, baseline measures of risk attitudes are estimated from the subset of risky questions, under the alternative assumptions regarding linearity in probabilities assumed by the SEU and rank dependent expected utility (RDEU) models.[5] Second, ambiguity attitudes measured as deviations from additive probability measures are estimated for each of the theoretical benchmarks.

The paper proceeds as follows: Section 2 clarifies what we mean by ambiguity in the experimental environments, by discussing the links between our operationalization of ambiguity and measurement methods and alternative conceptions of ambiguity in the theoretical literature. Section 2 also provides an overview of experimental findings that

[5] RDEU accommodates probabilistically sophisticated preferences that are not SEU and thus provide our baseline for ambiguity neutrality in Epstein's model. In RDEU individuals have well-formed probabilities, but these probabilities are not used in decision-making. Instead transformations of the probabilities that also take into account the rank order of outcomes are used as decision weights. Quiggin (1982) introduced the decumulative transformation that insures that RDEU preferences satisfy first order stochastic dominance. Decumulative weights applied to ambiguity yields the CEU model (Schmeidler 1989) which is discussed in greater detail below.

shed light on the underlying causes of behavior that is sensitive to ambiguity. Section 3 presents the experimental design, and Section 4 provides details of the methods and results in the individual choice setting for both the parametric and nonparametric approaches to the measurement of ambiguity aversion. Section 5 reports on tests of the comparative ignorance hypothesis in the asset market setting. Section 6 examines the conservation of ambiguity attitudes across the choice and market environments. Section 7 concludes.

1.2 Ambiguity Defined

Ambiguity has been defined as "a quality depending on the amount, type, and 'unanimity' of information, and giving rise to one's degree of 'confidence' in an estimate of relative likelihoods" (Ellsberg (1961) or more concisely, as "known-to-be-missing information" (Frisch and Baron (1988), see also Camerer (1999) and Appendix 1.3 below). These verbal definitions have been supplemented by axiomatic approaches that modify the SEU frameworks developed by Savage (1954) or Anscombe and Aumann (1963).

One way to reconcile Ellsberg type behavior theoretically assumes that ambiguity has an effect directly on utility (Smith 1969). This approach, which is compatible with our nonparametric approach to measuring ambiguity attitudes, assumes that there are utility implications from the choice process as well as from the monetary outcomes of the draws from the Ellsberg urn. Under this assumption the response to ambiguity can be treated as a fixed effect in the utility elicitation.[6]

[6] Smith conjectures that it would matter little whether decision makers were completely in the dark regarding the contents of the Ellsberg urn or were informed of a second order probability. While we do not test this conjecture, our operationalization of ambiguity does provide the second order probability in a

An alternative and at present, more common approach to modeling ambiguity assumes that utility is state independent, but adopts the view, common also to RDEU, that beliefs and tastes are not completely independent. As a result attitudes towards uncertainty are captured jointly through decision weights and the utility function (Schmeidler (1989); see also Diecidue and Wakker (2001)). Consider a finite set of states, *S*, and outcomes, *X*, with the true state unknown. Events, *E*, are disjoint subsets of states, with acts of the form $(E_1, x_1; ...; E_n, x_n)$ mapping events to outcomes so that x_i results if the true state is in E_i. The dependence between beliefs and tastes arises by assuming that decision weights depend not only on the likelihood of events, but also on the relative magnitude of the outcomes. Assume $x_1 > ... > x_n$. The decision weight associated with the outcomes $x_1, ..., x_i$ is given by $\sum_{k=1}^{i} \pi_i = v(E_1 \cup ... \cup E_i)$, and that of x_i by $\pi_i = v(E_1 \cup ... \cup E_i) - v(E_1 \cup ... \cup E_{i-1})$, with $v(\varnothing) = 0$, and $v(S) = 1$, and v monotonic. In the Choquet Expected Utility (CEU) model of Schemeidler (1989), the independence axiom is relaxed to apply only to comonotonic events. As a result ambiguity aversion is captured by convexity of *v*, which implies that for disjoint events $v(E_1 \cup E_2) \geq v(E_1) + v(E_2)$. While a convex capacity can rationalize behavior associated with the Ellsberg paradox, as demonstrated in Appendix 1.3, a number of counterexamples have been developed that show convexity is neither necessary nor sufficient for ambiguity aversion more generally. Alternative definitions based on

---

transparent way so as to reduce concerns that subjects may be responding to the potential for experimenter deception. This implementation is consistent with the notion that ambiguity can arise through weighting of multiple probability distributions (Gilboa and Schmeidler (1989), Camerer (1999)). As a result we conjecture that the effects observed may be considered a lower bound on ambiguity responses when the underlying distribution is not provided.

comparative foundations are presented in Appendix 1.3 and examined empirically in Section 1.4.2.

With regard to the causes of ambiguity attitudes, individual choice experiments have demonstrated that in situations characterized by lack of competence about a judgment task, an ambiguity effect is likely (Heath and Tversky (1991). However, when individuals feel competent in a situation in which probabilities are vague they are likely to prefer this vagueness, contrary to the predictions of ambiguity aversion. Fox and Tversky (1995) extend investigation of this result by examining the conditions under which competence may arise. They find that missing information is made more salient when a simultaneous comparison with more complete information is possible. In contrast, in between subject tests, where valuations of ambiguous lotteries and events are made in isolation Fox and Tversky find no evidence of ambiguity aversion.[7]

Wakker (2000) notes that "Fox and Tversky's finding seems to place the Ellsberg paradox in an entirely new light," arguing that the importance of ambiguity aversion is in doubt if the results are "merely…a contrast effect." While there is mixed evidence as to whether ambiguity aversion is solely a contrast effect, it clearly is important behaviorally.[8] Similar contrast effects have been shown to cause preference reversals in other contexts (Hsee et al. (1999), List (2002), Alevy, List, and Adamowicz (2003)). Further, understanding the types of contrasts that may generate an ambiguity response is of some interest in practical applications. In the asset market portion of this experiment

[7] Chow and Sarin (2001) present contrasting evidence on the comparative ignorance hypothesis in individual choice settings. They find that while ambiguity responses are smaller, between subjects, the differences remain statistically significant.

[8] The question raised by the context results is whether paradoxical behavior in the settings investigated by Ellsberg should be explained by underlying preferences or by differences in the ability to evaluate information when the choices are presented jointly or in isolation.

we shed additional light on the importance of contrast effects, by examining sequential exposure to risky and ambiguous assets.

1.3 Experimental Protocol

An experimental session contained two protocols. In each session subjects completed a computerized questionnaire that elicited preferences over pairs of lotteries, some with known and some with unknown probabilities. Following this task, subjects became traders in a double auction asset market for ten rounds. After each round of trade, a random draw determined the underlying state on which the asset value depended. In five rounds of each session the probabilities of each of the three possible states were known and in five rounds two of the three state probabilities were unknown. The underlying probabilities changed with each round. The sessions were distinguished by the order of presentation of the risky and ambiguous assets. Session 1 consisted of six subjects who traded risky assets for the first five rounds and ambiguous assets in the last five. Session 2 consisted of seven subjects who traded first ambiguous and then, risky assets also for five rounds each. We denote these sessions “RA” and “AR” consistent with the order of exposure to the risky and ambiguous assets.

Subjects were paid for the individual choice portion, through the random selection of an elicitation question and then by resolving the uncertainty with respect to the preferred lottery for that question. The probability distribution for ambiguous lotteries for payoff purposes was determined by drawing a ball numbered from 1 to 100 from a bingo cage. The cage was spun repeatedly before releasing a ball in order to approximate a uniform distribution. This procedure was conducted after the double auction was completed but subjects were informed of the process before making their choices.[9] Payments for the elicitation section ranged between \$0 and \$15 with an average payment of \$3.75. Subjects received an average of \$27.33 for the market segment, and an additional \$5 for timely arrival. Total earnings ranged from \$5.00 to \$101.90 with an average of \$36.08 for sessions that were about 2 hours and 30 minutes long. Subjects were recruited from undergraduate economics classes at the University of Arizona.

1.4 Individual Choice Experiment: Methods and Results

1.4.1 Elicitation Method

The elicitation procedure implemented in the lab was a variant of the lottery tradeoff method introduced by Wakker (1994). Rather than eliciting certainty or probability equivalents repeatedly for single lotteries, the tradeoff method asks the subject to compare two lotteries directly.[10] The sequence of lottery comparisons a subject faces is generated endogenously based on their initial responses. Lotteries under risk are of the form $(x, p; y)$, where p is the probability of outcome x and $1-p$ the probability of outcome y. The subject is asked to change one of the outcomes, the "choice outcome" denoted C_1 below, so that they are indifferent between the two lotteries.

[9] To mitigate concern with deception in the resolution of ambiguous probabilities, one of the subjects was selected at random to come to the front of the room at the time of the draw and confirm for the others that the distribution of the balls was as described in the instructions.

[10] Wakker and Deneffe (1996) compare the trade-off method to probability-equivalent and certainty-equivalent methods and show that the trade-off method yields measures of utility that are robust to probability distortions.

In the "inward"[11] tradeoff method, the subject is presented with two lotteries $(C_1, p; R)$ and $(x_0, p; r)$, with $R > r$ the reference outcomes, and x_0 the maximal payment fixed by experimental design. The subject is asked to replace C_1 with a value, x_1, that leaves them indifferent between the two lotteries. Once $(x_1, p; R) \sim (x_0, p; r)$ is elicited, the process is repeated with the subject presented with two new lotteries $(C_2, p; R)$ and $(x_1, p; r)$. The questionnaire is endogenous since x_1, the value elicited from the first lottery comparison, becomes a potential payoff in the next question. Expressing the elicited indifference relations from the first two questions in an expected utility formulation yields the following two equalities:

$$pU(x_o) + (1-p)U(r) = pU(x_1) + (1-p)U(R)$$

$$pU(x_1) + (1-p)U(r) = pU(x_2) + (1-p)U(R).$$

These equations imply that $U(x_0) - U(x_1) = U(x_1) - U(x_2)$ with the result independent of p, or more generally, of a decision weight π which may be either a probability weight $w(p)$ under risk, or a capacity $v(E)$ under uncertainty. Continued elicitation yields utility differences of equal magnitude which are normalized over the unit interval by setting $U(x_0) = 1$ and $U(r) = 0$.

A pilot session with hypothetical payments was conducted to test the utility elicitation software on four subjects. Three of the four subjects consistently violated the minimal rationality constraint and exhibited utility curves with substantial downward sloping portions. Since $R > r$ by design, the violation of dominance occurred when

[11] The inward (outward) method sets a maximum (minimum) on the potential payoff to the subjects if their choices are consistent with first order stochastic dominance. For this reason the inward method was used in

indifferences $(x_o, p; r) \sim (x_1, p; R)$ were indicated with $x_1 > x_o$. This result demonstrates the difficulties introduced by an elicitation procedure that uses an endogenous questionnaire. If subjects recognize that their responses will alter the subsequent questions and thus their potential payoffs, there is no incentive for truthful revelation. Since subjects in the later sessions were to be motivated by salient payments, it was necessary to modify the tradeoff method so that it could not be manipulated for monetary gain.[12]

The modified tradeoff method used to generate the results reported below, asks the subjects only for their preference over lotteries and not for a value that makes them indifferent between the two. All questions are fixed in advance and a menu of four or five lottery pairs is presented with the C value systematically increasing so that a "switching point" in which preference changed from one lottery to the other was observed. Indifference between lottery pairs could also be expressed directly. A series of eight of these lists, four each for risk and ambiguity, was used to elicit utility measures. The first list is shown in Table 1.1 with the parameters $x_o = 15,\ p = .33,\ r = 0,\ R = 4$, and $C_1^1 = 6, C_1^2 = 8, C_1^3 = 10,$ and $C_1^4 = 11$. This series of four questions substituted for the single initial comparison proposed by Wakker (1994) and the elicitation method required that the subjects express a preference, or indifference, on alternative a) before observing

the experimental sessions. Inward and outward tradeoff methods are discussed more fully by Fennema and van Assen (1999).

[12] An alternative approach to the modification described below is to constrain responses to satisfy stochastic dominance. See for example Tversky and Kahneman (1992).

question b) and so on.[13] Appendix 1.2 contains a sample screen and the full list of questions.

Table 1.1: Sample menu of lottery choices

	Lottery A	Lottery B
1	(15,.33;0)	(6,.33;4)
2	(15,.33;0)	(8,.33;4)
3	(15,.33;0)	(10,.33;4)
4	(15,.33;0)	(11,.33;4)

Lotteries are of the form (x,p;y) where p is the probability of receiving x and 1-p the probability of receiving y. The experimental protocol asks for a preference between lottery A and lottery B in each of the four rows

The responses to the first menu yielded C_1^*, either directly if indifference was indicated, or as $\left(C_1^j + C_1^{j+1}\right)/2$ with the switch in lottery choice occurring from menu point *j*, to *j*+1. For the next series of lotteries the subjects again faced a predetermined set of questions so that in general $x_1 \neq C_1^*$ as would be the case in the endogenous tradeoff method. As an example, suppose a subject indicated indifference for the lottery pair 2 in Table 1.1. With the endogenous tradeoff method this would imply that $x_1 = 8$ and a C_2 would be elicited so that $(8,.33;0) \sim (C_2,.33;4)$. With the modified method we have learned that $U(15) - U(8) = 2U(4)$.

In the questionnaire implemented in the lab, $x_1 = 12$ by design and the second set of lottery pairs for risky utility may yield, for example with $C_2^* = 5$, the result that

[13] The method is similar to that of Holt and Laury (2002) which has become a popular method for measuring or controlling for risk attitudes in the lab and in the field. The major differences between our procedure and theirs is that we present multiple (though shorter) lists and they maintain equal payoffs and alter probabilities within the list, while we alter the payoff structure holding probabilities constant.

$U(12) - U(5) = 2U(4)$. The modified tradeoff method therefore yields a series of equal utility differences, as in Wakker's original formulation, but the endpoints of the segments need not coincide as they would with the endogenous questionnaire. To generate a piecewise linear utility function a linear program was implemented that minimized the utility differences of the end point of a segment with its interior point on the next segment, subject to the constraint of equal utility differences for each segment. Details of the technique and the calculation of *Av* are in Appendix 1.4.

The tradeoff method was conducted for each subject under both risk (R) and ambiguity (A) in order to generate state dependent measures of utility. For the risk elicitation the probability of each lottery outcome was made explicit on the questionnaire. For the ambiguity elicitation the outcomes of one of the lotteries was associated with its probabilities and the other the probabilities were replaced by a "?" on the form. The nonparametric measures of ambiguity aversion, *Av,* derived from this procedure were calculated for each subject by finding the area under the utility curves in both the risky and ambiguous state. Subtracting the value of the ambiguous utility from that of the risky utility yielded *Av*, so that a positive measure implies ambiguity aversion and a negative *Av*, ambiguity seeking.

Alternative measures of ambiguity attitudes are generated parametrically, using maximum likelihood techniques, in a two-stage process. In stage one, cardinal utility is estimated using only the risk-based lottery pairs. Separate measures of risk attitudes are obtained for SEU and RDEU preference functionals.[14] In the second stage, weighting

[14] The two-step procedure is motivated by the models of Epstein (1999) and Ghirardhato and Marinacci (2001) who develop alternative conceptions of ambiguity neutrality based on probabilistically sophisticated and SEU preferences (see also Appendix 1.3). The estimation results do not constitute a test of these

functions derived from the ambiguous lists are estimated, using the predicted utilities from stage one. The functional form is used to estimate the expected utility of the lotteries in stage one satisfies constant relative risk aversion (CRRA), and is given by

$$U(X) = \pi_1 \frac{x_1^{1-r}}{1-r} + \pi_2 \frac{x_2^{1-r}}{1-r},$$

where $X = (x_1 \quad x_2)$ are the monetary payoffs with $x_1 > x_2$. The parameter r is a measure of risk attitude with r = 0 indicating risk neutrality and r> (<) 0 characterizing risk aversion (risk seeking). The decision weights for SEU preferences are $\pi_i = p_i$, $i = 1,2$, and for RDEU $\pi_1 = w(p_1)$ and $\pi_2 = w(p_1 + p_2) - w(p_1) = 1 - w(p_1)$, with $w(0) = 0$, $w(1) = 1$. The functional form used to estimate the RDEU decision weights is a one-parameter model introduced by Tversky and Kahneman (1992), which can reconcile optimistic and pessimistic distortions over the range of outcomes and is given by $w(p) = \frac{p^\gamma}{\left[p^\gamma + (1-p)^\gamma\right]^{1/\gamma}}$. Maximum likelihood estimates of r and γ are derived from a probit model that estimates the probability of individual i choosing lottery A in question t. The probability is given by $\Pr_{it}(A \mid \beta) = \Pr\left(U_i(A^t) - U_i(B^t) > \varepsilon_{it}\right)$ where $\beta = r$ $(\beta = [r, \gamma])$ for the SEU (RDEU) models.

Let $\left[U_i(A^t) - U_i(B^t)\right] = \Delta_{it}$, $\Phi(\bullet)$ be the cumulative normal distribution with $\varepsilon_i \sim N(0, \sigma^2)$, and $y_{it} = 1$ when lottery A is chosen and zero otherwise. The likelihood function is given by

models, however, since the elicitation procedure made use of objective probabilities while the models are developed for situations in which objective probabilities are not available.

$$L(\beta \mid Y, X \mid) = \prod_{y_i} \left[\Phi(\Delta_i)^{y_i} \right] \left[\Phi(\Delta_i)^{1-y_i} \right]$$

where Y is the vector of responses and X a matrix containing the lottery payoffs and probabilities.

1.4.2 Elicitation Results

In this section we report on both the parametric and non-parametric approaches to measuring ambiguity aversion. The non-parametric results reveal that both ambiguity aversion and ambiguity seeking existed in the student population. Of the 13 subjects who used the modified trade-off method, 7 exhibited ambiguity aversion and 6 ambiguity seeking. The mean score was not different from zero at the 5% level. Figure 1.1 and Table 1.2 present the scores for subjects participating in the auction market. A positive number reflects an aversion to ambiguity. The ordinal and cardinal scores are used below to assess the consistency of ambiguity attitudes in choice and market environments.

Table 1.2: Non-parametric Ambiguity Aversion Scores (Av)

ID	1	2	3	4	5	6	7
Session 1	1.130	1.015	0.455	0.258	-0.810	-1.691	
Session 2	0.913	0.568	0.375	-0.075	-0.287	-0.388	-0.918

Ambiguity scores for each individual are measured as total utility in risky state minus total utility in ambiguous state.

Figure 1.1: Ambiguity Scores Derived from the Lottery Tradeoff Elicitation

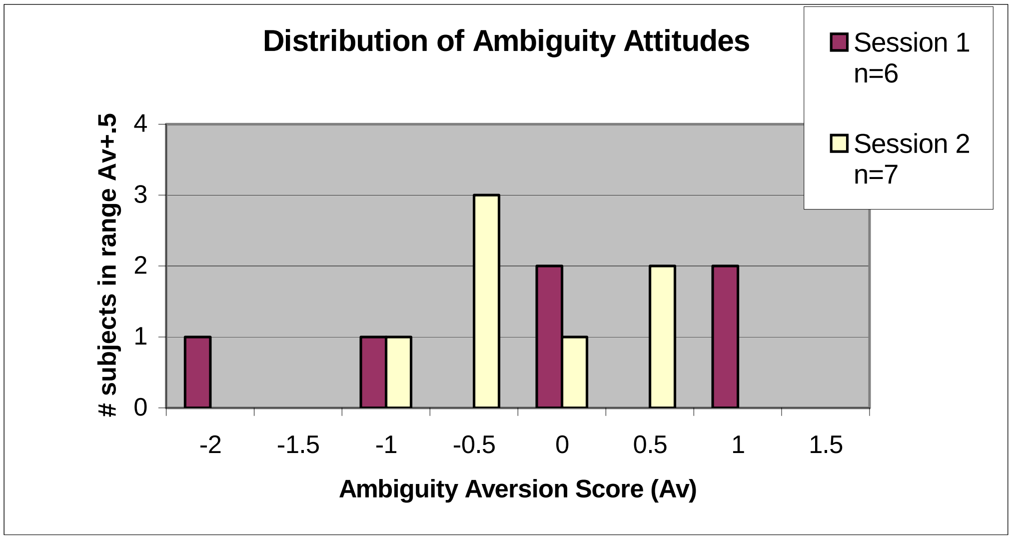

Ambiguity scores for each individual are measured as total utility in risky state minus total utility in ambiguous state.

The parametric results are reported in Table 1.3 pooled over all respondents. In the pooled data there is evidence of significant risk aversion in the first stage estimates, conducted over the risky lottery choices, with the *r* parameter = .720 (.625) for the SEU (RDEU) specification. There is evidence of probabilistic risk aversion in the RDEU model, since RDEU reduces to SEU with $\gamma = 1$. The estimated $\gamma = .367$ is significantly less than this value (p = .001). Stage two estimates under ambiguity, yielded subadditive decision weights consistent with ambiguity aversion in the aggregate, with additional weight placed on the low-valued outcome in the RDEU model. The Aikake Information Criteria (AIC) suggests that the RDEU models are slightly more informative.

Table 1.3: Pooled Parametric estimates of risk and ambiguity attitudes

	Stage one n = 252		Stage two n = 351	
Model	SEU	RDEU	SEU	RDEU
r	.720 (0.000)	.625 (0.001)		
γ		.367 (0.001)		
π_1			.434 (0.023)	.601 (0.039)
π_2			.384 (0.030)	.202 (0.427)
AIC	259.82	259.35	433.09	431.84
Log likelihood	-128.91	-123.28	-214.55	-213.92

r is the CRRA risk parameter, and γ the RDEU probabilistic risk aversion parameter. π_1 (π_2) is the estimated weight on the low (high) outcome under ambiguity in stage two. AIC is the Akaike information criteria which is calculated as -2ln(likelihood)+2(number of parameters). P-values are in parentheses below the parameter estimates. For γ the null hypothesis is γ =1. For all other parameters the null hypothesis is that the value is zero.

1.5 Asset Market

1.5.1 Asset Market Design and Hypotheses

The asset market implemented in this study adopted many features of the market used in studies of the informational efficiency of markets under risk (Plott and Sunder (1988); Sunder (1992)). It is also quite similar to the design used by Sarin and Weber (1993). Key features of these markets are the following:

1) Subjects are in the role of traders, able to either buy or sell assets.

2) Traders are endowed with a single asset and with cash in each round.

3) The asset expires at the end of each round and pays a state contingent dividend.

4) Traders pay a fixed fee of their endowed cash in each round so that each trader's profits are the sum of dividends and trading profits.

This design differs from that of Sarin and Weber in that it is an electronic instead of an oral double auction, and the asset has a three-state dividend instead of two. In addition, short sales are not allowed in this market but were permitted by Sarin and Weber.

The three-state asset paid \$0.00, \$0.50, or \$5.00 depending on the draw of a random number. All traders faced the same draw so dividend earnings per asset were the same for all traders. There was 40% probability of receiving \$0.50 in all rounds. The distribution of the balance of the probability in each round was determined by a pseudo-random number generator. In the rounds with risky assets, the probability of all states was made known prior to trade. In the rounds with ambiguous assets, only the probability of the \$0.50 payoff, uniformly 40%, was known to the subjects. For both risky and ambiguous assets independent draws were made for each round of trade.

The trading sessions exposed subjects to risky and ambiguous assets sequentially, with the two sessions controlling for order effects. This treatment was implemented rather than the direct comparison with simultaneous trade of the different assets, because the results of Sarin and Weber's simultaneous treatments demonstrated significant, persistent evidence of an ambiguity response. Although they found ambiguous prices lower than risky ones in 13 out of 14 markets, price differences in the sequential markets were smaller, and in one treatment, equivalent to session 2 in this study, the ambiguity aversion disappeared completely. The sequential study of ambiguity exposure sheds light on the plausibility of extending the comparative ignorance hypothesis to market settings. These hypotheses combine within and between subject tests in the two market sessions.

Comparing the first five rounds across sessions yields a market analog of the between group test of Fox and Tversky (1995). In the individual choice setting, Fox and Tversky failed to reject the null hypothesis of no difference between the groups, lending support to the claim that comparability was important in generating a response to ambiguity. In the market we consider if the order of exposure to the different types of assets makes a difference in market prices. Do prices for ambiguous assets differ depending on previous exposure to an unambiguous asset? We test the null hypothesis that prices for ambiguous assets do not vary across the treatments. Also, within each session do prices differ for the different types of assets? Does the market reflect ambiguity aversion even in this setting where assets are not compared simultaneously? These questions are formalized in the following hypotheses:

1. Between subject test of comparative ignorance. (Rounds 1-5 of each session.)
   H_0: Session RA risky assets are indistinguishable from Session AR ambiguous.
   H_1: Ambiguous assets will trade at lower prices than risky assets.

2. Between subject test of order effects in exposure to ambiguous assets.
   H_0: Ambiguous assets are indistinguishable across the two sessions.
   H_1: Ambiguous assets in Session RA will trade at lower prices than those in Session AR because of sequential comparison of exposure to risky assets. (Session RA:Rounds 6-10 and Session AR: Rounds 1-5)

3. Within subject test of the existence of an ambiguity price effect.
   (a) H_0: Session RA risky and ambiguous assets are indistinguishable.
   H_1: Ambiguous assets trade at a discount to risky assets.
   (b) H_0: Session 2 AR risky and ambiguous assets are indistinguishable.
   H_1: Ambiguous assets trade at a discount to risky assets.

1.5.2 Asset Market Results

Table 1.4 presents the mean prices and the mean deviations from expected value for each round of trade. Charts of the prices and the deviations are in Figure 1.2 and Figure 1.3 respectively. The expected value of the dividend changes with each round for the risky assets. When ambiguous assets are traded the expected value remains constant at $1.70 if the unknown probability is assigned equally to each of the remaining states. In all cases, statistical tests of the hypotheses are conducted on the data derived from the deviations from expected values. The test used is a randomization test that pairs the mean deviations by round with those in other sessions in the order traded (Siegel (1956)).

Table 1.4: Prices, Expected Value Deviations, and Bid/Ask Ratio

Session	Round	Mean Price	EV Deviation	Asset Type	Bid/Ask Ratio
1	1	797.67	482.67	R	0.929
1	2	680.17	540.17	R	1.053
1	3	652.82	367.82	R	1.111
1	4	667.22	472.22	R	0.654
1	5	606.43	446.43	R	0.773
1	**6**	536.67	**366.67**	**A**	**0.297**
1	**7**	503.33	**333.33**	**A**	**0.355**
1	**8**	521.78	**351.78**	**A**	**0.278**
1	**9**	495.13	**325.13**	**A**	**0.233**
1	**10**	426.67	**256.67**	**A**	**0.458**
2	1	584.00	414.00	A	2.625
2	2	618.50	448.50	A	2.083
2	3	604.67	434.67	A	1.077
2	4	601.00	431.00	A	0.773
2	5	615.17	445.17	A	1.118
2	6	622.87	402.87	R	1.529
2	7	611.37	466.37	R	0.926
2	8	622.31	357.31	R	0.955
2	9	625.15	430.15	R	1.077
2	10	616.47	511.47	R	1.250

By Session and Round of trade the mean prices and deviations are listed along with the asset type, either risky (R) or ambiguous (A). The ratio of bids to asks in a trading period are also listed. Session 1, periods 6 through 10, in bold type, reveal an effect of ambiguity through lower prices and many fewer bids per offer.

Figure 1.2: Mean Prices

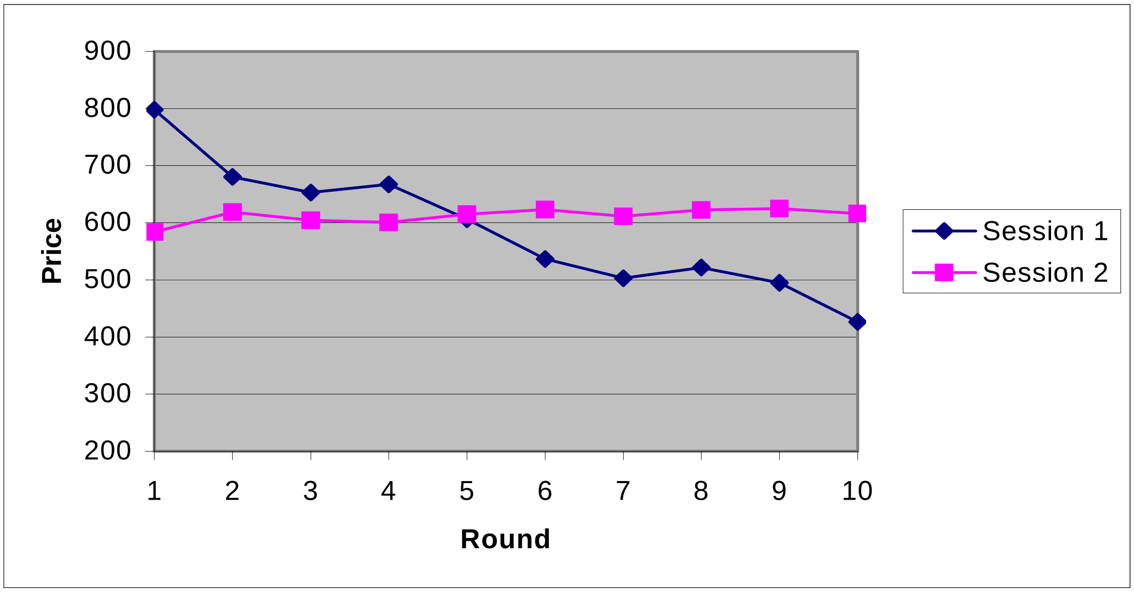

Figure 1.3: Ambiguity Aversion and Asset Allocation

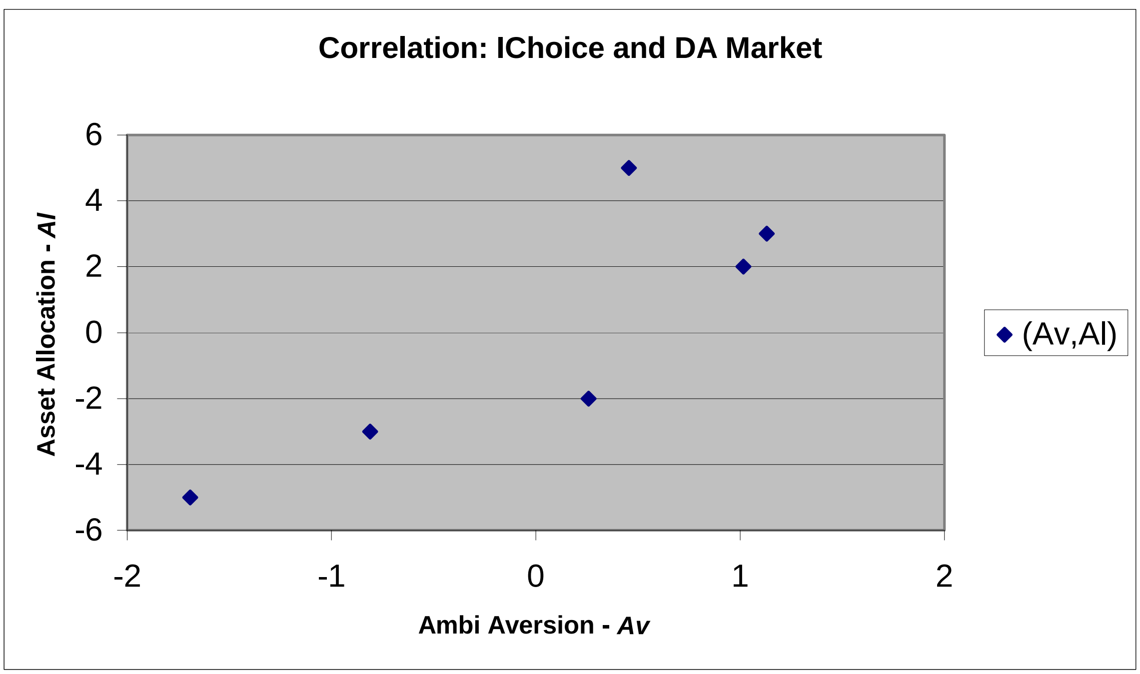

Av is the ambiguity attitude in the individual choice setting and is measured as the difference between utility under risk and under ambiguity. *Al* is the ambiguity attitude in the market environment and is measured as the difference between risky and ambiguous assets accumulated by the individual during the session.

