THREE ESSAYS ON MICROECONOMIC THEORY AND EXPERIMENT

by

Huan Xie

B.A. in Economics, Fudan University, China, 2002

Submitted to the Graduate Faculty of

Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2008

UNIVERSITY OF PITTSBURGH

ARTS AND SCIENCES

This dissertation was presented

by

Huan Xie

It was defended on

May 1st 2008

and approved by

Andreas Blume, Professor, Department of Economics, University of Pittsburgh

Lise Vesterlund, Andrew Mellon Associate Professor, Department of Economics, University of Pittsburgh

John Duffy, Professor, Department of Economics, University of Pittsburgh

Roberto Weber, Associate Professor, Department of Social and Decision Sciences, Carnegie Mellon University

Dissertation Advisors: Andreas Blume, Professor, Department of Economics, University of Pittsburgh,

Lise Vesterlund, Andrew Mellon Associate Professor, Department of Economics, University of Pittsburgh

THREE ESSAYS ON MICROECONOMIC THEORY AND EXPERIMENT

Huan Xie, PhD

University of Pittsburgh, 2008

My dissertation uses game theoretical and experimental approaches to study how individual's behavior in different informational environments affects economic outcomes and motivations for charitable giving. Chapter 2 "Bargaining with Uncertain Value Distributions" studies a bargaining model in which the seller is uncertain not only about the buyer's value but also about which distribution the buyer's values are drawn from. Different from the classical models, the distribution of the buyer's values is fixed across periods, while the buyer's values are drawn independently from the distribution each period. I find that adding this additional layer of uncertainty improves the seller's profit when her ex ante beliefs are sufficiently optimistic. Chapter 3, "Social Norms, Information, and Trust among Strangers: An Experimental Study" (with John Duffy and Yong-Ju Lee), investigates whether norms of trust and reciprocity arise in response to different reputational mechanisms. We conduct an experiment where anonymous subjects are randomly matched each period and play a series of indefinitely repeated trust games. We find that the social norm of trust and reciprocity is difficult to sustain without reputational information, although it is supported as an equilibrium by the parameters. The provision of information on players' past decisions significantly increases trust and reciprocity. Furthermore, making such information available at a small cost also leads to a significant improvement, despite that most subjects do not choose to purchase this information. Finally, Chapter 4 "Motives for Charitable Giving" (with Lise Vesterlund and Mark Wilhelm) reports an experiment which tests the pure altruistic and the impure altruistic explanations for charitable giving. We focus on the comparative static predictions of both models and quantify the relative weight attached to the warm-glow component of giving in the impure altruism model. A methodological innovation is to create the equivalent of a series of real-world charities. Each participant is paired with a child who has suffered a severe fire and informed that they single-handedly determine the size of a gift given to the child. Our results show that participants behave as predicted by the impure altruism model. However the relative weight attached to the warm-glow of giving is very small.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0 INTRODUCTION

Microeconomics studies individual's behavior under different economic circumstances. Using both game theoretical and experimental approaches, this dissertation in particular contribute to two topics: individuals' strategic behavior with response to the information environment and motives for altruistic behavior such as trust and charitable giving.

The first chapter, "Bargaining under Uncertain Value Distributions," studies a bargaining procedure where a durable-good monopolist offers a price to rent a durable good to a buyer repeatedly in multiple periods. The buyer is assumed to have private information. Since bargaining proceeds between the two parties repeatedly, the players have the consideration that information revealed in early periods may affect the bargaining outcome in later periods and take this into account from the beginning of the game. I ask how the buyer's private information is revealed or concealed over time and how it affects the distribution of economic surpluses between the two parties.

In the classical models, it is usually assumed that the distribution of the buyer's value is common knowledge, only the buyer's value is private information, and the buyer's value is fixed across periods. In this framework, a stark and established result is that the seller has a large disadvantage and loses most of her monopoly power. The literature on Coase conjecture finds that if the durable-good monopolist sells over time and can quickly lower prices, the seller can hardly achieve profits greater than the lowest buyer valuation and the buyer obtains the entire surplus from trade in excess of his lowest valuation (Coase 1972, Fudenberg et al. 1985). When the monopolist rents the durable, Hart and Tirole (1988) show that the seller always offers a low price until the end of the game if the horizon is long enough.

The purpose of this paper is to examine whether the seller can improve her profit in a rental model when there is an additional layer of uncertainty about the distribution of the buyer's value. It is assumed that the buyer has private information not only about his value in each period but also about the distribution which his values are drawn from. The distribution of the buyer's values is fixed across periods, while the buyer's values are assumed to be drawn from the distribution independently in each period. This information structure characterizes a situation where the buyer has one-sided private information about his long-term preference but also has some uncertainty about the short-term values in the future.

The main result of the paper is that the seller is indeed better off when she has sufficiently optimistic ex ante beliefs of the favorable distribution, compared to the model of Hart and Tirole (1988). Intuitively, given that the seller's prior belief is large and there is uncertainty on the buyer's value distribution, the seller's posterior belief does not change critically and she offers a same price in the second no matter whether the first offer is accepted or rejected by the buyer. Thus the buyer with a high value in the first period is more willing to accept a high price, and this gives the seller a leeway to improve her profit.

In the second chapter "Social Norms, Information, and Trust among Strangers: An Experimental Study," we examine how the social norms of trust and reciprocity emerge among a population of anonymous strangers who do not meet each other frequently. For instance, why the online market with essentially anonymous buyers and sellers functions efficiently? Why the credit card companies lend money to customers without credit history? Why restaurants in vacation areas may also provide good services to tourists who never come back?

We first examine the hypothesis that trust might be attached to a society as a whole; the fear of the destruction of that trust might suffice to enforce trustworthy behavior by all members of the society as shown by Kandori (1992). On the other hand, such a mechanism might be too fragile and so we examine the second possibility that trustworthiness resides at the individual rather than the societal level. In particular, we ask whether the provision of information on individual reputations for trustworthiness engenders greater trust than in the case where such information is absent. We further explore whether the free provision of reputational information is responsible for our findings or whether the availability of acquiring such information (at a small cost) suffices to sustain greater trust and reciprocity.

We find that, although the social norm of trust and reciprocity is theoretically sustained as a sequential equilibrium without any reputational information, there is very little trust and reciprocity found in this case in the experiment. Providing information on the trustee's previous behavior improves matters. When the amount of information about trustees is increased, it leads to more significant increases in trust and reciprocity relative to the absence of such information. Finally, when investors must decide whether to purchase the information concerning their current matched trustee, we find that on average, only one-fourth of investors choose to purchase this information, so that the other three-fourths are in the dark about the prior behavior of their current trustee. Nevertheless, trust and reciprocity is significantly higher in this costly information treatment as compared with the baseline no-information treatment.

In the third chapter "Motives for Charitable Giving," we present an experimental study on why people contribute to charities. Economic theory on charitable giving distinguishes between the pure and impure altruism model. The pure altruism model argues that the sole motive for giving is a concern for securing the charity's output, whereas the impure altruism model allows for the possibility that a donor also get a warm glow from being the one who secures the contribution to the charity. The literature on motives for giving has traditionally taken the pure altruism model as the null hypothesis, and interpreted rejections as evidence in favor of impure altruism.

In contrast to previous studies we test the comparative static predictions of both the pure and impure altruism model, and account for the possibility that consistent with the impure altruism model the support for pure altruism may be sensitive to the point at which the motive for giving is evaluated. In particular in the impure altruism model increases in the contribution of others will shift the marginal motive for giving from a concern for altruism to a concern for the private benefit from giving. From a methodological viewpoint we develop an environment that closely mirrors those of the theoretical models. We examine contributions to an actual charity where each participant is informed of an initial contribution amount and singlehandedly determines the final dollar amount to be transferred to a recipient of the charity.

In examining charitable contributions across several budgets, we consistently find behavior in line with the comparative static predictions of the impure altruism model. However our data also make clear that a substantial weight is attached to the altruistic component of preferences. When estimating a representative utility function we find, consistent with our comparative static results, that participants get a private benefit from giving, however it accounts for but a fraction of the weight attached to the public benefit associated with providing funds for the charity's recipient. Thus we demonstrate that there are environments for which it would be incorrect to assume that donor's charitable contributions primarily are made because of the private benefit one may experience from giving.

2.0 BARGAINING WITH UNCERTAIN VALUE DISTRIBUTIONS

2.1 INTRODUCTION

In this paper we analyze a two-period bargaining model in which a durable-good monopolist rents the durable good to a buyer in each period.[1] The buyer has private information about his values realized in the current and previous periods and the distribution which his values are drawn from, but is uncertain about his future values. The distribution of the buyer's values is assumed fixed across periods, while the buyer's values are assumed to be drawn from the distribution independently in each period. As early revelation of the buyer's private information affects future outcomes when the players interact repeatedly, we focus on the question of how information is revealed over time. Related to that, we examine how economic surplus is distributed between the buyer and the seller. In the classical models, it is usually assumed that the distribution of the buyer's value is common knowledge, only the buyer's value is private information, and the buyer's value is fixed across periods. In this framework, a stark and established result is that the seller has a large disadvantage and loses most of her monopoly power, since the high-value buyer has a large incentive to conceal his value given that his value is fixed and bargaining proceeds in multiple periods. This paper asks whether the seller's standing may be improved by introducing a layer of uncertainty about the distribution of the buyer's values and allowing the buyer's values to randomly change each period.

The literature on Coase conjecture finds that if the durable-good monopolist sells over time and can quickly lower prices, the seller can hardly achieve profits greater than the lowest buyer valuation and the buyer obtains the entire surplus from trade in excess of his

[1] We will use she to denote the seller and he to denote the buyer.

lowest valuation (Coase 1972, Fudenberg et al. 1985).[2] The intuition is that the monopolist is induced to reduce the price when facing the residual demand after having sold some quantity to high-value buyers, and rationally anticipating falling prices causes most potential buyers to wait for a lower price in the future.

The Coase conjecture, however, may fail to hold if we relax the assumption that the buyer's valuation is fixed across periods.[3] Sobel (1991) shows that when there is a constant flow of new buyers, a Folk theorem holds, that is, any positive average profit less than the maximum feasible level can be attained. Blume (1990) examines a bargaining model where the low buyer type's value varies over time and the high buyer type's value stays fixed, and demonstrates that both uniqueness and Coase conjecture may fail to hold when valuations are allowed to vary randomly.

Another approach that tries to ameliorate the seller's position in the Coase conjecture is to allow the seller to rent the durable good instead of selling it. Bulow (1982) argues that the durable-good monopolist may be better off if she chooses to rent the durable good rather than sell it, however, he assumes that the buyer is anonymous, that is, the seller cannot identify the buyer nor his past behavior.

When the monopolist bargains over renting the durable good to a *non*-anonymous buyer with private value, Hart and Tirole (1988) show that the seller always offers a low price until the end of the game given any prior belief, if the horizon is long enough. Intuitively, when the time horizon is long, the high-value buyer will not accept any price rejected by the low-value buyer, in order to avoid being charged with a high price in all the later periods. So the seller is not able to price discriminate and she charges a low price to both the low-value and high-value buyer, until close to the end of the horizon. Therefore, if the durable-good monopolist rents the durable good to a *non*-anonymous buyer, the seller is again caught in an unfavorable position. Furthermore, Hart and Tirole (1988) show that renting is even no better than selling, and it is strictly worse if the horizon is long enough.

The main purpose of this paper is to examine whether the seller can improve her profit

[2] This result only holds under the assumption that the seller's marginal cost is lower than the buyer's lowest value, which is called the "gap" case in the literature.

[3] Failures of the Coase conjecture are also found when the lowest buyer valuation does not exceed the seller's cost, which is referred as the "no-gap" case in the literature (Gul et al. 1986, Ausubel and Deneckere 1989). Ausubel and Deneckere (1989) prove a Folk theorem similar to Sobel (1991).

in a rental model with a *non*-anonymous buyer if the buyer's value is allowed to change each period. It is not clear how much Hart and Tirole's (1988) result depends on the assumption that the buyer's value is fixed across periods. Intuitively, when the buyer's value is invariant over time, the high-value buyer reveals all the information about his future values once he takes an action different from the low-value buyer. On the contrary, if there always remains some degree of uncertainty about the buyer's future values, a buyer whose current value is revealed to be high does not necessarily have a high value in later periods. Therefore, he may be more willing to reveal his current value and realize a positive payoff in early periods without losing all potential future surplus. Thus, the uncertainty of the buyer's future values may provide an additional leeway to solve the problem of the durable-good monopolist.

In the paper we introduce an additional layer of uncertainty about the distribution of the buyer's value. The distribution of the buyer's value may be either favorable or unfavorable, with the favorable distribution generating a high value with a higher probability. At the beginning of the game, the buyer privately observes the distribution. At the beginning of each period, the buyer's value is randomly drawn from the distribution. Our information structure and the assumption of uncertain value distributions characterize a situation where the buyer has one-sided private information about his long-term preference but also has some uncertainty about the short-term values in the future.

The assumption that the buyer's value distribution is uncertain can be illustrated in the following examples. Imagine that the buyer rents an apartment in a city from the landlord (seller) in multiple periods. The buyer is not completely sure about his value in future periods since it may depend on how much time he will spend in the city. So it is reasonable to assume that the buyer's value of renting the apartment in each period is drawn from a distribution. The buyer's value distribution, however, is decided by some private information of the buyer, for instance, whether the buyer's family lives close to the city. If the buyer's family lives close to the city, the buyer may spend more time in the city, and the value distribution generates a high value with a higher probability.

Another example is that an intermediate producer repeatedly rents a durable good from a monopolist, for example, a construction company rents big equipments every time when a new project begins. The producer's value of consuming the durable good in each period depends on the quality of his final products in that period. The quality of the producer's final products depends on both the producer's technology and some random effects. The producer may have a superior technology or an inferior technology, and the probability of generating a good final product is higher with a superior technology. Both the technology and the quality of the final products are private information of the producer.

The main result we find is that the seller is indeed better off when she has sufficiently optimistic ex ante beliefs of the favorable distribution, compared to the model of Hart and Tirole (1988) with the same ex ante probability of high-value buyer. The unique equilibrium outcome is for the seller to offer a high price and for the buyer with a high value to accept the offer in each period.[4] Intuitively, if the seller always offers a high price in the second period no matter whether the first-period offer is accepted or rejected, the buyer has no incentive to play strategically in the first period and will simply adopt a strategy of truth-telling, i.e., accepting the offer if and only if it is less than or equal to his value. Thus, buyer types who have the same value but draw from different distributions will behave the same. Given that the favorable distribution has a higher probability of drawing a high value, the seller becomes more optimistic after acceptance and more pessimistic after rejection. However, since the seller's ex ante belief of the favorable distribution is sufficiently optimistic, the seller's posterior beliefs after both acceptance and rejection will still be optimistic enough for her to offer a high price in the second period.

The equilibrium outcome described above, however, cannot hold when the seller has moderate ex ante beliefs. In this case, all the buyer types still truthfully reject any offer greater than their value. But the high-value buyer from the favorable distribution strategically randomizes to reject a range of prices less than but close to the high value, in order to conceal information about his type. For the low-value types, multiple equilibrium strategies are found for a range of prices less than but close to the low value. It is an equilibrium strategy for both low-value types to accept the offer or to reject the offer, or for the low-value buyer from the unfavorable distribution to randomize and the low-value buyer from the favorable distribution to reject the offer. When the seller has moderate ex ante beliefs,

[4] The equilibria in this paper refer to those that survive a refinement which is a variant of the D_1 criterion in signaling games (Cho and Kreps 1987, Banks and Sobel 1987). We discuss the equilibrium concept in more detail in section 2 and Appendix A.

the buyer is found more strategic since it is easier to affect the seller's posterior belief. The seller's revenue, however, can still be higher than that in Hart and Tirole (1988) in this case. Sufficient conditions for the seller to be better off are provided.

Two other papers also examine a rental model in which a non-anonymous buyer's value randomly changes over time.[5] Kennan (2001) analyzes infinitely repeated contract negotiations where the buyer has persistent (but not permanent) private information. The buyer's value is assumed to change according to a two-state Markov chain. Kennan (2001) focuses on the cyclic screening equilibria in which several pooling offers in sequence make the seller more and more optimistic and the seller makes an aggressive screening offer eventually.

The paper closest to our study is Loginova and Taylor (2007). They investigate a two-period model where the monopolist employs price experimentations to learn the permanent demand parameter of the buyer. Although we have benefited a lot from reading their paper, the two papers were developed independently and differ significantly in the modeling and results.

