

Choices in an Interdependent Economic Environment:
Inequity Aversion and Bargaining Games

Andy Galbraith

February 2013

*Senior thesis submitted in partial fulfillment of the requirements for a
Bachelor of Science degree in Economics
at the University of Puget Sound*

Abstract

Experimental economics has revealed an underlying tension between preferences for fairness and the purely self-interested behavior assumed throughout economic theory. Though many models attempt to explain observed behavior that is inconsistent with theory, the motivations behind such deviations from theoretic equilibrium are still largely unknown.

This paper investigates preferences for fairness in a controlled laboratory setting using an ultimatum game variant. Payoffs to participants are altered over several treatments to examine if subjects exhibit behavior that is consistent with a preference model which incorporates envy and guilt into an individual's utility function. The results indicate evidence of both guilt and envy in participant behavior, and shed further light on the tension between fairness preferences and self-interest in the ultimatum bargaining game.

Keywords

behavioral economics, experimental economics, other regarding preferences, inequity aversion, ultimatum bargaining game

Economics has many long-held assumptions about the way individuals make decisions. One of these assumptions is that individuals seek to maximize their own utility, and that their utility is a function of purely self-interested factors. All of game theory is built with this assumption at its core, and its applicability and predictive power is contingent on pure self-interest accurately representing how people actually behave. Experimental evidence has put the validity of this assumption under serious scrutiny over the past few decades. Behavioral and experimental economics, growing subfields of the economics discipline, are the result of this systematic questioning of long-held assumptions like that of pure self-interest. The focus of these subfields is to reach a more accurate understanding of what motivates people's behavior.

Experimental economics has revealed an underlying tension between preferences for fairness and the game theoretic equilibrium behavior assumed by classic economic theory. Simple game structures such as the ultimatum and dictator games are commonly used to examine behavioral deviations in different scenarios. In these games, discussed in more detail in Section 1, a proposer and responder must arrive at some division of a sum of money in a simple bargaining scenario. Proposers generally offer more to the responder than theory predicts, and some positive offers are rejected by the responder, leaving both players with nothing. These actions are inconsistent with the equilibrium outcome with the purely self-interested equilibrium.

Alternate models of preferences incorporate factors other than pure self-interest into an individual's utility function, such as a preference for equitable outcomes, a desire to reciprocate a partner's behavior, or a general concern for the well-being of others. These factors are referred to broadly as other regarding preferences. A prominent

model of other regarding preferences is the Fehr-Schmidt inequity aversion model, which posits that both guilt and envy, resulting from a difference between payoffs to subjects, can negatively impact a subject's utility (Fehr and Schmidt, 1999). Much experimental work seeks to determine which of these models best predict an individual's behavior under different circumstances, but the work is ongoing.

This paper tests the predictive accuracy of the inequity aversion model in an ultimatum game variant, the structure of which is inspired by a classic paper that exhibits many simple deviations from equilibrium behavior. Data collected from student subjects participating in incentivized treatments is used as the basis of this analysis. Payoffs in these treatments are designed to reveal evidence of inequity averse behavior if some portion of the subjects hold such preferences. This investigation contributes to the growing body of literature focused on examining and modeling the behavior of individuals in different circumstances. Specifically, the data generated serve as evidence about the applicability of the inequity aversion model in a one-shot ultimatum bargaining game. Observed participant behavior is compared with the results predicted by both game theory and the inequity aversion model to examine its consistency with either theoretical model. This paper argues both that the assumption of pure self-interest is not a representative model of an individual's behavior, and that, though some inconsistencies in the data exist, the behavior of a much greater portion of the sample population is consistent with the inequity aversion model of other regarding preferences.

This paper is divided into four sections. Section 1 provides an overview of relevant examples of behavioral deviations from equilibrium and common models of other regarding preferences. Section 2 outlines the details of the ultimatum game

variant structure and inequity aversion model, the connection between the two, and the treatments used to test the model's predictive capabilities. Section 3 presents and discusses the results of the experiment. Section 4 provides a broader summary of the experimental findings, their implications, and possible avenues for future research.

1. Literature Review

There is substantial controversy about what, if anything, connects these observations. The issue goes to the heart of what it is that experimental economics can hope to accomplish. If no connections can be found, we are left with a set of disjoint behavioral charts, each valid on a limited domain. (Bolton and Ockenfels, 2000, 166)

The literature reviewed falls generally into two categories. The first comprises the literature on behavioral deviations. The circumstances in which observed behavior differs from the theoretic self-interested equilibrium, the outcome reached by two rational agents seeking only to maximize their own earnings, are often studied. Much work has been done in an attempt to reveal systematic patterns of such deviations. The second category encompasses models that attempt to account for observed deviations from equilibrium by incorporating preferences other than self-interest into an individual's utility function. The body of literature on these subjects is both large and growing, and the articles presented in this review represent core examples of commonly referenced deviations and models.