Table 1.5 reports the results on the effects of trading sequence on pricing for all the hypotheses. The null is rejected in two cases, for hypotheses 2 and 3.a., implying that the prices of ambiguous assets differ across sessions while within sessions ambiguity has an effect on prices only when there is prior exposure to risky assets. In fact, the ambiguous assets in session RA, contain the only prices that differ significantly from any of the others. Thus, an ambiguity response is generated by the sequential treatment only when the order of the treatments provides previous exposure to an unambiguous standard. Without any prior exposure to risky assets the ambiguous assets do not differ from the risky in price either within (Hyp. 3b) or between subjects (Hyp. 1).

Table 1.5: Hypotheses on Sequential Comparability Effects

Hypothesis	**Rounds Tested**	**P value**
1. Between Subjects: Comparative Ignorance	S1:R1-5,S2:R1-5	0.15625
2. Between Subject: Ambiguous	**S1:R6-10**,S2:R1-5	**0.03125**
3. a) Within Subjects: Session 1	S1:R1-5,**S1:R6-10**	**0.03125**
3. b) Within Subjects: Session 2	S2:R1-5,S2:R6-10	0.46875

Note: S stands for session and R for round. Bold indicates the rounds in which a significant ambiguity effect was observed. Risky (ambiguous) assets are traded in S1:R1-5 & S2:R6-10 (S2:R1-5 & S1:R6-10).

A parametric random-effects model complements the randomization test results. Rather than looking at mean trading prices per round, observations are at the level of specific trades, with the error components model controlling for individual-specific errors that are constant over time. The dependent variable, *price-diff*, is calculated as the difference between the trade price and the expected value of the asset. Independent variables include *ambi* which is zero (one) when the asset traded is ambiguous (risky); a dummy variable for the session which is zero (one) for the session that trades ambiguous (risky) assets first. Also included is the interaction between the two variables and the trend variable *round*.

The regression results are presented in Table 1.6 and provide further evidence of the importance of the contrast effects associated with the comparative ignorance hypothesis. While the asset type and session are not individually significant, the interaction between the two is both statistically and economically significant. The coefficient value of -160.32 indicates that in the session in which the order of exposure was first to risky and then ambiguous assets the prices are \$1.60 lower than the other trading rounds. These findings extend the findings of Fox and Tversky (1995) and Sarin and Weber (1993) and demonstrate that in the market environment sequential trade is sufficient to generate the ambiguity response, contingent on the order of exposure.

Table 1.6: Asset pricing - random-effects model

	Coefficient	Standard error	p-value
ambi	24.32	18.34	0.19
session	25.49	19.48	0.19
ambiXsession	-160.32	33.79	0.00
round	4.99	3.32	0.13
Constant	401.50	25.54	0.00
n=176	*rho*=0.12	$\chi^2_{(4)} = 90.95$	R^2=0.38

The random-effects model estimates the difference between the expected-value and trade price for each completed transaction as a function of the asset type (risky or ambiguous), session, and the interaction of the two, as well as a trend variable for the round. *Rho* represents the proportion of the variance associated with individual-specific effects across rounds.

To summarize, the rounds in bold type in Table 1.5 and, which are the ambiguous rounds six through ten in session RA, are the only rounds that differ significantly from any of the others in terms of price and also with regard to bidding behavior. The ratio of the number of bids and offers in the markets is also presented in Table 1.4. These ratios provide a further indication of the selling pressures peculiar to these rounds and are summarized in Table 1.7.

Table 1.7: Average Bid/Ask ratios by treatment

Session	Rounds	Asset	Avg. Bid/Ask
1	1-5	R	0.90
1	6-10	A	0.32
2	1-5	A	1.53
2	6-10	R	1.15

This table summarizes the bids per ask that are listed by round in Table 1.2.

1.6 Attitudes Across Market and Individual Choice Institutions

In this section we investigate whether ambiguity attitudes are conserved across the individual choice and market institutions. The descriptor used in the individual choice environment is the ambiguity score *Av* which measures the difference in utility in the risky and ambiguous settings derived from the non-parametric method. For the asset market, the descriptor depends on the final allocation of risky and ambiguous assets for each individual. The allocation variable, Al_j, is the difference between the number of risky and ambiguous assets accumulated by trader *j* in the different market environments.[15] Thus, $Al_j = \sum_{i=1}^{5} r_{ij} - \sum_{k=1}^{5} a_{ij}$ with *i,k=1,…,5* represent the five risky and ambiguous trading rounds, respectively. The test statistic for the correlation between the ranks of the *Al* and *Av* measures is the Spearman Rank Correlation Coefficient r_S.

Table 1.8 presents the ambiguity aversion scores and their ranks and the asset accumulation differences and their ranks. Figure 1.3 charts the raw data of the

[15] Plott and Sunder (1988) and Sunder (1992) use final allocations to test hypotheses on information dissemination and aggregation under risk.

descriptors for the statistical test. The correlation between the ranks of the two measures, r_s=.83, is significant at the 5% level, implying that ambiguity attitudes are conserved across the individual choice and market institutions.

Table 1.8: Measures of Ambiguity Effects in Individual Choice and Market Settings

Trader #	Av Score	Av Rank	R	A	Al (r-a)	Al Rank
1	0.45464	3	8	3	5	1
2	1.01531	2	8	6	2	3
3	1.12987	1	9	6	3	2
4	0.25782	4	0	2	-2	4
5	-0.81028	5	4	7	-3	5
6	-1.69142	6	2	7	-5	6

Individual choice ambiguity scores (*Av*)are derived as the difference in total utility under risk and ambiguity. Marketplace ambiguity scores (*Al*) are the difference in the final allocation of assets under risk and ambiguity. The randomization test and spearman correlation indicate a significant correlation of the two ambiguity scores.

1.7 Conclusion

This study implements tests to discover the role of ambiguity attitudes in individual choice and market environments. The evidence of conservation of ambiguity attitudes across institutions appears to be a new finding. This result suggests that either attitudes towards ambiguity are more stable than those for risk, or that we have implemented an individual choice elicitation technique that happens to be congruent with the behavior in the asset market environment. Given the importance of ambiguity in field environments either result is of interest and these findings deserve further study. The results in the asset market in this study complement and extend the work of Sarin and Weber (1993). Together they demonstrate that responses to ambiguity in markets can be found in both simultaneous and sequential treatments. Thus the notion that comparative ignorance motivates ambiguity responses, first discussed by Fox and Tversky (1995) with respect to individual choice behavior, receives support also in the market environment.

With regard to alternative approaches to measuring ambiguity, our results are not conclusive. In the pooled data we observe that measures of ambiguity attitudes are slightly more informative if measured from a baseline that assumes probabilistic sophistication but not SEU preferences. This finding is due to the subject pool exhibiting probabilistic risk aversion in the questions in which they encountered risk but not ambiguity.

Appendix 1.1: The Ellsberg Paradox

This appendix provides an overview of the two-color and three-color decision problems, introduced by Ellsberg (1961). For many individuals, decisions in these settings give rise to behavior that is paradoxical in the SEU framework, since they imply that well-formed probabilities do not exist. In the two-color problem two urns with 100 balls in each are presented to the subjects. The subjects know that urn 1 contains 50 red and 50 black balls, while Urn 2 contains red and black balls in an unknown proportion. Subjects win X>0 if they predict the ball that is drawn and 0 if they are incorrect. Behavior that violates the axioms of expected utility occurs when the following two choices are made:

1.) Subjects are indifferent between choosing red or black in both urns, which implies that the subjective probability of drawing either color are identical within each urn. That is, $p_1(R) = p_1(B) = p_1$ and $p_2(R) = p_2(B) = p_2$ where the subscript indicates the urn.

2.) Subjects prefer to have a ball drawn form urn 1 with known probabilities which implies that the expected utility of urn 1 is greater than that of urn 2.

$$p_1U(X) + p_1U(0) > p_2U(X) + p_2U(0)$$

Normalizing so that $U(X) = 1$ and $U(0) = 0$ yields the result that $p_1 > p_2$, a result inconsistent with additive subjective probabilities over the events in both urns. For example, if the subjective probabilities in urn 1 are equivalent to the verifiable proportions of the balls in the urn, so that $p_1(R) = p_1(B) = .50$, this behavior yields the result that the sum of the probabilities in urn 2 is less than 1.

In the three-color problem there is one urn which contains 90 balls. Thirty are red and 60 are black or green in an unknown proportion. The decision maker faces choices over two sets of acts. The first set of is a choice between Act X and Act Y and the second a choice between Act X′ and Act Y′.

Act X: Win on draw of red.

Act Y: Win on draw of black.

Act X′: Win if red or green is drawn.

Act Y′: Win if black or green is drawn.

Paradoxical behavior with respect to SEU results when, X$\succ$Y and Y′$\succ$X′. The first choice implies that $p(R) > p(B)$ and the second that $p(B) + p(G) > p(R) + p(G)$ yielding the contradictory implication that $p(B) > p(R)$.

The reversal of the probability magnitudes between the original and primed problems can be shown to result from a failure of Savage's P2 axiom, now known as the *sure-thing* principle. The sure-thing principle requires that $R \succ B \Rightarrow R + G \succ B + G$, and it fails in this setting when the act with the unambiguous number of balls is always chosen. The paradoxical behavior is explained by the fact that the green balls act as a hedge on the uncertainty, and so their importance in the decision problem depends on what else is known. Decision theorists have reconciled this behavior by weakening the sure thing principal so that it applies only to acts that cannot serve as a hedge in this way. A more formal examination of this issue is found in Appendix 1.3 which discusses several alternative theories that reconcile the Ellsberg paradox.

Appendix 1.2: Individual Choice Sample Screen and Questions

Sample screen print

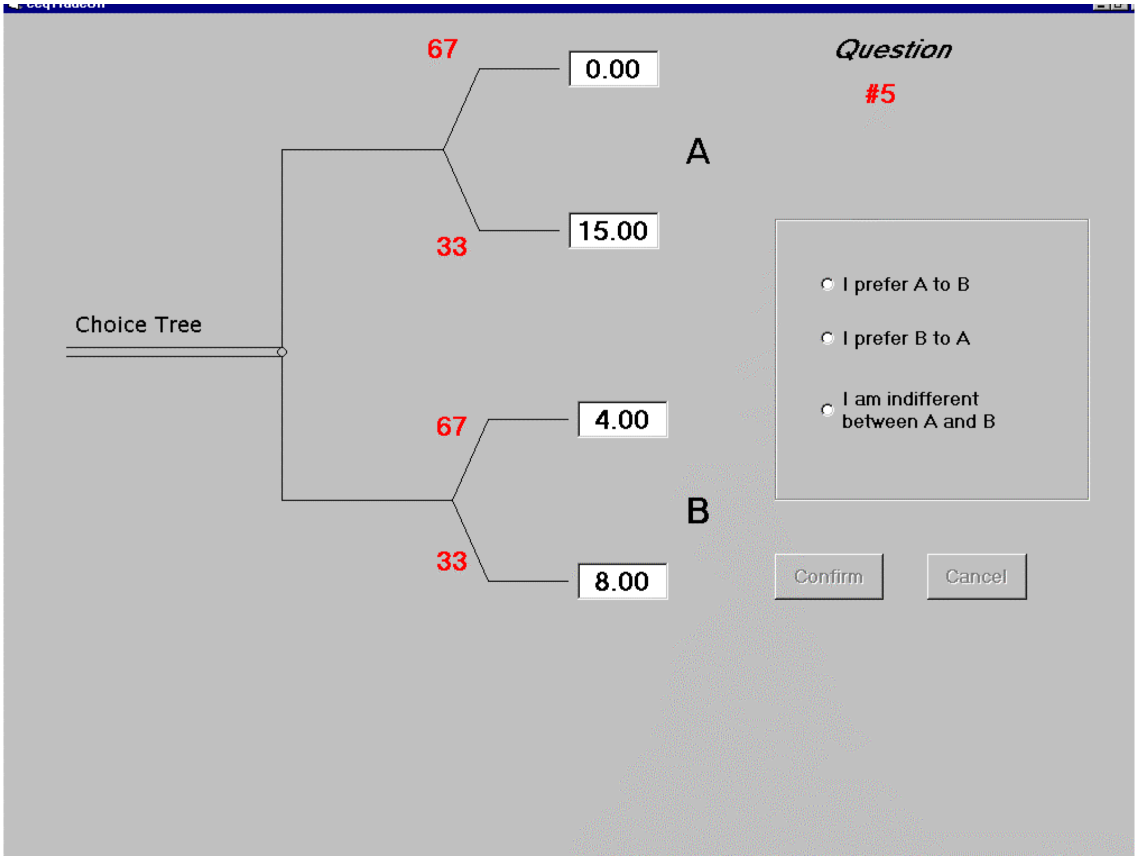

Elicitation Questions:

The lottery pairs used to elicit risky and ambiguous utilities are below in sequence. Lotteries are of the form *(x,p;y)*, where *x* is a payoff received with probability *p*, and *y* a payoff received with probability *(1-p)*. "?" implies an ambiguous lottery where the second probability is also presented as "?".

Question	Lottery A	Lottery B
1	(15,33;0)	(6,33;4)
2	(15,33;0)	(8,33;4)
3	(15,33;0)	(10,33;4)
4	(15,33;0)	(11,33;4)
5	(6,33;0)	(0,?;4)
6	(6,33;0)	(1,?;4)
7	(6,33;0)	(2,?;4)
8	(6,33;0)	(3,?;4)
9	(12,33;0)	(3,33;4)
10	(12,33;0)	(5,33;4)
11	(12,33;0)	(6,33;4)
12	(12,33;0)	(7,33;4)
13	(12,33;0)	(8,33;4)
14	(9,33;0)	(2,?;4)
15	(9,33;0)	(3,?;4)
16	(9,33;0)	(4,?;4)
17	(9,33;0)	(5,?;4)
18	(9,33;0)	(6,?;4)

Question	Lottery A	Lottery B
19	(9,33;0)	(7,?;4)
20	(9,33;0)	(2,33;4)
21	(9,33;0)	(3,33;4)
22	(9,33;0)	(4,33;4)
23	(9,33;0)	(5,33;4)
24	(9,33;0)	(6,33;4)
25	(9,33;0)	(7,33;4)
26	(12,33;0)	(3,?;4)
27	(12,33;0)	(5,?;4)
28	(12,33;0)	(6,?;4)
29	(12,33;0)	(7,?;4)
30*	(9,33;0)	(8,?;4)
31	(6,33;0)	(0,33;4)
32	(6,33;0)	(1,33;4)
33	(6,33;0)	(2,33;4)
34	(15,33;0)	(6,?;4)
35	(15,33;0)	(8,?;4)
36	(15,33;0)	(10,?;4)
37	(15,33;0)	(11,?;4)

Appendix 1.3: Modeling Ambiguity

In this appendix we provide additional details on alternative approaches to modeling ambiguity; an area of active research. As discussed in Appendix 1.1, the violation of SEU inherent in the Ellsberg paradox is associated with a failure of the sure-thing principal. A seminal paper modifying SEU to resolve the paradox is that of Schmeidler (1989) who restricted the applicability of independence to comonotonic acts and then derived a measure of ambiguity aversion based on sub-additive capacities (non-additive measures). Schmeidler's axiom of comonotonic independence states that for $\alpha \in (0,1)$, $f \succ g$ implies that $\alpha f + (1-\alpha)h \succ \alpha g + (1-\alpha)h$ for pairwise comonotonic acts *f*, *g*, and *h*. Acts are comonotonic if for no states *s* and *t* is $f(s) \succ f(t)$ and $g(t) \succ g(s)$. The implications of comonotonicity are also discussed by Yaari (1987) who views it as a "no-hedge" condition. In the Ellsberg three-color problem the behavior that is paradoxical from the perspective of SEU is attributed to the hedge on ambiguity provided by the state G in Y'. That is, while B and G are individually ambiguous, $B \cup G$ is unambiguous. Convex capacities rationalize the ambiguity averse preferences by allowing for sub-additivity in beliefs such that for A and B disjoint, $v(A \cup B) \geq v(A) + v(B)$. Thus in the three ball problem $v(B \cup G) > v(B) + v(G)$ and a ranking of the weights such that $v(B \cup G) > v(R \cup G) = v(R) + v(G) > v(B) + v(G)$ rationalizes the observed choices.

Following Schmeidler's fundamental contribution, additional research has investigated whether convexity of the capacity is an appropriate way to define ambiguity. Both Epstein (1999) and Ghirardato and Marinnaci (2002; hereafter GM) demonstrate that convexity of a capacity is neither necessary nor sufficient for ambiguity-averse preferences with examples that increase the number of ambiguous states beyond the two found in the Ellsberg paradox and in the experiments reported here.[16] Common to both of these critiques is an attempt to ground the definition of ambiguity aversion using a comparative foundation in a way that is similar to Yaari's (1969) contribution with regard to risk. The comparative approach requires choosing an ambiguity neutral benchmark, just as Yaari used expected value maximization as the risk neutral referent for defining risk aversion.

GM argue that constant acts can play a similar role as intuitive benchmarks for ambiguity neutrality, as in the theory of risk and reflect the "weakest prejudgment" regarding acts that can be considered unambiguous. For two preference relations $\succ_1$ and $\succ_2$, GM define $\succ_2$ as more uncertainty averse than $\succ_1$ if $x \succ_1 f \Rightarrow x \succ_2 f$, where *x* (*f*) is a constant (non-constant) act. The further assumption that the two preferences share the same cardinal risk attitude implies that $\succ_2$ is more ambiguity averse than $\succ_1$. GM show that for the class of preferences under consideration SEU preferences ($\succ_{SEU}$) are the only ones that are ambiguity neutral.

The comparative notion of ambiguity aversion developed by Epstein (1999) makes use of a larger set of acts, not all of which are constant, to serve as unambiguous referents. This choice implies that the absolute measure of ambiguity aversion is made relative to probabilistically sophisticated preferences, $\succ_{PS}$. Probabilistic sophistication implies that acts are evaluated with respect to a (subjective) probability distribution over outcomes, and thus that acts are lotteries over pure risk (Epstein (1999); p.585). The set

[16] Ghirardato and Marinacci (2002) argue that Epstein's critique also has some unintuitive characteristics. In particular the sum of unambiguous events in his example do not add up to one.

of probabilistically sophisticated preferences is larger than that of the SEU preferences considered by GM. It includes preferences in which probabilities exist, but may be transformed, for example to reflect optimism or pessimism. Thus, preferences explainable by Quiggin's (1982) rank dependent expected utility model are consistent with ambiguity neutrality in Epstein's formulation but not in GM's.

Appendix 1.4: Calculating Non-parametric Ambiguity Scores

Ambiguity scores in the individual choice setting were calculated by generating piecewise linear utility functions for each subject for both the risky and ambiguous lottery sequences. Consider the two segment example presented in the text in which indifference between $U_1(15, .33; 0) = U_1(8, .33; 4)$ and $U_2(12, .33; 0) = U_2(5, .33; 4)$ are postulated, implying $U_1(15) - U_1(8) = U_2(12) - U_2(5) = U_1(4) + U_2(4)$ where $U_i \;, i = 1{,}2$ represents the utility along the first and second elicited segments respectively.. The piecewise linear utility function is calculated by minimizing the difference between the endpoints that are common to both segments as follows

$$\min \left|U_1(12) - U_2(12)\right| + \left|U_1(8) - U_2(8)\right| \; s.t. \, U_1(15) - U_1(8) = U_2(12) - U_2(5);$$

$U_i \;, i = 1{,}2$ represents the utility along the first and second elicited segments respectively. More generally, over the four segments the following linear program was implemented.

$$\min_{U(C_i^*),U(x_i)} \quad \sum_{i=1}^{n} \left|U_i(x_{i+1}) - U_{i+1}(x_{i+1})\right| + \left|U_i(C_i^*) - U_{i+1}(C_i^*)\right|$$

subject to:

$$U_1(x_1) = 1,$$
$$U_4(0) = 0,$$
$$U_i(x_i) - U_i(C_i^*) = U_{i+1}(x_{i+1}) - U_{i+1}(C_{i+1}^*) \qquad i = 1,\ldots,3$$

2 Information Cascades: Evidence from a Field Experiment with Financial Market Professionals[17]

2.1 Introduction

In economic and financial environments in which decision makers have imperfect information about the true state of the world, it can be rational to ignore one's own private information and make decisions based upon what are believed to be more informative public signals. In particular, if decisions are made sequentially and earlier decisions are public information, "information cascades" can result. Information cascades arise when individuals rationally choose identical actions despite having different private information.[18] Cascades may arise in myriad settings, including technology adoption, medical treatment, and environmental hazard response. Arguably, however, the most well-known herds or cascades occur in financial markets, where bubbles and crashes may be examples of such behavior.[19]

Since the private information of cascade followers is not revealed, information cascades can be suboptimal. Moreover, because the small amount of information revealed early in a sequence has a large impact on social welfare, cascades can be fragile, with abrupt shifts or reversals in direction when new information becomes available

[17] This chapter is authored jointly with Michael Haigh and John List.

[18] Herding is a more general phenomenon than an informational cascade though both result in behavioral conformity. The homogeneity of a herd may arise through other than informational means such as payoff externalities, preferences for conformity, or sanctions. A comprehensive taxonomy of herd behavior is developed by Hirshleifer and Teoh (2003) and Smith and Sorenson (2000). Devenow and Welch (1996) and Bikhchandani and Sharma (2000) also discuss alternative sources of herd behavior and review the extant literature.

[19] It has been argued, also, that information cascades can explain a large variety of social behaviors such as fashion, customs, and rapid changes in political organization. Anderson (1994), Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992, 1998), and Welch (1992) discuss a variety of interesting examples. A number of historical anecdotes can be found in MacKay (1980) and Garber (2000).

(Bikhchandani, Hirshleifer, and Welch (1992, 1998; hereafter BHW), Gale (1996), Goeree et al. (2004)). Indeed, some argue that the volatility induced by herding behavior can increase the fragility of financial markets and destabilize the broader market system (Eichengreen et al. (1998), Bikhchandani and Sharma (2000), Chari and Kehoe (2004)).

Previous empirical approaches that examine cascade behavior can be divided into two classes; regression-based tests that use naturally occurring data and laboratory experiments that use data gathered from student subjects. In a review of the extant regression-based results for herding in financial markets, Bikhchandani and Sharma (2000) note the difficulty of controlling for underlying fundamentals, and argue that as a result of this difficulty there is often "a lack of a direct link between the theoretical discussion of herding behavior and the empirical specifications used to test for herding."[20] The laboratory environment, in contrast, allows one to control for public and private information and thus to make explicit tests of theoretical predictions more easily. Yet an important debate exists about the relevance of experimental findings from student subjects for understanding phenomena in the field. For example, professional behavior in the field might differ from student behavior in laboratory experiments due to training or regulatory considerations, which may affect the development of decision heuristics, as well as the overall naturalness of the experimental environment (see, for example, Harrison and List (2004)). Locke and Mann (2005) argue that financial market research that ignores the effect of professional expertise is likely to be received passively because "ordinary" individuals, as opposed to professional traders, are too far removed from the price discovery process. Bikhchandani and Sharma (2000, p. 13) also argue that "to

[20] Fama (1998) discusses the interpretation of empirical results as evidence of irrational behavior.

examine herd behavior, one needs to find a group of participants that trade actively and act similarly."

We find these arguments compelling and therefore combine the most attractive aspects of these two classes of empirical research, that is, we observe professionals in a controlled environment, and extend the literature in several new directions. First, we compare the behavior of market professionals from the floor of the Chicago Board of Trade (CBOT) with that of college students in an experimental setting in which the underlying rationality of herd behavior can be identified. Second, given the vast normative implications of work that has established the importance of the domain of earnings for decision making under risk (Kahneman and Tversky (1979), Shefrin and Statman (1985), Odean (1998)), we examine the behavior of each group in the gain and loss domain. We further examine whether, and to what extent, cascade formation is influenced by both private signal strength and the quality of previous public signals, as well as decision heuristics that differ from Bayesian rationality. Finally, within the group of market professionals, we examine the extent to which differences in cascade formation are associated with individual characteristics such as whether the participant is a day trader.

Empirical findings gained from an examination of more than 1,500 individual decisions lend some interesting insights into cascade behavior. A key finding is that market professionals tend to make use of their private signal to a greater degree and base their decisions on the quality of the public signal to a greater extent, than do students. As a result, the professionals are involved in weakly fewer overall cascades and significantly fewer *reverse cascades* (cascades that lead to inferior outcomes). This result is novel to the literature and has important implications for financial markets.[21] Further, while the behavior of students is consistent with the notion that losses loom larger than gains, market professionals are unaffected by the domain of earnings. This finding is consistent with Locke and Mann (2005), Genesove and Mayer (2001), and List (2003, 2004), who find, in varying environments, that market experience is associated with a decline in deviations from classical assumptions.

Note that we observe behavioral differences not only across subject pools, but also within the market professional group. For example, Bayesian play is correlated with market experience and day traders are much more likely to join an informational cascade than are non-day traders. Finally, we present data on the prevalence of non-Bayesian decision heuristics, an area in which the two subject pools demonstrate similarities.

The remainder of the study is crafted as follows. Section II outlines the basic theory and experimental design. Section III presents our empirical results. Section IV considers implications of our results for financial markets and briefly discusses the use of professionals in experimental practice more broadly. Section V concludes.

2.2 Theory and Experimental Design

Imitative behavior associated with herding has often been viewed as the product of irrational decision-making (Keynes (1936); Shleifer and Summers (1990); Hirshleifer (2001)). Alternatively, models such as Banerjee (1992), BHW (1992), and Welch (1992) consider the conditions under which it is rational to join a cascade. The model we

[21] Combined with the insights gained from the models of Barberis, Shleifer, and Vishny (1998), Daniel, Hirshleifer, and Subrhamanyam (2001), and from Hirshleifer (2001), our results indicate that the ability of the strength and weight of the evidence to have a differential impact on asset pricing is a potentially powerful phenomenon.

present below, and the experimental environment we implement, is consistent with the work of this second set of authors in that it is predicated on Bayesian updating of beliefs, given private signals and a history of observable actions.[22] The empirical investigation of the cascade phenomenon raises interesting questions beyond whether agents update information in a manner that is consistent with Bayes' rule.[23] Since the formation of informational cascades is a social phenomenon, individual behavior may depend on how agents view the rationality of others. Accordingly, we examine how our two subject pools respond to uncertainty about the quality of information that arises due to potential deviations from Bayesian rationality by others. We adopt two approaches. First, we use a model in which the null hypothesis is that Bayesian rationality is universally applied and is common knowledge. Second, we estimate a quantal response equilibrium (QRE) model that assumes decision error (McKelvey and Palfrey (1998), Goeree, Holt, and Palfrey (2005)).

2.2.1 Theoretical Model and Predictions

Consider an environment in which there are two possible underlying states of nature $\Omega = \{A, B\}$, with the true state denoted by $\omega \in \Omega$. Each of a set of $I = \{1,2,..,n\}$ agents receives an independent private signal, $s_i \in \{a,b\}$, that is informative in the sense

[22] As we discus below, our experimental environment makes use of a binary signal, binary state, and fixed payoff regardless of the history of announcements. Avery and Zemsky (1998), Lee (1998), Chari and Kehoe (2004), and Cipriani and Guarino (2005a) explore more general settings in which variable pricing reduces but does not eliminate the potential for information cascades. Chamley (2004) provides a comprehensive review of rational herding models.

[23] The ability of humans to reason in a Bayesian manner seems to depend on how information is presented. Studies that present base rates as percentages often show that we are poor "intuitive statisticians" (Tversky and Kahneman (1974)). Decisions tend to be more consistent with Bayesian rationality when individuals experience probability distributions through repeated exposure (see Gigerenzer and Murray (1987)). Our experiment is consistent with protocols that have been shown to give Bayesian decision making its best chance.

that $\Pr(A|a) > \Pr(B|a)$ and $\Pr(A|b) < \Pr(B|b)$. Signal precision, given by $\Pr(s = \omega | \omega)$, is identical for all agents. After receiving their signal, each agent chooses either *A* or *B* with their choice, c_i. If $c_i = \omega$, individual *i* receives a reward normalized to one. , If $c_i \neq \omega$, individual *i* receives zero. Each individual receives their signal in an exogenously determined choice order. Along with their private signal s_i, each agent observes the history of choices, $H_i = \{c_1, ..., c_{i-1}\}$. The prior probability of an underlying state, given by $\Pr(\omega = A) = p$ and $\Pr(\omega = B) = 1 - p$, is common knowledge. If all individuals update beliefs according to Bayes' rule and this updating is common knowledge, the posterior probability $\Pr(\omega | H_i, s_i)$ is easily derived. We demonstrate the formation of an information cascade in this setting via a simple example, parameterized with the probabilities from one of our experimental treatments.

Let $\Pr(\omega = A) = \Pr(\omega = B) = p = 1/2$ be the prior probability, with the precision of the symmetric signal given by $\Pr(a | A) = \Pr(b | B) = 2/3$, with complementary probabilities $\Pr(b | A) = \Pr(a | B) = 1/3$. Suppose that $s_1 = a$. Bayes' rule implies that

(2.1) $$\Pr(\omega = A | s_1 = a) = \frac{\Pr(a | A)\Pr(A)}{\Pr(a | A)\Pr(A) + \Pr(a | B)\Pr(B)} = \frac{2}{3}$$

An expected utility maximizer would therefore predict *A* as the state of nature since expected profits for announcing *A*, π_A, exceed those for announcing *B*, π_B.[24] If the second subject also receives an *a* signal, updating according to Bayes' rule yields

[24] In the gain treatments, $\pi_A - \pi_B = \frac{\$W}{3}$ after an initial *a* signal, where \$W is the win amount. Treatments over gains and losses yield identical predictions (i.e., expected losses are minimized by picking the most probable urn).

$$(2.2) \quad \Pr(\omega = A \mid H_2 = A, s_2 = a) = \frac{\Pr(a \mid A)^2}{\Pr(a \mid A)^2 + \Pr(a \mid B)^2} = \frac{4}{5}$$

That is, two consecutive identical announcements yield a posterior probability of 0.80 in favor of the indicated urn.[25] As a result, the third decision maker should "follow the herd" and choose $\omega = A$ regardless of her signal, as can be seen by examining the posterior in which an opposing *b* signal is the private draw of the third player after two consecutive *A* announcements:

$$(2.3) \quad \Pr(\omega = A \mid H_3 = A, A, s_3 = b) = \frac{\Pr(a \mid A)^2 \Pr(b \mid A)}{\Pr(a \mid A)^2 \Pr(b \mid A) + \Pr(a \mid B)^2 \Pr(b \mid B)} = \frac{2}{3}$$

We classify a decision of this type—consistent with Bayesian rationality, but inconsistent with one's own private signal—as a *cascade decision*. In this example, the decision maker in the third position reveals nothing about their private information and thus the positive externality associated with learning from other's choices is blocked by a cascade. The analysis implies that, with this parameterization, public announcements are uninformative whenever the number of public signals of one type exceeds the other by two or more. As a result, if a cascade has not started, two consecutive low probability draws can result in a *reverse cascade* whereby everyone rationally herds on the incorrect state.