First, we assume that the value distribution may either be favorable or unfavorable, with the favorable distribution generating a high value with a higher probability. Loginova and Taylor (2007) assume that the value distribution is represented by λ, which is a continuous random variable distributed on $[0, 1]$, and a distribution represented by λ generates a high value with probability λ and a low value with probability $1 - \lambda$. We keep our model simpler so that we can completely characterize the equilibria and compare the seller's revenue with that in Hart and Tirole (1988).

Second, Loginova and Taylor (2007) assume that there is no discounting, while the discount rate is between 0 and 1 in our model. This difference has several effects on the results. First, Loginova and Taylor (2007) show that the buyer has a unique equilibrium strategy when the first-period offer is greater than the low value, by assuming the continuity of λ and no discounting. We achieve the same result after refinement. Multiple equilibrium strategies, however, reappear in Loginova and Taylor (2007) if discounting is introduced into the model. Second, similar to our results when the seller has a moderate prior, Loginova and

[5] Several other papers also allow the buyer's valuations to vary over time. Blume (1998) and Battaglini (2005) study long-term contracting. Biehl (2001) analyzes a durable-goods model with anonymous buyers.

Taylor (2007) find multiple equilibrium strategies for the low-value types when the first-period offer is less than but close to the low value. They focus on two equilibria: the Good Equilibrium and the Bad Equilibrium. In the Good Equilibrium, all the low-value types accept the offer. In the Bad Equilibrium, a low-value buyer with high λ strategically rejects the offer. The main finding in the Good Equilibrium is that the seller never offers a first-period price that yields her valuable information about the buyer's permanent demand parament λ. In our model, the seller offers a price that yields valuable information in the Good Equilibrium, if the discount rate is low enough or the seller's prior is high enough.

The remainder of the paper is organized as follows. Section 2 describes the model. Section 3 presents players' equilibrium strategies in the second period. We discuss some preliminary results in Section 4. Section 5 presents the set of equilibria. Section 6 compares the seller's revenue in this model with that in Hart and Tirole (1988). Section 7 concludes. Appendix A provides a detailed discussion about equilibrium concept. All the proofs are in Appendix B.

2.2 THE MODEL

One seller and one buyer bargain over renting a durable good in two periods $t = 1, 2$. The seller's cost is assumed to be 0. The buyer has private information about both his value in each period t, v_t, and the distribution which his values are drawn from. The buyer's distribution d can be the F distribution or the G distribution. The buyer's value v_t equals h with probability q^d and equals l with probability $1 - q^d$ for a given d. Assume that $0 < q^F < q^G < 1$ and $0 < l < h$. The G distribution is more favorable since it has a higher probability of generating a value h. The ex ante probability is α for the G distribution and $1 - \alpha$ for the F distribution.

At the beginning of the game, the buyer privately observes the realization of distribution d, which is fixed throughout the game. At the beginning of each period t, the buyer's valuation v_t is drawn from the realized distribution d independently across time periods. After the buyer privately observes v_t, the seller offers a price $p_t \in \mathbb{R}$, and then the buyer chooses an action $a_t \in \{0, 1\}$, where $a_t = 1$ means acceptance and $a_t = 0$ means rejection.

Both the seller and the buyer are assumed to be risk-neutral. If the buyer accepts the seller's offer in period t, the buyer's payoff is $v_t - p_t$ and the seller's payoff is p_t in period t. They both gain nothing in period t if p_t is rejected. The two players share a common discount factor δ, and both of them maximize the discounted present value of expected payoffs. All of the information above is common knowledge.

$\theta_1 = (d, v_1)$ is referred as the buyer's type in period 1 and $\theta_2 = (d, v_1, v_2)$ as the buyer's type in period 2. Since we focus on the buyer's behavior in period 1, it is helpful to notice that in period 1 there are four buyer types: (G, l), (F, l), (G, h), and (F, h). Denote h_t^S as the history observed by the seller before she announces p_t and h_t^B as the history observed by the buyer before he chooses a_t. Specifically, $h_1^S = \varnothing$, $h_1^B = (\theta_1, p_1)$, $h_2^S = (p_1, a_1)$ and $h_2^B = (\theta_2, p_1, a_1, p_2)$. A behavioral strategy for the seller, σ^S, assigns probability (or density) $\sigma^S(p_t \mid h_t^S)$ to p_t given any history h_t^S for $t = 1, 2$. A behavioral strategy for the buyer, σ^B, assigns probability $\sigma^B(a_t \mid h_t^B)$ to a_t given any history h_t^B for $t = 1, 2$. For convenience, let $\sigma^B(h_t^B) \equiv \sigma^B(a_t = 1 \mid h_t^B)$ denote the probability that the buyer accepts p_t given history h_t^B, since the buyer can only choose to accept or reject an offer.

Let $\gamma(h_t^S)$ denote the probability that the seller's belief assigns to the G distribution at the beginning of period t given history h_t^S. The seller's ex ante belief of the G distribution is α. Based on α, the seller forms her ex ante beliefs over the buyer's type θ_1. After offering p_1 and observing a_1, the seller updates her belief of θ_1, using Bayes' rule whenever possible. Then the seller's posterior belief of the G distribution is formed based on her posterior belief of θ_1. Notice that $\gamma(p_1, 0)$ and $\gamma(p_1, 1)$ denote the seller's belief of $d = G$ given that p_1 is rejected and accepted respectively.

The equilibrium concept used in this paper is strong Perfect Bayesian equilibrium. Bayes' rule is used to update the seller's belief conditional on reaching any price p_1, even if p_1 is off the equilibrium path. We also employ a refinement which is a variant of criterion D_1 in the signalling game (Cho and Kreps 1987, Banks and Sobel 1987). In Appendix A, we formally define criterion D_1 and give an example on how criterion D_1 can help select an equilibrium.

2.3 THE SECOND-PERIOD EQUILIBRIUM STRATEGIES

We start the analysis from the second period. Since it is the last period, the buyer accepts p_2 if and only if p_2 does not exceed v_2.[6] Then the optimal p_2 for the seller is either l or h. The seller offers $p_2 = l$ if her posterior belief of $v_2 = h$ is less than the cutoff belief l/h and offers $p_2 = h$ if it is greater than l/h. Notice that the seller always offers $p_2 = h$ if the probability of drawing a high value from the F distribution, q^F, is greater than l/h, and the seller always offers $p_2 = l$ if the probability of drawing a high value from the G distribution, q^G, is smaller than l/h, regardless of history h_2^S. To make the problem more interesting, we assume $q^F < l/h < q^G$ in this paper. Since the seller's posterior belief of $v_2 = h$ is $q^G\gamma + q^F(1-\gamma)$ if her posterior belief of $d = G$ is γ, the seller offers $p_2 = l$ if her posterior belief of the G distribution is less than γ^* and offers $p_2 = h$ if her posterior belief of the G distribution is greater than γ^*, where γ^* satisfies the equation $q^G\gamma^* + q^F(1-\gamma^*) = l/h$. Since the seller either offers $p_2 = l$ or $p_2 = h$ in period 2, let $x(h_2^S)$ denote the probability that $p_2 = l$ and $1 - x(h_2^S)$ denote the probability that $p_2 = h$ after history h_2^S. Lemma 1 formally states the discussion above.

Lemma 1. *In any PBE, the buyer's strategy in the second period is*

$$\sigma^B(h_2^B) = \begin{cases} 1, & \textit{if } p_2 \leq v_2; \\ 0, & \textit{if } p_2 > v_2, \end{cases}$$

and the seller's strategy in the second period is

$$x(h_2^S) = \begin{cases} 1, & \textit{if } \gamma(h_2^S) < \gamma^*; \\ 0, & \textit{if } \gamma(h_2^S) > \gamma^*; \\ \in [0,1], & \textit{if } \gamma(h_2^S) = \gamma^*, \end{cases}$$

where $\gamma^* = (l/h - q^F)/(q^G - q^F)$.

[6] As noted by Fudenberg and Tirole (1983), the buyer is indifferent between acceptance and rejection if $p_2 = v_2$, but on the equilibrium path the buyer accepts $p_2 = v_2$ with probability one.

2.4 PRELIMINARY RESULTS

In this section we first present two preliminary results that provide us the intuition of equilibria. Roughly speaking, the first observation is that no separation between one buyer type and the other three types is possible in equilibrium. Second, the seller is more pessimistic after rejection of a first-period offer rather than after acceptance of the offer.

We start with defining the cutoff value for each buyer type, which the buyer type is indifferent between accepting and rejecting. From the players' equilibrium strategies in the second period, the buyer's expected payoff from accepting p_1 is $v_1 - p_1 + \delta q^d x(p_1, 1)(h - l)$, where $v_1 - p_1$ is the buyer's gain in the first period from accepting p_1, and $q^d x(p_1, 1)$ is the product of the probability for the buyer type to draw an h value in period 2 and the probability for the seller to offer $p_2 = l$ after acceptance of p_1. Correspondingly, the buyer's expected payoff from rejecting p_1 is $\delta q^d x(p_1, 0)(h - l)$, where $x(p_1, 0)$ is the probability for the seller to offer $p_2 = l$ after rejection of p_1. By comparing the payoffs from accepting and rejecting p_1, the buyer type (d, v_1) accepts p_1 with probability one if p_1 is smaller than $v_1 + \delta q^d[x(p_1, 1) - x(p_1, 0)](h - l)$ and rejects p_1 with probability one if it is greater than $v_1 + \delta q^d[x(p_1, 1) - x(p_1, 0)](h - l)$. Therefore, the buyer's cutoff value is defined as follows.

Definition 1. *Denote* $C(d, v_1) \equiv v_1 + \delta q^d[x(p_1, 1) - x(p_1, 0)](h - l)$ *as the* Cutoff Value *for buyer type* $\theta_1 = (d, v_1)$ *given* $x(p_1, 0)$ *and* $x(p_1, 1)$.

Lemma 2. *In any PBE, the probability for buyer type* $\theta_1 = (d, v_1)$ *to accept* p_1 *is*

$$\sigma^B(\theta_1, p_1) = \begin{cases} 1, & \text{if } p_1 < C(d, v_1); \\ 0, & \text{if } p_1 > C(d, v_1); \\ \in [0, 1], & \text{if } p_1 = C(d, v_1). \end{cases}$$

By definition the buyer's cutoff value depends not only on his type but also on the seller's second-period strategy. Figure 2.1 below describes how the order of all buyer types' cutoff values depends on the seller's strategies in the second period.

From Figure 2.1 we can see how the two elements of the buyer's type v_1 and d, along with the seller's second-period strategy, affect the buyer's behavior. First, the buyer types with $v_1 = l$ always have a smaller cutoff value than types with $v_1 = h$. This is robust with

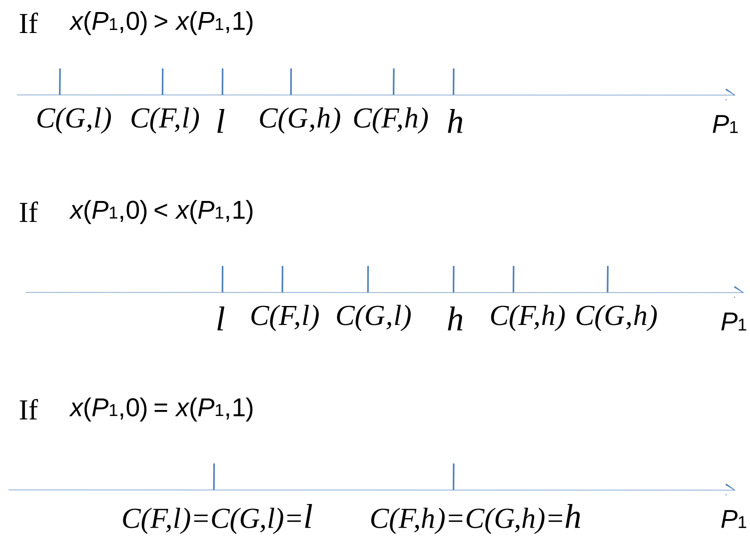

Figure 2.1: The Order of Cutoff Values

the seller's strategy in the second period. Loosely speaking, a low-value type is more likely to reject an offer than a high-value type, no matter which distribution their values are drawn from.

Second, the effect of the distribution d on the order of cutoff values depends on the seller's strategy in the second period. If the seller offers $p_2 = l$ with a larger probability when p_1 is rejected than accepted (i.e., $x(p_1, 0) > x(p_1, 1)$), then the cutoff value for buyer type (d, v_1) is smaller than his value v_1. The intuition is that, since the potential payoff in the second period is larger after rejection than acceptance of p_1, the buyer is willing to offset some current benefit and reject an offer less than his current value in order to gain a larger payoff in the future. Furthermore, the cutoff value of buyer type (G, v_1) is smaller than the cutoff value of buyer type (F, v_1), since buyer type (G, v_1) has a larger probability of drawing an h value in period 2 and therefore has a larger potential payoff after rejection of p_1. On the contrary, if the seller offers $p_2 = l$ with a larger probability when p_1 is accepted than rejected (i.e., $x(p_1, 0) < x(p_1, 1)$), the cutoff value for buyer type (d, v_1) is greater than his value v_1, and the cutoff value of buyer type (G, v_1) is greater than that of buyer type (F, v_1). The intuition is similar to the last point. Since acceptance of p_1 gives the buyer a potentially larger payoff in the second period than rejection of p_1, the buyer is willing to incur a loss and accept p_1 which exceeds his value v_1. Finally, if the seller offers $p_2 = l$ with the same probability when p_1 is rejected and accepted, then the distribution d does not have any effect on the buyer's first-period behavior, and the buyer's cutoff value coincides with his value v_1, thus the buyer truthfully reveals his value in the first period.

Next we discuss the two observations mentioned before.

Observation 1: No separation between one buyer type and the other three types is possible in equilibrium.

The key point of Observation 1 is that the seller cannot separate one buyer type from the others if she offers $p_2 = l$ with a different probability after acceptance and rejection of p_1 in equilibrium. Intuitively, the seller offers $p_2 = l$ with a larger probability after one action than the other only when she is more pessimistic after the first action, however, since the buyer types from the G distribution have a larger incentive to take the action that makes the seller more pessimistic, the seller then becomes more optimistic if that action is taken by a single type from the G distribution. So it easily reaches a contradiction.

It is intuitive to see this from Figure 2.1. Suppose the seller offers $x(p_1, 0) > x(p_1, 1)$. The seller can separate one buyer type from the other three types if p_1 is rejected only by type (G, l) (following which the seller should offer $x(p_1, 0) = 0$), or p_1 is accepted only by type (F, h) (following which the seller should offer $x(p_1, 1) = 1$). Both cases contradict with the assumption $x(p_1, 0) > x(p_1, 1)$. It is easy to verify that the similar contradiction is reached when the seller offers $x(p_1, 0) < x(p_1, 1)$.

The remaining case is that the seller offers $p_2 = l$ with the same probability after acceptance and rejection of p_1 (i.e., $x(p_1, 0) = x(p_1, 1)$). Figure 2.1 shows that two buyer types with the same v_1 behave similarly: both accept p_1 less than their value v_1 and reject p_1 greater than v_1. Therefore, the seller cannot distinguish the buyer types who have the same v_1 but draw from different distributions.

From the discussion above, the seller cannot separate one buyer type from the other three given any second-period strategy of the seller. Observation 1 have two important and related implications. First, the seller is not able to identify the buyer's distribution in equilibrium. Second, the screening between two buyer types with $v_1 = l$ and two buyer types with $v_1 = h$ is an important feature of equilibria.

Observation 2: The seller is more pessimistic after rejection of p_1 than acceptance of p_1.

The intuition for Observation 2 is as follows. Suppose that the seller is more optimistic after rejection of p_1 and offers $x(p_1, 0) < x(p_1, 1)$ in the second period. From Observation 1, no p_1 separates one buyer type from the other three in equilibrium. Suppose then p_1 separates two buyer types with $v_1 = l$ from the types with $v_1 = h$. From Figure 2.1 it must be the case that the buyer types with $v_1 = l$ reject the offer and the buyer types with $v_1 = h$ accept the offer. Then the seller must be more pessimistic after rejection than acceptance of p_1, given that the G distribution has a higher probability of drawing an h value. The seller then should offer $x(p_1, 0) \geq x(p_1, 1)$, which reaches a contradiction.

Let $\Psi(p_1, a_1)$ denote the probability that action a_1 is taken in the continuation game following p_1. Observation 2 can be expressed more formally in the following lemma.

Lemma 3. *If $\Psi(p_1, 1) \in (0, 1)$ for a given p_1 in a PBE, then $x(p_1, 0) \geq x(p_1, 1)$.*

As a summary, the preliminary results in this section imply that the seller cannot separate one buyer type from the other three types in equilibrium. Related to that, nor can the seller learn perfectly the buyer's distribution. However, the seller updates her belief about the buyer's distribution, when she separates buyer types with $v_1 = l$ and buyer types with $v_1 = h$. Finally, the seller is always more pessimistic when p_1 is rejected than accepted, if both acceptance and rejection of p_1 occur with a positive probability in the continuation game following p_1. A more rigorous description of Observation 1 can be found in Lemma 16 and Lemma 17, which are delegated to Appendix B.