1.1. Experimental Evidence of Behavioral Deviations

One of the most commonly referenced, studied, and modeled deviations from equilibrium behavior occurs in the ultimatum bargaining game presented by Guth,

Schmitterberger and Schwarze (1982). The basic ultimatum game has two players, a proposer and a responder, who must split a fixed sum of money. The proposer offers some division of the money to the responder. The responder may either accept the offer as it was made, or reject it. Rejecting an offer causes both players to earn nothing.

Equilibrium for the game occurs when the proposer offers the minimum positive amount, and that the responder always accepts the non-zero offer. This leaves the responder better off than if he¹ rejects the offer and receives nothing, while the proposer maximizes his own earnings. The authors find that such behavior is rarely observed in practice. Instead, most proposers offer more than the minimum positive amount, and offers of a relatively small portion of the sum to be divided are commonly rejected. The conclusion states that these results indicate a thought process for the responder along the lines of, “if [the proposer] left a fair amount to me, I will accept. If not and if I do not sacrifice too much, I will punish him by choosing conflict,” (Guth et al., 1982, 384). The exact meaning of “fair” and “too much” are not discussed or investigated any further in this initial work, but have been the focus of much followup research.

Camerer (2003) presents the stylized facts from many subsequent studies that explore behavioral deviations in ultimatum and dictator bargaining games. There is evidence from studies conducted with different cultures, game structures, and payoff magnitudes, all of which produce results that are generally consistent with the early findings of Guth et al. (1982). Proposers usually offer more than the minimum positive amount predicted in equilibrium, at times as much as half of the sum to be divided.

Similarly deviating from theory, responders reject offers of less than 20% of the sum

¹ The singular male pronoun will be used throughout this paper for consistency when referring to any generic individual.

about half of the time (Camerer, 2003). Non-equilibrium behavior is also observed in most dictator game experiments. The dictator game is similar to the ultimatum game, except the responder has no option to reject the proposer's offer, allowing the proposer to decide the division of the funds unilaterally without concern of rejection. The game theoretic equilibrium is that the proposer keeps the entire sum, maximizing his own earnings. Offers in games of this sort are still larger than predicted by self-interest, though they usually decrease somewhat relative to ultimatum game levels, with the strategic concerns of an offer being rejected removed in this form of the game (Camerer, 2003). The exact motivation for this behavior is still unknown.

Goeree and Holt (2001) discuss ten behavioral deviations from the outcomes predicted by traditional economic theory. In each of the ten cases, a simple experiment with a clear equilibrium is conducted in which observed behavior is generally consistent with self-interest. Then, with only a slight alteration to some component of the experiment that does not change the self-interested outcome, many of the same participants behave in ways entirely inconsistent with self-interest. These examples demonstrate that even a subtle change in the payoffs or structure of a game can affect participant behavior in ways that traditional theory does not anticipate. One example looks at participant behavior in an ultimatum game variant, which is discussed in more detail in Section 2. Investigating the preferences that motivate the behavior in this game is the focus of this paper.

Other studies show how factors that have nothing to do with the explicit payoffs or structure of a game can still influence behavior. The physical attractiveness or amount of personal information known about one's partner in a game can significantly

impact observed behavior. Simply being told a partner's name has been shown to increase offers in the dictator game (Charness and Gneezy, 2003). The level of public visibility of monetary rewards received can also impact individual behavior (Ariely, Bracha and Meiser, 2009). Results from such studies imply that a wide array of factors can influence an individual's behavior.

1.2. Models of Other Regarding Preferences

One of the more commonly referenced models of other regarding preferences is Fehr and Schmidt's (1999) inequity aversion model. Inequity aversion models utility as a function of both a player's own payoff and the payoff of his partner. When there is a difference between the players' payoffs, each individual's utility is decreased to some extent by envy or guilt felt as a result of the disparity. This model has been shown to have strong descriptive power in some contexts, and its simplicity has helped it gain prominence. Because of this model's academic importance and mathematic simplicity, it is the model that this investigation will test in the context of an ultimatum game variant. Its application will be discussed in more detail in Section 2.

Croson (2007) summarizes and tests models of commitment, altruism, and reciprocity as possible explanations for behavioral deviations from equilibrium in a public good game. Generally, commitment theories predict that individuals are self-interested, that they assume others will behave as they do, and that they will act to maximize their own earnings given this belief. Altruism theories incorporate the amount that others earn positively into an individual's utility function. Reciprocity theories claim that individuals attempt to match the amount contributed by the others in their group.

Reciprocity of the median amount contributed by other group members is found to be the most apt predictor of individual behavior of the three models considered in the public good game (Croson, 2007).

Kahneman and Tversky (1979) present prospect theory as an alternate model for explaining the behavior of individuals facing choices with uncertain outcomes. The authors cite many commonly observed behavioral deviations from what is predicted by standard expected utility calculations. Outcomes that are certain or have very low probabilities of occurring are often overweighted, and preferences for risk often reverse depending on whether the same choices are framed as gains or losses. The model's assumptions allow preference reversals like the Allais Paradox (Allais, 1953) and the appeal of both gambling and insurance policies to be explained in one theoretic framework (Kahneman & Tversky, 1979).