[25] A second A announcement could arise in this setting if the second subject receives a *b* signal. We consider an announcement of A given the history Ab to be inconsistent with Bayesian rationality, although alternative interpretations are possible. Since the posterior probability is 0.5 in this case, a tie-breaking rule must be invoked. We follow Anderson and Holt (1997) in assuming that individuals who are indifferent announce their own signal. This is sensible if individuals recognize the possibility of decision error in previous announcements. Alternative tie-breaking rules include random choice as in BHW (1992) and a "nonconfident" rule in which one ignores one's own information (Koessler and Zieglemeyer (2000)). In our treatments the Anderson and Holt rule is followed 81% of the time, with most of the deviations occurring in the early rounds of play.

2.2.2 Experimental Design

Anderson and Holt (1997) present a seminal experimental investigation of cascade formation using a subject pool of undergraduates. To ensure comparability of our results to the extant literature, we use experimental protocols that are closely related to those of their work.[26] The parameterization in the example above is consistent with their symmetric treatment $(\Pr(a \mid A) = \Pr(b \mid B))$. The experimental sessions we conduct comprise 15 *rounds* of the basic game for a group of either five or six players whose choice order in each round — either first, second, third, …., sixth — is determined by a random draw.

A round begins with the experimental monitor selecting the state of nature with a roll of a die that is unobserved by the subjects. Subjects gain information about the state by drawing a single ball out of an unmarked bag into which the contents of the selected urn have been transferred. The draw is made while the subject is isolated from the other players. The monitor is informed of the choice of the state, and announces it publicly. After all subjects have made their choices, the true state is revealed.

To provide exogenous variation in the informational content of the private signal across treatments, we use two urn types. In the symmetric treatment, Urn A contains two type-*a* balls and one type-*b* ball, while Urn B contains two type-*b* balls and one type-*a* ball. To create the *asymmetric* treatment, we add four *a* balls to both urns, yielding 6 (5)

[26] Our experimental instructions are available upon request. Note that Anderson and Holt (1997) find that cascades form in roughly 70 percent of the rounds in which they are possible. Deviations from Bayesian cascade formation occur most often when a simple counting rule gives a different indication of the underlying state. Extensions to the experimental literature introduce relevant complications to the cascade process that include costly information, endogenous sequencing of choice order, collective decision making, expanded signal spaces, and payoff externalities (Celen and Kariv (2004, 2005), Cipriani and Guarino (2005b); Drehmann, Oechssler, and Roider (2005), Huck and Oechssler (2000), Hung and Plott (2001), Kubler and Weizsacker (2004), Noth and Weber (2003), Sgroi (2003), Willinger and Ziegelmeyer (1998)).

a signals and 1 (2) b signal in the A (B) state. This modification results in a significant dilution of the strength of an a signal, the relative weakness of which can be observed in Table 2.1, which provides posterior probabilities for all possible signal histories for both the symmetric and asymmetric urn types. As an example, the two-thirds probability that arises after a single a draw in the symmetric treatment arises after four consecutive a draws in the asymmetric setting. One consequence of the change in signal strength is that in the asymmetric treatment, a cascade on the B state should take place after one b signal even with either one or two a signals in the game's history.

The difference in signal strength across urn types allows us to investigate the relationship between Bayesian updating and a choice heuristic based on a counting rule. In the symmetric treatment, the optimal decision is always consistent with choosing the state with the most informative signals. In the asymmetric case, four sequences violate this counting rule in that it is optimal to choose B even when there are fewer b signals; these *noncounting rule* sequences are $(a,b) \in \{(2,1),(3,1),(3,2),(4,2)\}$, as indicated by bold type in Table 2.1. Thus, the asymmetric treatment allows us to gain insights into the extent to which decisions are better characterized as following a counting heuristic rather than Bayesian updating.

Table 2.1: Posterior Probabilities-Symmetric (upper) and Asymmetric (*lower*) Urns

a \ b	0	1	2	3	4	5	6
0	0.500	0.330	0.200	0.110	0.060	0.030	0.020
	0.500	*0.333*	*0.200*	*0.111*	*0.059*	*0.030*	*0.015*
1	0.670	0.500	0.330	0.200	0.110	0.060	
	0.545	*0.375*	*0.231*	*0.130*	*0.070*	*0.036*	
2	0.800	0.670	0.500	0.330	0.200		
	0.590	***0.419***	*0.265*	*0.153*	*0.083*		
3	0.890	0.800	0.670	0.500			
	0.633	***0.464***	***0.302***	*0.178*			
4	0.940	0.890	0.800				
	0.675	*0.509*	***0.341***				
5	0.970	0.940					
	0.713	*0.554*					
6	0.980						
	0.749						

Entries represent the posterior probabilities for all possible sequences of draws for both symmetric (upper) and asymmetric (lower) treatments based on choice histories (a, b). The prior probability of an urn is 0.5 in (0,0). Bold entries for the asymmetric urn are those in which counting and the posterior probability make different predictions about the state.

To provide exogenous variation in the earnings domain, we randomly place subjects in either a gain or a loss treatment for all 15 rounds. The treatment is implemented so that in gain (loss) space a correct (incorrect) inference about the underlying state results in positive (negative) earnings of $1 for students and $4 for market professionals.[27] An incorrect (correct) choice in gain (loss) space results in no earnings. To generate similar monetary outcomes across treatments, in the loss treatments, students and market professionals are endowed with $6.25 and $25.00, respectively.[28] We believe that this is the first study to vary the gain/loss domain in cascade games.

[27] CBOT officials suggest that designing a 30-minute game with an expected average payout of approximately $30 is more than a reasonable approximation of an average trader's earnings for an equivalent amount of time on the floor. In our experiments the median earnings for the market professionals are slightly in excess of this amount and therefore likely to be salient.

[28] To ensure that subjects depart with positive money balances we have both subject pools participate in other unrelated games during the experimental session.

Experimental subjects in a particular session consist entirely of one of the two subject types, students or market professionals. The experimental sessions with market professionals are conducted at the Chicago Board of Trade (CBOT) and the student data are gathered from undergraduates at the University of Maryland in College Park. The CBOT (student) subject pool includes 55 (54) subjects recruited from the floor of CBOT (the university). The resulting experimental design is a 2x2x2 factorial across urn type (symmetric (S) or asymmetric (A)), domain type (gains (G) or losses (L)), and subject type (college undergraduates (C) or market professionals (M)). Each experimental session consists of a group of either five or six participants making decisions within the same treatment type over 15 rounds. Table 2.2 summarizes our experimental sessions.

Table 2.2: Experimental Design

	Symmetric Urn		Asymmetric Urn	
	Gains	Losses	Gains	Losses
Panel A: Ten Market Professional Sessions				
Number of Sessions	3	1	3	3
Participants in Session	5	5	One with 5, two with 6	6
Total Decisions	225	75	255	270
Average Earnings	$43.20	-$20.80	$39.06	-$22.89
Panel B: Ten Student Sessions				
Number of Sessions	3	1	3	3
Participants in Session	One with 5, two with 6	5	One with 5, two with 6	5
Total Decisions	267	75	255	225
Average Earnings	$11.61	-$2.80	$11.00	-$6.40

Panel A (B) shows that Market Professionals (Students) are exposed to either the Symmetric or Asymmetric urn and play the game in either the gain or the loss domain. The symmetric urn consists of three balls — two *a* and one *b* in Urn A, and one *b* and two *a* in Urn B. The Asymmetric urn consists of seven balls — six *a* and one *b* in Urn A, and five *a* and two *b* in Urn B. The number of decisions is a function of the number of players, the number of games, and the number of rounds in each game.

2.3 Experimental Results

Table 2.3, Panel A presents descriptive statistics from the experiment. We report the rate of Bayesian decision making and the rate of cascade formation, with a Bayesian decision defined assuming common knowledge of Bayesian rationality (no decision error). Pooled, the 20 experimental sessions yield a total of 1,647 decisions, 1,284 (78%) of which are consistent with a perfect Bayesian equilibrium.[29] Cascade decisions (i.e., Bayesian decisions inconsistent with the private signal) occur in 15% of the choices. Of these, just under one-quarter (55 out of 245) are "reverse" cascades, resulting in the wrong inference about the underlying state.

Perhaps more revealing than the aggregate number of cascades is the proportion of cascade decisions made when the opportunity arises. Recall that a cascade decision is possible only when the private draw is inconsistent with the probability weight derived from the choice history and one's own private signal. In our data, cascade formation is possible in 441 of the decisions, representing 27% of the total; cascades are realized in 245 (56%) of these cases. These results are presented in the *potential* and *realized* cascades columns of Table 2.3, Panel A.

Table 2.3, Panel A also reports statistics disaggregated by subject and treatment type. In aggregate, 81% (75%) of the students' (market professionals') decisions are consistent with Bayesian Nash equilibrium. Decisions of individual subjects range from 38% to 100% Bayesian (these results are not shown to conserve space), and of the 14

[29] In the discussion that follows we use the term "Bayesian decision" to mean that the decision is consistent with the predictions of perfect Bayesian equilibrium.

subjects perfectly consistent with Bayesian rationality, 10 were students.[30] In situations in which Bayesian behavior requires that one ignore private information, fewer agents are Bayesian: The final column of Table 2.3, Panel A shows that 61% (49%) of students (market professionals) ignore their signal when doing so leads to a cascade. Interestingly, rates of cascade formation and Bayesian decision making are lower in the asymmetric treatments for both subject pools.

Table 2.3: Aggregate Decision Making

Treatment	Bayesian	Cascades (total)	Reverse Cascades	Potential Cascades	Realized Cascades
Panel A: Decision Making Pooled and by Treatment					
Pooled Data					
C & M	0.780	0.149	0.033	0.268	0.556
n=1,647	1,284	245	55	441	245/441
College Student Treatments (C)					
C	0.814	0.178	0.045	0.292	0.608
n = 822	669	146	37	240	146/240
SGC	0.940	0.157	0.041	0.172	0.913
n = 267	251	42	11	46	42/46
SLC	0.960	0.067	0.013	0.080	0.833
n = 75	72	5	1	6	5/6
AGC	0.682	0.251	0.051	0.451	0.557
n = 255	174	64	13	115	64/115
ALC	0.764	0.155	0.053	0.324	0.480
n = 225	172	35	12	73	35/73
Market Professional Treatments (M)					
M	0.745	0.120	0.021	0.244	0.493
n = 825	615	99	18	201	99/201
SGM	0.818	0.098	0.022	0.142	0.688
n = 225	184	22	5	32	22/32
SLM	0.867	0.147	0.067	0.213	0.688
n = 75	65	11	5	16	11/16
AGM	0.714	0.133	0.008	0.275	0.486
n = 255	182	34	2	70	34/70
ALM	0.681	0.133	0.022	0.307	0.385
n = 270	184	32	6	83	32/83

[30] Thirteen of the 14 who are perfectly consistent with Bayesian rationality are in the symmetric urn treatment. One market professional is perfectly Bayesian in the asymmetric setting.

Table 2.3 (*cont.*): Aggregate Decision Making

Treatment	Bayesian	Cascades (total)	Reverse Cascades	Potencial Cascades	Realized Cascades
Panel B: Decision Making by Counting Rule Predictions (Asymmetric Treatments)					
Pooled Data					
C & M	0.709	0.164	0.033	0.339	0.477
n = 1,005	712	165	33	341	165/341
Count = Baye	0.759	0.152	0.024	0.267	0.565
n = 843	640	128	20	225	128/225
Count ≠ Baye	0.444	0.228	0.080	0.716	0.313
n = 162	72	37	13	116	37/116
College Student Treatments (C)					
C	0.721	0.206	0.052	0.392	0.527
n = 480	346	99	25	188	99/188
Count = Baye	0.760	0.189	0.036	0.325	0.582
n = 412	313	78	15	134	78/134
Count ≠Baye	0.485	0.309	0.147	0.794	0.389
n = 68	33	21	10	54	21/54
Market Professional Treatments (M)					
M	0.697	0.126	0.015	0.291	0.431
n = 525	366	66	8	153	66/153
Count = Baye	0.759	0.116	0.011	0.211	0.550
n = 431	327	50	5	91	50/91
Count ≠Baye	0.415	0.170	0.032	0.660	0.258
n = 94	39	16	3	62	16/62

The *Bayesian* column represents the total proportion and number of decisions consistent with a perfect Bayesian equilibrium. *Cascade* decisions (those that are Bayesian but for which private information is ignored) and *reverse* cascades (cascades in which the wrong inference of the underlying state occurs) occupy the next two columns. The *potential* cascades category represents the proportion (and number) of cascades that could have occurred when it was possible to make one, and the *realized* cascades category represents the proportion of those potential cascades that were actually realized. "n" = number of decisions. Treatment codes are S = symmetric, A = asymmetric, G = gain, L = loss, C = college student, and M = market professional. Panel A includes all decisions and Panel B restricts attention to those sequences in which the Bayesian posterior and a counting rule make different predictions.

The final set of descriptive statistics is presented in Table 2.3, Panel B, which displays results from the asymmetric treatments first pooled and then parsed by subject pool and sequence type, where the type is either a counting rule or a noncounting rule sequence.[31] Table 2.3, Panel B demonstrates that both Bayesian decision making and

[31] We will see below that there are differences between the symmetric and asymmetric treatments even after controlling for the counting rule sequences. As a result, we do not pool the symmetric results in this table.

cascade formation decline when the rules are not reinforcing: The proportion of Bayesian decisions by students (market professionals) declines from 76% (76%) to 49% (42%), and the rate at which cascades obtain declines from 58 percent (55%) to 39% (26%). These results suggest that the non-counting rule sequences pose a challenge for both subject pools.[32]

To permit more formal inference, we apply a variety of parametric and non-parametric statistical techniques and group our results into five categories. Three of the categories compare students and market professionals to consider differences in (1) Bayesian decision making, (2) cascade formation, and (3) behavior across the gain/loss domain. A fourth category concentrates on data from market professionals by making use of additional demographic data collected during the experiment. The fifth category considers the exogenous alteration of signal strength through the use of the symmetric and asymmetric urns. Our analysis leads to the following insight:

> *Result 1: Market professionals are less Bayesian than students. Despite this behavioral discrepancy, earnings are not significantly different across subject pools.*

To provide evidence of this result we employ both unconditional and conditional statistical tests. When using unconditional tests, we account for the data dependencies within an experimental session by using session-level aggregates to yield the most conservative statistical tests. Our unconditional test used to support Result 1 is a non-

[32] Anderson and Holt (1997) find that the rate of Bayesian behavior in the noncounting rule sequences is 50%, comparable to our student population rate of 49% and close to the pooled rate of 44%.

parametric Mann-Whitney U test, which indicates that the rate of Bayesian decision making differs across subject pools at a level of significance of $p = 0.052$.[33]

To complement this analysis, we employ conditional tests that recognize the panel nature of our data; in particular, we use a random effects probit specification of the form

(2.4) $Baye_{it} = \beta` X_{it} + e_{it}, \qquad e_{it} \sim N[0,1],$

where $Baye_{it}$ equals unity if agent i is a Bayesian in round t under the assumption of no decision error by preceding players, and zero otherwise, and X_{it} includes treatment effects (*gain*, *sym*, and *trader*) and other variables predicted to influence play (*order*_x, *diff*, and *heurist*). The treatment variables are as defined above: *gain* equals one (zero) for sessions in the domain of gains (losses), *sym* equals one (zero) for the symmetric (asymmetric) sessions, and *trader* equals one (zero) for market professionals (students).

The remaining variables are defined as follows. The categorical variable *order_x (x=2,..,6)* indicates the positional order in which the individual choice is made. The posterior probability is incorporated in the variable *diff,* which is calculated as $\left|\Pr\left(\omega = A \mid H, s\right) - 0.5\right|$ and measures the accrued public and private information at the disposal of each decision maker; note that *diff*, therefore, varies from zero to one-half, increasing with evidence of the underlying state.[34] The variable *heurist* is equal to one (zero) for noncounting rule (counting rule) sequences. In a perfect Bayesian equilibrium, the coefficients of these latter two variables should not differ from zero.

We specify $e_{it} = u_{it} + \alpha_i$, where the two components are independent and normally distributed with mean zero: $\mathrm{Var}(e_{it}) = \sigma_u^2 + \sigma_\alpha^2$. We estimate equation (2.4) using the

[33] There are 10 session-level observations for each subject pool as summarized in Table II.

[34] The posterior, and thus the *diff* variable, remains constant once a cascade has formed, unless a decision breaking the cascade is observed.

maximum likelihood approach derived in Butler and Moffitt (1982). Estimation of this model is amenable to Hermite integration. To estimate the model, we use a 12 point quadrature and the method of Berndt et al. (1974) to compute the covariance matrix.

Empirical results are reported in Table 2.4, which presents the marginal effects associated with a change in each of the regressors computed at the overall sample means.[35] Concerning subject pool effects, results from both a likelihood ratio test and the *trader* dummy variable in the pooled regression model (Panel 4a) support the non-parametric finding that market professionals are less Bayesian than students.[36] The estimated marginal effect in the pooled model suggests that traders are 6% less likely to be Bayesian, and this effect is significant at the $p < 0.05$ level.

Despite the noisier environment (fewer professionals are Bayesian), market professionals and students choose the correct underlying state at similar rates. Indeed, using a Mann-Whitney U-test, we find that we cannot reject the homogenous null that success rates are similar at conventional levels ($p = 0.29$), leading to the result that earnings are similar across the subject pools. To dig a level deeper into this finding, we estimate a model similar in spirit to equation (2.4), but make the dependent variable *win* be dichotomous and equal to unity (zero) if the individual chooses correctly (incorrectly).

We include an additional independent variable, *round*, to identify learning during the course of the session; *round* is a time trend and increases from 1 to 15 within a session.[37]

---

[35] The alternative approach of computing the marginal effects for each observation and taking the means yields very similar results. Results are also robust to the inclusion of a time trend for round or time dummies (categorical time dummy variables for each round of play). We discuss our evidence of learning further, below.

[36] A Chow test rejects the null hypothesis of no differences across the subject pools at the $p < 0.01$ level.

[37] We test several specifications of the model for learning and find no such effect.

The empirical results summarized in Table 2.5 support the nonparametric finding concerning earnings and provide more formal evidence of the second half of Result 1. In particular, the *trader* variable in Table 2.5, Panel 5a is not significantly different from zero at conventional levels (p = 0.27). This result suggests that traders and students choose the correct urn at similar rates. The two groups differ in their temporal play, however, as evidenced by the significant (insignificant) and positive marginal effect of *round* for the traders (students), consistent with learning effects among traders.

Besides providing empirical support for Result 1, the models in Table 2.4 and Table 2.5 reveal some of the important effects of the other independent variables. For example, the *diff* and *heurist* coefficient estimates in the pooled model of Table 2.4 indicate that a marginal change in the posterior probability has a large positive effect (66%), while decisions in the *counting rule* sequences are 23% less likely to be Bayesian than those in which counting and Bayesian posterior imply the same result. Similar insights arise when we split the sample by subject type, as summarized in Panels 4b and 4c of Table 2.4. In addition, the effect of *diff* is statistically significant for both subject pools in the Table 2.5 *win* models.

Interestingly, urn symmetry, as captured by the *sym* dummy variable, is not significant for the market professionals in either model, implying that, for the traders, the difference across urn types is captured by the counting rule distinction. In contrast, the urn difference has a significant influence on students, who are much more likely to be Bayesians in the symmetric treatment (see Table 2.4, Panel 4b).

Table 2.4: Bayesian Decisions- Probit Model

Dependent variable: *baye*	4a: Pooled Model n = 1,647			4b. Student Model n = 822			4c. Market Professionals Model n = 825		
	Pr(*baye*=1)=0.818			Pr(*baye*=1)=0.868			Pr(*baye*=1)=0.772		
Ind. Variables:	Marginal Effect	z stat	P>|z|	Marginal Effect	z stat	P>|z|	Marginal Effect	z stat	P>|z|
Diff	0.655	5.53	0.000	0.769	5.06	0.000	0.546	3.11	0.002
Heurist	-0.232	-4.95	0.000	-0.161	-2.51	0.012	-0.284	-4.47	0.000
Gain	-0.030	-1.13	0.259	-0.060	-2.19	0.028	0.015	0.34	0.737
Sym	0.102	3.64	0.000	0.145	4.88	0.000	0.037	0.79	0.430
Trader	-0.060	-2.32	0.020	-	-	-	-	-	-
order_2	-0.023	-0.66	0.507	0.019	0.54	0.590	-0.084	-1.42	0.157
order_3	-0.041	-1.06	0.291	0.017	0.42	0.673	-0.120	-1.86	0.063
order_4	-0.120	-2.91	0.004	-0.052	-1.13	0.261	-0.205	-3.12	0.002
order_5	-0.035	-0.95	0.343	0.017	0.46	0.649	-0.107	-1.71	0.087
order_6	-0.040	-0.83	0.408	-0.035	-0.56	0.577	-0.080	-1.05	0.294
	Log Likelihood: -766.49, Wald $\chi^2_{(10)}$=141.03, Prob > $\chi^2_{(10)}$ = 0.000			Log Likelihood: -328.95, Wald $\chi^2_{(9)}$= 87.77, Prob > $\chi^2_{(9)}$ = 0.000			Log Likelihood: -427.28, Wald $\chi^2_{(9)}$= 68.91, Prob > $\chi^2_{(9)}$ = 0.000		

The dichotomous dependent variable in all three probit models (pooled, student, and market professional) is coded one for a decision consistent with the Bayesian posterior and zero otherwise. Independent variables include *diff*, which is $\left|prob\ (urn = A) - 0.5\right|$, where $prob\ (urn = A)$ is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain*, *sym*, and *trader* (in the case of the pooled model) are dichotomous and distinguish the treatments. *Heurist* is a dummy variable equal to one for the noncounting rule sequences and zero for all others. *Order_x (where x=2,..6)* is a categorical variable indicating where in the round of play the decision was made. The Wald statistic tests the null hypothesis that all coefficients are zero.

Table 2.5: Winning Decisions - Probit Model

Dependent variable: *Win*	5a: Pooled Model n = 1,647			5b. Student Model n = 822			5c. Market Professionals Model n = 825		
	Pr(*win*=1)=.702			Pr(*win*=1)=.729			Pr(*win*=1)=.670		
Ind. Variables:	Marginal Effect	z stat	P>|z|	Marginal Effect	z stat	P>|z|	Marginal Effect	z stat	P>|z|
Diff	1.070	7.85	0.000	1.039	5.51	0.000	1.053	5.55	0.000
Heurist	-0.008	-0.19	0.846	-0.009	-0.14	0.886	-0.050	-0.31	0.754
Gain	0.083	2.65	0.008	0.113	2.52	0.012	0.158	1.26	0.207
Sym	-0.028	-0.81	0.418	-0.013	-0.31	0.759	-0.122	-0.86	0.388
Trader	-0.033	-1.10	0.272	-	-	-	-	-	-
Round	0.005	1.95	0.051	0.003	0.68	0.499	0.022	2.06	0.040
order_2	0.008	0.22	0.824	0.014	0.29	0.774	0.008	0.05	0.960
order_3	0.019	0.49	0.627	0.016	0.30	0.767	0.073	0.45	0.656
order_4	0.048	1.26	0.206	-0.001	-0.02	0.988	0.293	1.79	0.074
order_5	0.057	1.53	0.126	0.081	1.64	0.102	0.010	0.62	0.534
order_6	0.063	1.33	0.184	0.046	0.63	0.528	0.192	0.96	0.337
	Log Likelihood: -964.06, Wald $\chi^2_{(11)}$ = 99.52, Prob > $\chi^2_{(11)}$ = 0.000			Log Likelihood: -462.25, Wald $\chi^2_{(10)}$ = 54.39, Prob > $\chi^2_{(10)}$ = 0.000			Log Likelihood: -498.55, Wald $\chi^2_{(10)}$ = 48.30, Prob > $\chi^2_{(10)}$ = 0.000		

The dichotomous dependent variable in all three probit models (pooled, student, and market professional) is coded one for a decision that correctly predicts the underlying state and zero otherwise. Independent variables include *diff*, which is $\left|prob\left(urn = A\right) - 0.5\right|$, where $prob\left(urn = A\right)$ is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain*, *sym*, and *trader* (in the case of the pooled model) are dichotomous and distinguish the treatments. *Heurist* is a dummy variable equal to one for the noncounting rule sequences and zero for all others. *Round* represents a time trend that increases from 1 to 16 with each completed play of the cascade game. *Order_x (where x=2,..6)* is a categorical variable indicating where in the round of play the decision was made. The Wald statistic tests the null hypothesis that all coefficients are zero.

A final important difference across subject pools is that the *order_x* variables indicate a decline in Bayesian behavior among market professionals who choose in the third through fifth positions. The magnitude of the effect is rather large, having from one-third to two-thirds of the effect of the counting rule sequences as represented in the *heurist* variable (Table 2.4, Panel 4c). In contrast, the students show no such effect. The behavior reflected in this finding is consistent with the idea that the market professionals recognize that no new additional information is added by choices once a herd has been formed.

Given the significance of the *diff* and *heurist* variables, we explore the individual data further in a QRE model, which examines the degree to which incentives affect error rates in decision making. Following Anderson and Holt (1997), we focus on data from our symmetric sessions and make use of the QRE model developed by McKelvey and Palfrey (1995, 1998); (see also Goeree et al. (2004), Goeree, Holt, and Palfrey (2005)). The QRE model assumes that the probability of choosing an urn is increasing in its expected value. Given the positive and significant coefficient on the *diff* variable in Table 2.4, the usefulness of such a model for both the students and market professionals appears evident. For parsimony, we reserve detailed discussion of the QRE model for the appendix, however, we briefly describe the results below.

Table 2.10 in Appendix 2.1 reports estimates of the lambda parameter in the QRE model. The lambda parameter indicates the extent to which noise affects decision outcomes; as $\lambda \to \infty$, the choice converges to the Bayesian outcome; as $\lambda \to 0$, the decisions become purely random. Significant differences in lambda across the subject pools are observed at choice orders one, two, and five as reflected in the *p*-values in column "p" of Table 2.10. Particularly notable is the difference at choice order two, where the students exhibit few errors. The differences in noise in the first two choice orders lead to quite different behaviors in choice order three, despite the fact that estimates of the lambda parameter are indistinguishable.

The lambda estimates imply that the two subject pools have similar deviations from Bayesian rationality at choice order three. Thus, the market professionals' tendency to rely on their own signal due to errors in earlier rounds is as rational as the students' decision to ignore theirs and join the cascade. Table 2.11 clarifies the meaning of this result by examining in detail the impact of the noisy decision process on revealed public information and choice probabilities for the first three rounds of play. For comparison, we present the posteriors and choice probabilities assuming a perfect Bayesian equilibrium, as well as the actual individual decisions.

Consider the posterior probability for choice order three in Table 2.11, the first choice at which a cascade may form in the symmetric treatment, when the signal history is AAb (or BBa). In this case, the posterior probability of urn A has dropped from 0.67 for the most likely urn to 0.51 (0.59) for the market professionals (students). Thus, while ignoring one's private information is optimal for both groups, the noise in prior decisions dilutes the strength of the signals, with the market professionals facing essentially a random choice. The probability that urn A is chosen is 0.54 (0.83) for the market professionals (students). The differences across the sequences in choice order three highlights the fact that noise in the decision-making process dilutes the value of the public signal.

Despite the evidence from the QRE estimation of the noisier environment for the market professionals we find that the two groups do not differ significantly in their earnings,. Further exploration into this observation leads to the following two results:

Result 2a: In aggregate, the rate of cascade formation is not significantly different for students and market professionals; however market professionals enter into fewer reverse cascades in the asymmetric treatments.

Result 2b: Market professionals are better able to discern the quality of the signal associated with other players' announcements than are students.

Evidence in favor of Results 2a and 2b follows from both nonparametric and parametric statistical tests. Even though the rate at which cascades are realized is roughly 60% for the students and only 50% for the market professionals (see Table 2.3, Panel A), using a Mann-Whitney test the homogeneous null cannot be rejected at conventional levels (Mann-Whitney p=0.33).

While the rate of cascade formation indicates that there is only weak evidence that students enter into a greater number of cascades than do professionals, there are significant differences across subject pools in the rate of cascade formation in the asymmetric urn treatment. Table 2.3, Panel A reveals that in the asymmetric treatment only 12% (8 of 66) of the cascades entered by market professionals are reverse cascades. This is roughly half of the rate observed for students (25 of 99), a difference that is statistically significant at the $p < 0.05$ level using a Mann-Whitney test.

To complement these nonparametric insights, we estimate models similar to equation (2.4), but set the dependent variable equal to one when a cascade is formed and zero otherwise. To conserve space, we do not formally tabulate these results since they reinforce the nonparametric insights gained above. We find that in the model that pools the symmetric and asymmetric data cascade formation is similar across the students and market professionals. When we focus, instead, on reverse cascades and use only the data from the asymmetric urn treatments, we find that students enter significantly more reverse cascades than do professionals.

These results cannot be explained by our model of decision making based on posterior probabilities derived from signals and actions. We therefore investigate the hypothesis that market professionals use auxiliary information that the students ignore in order to avoid reverse cascades. To do so, we augment the cascade formation model discussed above by considering whether subjects use information specific to individuals selecting prior to them in the current round.

Specifically, we construct two variables that each provides an indication of the Bayesian decision making of subjects who preceded each player in a particular round: *othb_max* (*othb_min*) measures the extent of previous Bayesian decisions by the most (least) Bayesian players. For example, for player *i* whose choice order is *x* in round *t*, we calculate *othb_min* as

$$(2.5) \quad othb\_min_{it}^{x} = \min\left[\frac{\sum_{t=1}^{t-1} baye_{jt}}{t-1} \quad \forall\ j : x_{jt} < x_{it}\right].$$

In this case, the proportion of Bayesian decisions for the individual with the lowest proportion among all *j* agents preceding the current decision maker is used as the independent variable, although the empirical results are robust to other specifications including replacing the min operator with the mean or max. In the case of *othb_max*, we simply replace “min” with “max” in equation (2.5). Note that these variables are calculated for each *t* (round of the game) so that they include only those decisions that have already occurred. The variables *diff*, *heurist*, and *gain* are also included, and are defined as in the previous models.