2.5 THE EQUILIBRIA

In this section we present the equilibria of the game. The equilibrium outcome and players' strategy are found to greatly depend on the seller's ex ante beliefs. So we first divide the seller's ex ante beliefs into different ranges and then discuss the set of equilibria correspondingly. When the equilibrium price in the first period is accepted by all buyer types, we call the equilibrium pooling. When the first-period offer is accepted and rejected by more than one buyer type, we call the equilibrium semi-separating.

The results in the last section suggest that, the seller cannot separate one buyer type from the other three types, but it is more likely for the seller to distinguish buyer types with $v_1 = l$ and buyer types with $v_1 = h$. So we classify the seller's ex ante beliefs according to whether the seller's posterior beliefs conditional on $v_1 = l$ and $v_1 = h$ are greater than or less than the cutoff belief γ^*.

Since the favorable distribution G has a larger probability of generating a high value, the seller's posterior belief conditional on $v_1 = l$ is lower than her ex ante belief and her posterior belief conditional on $v_1 = h$ is higher than her ex ante belief. Furthermore, since the F distribution and the G distribution have the same support, the seller's posterior beliefs conditional on $v_1 = l$ or $v_1 = h$ are always between 0 and 1. Recall that whether the seller offers the low price or the high price in the second period depends on whether her posterior belief is greater or smaller than the cutoff belief γ^*. Therefore, when the seller's ex ante belief is small enough, her posterior belief is smaller than the cutoff belief γ^* even conditional on $v_1 = h$. On the contrary, when the seller's ex ante belief is large enough, her posterior belief can be greater than the cutoff belief conditional $v_1 = l$. Only when the seller has moderate ex ante beliefs, her posterior belief is above γ^* conditional on $v_1 = h$ and below γ^* conditional on $v_1 = l$. The following equations and graph give us a more clear illustration.

Define functions

$$\overline{\gamma}(\alpha) \equiv \frac{\alpha q^G}{\alpha q^G + (1-\alpha) q^F},$$

and

$$\underline{\gamma}(\alpha) \equiv \frac{\alpha(1-q^G)}{\alpha(1-q^G) + (1-\alpha)(1-q^F)}.$$

Define $\widetilde{\alpha} \equiv \overline{\gamma}^{-1}(\gamma^*)$ and $\widehat{\alpha} \equiv \underline{\gamma}^{-1}(\gamma^*)$.[7] $\overline{\gamma}(\alpha)$ and $\underline{\gamma}(\alpha)$ are the seller's posterior beliefs of the G distribution conditional on $v_1 = h$ and $v_1 = l$ respectively.

[7] Both $\overline{\gamma}(\alpha)$ and $\underline{\gamma}(\alpha)$ are continuous and increasing in α; $\underline{\gamma}(\alpha) < \alpha < \overline{\gamma}(\alpha)$ for $\alpha \in (0,1)$; $\underline{\gamma}(\alpha) = \overline{\gamma}(\alpha) = \alpha$ for $\alpha \in \{0,1\}$. So $\widetilde{\alpha}$ and $\widehat{\alpha}$ are well-defined and $\widetilde{\alpha} < \gamma^* < \widehat{\alpha}$.

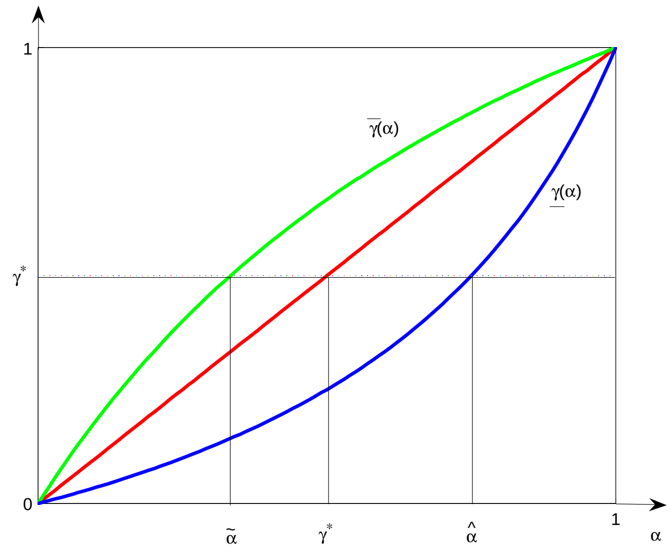

Figure 2.2: Seller with Different Prior Beliefs

In Figure 2.2, the x axis is the seller's ex ante belief α. The 45^0 line represents the seller's belief when she has no more information than at the beginning of the game. $\overline{\gamma}(\alpha)$ and $\underline{\gamma}(\alpha)$ are generated with $q^F = 0.4$, $q^G = 0.8$, and $l/h = 0.6$. The curve $\overline{\gamma}(\alpha)$ is above the 45^0 line since the seller's belief conditional on $v_1 = h$ is larger than her ex ante belief. On the contrary, the seller becomes more pessimistic conditional on $v_1 = l$ relative to her ex ante belief, i.e., the curve $\underline{\gamma}(\alpha)$ is below the 45^0 line. Furthermore, when the seller's ex ante belief is smaller than $\widetilde{\alpha}$ (greater than $\widehat{\alpha}$), her posterior beliefs conditional on $v_1 = h$ and $v_1 = l$ are both below (above) γ^*. When the seller's ex ante belief is between $\widetilde{\alpha}$ and $\widehat{\alpha}$, her posterior belief conditional on $v_1 = h$ is above γ^* and her belief conditional on $v_1 = l$ is below γ^*.

According to the seller's ex ante belief of the G distribution, we define a seller *Pessimistic* if $0 < \alpha < \widetilde{\alpha}$, *Moderately Pessimistic* if $\widetilde{\alpha} < \alpha < \gamma^*$, *Moderately Optimistic* if $\gamma^* < \alpha < \widehat{\alpha}$, and *Optimistic* if $\widehat{\alpha} < \alpha < 1$. As in the previous literature, the knife-edge cases are omitted in this paper.

In Lemma 3 we have shown that the seller offers $p_2 = l$ after rejection of p_1 with a probability at least as large as after acceptance of p_1, if p_1 is both accepted and rejected with a positive probability in the continuation game. The following lemma states more specifically that the seller offers $p_2 = l$ with a strictly larger probability after rejection of p_1 than acceptance of p_1 only if the seller has a moderate ex ante belief. With an extreme ex ante belief, the seller offers $p_2 = l$ with the same probability after rejection and acceptance of p_1. These results hold when both acceptance and rejection of p_1 are reached with a positive probability in the continuation game.

Lemma 4. *If $\Psi(p_1, 1) \in (0, 1)$ for a given p_1 in a PBE, then*

(*i*) $x(p_1, 0) > x(p_1, 1) \Rightarrow \alpha \in [\widetilde{\alpha}, \widehat{\alpha}]$;

(*ii*) $\alpha \in (0, \widetilde{\alpha}) \cup (\widehat{\alpha}, 1) \Rightarrow x(p_1, 0) = x(p_1, 1)$.

2.5.1 Seller with Extreme Ex Ante Beliefs

When the seller has a pessimistic or an optimistic belief, the buyer's strategy is to truthfully reveal his value in both periods: accept an offer no greater than his value and reject an offer otherwise. Given the buyer's strategy, the two buyer types who have the same v_1 but draw from different distributions behave the same, so the seller can only distinguish the buyer's value v_1 but cannot tell the buyer's distribution d. As demonstrated in Figure 2.2, the seller's posterior beliefs conditional on $v_1 = l$ and $v_1 = h$ are both below (above) the cutoff belief γ^* when she has a pessimistic (an optimistic) ex ante belief. Therefore, a pessimistic seller always offers $p_2 = l$ and an optimistic seller always offers $p_2 = h$ given the buyer's equilibrium strategy. Since the seller's second-period offer does not depend on the buyer's acceptance/rejection action in the first period, truth-telling is the buyer's equilibrium strategy in period 1.

Given the buyer's first-period strategy, a pessimistic seller offers equilibrium price $p_1 = l$ and all buyer types accept p_1.

Proposition 5 (Pessimistic Seller). *When the seller is pessimistic* $(0 < \alpha < \widetilde{\alpha})$, *there is a unique D_1 equilibrium outcome: the seller offers $p_t = l$ and all buyer types accept p_t, for $t = 1, 2$.*

When the seller has an optimistic ex ante belief, she always offers $p_t = h$ on the equilibrium path and the h-value types accept the offer.

Proposition 6 (Optimistic Seller). *When the seller is optimistic* ($\widehat{\alpha} < \alpha < 1$)*, there is a unique* D_1 *equilibrium outcome: the seller offers* $p_t = h$*, the buyer types with* $v_t = h$ *accept* p_t *and the buyer types with* $v_t = l$ *reject* p_t*, for* $t = 1, 2$*.*

Remark: Although the buyer has a unique equilibrium strategy when the seller has extreme ex ante beliefs, multiple equilibria arise since different beliefs can be assigned after all buyer types reject $p_1 \leq l$ or after all buyer types accept $p_1 > h$ to support the unique D_1 equilibrium outcome presented in Proposition 5 and 6.

The outcome in Proposition 6 is of particular interest to us. In the model of Hart and Tirole (1988), in which the buyer's value is private information but the value distribution is common knowledge, it does not happen in any equilibrium that the h-value buyer accepts $p_1 = h$ with probability one, even if the seller has a very optimistic ex ante belief of the buyer's value. The intuition is that, the seller will offer $p_2 = l$ after rejection if the h-value buyer accepts $p_1 = h$ with probability one, and then the h-value buyer has an incentive to deviate to reject $p_1 = h$. Therefore, when the seller has a sufficiently optimistic ex ante belief, introducing the uncertainty about the buyer's value distribution improves the seller's revenue. We will discuss the revenue comparison between our model and Hart and Tirole's (1988) in more detail in Section 6.

2.5.2 Seller with Moderate Ex Ante Beliefs

When the seller has a moderate ex ante belief, the buyer's strategy is quite different from when the seller has an extreme ex ante belief. First, the buyer does not always truthfully reveal his value anymore. Since the separation between l-value buyer types and h-value buyer types is an important feature of the equilibria, the seller offers high price in the second period conditional on $v_1 = h$ and low price conditional on $v_1 = l$. This gives the l-value buyer types an incentive to signal their current values, and at the same time, gives the h-value buyer types an incentive to conceal their first-period value. Second, the buyer has multiple equilibrium strategies instead of unique equilibrium strategy. In particular, two l-value buyer types may reject or accept p_1 less than l but close to l in equilibrium. Intuitively, l-value types can signal their value by rejecting such a p_1 and have a low offer in the second period. On the other hand, it is also optimal for the l-value types to accept such a p_1 if the buyer believes that the seller offers a same price after acceptance and rejection of p_1. Finally, some buyer types sometimes use a mixed strategy instead of a pure strategy, which will be elaborated in the discussion below. Next we present two equilibrium strategies of the buyer in Figure 2.3 and 2.4 and provide the general results in Lemma 7.

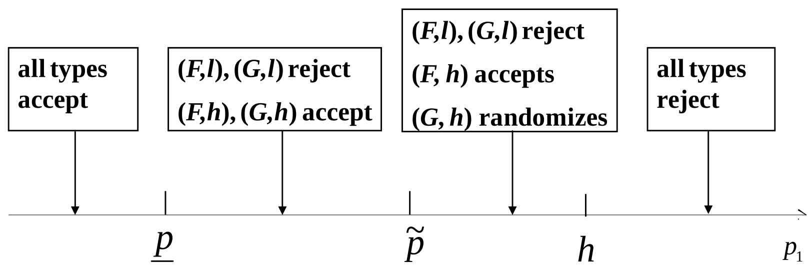

Figure 2.3: One Equilibrium Strategy of the Buyer for Moderate Seller

In Figure 2.3, it is part of an equilibrium strategy for buyer types with $v_1 = l$ to reject and buyer types with $v_1 = h$ to accept $p_1 \in (\underline{p}, \widetilde{p}]$, where $\underline{p} = l - \delta q^F(h-l)$ and $\widetilde{p} = h - \delta q^G(h-l)$. Given the buyer's behavioral strategy for $p_1 \in (\underline{p}, \widetilde{p}]$, the seller is able to distinguish a high-value buyer type from a low-value buyer type, so she offers $p_2 = l$ after rejection of p_1 and $p_2 = h$ after acceptance of p_1. Therefore, the buyer types with $v_1 = l$ have an incentive to reject an offer less than l in order to signal their current value, since there is a positive probability for them to draw a high value in the second period and gain a positive payoff if the seller offers $p_2 = l$. The buyer types with $v_1 = h$, however, still accept an offer in this range since it is still much less than their current value h. The cutoff prices $\underline{p} = l - \delta q^F(h-l)$ and $\widetilde{p} = h - \delta q^G(h-l)$ are the highest prices that buyer type (F, l) and (G, h) are willing to accept respectively given that the seller offers $p_2 = l$ after rejection and $p_2 = h$ after acceptance.

For p_1 less than $\underline{p}$, all buyer types including those with $v_1 = l$ accept p_1 since the payoff in the second period cannot compensate the loss from rejecting such a low price in the first period. Correspondingly, all buyer types reject p_1 greater than h.

For p_1 greater than $\widetilde{p}$ but less than h, the buyer types with $v_1 = l$ reject p_1, buyer type

(F, h) accepts p_1, and buyer type (G, h) plays mixed strategy. To see the intuition, suppose all the buyer types reject the offer, then buyer type (F, h) has an incentive to deviate to accept p_1 since he gains an positive payoff in the first period and gets the lowest offer in the second period by revealing his distribution. Given that buyer type (F, h) accepts p_1, buyer type (G, h) has an incentive to accept p_1 as well. However, the seller's posterior belief after acceptance of p_1 is then above the cutoff belief γ^* since her ex ante belief is moderate. The seller then offers $p_2 = l$ after rejection and $p_2 = h$ after acceptance of p_1. Since p_1 is greater than $\widetilde{p}$, buyer type (G, h) then has an incentive to reject p_1. Therefore, buyer type (G, h) randomizes to accept and reject $p_1 \in (\widetilde{p}, h]$ in equilibrium.

The equilibrium strategy in Figure 2.3, however, is not the unique equilibrium strategy. Figure 2.4 shows another equilibrium strategy of the buyer.

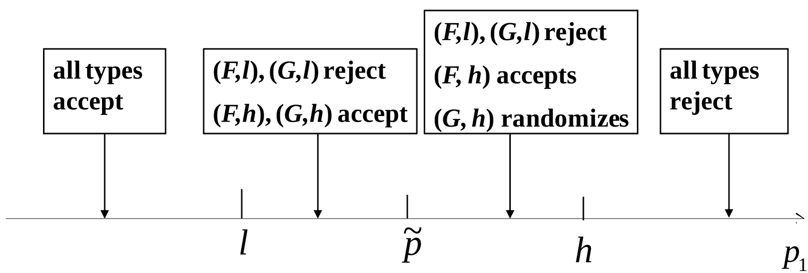

Figure 2.4: Another Equilibrium Strategy of the Buyer for Moderate Seller

In Figure 2.4, it is part of an equilibrium strategy for all buyer types to accept $p_1 \leq l$. Given that all buyer types accept p_1, the seller can always assign the same posterior belief after rejection as after acceptance of $p_1 \leq l$ by the definition of PBE, then all buyer types should accept p_1 less than v_1.

Figure 2.5 and Lemma 7 summarize the buyer's strategy. Besides the two strategies discussed above, it is also part of an equilibrium strategy that buyer type (G, l) rejects $p_1 \in (\underline{p}, l]$, the buyer types with $v_1 = h$ accept $p_1 \in (\underline{p}, l]$, and buyer type (F, l) plays mixed strategy.