Commonly, models of other regarding preferences will incorporate several of these preference theories into a hybrid model. Bolton and Ockenfels (2000) include the ideas of both equitability and reciprocity into a single model of preferences. It is similar in spirit to the Fehr Schmidt (1999) inequity aversion model, in that it takes the difference between monetary payoffs to players into account. It also integrates ideas of reciprocity similar to those described and tested by Croson (2007). The inequity and reciprocity arguments taken together lead to the prediction of an equal division of the sum in ultimatum games, which is consistent with much of the experimental evidence.

Erlei (2008) poses a similar heterogeneous social preferences model that accounts for individuals who have different preferences for fairness and welfare. The model assumes that individuals are either inequity averse, have "welfare preferences,"

or are purely self-interested. Utilities of inequity averse and purely self-interested types are modeled in ways similar to the previously discussed models. Individuals with welfare preferences are said to behave somewhat altruistically, with the payoffs to other players factoring positively into their utility. Erlei also incorporates a negative reciprocity component in the utility function of each type that makes individuals of all types inclined to punish perceived misbehavior of individuals with whom they are paired.

In most cases, these models predict behavior more accurately than pure self-interest. However, none of them are fully accurate at predicting observed behavior in all of the contexts in which they are tested. Such inconsistencies highlight that even complex and seemingly robust models are not yet fully reflective of the preferences that drive individual decision-making. Much work is still required to reach a more complete understanding of what motivates people across a wide range of circumstances.

2. Methods

This project investigates the underlying tension between self-interest and other regarding preferences by testing the predictive power of the Fehr-Schmidt inequity aversion model in the framework from Goeree and Holt's ultimatum game variant. This section provides a more detailed discussion of the structure of the ultimatum game variant, the inequity aversion model, and how inequity aversion can be used to analyze participant behavior in the game structure.

2.1. *Experimental Background*

Goeree and Holt (2001) present a two stage sequential move game, in which the proposer may choose between a safe option (S) or a risky option (R). The safe option ends the game and prevents the responder from making a move. If the proposer chooses the risky option, the responder can pick the Nash equilibrium outcome (N), or a “punish” option (P), in which both players get a lower payoff than in outcome N.

In each game tree, the proposer can choose between moves of S (safe) or R (risky), and the responder can choose between moves of P (punish) or N (Nash equilibrium). The choices made by the two players allow them to jointly reach outcome S, P, or N. The original game structures and percentage of subject pairs that ended in each outcome is displayed in Figure 1.

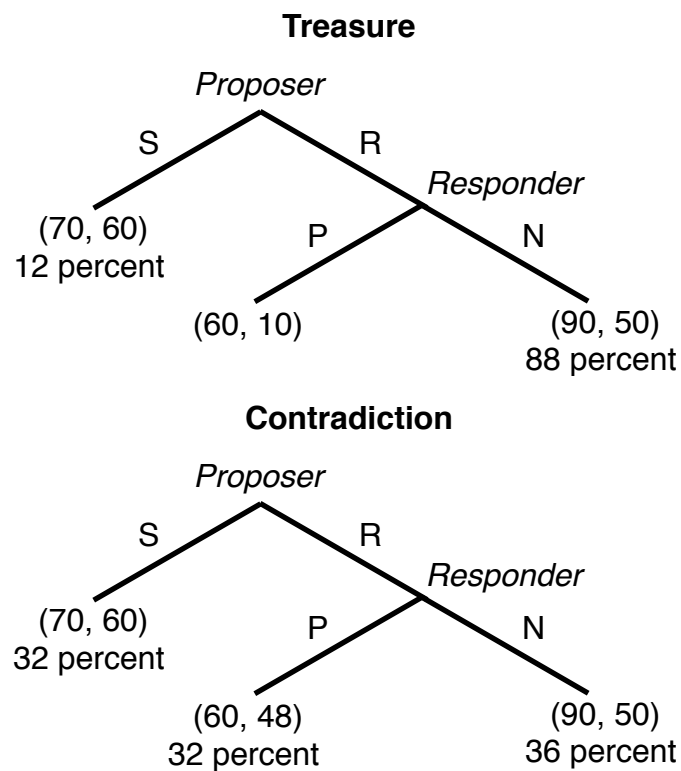


Figure 1. Goeree & Holt’s original ultimatum game variant and results

The payoffs listed for each of the three possible outcomes are for the proposer and responder, respectively. For example, if in the Treasure treatment the proposer picks R and the responder picks P, then the game ends in outcome P, with the proposer earning 60 and the responder earning 10.

Rollback equilibrium logic employed in game theory reveals that two purely self-interested players will always end in outcome N. Given the chance, a rational responder will always choose N over P to maximize his own earnings. Knowing a responder will always pick N in the risky subgame, a rational proposer will always pick R, as he earns more in outcome N than in outcome S.