Empirical results are presented in Table 2.6. Since the results across models yield similar insights concerning the nature of interpreting signals, we focus on the *othb_min* results. Although *othb_min* is insignificant in the pooled specification in Panel 6a, this result masks a difference in how the two subject pools respond to the announcements of others. Results in Table 2.6, Panel 6c suggest that cascade formation for the market professionals is significantly and substantially associated with the *quality* of the others' signals. The marginal effect of a higher minimum in the preceding players' share of Bayesian decisions is 47%, which is the largest of the variables that are statistically significant and is an indication of the impact of the inferred signal quality on the willingness to make a decision that relies on others. This variable is significant and negative in the student sample (Panel 6b).

Table 2.6: Cascade Formation: Probit Model

Dependent variable: *cascade*	6a: Pooled Model n = 416			6b. Student Model n = 226			6c. Market Professionals Model n = 190		
	Pr(*cascade*=1)=.588			Pr(*cascade*=1)=.676			Pr(*cascade*=1)=.493		
IND. VARIABLES:	MARGINAL EFFECT	Z STAT	P>|Z|	MARGINAL EFFECT	Z STAT	P>|Z|	MARGINAL EFFECT	Z STAT	P>|Z|
DIFF	0.861	0.75	0.453	-0.136	-0.08	0.939	1.689	1.13	0.259
Othb_min	-0.014	-0.09	0.924	-0.572	-2.54	0.011	0.469	2.11	0.035
Heurist	-0.354	-4.35	0.000	-0.331	-2.49	0.013	-0.394	-3.95	0.000
Gain	0.133	1.67	0.095	0.021	0.19	0.846	0.120	1.12	0.261
Sym	0.126	1.03	0.303	0.389	4.09	0.000	-0.200	-1.16	0.246
Trader	-0.111	-1.44	0.151	-	-	-	-	-	-
order_2	-0.146	-1.13	0.260	0.064	0.39	0.696	-0.332	-2.33	0.020
order_3	0.078	0.72	0.469	0.167	1.24	0.214	0.039	0.24	0.810
order_4	0.042	0.39	0.696	0.048	0.33	0.744	0.066	0.44	0.658
order_5	0.233	2.35	0.019	0.169	1.21	0.225	0.308	2.13	0.033
	Log Likelihood: -245.22, Wald $\chi^2_{(10)}$ = 46.24, Prob > $\chi^2_{(10)}$ = 0.000			Log Likelihood: -125.20, Wald $\chi^2_{(9)}$ = 27.78, Prob > $\chi^2_{(9)}$ = 0.001			Log Likelihood: -111.06, Wald $\chi^2_{(9)}$=27.42, Prob > $\chi^2_{(9)}$ = 0.0012		

The dichotomous dependent variable in all three probit models (pooled, student, and market professional) is coded one for a cascade decision and zero otherwise. Independent variables include *diff*, which is $\left|prob\,(urn = A) - 0.5\right|$, where $prob\,(urn = A)$ is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain* and *trader* (in the case of the pooled model) are dichotomous and distinguish the treatment/subject type. *Othb_min* is the proportion of Bayesian decisions by the individual with the lowest proportion among all agent's preceding the decision maker, and is calculated in each round of the game to include only those decisions that have already occurred. *Heurist* is a dummy variable equal to one for the noncounting rule sequences and zero for all others. *Order_x* is a categorical variable indicating where in the round of play the decision was made. Note: Because the othbys variable is not applicable for those in the first round or first in choice order in subsequent rounds (they do not observe others' decisions in the current round), these observations are excluded. This results in the exclusion of 25 of the 441 potential cascades. The order_6 dummy variable is also excluded and choice order two serves as the baseline to which others are compared. The Wald statistic tests the null hypothesis that all coefficients are zero.

Using *othb_max* in the regression yields an insignificant effect for the students, while the market professionals again respond positively, with a marginal effect of 57% (detailed results omitted).[38] We therefore conclude that the market professionals make better use of available public information, incorporating evidence on others' rationality in their decision making in a way that is payoff relevant.[39] Note also that, in contrast with what we found with respect to all decisions (Table 2.4), the *diff* variable is not significant for either group when we restrict our attention to the subset on cascade formation.

One may wonder whether the result on signal quality is due to market professionals having a greater level of previous interaction with one another than students, or, alternatively, whether there is evidence of learning in the experiment. To explore this issue, we again examine changes in behavioral patterns during an experimental session. The evidence is consistent with the view that market professionals learn over these 15 rounds. Comparing behavior from the first and last three rounds of a session, we find that market professionals: a) significantly reduce the rate at which they join reverse cascades (from 13% to 2%), and b) increase the rate at which they join cascades with good outcomes (from 24% to 46%). Both results are statistically significant in probit specifications that include the cascade type as the dependent variable and the temporal variable along with the control variables as independent variables (full

---

[38] We estimated six models that included variables designed to measure the quality of previous agent's decision-making on cascade formation. In addition to the three that used other Bayes variables (*othb_min*, *othb_max*, and *othb_mean*) we considered whether individuals who had previously revealed their private signal were followed when cascades were possible. These *other_reveal* models also tested the min, max, and mean operators. In all six cases market professionals followed those with higher levels of reliability into cascades. Among the students, in five of six cases there was no significant effect of signal quality, with *othb_min* the sole exception as reported in Table VI, Panel 6b.

[39] Support for the significant differences between subject pools found in the parametric results is also found in nonparametric (Mann-Whitney) tests.

results omitted to conserve space). By contrast, there are no significant changes in the rate of cascade formation for either type of cascade for the student subjects.

Our final insight concerning the comparison between students and professionals concerns the domain of earnings of the game:

Result 3: Bayesian behavior of the student population is affected by whether earnings are in the gain or loss domain, while market professionals are unaffected.

Summary evidence in favor of this result can be found in Table 2.3, Panel A, where we observe that professionals exhibit a similar degree of Bayesian decision making across the gain and loss domains (roughly 75%), whereas for students Bayesian play increases in the loss domain. For example, considering the asymmetric treatments, we find that a Mann-Whitney test indicates that college students are less Bayesian in the gain treatment than in the loss treatment, while market professionals are unaffected by the domain of earnings (students: $p < 0.08$; traders: $p = 0.61$).[40]

Empirical estimates in Table 2.4 provide additional evidence of this result. In the pooled data (Panel 4a), the dummy variable *gain* is not significant at conventional levels, and it remains insignificant for the market professionals' specification (Panel 4c). For the students, however, the parameter estimate is both significant (p=0.028) and negative, indicating a 6% increase in Bayesian behavior in the loss domain. This result is consistent with the notion that, for the student population, losses loom larger than gains. This result is consonant with results in List (2003, 2004), who explore loss aversion in a

[40] Due to the small number of sessions at the individual treatment level, p-values for the Mann-Whitney test are reported for observations aggregated at the individual participant level.

much different environment. Nevertheless, consistent with the notion that repetition might attenuate such anomalies (see, for example, Knez et al. (1985); Coursey et al. (1987)), analysis of the data from the student sessions provides some evidence that the effect of the domain is mitigated via repetition.

While Results 1-3 highlight differences between the professional and student subjects, we also find important differences within the group of market professionals that are relevant for understanding their decision processes. We supplement our data with a survey implemented at the end of the experimental session. Upon exploring these data more closely, we find:

Result 4: Behavioral differences exist within the professional subject pool.

Evidence of this result can be obtained by augmenting equation (2.4) using the additional demographic data collected from the CBOT floor personnel after the experiment. We focus on data collected from a group of 28 of the 55 traders who reported information on *intensity* (the average number of contracts traded per day), *gender* (one for female, zero otherwise), *yrs* (years of experience), *income,* and *overnight*, (a dichotomous variable that equals one if the trader takes overnight positions and zero otherwise). Panel A. of Table 2.7 reports on the Bayesian decision making and Panel B reports on the cascade formation for these traders.

Table 2.7: Bayesian and Cascade Behavior of Traders

	7a: Trader subset of CBOT Market Professionals n = 227 Dependent Variable: *baye* Pr(*baye*=1)=.745			7b.Trader subset of CBOT Market Professionals n = 66 Dependent Variable: *cascade* Pr(*cascade*=1)=.388		
Ind. Variables:	Marginal Effect	z stat	P>|z|	Marginal Effect	z stat	P>|z|
Diff	0.467	1.34	0.181	3.710	0.65	0.517
Heurist	-0.391	-3.05	0.002	0.001	0.05	0.997
Gain	-0.023	-0.24	0.81	0.163	0.59	0.552
Sym	0.009	0.08	0.938	0.498	0.79	0.432
order_2	0.061	0.64	0.523	n/a	n/a	n/a
order_3	-0.029	-0.28	0.778	-0.185	-0.48	0.629
order_4	-0.095	-0.86	0.392	-0.219	-0.68	0.498
order_5	-0.024	-0.22	0.828	0.474	1.22	0.221
order_6	0.089	0.67	0.504	0.229	0.36	0.721
Intensity	0.004	2.44	0.015	-0.029	-1.91	0.056
Gender	0.069	0.58	0.561	-0.955	-0.2	0.838
Experience (yrs)	-0.001	-0.18	0.859	0.011	0.34	0.735
Income	0.013	0.55	0.582	0.394	1.58	0.115
Overnight	-0.173	-2.03	0.042	-0.804	-2.24	0.025
	Log Likelihood: -93.93, Wald $\chi^2_{(14)}$ = 52.74, Prob > $\chi^2_{(14)}$ = 0.0000			Log Likelihood: -15.62, Wald $\chi^2_{(13)}$ = 34.20, Prob > $\chi^2_{(13)}$ = 0.0011		

The dichotomous dependent variable in Panel A is coded one for a decision consistent with the Bayesian posterior and zero otherwise. For Panel B cascade formation is indicated by a one and cascade failure by a zero. Independent variables include *diff*, which is $\left|prob\,(urn = A) - 0.5\right|$, where $prob\,(urn = A)$ is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain* and *sym* are dichotomous and distinguish the treatments. *Heurist* is a dummy variable equal to one for the non-counting rule sequences and zero for all others. *Order_x (where x=2,..6)* is a categorical variable indicating where in the round of play the decision was made. *Intensity* reflects the level of trading intensity among participants, measured as the number of contracts traded per day. *Gender* is one for female and zero for male. *Experience* (years), *income* (dollars), and *overnight* (one for holding overnight positions, zero for daytrader) are additional control variables.

Concerning Bayesian decision making, we find that *diff* is not significantly different from zero. Indifference to the magnitude of the posterior, for the Bayesian models, does not occur elsewhere in our study, and as we discuss previously is consistent with Bayesian rationality and inconsistent with theories of decision error. Variables that are significant include *heurist*, *intensity*, and *overnight*. As with the previous results reported in Table 2.4, *heurist* has a strong negative effect (-39.1%). Trading intensity increases Bayesian behavior slightly (0.4%) and overnight trade has a significantly negative impact on the rate of Bayesian decision making (-17.8%). The probit estimates in Panel B reveal that day traders are much more likely to join an informational cascade, as are traders with lower trading intensity, with marginal effects of -80% on *overnight* and -2.9% on *intensity*.

For those making consequential trading decisions, the link between trading intensity and Bayesian rationality is consistent with the empirical results of Locke and Mann (2005), Genovese and Mayer (2001), and List (2003, 2004), who find similar results in diverse settings that include financial, housing, and memorabilia markets. We believe that the result on trading style is novel, and we offer some thoughts on its implications in the discussion section below.

Results 1-4 highlight differences in cascade formation and Bayesian decision making across subject types, and include the exogenous alteration of signal strength due to urn type through the *heurist* variable. Our final result looks more closely at the impact of signal strength:

Result 5: Deviations from Bayesian norms are greatest when the counting rule and Bayesian updating make different predictions.

Our probit specifications reveal that when counting and Bayesian rationality yield different predictions, both market professionals and students are less Bayesian. Table 2.8 presents all of the observed signal patterns for the asymmetric treatment. Those in which the counting rule and Bayesian posteriors yield different predictions are in bold type. Statistical tests confirm what a visual scan of the data suggests: Bayesian behavior is significantly reduced in the non-counting rule sequences.[41] In fact, the four non-counting rule sequences have lower rates of Bayesian decision making than any of the other sequences, despite the fact that others have smaller *diff* values.

Table 2.8: Posterior Probability Urn is *A* and Proportion of Bayesian Decisions

a\ b	0	1	2	3	4	5	6
0	0.50	0.33	0.20	0.11	0.06	0.03	0.02
		0.85	*1.00*	*1.00*	*1.00*	*0.89*	*1.00*
1	0.55	0.38	0.23	0.13	0.07	0.04	
	0.76	*0.56*	*0.72*	*0.64*	*0.95*	*0.63*	
2	0.59	**0.42**	0.26	0.15	0.08		
	0.87	***0.46***	*0.63*	*0.69*	*0.79*		
3	0.63	**0.46**	**0.30**	0.18			
	0.84	***0.30***	***0.52***	*0.58*			
4	0.67	0.51	**0.34**				
	0.76	*0.76*	***0.29***				
5	0.71	0.55					
	0.87	*0.78*					
6	0.75						
	0.80						

The amount of information associated with urns A and B are given in the first row and first column, respectively. The pairs of numbers within an (a,b) pair represent the Bayesian posterior (upper number) and the proportion of Bayesian decisions (*lower number*). Those in bold type are the sequences in which counting and the Bayesian posterior make different predictions. Thus (2,1) has a posterior probability of 42% that the urn is *A* (diff=0.08). Forty-six percent made the Bayesian decision in this case. By contrast the (2,0) sequence (in which diff=0.09) has a posterior probability of 0.59, and 87% of those decisions were Bayesian.

[41] We use a Wilcoxon matched pairs test with the variable of interest equal to the proportion of Bayesian decisions aggregated at the session level. The *diff* variable for the *counting rule* sequences is in the range from 0.0 to 0.2, and all other sequences with *diff* variables in this range are included for the paired comparison. Using data from the 12 asymmetric sessions we find that the *counting rule* sequences reflect less Bayesian decision making despite roughly equivalent diff scores at $p<.01$.

Figure 2.1 illustrates this insight by presenting the proportion of Bayesian decisions for all observed histories of play as a function of the posterior probability.

Figure 2.1: Counting Rule Heuristic, Signal History and Bayesian Behavior

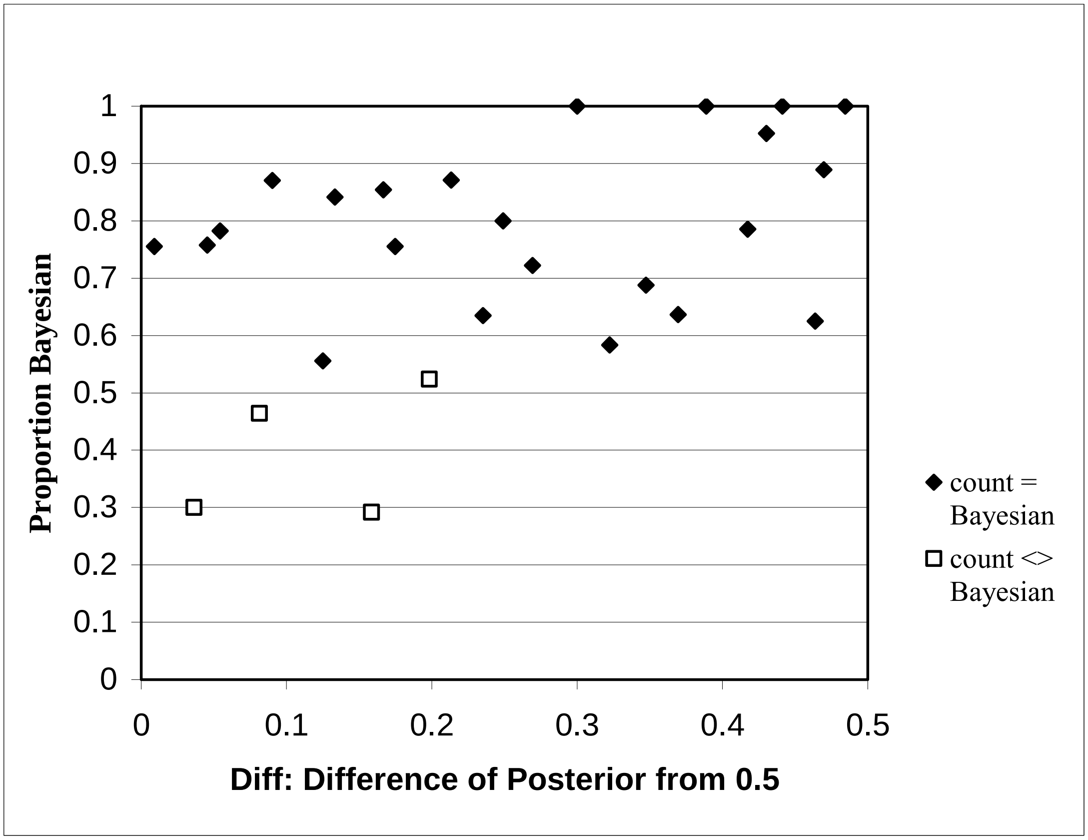

The proportion of Bayesian decisions for every realized posterior probability is presented as a data point. The choice histories in which the counting rule and Bayesian posterior yield different predictions are presented as open squares. All other sequences are presented as black diamonds. Note that the sequences in which Bayesian behavior and the counting rule heuristic make different predictions have a uniformly lower proportion of Bayesian decisions than the others.

The non-counting rule sequences (square entries) are uniformly lower than the other choice histories, represented as black diamonds. Compiling the results from Figure 2.1, we find that Bayesian behavior occurs at a rate of 44% in the non-counting rule choice histories and at a rate of 81% in the remaining choice histories in the asymmetric treatments. There is an important difference in the rate of Bayesian behavior in non-counting rule sequences that depends on whether one's decision involves choosing to join a cascade. The difference is best explained by considering whether individuals rely on their private signal. Restricting attention to non-counting rule sequences, we find that individuals are Bayesian in 31% of the cases when the decision involves choosing to enter a cascade. Therefore, 69% follow their own signal. By contrast 74% of decisions are Bayesian when there is no potential cascade and the decision is consistent with one's private information (see Table 2.9). Thus, when the signal history requires that Bayesian agents ignore their own signal, agents generally fail to do so. As a result, the failure of cascade decisions implies that 69% rely on their own information – a result statistically indistinguishable from the 74% who rely on their own signal when doing so is optimal. We conclude that for the noncounting rule sequences, a Bayesian perspective provides a less accurate description of decision making than the simple rule of using private information.

Table 2.9: Bayesian Behavior According to Cascade Potential - Asymmetric Treatments

	No Potential Cascade	Potential Cascades	Total
All n = 1,005	.82 541/657	.48 166/348	.70 707/1,005
Noncounting Rule n = 843	.83 512/618	.57 128/225	.76 640/843
Counting Rule n = 162	.74 29/39	.31 38/123	.41 67/162

The proportion of Bayesian decisions both when a cascade is possible and when one is not for both counting rule and non-counting rule sequences in the asymmetric treatments are provided in the table. When there is no potential cascade the proportion of Bayesian decisions (.74) is the proportion in which one follows the private signal. When there is a potential cascade (1-.31=.69) is the proportion of decisions that follow the private signal.

2.4 Discussion

Our cascade game data yield interesting evidence of heterogeneity both across the two subject pools and within the market professional group. Simple measures of performance indicate that the students outperform the market professionals. Controlling for learning about signal quality, however, makes clear that the market professionals use a more sophisticated decision process, more finely parsing the quality of public information and relying on their own signal more frequently. Within the market professional group, trading style has a strong effect on behavior, with those taking overnight positions entering cascades much less frequently.

We view these results as having potentially interesting implications for financial markets, although care must be taken with the interpretation, in part because of the fixed payoff that subjects received in our experiment.[42] However, fixed prices are not irrelevant in financial markets as variability in order size means that prices need not change with each transaction. Thus, it is reasonable to study cascade decisions occurring at a constant price as well as those that lead to a change in price.

We believe it is plausible that the heterogeneity among traders regarding cascade formation may be related to differences in their trading practices, including those around fixed prices. Local floor traders who do not take overnight positions typically specialize as market makers and are more likely to face situations in which herding, including herding at a constant price, is part of their trading practice. This type of herding may

[42] There is a long and important debate on the relevance of cascade models for financial markets (Vives (1996)). Avery and Zemsky (1998) show that the introduction of variable prices to the BHW model can eliminate informational cascades (herding in their terminology) under certain conditions. Lee (1998), Chari and Kehoe (2004) and Cipriani and Guarino (2005a) demonstrate the potential for informational cascades in the variable price setting by introducing transaction costs, endogenous timing, and preference heterogeneity.

occur, for example, when several floor traders each take a portion of a large institutional order. Manaster and Mann (1999) provide evidence that market makers are willing to give up their advantage in executions, narrowing or eliminating the bid-ask spread, when they have an informational advantage over the outside order. If information is dispersed among traders heterogeneously, the situation is similar to the cascade environment we study here. A crucial difference is that timing and transaction size in the market is endogenous, and ultimately, of course, prices do change.[43]

Avery and Zemsky (1998) introduce flexible pricing into the BHW model and find that for cascades to form, the *value uncertainty*, which we implement in our experimental protocol, needs to be accompanied by *event uncertainty* (the possibility of a change in asset value) and *composition uncertainty* (which implies that the distribution of trader types is not common knowledge). Our results on the discernment of the quality of public announcements suggest that experienced professionals are better able to estimate the composition of the distribution of trader types, and so may act to mitigate price bubbles and crashes.[44] Clearly, while additional research regarding the impact of trader specialization is warranted, our findings highlight the benefits of controlled experimentation with nonstudent subject pools.

---

[43] One mechanism through which cascades might arise is, in the jargon of the trading floor, when local traders "lean on" large orders by trying to enter the market on the same side and at the same price. Locals who trade alongside an institutional order accumulate a position knowing that they can transact with the institution and avoid a loss. The decision process associated with deciding to trade with the institution has the character of a fixed price cascade. In the context of option markets Berkman (1996) discusses how market makers supply liquidity in the presence of large fixed price orders. Chamley and Gale (1994) introduce endogenous timing in a cascade model that predicts the least informed would trade later, and potentially face adverse prices.

[44] Drehmann, Oechssler, and Roider (2005) test experimentally a version of the Avery and Zemsky (1998) model that omits event and composition uncertainty and find behavior fairly consistent with its predictions, though subject to decision error and contrarian behavior.

We believe that our findings may also shed light on other types of cascade behavior. Consider Welch's (1992) interesting model of initial public offerings (IPOs), for example, which addresses cascade formation at a fixed price due to regulatory requirements for IPOs. Welch finds that issuing firms have an interest in pricing to generate an informational cascade in order to increase the probability of a successful offering. Our results that emphasize the potential for cascade fragility arising from variation in the ability to interpret signal quality may be important in this context. One possible implication is that when underpricing of offerings is optimal in the Welch model, heterogeneity in signal strength and interpretation might play an instrumental role since reverse cascades in which no investment occurs will be fragile. The welfare implications, however, are not immediately obvious given that the resulting cascades are of shorter duration. Further, the importance of the effect may differ across firms or industries depending on the economies of scale of the investment and thus the need to have full or only partial subscription (Welch (1992), p. 709).

Both the differences due to specialization and the heterogeneity in signal quality and processing abilities suggest fruitful directions for future research. How the specialized skills of market participants interact in price discovery could be explored in experiments that move towards a full market setting, but in which liquidity and informational conditions are varied in a controlled manner. A natural part of this research program would be to extend the current environment to study the impact of heterogeneity on the IPO model of Welch (1992). In a recent study that provides evidence from asset market experiments with student subjects Dufwenberg, Lundqvist, and Moore (2005) find that mixed experience levels can reduce the incidence of bubbles and crashes. Heterogeneous subject pools that include professionals would shed crucial light on this issue, and help to identify the mechanisms underlying cascade formation and fragility in settings that mix fixed and variable prices.

2.5 Concluding Comments

In this study, we introduce market professionals from the CBOT floor to a controlled experimental environment. Making use of information cascades games, we report several insights. While student subjects more closely follow Bayes' rule, they do not perform significantly better than the market professionals along the important dimension of earnings. This puzzle is explained by the fact that professionals are more sophisticated in their use of public information, as manifested over the course of the decision process: Market professionals are less Bayesian when making decisions later in the choice order in a cascade game, consistent with recognizing that the quality of initial announcements is variable, altering the payoffs of joining cascades.

While market professionals learn over the course of an experimental session to account for the quality of others' decisions, student subjects fail to do so. A further insight is that market professionals are consistent in behavior over the gain and loss domains, while in aggregate, students' behavior is consistent with the notion of loss aversion. Perhaps most provocatively for the operation of markets, we find an important heterogeneity among the market professionals that depends on their trading style. In summary, our data reveal that the decisions of market professionals are consistent with behaviors that may mitigate informational externalities in market settings, and thus reduce the severity of price bubbles due to informational cascades.

Besides revealing both positive and normative insights, our work also offers a methodological contribution. For example, it highlights the potential for experiments with students and professionals to be complementary inputs to research when field data is suggestive but inconclusive. Indeed, in transferring the insights gained in the laboratory with student subjects to the field, a necessary first step is to explore how market professionals behave in strategically similar situations. In this spirit, we focus on the representativeness of the sampled population to lend insights into which empirical results are similar across subject pools. A related issue concerns the representativeness of the environment, which also merits serious consideration. For example, before we can begin to make reasonable arguments that behavior observed in the lab is a good indicator of behavior in the field, we must explore whether the other dimensions of the laboratory environment might cause differences in behavior, including the abstract task, the stakes, the good, and the institution. While our research represents a necessary first step in the discovery process, we hope that future efforts will explore more fully other potentially important dimensions of the controlled laboratory experiment.

Appendix 2.1: QRE Estimation Results

Results in Table 2.4, indicate that deviations from a perfect Bayesian equilibrium are associated with the payoff consequences of a decision. This result leads us to estimate the quantal response equilibrium (QRE) developed by McKelvey and Palfrey (1995, 1998). By accounting for how decision error affects those later in the choice order the QRE yields alternative measures of the public belief. While the QRE has been almost universally applied to experimental results on information cascades, there have been two significant criticisms of the model. Haile, Hortaçsu, and Kosenok (2004) argue that when the assumption that errors are i.i.d. is relaxed, the QRE can rationalize choices that violate monotonicity.[45] Goeree, Holt, and Palfrey (2005), however, show that economically sensible properties, including monotonicity and responsiveness, can be obtained with less restrictive assumptions.[46]

A further criticism of the QRE questions the plausibility of the assumption that players have rational expectations about the others errors. Two approaches have been used test this hypothesis, yielding mixed results. In a two parameter model that separately measures the rationality of beliefs and actions, Goeree et al. (2004) find support for rational expectations. Kubler and Weizsacker (2004)) estimate a more complicated model that generalizes the belief parameter for different depths of reasoning and reject the hypothesis of rational expectations. To our best knowledge these models have not been tested on the same dataset, however an evaluation of the two approaches is in

[45] That is, the QRE can assign to outcomes with low payoffs probabilities that are arbitrarily close to one.

[46] Monotonicity implies that in a comparison across strategies those with higher payoffs are played with greater probability. Responsiveness implies that when the payoff of a given strategy increases the probability with which it is played does not decline (Goeree, Holt, and Palfrey (2005)).

process using the dataset from this paper (Alevy 2006). Preliminary results support the rational expectations hypothesis using the model of Goeree et al. (2004).

The current results are derived from the following model which retains the i.i.d assumption.[47] Let the probability of choosing urn A be given by

$$pr(c_i = A \mid H_i, s_i) = pr\left(\pi_i^A + \varepsilon_i^A > \pi_i^B + \varepsilon_i^B\right) = pr\left(\varepsilon_i > 1 - 2\pi_i^A\right),$$

where $\pi_i^A = pr(A \mid H_i, s_i) * \$W = \$W - \pi_i^B$, and $\varepsilon_i = \varepsilon_i^A - \varepsilon_i^B$. For comparability across subject pools we normalize earnings so that *$W* = 1 for both subject pools. If the errors have an extreme value distribution, then the conditional probability of the urn choice is given by the logistic choice rule

$$pr(c_i = A \mid H_i, s_i) = \frac{1}{1 + \exp\left(\lambda_i \left(1 - 2\pi_i^A\left(\lambda^{1 \ldots i-1}\right)\right)\right)}.$$

The lambda parameter indicates the extent to which noise affects decision outcomes. As $\lambda \to \infty$, the choice converges to the Bayesian outcome; as $\lambda \to 0$, the decisions become purely random. Note that the posterior probability that the urn is *A*, π_i^A, is a function of the lambda estimates from previous choice orders, with $\left(\lambda^{1 \ldots i-1}\right)$ representing the vector of previous estimates.

Embedding the earlier errors in the choice probabilities that follow is what generates interesting insights from this model. As shown below in a comparison of the error rates of students and market professionals, equal levels of rationality, as reflected in

[47] In addition to monotonicity and responsiveness, the i.i.d. assumption imposes strong substitutability and translation invariance. Strong substitutability requires that when the probability of a choice of one outcome increases the probability for all other outcomes declines. This is not objectionable in the two-state case considered here. Translation invariance implies that when payoffs are multiplied the probability of a choice does not change. Given that payoffs are constant within a session this assumption does not appear restrictive.

comparable magnitudes of λ_i can predict very different behaviors that depend on the errors in previous rounds.

In our estimation, we follow Anderson and Holt (1997) and focus on the symmetric data. The QRE results using these data are displayed in Table 2.10 and Table 2.11.

Table 2.10: Lambda estimate for Quantal Response Equilibrium, Symmetric Gain Treatment

Choice Order	M	C	p
1	4.59	7.12	0.094
2	4.56	27.75	0.012
3	8.67	8.62	0.505
4	3.90	4.99	0.258
5	2.48	6.34	0.026

Columns M and C report the lambda parameter for market professionals and college students. Column p reports the one-tailed p-value for the null hypothesis that the lambda parameter does not differ across the two groups. All lambda estimates differ significantly from zero.

Our results emphasize the fact that not only the numbers of each signal, but also the order in which they are revealed have an important impact on behavior. For example, note the posterior probability for order choice three in Table 2.10, the first choice where a cascade may form in the symmetric treatment, when the signal history is AAb.[48] In this case the posterior probability of urn A has dropped from 0.67 for the most likely urn to 0.51 (0.59) for the market professionals (students). Thus, while ignoring one's private information is optimal for both groups, the noise in prior decisions dilutes the strength of the signals, with the market professionals facing essentially a random choice with the probability that urn A is chosen being 0.539 (0.833) for the market professionals (students). In comparison the ABa sequence, which has an identical posterior probability

[48] Due to the symmetry of the game, in Table 2.10 the AAb sequence also includes the BBa results. All other symmetric choice sequences are treated similarly.

when there is no noise, the posterior 0.64 (0.65) for market professionals (students), and the optimal decision is made uniformly by both subject pools. This difference across the sequences in choice order three highlights the fact that noise in the decision making process dilutes the value of the public signal. In particular, the noise associated with the market professionals at choices 1 and 2 results in a dependence on the private signal at choice 3. The decisions revealing the private signal provide those that follow with a richer information set.