Lemma 7. *When the seller has a moderate prior belief* $(\widetilde{\alpha} < \alpha < \widehat{\alpha})$, *the buyer's strategy in*

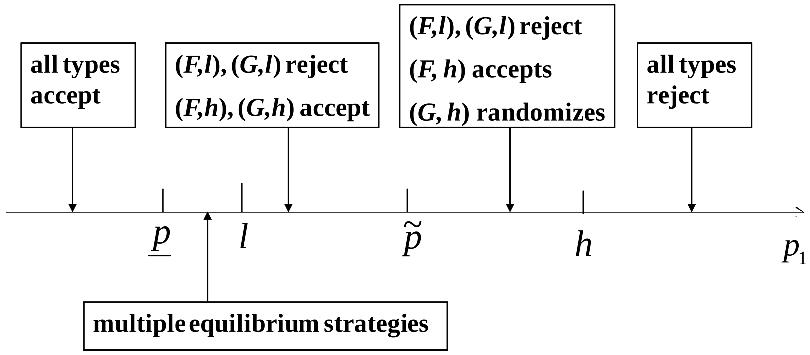

Figure 2.5: Buyer's Strategy for Moderate Seller

a D_1 equilibrium is as follows[8]*:*

(i) if $p_1 \leq \underline{p}$, all buyer types accept p_1;

(ii) if $\underline{p} < p_1 \leq l$, (1) all buyer types accept p_1; or (2) types (F, l) and (G, l) reject p_1, and types (F, h) and (G, h) accept p_1; or (3) type (G, l) rejects p_1, type (F, l) randomizes, and types (G, h) and (F, h) accept p_1;

(iii) if $l < p_1 \leq \widetilde{p}$, types (F, l) and (G, l) reject p_1, and types (F, h) and (G, h) accept p_1;

(iv) if $\widetilde{p} < p_1 \leq h$, types (F, l) and (G, l) reject p_1, type (G, h) randomizes, and type (F, h) accepts p_1;

(v) if $p_1 > h$, all buyer types reject p_1.

There are several points we find important about the equilibrium strategies presented above. First, although some buyer types strategically reject an offer less than their first-period value, all buyer types truthfully reject $p_1 > v_1$, and the buyer never incurs a loss in any period in equilibrium. As shown in Lemma 3, the seller offers $p_2 = l$ after rejection with a probability at least as large as after acceptance of p_1, so no buyer type has an incentive to accept an offer larger than his value.

Second, in the equilibrium strategy in Figure 2.3, the highest first-period offer accepted by all buyer types, $\underline{p}$, is less than the buyer's lowest value l. This feature is also found by Blume (1990), Kennan (2001), and Loginova and Taylor (2007). In all these models including

[8]We require that the buyer's strategy is left continuous for the cutoff prices $p_1 \in \{\underline{p}, l, \widetilde{p}, h\}$, that is, the behavioral strategy following the cutoff prices are the same as the strategy following $p_1 - \epsilon$.

ours, a buyer type with a low value in the current period has a positive probability of drawing a high value in the next period, so the low-value type will reject an offer less than but close to the low value if rejection can help the seller price discriminate in the next period.

Finally, the equilibrium strategies in this paper are different from those in Kennan (2001). In Kennan (2001), the buyer's value changes according to a Markov process, so the seller's posterior belief becomes more optimistic when all buyer types accept a pooling offer, and the seller offers an aggressive screening offer following acceptance of several pooling offers when her posterior belief grows beyond some cutoff. This pattern is described as a cyclic equilibrium. In our model, the seller's posterior belief is the same as her ex ante belief after acceptance of a pooling offer. So we do not expect that the same pattern as in the cyclic equilibrium emerges in this model, even in a longer horizon.

Next we discuss the seller's optimal p_1 and then conclude by describing the equilibria of the game. For moderately pessimistic and moderately optimistic seller respectively, we first provide the seller's payoffs from offering the cutoff prices $p_1 \in \{\underline{p}, l, \widetilde{p}, h\}$, and then discuss the conditions for there to exist pooling and semi-separating equilibria.

From Lemma 7, buyer types (F, l) and (G, h) may play a mixed strategy for some range of p_1. We use X^* and Y^* to denote the probabilities that buyer types (F, l) and (G, h) randomize to reject p_1 respectively when the seller has a moderately pessimistic ex ante belief, and use X^{**} and Y^{**} to denote the corresponding probabilities when the seller has a moderately optimistic ex ante belief. The explicit expressions for X^*, Y^*, X^{**}, and Y^{**} are independent of p_1 and calculated in the proof of Lemma 7. All cases presented below in Proposition 8-11 arise for a non-negligible set of parameters.[9]

2.5.2.1 Moderately Pessimistic Seller ($\widetilde{\alpha} < \alpha < \gamma^*$) Define the following payoffs from offering $p_1 \in \{\underline{p}, l, \widetilde{p}, h\}$ for the seller with moderately pessimistic ex ante beliefs. Given Lemma 7, the seller can always guarantee payoff U_1 by offering $p_1 = \underline{p}$ and $p_2 = l$, with p_1 and p_2 accepted by all buyer types. U_1 is the seller's lowest payoff from a pooling offer.

$$U_1 = \underline{p} + \delta l;$$

[9] This is proved using Mathematica. The program is available upon request.

Since there are multiple equilibrium strategies for the buyer given $p_1 \in (\underline{p}, l]$, the seller's payoff from offering $p_1 = l$ depends on the buyer's strategy. Suppose that all buyer types choose to accept $p_1 = l$, then the seller can achieve the highest payoff from a pooling offer U_2 by offering $p_1 = p_2 = l$, with p_1 and p_2 accepted by all buyer types.

$$U_2 = l + \delta l;$$

However, if given $p_1 = l$, buyer type (G, l) rejects p_1, buyer type (F, l) randomizes, and buyer types (F, h) and (G, h) accept p_1, then the seller's payoff from offering $p_1 = l$ is U_3.

$$U_3 = [\alpha q^G + (1-\alpha) q^F + (1-\alpha)(1-q^F)(1-X^*)]l + \delta l;$$

Notice that if both l-value buyer types choose to reject and both h-value buyer types choose to accept $p_1 = l$, then offering $p_1 = l$ is dominated by offering $p_1 = \widetilde{p}$. Payoff U_4 is achieved if the seller offers $p_1 = \widetilde{p}$. By Lemma 7, buyer types (F, l) and (G, l) reject p_1, buyer types (F, h) and (G, h) accept p_1, and the seller offers $p_2 = l$ if p_1 is rejected and $p_2 = h$ if p_1 is accepted.

$$\begin{aligned} U_4 &= [\alpha q^G + (1-\alpha) q^F]\widetilde{p} + \delta[\alpha (q^G)^2 + (1-\alpha)(q^F)^2]h \\ &\quad + \delta[\alpha(1-q^G) + (1-\alpha)(1-q^F)]l; \end{aligned}$$

Finally, if the seller offers $p_1 = h$, buyer type (G, h) randomizes, buyer type (F, h) accepts p_1, and buyer types (F, l) and (G, l) reject p_1, then the seller's payoff is U_5.

$$U_5 = [\alpha q^G (1-Y^*) + (1-\alpha) q^F]h + \delta l.$$

When the seller has moderately pessimistic ex ante beliefs, we find that there always exists a pooling equilibrium with $p_1 = l$ on the equilibrium path, since the highest payoff from a pooling offer, i.e. U_2, is always greater than the highest payoff from a semi-separating offer, i.e. $\max\{U_3, U_4, U_5\}$, for a moderately pessimistic seller. Surprisingly, this result implies that the payoffs of a moderately pessimistic seller is no better than those of a pessimistic seller. The equilibrium outcome with all buyer types accepting $p_1 = p_2 = l$ is the best outcome for the moderately pessimistic seller. Intuitively, although a moderately pessimistic seller has a larger ex ante belief than a pessimistic seller, the buyer's strategic action makes the seller even worse off.

Recall that the buyer has multiple equilibrium strategies for $p_1 \in [\underline{p}, l]$: all buyer types accept p_1, or the h-value types accept p_1 and the l-value types reject p_1. When the lowest payoff from a pooling offer, i.e. U_1, is greater than $\max\{U_4, U_5\}$, any $p_1^* \in [\underline{p}, l]$ can arise in a pooling equilibrium, with the l-value buyer types accepting $p_1 \leq p_1^*$ and rejecting $p_1 > p_1^*$. When U_1 is less than $\max\{U_4, U_5\}$, we can find a pooling offer $p' \in [\underline{p}, l]$ which gives the seller the same payoff as $\max\{U_4, U_5\}$. Then any $p_1^* \in [p', l]$ can arise in a pooling equilibrium, with the l-value buyer types accepting $p_1 \leq p_1^*$ and rejecting $p_1 > p_1^*$.

Proposition 8 (MP Seller: Pooling Equilibrium). *When the seller is moderately pessimistic, there always exists a pooling D_1 equilibrium with $p_1 = l$.*

(i) If $U_1 > \max\{U_4, U_5\}$, any $p_1 \in [\underline{p}, l]$ can arise in a pooling equilibrium;

(ii) If $U_1 < \max\{U_4, U_5\}$, any $p_1 \in [p', l]$, with $\underline{p} < p' < l$, can arise in a pooling equilibrium.

Proposition 9 presents the conditions for semi-separating D_1 equilibria. If the lowest payoff from a pooling offer U_1 is greater than the highest payoff from a semi-separating offer, then there is no semi-separating equilibrium. On the contrary, if U_1 is less than the highest payoff from a semi-separating offer, then semi-separating equilibria exist. Furthermore, if $p_1 = \widetilde{p}$ or $p_1 = h$ gives the highest payoff among all semi-separating offers, the equilibrium path of the semi-separating equilibria is unique. If $p_1 = l$ gives the highest semi-separating payoff, then a continuum equilibrium price $p_1 \in [p'', l]$, with $\underline{p} < p'' < l$, arises.

Proposition 9 (MP Seller: Semi-separating Equilibrium). *When the seller is moderately pessimistic, the semi-separating D_1 equilibria are characterized as follows.*

(i) If $U_1 > \max\{U_3, U_4, U_5\}$, no semi-separating equilibrium exists;

(ii) If $U_1 < \max\{U_3, U_4, U_5\} = \max\{U_4, U_5\}$, semi-separating equilibria exist and the path is unique, with $p_1 = \widetilde{p}$ or $p_1 = h$;

(iii) If $U_1 < \max\{U_3, U_4, U_5\} = U_3$, any $p_1 \in [p'', l]$, with $\underline{p} < p'' < l$, can arise in a semi-separating equilibrium, so does $p_1 = \widetilde{p}$ or $p_1 = h$ if $\max\{U_4, U_5\} > U_1$.

2.5.2.2 Moderately Optimistic Seller ($\gamma^* < \alpha < \widehat{\alpha}$) In this subsection we discuss the pooling equilibria and semi-separating equilibria for a seller with moderately optimistic ex ante beliefs. Similar to last subsection, we start with the seller's payoffs from offering the cutoff prices $p_1 \in \{\underline{p}, l, \widetilde{p}, h\}$. Each payoff discussed below is corresponding to the payoff discussed for the moderately pessimistic seller. V_1 is the seller's payoff from offering $p_1 = \underline{p}$ and $p_2 = h$, with p_1 accepted by all buyer types and p_2 accepted by types with $v_2 = h$. Notice that a moderately optimistic seller offers $p_2 = h$ when all buyer types accept p_1, which is different from a moderately pessimistic seller. V_1 is the seller's lowest payoff from a pooling offer.

$$V_1 = \underline{p} + \delta[\alpha q^G + (1-\alpha)q^F]h;$$

V_2 is the seller's payoff from offering $p_1 = l$ and $p_2 = h$, with p_1 accepted by all buyer types and p_2 accepted by types with $v_2 = h$. V_2 is the seller's highest payoff from a pooling offer.

$$V_2 = l + \delta[\alpha q^G + (1-\alpha)q^F]h;$$

V_3 is the seller's payoff from offering $p_1 = l$, buyer type (G, l) rejects p_1, buyer type (F, l) randomizes, and buyer types (F, h) and (G, h) accept p_1.

$$V_3 = [\alpha q^G + (1-\alpha)q^F + (1-\alpha)(1-q^F)(1-X^{**})]l + \delta[\alpha q^G + (1-\alpha)q^F]h;$$

V_4 is the seller's payoff from offering $p_1 = \widetilde{p}$, buyer types (F, l) and (G, l) reject p_1, buyer types (F, h) and (G, h) accept p_1, and the seller offers $p_2 = l$ if p_1 is rejected and $p_2 = h$ if p_1 is accepted.

$$\begin{aligned} V_4 &= [\alpha q^G + (1-\alpha)q^F]\widetilde{p} + \delta[\alpha(q^G)^2 + (1-\alpha)(q^F)^2]h \\ &\quad + \delta[\alpha(1-q^G) + (1-\alpha)(1-q^F)]l; \end{aligned}$$

Finally, V_5 is the seller's payoff from offering $p_1 = h$, buyer type (G, h) randomizes, buyer type (F, h) accepts p_1, and buyer types (F, l) and (G, l) reject p_1.

$$V_5 = [\alpha q^G(1-Y^{**}) + (1-\alpha)q^F]h + \delta[\alpha q^G + (1-\alpha)q^F]h.$$

The proof of next two propositions are omitted since it is similar to the proof of Proposition 8 and 9.

Proposition 10 (MO Seller: Pooling Equilibrium). *When the seller is moderately optimistic, the pooling D_1 equilibria are characterized as follows.*

(i) If $V_1 > \max\{V_4, V_5\}$, any $p_1 \in [\underline{p}, l]$ can arise in a pooling equilibrium;

(ii) If $V_1 < \max\{V_4, V_5\} < V_2$, any $p_1 \in [p''', l]$, with $\underline{p} < p''' < l$, can arise in a pooling equilibrium;

(iii) If $V_2 < \max\{V_4, V_5\}$, no pooling equilibrium exists.

Different from the results for a moderately pessimistic seller, case (*iii*) in Proposition 10 implies that it is possible for a semi-separating equilibrium to emerge even if all buyer types accept $p_1 \in (\underline{p}, l]$. That is, when the seller's ex ante belief is sufficiently optimistic, the best pooling offer does not necessarily arise as an equilibrium price. This finding is different from that of Loginova and Taylor (2007). They argue that the seller never offers a first-period price that yields valuable information about the buyer's distribution in a Good equilibrium where all buyer types accept p_1 less than l. But that conclusion depends on the assumption of no discounting. We find that case (*iii*) of Proposition 10 arises for a non-negligible set of parameters when the discount factor is sufficiently low.

When the seller is moderately optimistic, the conditions for the semi-separating D_1 equilibria are similar to those when the seller is moderately pessimistic.

Proposition 11 (MO Seller: Semi-separating Equilibrium). *When the seller is moderately optimistic, the semi-separating D_1 equilibria are characterized as follows.*

(i) If $V_1 > \max\{V_3, V_4, V_5\}$, no semi-separating equilibrium exists;

(ii) If $V_1 < \max\{V_3, V_4, V_5\} = \max\{V_4, V_5\}$, a semi-separating equilibrium exists and the path is unique, with $p_1 = \widetilde{p}$ or $p_1 = h$;

(iii) If $V_1 < \max\{V_3, V_4, V_5\} = V_3$, any $p_1 \in [p'''', l]$, with $\underline{p} < p'''' < l$, can arise in a semi-separating equilibrium, so does $p_1 = \widetilde{p}$ or $p_1 = h$ if $\max\{V_4, V_5\} > V_1$.

2.6 REVENUE COMPARISON

The most important question that this paper is concerned with is whether the seller improves her revenue and gains more monopoly power when an additional layer of uncertainty is associated with the buyer's value distribution. In this section, we address this issue by comparing the seller's revenue in our model with that in the two-period version of Hart and Tirole's (1988) rental model in which the buyer's value distribution is common knowledge.

The two-period version of Hart and Tirole's (1988) rental model is as follows. The buyer has private information about his value, which can be either high or low. The buyer's value is drawn at the beginning of the game and fixed once realized. In each period $t = 1$ or 2, the seller offers a rental price and the buyer decides to accept or reject the offer. Let μ denote the seller's ex ante belief that she is facing a high-value buyer.

In order to make a fair comparison, we require that the ex ante probabilities of the high-value buyer in both models be equal, that is, $\mu = \alpha q^G + (1-\alpha)q^F$, where α is the seller's ex ante belief of the G distribution in our model.

The following proposition compares the revenues in the equilibria of the two models for any ex ante belief the seller can have. When the seller has an optimistic ex ante belief, her revenue in our model is higher than that in Hart and Tirole's (1988). As shown in Proposition 6, the buyer types with $v_1 = h$ accept $p_1 = h$ with probability one since the seller offers $p_2 = h$ independent of whether p_1 is accepted or rejected. In contrast, in the two-period version of Hart and Tirole's (1988) rental model, the high-value buyer rejects $p_1 = h$ with a positive probability even if the seller is optimistic enough to offer equilibrium price $p_1 = h$, since otherwise the seller offers $p_2 = l$ after rejection of p_1 and the high-value buyer has an incentive to deviate to reject p_1.

When the seller has a pessimistic or moderately pessimistic ex ante belief, there always exists a pooling equilibrium in our model where the seller offers $p_1 = p_2 = l$ and all buyer types accept the offers. This equilibrium yields the seller the same revenue as in Hart and Tirole (1988).