In the Treasure treatment, when the responder must sacrifice a large amount to choose P instead of N, observed behavior is largely consistent with pure self-interest. The original results show that 88% of trials end in equilibrium, and 100% of responders who are given a chance to move pick the equilibrium outcome N. However, this is not the case in the Contradiction treatment. When the responder must give up only a small amount, 32% of all trials end in outcome P, with 47% of responders who are allowed to move choosing P over N. Safe moves by proposers also increase dramatically, with the number of trials ending in outcome S rising from 12% in the Treasure treatment to 32% in the Contradiction (Goeree & Holt, 2001).

The Fehr-Schmidt inequity aversion model (1999) can be applied in the context of this ultimatum game variant. Inequity averse utility functions for the case of a two player game are displayed in Figure 2.

Figure 2. Separated envy and guilt inequity aversion equations

Envy: $U_i(x) = x_i - \alpha_i (x_j - x_i)$	Guilt: $U_i(x) = x_i - \beta_i (x_i - x_j)$
where $i \neq j$, $0 \leq \alpha$, $0 \leq \beta \leq 1$	

The utility of a player i is calculated by taking his monetary payoff, x_i , and subtracting some amount from it. If the other player has earned more, player i 's is assumed to experience some amount of guilt, and his payoff is reduced by the disparity between the players' payoffs multiplied by a scalar α_i , which represents player i 's marginal propensity for envy. If player i earns more than the other player, he experiences guilt. In this case, his payoff is instead decreased by that difference multiplied by a scalar β_i , player i 's marginal propensity for guilt. Larger values of α and β imply a stronger aversion to either type of inequity (Fehr & Schmidt, 1999).

In Goeree and Holt's original work, the increased percentage of trials that end in non-equilibrium outcomes in the Contradiction treatment could be explained by inequity averse preferences. Inequity aversion accounts for the observed increase in P outcomes by assuming a responder's utility is influenced negatively by envy over the disparity between his earnings and the proposer's. Depending on how envious a particular responder is, he may be willing to pick P, sacrificing some monetary gain in order to reduce both the disparity between his own earnings and his partner's, and consequently the negative impact of envy on his utility. However, the original data from these two treatments alone is insufficient to rigorously test the applicability of the inequity aversion model in this context.

This paper investigates more closely the potential of the inequity aversion model to explain behavior in this game structure. A series of experimental treatments using student participants test for evidence of inequity averse behavior to see if guilt and envy are driving non-equilibrium choices of S and P by proposers and responders. Specifically, participant behavior will be observed as the magnitude of the difference between payoffs to the proposer and responder is carefully manipulated. Payoffs in each set of treatments are designed to reveal patterns of guilt or envy, should such factors be motivating individual behavior. The resulting data is then analyzed both in aggregate and on an individual behavioral level. Results are compared with the predictions of the Fehr-Schmidt inequity aversion model and pure self-interest to assess the predictive validity of both preference models in the context of this one shot ultimatum bargaining game variant.

2.2. Experimental Design

This research examines payoffs from Goeree and Holt's ultimatum game variant using the Fehr-Schmidt inequity aversion model. Monetary payoffs are then manipulated carefully to create nine treatments that test for indications of inequity averse behavior.

Taking the equations for envy and guilt in Figure 2 and replacing i and j with P and R for proposer and responder where appropriate, the game trees from Goeree and Holt can be rewritten with all payoffs expressed in terms of inequity averse utility, as seen in Figure 3.

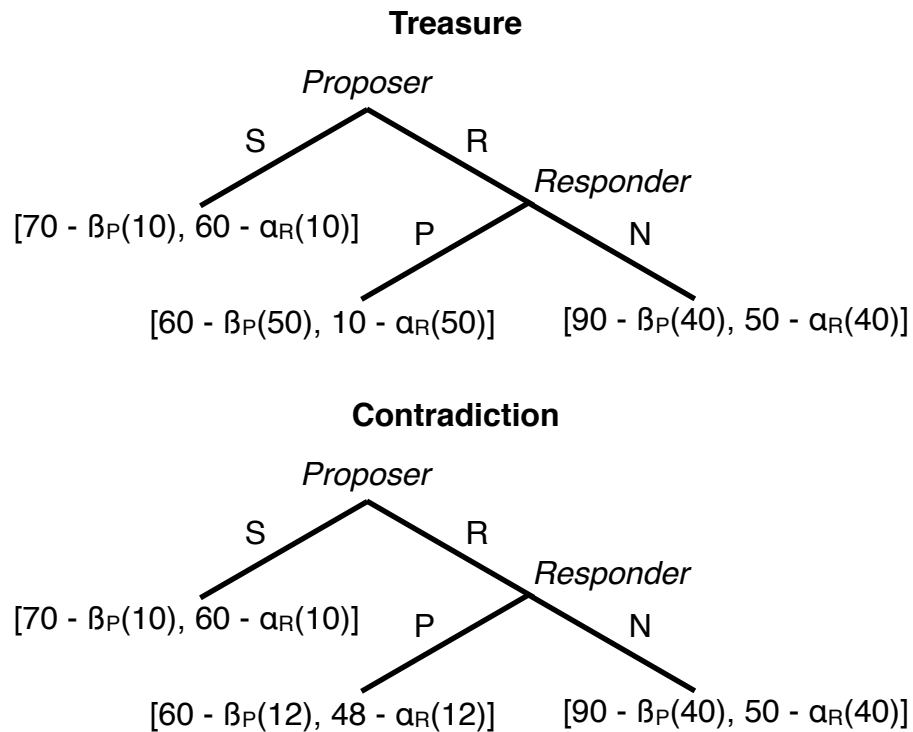


Figure 3. Goeree & Holt’s original ultimatum game structure with inequity averse payoffs