Table 2.11: Posterior Probabilities and Choice Probabilities with QRE Decision Error

Choice Order	History & Signal	Choice Probability $pr(c = A \mid H, s, \lambda)$ Bayes	QRE M	QRE C	Posterior Probability $pr(\omega = A \mid H, s, \lambda)$ Bayes	QRE M	QRE C	Actual Decisions M A	M B	M Share A	C A	C B	C Share A
1	A	1.00	0.82	0.92	0.67	0.67	0.67	37	8	0.82	43	4	0.92
2	Aa	1.00	0.91	0.99	0.80	0.76	0.78	23	3	0.89	27	0	1.00
2	Ab	0.00	0.36	0.15	0.50	0.44	0.47	4	15	0.21	3	17	0.15
3	AAa	1.00	0.99	0.99	0.89	0.81	0.85	14	0	1.00	17	1	0.94
3	AAb	1.00	0.54	0.83	0.67	0.51	0.59	7	6	0.54	12	0	1.00
3	ABa	1.00	0.94	0.92	0.67	0.65	0.64	10	0	1.00	8	0	1.00
3	ABb	0.00	0.05	0.03	0.33	0.32	0.31	1	7	0.13	0	9	0.00

Calculations are for the first three choices of the symmetric gain treatment for market professionals (M) and college students (C), with the choice probability and the posterior probability adjusted for decision error. For comparison, the probabilities assuming a perfect Bayesian equilibrium (Bayes) are also presented as are the actual decisions. Due to the symmetry of the treatment, the history and signal combination also represents its complement. For example the row reporting history and signal "ABa" also includes the "BAb" sequences.

Appendix 2.2: Experimental Instructions – SGS Treatment

Instructions:

In this experiment, you will be asked to decide from which of two urns balls are being drawn. We will begin by rolling a six-sided die. If the die roll yields a 1,2, or 3, we will draw from Urn A. If the roll of the die yields a 4,5, or 6, we will draw from Urn B. However, the roll of the die will be done behind a screen so that you will not know which urn has been chosen.

The urns differ in the following way:

Urn A (used if die is 1,2,or 3)	Urn B (used if die is 4,5,or 6)
2 Striped Balls	1 Striped Ball
1 Terp Ball	2 Terp Balls

Once an urn is determined by the roll of the die we will empty the contents of that urn into a container. (The container is always the same, regardless of which urn is being used.)

After the urn has been chosen each of you will come behind the screen one at a time and draw a ball from the container. The order in which you will draw has been determined randomly.
The result of your draw is your private information and **MUST NOT** be shared with other participants.

After each draw, we will return the ball to the container before making the next private draw. Each person will have one private draw, with the ball being replaced after each draw.

After each person has seen the results of their own draw, we will ask them to record the letter of the urn (A or B) that they think is more likely to have been used. When the first person to draw has indicated a letter, we will display that letter. After displaying the first person's decision, we will call out the next registration number, and the person with that number will draw a ball and record a letter (A or B). Again, their decision will be displayed on the overhead projector. This process will be repeated until everyone has made a draw and made a decision about which urn they believe is being used. After everyone has made a decision, the monitor will announce which of the urns was actually used. Everyone who chose the correct urn earns $1. All others earn nothing.

This session will consist of 15 periods of the procedure just described.

Now I will describe the use of the record sheet, which is at the back of these instructions.

The results for each period are recorded on a separate row on the record sheet. Period numbers are listed on the left side of each row. Next to the period number record your draw (S or T) in column "***Own Draw***". In columns "***choice1***" through "***choice10***" record each participants decisions (A or B) as they are displayed. (If there are less than 10 players the last choice columns remain blank.) This means that when you are asked to make a decision about which Urn is being used the decisions of participants who have drawn before you will be available. Write your decision in the appropriate column depending on the order in which you draw, and *circle your decision to distinguish it from other's decisions.*

When all participants have made their choices, the monitor will announce the letter of the Urn that was actually used. Record this letter in the column headed "***Urn***" for that period. If your circled decision matches the letter of the urn used, record your earnings of $1 in the "***Payoff***" column. If your choice does not match the urn used record your earnings of $0. You should keep track of your cumulative earnings in column "***Total Payoff***".

Before we begin we will conduct a demonstration. During the demonstration, the roll of the die and the draw of the ball from the container will be publicly visible. When we move to Period 1 the roll of the die will be visible only to the monitor, and the draw will be visible only to the monitor and the person called behind the screen. Remember that urn A contains 2 striped balls and 1 terp ball. It is used if the throw of the die is 1,2, or 3. Urn B contains 1 striped ball and 2 terp balls, and is used if the throw of the die is 4,5, or 6.

Before we begin, be sure that your registration number is on the Record Sheet.

Here is an overview of the procedure that will be followed in each period:

1. The monitor rolls the die to determine which urn is used and transfers balls from that urn to the container.

2. The monitor calls on a participant.

3. The participant goes behind the screen: (Be sure to bring your record sheet)
    a) Makes a draw from the container
    b) records the draw on record sheet in "***Own Draw***" column
    c) records urn choice *and circles their choice*

4. The monitor displays the participant's choice and the other participants record the urn choice on their record form.

5. Repeat steps 2 – 4 until all participants have made their choice.

6. The monitor reveals the urn used in that period by displaying the balls in the container.

7. Subjects record the urn used in that period and record their earnings, and their cumulative earnings.

Please refrain from conversation during all periods of play, and keep the information on your record sheet confidential.

3 Common Agency with Other-regarding Preferences

3.1 Introduction

The insights of agency theory have broadened as theoretical and empirical work has explored more realistic and complex environments, such as those with many agents, many tasks, and many principals. Because of the importance of asymmetric information and hidden action in principal-agent models, they are amenable to testing in laboratory environments, where information and preferences can be more readily controlled. Efforts in this area, as well as in simpler proposer-responder games have led to the creation of the sub-field of behavioral contract theory (Charness and Dufwenberg 2003).

In this chapter we extend the scope of behavioral contract theory by examining a model of common agency in which the agent performs multiple tasks. We consider the work "behavioral" since theoretical predictions are developed and tested that explicitly account for the role of inequity, intentions, and reciprocity on economic outcomes. Exploring the implications of distributional concerns and intentions has been identified as one area in which experimental methods can inform economic theory (Samuelson 2005). Agency theory is an area where this investigation is particularly relevant, since many principal-agent models include unequal bargaining settings similar to ultimatum games.

The paper proceeds by first formalizing in a single-principal setting a simple model that incorporates inequity-averse preferences. Contracts can be composed of piece rates, fixed fees, or some combination of the two instruments. The results are intuitive, with inequity aversion leading to more equal splits of the resulting surplus.[49] Extending

[49] The model formalizes much of the intuitive discussion by Anderhub, Gachter, and Konigstein (2002), whose experiment in a single principal setting is similar to our baseline treatment. The theoretical

the model to the case of common agency, results in interesting complications. In contrast to the single-principal case, principals consider the agent's level of inequity aversion when choosing incentives, which may be higher or lower than under the self-interested baseline. In addition to considering the impact of distributional concerns on efficiency, an alternative model of the common agency problem is developed that has differing implications for the potential efficiency of the contracting environment. The theoretical results are tested experimentally in one- and two-principal settings.

The chapter continues in section 3.2 with a review of existing models that identify the diversity of the settings in which common agency comes into play, as well as previous experimental work that informs the current study. Section 3.3 presents the theoretical framework and section 3.4 summarizes the theoretical results. Section 3.5 introduces the experimental protocol and 3.6 presents the experimental results. Section 3.7 concludes.

3.2 Motivation and Literature Review

The analytical lens of common agency has been used to study a variety of economic and political processes as disparate as manufacturer/wholesaler relationships, joint federal/state regulation of firms, privatization, aid effectiveness, debt and equity as tools for controlling corporate behavior, and separation of powers in a constitutional democracy. Common to these diverse issues is the question of how the addition of one or more principals to an agency relationship affects the provision of incentives, agent effort, and the allocation of the surplus among players. Interesting questions arise when the

predictions of Anderhub et al. (2002) were based on a model with self-interested players. Additional discussion of their work is in the section 3.2.

principals' objectives differ or, in the case of multiple tasks, when there are interactions among them in the agent's cost function.

A seminal contribution to a general formulation of the problem of common agency is Bernheim and Whinston's (1986a) model which extends the basic insights of the Grossman and Hart (1983) single principal model to the case of multiple principals.Their contribution focuses on an environment with hidden action, although discussion of the "truthful equilibrium" that is their solution concept, is developed in a complete information setting (1986b). The complete information approach allows the authors to distinguish the effects of common agency from second-best results arising in bilateral agency due to risk sharing. Our experiment follows a similar strategy due to the complexity of the common agency environment relative to previous experimental investigations of contracting relationships. As a result, theoretical predictions are developed in a setting where actions are observable though not enforceable.

The framework employed below is most closely associated with the work of Dixit (1996, 1997) who employs linear contracts to explore some general questions in political economy. Dixit demonstrates, in a setting with hidden action, that common agency can have equilibria similar in character to a prisoner's dilemma. Dixit's model, in the terminology of Laussel and Le Breton (1996), is a game of *private common agency*, since each principal has a direct interest in only one of the agent's tasks. Two mechanisms give rise to incentive problems in this model. First, the structure of the cost function matters. Dixit considers the case of efforts that are strategic substitutes so that each principal is affected by all efforts regardless of his benefit function. He shows that this case gives rise to tight competition between the principals and thus to low powered incentives.

In addition to interactions through the cost function, the strategy space in which contracts are created can give rise to conflict between principals since each can devise incentives for the other's tasks. Imposing a negative incentive on the tasks that benefit the other principals leads to reduced incentives overall. Dixit shows that, in equilibrium, negative incentives for these "off tasks" are imposed. While a unique equilibrium in incentives arises in the linear contracting model, multiple equilibria arise from the use of the fixed component that determines the final distribution.[50]

In the model that follows we extend Dixit's approach to modeling common agency by incorporating other-regarding preferences directly in the utility specification of the agent. This extension is motivated by earlier experimental work that suggests monetary payoffs cannot entirely explain behavior in simpler games of proposal and response (see Camerer 2003 for a review). In an agency setting, Anderhub, Gachter, and Konigstein (2002; AGK hereafter) create a single principal environment in which both incentives and fixed contract components are available. The authors find that while incentive constraints were often recognized, participation constraints generally did not bind. In addition, agent effort choices were often not optimal for a money maximizing agent and were shown to be related to the division of the surplus suggested by the principal's contract offer. These two deviations, AGK argue, imply the existence of behaviorally important "equity" and "reciprocity" constraints. The single principal model, presented in Appendix 3.1 formalizes these results.

[50] In addition to Dixit's work in the hidden action setting, Olsen and Torsvik (1995) explore linear contracting under hidden information. Both investigations acknowledge but do not address the distributional issues arising from the multiple-equilibria in the two-principal setting that we discuss below.

Also relevant to our modeling exercise is the series of experiments initiated by Ernst Fehr and his colleagues on gift exchange. These experiments have demonstrated a relationship similar to that found in AGK. High wage offers are reciprocated with high effort levels, when shirking is the optimal response for self-interested agents. By backwards induction, high wages should not be observed. Of the gift exchange experiments, perhaps the most relevant to this study is that of Gachter and Falk (2002) which investigates the interaction between other-regarding preferences and reputation effects, showing that reciprocal motivations are at work in both one-shot and repeated game environments. Gachter and Falk's design includes a treatment similar to the one implemented here in that players are exogenously assigned, and remain assigned, to interacting groups. The critical difference between our protocol and the gift exchange settings is that, in our setting (as in AGK's) contracting opportunities are available that can yield high output without relying on gift exchange, which enables a more direct study of intentions.

Experimental investigations of common agency are few. To our best knowledge, multiple principals were first studied by Kirchsteiger and Prat (1999, 2001), although their concerns were somewhat different than those motivating the current research. In particular, Kirchsteiger and Prat limit their investigation to the case of a robotic agent programmed to maximize its monetary income over a single task. Thus the iinvestigation of effort choices and their relation to inequity, reciprocity, or reputation formation is not germane to their work. Instead, their goal was to investigate the behavioral validity of Bernheim and Whinston's (1986b) truthful equilibrium. They found that the truthful equilibrium is rarely chosen by the principals, and instead subjects appear to play a less efficient 'natural' equilibrium, which is computationally less complex.

We build on this previous work first by developing equilibrium predictions in the single principal case under both self-interested and inequity-averse preferences and show that in the linear contracting setting it is optimal in both cases for principals to offer high powered incentives. We show also that when reciprocating agents infer intentions from the structure of the contract, low powered incentives can yield higher rents for the principal than the full incentive benchmark. We also consider issues related to repeated play, and show that some inefficiency in the form of low-powered incentives may be optimal so that the principal can gather information on the agent's type. Thus information elicitation in repeated play may be an alternative explanation for the contracts that have the "gift-exchange" structure.

Extending the investigation to the setting of private common agency we consider how the power of incentives can be affected by the additional principal. Two alternative solutions to the principals' problem are developed, each implying different beliefs about the other principal's decision process. The solution concept chosen has important implications, since in one, low powered incentives are part of any equilibrium. In the alternative solution high-powered incentives equivalent to the first-best in the single principal setting can be achieved. Inequity-averse preferences in the two-principal setting result in different contracting strategies than in the single-principal benchmark. In particular, the optimal incentive is sensitive to the inequity aversion of the agent under common agency. Further, principal competition causes the efforts associated with an optimal response to contract incentives to remain constant over a broad range of agent types.

3.3 A Model of Common Agency

This section introduces two finite proposer-responder games, Γ^1, with one principal and one agent, and Γ^2, with two principals and one agent. Actions are completely observable, though not enforceable. Thus, while the agent's effort choices are known to the principal(s) he (they) can neither prescribe effort levels ex-ante, nor alter the agent's compensation after effort choices are made.

In Γ^1, $N = \{0,1\}$, and the set of feasible actions available to player 1, the principal, is finite and consists of a contract choice $a_1 = (\alpha_1, \alpha_2, \beta)$, which consists of piece rate incentives $\alpha_i \in \{A_{i\alpha} : \underline{\alpha}_i \le \alpha_i \le \overline{\alpha}_i\}$ $i = 1,2$ for each task and $\beta \in \{A_{1\beta} : \underline{\beta} \le \beta \le \overline{\beta}\}$ is a fixed payment, with $A_{1\alpha}$ and $A_{1\beta}$ partitions of A_1.[51] The set of feasible actions for player 0, the agent, consists first of an option to accept or reject the principal's contract offer. An indicator $Z_0 \in A_{0Z} \equiv \{0,1\}$ takes on the value of 1 when the contract is accepted. Conditional on acceptance, the agent makes effort choice $e_i \in A_{0e} \equiv \{\underline{e}_i,..\overline{e}_i\}$ $i = 1,2$, with $\underline{e}_i > 0$. The agent's effort cost function is increasing and convex. To generate precise predictions for the experimental tests of the theory I assume the cost function is quadratic and is given by $c_0(e) = e'Ce = \frac{e_1^2 + e_2^2}{2}$ for $e' = [e_1 \quad e_2]$.

[51] The sign of the parameters are not restricted. In the experimental implementation, $\underline{\alpha}_i = -\overline{\alpha}_i$, $i = 1,2$, and $\underline{\beta} = -\overline{\beta}$ and so negative and positive incentives and fixed fees are allowed.

The timing of players' moves in the extensive form are identical for Γ^1 and Γ^2. At $\tau = 0$, the principal or principals make their contract selection, a_i. At $\tau = 1$ the agent decides whether to accept or reject the offered contract, and conditional on acceptance the agent makes their effort choices. At $\tau = 2$, payoffs are allocated. If the contract is rejected all parties receive their reservation payments of zero. If accepted, the agent's efforts and the principals' contract choices determine the payoffs.

3.3.1 The Case of One Principal

The material payoff functions for the players in Γ^1 are given by

(3.1) (a) $$\pi_o = Z_0\left[\alpha_1 e_1 + \alpha_2 e_2 + \beta - \left(\frac{e_1^2 + e_2^2}{2}\right)\right]$$

(b) $$\pi_1 = Z_0\left[(b_1 - \alpha_1)e_1 + (b_2 - \alpha_2)e_2 - \beta\right].$$

In the agent's payoff function the $\alpha_i e_i$, $i = 1,2$ terms represent the payoff from the incentive portion of the contract. Through these instruments the principal chooses the agent's marginal benefit of effort. The amount of the fixed payment is β. The agent also incurs costs of effort. For the principal, $b' = [b_1 \quad b_2]$ are the marginal benefits of the agent's effort which are exogenous to the model.

The analysis begins by exploring, in Γ_1, the effects of money-maximizing (*M*) and inequity averse (*I*) payoff functions on contract choice, effort levels, and size and distribution of the surplus. The payoff function for a money-maximizing player, *M*, is simply the monetary payoff, $u_i^M(a) = \pi_i$, $a \in A, i = 0,1$, where π_i is defined for each player in equation (3.1). Preferences for an inequity-averse player are captured by the utility function

(3.2) $u_i^I(a) = \{Z_0[\pi_i - \delta_i \max(\pi_j - \pi_i, 0) - \gamma_i \max(\pi_i - \pi_j, 0)] \quad \forall\, i, j = 0,1$

The specification of inequity aversion is consistent with the model developed by Fehr and Schmidt (1999). The first component of the utility function for the inequity-averse player is the monetary payoff. The weighted difference of the inequity in payoffs is subtracted from the monetary outcome with the weight given to inequity associated with the parameters δ_i and γ_i. The *max* operators in the inequity-component imply that only one of the weights is operative for a payoff comparison between two players, with δ_i (γ_i) utility relevant when player *i* is behind (ahead of) player *j* in monetary terms. Following Fehr and Schmidt we define $0 \leq \gamma_i < 1$ and $\gamma_i \leq \delta_i$ which reflects the intuition that inequity is more disadvantageous when one is behind in monetary terms than when one is ahead.[52] The inequity-averse formulation captures the intuition that for a given level of surplus, utility is maximized when inequity is absent.

Inequity-averse preferences incorporate the material outcomes of other players in utility. An alternative is to consider intention-based reciprocity (R preferences). The intention-based model implies that kind actions are rewarded and unkind actions punished through effort choices that deviate from the agent's optimal choice under *M* preferences. Effort choices greater (less) than optimal are costly to the agent and reward (punish) the principal in material terms.[53] The payoff function that accommodates reciprocal motivations must incorporate beliefs about the intentions of others, and not only the concern for material outcomes as in u_i^M and u_i^I. As a result, identical monetary

[52] Alternatively, $\gamma_i < 0$ suggests competitive or status-seeking preferences in which utility is increasing with one's own advantage.

[53] Inequity aversion can also motivate these effort deviations as will be discussed below.

outcomes can yield different utilities to the players depending on the process through which the outcomes are achieved. The specification of reciprocity based preferences employed here is very similar to that developed by Falk and Fischbacher (2006) and is introduced in Appendix 3.2. Since the functional form for reciprocal preferences incorporates optimal responses for money-maximizing players as a reference from which intentions are inferred, and in order to present our baseline results, we turn next to the model with *M* preferences which will yield expressions for these optimal actions.

3.3.1.1 Equilibrium Predictions: Γ^1

Equilibrium predictions for Γ^1 for *M* and *I* preferences are developed, and a numerical example is considered that clarifies the relationship between *I* and *R* preferences. The predictions are generated from a two-stage process. In the first stage the agent's optimal response to all feasible contracts is determined, and in the second the principal maximizes his own objective conditional on the first-stage result. Results on the efficiency of the two-part contract, on the impact of other-regarding preferences on distribution and on the optimal contract form are presented. We begin with the baseline case of a self-interested principal and agent.

In stage one of the solution, a self-interested agent chooses an effort level that maximizes monetary payoffs subject to a participation constraint requiring that payoffs are nonnegative.

$$(3.3) \qquad e^* \equiv \arg\max_{A_0} \pi_o\left(a_0;a_1,C\right) = \max_{e,Z} Z_0\left[\alpha_1 e_1 + \alpha_2 e_2 + \beta - \left(\frac{e_1^2+e_2^2}{2}\right)\right]$$

$$(3.4) \qquad \pi_0\left(e^*;a_1,c\right) \geq 0$$

The solution $e*(a_1) = [\alpha_1, \alpha_2\ ]$defines a money-maximizing best reply to the principal's contract.

The principal maximizes his own payoffs by maximizing social welfare and then using the fixed fee to keep the agent at his reservation utility. The principal's problem is therefore to maximize his own objective net of the agent's costs

(3.5) $\max_{A_1} \pi_1(e^*) = \max_{\alpha,\beta}\left[(b_1 - \alpha_1)e_1^* + (b_2 - \alpha_2)e_2^* - \frac{e_1^* + e_2^*}{2}\right]$
subject to $e*(a_1) = [\alpha_1, \alpha_2\ ]$

First order conditions reveal the $\alpha_1^* = b_1$ and $\alpha_2^* = b_2$ and implementing these incentives yields the social surplus $W^* = \frac{b_1^2 + b_2^2}{2}$. Given the contract structure the initial distribution of the surplus is $D = (\pi_0, \pi_1) = \left(\frac{b_1^2 + b_2^2}{2}, 0\right) = (W^*, 0)$. The principal can improve his return by using the fixed component to extract the entire surplus, thus the optimal fixed fee is given by $\beta^* = -\frac{b_1^2 + b_2^2}{2}$. Result 1 follows.

Result 1: Both contract instruments are used in equilibrium. The principal maximizes the total surplus and sets the fixed fee so that the agent is held to his reservation value.

Given this behavior, the optimal contract, agent response, and final distribution of surplus are given by

(3.6) $a_1^* = (\alpha^*, \beta^*) = \left(b, -\frac{b^2}{2}\right);$

$a_0^* = (Z^*, e^*) = (1, b),$

$$D^* = (\pi_0, \pi_1) = \left(\frac{b^2}{2} + \beta^*, -\beta^*\right) = (0, W^*)$$

Incorporating inequity aversion in the utility of the agent yields a change in the distribution of surplus that is easily understood. The aversion to inequity requires the principal to leave some of the monetary surplus with the agent, in order to avoid contract rejection. It remains optimal to maximize social surplus and the fixed fee is used to alter the distribution so that the participation constraint is not violated,

Result 2: Optimal incentives and efforts under inequity aversion are unchanged from the case of M preferences. A proof is in Appendix 3.1

In considering the effect of *I* preferences we restrict attention to the case of disadvantageous inequality for the agent so that the relevant portion of the other-regarding functional is associated with the parameter δ_0. This restriction is intuitive since an inequity-averse agent will always accept a contract resulting in an equal split of the surplus and so the principal need never offer the agent a share greater than one-half. Result 2 on the level of effort and the power of incentives follows immediately from the solution to the game that incorporates inequity-averse preferences for the agent. The inequity-averse agent must solve:

$$(3.7) \qquad u_0^I * \equiv \max_{Z,e} u_0^I(a) = \max_{Z,e} Z\left[\alpha e - \frac{e^2}{2} + \beta - \delta_0\left[(b-\alpha)e - \beta - \left(\alpha e - \frac{e^2}{2} + \beta\right)\right]\right]$$

To understand the impact, consider how the participation constraint of an inequity-averse agent affects the final distribution. A binding participation constraint for the inequity-averse agent requires that $[\pi_0 - \delta_0(\pi_1 - \pi_0) = 0]$ which, rearranged, yields

(3.8) $\pi_0 = \left(\dfrac{\delta_0}{1+\delta_0}\right)\pi_1$.

Result 2 implies that the available surplus remains the same as in the case of a self-interested agent so that $W^I = W^* = \dfrac{b^2}{2} = \pi_0 + \pi_1 = \left(\dfrac{1+2\delta_0}{1+\delta_0}\right)\pi_1$. When the participation constraint binds, the optimal contract is $a_1^{*I} = \left(b, -\left(W^*\right)\left(\dfrac{1+\delta_0}{1+2\delta_o}\right)\right)$ yielding the distribution

(3.9) $D^{*I} = \left(\pi_0, \pi_1\right) = \left(W^* \dfrac{\delta_0}{1+2\delta_0}, W^* \dfrac{1+\delta_0}{1+2\delta_0}\right)$.

Note that the distribution approaches an equal split of the surplus as δ_0 grows large.

3.3.1.2 Reciprocity and Gift Exchange

For agents with *R* preferences it can be optimal for principals to offer *gift-exchange* contracts which are interpreted by the agent as 'kind' offers. As a result the agent may reciprocate with efforts, conditional on the contract offer, greater than those under either *M* or *I* preferences. Appendix 3.2 presents the details of the intention based model which closely parallels that of Falk and Fishbacher (2006). The key implication of the model is presented in Result 3.

Result 3: If agents have R preferences it is optimal to offer reduced incentives and more generous fixed fees than in the baseline result for M preferences.

To understand the gift exchange contract first consider the implications of a non-optimal effort relative to the *M* preference benchmark which occurs when the first order condition is violated and $e \neq \alpha$.[54] A positive deviation, with $e > \alpha$, rewards the principal by increasing his monetary payoff, while a negative deviation lowers his monetary rewards. Effort levels that are not optimal are, of course, always costly for the agent. Thus, while low powered incentives are inefficient when the agent responds optimally, an agent who reciprocates intentions may reward the principal through high effort if the fixed component of the contract compensates.

The intention-based model can alter equilibrium predictions as demonstrated by the following "mini-game" for which numerical utilities are calculated. The choice sets in Γ^1 are restricted so that $A_{1\alpha}^m \equiv \{4,6\}$, $A_{1\beta}^m \equiv \{2,-9\}$, and the possible contract pairings of incentives and fixed fees are limited so that $A_1^m \equiv \{\alpha, \beta : \{(\alpha, \beta)\} \in \{(4,2),(6,-9)\}\}$, where the *m* superscript indicates the set of feasible actions in the mini-game. The agent can respond with $e \in \{4,6\}$. The mini-game is parameterized so that *b*=6, *c*=.5, $\delta_0 = .5$, $\gamma_0 = .25$, and for the intention-based utility specification $\rho_0 = 1$.[55]

Given these parameters, define $FI \equiv (\alpha, \beta) = (6,-9)$ the full incentive contract, since $\alpha = b = 6$. With an optimal effort choice *FI* yields an equal split of the surplus between the principal and agent. Define the contract $GE \equiv (\alpha, \beta) = (4,2)$ the gift-exchange contract since it contains a reduced incentive that leads to inefficient effort for a self-interested agent, and a positive fixed component that, given an optimal response, yields greater monetary payoff than the *FI* contract for the agent. To fix beliefs for the reciprocal preferences we assume that the agent believes the principal will offer the *FI*

[54] Due to the symmetry of the tasks this discussion restricts attention to a single task, dropping the identifying subscripts.

contract which implies an equal share of the surplus to each player when the efficient effort level is chosen.[56] This belief implies that kindness is indicated when the principal offers more than half of the resulting surplus arising from the *FI* contract.

Figure 3.1 presents the extensive form of the game and highlights equilibrium outcomes for the different preference structures. Table 3.1 provides additional detail on the components of the payoff functions. The *GE* contract is optimal for the self-interested principal when the agent has *R* preferences. The principal achieves a payoff of 10 with *GE* since the agent responds with effort of 6 that rewards the principal rather than the optimal effort of 4. For the inequity-averse and money-maximizing agent the *GE* contract yields low effort and the principal therefore chooses the *FI* contract when the preference structure is known.

In the mini-game when the *GE* contract is offered the inequity-averse agent's utility involves the γ_0 parameter since $\pi_0 > \pi_1$. In determining the optimal contract for an inequity-averse agent this parameter did not come into play since it is never optimal for the principal to offer more than one-half of the surplus. Note that the inequity aversion parameter does not have an impact on effort choices in the mini-game. The numerical results show that agents with *M* and *I* preferences make the same effort choice of 4. A simple alteration in the choice set of the agent, however, yields an additional result.

---

[55] The behavioral parameters, δ, γ, and ρ are within the ranges inferred from previous experimental results (Falk and Fischbacher (2006), Fehr and Schmidt (1999))

[56] Falk and Fischbacher (2006) present survey evidence that equality is a focal distribution for weighing the kindness of offers.

Figure 3.1: Single Principal "mini-game"

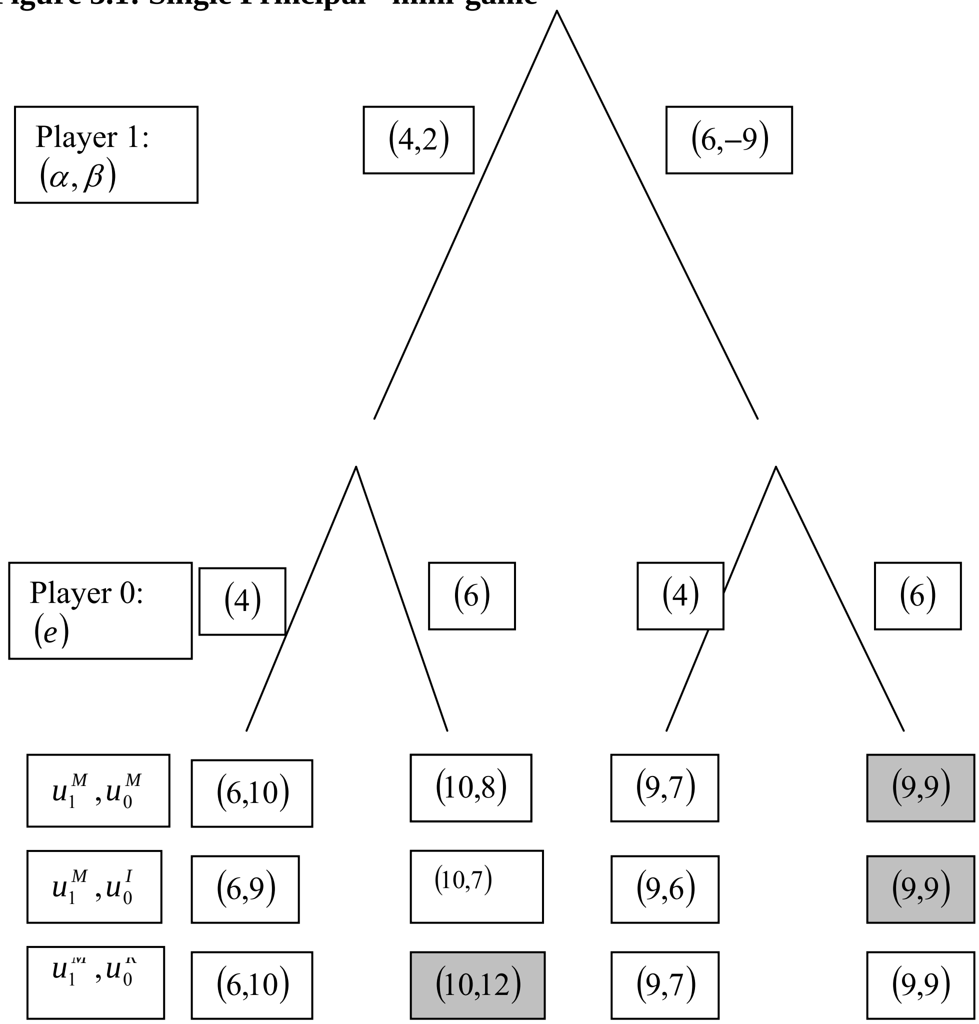

In this extensive form of the mini-game for a single principal, the principal chooses one of two contracts, $A_1^m \equiv (\alpha, \beta) \in \{(4,2),(6,-9)\}$, and the agent one of two efforts $A_{0e}^m \equiv (e) \in \{4,6\}$. If the agent's behavior is based on reciprocal preferences he will choose the high effort despite the low incentives, motivating a gift exchange contract by the principal. Recall that "*M*" refers to money-maximizing preferences and "*m*" to the mini-game restriction on the feasible actions. *I* preferences are inequity-averse and *R* preferences are reciprocal.