When the seller has a moderately optimistic ex ante belief, she can still be better off than in Hart and Tirole (1988) if q^F is small enough, q^G is big enough, and the seller's ex ante belief is sufficiently optimistic. However, if the seller's ex ante belief is close to the lower bound of moderately optimistic beliefs, γ^*, then the seller is worse off than in Hart and Tirole (1988).

Proposition 12 (Revenue Comparison). *If the ex ante probability of high-value buyer in* *the two-period version of Hart and Tirole's (1988) rental model is the same as in this model, then*

(i) For an optimistic seller, the seller's revenue is higher than in Hart and Tirole's (1988);

(ii) For a moderately optimistic seller, if q^F is small enough and q^G is big enough, there exists $\overline{\alpha} \in (\gamma^*, \widehat{\alpha})$ such that, for all $\alpha \in (\overline{\alpha}, \widehat{\alpha})$, the seller's revenue is higher than in Hart and Tirole's (1988);

(iii) For a pessimistic and moderately pessimistic seller, there always exists an equilibrium in this model which yields the same revenue as in Hart and Tirole's (1988).

From Proposition 12 we conclude that, when the seller has sufficiently optimistic ex ante beliefs, the seller is better off compared to the case that the distribution of the buyer's value is common knowledge.

2.7 CONCLUSION

In this paper we have considered a two-period repeated bargaining model where the seller offers a price to rent a durable good in each period. The buyer's value of consuming the durable good is drawn from a distribution in each period. The buyer has private information not only about his value, but also about the distribution which his values are drawn from.

We compare the seller's revenue in our model with that in the two-period version of Hart and Tirole's (1988) rental model where the distribution of the buyer's value is common knowledge, under the assumption that the ex ante probabilities of high-value buyer are the same in the two models. We find that the seller is better off with the additional layer of uncertainty about the buyer's value distribution, when she has sufficiently optimistic ex ante beliefs.

The results we found may cast some light on the longer horizon. In the current two-period model, the seller cannot perfectly learn the buyer's value distribution. It is interesting to examine whether the seller is able to learn the buyer's distribution eventually if she is allowed to employ price experimentation in a finite or an infinite horizon.

On the other hand, this model only allows the seller to rent the durable good. For future research, we are interested in investigating the case where the seller is able to adopt a more general strategy, such as selling the durable good or providing both options to the buyer of selling and renting the durable good.

3.0 SOCIAL NORMS, INFORMATION, TRUST AMONG STRANGERS: AN EXPERIMENTAL STUDY

3.1 INTRODUCTION

Trust is a key element in specialization and trade. Yet an understanding of how trust emerges among essentially anonymous agents who have little recourse to punishment – as is typically the case in many economic transactions – has been slow to come. In this paper we examine two mechanisms by which trust and the reciprocation of trust might be sustained in a population of anonymous strangers. We first examine the hypothesis that trust might be attached to a society as a whole; the fear of the destruction of that trust might suffice to enforce trustworthy behavior by all members of the society as shown by Kandori (1992). On the other hand, such a mechanism might be too fragile and so we examine the second possibility that trustworthiness resides at the individual rather than the societal level. In particular, we ask whether the provision of information on individual reputations for trustworthiness engenders greater trust than in the case where such information is absent. We further explore whether the free provision of reputational information is responsible for our findings or whether the availability of acquiring such information (at a small cost) suffices to sustain greater trust and reciprocity.

To explore these issues we conduct an experiment that makes use of the two-player sequential trust game (Berg et al., 1995). In this game, the first mover or "investor" decides whether to invest his endowment with the second mover, the trustee, resulting in an uncertain payoff or decides to keep his endowment. If the investor invests, the endowment is multiplied by a fixed factor > 1 and it falls to the trustee to decide whether to keep (abscond with) this amount or return some fraction of it to the investor, keeping the rest for himself. Subjects are asked to play this game for several indefinite sequences, each consisting of a number of rounds. In each round, they are randomly and anonymously matched with one another. We examine several different treatments. In our baseline treatment (and in fact, in all of our treatments), the trust game is parameterized in such a way that, given the number of subjects we have and random anonymous matching, a social norm where all investors invest (trust) and all trustees return part of the investment (reciprocate) constitutes a sequential equilibrium. In a second treatment, everything is the same as in the baseline treatment except that, prior to making a decision, the investor can observe the trustee's action choice in the prior round (keep or return). In a third treatment, everything is the same as in the second treatment except that, prior to making a decision, the investor can observe a longer history of the trustee's most recent previous choices (up to 10 rounds) in all prior rounds of the current supergame. Finally, in a fourth treatment condition, everything is the same as in the third treatment, except that the investor must first choose whether to pay a small cost to view the trustee's history of actions for the current supergame. If the investor does not pay, then the game is identical to our first, baseline treatment where the investor has no knowledge of the prior actions of the trustee with whom he is matched.

In the first treatment, where no individual information is available, we are able to test whether a social norm of trust and reciprocity can be sustained by anonymous agents out of the fear that deviating from such a norm would precipitate a contagious wave of distrust and retaliatory non-reciprocation. We find that there is very little trust and reciprocity in this baseline treatment. Our second treatment asks whether "minimal" reputational information at the individual level can improve matters, specifically whether additional information on the prior-round behavior of trustees (second-movers) causes these players to reciprocate (return) more often and if so, whether this change in trustees' behavior engenders greater trust on the part of investors who move first. We find that, when minimal information on the trustee's prior-round choice is provided following the absence of such a reputational mechanism (treatment 1 to treatment 2), it leads to a large and significant increase in both trust and reciprocity. However, reversing the order, when minimal information about trustees is initially provided and then removed (treatment 2 to treatment 1) we find no significant difference in the level of trust and reciprocity. When the amount of information about trustees is increased (in our third treatment) to include up to 10 most recent rounds of trustee's actions in the current supergame, we find that such order effects disappear: the provision of a longer history leads to significant increases in trust and reciprocity relative to the absence of such information. Finally, in our fourth treatment, where investors must decide whether to purchase this longer history concerning their current matched trustee, we find that on average, only one-fourth of investors choose to purchase this information, so that the other three-fourths are in the dark about the prior behavior of their current trustee. Nevertheless, trust and reciprocity is significantly higher in this costly information treatment as compared with the baseline no-information treatment.

We conclude that the emergence of trust and reciprocity resides with the availability of information at the individual level as provided, for example, by a credit bureau and not through society-wide enforcement of a social norm of good behavior. We further conclude that longer histories are more beneficial than shorter histories in the promulgation of reputational concerns.

3.2 RELATED LITERATURE

We are not the first to explore the mechanisms supporting trust and reciprocity among anonymous strangers. Our research draws upon several prior theoretical and experimental studies.

3.2.1 Societal Cooperation under Random Matching in the Infinitely Repeated Prisoner's Dilemma Game

With anonymous random matching, it is impossible to maintain cooperation in a repeated game simply by punishing players who deviate. Kandori (1992) shows that cooperation may be possible if players employ a "contagious strategy" in which individuals who have not experienced a defection choose "Cooperation," and individuals who have either experienced a defection or has defected in the past chooses "Defection." Specifically, he models the infinitely repeated Prisoner's Dilemma with anonymous random matching and shows that, for any fixed population size, we can define payoffs for the Prisoner's Dilemma that sustain cooperation in a sequential equilibrium.

As pointed out by Kandori (1992), there are two substantial problems associated with a "contagious equilibrium." First, when the population is large, the argument applies only to games with extreme payoff structures. Second, a single defection causes a permanent end to cooperation and comments that this fragility may make the equilibrium inappropriate as a model for trade.

Ellison (1994) extends Kandori's work and remedies these problems by introducing a public randomization device which adjusts the severity of the punishment. Compared to Kandori's (1992) results, the equilibrium in Ellison (1994) does not require excessive patience of players and applies to more general payoff structures. Furthermore, given public randomizations, the equilibrium strategy supports nearly efficient outcomes even when players make mistakes with a small probability.

Duffy and Ochs (2007) conduct an experimental test of Kandori's (1992) contagious equilibrium using groups of subjects who play an indefinitely repeated two-person Prisoner's Dilemma under different matching protocols and different amounts of information transmission. Their results show that, under fixed pairings there appears to develop a social norm of cooperation as subjects gain experience, while under random matching, experience tends to drive groups toward a far more competitive norm, even when some information is provided about the prior choices of opponents. Thus they conclude that random matching works to prevent the development of a cooperative norm in the laboratory. By contrast, in this study we examine the indefinitely repeated, sequential move "trust" game and only consider the case of random matching. Contrary to Duffy and Ochs, we find that information on the prior behavior of others plays an important role in the trust game with random matching.

3.2.2 Social Norms under Random Matching in the Trust Game

Unlike the prisoner's dilemma game, the trust game (Berg et al., 1995) has 1) sequential moves and 2) no dominant strategy. Also, the trust game (unlike the Prisoner's Dilemma game), is more closely related to many real-life one-sided incentive problems that are found in credit markets, or in transactions between buyers and sellers in cyberspace (e-Commerce), and other trading situations (including Greif's, 1989, remarkable analysis of medieval trade).[1] Citing these reasons, Lee and Xie (2007) theoretically extend Kandori's (1992) argument to the development of trust and reciprocity among anonymous, randomly matched players in the infinitely repeated trust game. In particular, Lee and Xie (2007) provide sufficient conditions on trust game parameters that support a social norm of trust and reciprocity as a sequential equilibrium among random and anonymously matched players in the absence of information about other players. The trust game experiment we report on in this paper satisfies the Lee and Xie conditions in all treatments, so that in the absence of any information about one' randomly determined opponents, a social norm of trust and reciprocity may be sustained by the threat to move to a contagious wave of distrust and confiscation. However, we also explore the notion that some information about opponents' prior behavior may help to sustain social norms of trust and reciprocity, as such information makes it easier for players to discern player types thus enabling reputational considerations.

There is some experimental literature on repeated trust games that relates to this study. Bolton et al. (2004) report an experiment that evaluates the effectiveness of electronic reputation mechanisms. A trust game with binary choices (buyer-seller game) is played repeatedly for 30 rounds in each session. They compare the results from three treatments: a stranger market, where individual buyers and sellers meet no more than once and the buyer has no information about the seller's transaction history; a feedback market, which has the same matching rule as the stranger market and provides the seller's histories of shipping decisions to the buyer; finally, a partners market, where the same buyer-seller pairs interact repeatedly in every round. Not surprisingly, transaction efficiency, trust and trustworthiness (reciprocity) are smallest in the stranger market, greater in the feedback market, and greatest in the partners market. Their results imply that the information

[1] Kandori (1992) has a formal definition for one-sided incentive problem (Definition 4 on page 73). The concept requires that, only one party of two players has an incentive to deviate from the cooperative outcome, and there is a Nash equilibrium such that the payoff from the equilibrium is less than the payoff from the cooperative outcome for the party who has the incentive problem. The trust game we use in this paper is a representative of one-sided incentive problems since, first Invest by the first mover is a best response to Return by the second mover and only the trustee has an incentive to deviate from the outcome Invest&Return, and second, the payoff for the trustee from the unique equilibrium is 0, which is less than the payoff from the cooperative outcome Invest&Return.

from other's feedback and the information from one's own experience have different effects on behavior. However, their environment is of a finite duration and is not one that can rationalize trust and trustworthiness as an equilibrium phenomenon (as is the case in our study).

Engle-Warnick and Slonim (2006) examine how exogenously determined lengths of past relationship affect trust and trustworthiness in new relationships. The participants of the experiment play several supergames. A supergame is a sequence of trust games played between the same two players. The lengths of supergames were drawn prior to the first session. The treatments focus on whether an initial sequence of short- or long- supergames impacts on the extent of trust and trustworthiness later. They find that initial short-supergame relationships have an immediate negative impact on both trust and trustworthiness in the relationships that immediately follow, while longer-lasting relationships have the opposite effect. In the long run, the effect declines for trustworthiness but not for trust as subjects gain experience.

The literature on economic institutions is also related to the random matching model (Greif, 1989; Milgrom et al., 1990). These papers model a large number of traders who are randomly paired with each other in each period. Each pair is presumed to play a game similar to the trust game, where one party has an incentive to cheat the other by supplying goods of inferior quality or reneging on promises to make future payments. In this literature, institutions are seen as a way of avoiding the inefficiency of noncooperative equilibria. Greif (1989) and Milgrom et al. (1990) argue that the exchange of information on the identity of cheaters or the development of a mechanism which strengthen the power of enforcement can help sustain cooperation.

None of these studies directly addresses the question we pose here: whether the mechanism that supports trust and reciprocity comes about through community enforcement (fear of a contagious wave of distrust and confiscation) or from provision of information on individual behavior (that affects the behavior of both the observed and those deciding whether to trust). For this reason, we designed the experiment we report on in this paper.[2]

[2]Many experimental studies find that trust and reciprocity prevail under the conditions of complete anonymity and one-shot interaction. As these behaviors are inconsistent with all participants being payoff maximizers, they are often explained by psychological factors such as fairness, altruism, and inequality

3.3 THE MODEL

In this section we briefly describe the model and results of Lee and Xie (2007). The set of players $N = \{1, 2, \ldots, 2n\}$ is partitioned into two sets of equal size, the set of investors $N_I = \{1, 2, \ldots, n\}$ and the set of trustees $N_T = \{n+1, n+2, \ldots, 2n\}$. In each period, each investor is matched with a trustee according to the uniform random matching rule, and they play the trust game as a stage game. This procedure is repeated infinitely and each player's total payoff is the expected sum of his stage payoffs discounted by $\delta \in (0, 1)$.

The stage game is as follows. At the beginning of the game, the investor is endowed with one unit of capital and decides whether to invest the capital in the trustee's business or not. If the investor decides not to invest, the game ends and she gains $a < 1$ from the outside option, and the trustee gets nothing. If the investor chooses to invest, the capital grows into 1, which is the return from the trustee's business. Then the trustee decides whether to keep all gains or to return the amount b to the investor. If the investor chooses to invest and the trustee chooses to return, the payoff is b for the investor and $1 - b$ for the trustee. If the investor chooses to invest but the trustee chooses to keep, then the investor gets nothing and the trustee gets 1. We assume $0 < a < b < 1$. The trust game and its payoff structure is described in Figure 3.1.

If the game is played once, the unique subgame perfect equilibrium is for the investor not to invest and for the trustee to keep all gains. Since the return from capital in the trustee's business is bigger than the outside option a of the investor, it is efficient for the investor to invest. Although the efficient outcome can not be achieved in the one-shot trust game, we will show below that it can be achieved in the "contagious equilibrium" when the trust game is infinitely repeated, even if the opponents are randomly rematched after each period.

We define No Invest as the defection of an investor and Keep as the defection of a trustee. Define *d-type* investors or trustees as those whose history include defection of themselves or their partner, otherwise the players are *c-type*.

Definition 2. *The "contagious strategy" is defined as follows: An investor invests if she is*

aversion etc. (Berg et al. (1995), Bolton and Ockenfels (2000), Fehr and Schmidt (1999), see Camerer (2003) for a survey).

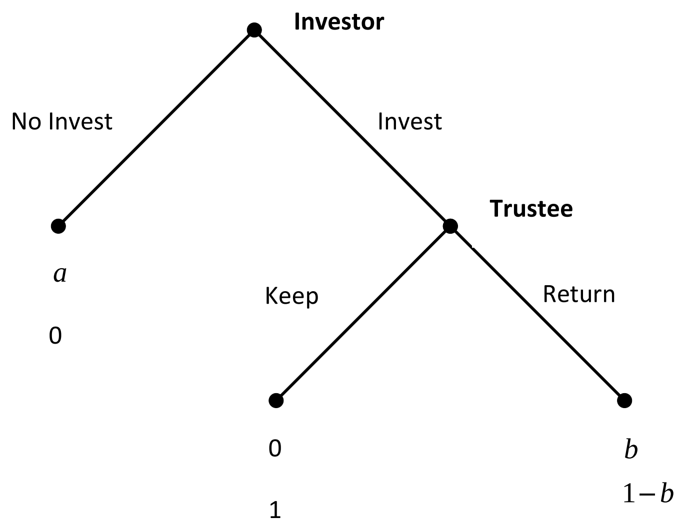

Figure 3.1: The Trust Game

c-type and does not invest if she is d-type. A trustee returns if he is c-type and keeps if he is d-type.

The idea of the contagious strategy is that trust is applied to the community as a whole but cannot be applied to each individual player due to anonymity. Therefore, a single defection by a member means the end of trust in the whole community, and a player who experiences dishonest behavior starts defecting all of his or her opponents (Kandori, 1992). Now we show below that we can define payoffs for the trust game which allow trust and reciprocity in a sequential equilibrium for any finite population.