Examining these two original treatments in light of inequity aversion’s impact on utility allows some forms of non-equilibrium play to be explained. For example, in the Treasure treatment, $\alpha_R \leq -4$ would induce a responder to choose P over N to maximize his own utility. However, inequity aversion bounds $\alpha \geq 0$, with the justification that a negative value of α implies a preference both for lower earnings and greater inequity, which does not make intuitive sense. That no responders picked P in the original Treasure treatment is then consistent with behavior predicted by inequity aversion. An inequity averse proposer may find it preferable to choose S instead of R. If a proposer has $\beta_P \geq 2/3$, he would maximize his own utility by choosing S, reducing the guilt he would experience from the relatively large disparity between payoffs to each player in outcome N by enough to offset the monetary loss relative to his earnings in S. Assuming

β values for individuals range between $0 \leq \beta \leq 1$, $\beta_P \geq 2/3$ encompasses somewhat high values and would be uncommon. This is consistent with Goeree and Holt's original results, in which only 12% of proposers chose S.

The Goeree and Holt game structure, conceptualized in the context of the inequity aversion model, is the foundation for the design of the treatments used in this study. The minimum values of α and β that individuals must have to invoke non-equilibrium behavior are referred to from here on as "critical" values of either parameter. These critical values can be calculated for any set of monetary payoffs in this game structure. The treatments vary the disparity of monetary payoffs between players such that a range of critical levels of α and β are required to induce non-equilibrium play. An individual with an α or β greater than these critical values finds his non-equilibrium choice preferable to what pure self-interest predicts. If the inequity aversion model is an accurate predictor of behavior in this setting, one would expect that when lower critical values of α and β are required to alter the utility maximizing outcome away from traditional equilibrium play, such outcomes will occur more frequently.

The general intent behind the payoff adjustments used in each treatment is to increase the disparity between monetary payoffs to players such that varying levels of α and β are required to induce non-equilibrium play, assuming individuals are inequity averse. The frequency with which each outcome is reached is recorded in each treatment, and shifts in the outcome distribution can be analyzed for consistency with predictions of the inequity aversion model as well as traditional self-interest. The behavior of individuals across rounds is also examined to determine what portion of subject behavior is explained by both models.

2.3. Treatments and Expected Outcomes

The nine experimental treatments used in this study fall into three categories— benchmark, envy, and guilt treatments. The outcomes in each set of treatments that one would expect to observe if inequity aversion is a good predictor of behavior are also discussed for each treatment group.

Benchmark

Table 1. Payoffs and critical α for each Benchmark treatment

Benchmark Treatments	P if $\alpha \geq$	S		P		N	
		Prop.	Resp.	Prop.	Resp.	Prop.	Resp.
Treasure	-4	70	60	60	10	90	50
Contradiction	0.07	70	60	60	48	90	50
Scaled	0.07	210	180	180	144	270	150

The benchmark treatments allow this study’s results to be compared with those found by Goeree and Holt (2001) to see if this subject pool behaves similarly to the original group of participants. The payoffs, found in Table 1, from the original Treasure and Contradiction treatments are recreated exactly. The expectation is that participants should match the behavior observed in the original paper if this study’s subjects have preferences similar to those seen in the original study. A third “scaled” version of the Contradiction treatment is also included in the experiment. To manipulate payoffs in other treatments with fine enough granularity, it is convenient to roughly triple all payoffs from the original games. The inequity aversion model predicts no change in behavior as a result of scaling payoffs. To ensure that scaling on its own does not impact behavior,

an exactly scaled version of the Contradiction treatment is included. Inequity aversion predicts that behavior should not change between the Contradiction and Scaled treatments. The critical α level above which an inequity averse responder finds P preferable to N is also listed for each treatment.

Envy

Table 2. Payoffs and critical α for each Envy treatment

Envy Treatments	P if $\alpha \geq$	S		P		N	
		Prop.	Resp.	Prop.	Resp.	Prop.	Resp.
High α	1.25	210	180	180	100	270	150
Mid α	0.50	210	180	180	120	270	150
Low α	0.20	210	180	180	135	270	150

Three treatments are used to gather evidence on how responders react to changing costs for punishing their opponents relative to their equilibrium payoff. Payoffs for each are shown in Table 2. In the High α treatment, only responders with a marginal propensity for envy that is greater than the critical α value would pick P over N if their behavior is motivated by inequity aversion. As the critical α value is lowered from the High α to the Low α treatment, one would expect a greater portion of the population to find their α values higher than the critical α in a round. One would then expect a corresponding rise in the occurrence of P choices by responders as a percentage of total outcomes.