Table 3.1 Mini-Game Utilities

	Gift Exchange $a_1^m = (4,2)$		Full incentive $a_1^m = (6,-9)$	
Effort (a_0)	*4*	*6*	*4*	*6*
Inequity*	-2	-1	-1	0
Kindness**	1	1	0	0
Reciprocity***	0	4	0	0
u_1^M, u_0^M	6,10	10,8	9,7	**9,9**
u_1^M, u_0^I	6,9	10,7	9,6	**9,9**
u_1^M, u_0^R	6,10	**10,12**	9,7	9,9

Utilities for money-maximizing (M), inequity-averse (I) and reciprocal (R) preferences. Equilibrium choices for each pairing (row) are in bold. The equilibrium predictions are also identified as the shaded squares in Figure 3.1.

*Inequity = $-.5[\pi_1(a_1, a_0) - \pi_0(a_1, a_0)]$ when $u_1^S > u_0^S$ and

Inequity = $-.25[\pi_0(a_1, a_0) - \pi_1(a_1, a_0)]$ when $u_1^S < u_0^S$.

**Kindness = $[\pi_0^*(a_1) - \pi_0^*(a_0')]$

***Reciprocity = $[\pi_1(a_1, a_0) - \pi_1^*(a_1)]$

Result 4: Given a sufficiently dense set of feasible efforts, inequity-averse agents will reward the principal with effort greater than e^{*M} when offered gift-exchange contracts. A proof is in Appendix 3.1.

Consider adding the intermediate effort level, 5, to the agent's set of feasible actions to create the mini-game m' so that $A_{0e}^{m'} \equiv \{4,5,6\}$. With these alternatives the inequity-averse agent chooses effort level 5 in response to the *GE* contract. Monetary payoffs and utility for m' are presented in Table 3.2, which shows that the agent prefers the monetary outcomes $(\pi_0 \quad \pi_1) = (9.5 \quad 8)$ resulting from an effort level of five to the payoffs $(\pi_0 \quad \pi_1) = (10 \quad 6)$ due to the reduction in inequity. Thus, although it is not optimal to offer an inequity-averse agent the *GE* contract, the comparative static effect is the same as would be found for the reciprocal agent when it is offered. In both cases maximizing utility causes the agent to make additional, costly effort which benefits the principal. Result 4 demonstrates that the agent's response to the *GE* contract is similar for *I* and *R* preferences. As a result our theoretical examination of other-regarding preferences in Γ^2 focuses on the the case of inequity aversion.

Table 3.2 Optimal Efforts in Expanded Mini-Game

	Gift Exchange $a_1^m = (4,2)$		
Effort (a_0)	*4*	*5*	*6*
u_1^M, u_0^M	**6,10**	8, 9.5	10,8
u_1^M, u_0^I	6,9	**8, 9.1**	10,7

The inequity-averse agent chooses effort 5, when the feasible set is expanded to {4,5,6} in game *m′* sacrificing own payoff for increased equity.

3.3.1.3 Some Rudimentary Dynamics

The results so far have focused on the stage game solutions for Γ^1. Given the assumptions of common knowledge of game and utility parameters, in a finitely repeated version of the game the results for the stage game extend naturally to a consideration of final period play. An argument based on backwards induction suggests that anticipation of the final stage game should cause potential deviations in all previous rounds to "unravel" and so the results presented above should hold in all rounds (Selten (1978); see also Anderhub, Gachter, and Konigstein (2002)).

A drawback of this model is the strong assumption that utility parameters associated with inequity aversion are common knowledge. Rather than developing a complete model of reputation formation, in the empirical section, the behaviors of the agent are examined to distinguish self-interested play, which may be influenced by reputational concerns, from that associated with other-regarding preferences. Effort levels that are correlated with the fixed component of the contract or with the equity of payoffs are indicators of other-regarding motives, however, the repeated game gives self-interested players the opportunity to earn surplus, by imitating other-regarding behavior. This strategy would be costly in the last round of play, which comprises the only one-shot game in the session, and so we should observe low-effort "defections" in the last period by money-maximizing players. In addition, since the payoffs to cooperation diminish with the time remaining in a repeated relationship, a reputational model predicts smaller rewards to the principal over time, and not just in the last round of play (Fudenberg and Tirole (1991)).

The empirical investigation of these alternatives is similar to the approach taken by Falk and Gachter (2002) who analyze one-shot and repeated interactions in a gift-exchange game without developing a model of reputation. Our analysis is complicated due to the dual contracting instruments available in this game. In particular, the principal can create gift-exchange contracts early in the game to observe the agent's efforts and so estimate their inequity aversion parameter. High powered incentives and a fixed fee that arranges a distribution consistent with this behavior reduce the returns to defection in the final period. As a result under uncertainty about agent type we expect to see increasing incentives and a falling fixed fee over time in order to lock in desired behaviors while avoiding the risk of contract rejection.

Result 5: With uncertainty regarding the extent of inequity aversion we expect to observe an increase in incentives and a falling fixed fee over time.

3.3.2 The Case of Two Principals

The addition of a second principal defines Γ^2. Player 1 and Player 2 have actions defined identically to those of player 1 in Γ^1. The agent thus faces the aggregate contract derived from both principal's choices, $a^{\Gamma_2} = (\alpha_{11} + \alpha_{21}, \alpha_{12} + \alpha_{22}, \beta_1 + \beta_2)$, with α_{ij} principal i's choice of incentive for task j, $i, j = 1,2, \quad i \neq j$.

The critical difference between Γ^1 and Γ^2 is that the principals must consider not only the effort and acceptance behavior of the agent, but also the contract choices of the other principal. The solution concept developed by Dixit (1996) in the multiple principal setting mirrors that of the single principal example just examined. The principals derive the optimal response through backwards induction, incorporating the fact that the other principal is also setting incentives non-cooperatively. In our baseline case with C diagonal, there is a dynamic similar to the prisoner's dilemma that yields low powered incentives relative to the single principal benchmark. The final allocation is then determined through the fixed fee which contains a component set by each of the two principals.

One difficulty with this solution is that, while there is a unique equilibrium in incentives, there are multiple equilibria with regards to the fixed fee that may yield uncertainty for the players. An alternative solution concept consists of the simultaneous choice of incentives and fixed components, and we show that players who anticipate the non-cooperative solution in incentives can unilaterally improve their position with cooperative incentives and an altered fixed fee. As a result the non-cooperative equilibrium unwinds and something close to the fully cooperative solution may be implemented.

As with Γ^1 we begin solving a simple version of Γ^2 for M preferences, restricting the parameters of the cost function so that $c_{11} = c_{22} = 1/2$, and $c_{12} = c_{21} = 0$. Following Dixit (1997), we model private common agency in which each principal has an interest in a single task. In the extension to two principals, benefits are distributed across the principals so that $b_1' = [b_{11} \quad b_{12}]$, and $b_2' = [b_{21} \quad b_{22}]$, with the first subscript referring to the principal, the second to the task. Private common agency implies that $b_{12} = b_{21} = 0$ and marginal benefits are therefore $b_1' = [b_{11} \quad 0]$ and $b_2' = [0 \quad b_{22}]$, with $b_{11}, b_{22} > 0$. Interactions between principals are still relevant in the case of private common agency, due to instruments that allow principals to penalize or reward the agent for effort on the other's task.

The extensive form of Γ^2 is as in the single principal case. Both principals, create components of the aggregate contract for the agent at $\tau = 0$. Each principal can create incentives for both of the tasks and so that principal i selects from $\alpha_{ij} \in \{\underline{\alpha}_{ij}, \underline{\alpha}_{ij} + 1, \ldots, \overline{\alpha}_{ij} - 1, \overline{\alpha}_{ij}\}$ $i, j = 1,2$, where $\underline{\alpha}_{ij} = -\overline{\alpha}_{ij}$ $i, j = 1,2$. The agent's acceptance and effort decision solves

$$(3.10) \quad e^{*\Gamma^2} \equiv \arg\max_{Z_0, e} \pi_o(a) = Z_0\left[(\alpha_{11} + \alpha_{21})e_1 + (\alpha_{12} + \alpha_{22})e_2 + \beta_1 + \beta_2 - \left(\frac{e_1^2 + e_2^2}{2}\right)\right].$$

The solution to the first stage problem takes into account the aggregate incentives provided by both principals, and is given by

$$(3.11) \quad e_i^*(\alpha) = \alpha_{ii} + \alpha_{ji}, \; i, j = 1,2, \quad i \neq j.$$ [57]

In thinking about the principal's problem, consider first that a cooperative agreement to provide full incentives leaves opportunities for defection. One way to see this is to consider the aggregate contract in which each principal offers full incentives for his own task and ignores the other. This contract yields the same results for each task as

[57] In what follows we restrict attention to Γ^2 and so simplify notation by eliminating the identifying superscript unless a comparison is made to Γ^1 .

in the single-principal benchmark. That each principal has a potentially profitable deviation from this cooperative full incentive contract can be seen in the formulation of their maximization problem:

(3.12) $\max_{\alpha_{ii}\alpha_{ij}} \quad (b_{ii} - \alpha_{ii})e_i^* + (0 - \alpha_{ij})e_j^* - \beta_i \quad i, j = 1,2 \;\; i \neq j$ **.**

Equation (3.11) shows that *e** is the response to the aggregate incentives. That defection from the cooperative solution may be beneficial can be inferred from the second term of equation (3.12) since the product $(-\alpha_{ij})e_j^* > 0$ when $\alpha_{ij} < 0$ and the contract is accepted. This increase is mitigated by lower agent efforts since $e_j^* < \alpha_{jj}$ in this case. The tradeoffs involved in the choice of incentives are clarified by examining the first order conditions from equation (3.12) which yield $\alpha_{11}^* = \frac{2b_{11}}{3}, \alpha_{12}^* = -\frac{b_{22}}{3}$, $\alpha_{22}^* = \frac{2b_{22}}{3}$, and $\alpha_{21}^* = -\frac{b_{11}}{3}$. The resulting optimal aggregate incentives are therefore, $\alpha_1^* = \alpha_{11}^* + \alpha_{21}^* = \frac{b_{11}}{3}$ and, symmetrically, $\alpha_2^* = \alpha_{22}^* + \alpha_{12}^* = \frac{b_{22}}{3}$, which are one-third of the power of those in the cooperative first-best.[58] The resulting low-powered incentives constitute Result 6.

Result 6: With M preferences and C diagonal, incentives are reduced from the cooperative first best. A consequence is that the resulting surplus per task is

[58] When the cost function is not diagonal aggregate incentives are given by $\alpha_1 = \alpha_{11} + \alpha_{21} = -\left[\frac{8c_{11}(c_{12}+c_{21})(b_{22}) + \left[12c_{11}c_{22} + (c_{12}+c_{21})^2\right](b_{11})}{(c_{12}+c_{21})^2 - 36c_{11}c_{22}}\right]$. This expression is increasing and convex and so generates an increase (decrease) in incentives from the baseline if $(c_{12} + c_{21})$ is greater than (less than) zero.

smaller than under the single principal regime and is given by $W^{*\Gamma^2} = \frac{5b^2}{18} < \frac{b^2}{2} = W^{*\Gamma^1}$.

This solution yields a unique equilibrium in incentives and the preliminary distribution is $D = \{M_0, M_1, M_2\} = \{(b_{11}^2 + b_{22}^2)/18, (b_{11}^2 + b_{22}^2)/9, (b_{11}^2 + b_{22}^2)/9\}$ prior to reallocation of the surplus due to the fixed component of the contract. While the optimal aggregate fixed fee is also uniquely defined as $\beta = -(b_{11}^2 + b_{22}^2)/18$ there are multiple equilibria with respect to the contribution of each principal to the aggregate β. The equilibrium contract offer is $a_i = (\alpha_{ii}, \alpha_{ij}, \beta_i) = (2b_{ii}/3, -b_{jj}/3, -(b_{11}^2 + b_{22}^2/18) - \beta_j)$, for principal *i*.

The dependence of β_i on β_j poses a problem for equilibrium selection that can be resolved in a number of ways. The simplest, which may be sensible in a symmetric situation, is to assume a focal point with each principal demanding one-half of the remaining surplus leaving the agent at his reservation utility. Alternatives to this scenario are also plausible. Consider for example the reasoning of a player who thinks that a tremble or some other violation of rationality by the other principal may result in a contract that violates the agent's participation constraint. In this setting it is plausible to leave a portion of the surplus with the agent. A contract of this type, with low powered incentives due to principal competition, and with a rent to the agent, is from the perspective of the agent qualitatively similar to the *GE* contract considered in the single principal setting. In addition to the complication arising from the multiple equilibria, the possibility that the agent is inequity-averse may cause additional surplus to be left with the agent.

Before examining Γ^2 with an inequity-averse agent, we consider an alternative solution concept in which contract components are chosen simultaneously by each principal rather than through the two-stage process examined above.[59] To examine this alternative a numerical example extends the mini-game presented for Γ^1 to include the second principal, and a larger set of feasible actions. The feasible actions for each principal are $\beta_i \in \{-18,-17,...,0,...17,18\}$, $\alpha_{ij} \in \{-6,-5,...,0,...,5,6\}$ for *i*,*j*=1,2. Agent effort choices are restricted so that $1 \le e_i \le 6$.[60] One would be hard-pressed to call this a "mini-game" since it yields 37x13 = 481 possible contract choices for each principal, and 6 possible effort responses to each contract choice. This game can be presented in normal form as six matrices with 481^2 = 231,361 cells in each, a total of 1,388,166 possible outcomes. The iterated elimination of dominated strategies yields 830 pure strategy Nash equilibria, 222 of which leave all players with profits greater than or equal to zero. Interestingly, the incentive pair $\alpha_{ii} = 4, \alpha_{ji} = -2$, the solution derived from equation (3.12) for the parameters used in this example is not among the strategies that survive. A small number of cells from the normal form are presented in Table 3.3 to clarify this result.

Result 7: A solution that has the potential to yield high powered incentives in the common agency game arises from the iterated elimination of dominated strategies.

Table 3.3 presents, symmetrically, three strategies for each principal in normal form. Principal 1 is the row player and principal 2 the column player and strategy

[59] The simultaneity is conceptual in that with either solution concept the extensive form timing does not change.

choices are numbered $s_i \in \{1_i, 2_i, 3_i\}$ for player i = 1, 2. Each player's numbered strategy is associated with a contract choice that consists of $a_i = (\alpha_{ii}, \alpha_{ij}, \beta_i)$ $i, j = 1,2$, and the payoffs in the cells assume optimal responses by a self-interested agent. The strategy pair (1_1,1_2) is the equilibrium found when incentives are chosen non-cooperatively, and the fixed component holds the agent to his reservation utility while yielding the principals an equal split of the surplus. Each principal has a profitable deviation through strategy (2) which encourages output by increasing incentives and retains a larger share through the fixed fee. Note that unilaterally shifting strategy 2 increases the surplus of both principals and so is interpreted as a kind act in the reciprocity model. The strategy pair (2, 2) is not an equilibrium, however since each can deviate profitably through strategy (3). Thus the two solution concepts, offer competing hypotheses about the outcome of Γ^2. When contract components are chosen simultaneously, the incentive portion includes the cooperative solution $a_{i\alpha} = (\alpha_{ii}, \alpha_{ij},) = (6,0)$ $i, j = 1,2$ among the equilibria.[61]

Table 3.3 Monetary Payoffs to Selected Strategies

	(1_2) 4, -2, -2	(2_2) 6, 0, -14	(3_2) 6, 0, -18
(1_1) 4, -2, -2	0,10, 10	**0,18, 14**	0, 0, 0
(2_1) 6, 0, -14	**0,14,18**	8,14,1 4	4, 14,18
(3_1) 6, 0, -18	0, 0, 0	4, 18, 14	**0, 18, 18**

Principal 1 is the row player and principal 2 the column player. The agent responds optimally to incentives. Pure strategy Nash Equilibria are in bold type.

Under common agency the impact of inequity aversion is not as straightforward as in the single principal case. For one, the reference group to which an inequity-averse

[60] These parameters are used in the experiment below.

[61] With *M* preferences the incentive principal *i* sets for task *j* is, in equilibrium, always less than or equal to zero. With reciprocal preferences incentives for the off-task may be greater than zero.

player compares outcomes cannot be known a priori; an agent may be concerned with the principals as a unit, or comparisons with each individually. In what follows we consider the case in which the comparisons are made relative to each principal individually since the group comparison yields the single principal results.

When agent payoffs are compared relative to each principal, the agent's problem is to choose effort to maximize

$$(3.13) \quad u_0(a) = Z_0\left[\pi_0 - \frac{\delta_0}{2}\sum_i \max(\pi_i - \pi_0, 0) - \frac{\gamma_0}{2}\sum_i \max(\pi_0 - \pi_i, 0)\right] \quad i = 1,2$$

The solution to the inequity aversion model is then derived by solving equation (3.12) which yields result 8.

Result 8: In contrast to the single principal case, optimal incentives depend on the agent's level of inequity aversion.

The columns of Table 3.4 contain the effort levels and incentives associated with the solution to this problem for the utility specifications $M, I\gamma$, and $I\delta$, with the latter two specifications identifying the term in the *I*-preference functional that is utility relevant for the payoff comparison. Optimal incentives increase (decrease) when the agent is behind (ahead of) the other two players. This can be seen by comparing incentives in the last row of Table 3.4 and observing that $\alpha^{\gamma} < \alpha^{M} < \alpha^{\delta}$.

Table 3.4 Efforts and Incentives in Common Agency

	M	γ	δ
e_i	α	$\dfrac{\alpha(-2+3\gamma)-b_{ii}\gamma}{2(-1+\gamma)}$	$\dfrac{\alpha(2+3\delta)-b_{ii}\delta}{2(1+\delta)}$
α_{ii}	$\dfrac{2b_{ii}}{3}$	$\dfrac{b_{ii}(4-7\gamma)}{3(2-3\gamma)}$	$\dfrac{b_{ii}(4+7\delta)}{3(2+3\delta)}$
α_{ji}	$-\dfrac{b_{ii}}{3}$	$-\dfrac{2b_{ii}(1-\gamma)}{3(2-3\gamma)}$	$-\dfrac{2b_{ii}(1+\delta)}{3(2+3\delta)}$
α	$\dfrac{b_{ii}}{3}$	$\dfrac{b_{ii}(2-5\gamma)}{3(2-3\gamma)}$	$\dfrac{b_{ii}(2+5\delta)}{3(2+3\delta)}$

The first row identifies the preference structure, with M preferences representing those of a money-maximizing agent and γ and δ those of an inequity-averse agent when ahead or behind in monetary terms. Subsequent rows identify the efforts and incentives for task i created by the principals i and j.

For delta contracts the optimal aggregate incentive is $a_i^{\delta} = \alpha_{ii}^{\delta} + \alpha_{ji}^{\delta} = \dfrac{b_{ii}(4+7\delta)}{3(2+3\delta)} + \dfrac{-2b_{ii}(1+\delta)}{3(2+3\delta)} = \dfrac{b_{ii}(2+5\delta)}{3(2+3\delta)}$, $i = 1, 2$, which is increasing and concave since $\dfrac{\partial a_i^{\delta}}{\partial \delta} = \dfrac{4b_{ii}}{3(2+3\delta)^2} > 0$ and $\dfrac{\partial^2 a_i^{\delta}}{\partial \delta^2} = -\dfrac{8b_{ii}}{3(2+3\delta)^2} < 0$. Similarly, for gamma contracts the optimal aggregate incentive is $a_i^{\gamma} = \alpha_{ii}^{\gamma} + \alpha_{ji}^{\gamma} = \dfrac{b_{ii}(4-7\gamma)}{3(2-3\gamma)} + \dfrac{-2b_{ii}(1-\gamma)}{3(2-3\gamma)} = \dfrac{b_{ii}(2-5\gamma)}{3(2-3\gamma)}$, $i = 1, 2$, which is decreasing and concave since $\dfrac{\partial a_i^{\gamma}}{\partial \gamma} = -\dfrac{4b_{ii}}{3(2-3\gamma)^2} < 0$ and $\dfrac{\partial^2 a_i^{\gamma,\delta}}{\partial \gamma^2} = -\dfrac{8b_{ii}}{3(2-3\gamma)^2} < 0$.

Figure 3.2 depicts how incentives change across a relevant range of γ, and δ parameters for both the components of the incentive for each principal and the aggregate. Result 8 makes intuitive sense since, as was seen in the single-principal case, the inequity-averse agent punishes through low efforts when behind and rewards with high effort when ahead. Thus additional incentives are required to maintain effort in the case of disadvantageous inequality, but may be reduced when the agent desires to reduce inequality when ahead materially.[62] Under the γ specification the incentives rapidly approach zero as gamma increases and the agent responds only to the fixed fee. In this extreme case, the motivation for effort arises only from the desire to reduce inequity. Further, as the net incentives approach zero, the two-principal non-cooperative solution implies that each principal creates contract incentives that exactly offset each other. Thus, when gamma is equal to .4, the incentives offered for each task by each principal are equal to plus and minus 3. The effort level obtained in all the inequity-averse cases is the same, however, and is equal to that under *M* preferences. This can be seen by substituting the appropriate α from Table 3.4 in the expression for *e*, which in all cases yields $e = \frac{b}{3}$.

---

[62] For *I2* the choice of incentives depends on the relative size of δ and γ, however the optimal effort is the same as in the three other cases.

Figure 3.2 Optimal Incentives for Inequity-Averse Agents

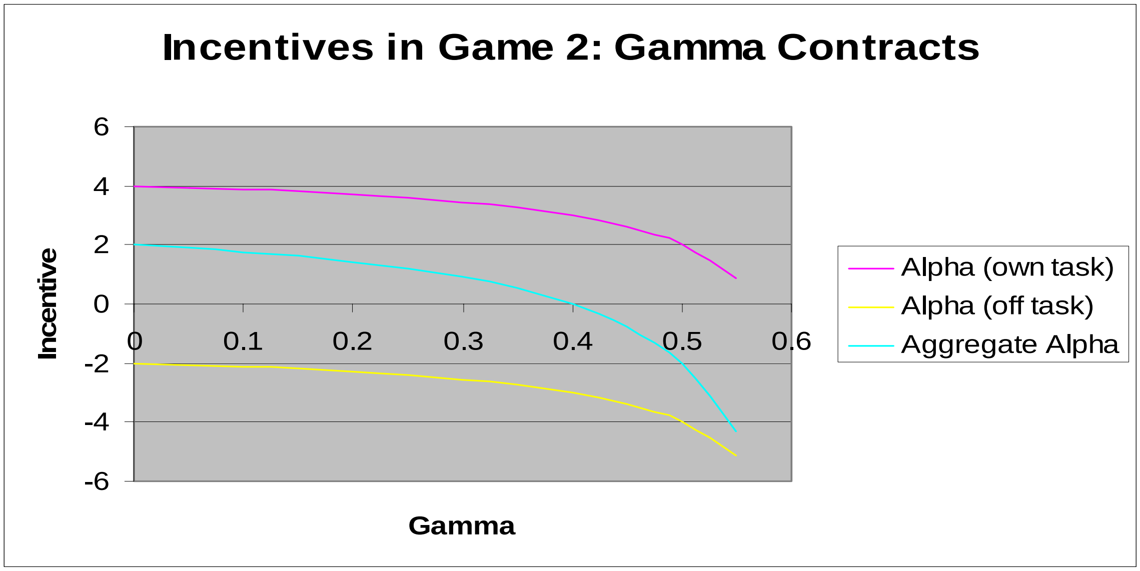

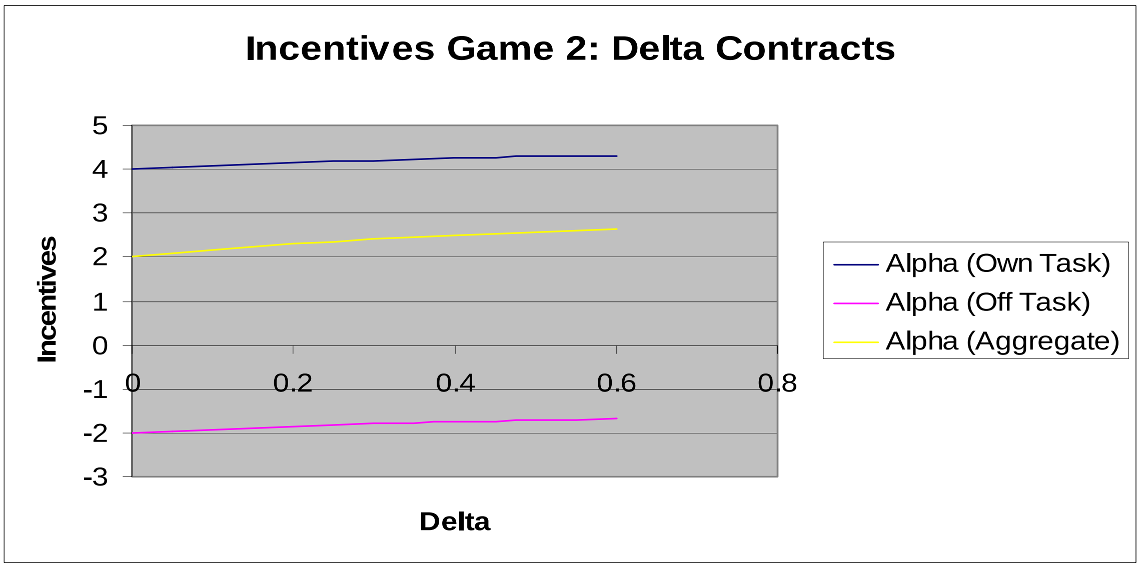

Aggregate incentives decrease (increase) with inequity aversion when the agent is ahead of (behind) the principals in monetary terms.

To determine the final distribution the relevant efforts and incentives from Table 3.4 are substituted into the expressions for monetary returns. For delta preferences, the monetary reward for each principal is then given by

$$(3.14)\quad \pi_i(a) = \left[\frac{\left(2b_{11}^2 + 2b_{22}^2\right)(1+\delta)}{9(2+3\delta)}\right] - \beta_i, \quad i = 1,2$$

which is identical for each principal if the fixed fees are the same. The agent's monetary profit is given by

(3.15) $\pi_0(a) = \left[\frac{(2b_{11}^2 + 2b_{22}^2)(1+\delta)(1+3\delta)}{9(2+3\delta)^2}\right] + \beta_1 + \beta_2$

If $u_0(a) = \pi_0 - \frac{\delta_0}{2}(\pi_1 + \pi_2 - 2\pi_0) = 0$ then the sum of the fixed fees is given by $\beta_1 + \beta_2 = \frac{\delta(\breve{\pi}_1 + \breve{\pi}_2) - 2\breve{\pi}_0(1+\delta)}{2+3\delta}$ where the $\breve{\pi}_i$ represents the profits net of the fixed fees, the portion within square brackets in equations (3.14) and (3.15). The resulting optimal contract is given by

(3.16) $a_{i\delta} = (\alpha_{ii}, \alpha_{ij}, \beta_i) = \left(\frac{b_{ii}(4+7\delta)}{3(2+3\delta)}, \frac{-2b_{jj}(1+\delta)}{3(2+3\delta)}, \left(\frac{\delta(\breve{\pi}_1 + \breve{\pi}_2) - 2\breve{\pi}_0(1+\delta)}{(2+3\delta)}\right) - \beta_j\right).$

Identical reasoning yields an optimal gamma contract of

(3.17) $a_{i\gamma} = (\alpha_{ii}, \alpha_{ij}, \beta_i) = \left(\frac{b_{ii}(4-7\gamma)}{3(2-3\gamma)}, \frac{-2b_{jj}(1-\gamma)}{3(2-3\gamma)}, \left(\frac{\gamma(\breve{\pi}_1 + \breve{\pi}_2) - 2\breve{\pi}_0(1-\gamma)}{(2-3\gamma)}\right) - \beta_j\right).$

3.4 Summary of Theoretical Results

The model developed in the previous section, and the experimental treatment that follows, presents the subjects with a simpler agency problem than is found in the model of moral hazard developed by Dixit (1996, 1997). Since risk-sharing motivations are absent, the optimal action of a single principal is to "sell the firm" at the price of the entire surplus. This requires the provision of full incentives to maximize the aggregate surplus and the use of a negative fixed fee that holds the agent to their reservation utility to implement the equilibrium distribution.

The possibility that the agent is inequity-averse introduces complications for the principal that are similar to those encountered in simpler games of proposal and response. When rationality and utility parameters are common knowledge, efficiency is preserved and the final distribution of surplus is more favorable to the inequity-averse agent. When the inequity aversion parameter is uncertain, there are plausible conditions under which efficiency will be sacrificed to gain information about the agent's parameter in a repeated game setting. The contract that can result from this uncertainty can look like a gift exchange contract in which low-powered incentives and a high fixed fee is observed. The power of incentives will increase over time if the principal tries to constrain opportunistic behavior by reciprocity imitators.

We examine arguments for two alternative solutions to the common agency problem. Solving a two-stage optimization problem through backwards induction yields the result that efficiency is undermined as each principal punishes the agent for effort devoted to the other's task. The iterated elimination of dominated strategies suggests however that the first best remains a potential equilibrium. Under principal competition with an inequity-averse agent optimal incentives increase (decrease) when the agent is behind (ahead) materially, but the competition between principals implies that overall efforts and thus efficiency remain at the low level found under *M* preferences.

3.5 Experimental Protocol

The following sections report on the results of a controlled laboratory experiment that tests key elements of the theory developed in the previous sections. The research subjects were undergraduate students at the University of Maryland who received cash payments based on their performance during the experimental sessions. Interactions were anonymous with subjects communicating their decisions through a computer network. The critical element of the investigation is the exogenous variation in the number of principals from one to two in the principal-agent setting. Interacting groups in all treatments contained a single agent. Treatment 1 contained a single principal and treatment 2, two principals.

In the experimental sessions subjects were randomly assigned to variants of Γ^1 and Γ^2. Within each session, random assignments to the roles of either principal or agent were made and subjects were matched for the duration of the experiment. The roles were given the neutral labels ("Player X" and "Player Y" in the two person treatments, and "Player X", "Player Y", and "Player Z" in the three player treatments). Contract components were labeled as "multiplier" and "fixed" components, and the agent's effort levels as "efforts".

The experimental environments were parameterized as in the mini and maxi-game examples discussed earlier, with b_{ii} = 6, $c_{11} = c_{22} = .5$, $c_{12} = c_{21} = 0$. Effort choices were integers with $\underline{e_i} = 1$. Efficient effort for each task is given by $e_i = 6,\ i = 1, 2$ and the maximum output is therefore W^* = 36. Each point earned was worth 25 (37.5) cents in treatments with 2 (3) players.[63] As a result the maximum earnings in each of eight rounds of play were $4.50 per player yielding players who achieved efficiency average earnings of $36 each for the session.

Each experimental session consisted of an introductory portion that began with the reading of the experimental instructions to the group.[64] The introduction also included, two practice periods of five minutes during which the players became familiar with the mechanics of the experimental program. The program included a profit calculator so that both principals and agents could map strategies to outcomes for any of

[63] The difference in point values across two and three player games equalized the aggregate potential earnings across treatments per player.

[64] Instructions for the treatments are in Appendix 3.2.

the strategy combinations that were available.[65] During the practice session the experimenter responded in private to questions raised by any of the student subjects. At the conclusion of the trial periods subjects were asked to fill out a questionnaire to demonstrate their understanding of how contract and effort choices affected outcomes.