Theorem 13. *Consider the random matching model described above where $2n \geq 4$ players play a trust game. Then for any δ and n, there exist a and b such that (i) $0 < a < b < 1$; and (ii) the contagious strategy constitutes a sequential equilibrium in which (Invest, Return) is the outcome in every period along the equilibrium path under uniformly random matching.*

From the literature of repeated games, it is sufficient to show that one-shot deviations are not profitable after any history in order to show that the contagious strategy constitutes a sequential equilibrium. In particular, Lee and Xie (2007) show that the conditions in the following lemma control the investors and the trustees' incentive of one-shot deviations from the contagious strategy both on-the-equilibrium-path and off-the-equilibrium-path, and given any δ and n they can always find a proper payoff profile a and b such that $0 < a < b < 1$ for the binary trust game to satisfy these conditions.

Before moving onto the lemma, we first introduce $f(\delta)$ and $g(\delta)$ as functions of δ. Conceptually, $f(\delta)$ is the gain for a trustee from deterring starting a defection when all the other players in the community are c-type, and $g(\delta)$ is the gain for the d-type trustee from deviating from defection (i.e., continuing to return) given that there are one d-type investor and one d-type trustee (himself) in the current period. So $f(\delta)$ and $g(\delta)$ are the payoffs in the future periods for a d-type trustee from slowing down the contagious procedure in the current period at different states of the world (i.e., when there are different numbers of d-type investors and d-type trustees currently in the community).

To provide a formal expression of $f(\delta)$ and $g(\delta)$, more notations are necessary. Let X_t be the total number of d-type investors and Y_t be the total number of d-type trustees at the beginning of period t. The state of the world in period t, Z_t, contains information about the number of d-type investors and d-type trustees in the current period and is defined as a one-to-one and onto function from (X_t, Y_t) to the set of natural numbers $\{1, 2, ..., n(n+2)\}$:

$$Z_t = (n+1)X_t + Y_t \quad \text{for} \quad X_t + Y_t > 0.$$

Let A be an $n(n+2) \times n(n+2)$ transition matrix when all players follow the contagious strategy. It has elements

$$a_{ij} = \Pr\{Z_{t+1} = j \mid Z_t = i\}.$$

For example, $a_{12} = \Pr\{Z_{t+1} = 2 \mid Z_t = 1\} = \Pr\{(X_{t+1}, Y_{t+1}) = (0, 2) \mid (X_t, Y_t) = (0, 1)\}$ denotes the probability that there are two d-type trustees and no d-type investors in next period given one d-type trustee and no d-type investors in the current period. Similarly, let B be an $n(n+2) \times n(n+2)$ transition matrix when the d-type trustee in consideration deviates from the contagious strategy while all other players still follow the strategy, with elements $b_{ij} = \Pr\{Z_{t+1} = j \mid Z_t = i \text{ and one } d\text{-type trustee deviates to Return}\}$. So matrix $B - A$ characterizes how the diffusion of d-type players is delayed if one d-type trustee unilaterally deviates from the contagious strategy. Define ρ as an $n(n+1) \times 1$ column vector with the ith element equal to the conditional probability for the trustee to meet a c-type investor when the state is i in period t. Finally, let e_i be a $1 \times n(n+2)$ row vector with the ith element equal to 1 and all other elements equal to 0. So $f(\delta) = \delta e_1 (B - A)(I - \delta A)^{-1} \rho$ is the increase in the sum of the expected probability to meet a c-type investor (and also the increase in the sum of the expected payoff) for the d-type trustee in all the future periods when this d-type trustee chooses to deviate from defection (i.e. to return) given that the d-type trustee is the only d-type player. Similarly, $g(\delta) = \delta e_{n+2}(B - A)(I - \delta A)^{-1} \rho$ is the increase in the sum of the expected payoff (and the probability to meet a c-type investor) in future periods for the d-type trustee from slowing down the contagious procedure given currently there are one d-type trustee and one d-type investor.[3]

Lemma 14. *The contagious strategy constitutes a sequential equilibrium if*

$$a \geq \frac{n-1}{n} b, \tag{3.1}$$

and

$$g(\delta) \leq b \leq f(\delta). \tag{3.2}$$

Condition (3.1) in the lemma means that the loss of deviation from the contagious strategy off-the-equilibrium-path for an investor (i.e., deviation from defection, that is, continuing to invest for a d-type investor) is greater than the benefit of deviation, even if there is only one d-type trustee. The left hand side of inequality (3.1), a, is the investor's opportunity cost to invest, and the right hand side of inequality (3.1) is the expected payoff of return given there is only one d-type trustee. So condition (3.1) controls the investor's incentive to deviate from the contagious strategy off-the-equilibrium-path. Due to the nature of one-sided incentive problem, Invest is the best response to Return, so the investor has no incentive to deviate on-the-equilibrium-path neither. Therefore, the investor has no incentive to deviate from the contagious strategy given any history.

The implication from condition (3.1) is that the existence of the contagious equilibrium requires high outside option. For the development of a cooperative social norm, the concept of contagious equilibrium requires harsh punishment scheme. Not only are deviators from the desired behavior punished, but a player who fails to punish is in turn punished (Kandori,

[3]The formal derivation of $f(\delta)$ and $g(\delta)$, as well as the formula for each element of matrix A and B can be found in Lee and Xie (2007).

1992). So an investor must defect forever once she is cheated before. To control the d-type investor's incentive from investing again off-the-equilibrium-path, the outside option a must be high enough.

Condition (3.2) controls the trustee's incentive to deviate from the contagious strategy both on-the-equilibrium-path and off-the-equilibrium-path. Notice that b is the trustee's payoff of defection and the loss of not to defect at the same time. The first part of condition (3.2), $f(\delta) \geq b$, means that the trustee's gain from defection, i.e., b, must be less than the gain from deterring starting a defection, i.e., $f(\delta)$. So the trustee (also the only d-type player) will not start a defection in the current period. The same condition guarantees that the trustee will not start a defection in all the later periods. So it is sufficient to control the trustee's incentive of deviation on-the-equilibrium-path. The second part of condition (3.2), $g(\delta) \leq b$, implies that the loss of slowing down the contagious procedure, i.e., b, must be greater than the gain of slowing down the procedure when there is already defection in the community. So it controls the trustee's incentive of deviation off-the-equilibrium-path. Finally, to show that there always exists b between $g(\delta)$ and $f(\delta)$, Lee and Xie (2007) show that $g(\delta)$ is less than $f(\delta)$ for any δ greater than 0 given any finite population size. The intuition is that the trustee's payoff from slowing down the contagious procedure is greater when the trustee in consideration is the only d-type player in the community, in which case the contagious procedure stops completely for the current period if he chooses not to defect, than when there are other d-type players.

3.4 EXPERIMENTAL DESIGN

The main treatment variable in this paper concerns information available to the investor in advance of his/her decision. We investigated four different informational mechanisms. In the "no information" treatment (denoted as No), investors only know their own history of play and payoff in each round. In the "minimal information" treatment (denoted as Min), investors are informed of the prior-round decision of their current paired trustee, i.e., whether they chose Keep or Return in the event that the trustee had the opportunity to make a choice in the prior round; if the trustee did not have an opportunity to make a decision in the prior round, the information reported to the investor is "no choice." In the "information" treatment (denoted as Info), investors are shown their current paired trustee's most recent previous decisions up to 10 rounds as well as the total numbers that the trustee chose Keep and Return out of the total number that the trustee had the opportunity to make a choice in the history. Finally, in the "costly information" treatment (denoted as Cost), the system does not automatically provide investors information on their paired trustee's previous choices; in stead, the investors can choose to purchase the same information as in the "information" treatment at a small cost.

The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). All of our experimental sessions involve groups of 6 subjects, who were randomly assigned the role of Investor or Trustee at the start of each indefinitely repeated supergame.[4] This design gives subjects experience with both roles.

An indefinitely repeated supergame was implemented as follows. The 6 subjects were randomly matched and played the stage game. At the end of play of the stage game, a 10-sided die was rolled. If the die came up 8 or 9, the supergame was declared over; otherwise the game continued with another round. Subjects were randomly rematched before playing the next round, though they remained in the same role in all rounds of the supergame. We told them that we would play a number of sequences (i.e., indefinitely repeated supergames) but did not specify how many. For transparency and credibility purposes, we had the subjects take turns rolling the 10-sided die themselves and calling out the result. Our design thus implements random and anonymous matching, a discount factor $\delta = 0.8$, and the stationarity associated with an infinite horizon. All the informational mechanisms discussed above apply to each indefinitely repeated supergame. That is, when a new supergame begins, information on the trustee' behavior in previous supergames does not carry over to the new supergame. In the treatments where information is available, it is available from the start of the second round of each supergame.

The parameterization of the stage game used in all experimental sessions is given in Figure 3.2. This parameterization of the game was chosen so that conditions (3.1) and (3.2)

[4] In the instructions, we use neutral word "First Mover" for investor, "Second Mover" for trustee, and "sequence" for indefinitely repeated supergame. We also use "A" "B" "C" "D" to denote the investor and trustee's choices. See Appendix C for instructions.

are satisfied given our values $n = 3$ and $\delta = 0.8$. (The program for details on verification of condition (3.2) is available upon request). While other parameterizations are possible, we chose a parameterization that is not at the boundary of the conditions (3.1)-(3.2), but are well within the region supporting trust and reciprocity among randomly matched players. The cost to purchase information in the "costly information" treatment is 2 points.

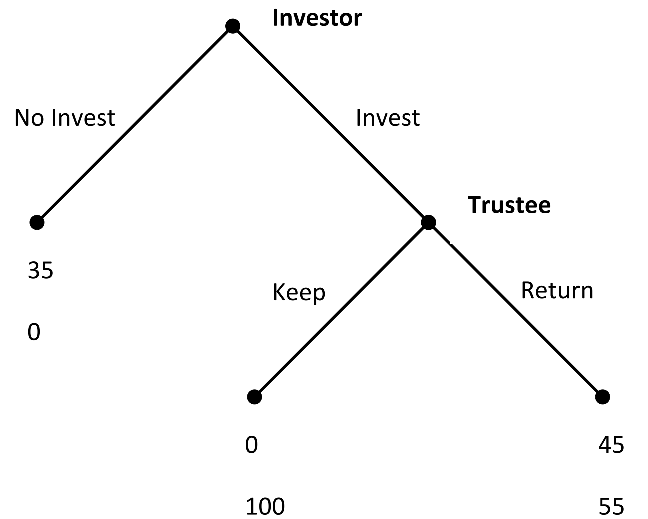

Figure 3.2: Parameterization of the Stage Game

We use a within-subjects design in all the sessions. Subjects begin to play under one information condition and switch to the second information condition (and then to the third condition in some sessions). Subjects are only informed of the change of information condition when the switch took place.

The first set of 8 experimental sessions examines the effect of "minimal information" compared to "no information." We conducted 4 No_Min sessions (sessions which began with "no information" and switched to "minimal information", similar notations afterward) and 4 Min_No sessions. Another set of 6 sessions investigates the effect of the longer history (up to 10 most recent rounds) of the trustee's behavior and the possibility to purchase information at a cost. We conducted 3 No_Info_Cost sessions and 3 Info_No_Cost sessions. We reverse the order of the first two treatments to examine whether the timing when information is available matters. The treatment of "costly information" is always the last treatment since we want to have subject experience the information for free before they decide whether to purchase it at a cost.

In practice, there are at least 30 rounds under each information condition in all sessions. When the total number of rounds under one information condition is over 35 rounds, the sequence is the last sequence under that condition. In some cases the total number of rounds under one condition is between 30 and 35.

The motivations for this experimental design are as follows. First, theory allows a social norm of trust and reciprocity among randomly matched anonymous players to emerge under the chosen parameterization even when no information is available. However, we cannot exclude other equilibria, e.g., the social norm of no trust and no reciprocity. So it is of our interest to test, in an empirical approach, first, whether the community-wide enforcement is enough to support the social norm of trust and reciprocity, and second, whether different informational mechanisms can help select different social norms.

Second, since the collection and dissemination of information is always costly to the society, a question of practical interest is how much reputational information is enough in order to significantly enhance the frequencies of trust and reciprocity. So we are not only interested in comparing the difference when information is available or not, but also the difference when different amount of information is available. That is our motivation to have both the "minimal information" and "information" treatments. Notice that by our design the information shown in the "minimal information" treatment is nested in the "information" treatment.

Finally, the "costly information" treatment addresses the question of whether the content of information or the availability of information matters if reputational information is found to improve the frequencies of trust and reciprocity. The trustees are not informed of whether their paired investor purchases information or not, and this fact is made public information among all the subjects. On one hand, when a portion of the investors purchase information, there is a positive externality among the whole community due to the anonymously random matching. On the other hand, if trustees believe that most of the investors do not purchase information, they may behave similarly as in the "no information" treatment. So the answer is not clear whether the availability of information is sufficient to significantly enhance trust and reciprocity.

3.5 RESULTS

All the subjects were recruited from the undergraduate population at the University of Pittsburgh and Carnegie Mellon University. No subject had any prior experience participating in our experiment. Subjects were given $5 for showing up on time and they received their earnings from all rounds of all sequences played. Subjects accumulated points given their stage game choices (points are shown in Figure 3.2). Points were converted into dollars at the rate of 1 point = 0.5 cent.

Table 3.1 below provides the basic characterization of all sessions. Subjects earned on average, $17.07 ($13.36 for the first and $22.02 for the second set of sessions) in addition to their $5 show-up fee. All the sessions in the first set are finished within 1.5 hours and the sessions in the second set are finished within 2 hours.

In the following subsections, we first report the results from the first set and the second set of sessions respectively, and then analyze how different informational mechanisms affect investors' behavior.

3.5.1 No_Min and Min_No Sessions

We first analyze whether the provision of the minimal information, the information on the prior-round behavior of the trustee, can significantly increase the frequencies of trust and reciprocity. We start with the within-subject comparison and then move to the between-subject comparison.

Tables 3.2, 3.3, 3.4 report respectively on the aggregate frequencies of 1) invest, 2) return conditional on investment, and 3) combined frequencies of invest&return for the No_Min and Min_No sessions.

The first important finding is that the social norm of full trust and reciprocity is not sustained in the absence of information, however, the frequencies of trust and reciprocity are significantly different from zero (one-tailed binomial test, $p = 1\%$).

Table 3.1: All Experimental Sessions

Session	No. of Supergames	No. of Rounds	Avg. Payoff	Avg. Payoff per Round
No_Min1	16	84	$12.12	$0.14
No_Min2	13	71	$14.39	$0.20
No_Min3	10	72	$14.80	$0.21
No_Min4	17	74	$13.63	$0.18
Min_No1	17	81	$17.11	$0.21
Min_No2	16	79	$12.55	$0.16
Min_No3	16	72	$12.75	$0.18
Min_No4	18	79	$9.51	$0.12
No_Info_Cost1	22	119	$25.65	$0.22
No_Info_Cost2	24	106	$22.98	$0.22
No_Info_Cost3	23	97	$20.75	$0.21
Info_No_Cost1	21	119	$25.90	$0.22
Info_No_Cost2	20	113	$17.49	$0.15
Info_No_Cost3	24	108	$19.36	$0.18
Average	18.36	91	$17.07	$0.19

Table 3.2: Frequency of Invest for No_Min and Min_No Sessions

	1st treatment	2nd treatment
No_Min1	0.333	0.362
No_Min2	0.618	0.793
No_Min3	0.476	0.964
No_Min4	0.521	0.676
Avg. of No_Min	0.487	0.699
Min_No1	0.674	0.865
Min_No2	0.561	0.325
Min_No3	0.333	0.746
Min_No4	0.301	0.096
Avg. of Min_No	0.467	0.508

Table 3.3: Frequency of Return-given-Invest for No_Min and Min_No Sessions

	1st treatment	2nd treatment
No_Min1	0.526	0.760
No_Min2	0.587	0.841
No_Min3	0.780	0.972
No_Min4	0.525	0.831
Avg. of No_Min	0.605	0.851
Min_No1	0.820	0.854
Min_No2	0.797	0.725
Min_No3	0.588	0.706
Min_No4	0.649	0.364
Avg. of Min_No	0.714	0.662

Table 3.4: Frequency of Invest&Return for No_Min and Min_No Sessions

	1st treatment	2nd treatment
No_Min1	0.175	0.275
No_Min2	0.363	0.667
No_Min3	0.371	0.937
No_Min4	0.274	0.562
Avg. of No_Min	0.296	0.610
Min_No1	0.553	0.739
Min_No2	0.447	0.236
Min_No3	0.196	0.526
Min_No4	0.195	0.035
Avg. of Min_No	0.348	0.384

Second, there is an order effect when the minimal information is provided in the first half and the second half of the sessions. For the No_Min sessions, the provision of minimal information about the trustee's prior-round play in the second half of the sessions leads to significantly larger frequencies of invest, return, and invest&return (Wilcoxon signed ranks test, $p = 0.0625$ for all three tests). However, none of the frequencies of invest, return-given-invest, and invest&return is significantly different when the minimal information is provided in the first half of the Min_No sessions (Wilcoxon signed ranks test, $p > 0.4$ for all three tests).