Data from these treatments can also be analyzed at the individual level. According to inequity aversion, as the critical α value decreases, individuals will at some point switch from picking N to picking P if they find their own α greater than the critical α

for the treatment. They should not, however, pick P in one treatment and pick N in a treatment with a lower critical α value. Such a reversal of behavior would be inconsistent with the inequity aversion model.

Guilt

Table 3. Payoffs and critical α for each Guilt treatment

Guilt Treatments	S if $\beta \geq$	S		P		N	
		Prop.	Resp.	Prop.	Resp.	Prop.	Resp.
High β	0.75	180	180	180	30	270	150
Mid β	0.50	240	180	180	30	270	150
Low β	0.25	260	180	180	30	270	150

The guilt treatments presented in Table 3 are designed with similar logic to the envy treatments, but with a focus on observing proposer rather than responder behavior. A scaled version of the Treasure treatment is used instead of the Contradiction treatment. In the Contradiction treatment, it can be argued that proposers may be fully self-interested and still pick S if they fear the responder will pick P. Inequity aversion predicts that no responder would choose P over N in the Treasure treatment, as doing so would require a negative value of α , which is outside the domain of $\alpha \geq 0$ defined in the model. It can then be claimed that, if inequity aversion holds, a proposer would not need to factor fear of a responder picking P into his own decision of picking S or R. The choices facing the responder then are effectively between the S and N outcomes. The only reason a proposer would choose S over R is if he, for reasons of guilt, preferred outcome S over outcome N.

In the High β treatment, only individuals with a high marginal propensity for guilt would find it preferable to pick S over the payoff they would receive in N. The Mid β and

Low β treatments have S payoffs that are decreasingly costly for the proposer and have lower critical β values. One would expect that observed S decisions by proposers will increase in frequency between the High, Mid, and Low β treatments. Also, similar to the envy treatments, one expects consistency in an individual's choice to switch from picking R to S as the critical β value decreases. It would be inconsistent with inequity aversion if an individual picked R in one round if he had previously found picking S preferable in a round with a higher critical β value.

2.4. Experimental Logistics and Considerations

The intent behind this research is to see how individuals behave in the context of an incentivized, anonymous, one shot bargaining game. To accomplish this, many considerations were made beyond the payoffs in each treatment. Subjects are randomly rematched with a new partner after each round, and are made aware of the rematching at the beginning of the experiment. This prevents the session from effectively becoming a multi-stage game between two players and discourages participants from altering their behavior in one round in hopes of influencing their current partner's actions in following rounds.

Subject roles are randomly assigned at the outset of each session and are then fixed for all nine treatments within that session. Keeping player types constant allows data on the behavior of individuals to be analyzed. The behavior of a single subject in the proposer or responder role can be examined to see if his behavior in a group of treatments is consistent with either self-interest or inequity aversion.

The sample of outcomes in each treatment must be large in order to make any sort of data analysis statistically significant. The study uses data from sixty subjects, spread over four sessions of roughly equally-sized groups, generating thirty total outcomes for each treatment. Subjects are undergraduate students from the University of Puget Sound, recruited largely from courses in the Business and Economics departments. This study has approval from the University of Puget Sound Institutional Review Board.

Subjects in the experiment are paid for their participation based on the number of points they accumulate over the nine treatments. A .8% point to dollar conversion ratio is used to determine final payoffs, plus a \$3 “show-up” fee for each participant, rounded to the nearest dollar. Funding for this research is provided jointly by the University of Puget Sound’s University Enrichment Committee grant and the Economics Department. The commonly referenced induced value theorem posits that paying subjects helps make the results both monotonic and salient (Friedman & Sunder, 1994). In the case of this research, this means that subjects, other things equal, prefer more money to less, and that subjects understand that choosing the “safe” or “punish” outcome will reduce their monetary reward. Making such a choice then implies a willingness to incur a real cost due to some other preference.

The experiment is programmed and conducted with z-Tree software (Fischbacher, 2007). The software lets sessions be run in a fully automated manner, allowing subjects to interact anonymously and be randomly rematched quickly. This software also ensures consistency in data collection, with all data exported automatically into a spreadsheet for analysis.

The mechanics of the game and logistics of the experiment, including random rematching and consistency of player roles across treatments, are explained to all subjects at the outset of the experiment. All participants are given full information about the payoffs to both players at the beginning of each treatment.

3. Results & Discussion

Data from each group of treatments is analyzed below both in aggregate for patterns of individual behavior across the treatment group. Chi-squared tests are used to determine if differences between aggregate outcome distributions are significant². Self-interest predicts that all trials end in outcome N. Behavior in all treatments is significantly** different from this prediction.