The portion of the session that was consequential consisted of 8 rounds of play, with each round consisting of: a) the principal's contract creation, b) the agent's rejection or effort choices, c) profit calculation and display, and d) the recapitulation of session history. In the first four rounds the principals were allowed four minutes for part a) to calculate outcomes of alternative strategies, and to choose their binding offer. The time was reduced to two minutes in the last four rounds. In b), the agents' were given two minutes to choose a response, and the session proceeded to c), the profit display screen, when the last agent was finished. The profit display showed the actions and payoffs for all players in the same group, and in d) the session history displayed this information for the current and all preceding rounds. Sessions took approximately one and half hours and payments averaged $21.25 including a fixed payment of $5 for timely arrival. Table 3.5 provides additional detail on payments and the number of subjects and observations in each treatment.

Table 3.5 Experimental Session Data by Treatment

	Number of Subjects	Number of Stage Games	Average Earnings
Treatment 1	46	184	$27
Treatment 2	75	200	$18

Treatments are identified according to the number of principals, either one or two. In both treatments there was a single agent. Earnings include a fee of $5 for timely arrival.

[65] The experiment was programmed using Z-TREE (Fischbacher 1999).

3.6 Experimental Results

We present results of two types, first examining aggregate measures to discern overall treatment effects and then looking more closely at individual choices to discover the extent of heterogeneity in strategies both within and across experimental groups and treatments. We use both unconditional non-parametric statistical tests and conditional regressions in both parts of the analysis.[66] The results on aggregate behaviors examine the contract formation decisions of principals and the acceptance/rejection behavior and effort choices of agents, comparing behavior across treatments and relative to theoretical benchmarks. Uniting the analysis of the behavior of principal(s) and agent, we calculate equity and efficiency measures for the treatments which shed light on the other-regarding behaviors of the population as a whole. We gain additional insights by extending the study of the aggregate measures to capture changes over the time path of the repeated game.

The investigation of individual behavior focuses on identifying other-regarding preferences of agents as revealed in their effort choices, and in understanding the extent to which principals make use of this information. Player types are identified both by the relationship of effort to contract components, and by examining late round defections from predicted inequity-averse behavior. We identify the different types of situations in which principals may learn to contract more efficiently, and catalog their frequency within a group to see how different types of learning opportunities affect outcomes. The efficiency of learning is studied by developing an empirical model that compares

[66] We apply non-parametric tests since preliminary results suggest that the data is not normally distributed. For the non-parametric tests, independent observations are created by aggregating decisions at the level of principal-agent interaction group, which are either pairs or triples depending on the treatment. Due to the

observed responses to theoretical best responses. Contracting in the two-principal setting is further studied in light of the alternative predictions associated with iterated elimination of dominant strategies and the two-stage solution concepts.

3.6.1 Aggregate results

The fundamental prediction of the common agency model, as parameterized in this experimental environment, is that the power of incentives are reduced and the fixed component of the contract is increased when the number of principals increases from one to two. Table 3.6 presents the theoretical predictions and the key aggregate empirical results for contract components. Theoretical predictions are those for players who are myopic money-maximizers (*M* preferences) and are derived from the solution to the single and two-principal problems in section 3.3. The average aggregate incentives for both tasks of 5.95 and 2.45 in treatment 1 and 2 are lower than the predicted levels of baseline optimum for self-interested players of 12 and 4. The fixed fees are larger than the predicted values and in the single principal case the difference is dramatic, with the observed value of 2.1 much greater than theoretical prediction of -36. The observed value of .51 is modestly higher than the predicted value of -2, in treatment 2. Both low incentives and large fixed fees imply that a larger share of output than predicted by the model predicated on *M* preferences is retained by the agent. In addition, with *M* preferences, the low incentives imply inefficient output. Before examining the economic consequences in more detail we report on the statistical significance of the results.

Tests of the differences across treatments are made using the Mann-Whitney U test (Siegel 1956). A significant difference in the power of incentives is found with the

---

repeated interaction in the experimental design each of these interaction groups is independent of the others statistic yielding p = .001, while the comparison of the fixed components is not significant (p = .356). To compare the medians of the aggregated data with theoretical predictions we use the Wilcoxon one-sample test. We reject the theoretical predictions for treatment 1, with $p<.001$ for both the incentive and fixed fee components of the contract. Treatment 2 presents a mixed result with the incentive significantly below the theoretical prediction (p = .060) but the fixed component not significantly different (p = .236).

The preliminary analysis of aggregate contract formation suggests that in treatment 1, gift exchange contracts, with low-powered incentives and large fixed fees are likely to be prevalent. The contracts formed in treatment 2 present the same profile, and suggest that either the two principals are unable to cooperatively resolve the dilemma associated with their joint provision of incentives, or alternatively intend to provide motivations through gift exchange. The aggregate analysis, however, obscures the diversity of contract choices within and across treatments and across the session length. A first step to understanding this diversity is to differentiate the aggregate contract offers according to the agent's decision to accept or reject them.

---

for the entire session. There are 23 (25) independent observations in treatment 1 (2).

Table 3.6 Aggregate Contract Formation

	Incentive			Fixed Fee		
	All	Accepted	Rejected	All	Accepted	Rejected
Single Principal Theoretical Prediction	12.00	-	-	-36.00	-	-
Treatment 1 (91%)	5.95	6.08	4.63	2.10	3.20	-9.44
Two Principal Theoretical Prediction	4.00	-	-	-2.00	-	-
Treatment 2 (77%)	2.45	3.81	-1.99	0.51	4.05	-11.04

The observed contract components averaged over all rounds of play, and by acceptance decision of the agent. The share of contracts accepted in each treatment is included in column 1.

Contract components by acceptance decision are also presented in Table 3.6, and in column 1 the proportion of contracts accepted in each treatment is displayed. The rate of contract acceptance is significantly greater in treatment 1 (91%) than in treatment 2 (77%) with the Mann Whitney statistic yielding a p-value of .004. The difference in fixed fees is numerically large and statistically significant across acceptance in both treatments, while the incentives across accepted and rejected contracts differ only in treatment 2.[67] The contract acceptance decision is further clarified by a probit analysis that is presented in Table 3.7.

[67] The Mann-Whitney test yields significance at the level of $p = .953$ ($p < .001$) for incentives in treatment 1 (treatment 2), and significance at the level of $p = .088$ ($p < .001$) for fixed fees.

Table 3.7 Agent Acceptance Behavior - Random Effects Probit Model

Dependent variable: *Accept*	3.7a: All treatment 1 and treatment 2 decisions n = 384 Pr(*accept*=1)=.94			3.7b. Model with lagged profit variable. Only period 2 through 8 decisions are included. n = 336 Pr(*accept*=1)=.95		
Ind. Variables:	Marginal Effect	z stat	P>|z|	Marginal Effect	z stat	P>|z|
Apropt	0.007329	2.97	0.003	0.007037	3.10	0.002
Aprofit_lag1	-	-	-	0.002194	2.15	0.032
Ncntv	0.004797	1.04	0.298	0.002196	0.50	0.614
Ffee	0.002711	1.28	0.200	0.002460	1.31	0.191
Trt	-0.016010	-0.6	0.546	-0.017990	-0.79	0.430
	Log Likelihood: -96.73, $\chi^2_{(4)}$ = 149.34 Prob > $\chi^2_{(5)}$ = 0.0002			Log Likelihood: -74.95, $\chi^2_{(5)}$ = 156.07 Prob > $\chi^2_{(5)}$ = 0.000		

The dichotomous dependent variable in these two probit models is coded one for a decision to accept the contract offer and zero for a rejection. Independent variables include *apropt*, which is the agent's profit given the optimal effort level. *aprofit_lag1*, used in model b only, is the one-period lagged value of the agent's realized profits. *ncntv*, and *ffee* are the actual contract offers.

Table 3.8: Profits, Equity and Efficieny, Accepted Contracts

	Agent Profit		Principal Profit		Effort Gift	Equity		Efficiency	
	Offered	Observed	Offered	Observed		Offered	Observed	Offered	Observed
Treatment 1 (*n* = 168)	16.89	14.44	9.91	14.61	.99	.66	.48	.74	.81
Treatment 2 (*n* = 153)	17.32	11.90	3.75	11.60	1.40	.97	.52	.59	.65

Effort Gift represents the mean difference from the optimal effort with a positive number reflecting a benefit to the principal. Equity is the share, either offered or observed that is retained by the agent, and efficiency the proportion of the available total.

Two specifications are presented with the results reported in terms of the marginal effect of each independent variable on the change in the probability of contract acceptance. Each marginal effect is calculated at the mean value of the independent variable. The model in panel a) includes the contract components *ncntv* and *ffee* and a dummy variable for the treatment (*trtmnt*) as well as a variable *apropt* which measures the agent's profit from the current offer given an optimal (money-maximizing) effort choice. The specification in model b) adds to model a) the variable *aprof_lag1* which captures the one-period lagged realized earnings of the agent. The inclusion of the lagged variable causes the first period decisions to be omitted in model b).

In both models, the treatment variable is not significantly different from zero indicating that the non-parametric tests were conflating differences across treatments with an underlying cause or causes that are correlated with the treatments. Thus, the probit models reveal that contrary to what might be expected from the non-parametric results identical contracts across treatments have the same probability of being accepted. The probit results demonstrate that the agent's acceptance is strongly correlated with monetary income as reflected in the current contract offer. In addition, when the lagged variable is included there is evidence of a reciprocity effect, with higher agent earnings in the previous round increasing the probability of acceptance.[68] After controlling for the value of the contract offer through *apropt*, the contract components themselves are not significant in either model, indicating that the monetary outcome and not the process

[68] Additional lags as well as variable that accumulated profits and profit shares for all rounds turned out not to be significant in other specifications. Also, the period of play was not a significant factor after controlling for the offer value (results omitted).

through which that outcome was achieved, as reflected in the structure of the contract, was responsible for the contract acceptance decision.

Of the sixty-three contracts rejected in treatments 1 and 2, sixteen or roughly one-quarter satisfied the participation constraint for a myopic money-maximizing agent. If the agent had responded with the optimal effort level these contracts would have yielded an agent share of 12 percent and an efficiency of 75 percent, with the agent share measured as *apropt*/(*apropt* + *ppropt*) and efficiency measured as (*apropt* + *ppropt*)/36.[69] The corresponding efficiency figures for the accepted contracts in treatments 1 and 2 are, 81 percent, and 67 percent. The dramatically lower offered share is due primarily to the much lower fixed fee among the rejected contracts. As seen in Table 3.6, the fixed fees for accepted (rejected) contracts are 3.20 (-9.44) and 4.05 (-11.04) in treatment 1 and 2 respectively. Thus the accepted contracts do provide an additional "gift" to the agent beyond that revealed in the initial discussion of aggregate offers. By treatment, the offered shares are 66 percent (97 percent) for treatment 1 (treatment 2). The near total allocation to the agent in treatment 2 suggests that principals may do quite poorly in this treatment. We turn next to the agent's effort responses to accepted contracts to explore the relationship between these contract offers and the final allocations.

As discussed in section 3.3, conditional on acceptance, effort choices of money-maximizing players are equal to the incentive provided.[70] Table 3.8 presents both the optimal profits and the realized profits of principals and agents for treatments 1 and 2 as

[69] Given b=6, b^2 = 36 is the maximal output since $b^2/2$ is the efficient output for each of the two tasks. The variable *ppropt* is the principal(s') profit at the optimal effort level for a money-maximizing agent.

[70] An exception to this optimality result arises when the incentive for a task is less than $\underline{e} = 1$ and the fixed fee, or profits derived from the other task makes the contract acceptable to the agent. All variables that

well as the levels of efficiency and equity at the optimal and observed effort levels. In both treatments we see a difference between the optimal and observed profits that directs additional profit to the principal at the agent's expense. Note that, in aggregate the contract offer is a gift exchange contract which offers the agent significant surplus. These offers are reciprocated in treatments 1 and 2 by effort levels that exceed the agent's money-maximizing best reply.

Non-parametric tests strongly reject the hypothesis that effort levels are consistent with the theoretical optimal values, with the Wilcoxon test yielding p-value of .002 (.001) for treatment 1 (2). The efforts provide greater rewards to the principals than the optimal efforts and these rewards come at a monetary cost to the agent, shifting the split of the surplus to near a 50 percent share for the agent in both treatments.[71] The increases in equity in monetary terms, while costly to the agent do result in additional efficiency since incentives below the optimal level are provided in both treatments.

In aggregate, the contract offers place the agent's in the situation of advantageous inequality and so, if inequity-averse, the gamma parameter is utility relevant. Calculating the optimal gammas implied by the aggregate behaviors is straightforward given the results in equation (3.21). We find

(3.18) $\gamma_0 = \dfrac{e-\alpha}{b+e-2\alpha} = .23$ in treatment 1.

For treatment 2 two possible calculations depend on whether the principal's payoffs are considered in aggregate, or whether payoff comparisons are made with each

---

report optimal values such as *apropt* and *ppropt*, incorporate an adjustment to reflect that $e^* = \underline{e}$ in this case.

[71] This does not imply an even share in treatment 2 since two principals share the remaining 48 percent.

principal, as specified in equation (3.13). In the first case, the calculation is the same as for treatment 1, and we find $\gamma_0 = .25$. If the utility specification requires direct comparison with each principal the calculation yields,

(3.19) $\gamma_0 = \dfrac{2(e-\alpha)}{b+2e-3\alpha} = .40$ in treatment 2.

If the inequity aversion parameter is assumed to be a stable part of individual preferences, the random assignment suggests that gamma should be constant across treatments. Thus the calculations suggest that in treatment 2, the agents are comparing their own payoffs to the joint payoffs of the principals and not individually to each of the principal's payoffs as hypothesized in equation (3.13). This result is reflected in the final payoff shares which yield roughly one half of the surplus to the agent in both treatment 1 and treatment 2.

Further examination of agent effort choices is accomplished through a random-effects tobit model. In this model effort is explained by the contract components, and by an indicator variable to examine the possibility of changes in behavior in the last round of play. The tobit specification is used due to the bounds on effort in the experimental design. The model is specified as

(3.20) $\text{effort}_{it} = \beta_0 + \beta_1\,\text{ncntv}_{it} + \beta_2\,\text{ffee}_{it} + \beta_3\,\text{last}_{it} + u_i + e_{it}.$[72]

The error term consists of two components, with u_i a random disturbance that is constant through time for the *i*th group that interacts repeatedly. The contract components $ncntv_{it}$ and $ffee_{it}$ are as previously defined and our expectation is that the coefficient on

[72] A related model in which the dependent variable is in the payoff space of the principal was also estimated, using either *last* or a trend variable for all periods to look at how the principal was rewarded during the game. A reputation-based model would imply that the rewards to the principal through extra effort should decrease with time as the returns to cooperation diminish. Discussion of these results will accompany the discussion of the tobit estimates.

ncntv is equal to one and the coefficients on *ffee* and *constant* are equal to zero when players myopically maximize their monetary returns. The variable *last* is one in period eight and zero otherwise, and is expected to be negative if individuals defect from cooperative behavior in the last round of play.

The results of the tobit model are presented in Table 3.9 for both treatments. Likelihood ratio tests support the hypothesis of a structural difference within treatments that depends on whether the contract offer places the agent behind (delta) or ahead (gamma) of the principal with respect to monetary payoffs. Differences in agent responses across delta and gamma contracts are predicted by the model of inequity-aversion and, depending on beliefs, by the intention-based model as well. However, not all the coefficients are easily interpreted in light of these models.

In treatment 1 *ncntv* is significantly greater than one for the delta contracts (p = 0.055), which is contrary to the predictions of static models for money-maximizing and inequity-averse players. Since the monetary rewards to the principals' declines during the course of the game this behavior is consistent with a reputation based model.[73]

The gamma contracts in treatment 1 are inconsistent with the static money-maximizing predictions since effort is positively correlated with the fixed component of the contract. Examining the monetary returns to the principal, which are increasing over the first seven rounds, implies that the reputation model also cannot explain behavior in this setting which is therefore most consistent with the equity and reciprocity models. In addition, there is no evidence of defection in the last round of play, a result that is discussed further, below.

[73] The result on declining payoffs to the principal holds regardless of whether the last round of play is included in the reward regression.

In treatment 2 the coefficient on incentives is significantly greater (less) than one for the delta (gamma) contracts, contrary to expectations. One explanation for this counterintuitive result arises from the intercept term in the two models which is negative (positive) for the delta (gamma) contracts. This suggests that the large fixed components in the gamma contracts are reducing the marginal impact of incentives on effort. As in treatment 1, the agents respond positively to the fixed fee when faced with a gamma contract. In contrast to treatment 1 there is a significant reduction in effort in the last period of play for the delta contracts.

Table 3.9: Agent Efforts - Tobit Model

Dependent Var.: Effort	Treatment 1		Treatment 2	
Model	delta (n = 74)	gamma (n = 84)	delta (n = 69)	gamma (n = 83)
ncntv	1.118 (0.000)	1.152 (0.000)	1.320 (0.000)	0.747 (0.000)
fixed fee	-0.003 (0.876)	0.113 (0.000)	0.009 (0.800)	0.051 (0.010)
Last	.267 (0.271)	0.794 (0.897)	-0.800 (0.044)	-1.152 (0.271)
Constant	-0.789 (0.035)	-0.763 (0.485)	-1.377 (0.013)	1.817 (0.130)
	$\chi^2_{(3)} = 586.79$	$\chi^2_{(3)} = 81.14$	$\chi^2_{(3)} = 239.37$	$\chi^2_{(3)} = 39.37$

A tobit model to estimate effort levels uses the contract offer (*ncntv* and *fixed fee*) as explanatory variables, as well as dummy variable *Last*, that is one for the eighth round of play and zero otherwise.

Further insight into the differences in behavior across treatments and contract types can be gained by considering contract formation across time. Figure 3.3 and Figure 3.4 present the aggregate incentives and fixed fees for treatments 1 and 2 across the eight rounds of play. Statistical tests confirm what casual observation suggests; contract choices trend in treatment 1 with the power of incentives increasing and the fixed component declining with repetition so that the contract structure moves towards the stage game prediction for players with *M* preferences. In treatment 2 there is no discernible trend for either contract instrument. Statistical confirmation of this finding comes from random effects tobit models in which the dependent variable is the contract instrument, either *ncntv* or *ffee* and *period* is the independent variable which takes on values from 1 to 8. The tobit results for treatment 1 show that aggregate incentives on average rise .36 (.33) units per period from period 1 levels of 4.39 (4.70) for all (accepted) contracts. The fixed fee falls by 1.82 (1.59) per period from a baseline of 10.42 (10.39). As a result there are significantly fewer gamma contracts in which the agent is offered the advantage in the last four rounds of treatment 1 as compared to the first four rounds (Mann-Whitney U p=.0043). For treatment 2 both the tobit estimates and the Mann-Whitney test reflect that contract types are stable across time.

Figure 3.3 Power of Incentives by Treatment and Period

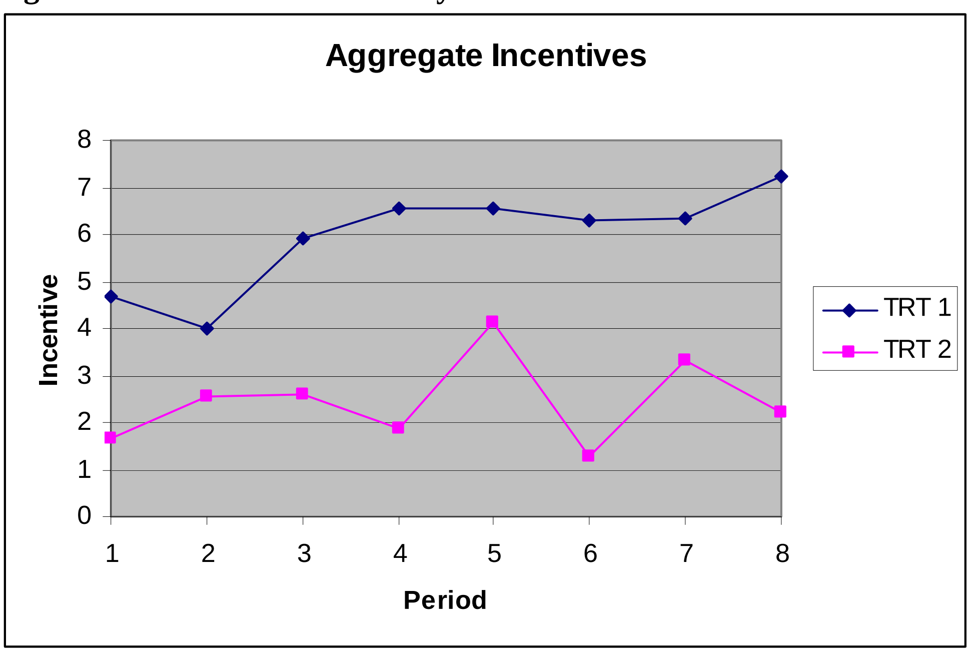

Figure 3.4 Fixed Fee by Treatment and Period

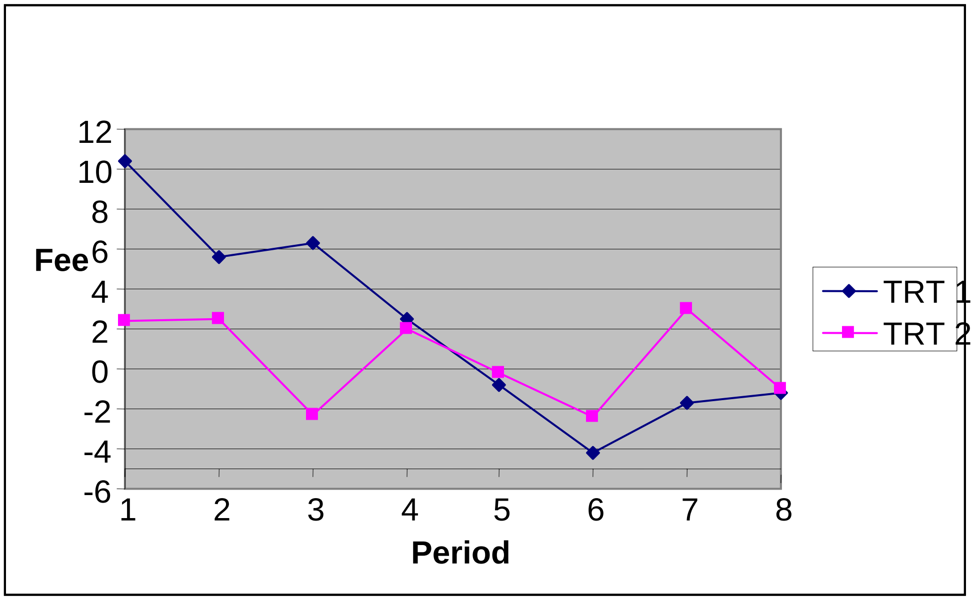

Despite the trends in contract creation in treatment 1 there are no discernible differences with regard to realized equity, efficiency, or agent effort levels when comparing the first and last four rounds of play. Thus the higher-powered contracts created in later rounds yield the same distribution of profits overall, but conditional on acceptance, constrain the flexibility of the agent's effort response, reducing the possibility of defection in later rounds.

Further insight into end of game effects results from an examination the agent's monetary reward to the principal in the final round of play. Wilcoxon tests that control for the dependence of observations by comparing the mean reward in the first seven rounds with the reward in the final round, within each interacting group indicate that rewards to the principal are significantly lower in the final round of play ($p = .021$). Consistent with the hypothesis that contracting in the single principal setting reduces the possibility of defection, the lower rewards to the principal in the final round attributable entirely to the common agency game, with the p-values equal .187 (.049) in the single-

(two-) principal games. As discussed in relation to the results reported in Table 3.9 the defection in treatment 2 is associated with the delta contracts.

3.6.2 Individual Contracts and Efforts

The aggregate results suggest that the challenge in treatment 1 was for the principal to assess the extent of inequity aversion of the agent, and that this was done with some success by offering gift exchange contracts, with the gift diminishing during the course of play. Principals in treatment 2 faced more severe challenges as evidenced by the lower earnings in these treatments. The need to coordinate both incentive provision and the distribution of the surplus through the fixed fee led to greater rejection rates, and when accepted, to lower efficiency and lower surplus shares for the principals. In this section we turn to a more detailed study of the contracting and effort choices of the groups that interacted repeatedly during the experimental sessions in order to understand how players addressed these challenges.

The first goal of this section is to examine in detail how the principals used the information revealed by agent's efforts to make inferences about inequity aversion parameters. We proceed in this task by estimating a model of the learning that takes place within an interacting group, by identifying best responses to previous period efforts, and examining the extent to which contracting behavior is consistent with the best response. We also investigate in more detail the challenges specific to treatment 2. By examining the underlying components of the aggregate contract we gain insight into the nature of the difficulty in generating full incentives. This analysis distinguishes *nash-like* contracts from other contract choices, where nash-like contracts are defined as those in which the principal creates positive incentives for his own task and disincentives for the other consistent with the theoretical predictions developed in section 3.3. We investigate the dynamics of these contract offers in order to generate a better understanding of the common agency setting.

The model we implement is based on identifying categories of contract/effort responses associated with three different modalities through which the principals' can learn about the agent's inequity aversion parameters. The *learning modalities* include optimal response, and non-optimal responses to either delta or gamma contracts.[74] We define a best response to any of these efforts as a response that imposes full incentives and reinstates the distribution consistent with the inequity aversion parameter revealed by the agent's effort choice. Thus if learning is complete and immediate, the agent should accept a best-response offer, and achieve efficiency.

For each contracting unit, the observed change in the principal's contract offer is the pair, $ncntv_t - ncntv_{t-1}$, and $ffee_t - ffee_{t-1}$, which are the dependent variables in a two-equation model that estimates best response dynamics. The independent variable in each equation is $opt\_ncntv_t – ncntv_{t-1}$, and $opt\_ffee_t – ffee_{t-1}$. The two equation model is

$$ncntv_t - ncntv_{t-1} = \beta_0^n + \beta_1^n (opt\_ncntv_t - ncntv_{t-1}) + e_t^n$$

$$ffee_t - ffee_{t-1} = \beta_0^f + \beta_1^f (opt\_ffee_t - ffee_{t-1}) + e_t^f$$

and we assume that there is contemporaneous correlation with $E\left(e_t^n e_t^f\right) = \sigma_{nf}$ and so estimate the parameters as a seemingly unrelated regression (SUR). A Breusch-Pagan test confirms the appropriateness of this specification.

[74] A fourth learning modality, contract rejection, provides information on the lower bound of inequity aversion parameters, but that information is not included in the best response model, since there is no well-defined best response to a contract rejection.

In Table 3.10 we present the evidence on the importance of the learning modalities for each interacting group by providing a count of the number of instances of each type. Also included in the table is the variable *learn* which can take the value of 0, 1, or 2 depending on the results of the SUR. When β_1^n and β_1^f are both equal to zero, *learn* = 0, and it is equal to 1 or 2 when one or both of those coefficients are different from zero at the 10 percent level of confidence, or better.

Table 3.10 Learning Modalities by type and Best Response Learning Outcomes

Treatment 1

Group ID	*Reject*	*Optim*	*Delta*	*Gamma*	*Learn*
1	1	2	1	4	1
2	0	2	2	4	0
3	1	3	1	3	1
4	1	3	0	4	1
5	0	6	2	0	1
6	2	3	2	1	2
7	0	0	0	8	0
8	1	6	1	0	1
9	2	5	1	0	0
10	0	1	4	3	2
11	0	3	0	5	2
12	0	6	1	1	2
13	0	1	0	7	2
14	0	6	1	1	1
15	0	2	2	4	0
16	0	2	0	6	2
17	2	3	0	3	2
18	1	7	0	0	1
19	0	8	0	0	0
20	0	8	0	0	0
21	3	3	0	2	0
22	2	5	1	0	0
23	0	0	0	8	0
Total	*16*	*85*	*19*	*64*	*21*

Treatment 2

Group ID	*Reject*	*Optim*	*Delta*	*Gamma*	*Learn*
24	1	4	2	1	1
25	4	1	1	2	2
26	2	2	1	3	1
27	0	8	0	0	2
28	2	5	1	0	1
29	3	3	0	2	1
30	2	3	2	1	1
31	2	6	0	0	1
32	2	6	0	0	1
33	1	3	0	4	1
34	1	2	0	5	2
35	2	6	0	0	0
36	1	1	0	6	1
37	3	0	1	4	2
38	0	2	2	4	1
39	2	0	3	3	1
40	6	2	0	0	0
41	4	0	1	3	2
42	0	2	6	0	1
43	3	2	1	2	0
44	3	3	0	2	1
45	0	4	2	2	1
46	3	2	0	3	1
47	0	0	0	8	1
48	0	3	2	3	2
Total	*47*	*70*	*25*	*58*	*28*

The table displays the number of times different agent behaviors are observed. *Reject* is contract rejection, *Optim* an optimal response to incentives, *Delta* a non-optimal response to disadvantageous inequality, and *Gamma* a non-optimal response to advantageous inequality. *Learn* identifies the number of times the parameter estimates from the best-response learning model are different than zero. *Learn* has a minimum of 0 and a maximum of 2.

Of the learning modalities in, the modal choice is to respond optimally to the contract offers. Neglecting the rejected contracts, which we have shown are more prevalent in treatment 2, a chi-square test of proportions fails to reject the null hypothesis of no difference in the distribution of the learning modalities across treatments (p=.393). Testing the amount of learning across treatments yields a different result, with learning in treatment 2 significantly greater than in treatment 1 (p=.038).

Table 3.11 presents the differences in *learn* by treatment and it can be seen that only three principals exhibit no learning in treatment 2 while nine do not learn in treatment 1. As reported in Table 3.12, a comparison of payoffs in periods 1-4 with those in periods 5-8 reveals that principals profits increased dramatically among those with *learn* =2 in treatment 2 (Mann Whitney, p=.047). By contrast, the three non-learning principals in treatment 2 all reduced their profits marginally over the course of the session, and in all other cases, there was no significant change.

Table 3.11 Learning Types by Treatment

		Learn	
	0	**1**	**2**
Treatment 1	9	7	7
Treatment 2	3	16	6

Learners are identified as 0,1, or 2 based on the number of coefficients that are significant from the individual SUR that measures best responses to agent choices.

Table 3.12 Learning Effect on Profits

	Learn		
	0	**1**	**2**
Treatment 1 (*n* = 23)	0.56	13.14	-0.57
Treatment 2 (*n* = 25)	-3.00	-1.31	28.83*

* Identifies a significant change in profits for the principals (p<.05) from the first four to the last four rounds of the session. All other changes are not statistically significant at conventional levels.

We next consider the adjustment process for the disaggregated contract choices in the two-principal treatment to assess how the response to the contract offer of the other principal affects profits. Because of the symmetric situation faced by principals in treatment 2, we present disaggregated information on the incentive choices of principals as *ncntv_own* and *ncntv_oth* with the reference in the variable label to one's own and to the other principal's task. To add some structure to the diverse contract strategies we partition the contracts according to whether the individual's contract choice is *nash-like*, where nash-like is defined as a contract in which *ncnt_own*>0 and *ncnt_oth*<0. An aggregate contract constructed from the choices of both players may consist of 0, 1, or 2 nash-like contract choices. Table 3.13 shows that the aggregate incentives are declining in the number of nash-like components. Table 3.13 also presents figures on the contract components and we see that when two players execute the nash-like strategy, the incentives are close to zero as each player obliterates the incentives provided by the other. While these strategies are inconsistent with the incentives under common agency for an agent with *M* preferences, they are consistent with a key prediction of the inequity aversion theory under common agency. As discussed in relation to Table 3.4 and illustrated in Figure 3.2 principals who compete through their incentive offers will reduce incentives as the agent's inequity aversion increases. Their *ncnt_own* component moves towards zero and *ncnt_oth* becomes more negative.