There may be two possible explanations for the finding that the minimal information has no significant effect in the Min_No sessions. One explanation is that, when the minimal information is provided in the first half of the sessions, it helps form the social norm of trust and reciprocity, and the social norm does not fall apart even when the information is removed in the second half of the sessions, i.e., the first half and the second half of the sessions have equally good outcomes. This explanation implies that providing information in the first place yields more efficient outcomes.

The second explanation is that, when the information is provided in the first half of the session, the subjects do not recognize the use of information very well, and they are slower in learning how to use it, i.e., the first half and the second half of the sessions have equally bad outcomes. This explanation has the different policy implication: it is more efficient to provide the informational mechanism to agents who have experienced the absence of the mechanism.

These two explanations also have different predictions. The first explanation predicts that the frequencies of trust and reciprocity of the "no information" treatment should be higher in the Min_No sessions than in the No_Min sessions, and the frequencies of trust and reciprocity of the "minimal information" treatment be indifferent in the Min_No sessions and the No_Min sessions. Differently, the second explanation predicts that the frequencies of trust and reciprocity of the "no information" treatment is indifferent in the Min_No sessions and the No_Min sessions, and the frequencies of trust and reciprocity of the "minimal information" treatment is higher in the No_Min sessions than in the Min_No sessions.

Next we move on to examine the between-subject tests for all the Min_No and No_Min sessions. In particular, we compare 1) the first treatment of all these 8 sessions (i.e., "no information" in the No_Min sessions and "minimal information" in the Min_No sessions); 2) the "no information" treatment of these 8 sessions; and 3) the "minimal information" treatment of these 8 sessions.

Since the subjects got to know the switch of treatment when it took place in the middle of the sessions, we can regard the first treatment of the sessions as independent observations which have not been influenced by other information treatments. We find that the frequencies of invest and invest&return in the first treatment are not significantly different when minimal information is provided or not. However, the frequency of return-given-invest is significantly higher when minimal information is available (Robust Rank Order test, $p = 0.05$). This implies that the provision of information has a more significant effect on trustees than investor.

For the between-subject tests on the "no information" treatment, we find that none of the frequencies of invest, return-given-invest, and invest&return is significantly different when the treatment is in the first half of the No_Min sessions or in the second half of the Min_No sessions. However, for the between-subject tests on the "minimal information" treatment, all the frequencies of invest, return-given-invest, and invest&return are significantly higher when the minimal information is provided in the second half of the No_Min sessions than in the first half of the Min_No sessions (Robust Rank Order test, $p = 0.05$ for all three tests).

The findings on the between-subjects tests support the second explanation, that is, when the minimal information is provided after subjects suffer from the absence of reputational information, subjects learn to use the minimal information better.

3.5.2 No_Info_Cost and Info_No_Cost Sessions

In this subsection we report the results from the second set of sessions. Tables 3.5, 3.6, 3.7 report the frequencies of invest, return-given-invest, invest and return respectively. Similar to the last subsection, we start with the within-subject analysis and then move to the between-subject comparison.

Table 3.5: Frequency of Invest for No_Info_Cost and Info_No_Cost Sessions

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	0.516	0.892	0.983
No_Info_Cost2	0.714	0.857	0.824
No_Info_Cost3	0.667	0.813	0.871
Avg. of No_Info_Cost	0.632	0.854	0.893
Info_No_Cost1	0.848	0.659	0.927
Info_No_Cost2	0.829	0.132	0.352
Info_No_Cost3	0.775	0.515	0.471
Avg. of Info_No_Cost	0.817	0.435	0.583

First of all, similar to the No_Min and Min_No sessions, the frequencies of trust and reciprocity in the "no information" treatment are significantly different from zero (one-tailed binomial test, $p = 5\%$), but the social norm of full trust and reciprocity is not supported.

Compared to the "minimal information" treatment, when the amount of information provided to investors becomes larger, the effect on increasing the frequencies of trust and

Table 3.6: Frequency of Return-given-Invest for No_Info_Cost and Info_No_Cost Sessions

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	0.800	0.970	0.992
No_Info_Cost2	0.680	0.867	0.775
No_Info_Cost3	0.603	0.923	0.889
Avg. of No_Info_Cost	0.694	0.920	0.885
Info_No_Cost1	0.982	0.741	0.989
Info_No_Cost2	0.931	0.471	0.838
Info_No_Cost3	0.911	0.471	0.625
Avg. of Info_No_Cost	0.941	0.561	0.817

Table 3.7: Frequency of Invest&Return for No_Info_Cost and Info_No_Cost Sessions

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	0.413	0.865	0.975
No_Info_Cost2	0.486	0.743	0.639
No_Info_Cost3	0.402	0.750	0.774
Avg. of No_Info_Cost	0.434	0.786	0.796
Info_No_Cost1	0.833	0.488	0.917
Info_No_Cost2	0.771	0.062	0.295
Info_No_Cost3	0.706	0.242	0.294
Avg. of Info_No_Cost	0.770	0.264	0.502

reciprocity becomes more significant. We find that the provision of the longer history on the trustee's previous behavior significantly increases the frequencies of invest, return, and invest&return, regardless of whether the "information" treatment is used after or before the "no information" treatment. (one-tailed Wilcoxon signed ranks test, $p = 0.05$ for all three tests, using six sessions as observations)

Even when information is not provided automatically, instead, investors are provided with the possibility to purchase information at a cost of 2 points, the frequencies of trust and reciprocity are still significantly increased compared to when no information is available. (one-tailed Wilcoxon signed ranks test, $p = 0.05$ for all three tests, using six sessions as observations)

The difference between the "information" treatment and the "costly information" treatment, however, is not very significant. We find the frequencies of invest and invest&return are not significantly different in the "information" treatment and "costly information" treatment. The frequency of return is significantly higher in the "information" treatment than in the "costly information" treatment. (one-tailed Wilcoxon signed ranks test, $p = 0.1$, using six sessions as observations) This evidence implies that trustees' behavior is more sensitive to the change in the treatment than investors' behavior, which is consistent with the findings in the No_Min and Min_No sessions.

By our design, there is no information available in the first round of a sequence in any treatment, so an alternative way to examine the treatment effect is to compare the frequencies of trust and reciprocity for the rounds excluding the first rounds. We find that the same results hold when the first rounds are excluded. The frequencies of invest, return-given-invest, and invest&return excluding the first rounds are in Appendix D.

Although the frequencies of trust and reciprocity are significantly increased when information is provided or available to purchase, one may be concerned whether the players' payoff or the efficiency of the society is also consistently increased with the same information mechanism. The following table presents the average payoff per round under each treatment. We find that, similar to the frequencies of trust and reciprocity, players' average payoff is significantly increased when the information is provided automatically or available to purchase compared to the case when the information is absent (one-tailed Wilcoxon signed ranks test,

$p = 0.05$ for both tests, using six sessions as observations).

Table 3.8: Players' Average Payoff per Round in Points

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	34	46	49
No_Info_Cost2	41	45	44
No_Info_Cost3	39	44	46
Avg. of No_Info_Cost	38	45	46
Info_No_Cost1	45	39	48
Info_No_Cost2	44	22	29
Info_No_Cost3	43	34	33
Avg. of Info_No_Cost	44	32	37

Now we move to the between-subject analysis for the second set of 6 sessions. First, we focus on the first treatment of all the 6 sessions. Different from the No_Min and Min_No sessions, we find that providing the larger amount of information significantly increases the frequencies of invest, return-given-invest, and invest&return (one-sided Robust Rank-Order test, $p = 0.05$). This evidence again shows that the "information" mechanism is more effective than the "minimal information" mechanism in enhancing trust and reciprocity. Second, we examine the third treatment of all the 6 sessions. We find that none of the frequencies of trust and reciprocity is significantly different. This finding implies that there is no significant order effect, that is, whether the subjects experience "no information" treatment or "information" treatment first and then switch to the other one does not significantly affect their behavior in the third treatment "costly information." Finally, we also check whether subjects' behavior of the "no information" treatment and "information" treatment is different when the same treatment is the first or the second treatment in the session. We find that there is no significant difference for most cases except that the frequency of invest in the "no information" treatment is slightly higher when "no information" is the first treatment than when it is the second treatment (Robust Rank-Order test, $p = 0.1$).

As a summary, the No_Info_Cost and Info_No_Cost sessions present several important findings. First, providing longer history on trustee's previous behavior has a larger effect in enhancing trust and reciprocity compared to just providing the trustee's prior-round choice. The effect of information is robust regardless of whether information is provided before or after the "no information" treatment, and also in both within-subjects and between-subjects sessions. Second, providing the possibility to purchase information at a small cost is sufficient to support the similar efficient outcome as in the case where information is provided automatically for free.

3.5.3 Use of Information

This subsection focuses on how information in different treatments affects investors' behavior. We first analyze the effect of minimal information, and then the larger amount of information and costly information.

3.5.3.1 Minimal Information Figure 3.3 below presents the frequency of invest conditional on the prior-round choice of the trustee (i.e., Keep or Return) in the "no information" treatment, the "minimal information" treatment of the No_Min sessions and the Min_No sessions respectively.

There are several interesting findings. First, the frequency of invest is higher when the trustee's prior-round choice is "Return" than "Keep" no matter whether or not the investors are informed of the trustee's prior-round behavior. In the "minimal information" treatment, this implies that investors respond in a correct direction with the information revealed. When no information is available, however, this evidence suggests that investors' own experience in early periods affects their belief over the reputation of the whole community, and consequently affects their trusting behavior in the current period.

Although investors' response conditional on trustee' prior-round choices is in the correct direction even when no information is shown to the investors, the minimal information still has its value. In particular, the frequency of invest conditional on "Keep" in the "no information" treatment is larger than in the "minimal information" treatment, furthermore, the frequency of invest conditional on "Return" in the "no information" treatment is smaller

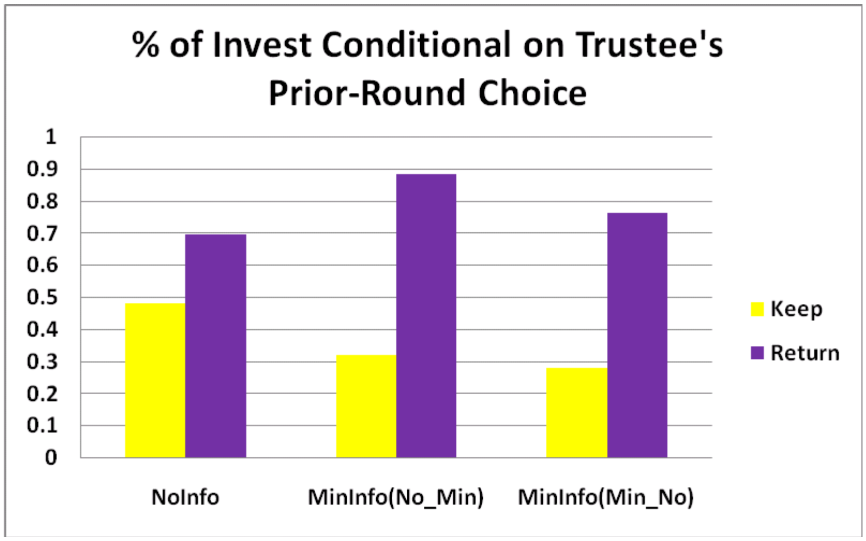

Figure 3.3: Frequency of Invest Conditional on Trustee's Prior-Round Choice

than in the "minimal information" treatment.

There is also an order effect in terms of whether the "minimal information" is the first or the second treatment of the sessions. The difference between the frequency of invest conditional on "Return" and "Keep" is larger when the "minimal information" is provided in the second half than in the first half of the sessions, which implies that the investors learn how to use information better after experiencing the absence of information. This finding is consistent with the aggregate data.

Finally, Figure 3.3 also suggests that investors' behavior does not completely reply on the information about trustees' previous behavior. Investors still invest with a positive probability even if the information shows that their current partners defected in the prior round, and they do not fully trust when the partner is revealed cooperative last period. There are several possibilities. One explanation is that there exist a small portion of investors who always cooperate (trust) or always defect (distrust). Another explanation is that subjects make mistakes or experiments to learn and subsequently adjust their behavior. It is also possible that, even in the presence of information, investors' behavior is still influenced by their own experience as suggested by the contagious hypothesis.

3.5.3.2 "Information" and "Costly Information" Treatment

Figure 3.4 shows the frequency of invest conditional on the trustee's aggregate frequency of return in the history of the current sequence in the "information" and "costly information" treatment respectively. The x axis is the frequency of return by the trustee (denoted as r), which is divided into three ranges ($0 <= r \leq 1/3$, $1/3 < r \leq 2/3$, $2/3 < r \leq 1$). For the "costly information" treatment, the frequency of invest is calculated when the information is purchased and shown to the investors.

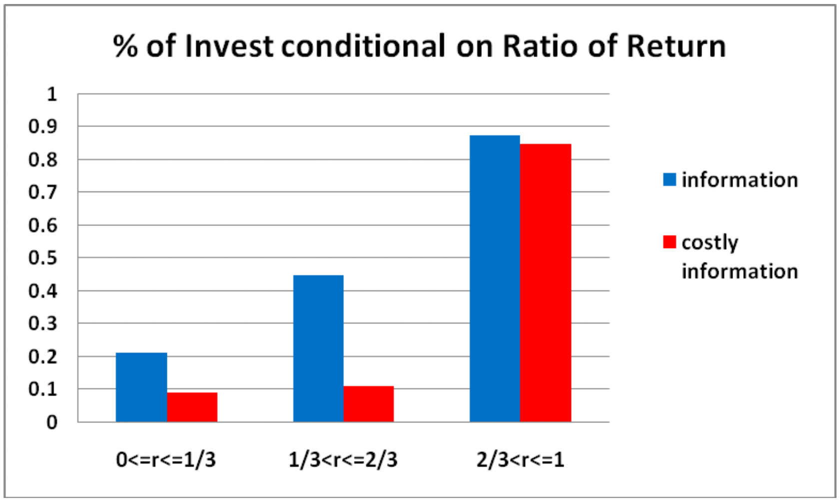

Figure 3.4: Frequency of Invest Conditional on Trustee's Aggregate Ratio of Return

From Figure 3.4, the frequency of invest is monotonically increasing in the ratio of return by the trustee in both "information" and "costly information" treatments. One strike difference between these two treatments, however, is that the frequency of invest conditional on a low ratio of return in the "information" treatment is much higher than in the "costly information" treatment. In particular, for the "information" treatment, the frequency of invest is more than 20% when the trustee's ratio of return is less than or equal to 1/3, and it is around 45% when the trustee's ratio of return is between 1/3 and 2/3. However, for the "costly information" treatment, the frequency of invest is less than 10% when the trustee's ratio of return is less than or equal to 1/3, and it is around 11% when the trustee's ratio of return is between 1/3 and 2/3.

For the "costly information" treatment, the average percentage of information purchase is 24.49%. Combined with the finding that the frequency of trust and reciprocity is not significantly different between the "information" and "costly information" treatment, it implies that the availability of information creates a large positive externality. The frequency of invest is 72.48% when information is not purchased, in comparison with 65.55% when information is purchased. (All analysis here does not include the first rounds, where information is not available to purchase.)

In the following table we take a closer look at which subjects bought information and how often they bought information. Each cell in Table 3.9 shows the ratio of the total number that each subject in each session purchased information over the total number that each subject was able to purchase information (i.e., the subject was assigned as an Investor in the "costly information" treatment excluding the first round of all sequences).

Table 3.9: Frequency of Information Purchase for Each Subject

Session	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6
No_Info_Cost1	0/14	0/9	0/7	**21/21**	0/18	0/24
No_Info_Cost2	4/27	**2/2**	5/24	2/6	**19/26**	0/2
No_Info_Cost3	0/8	2/20	**11/11**	**3/5**	**2/2**	0/23
Info_No_Cost1	0/23	N/A	0/4	0/23	1/9	0/16
Info_No_Cost2	0/19	0/14	0/19	4/12	**1/1**	**15/16**
Info_No_Cost3	1/7	5/20	1/3	0/7	**20/20**	0/24

In Table 3.9 ratios over 0.5 are in bold. Notice that 9/36=1/4 of subjects purchased more than half times when information is available. About half of subjects (17/36) never chose to buy information. One subject was never assigned as an Investor in the "costly information" treatment. The remaining 1/4 of subjects bought information at least once, but by less than half times when information is available.