3.1. Benchmark

Table 4. Benchmark outcome distribution by treatment

Benchmark Outcomes	All			Proposers		Responders	
	S	P	N	S	R	P	N
Treasure	30%	13%	57%	30%	70%	19%	81%
Contradiction	20%	40%	40%	20%	80%	50%	50%
Scaled	17%	37%	47%	17%	83%	44%	56%

The percentage of trials that ended in each outcome for the benchmark treatments are displayed in Table 4. Participant behavior in both the Treasure and Contradiction treatments is significantly* different from Goeree and Holt’s original findings from Figure 1, with more S and P outcomes in both cases. This suggests the

² Significance levels:

* = $p < .05$

** = $p < .01$

population in this study is somewhat less prone towards self-interested behavior than the subjects in Goeree and Holt's original study.

A significant** difference exists between behavior in the Treasure and Contradiction treatments. As inequity aversion predicts, the number of responders picking P increases dramatically in the Contradiction treatment when the cost of punishment and payoff disparity is decreased.

There is no significant difference in behavior in the Contradiction and Scaled treatments, a result that is consistent with predictions of the inequity aversion model. As long as all payoffs are multiplied by the same scalar, the critical α and β values required to induce non-equilibrium behavior are unchanged.

Though the observed behavior in this study is different from that reported in Goeree and Holt's original work, the Benchmark treatments nonetheless provide encouraging evidence in favor of inequity aversion's predictive power. The significance of the increase in punishment rates from the Treasure to Contradiction treatment and the consistency in behavior between the Contradiction and Scaled treatments both align with behavior predicted by the model.

3.2. Envy

Table 5. Envy outcome distribution by treatment

Envy Treatments	All			Proposers		Responders	
	S	P	N	S	R	P	N
High α	30%	20%	50%	30%	70%	29%	71%
Mid α	13%	33%	53%	13%	87%	38%	62%
Low α	17%	23%	60%	17%	83%	28%	72%

Table 6. Responder behavior across envy treatments

Responder Behavior	Consistent			Inconsistent
	Always P	Never P	IA	
# of Subjects	4	13	5	6
% of Subjects	14%	46%	18%	21%

The percentage of trials that ended in each outcome for the envy treatments are displayed in Table 5. In aggregate, the trends in envy outcomes are not fully consistent with the predictions of the inequity aversion model. The results of the High α and Mid α treatment are significantly** different from one another, with P outcomes increasing in the Mid α treatment as inequity aversion suggests. However, P outcomes fall significantly* from the Mid α to Low α treatment, behavior that is opposite the trend that would be expected if participants are inequity averse.

Though the trends in aggregate outcome distributions provide mixed evidence about inequity aversion’s predictive power, the observed individual behavioral patterns displayed in Table 6 are more consistent with the model. The behavioral patterns of responders who moved in at least two rounds are grouped into four categories: Always P, Never P, Inequity Averse (IA), and Inconsistent. The Always P behavioral pattern encompasses players with very high α values, who may find that all of the critical α values in the treatments are lower than their own. Such individuals would find it preferable to always pick P in these treatments. Self-interested individuals or those with

extremely small α values may never find P preferable to N. Those who never choose P are counted in the Never P category. The IA category counts participants who switch from picking N to picking P in one treatment, and did so consistently in all rounds with lower critical α values. The Inconsistent category captures the remainder of the population— those who exhibit a preference reversal by choosing P in one round and choosing N in a round with a lower critical α value.

Always P, Never P, and IA are all behavioral patterns consistent with the inequity aversion model. Seventy-nine percent of responders behaved in a way consistent with one of these behavioral patterns. On the contrary, Never P is the only behavioral pattern consistent with pure self-interest. Only 46% of responders employed a Never P strategy. These results suggest that inequity aversion's predictive power for responder behavior is more accurate than that of self-interest, and a proportions test reveals that the difference is statistically significant at a 95% confidence level.

3.3. Guilt

Table 7. Guilt outcome distribution by treatment

Guilt Outcomes	All			Proposers		Responders	
	S	P	N	S	R	P	N
High β	17%	7%	77%	17%	83%	8%	92%
Mid β	27%	10%	63%	27%	73%	14%	86%
Low β	43%	13%	43%	43%	57%	24%	76%

Table 8. Proposer behavior across guilt treatments

Proposer Behavior	Consistent			Inconsistent
	Always S	Never S	IA	
# of Subjects	2	11	10	7
% of Subjects	7%	37%	33%	23%

The guilt results displayed in Table 7 fit well with the behavior predicted by the inequity aversion model. The differences in outcome distributions between the High β , Mid β and Low β treatments are significant*. The percentage of proposers picking S in each treatment increased as the critical β value at which one would find S preferable to N decreased. This trend in S moves by proposers is consistent with what one would expect from a population with inequity averse preferences.

Patterns of individual proposer behavior in the guilt treatments shown in Table 8 are grouped and analyzed similarly to that of responders in the envy treatments. Always S, Never S, and IA are all behavioral patterns consistent with the inequity aversion model. Seventy-seven percent of proposers behaved in one of these three ways. Thirty-seven percent of proposers' behavior followed the Never S pattern, which is the only pattern consistent with pure self-interest. These results imply that the predictive power of the inequity aversion model is higher than that of pure self-interest. The difference is statistically significant at a 99% confidence level.