While the nash-like strategy is consistent with the theory, it is chosen for only 115 of the 400 contract choices or 29 percent of the time. The non-nash-like strategies that make up the balance of the contract provide additional incentives for the agent by making, in aggregate, choices of both *ncnt_own* and *ncnt_oth* that are greater than or equal to zero. These contracts lead to fewer rejections providing larger profits than the nash-like choices. Although these contracts are not nash-like they also do not have the same structure as equilibrium contracts arising from the solution associated with the iterated elimination of dominated strategies. The observed contracts incorporate *ncnt_own* (*ncnt_oth*) that are less than (greater than) predicted by this solution concept, but are consistent with offering reciprocal motivations. These represent an alternative and somewhat successful attempt to mitigate the inefficiency that results from the two-stage solution.

Table 3.13 Contract Components for Treatment 2

	ncnt_own		***ncnt_oth***		***ncntv***	**Share Accepted**
Nash-like Number	Nash-like	Not Nash-like	Nash-like	Not Nash-like		(%)
0 (n = 212)		.53		1.12	3.30	83
1 (n = 146)	2.59	.43	-2.56	1.39	1.84	73
2 (n = 72)	2.76		-2.60		0.33	57
Total (n = 400)	2.65	.50	-2.57	1.19	2.45	77

The contract components for nash-like and non-nash-like contracts are presented in aggregate according to the number of nash-like contracts in a given contract pair. Aggregate incentives and the share of accepted contracts are declining in the number of nash-like contractors.

3.7 Conclusion

The principal agent setting we have investigated provides a diverse strategy space that nests the typical gift exchange game as a special case. In the single principal case we observe inefficient gift exchange contracts at the outset with a clear trend towards more efficient contracting in the later periods. The usefulness of the gift exchange contracts is predicated on the fact that there is some underlying probability that agent's have other-regarding preferences, and in particular are inequity-averse. The common agency treatment by contrast appears trendless, but underlying the apparent stasis is a diversity of strategic choices and payoffs. A substantial minority of players (29 percent) choose strategies that have the general structure of incentives predicted by the common agency theory. These players do poorly in general and those pursuing an alternative strategy that avoids punishing the other principal fare better.

Appendix 3.1: Proof of Theoretical Results

Proof of Result 2: Equilibrium incentives and efforts are identical for *I* and *M* preferences. The first order condition for the agent under inequity aversion yields

$$(3.21)\quad e^{*I} = \frac{\alpha + 2\delta\alpha - \delta b}{1+\delta}.$$

Maximizing total social surplus requires that the principal choose incentives so that

$$(3.22)\quad \alpha^{*I} = \arg\max_{\alpha} be^{*I} - \frac{\left(e^{*I}\right)^2}{2}$$

Substituting for e^{*I}, the principal's first order condition yields

$$\frac{b+2\delta b}{1+\delta} - \left(\frac{\alpha + 2\delta\alpha - \delta b}{1+\delta}\right)\left(\frac{1+2\delta}{1+\delta}\right) = 0 .$$

Algebraic manipulation rapidly reduces this expression to $\alpha = b$ which was to be shown. The sufficient condition is satisfied since

$$\frac{\partial^2 u_1\left(e^{*I}\right)}{\partial\alpha^2} = -\frac{2\delta}{1+\delta}\left(\frac{1+2\delta}{1+\delta}\right) < 0 \ \ \forall\ \delta > 0 .$$

Substituting α^{*I} in (3.21) yields $e^{*I} = b$ proving the second part of result 2.

Proof of Result 4: An agent with *I* preferences will respond to a *GE* contract with $e > e^*$.

The utility of an agent with *I* preferences when his payoff exceeds the principal's is $u_0 = \pi_0 - \gamma_0\left(\pi_0 - \pi_1\right)$, and the optimal effort choice is given by

(3.23) $e_{\gamma}^{*I} = \dfrac{\alpha + b\gamma_0 - 2\alpha\gamma_0}{1-\gamma_0}$.

When the gift exchange contract is offered, α is less than b and so e_{γ}^{*I} can be reformulated as $e_{\gamma}^{*I} = \alpha^{GE} + \dfrac{\varepsilon\gamma_0}{1-\gamma_0}$, where $\varepsilon = b - \alpha^{GE} > 0$. Thus, $e_{\gamma}^{*I} > \alpha^{GE}$, since $\dfrac{\gamma_0}{1-\gamma_0} > 0$ for $0 < \gamma_0 < 1$. A similar argument reveals that when the agent suffers disadvantageous inequality the effort choice is

(3.24) $e_{\delta}^{*I} = \alpha^{GE} - \dfrac{\varepsilon\delta_0}{1+\delta_0}$

and the agent punishes the principal with a reduced level of effort if the participation constraint is satisfied.

Appendix 3.2: A Model of Reciprocal Preferences

Utility for reciprocity based preferences, *R*, is defined as follows:

$$u_i^R(a, a_i', a_i'') = Z_0\left[\pi_i(a_i, a_j) + \rho_i \varphi_i(a', a_i''; a_j)\sigma(a_i, a_i', a_i'')\right]$$

As with the model of inequity aversion, the utility functional contains a term representing the individual's monetary reward. Added to the material payoff, is an expression that captures the content of the reciprocal preferences - the product of perceived kindness, $\varphi_i(a_i', a_i''; a_j)$ and reciprocity, $\sigma_i(a_i, a_i', a_i'')$, weighted by $\rho_i \geq 0$ that captures the strength of reciprocal preferences. The multiplicative term implies that responding to kindness (unkindness) with reward (punishment) increases utility. The single and double primed terms represent first and second order beliefs about actions. First order beliefs, a_i', represent the beliefs of player *i* about *j*'s actions. Second order beliefs, a_i'', represent what player *i* believes player *j* believes about player *i*'s actions.

The *kindness* term $\varphi_i(a_i', a_i''; a_j)$ is the product of two components, $\Delta(a)$ associated with outcomes and ϑ, the intention factor, with $0 \leq \vartheta \leq 1$. The intention factor weights the outcomes according to whether the strategic setting allows the intentions of the other player to be assessed. The extensive choice alternatives in our setting imply $\vartheta = 1$ (see Falk and Fischbacher (2006) for additional discussion). The outcome term reflects the difference between the *i*'s maximum material payoff given player *j*'s actions, $\pi_i^*(a_j)$ and the material payoff offer based on one's beliefs about *j*'s intentions given the strategic setting $\pi_i^*(a_i', a_i'')$.

$$\Delta(a_i', a_i''; a_j) = \pi_i^*(a_j) - \pi_i^*(a_i', a_i'')$$

Positive kindness reflects an offer of material reward by *j* that exceeds player *i*'s expectations. The reciprocity term $\sigma_i(a_i, a_i', a_i'')$ reflects the material payoff to player *j* due to player *i*'s actions less the material payoff to *j* available from *i*'s optimal money-maximizing response to *j*'s offer.

$$\sigma_i(a_i, a_i', a_i'') = \pi_j(a_i, a_i', a_i'') - \pi_j^*(a_j)$$

Thus $\sigma_0(a_0, a_0', a_0'') > 0$ reflects an agent response that benefits the principal materially. This is a rational response for the agent only when $\Delta(a_0', a_0''; a_1) > 0$.

Appendix 3.3: Experimental Instructions

<u>Instructions</u>

Welcome. You are participating today in an experiment in economic decision-making. Depending on your decisions during the experiment you can earn a substantial amount of money that will be paid to you in cash at the end of the session.

During the experiment your income will be counted in points. The exchange rate of points into dollars is

1 point = $ .25

At the end of the experiment, all the points you have earned will be converted to dollars and paid to you in cash.

Please note that during the experiment talking is not allowed. If you do have questions please raise your hand. We will answer your questions in private.

What will you be doing during this experiment?

1. Introduction

   In this experiment you will participate in a task that involves two decision-makers. We will call the decision-makers participants X and Y.

   You belong to the group of participants Y.

   Thus, during the entire experiment, you will make decisions in the role of Y.

   At the beginning you will be **randomly** matched with a participant of group X. You will be paired with this participant throughout the session. Your decisions will be transmitted via the computer to participant X. This participant will only get informed about your decision. (S)he will never learn your name or your participant number etc. That is, your decisions remain **anonymous.**

2. An overview of the experiment.

   It may help to understand your task if you think of the following situation. Participant X will create a contract to compensate participant Y for producing "gizmos." There are two distinct tasks that Y must undertake to produce a gizmo. The tasks will be called **Task1** and **Task2.** The effort that participant Y chooses to exert on the tasks determines the total number of gizmos produced. The effort devoted to **Task1** is **E1** and the effort devoted to **Task2** is **E2**. There are also costs associated with Y's efforts. The contract created by X can affect how many gizmos are produced and also how much of the production is kept by X and how much by Y.

3. Creating a contract:

   Participant X has two tools to create the contract. X can offer a "**multiplier**" to Y for each of the tasks. We will call the multiplier for Task1, **M1** and the multiplier for Task2, **M2**. The multipliers can be positive, negative, or zero.

   The way the multiplier works is that the efforts chosen by Y times the multipliers chosen by X are paid to Y. Y is thus compensated for their effort through the multipliers.

   Multiplier amount paid to Y **M1*E1 + M2*E2**

   Note that X chooses the M's and Y chooses the E's.

   The second contract tool is a fixed payment or fee that is paid to Y no matter what effort Y chooses for the tasks, it can be positive negative or zero. The Fixed Fee is called **FF**

   Fixed Payment amount paid to Y **FF**

   The total payment to Y from X is the sum of the two portions:

   Contract: **M1*E1 + M2*E2+ FF**

   In a minute you will be able to try out some examples to see how your earnings are affected by your X's contract choice and by your effort choices. An online profit calculator will be provided so that you do not have to make these calculations yourself.

Session Parameters:

In this session the multipliers and fixed components must satisfy the following conditions

Multipliers, **M1** and **M2** can range from **–6 to 6**, including zero (integers only).
The Fixed Fee, **FF** can range from **-36 to 36**, including zero (integers only).

Y's effort levels, **E1** and **E2** can range from **1 to 6** (integers only)

The contracting game will be played a total of ten times.
Each round consists of the following **stages**:

1. X makes a contract offer in accordance with the rules above
2. Y evaluates the contract and chooses to accept or reject it.
3. a) If Y rejects the contract, both X and Y receive payments of zero points and the round is over.
   b) If Y accepts the contract Y makes an effort choice for each Task.
4. If the contract is accepted by Y, X's contract choice and Y's effort levels determine the payoffs for X and Y in that round.

Profits for X and Y are calculated as follows:

NOTE: An online profit calculator will do these calculations for you

==

Profit X = Total Output – Payment to Y

Total Output = **6*E1 + 6*E2**   Payment to Y = **M1*E1 + M2*E2+ FF**

Therefore X's payoff is:

6*E1 + 6*E2 – (M1*E1 + M2*E2+ FF)

which simplifies to

(6-M1)*E1 +(6-M2)*E2 - FF

==

Profit Y = Payment from X - Effort Costs

Payment from X = **M1*E1 + M2*E2+ FF**   Effort Costs = **C1 + C2**

So Y's payoff is :

M1*E1 + M2*E2+ FF - (C1 + C2)

C1 is the cost effort devoted to Task1 and C2 is the cost of effort devoted to Task2

Table 1 Effort Costs: Costs are the same Task1 and Task2

Effort	1	2	3	4	5	6
C1	.5	2	4.5	8	12.5	18
C2	.5	2	4.5	8	12.5	18

Summarizing the above:

Payoff to X = Total Output – Payoff to Y
= (6-M1)*E1 +(6-M2)*E2 - FF

Payoff to Y = Payment from X – Costs of effort
= M1*E1 + M2*E2 – C1 – C2 +FF

X chooses, **M1, M2**, and **FF**

$-6 \leq M1, M2 \leq 6$

$-36 \leq FF \leq 36$

Y chooses, **E1,** and **E2**

$1 \leq E1, E2 \leq 6$

Only integer numbers are allowed.

4. The Experimental Details

The contracting scenario will be repeated on the computer eight times. These rounds will be preceded by two practice rounds to familiarize you with the computerized environment and with the rules of the game. You will not be paid based on the outcome of the practice rounds. The points earned in each of the eight "paying rounds" will be added together, converted to dollars at the rate of **1 point = $.25.**

Participant X will have 5 minutes to create a contract and Y 3 minutes to choose efforts in the first rounds. This time will be reduced slightly as the experiment continues in order to allow you to play more rounds (and earn more points). There will always be an announcement when there is one minute left for X or Y.

We will now move to the computerized portion of the session which consists of three parts.

1. Two practice rounds in which will introduce you to contract creation and the profit calculator. In the first round you will work with a single task. In the second round there will be two tasks exactly as in the paying rounds. You will have 5 minutes in each of these practice rounds to explore different contracts and effort levels

2. Eight paying rounds in which your earnings will be recorded.

3. A short questionnaire at the end.

Good luck and thank you for your participation.

References

Alevy, J. E. "Depth of Reasoning among Financial Market Professionals." *Mimeo, University of Maryland*, 2006.

Alevy, J. E., J.A. List, W. Adamowicz. "More is Less Preference Reversals and Non-Market Valuation." *University of Maryland College Park Working Paper*, 2003,

Anderhub, V., S. Gachter, M. Konigstein. "Efficient Contracting and Fair Play in a Simple Principal-Agent Experiment." *Experimental Economics*, 2002, *5* (1), 5-27.

Anderson, Lisa R. Information cascades. 1994. Dissertation, Universty of Virginia.

Anderson, Lisa R., Charles A. Holt. "Information cascades in the laboratory." *American Economic Review*, 1997, *87* (5), 847-862.

Anscombe, F. J., R. Aumann. "A Definition of Subjective Probability." *Annals of Mathematical Statistics*, 1963, *34* 199-205.

Avery, Christopher, Peter Zemsky. "Multidimensional uncertainty and herd behavior in financial markets." *American Economic Review*, 1998, *88* (4), 724-748.

Banerjee, Abhijit V. "A simple model of herd behavior." *Quarterly Journal of Economics*, 1992, *107* 797-818.

Barberis, Nicholas, Andrei Shleifer, Robert Vishny. "A model of investor sentiment." *Journal of Financial Economics*, 1998, *49* 307-343.

Baron, D. P. "Noncooperative Regulation of a Nonlocalilzed Externality." *Rand Journal of Economics*, 1985, *16* (4), 553-568.

Berkman, Henk. "Large option trades, market makers, and limit orders." *Review of Financial Studies*, 1996, *9* (3), 977-1002.

Berndt, Ernst, Bronwyn H. Hall, Robert E. Hall, Jerry A. Hausman. "Estimation and inference in nonlinear structural models." *Annals of Economic and Social Measurment*, 1974, *3* 653-665.

Bernheim, B. D. and M. D. Whinston. "Common Agency." *Econometrica*, 1986, *54* (4), 911-30.

Bernheim, B. D. and M. D. Whinston. "Menu Auctions, Resource Allocation, and Economic Influence." *Quarterly Journal of Economics*, 1986, *CI* (1), 1-31.

Bikhchandani, Sushil, David Hirshleifer, Ivo Welch. "Learning from the behavior of others: Conformity, fads, and informational cascades." *Journal of Economic*

Perspectives, 1998, *12* (3), 151-170.

Bikhchandani, Sushil, David Hirshleifer, Ivo Welch. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy*, 1992, *100* (5), 992-1026.

Bikhchandani, Sushil, Sunil Sharma. "Herd behavior in financial markets: A review." *IMF Working Paper 48*, 2000,

Brown, M., A. Falk, E. Fehr. "Relational Contracts and the Nature of Market Interactions." *Econometrica*, 2004, *72* (3), 747-789.

Butler, John S., Robert Moffitt. "A computationally efficient quadrature procedure for the one-factor multinomial probit model." *Econometrica*, 1982, *50* 761-764.

Camerer, C. "Ambiguity Aversion and Non-additive Probability: Experimental Evidence, Models and Applications," Luini, L., *Uncertain Decisions: Bridging Theory and Experiments*. Boston, Ma.: Kluwer Academic, 1999, 53-79.

Camerer, Colin F. *Behavioral Game Theory*. Princeton, New Jersey: Princetion University Press, 2003.

Celen, Bogachan, Shachar Kariv. "Distinguishing informational cascades from herding behavior in the laboratory." *American Economic Review*, 2004, *94* 484-497.

Celen, Bogachan, Shachar Kariv. "An experimental test of observational learning under imperfect information." *Economic Theory*, 2005, *26* (3), 677-699.

Chamley, Christophe P. *Rational Herds: Economic Models of Social Learning*. Cambridge , U.K.: Cambridge University Press, 2004.

Chamley, Christophe P., Douglas Gale. "Information revelation and strategic delay in a model of investment." *Econometrica*, 1994, *62* (5), 1065-1085.

Chari, Varadarajan V., Patrick J. Kehoe. "Financial crises as herds: Overturning the critiques." *Journal of Economic Theory*, 2004, *119* 128-150.

Charness. Gary, Martin Dufwenberg. "Promises and Partnerships." *UCSB Working Paper*, 2003,

Chow, C. C., and R.K. Sarin. "Comparative Ignorance and the Ellsberg Paradox." *Journal of Risk and Uncertainty*, 2001, *22* (2), 129-139.

Cipriani, Marco, Antonio Guarino. "Herd behavior and contagion in financial markets." *Working Paper, George Washington University*, 2005a,

Cipriani, Marco, Antonio Guarino. "Herd behavior in a laboratory financial market." *American Economic Review*, 2005b, *95* (5), 1403-1426.

Coursey, Don L., John L. Hovis, William D. Schulze. "The disparity between willingness to accept and willingness t o pay measures of value." *Quarterly Journal of Economics*, 1987, *25* (2), 239-250.

Daniel, Kent., David Hirshleifer, Avandihar Subrhamanyam. "Mispricing, covariance risk, and the cross-section of security returns." *Journal of Finance*, 2001, *56* 921-965.

Devenow, Andrea, Ivo Welch. "Rational herding in financial economics." *European Economic Review*, 1996, *40* 603-615.

Diecidue, Enrico, Peter Wakker. "On the intuition of rank-dependent utility." *Journal of Risk and Uncertainty*, 2001, *23* (3), 281-298.

Dixit, A., G.M. Grossman, E. Helpman. "Common Agency and Coordination: General Theory and Application to Government Policy Making." *Journal of Political Economy*, 1997, *105* (4), 752-769.

Dixit, Avinash. "The Making of Economic Policy," *Appendix*. Cambridge: MIT Press, 1996, 157-71.

Dixit, Avinash. "The Power of Incentives in Private versus Public Organizations." *American Economic Review*, 1997, *87* (2), 378-82.

Drehmann, Mathias, Jorg Oechssler, Andreas Roider. "Herding and contrarian behavior in financial markets: An internet experiment." *American Economic Review*, 2005, *95* (5), 1403-1426.

Dufwenberg, Martin, Tobias Lindqvist, Evan Moore. "Bubbles and experience: An experiment." *American Economic Review*, 2005, *95* (5), 1731-1737.

Eichengreen, Barry, Donald Mathieson, Bankim Chandha, Anne Jansen, Laura Kodres, Sunil Sharma . "Hedge funds and financial market dynamics." *IMF Occasional Paper* , 1998, *172*

Ellsberg, D. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, 1961, *75* 643-69.

Epstein, L. "A Definition of Uncertainty Aversion." *Review of Economic Studies*, 1999, *66* (3), 579-608.

Falk, A., U. Fischbacher. "A Theory of Reciprocity." *Games and Economic Behavior*, 2006, *54* (2), 293-315.

Fama, Eugene F. "Market efficiency, long-term returns, and behavioral finance." *Journal of Financial Economics*, 1998, *49* 283-306.

Fehr, Ernst, Klaus Schmidt. "A Theory of Fariness, Competitin, and Cooperation."

Quarterly Journal of Economics, 1999, *108* (2), 817-868.

Fennema, H. and van Assen M. "Measuring the Utility of Losses by Means of the Tradeoff Method." *Journal of Risk and Uncertainty*, 1999, *17* (3), 277-95.

Fischbacher, Urs. "Z-tree toolbox for experimental economics." *University of Zurich: IEER Working Paper*, 1999,

Fox, C. R., Tversky, A. "Ambiguity Aversion and Comparative Ignorance." *Quarterly Journal of Economics*, 1995, *110* (3), 585-603.

Frisch, D., J. Baron. "Ambiguity and Rationality." *Journal of Behavioral Decision Making*, 1988, *1* 149-57.

Gachter, S., A. Falk. "Reputation and Reciprocity: Consequences for the Labour Relation." *Scandinavian Journal of Economics*, 2002, *104* 1-27.

Gale, Douglas. "What have we learned from social learning?" *European Economic Review*, 1996, *40* 617-628.

Garber, Peter. *Famous First Bubbles: The Fundamentals of Early Manias*. Cambridge, Mass.: MIT Press, 2000.

Genesove, David, Christopher Mayer. "Loss aversion and seller behavior: Evidence from the housing market." *Quarterly Journal of Economics*, 2001, *116* 1233-1260.

Ghirardato, Paolo, Massimo Marinacci. "Ambiguity Made Precise: A Comparative Foundation." *Journal of Economic Theory*, 2002, *102* (2), 251-289.

Gigerenzer, Gerd,|David J. Murray. *Cognition as Intuitive Statistics*. Hillsdale, NJ: Erlbaum, 1987.

Gilboa, Itzhak, David Schmeidler. "Maximin expected utility with a non-unique prior." *Journal of Mathematical Economics*, 1989, *18* (141-153),

Goeree, Jacob K., Charles A. Holt, Thomas R. Palfrey. "Regular quantal response equilibrium." *Experimental Economics*, 2005, *8* (4), 347-367.

Goeree, Jacob K., Thomas R. Palfrey , Brian W. Rogers, Richard D. McKelvey. "Self-correcting Information Cascades." *Working Paper: California Institute of Technology*, 2004,

Grossman, S. J., O.D. Hart. "An Analysis of the Principal-Agent Problem." *Econometrica*, 1983, *51* (1), 7-46.

Haile, P., A. Hortaçsu, and G. Kosenok. "On the Empirical Content of the Quantal Response Equilibrium." *Working Paper, Yale University*, 2004.

Harrison, Glenn, John A. List. "Field experiments." *Journal of Economic Literature*, 2004, *42* (4), 1009-1055.

Heath, C., Tversky, A. "Preference and Belief - Ambiguity and Competence in Choiceunder Uncertainty." *Journal of Risk and Uncertainty*, 1991, *4* (1), 5-28.

Hirshleifer, David. "Investor psychology and asset pricing." *Journal of Finance*, 2001, *LVI* (4), 1533-1597.

Hirshleifer, David, Siew Hong Teoh. "Herd behavior and cascading in capital markets: A review and synthesis." *European Financial Management*, 2003, *9* (1), 25-66.

Holt, C., S. Laury. "Risk Aversion and Incentive Effects." *American Economic Review*, 2002, *92* (5), 1644-1655.

Hsee, C. K, G.F. Loewenstein, S. Blount, M.A. Bazerman. "Preference Reversals Between Joint and Separate Evaluations of Options: A Review and Theoretical Analysis." *Psychological Bulletin*, 1999, *125* (5), 576-90.

Huck, Steffen, Jorg Oechssler. "Informational cascades in the laboratory: Do they occur for the right reasons? " *Journal of Economic Psychology*, 2000, *21* 661-671.

Hung, Angela A., Charles R. Plott. "Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions." *American Economic Review*, 2001, *91* (5), 1508-1520.

Kahneman, Daniel, Amos Tversky. "Prospect theory: An analysis of decision under risk." *Econometrica*, 1979, *47* (2), 263-291.

Keynes, J. M. *A Treatise on Probability*. London: Macmillan, 1921.

Keynes, John M. *The general theory of employment, interest and money*. London : Macmillan, 1936.

Kirchsteiger, G., A. Prat. "Common Agency and Computational Complexity." *Tilburg University : Working Paper*, 1999,

Kirchsteiger, G., A. Prat. "Inefficient equilibria in lobbying." *Journal of Public Economics*, 2001, *82* 349-375.

Knez, Peter, Vernon L. Smith, Arlington Willliams. "Individual rationality, market rationality, and value estimation." *American Economic Review*, 1985, *75* (2), 397-402.

Knight, F. H. *Risk, Uncertainty, and Profit*. Boston: Houghtoun Mifflin, 1921.

Koessler, Frederic, Anthony Ziegelmeyer. "Tie-breaking rules and informational cascades: A note." *University Louis Pasteur Working paper* , 2000,

Kubler, Dorothea, Georg Weizsacker. "Are longer cascades more stable? " *Journal of the*

European Economic Association , 2005, *3* 330-339.

Kubler, Dorothea, Georg Weizsacker. "Limited depth of reasoning and failure of cascade formation in the laboratory." *Review of Economic Studies*, 2004, *71* 425-441.

Laussel, D., M. Le Breton. "Complements and Substitutes in Common Agency." *Ricerche Economiche*, 1996, *50* 325-345.

Lee, In Ho. "Market crashes and informational avalanches." *Review of Economic Studies*, 1998, *65* (4), 741-760.

List, J. A. "Preference Reversals of a Different Kind: The "More is Less" Phenomenon." *American Economic Review*, 2002, *92* (5), 1636-1643.

List, John A. "Does market experience eliminate market anomalies?" *Quarterly Journal of Economics*, 2003, *118* (1), 41-71.

List, John A. "Neoclassical theory versus prospect theory: Evidence from the marketplace." *Econometrica*, 2004, *72* (2), 615-625.

Locke, Peter R., Steven C. Mann. "Professional trader discipline and trade disposition." *Journal of Financial Economics*, 2005, *76* (2), 401-444.

Machina, Mark, David Schmeidler. "A more robust definition of subjective probability." *Econometrica*, *60* (4), 745-780.

MacKay, Charles. *Extraordinary Popular Delusions and the Madness of Crowds*. New York: Three Rivers Press, 1980.

Manaster, Steven, Steven C. Mann. "Sources of market making profits: Man does not live by spread alone." *Working paper, Texas Christian University*, 1999,

McKelvey, Richard D., Thomas R. Palfrey . "Quantal response equilibria for extensive form games." *Experimental Economics*, 1998, *1* (1), 9-41.

McKelvey, Richard D., Thomas R. Palfrey. "Quantal response equilibria for normal form games." *Games and Economic Behavior*, 1995, *10* 6-38.

Mukerji, Sujojy, Jean-Marc Tallon. "An overview of economic applications of David Schmeidler's models of decision making under uncertainty," Gilboa, Itzhak, *Uncertainty in economic theory: Essays in honor of David Schmeidler's 65th birthday*. New York: Routledge, 2004, 283-302.

Noth, Markus, Martin Weber. "Information aggregation with random ordering: Cascades and overconfidence." *Economic Journal*, 2003, *113* 166-189.

Odean, Terrance. "Are investors reluctant to realize their losses?" *Journal of Finance*, 1998, *53* (5), 1775-1798.

Olsen, T. E., G. Torsvik. "Intertemporal Common Agency and Organizational Design: How Much Decentralization?" *European Economic Review*, 1995, *39* 1405-1428.

Plott, C., S. Sunder. "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets: An Appliction of Rational-Expectations Models." *Econometrica*, 1988, *56* 1088-118.

Quiggin, John. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization*, 1982, *3* 323-43.

Rabin, M. "Incorporating Fairness Into Game Theory and Economics." *American Economic Review*, 1993, *83* 1281-1302.

Samuelson, Larry. "Economic Theory and Experimental Economics." *Journal of Economic Literature*, 2005, *43* (1), 65-107.

Sarin, R. K., Weber, M. "Effects of Ambiguity in Market Experiments." *Management Science*, 1993, *39* (5), 602-615.

Savage, L. J. *The Foundations of Statistics*. New York: Wiley, 1954.

Schmeidler, D. "Subjective-Probability and Expected Utility Without Additivity." *Econometrica*, 1989, *57* (3), 571-587.

Selten, R. "The Chain Store Paradox." *Theory and Decision*, 1978, *9* 127-59.

Sgroi, Daniel. "The right choice at the right time: A herding experiment in endogenous time." *Experimental Economics*, 2003, *6* (2), 159-180.

Shefrin, Hersh M., Meir Statman. "The disposition to sell winners too early and ride losers too long." *Journal of Finance*, 1985, *40* 777-790.

Shleifer, Andrei, Lawrence H. Summers. "The noise trader approach to finance." *Journal of Economic Perspectives*, 1990, *4* (2), 19-33.

Siegel, S. *Nonparametric Statistics for the Behavioral Sciences*. New York: McGraw Hill, 1956.

Smith, Lones, Peter Sorenson. "Pathological outcomes of observational learning." *Econometrica*, 2000, *68* (2), 371-398.

Smith, V. L. "Measuring Nonmonetary Utilities in Uncertain Choices: The Ellsberg Urn." *Quarterly Journal of Economics*, 1969, *83* 324-9.

Stiglitz, J. E. "Credit Markets and the Control of Capital." *Journal of Money, Credit, and Banking*, 1985, *17* (2), 133-152.

Sunder, S. "Market for Information: Experimental Evidence." *Econometrica*, 1992, *60* (3), 667-95.

Tversky, Amos, Daniel Kahneman. "Judgment under uncertainty: Heuristics and biases." *Science*, 1974, *185* 1124-1131.

Tversky, A., D. Kahneman. "Advances in Prospect Theory: Cumulative Representations of Uncertainty." *Journal of Risk and Uncertainty*, 1992, *5* (4), 297-323.

Vives, Xavier. "Social learning and rational expectations." *European Economic Review*, 1996, *40* (3-5), 589-601.

Wakker, P. "Separating Marginal Utility and Probabilistic Risk Aversion." *Theory and Decision*, 1994, *36* 1-44.

Wakker, Peter. "Uncertainty Aversion: A Discussion of Critical Topics in Health Economics." *Health Economics*, 2000, *9* 261-263.

Wakker, Peter, D. Deneffe. "Eliciting von Neumann-Morgenstern Utilities When Probabities Are Distorted or Unknown." *Management Science*, 1996, *42* (8), 1131-50.

Welch, Ivo. "Sequential sales, learning, and cascades." *Journal of Finance*, 1992, *47* (2), 695-732.

Willinger, Marc, Anthony Ziegelemeyer. "Are more informed agents able to shatter information cascades in the lab?," Cohendet, Patrick, Patrick Llerna, Hubert Stahn, Gisele Umbhauer, *The Economics of Networks: Interaction and Behaviours*. Springer-Verlag, 1998, 291-305.

Yaari, Menahem. "Dual theory of choice under uncertainty." *Econometrica*, 1987, *55* 95-115.

Yaari, Menahem. "Some remarks on measures of risk aversion and on their uses." *Journal of Economic Theory*, 1969, *1* 315-329.