4. Conclusion

4.1. Evidence on Pure Self-Interest and Inequity Aversion

The results of this study are highly inconsistent with behavior that one would expect from purely self-interested individuals and provide encouraging evidence about the predictive ability of the inequity aversion model. As mentioned in Section 3, the outcome distribution in every treatment is significantly different from the distribution predicted by pure self-interest. Game theory predicts that 100% of trials in every treatment would end in the Nash equilibrium outcome N, but the actual data show that this is inaccurate. While the evidence against self-interest as a predictive model in this context is clear, the evidence in favor of inequity aversion's accuracy is mixed.

Though behavior in the benchmark treatments varied from Goeree and Holt's original results, behavioral changes between the benchmark treatments was signed and significant in ways consistent with the inequity aversion model. The number of P outcomes was higher in the Contradiction treatment than the Treasure treatment, and there was no significant difference between the Contradiction and Scaled treatments. Both of these observations are consistent with and provide evidence in favor of the inequity aversion model.

The envy treatment results provide mixed evidence of inequity averse behavior. Though the number of P outcomes rose from the High α to the Mid α treatment, they fell from the Mid α to Low α treatment when inequity aversion predicts P outcomes would occur with greater frequency. Despite this mixed evidence at the aggregate level, inequity aversion does account for 79% of behavioral patterns exhibited by responders

in these three treatments, which is a significantly higher percentage than can be accounted for by pure self-interest.

The guilt treatment results are the most encouraging evidence in favor of the inequity aversion produced by this research. The number of S outcomes rose significantly as the β value necessary to make S preferable to N decreased. These results affirm the predictions of the inequity aversion model. Proposer behavior at the individual level across all three treatments is also largely consistent with inequity aversion. The 77% of individual behavior explained by inequity aversion is significantly better than the 37% explained by self-interest. The inequity averse trends at both the aggregate and individual level imply that the inequity aversion model predicts proposer behavior in this ultimatum variant well, and significantly better than pure self-interest.

Taken together, evidence from this experiment largely favors the inequity aversion model. The trends in observed behavior in the benchmark and guilt treatments at an aggregate and individual level, and envy treatments at an individual level are all modeled well by Fehr and Schmidt's model of fairness preferences. Though the results from this study are not fully consistent with the inequity aversion model, it seems to be a reasonably accurate predictor of behavior, and does a significantly better job than the notion of pure self-interest assumed by classic economic theory.

4.2. A Practical Application

The applicability of these and other experimental findings is not limited to the economics discipline. Once the factors that motivate actions are more fully understood, an accurate model of an individual's preferences can be used by economists,

businesspeople, non-profit leaders and policymakers alike to guide decision making and analysis.

For example, the evidence, provided by this study, that guilt may factor into an individual's preferences could be used by a charity that supports the poor or homeless to increase fundraising efforts. Instead of soliciting funding by advertising the charity's name or prosocial cause, such an organization could leverage the effects of guilt by framing requests for donations as a dictator game of sorts between the potential donor and a single recipient of aid from the charity. If information were given about how much money the recipient lives off of per week, and a question were posed about the equivalent figure for the donor, this framing may be enough to induce a realization of the disparity between incomes and the resulting feelings of guilt. This sort of literal "guilt trip" could increase donations as a result of donors being made cognizant of the income disparity and attempting to reduce both it and their own guilt by donating more than they otherwise would have.

4.3. Suggested Research Extensions

Several steps can be taken to expand upon the results of this study. The envy component of inequity aversion is clearly in need of further investigation. More treatments with smaller increments for α levels that would make a player prefer P to N played by a larger pool of subjects could yield more thorough data than the three treatments used in this study. The results from this study may also be explained by other models of other regarding preferences, particularly those that incorporate negative reciprocity into an individual's utility function. The data from this study's treatments or

from additional treatments specifically designed to test for negative reciprocity may reveal a more accurate explanation for what motivates P moves by responders in this game.

Works Cited

- Allais, M (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. *Econometrica*, 21(4), 503-546.
- Bolton, G. & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *The American Economic Review*, 90(1), 166-193.
- Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. (pp. 1-113). Princeton, NJ: Princeton University Press.
- Charness, G. and Gneezy, U. (2003). What's in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization*, 68, 29-35.
- Croson, R. (2007). Theories of commitment, altruism, and reciprocity: Evidence from linear public goods games. *Economic Inquiry*, 45(2), 199-216.
- Erlei, M (2008). Heterogeneous social preferences. *Journal of Economic Behavior & Organization*, 65, 436-457.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 817-868.
- Fischbacher, U. (2007): z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics* 10(2), 171-178.
- Friedman, D., & Sunder, S. (1994). *Experimental economics: A primer for economists*. (pp. 1-207). New York, NY: Cambridge University Press.
- Goeree, J., & Holt, C. (2001). Ten little treasures of game theory and ten intuitive contradictions. *The American Economic Review*, 91(5), 1402-1422.
- Guth, W., Schmitterberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organizations*, III(1982), 367-388.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.