One possible explanation for the evidence in Figure 3.4 is that some investors rely more on the information than other investors, however, some investors tend to trust even when the information on the trustee's previous behavior is not favorable. When the information is shown automatically, the trusting behavior of the latter investors increases the frequency of invest corresponding to a low ratio of return by the trustees. In the "costly information" treatment, investors who rely more seriously on the information choose to purchase the information, while those investors who tend to ignore the information choose not to purchase the information.

In order to examine whether subjects behave differently regarding to information, we divide all the subjects into two sets, by whether the frequency that the subjects purchased information is over or less than1/2 (i.e. whether the ratio in Table 3.9 is bold or not). Then we go back to the "information" treatment and calculate the frequency of invest conditional on the ratio of return for these two sets of investors.

Table 3.10: Frequency of Invest for Two sets of Investors in the Information Treatment

Sets of investors	$0 \leq r \leq 1/3$	$1/3 < r \leq 2/3$	$2/3 < r \leq 1$
Freq. of Purchase>1/2	0.286	0.333	0.875
Freq. of Purchase<1/2	0.143	0.478	0.874

Surprisingly, for those investors who purchased information with a frequency over 50% in the "costly information" treatment, the frequency of invest corresponding to "$r \leq 1/3$" is larger than that of the other set of investors who did not purchase information very often. This evidence does not support the hypothesis that the finding in Figure 3.4 comes from the fact that some investors always rely more seriously on the information and tend to buy information more frequently than other investors. In stead, it seems to suggest that investors punish more severely when they pay to see the information on trustee's bad behavior than when the information is sent automatically.

3.6 CONCLUSION

We have studied the development of trust and reciprocity among strangers in the indefinitely repeated trust game with random matching. The main treatment variable is the amount of information on the trustee's previous behavior.

Although the parameters of the game are chosen to support the social norm of trust and reciprocity as an equilibrium in all treatments, we find that full trust and reciprocity is difficult to sustain when no reputational information is available. Providing reputational information increases the frequencies of trust and reciprocity, and the effect of information becomes more significant when the amount of information is larger. Furthermore, providing the possibility to purchase information at a small cost also significantly increases the frequencies of trust and reciprocity compared to the case when the information is absent.

Our findings help explain the emergence and prevalence of the reputation system in many important economic transactions between randomly and anonymously individuals, such as the online feedback system on eBay or the credit bureau. The significant contribution of this paper is that we identified the importance of individual information in sustaining the efficient outcome. This understanding is of obvious importance to the design and operation of economic institutions.

For future research, we are interested in designing and comparing different reputational mechanisms that improve the efficiency of the markets with random and anonymous players. For instance, one comparison is between the online feedback system and credit bureau. Online feedback system is a representative of decentralized information collection by voluntary contribution and with free dissemination. On the contrary, the traditional credit bureau collects information in a centralized and mandatory method. Although we have shown in this paper that providing reputational information significantly increases efficiency, we still have not solved the practical issue of how to provide or finance the reputational system. Comparing the efficiency of these two commonly used mechanisms may help answer the question.

4.0 MOTIVES FOR CHARITABLE GIVING

4.1 INTRODUCTION

Fifty-five percent of American households give to secular charities. The average given among donor households is just under one thousand dollars.[1] In trying to model and better understand this behavior economists have argued that most of the benefits derived from such contributions can be classified as being either public or private in nature. On the public side is the purely altruistic benefit one may get from knowing that the collected funds help feed a starving child, secure vaccines against malaria, provide shelter for the homeless etc. On the private side is the benefit one may get from being the one who contributed the funds, be it an internal warm-glow from having done your bit, or any other private rewards resulting from the contribution.

While the first economic models on charitable contribution simply focused on the purely altruistic motive for giving, economists have now accepted the 'impure altruism' model where in addition donors benefit from being the one who secures the gift. Empirical studies on the motives for giving suggest that pure altruism is an incorrect characterization of people's preferences for giving. However to date researchers have not subjected the impure altruism model to a direct empirical test, and we have yet to get a sense of the relative weights donors may attach to the two benefits from giving.

We present an experimental study which tests comparative static predictions of both the pure and impure altruism model. In contrast to previous studies we account for the

[1] Authors' calculation based on the 2005 *Center on Philanthropy Panel Study*, the philanthropy module in the *Panel Study of Income Dynamics*. "Secular" giving is giving to relieve poverty, to human services, to improve people's health, education, the arts, etc., that is, all giving except giving to places of worship.

possibility that consistent with the impure altruism model the support for pure altruism may be sensitive to the point at which the motive for giving is evaluated. In particular in the impure altruism model increases in the contribution of others will shift the marginal motive for giving from a concern for altruism to a concern for the private benefit from giving. From a methodological viewpoint we develop an environment that closely mirrors those of the theoretical models. We examine contributions to an actual charity where each participant is informed of an initial contribution amount and singlehandedly determines the final dollar amount to be transferred to a recipient of the charity.

In examining charitable contributions across several budgets, we consistently find behavior in line with the comparative static predictions of the impure altruism model. However our data also make clear that a substantial weight is attached to the altruistic component of preferences. When estimating a representative utility function we find, consistent with our comparative static results, that participants get a private benefit from giving, however it accounts for but a fraction of the weight attached to the public benefit associated with providing funds for the charity's recipient. Thus we demonstrate that there are environments for which it would be incorrect to assume that donor's charitable contributions primarily are made because of the private benefit one may experience from giving.

4.2 MOTIVATION

Giving to charity can be viewed as contributing to a public good. The reason is that as one donor contributes to a charity other individuals, who are concerned for the well-being of the recipients of the charity, cannot be prevented from benefitting from this increase in contributions (non-exclusion) nor does the benefit they experience influence that experienced by anyone else (non-rival). Thus denoting an individual's private consumption by x_i, and the sum of contributions to the charity by G, we may describe individual preferences by, $U_i(x_i, G)$. This model has historically been referred to as the pure altruism model.

Economists have been quick to point out that the comparative static predictions of this model make it unlikely that this is an appropriate description of what motivates giving. For example, individuals who only are concerned about the total contributions to the charity will view contributions by others as a perfect substitute to giving by self. An implication is that if the government imposes a lump-sum tax on a contributor and donates the tax to the charity, then the individual will in turn decrease his charitable contribution by the size of the tax. Thus the government's contribution completely crowds out that of the individual, and the total amount received by the charity will be as it was before the tax was imposed.

Empirical studies of actual charitable contributions have failed to find evidence of complete crowding out. In fact most studies find that an increase in government giving only causes a small decrease in giving by the individual. While Bergstrom, Blume, and Varian (1986) demonstrate that incomplete crowd-out may result when the government's tax exceeds the individual's contribution, Andreoni (1991) shows that the complete crowd-out result still holds in a large economy.[2]

In response to the limited private response to government giving it has been argued that donors not only contribute because they care about the recipients of the charity, but also because they get a private benefit from being the one who secured the contribution. That is, a donor may get, say, a warm glow from having done his share. This suggests that the individual's private contribution needs to be included as a separate argument in the utility function, i.e., individual preferences may be captured by $U_i(x_i, G, g_i)$, where g_i denotes the individual's private contribution to the public good. In this model the contribution by others is seen as an imperfect substitute for private giving, thus a forced contribution through taxes is not equivalent to a voluntary one, and the individual will not decrease his contribution by the amount of the tax. Given normality the impure altruism model therefore predicts incomplete crowd-out and the forced contribution will result in an increase in total giving to the public good.

In determining whether the impure altruism model is a better description of what motivates giving, researchers have examined the degree of crowd-out on actual contributions to charity. One of the disadvantages of identifying preferences from the response in charitable giving to a change in government giving is that the analysis relies critically on the assumption that donors know how much the government contributes to the charity. Thus,

[2] This prediction relies on the, perhaps not reasonable, assumption that the need for the public good is independent of the population size.

while these studies are informative in shedding light on the effect an increase in government giving has on individual contributions, inferences on preferences are more difficult. To better control the information available to donors, researchers have instead moved to the experimental laboratory. Just as for the field studies the focus has been on determining how private giving responds to a forced increase in giving. That is, the objective has been to determine whether behavior is consistent with complete crowd-out and thereby the pure altruism model. Rejecting the null has been seen as evidence of the impure altruism model.

Andreoni (1993) is the first to study crowd-out in the laboratory. Rather than examining contributions to a charity, he had participants make costly contributions to a financially induced public good. Using payoff induced preferences Andreoni examined the change in contributions when individuals were forced to contribute a certain amount to the shared public good (i.e., imposing a lump sum tax). Contradictory to the pure altruism model he finds a crowd-out of 71.5 percent. This leads him to reject pure altruism in favor of impure altruism. In a follow-up study Bolton and Katok (1998) comment that Andreoni's study relies on the assumption that individuals only care about their own payoffs and not that of others, indeed there maybe multiple equilibria of the game when individuals have other-regarding preferences. Instead they design a very simple dictator game where each individual is put in a position where they must share a predetermined amount of money with another participant in the laboratory. In one treatment the decision maker (dictator) is initially given \$18 and the recipient \$2, in the other, the dictator is initially given \$15 and the recipient \$5. In a between-subject analysis dictators are informed that they are free to contribute any amount they wish to the recipient, with the contribution for the latter budget being forced to be \$3 greater than in the first. As Andreoni (1993) they find incomplete crowding out (60%) and reject the pure altruism model in favor of the impure one.

Eckel, Grossman, and Johnston (2005) extend Bolton and Katok's dictator-game study to one where contributions are made to an actual charity. Using the same allocations as Bolton and Katok they replace the recipient with a charity of the participant's choice and find no crowd-out – in fact they see evidence of crowding in. They also examine a variant of this treatment where the dictator is informed that his or her initial endowment is \$20, but that a certain amount already has been taxed and will be given to the charity of their choice. A significant framing effect is found as in this case individuals completely crowd-out the forced contribution.

Common for the earlier studies on motives for charitable giving is that pure altruism is the null hypothesis and rejections of the null is seen as evidence for impure altruism. Ribar and Wilhelm (2002) however make clear that if the impure altruism model is the correct one, then the degree of crowd-out may depend on the particular point at which one is evaluating it. Consider for example the case where the charitable contributions by others increase, while the individual's income is held constant. If this increase in giving occurs at a point where the initial giving by others is low, then the donor's marginal motive for giving will be the concern for the charity's recipients. With altruism being the marginal motive for giving, contributions by others will be viewed as a close substitute for giving by self. Thus at low initial contributions an increase in contributions by others may result in a substantial reduction in individual giving. As initial contributions to the charity increase, the marginal motive for giving shifts from a concern for securing the charity's services to a concern for the private benefit or warm-glow from giving. Thus for a well-endowed charity the marginal motive for giving will be the concern for warm-glow. An increase in giving by others will therefore be a distant substitute for private giving, and the effect on the individual's private contribution will be limited. The sensitivity to the evaluation point suggests that to identify the warm-glow component of giving we need to examine contributions at multiple budgets.

In an attempt to better understand what motivates charitable giving, this paper reports on an experimental study where participants make decisions at multiple different budgets. Varying both the participant's endowment and the initial charitable contribution we can examine individual contributions at several different points and thus evaluate the comparative static predictions of the pure and impure altruism models. In addition we make a number of important methodological innovations to secure that we consistent with the theory are dealing with a public good where we have complete control over the provision of the good.

4.3 EXPERIMENTAL DESIGN

To characterize the motives for charitable giving we wanted to examine an environment where each participant had complete control over contributions to an actual charity. Thus we did not want to examine contributions to other participants in the laboratory, nor did we want to examine contributions to a general charity. Our only option was therefore to create several 'new' public goods, that is, a series of public goods that no one else provides, and where the provision of each one is determined solely by the sum of a fixed initial contribution and the participant's contribution. Thus we needed to find an environment with multiple individual public goods that could vary in size. Finally to easily identify responses to changes we wanted a cause that would trigger substantial contributions.

In selecting a public good that not only satisfied the criteria above, but also had a good reputation we were fortunate to get the American Red Cross of Southwestern Pennsylvania to help us create a series of new public goods. In the event of a fire in Southwestern Pennsylvania this chapter of the American Red Cross helps the affected families find temporary shelter, provide them with clothing, a meal, and give them essential toiletries. Prior to our study no items were given to the children affected by the fire, and we were given permission to make donations of books to the children. We paired each participant in the study with a child (1-12 years old) whose family home had suffered extensive fire damage. The participant was given an endowment and asked to allocate it between him- or herself and the child. They were informed that the total amount to be spent on the child was the sum of the participant's contribution and a predetermined contribution made by the foundation funding the study. The books purchased for the allocated money would be distributed to the child by the American Red Cross, immediately after the child had been affected by a severe fire.[3] Participants were informed that "Each participant in this study is paired with a different child. Only you have the opportunity to contribute books to the child, neither the American Red Cross nor any other donors provide these books to the child."

In explaining why the American Red Cross is seeking the participant's donation for these books, they were informed that their Emergency Preparedness Coordinator Sandi Wraith

[3] Each child was given three books of varying value, with younger children getting the lower value books.

had made the following statement: "Children's needs are often overlooked in the immediate aftermath of a disaster because everyone is concerned primarily with putting the fire out, reaching safety, and finding shelter, food and clothing...just the basics of life. So many times, I've seen children just sitting on the curb with no one to talk to about what's happening...for this reason I've found trauma recovery experts in the community to work with us to train our volunteer responders in how to address children's needs at the scene of a disaster.......being able to give the children fun, distracting books will provide a great bridge for our volunteers to connect with kids and get them talking about what they've experienced."

A total of 85 undergraduates at the University of Pittsburgh participated in one of six sessions. There were between 13 and 20 participants in each session. The experiment was conducted as follows. Participants were seated in a large class room. They were given a folder with a set of instructions, a quiz, an envelope, a calculator, and a pen. The instructions were then read out loud, and the participants were given a brief quiz to make sure that they could calculate the payoffs of a sample decision. Having received answers to the quiz participants proceeded to the decision task.

To help identify individual preferences the study was conducted as a within-subject design. Participants were given a set of budgets for which the individual's endowment and the child's fixed initial donation varied. For each budget the participant was told that she or he was free to allocate any portion of the endowment to the child. The child would receive the sum of the initial allocation and the contribution made by the participant.

Although the study was designed to be double-blind, participants had the option of relinquishing their anonymity in the event that they wanted to receive a receipt from the Red Cross. Each decision sheet was identified only by a Claim Check number, and this number was used for the subject's anonymous payment. Once the decision task was completed the participant placed the decision sheet in the envelope, and with the exception of the participants who requested an acknowledgement forms, the decisions were from that point onward only identified by a Claim Check number. While one set of experimenters prepared the participants' payments in sealed envelopes, an experimenter who did not oversee the payment was in charge of distributing the envelopes by Claim Check number.

A careful procedure was used to assure participants that the experimental procedures were followed.[4] During the instruction phase we randomly selected one participant to be the monitor. The monitor's job was to oversee the procedures of the experiment. He or she followed the experimenters throughout the study, oversaw that the payment procedures were as described in the instructions and secured that for each child a check was issued to the American Red Cross for the amount determined by sum of the participant's donation and the relevant initial contribution. At the end of the experiment the monitor made a statement indicating whether the experimenters had followed the procedures described in the instructions. Participants were then shown the acknowledgements and checks that were to be sent to the American Red Cross. These were shown from a distance where no details could be determined. Once the participants had received their payment and left the study, the monitor walked with the experimenter to the nearest mailbox, and dropped the envelopes with the checks in the mail box. The monitor then signed a statement to certify that all procedures had been followed and the statement was subsequently posted in the economics department at the University of Pittsburgh. Similarly a receipt from the American Red Cross was subsequently posted.

During the decision task participants were presented with the six budgets shown in Table 4.1. For example, for budget 1 the participant is informed that the foundation paying for the study has contributed $4 towards books for the child, and that the participant has been given an endowment of $40 which must be allocated between him or herself and the child. At the end of the decision task the monitor randomly selected a number between 1 and 6, and the decision for the selected budget was carried out.

These six budgets allow us to test the comparative statics predicted by the alternative motives for giving, and to determine whether the conditions that give rise to them hold. We will focus on the following three tests.

Income Effect: We can determine the response to changes in income by comparing total contributions on the budgets where the initial contributions to the child is held constant and income increases (budget 1 vs. 5 and 3 vs. 6).

[4] The procedure is similar to that by Eckel et al. (2005).

Table 4.1: Experimental Budgets

Budget	$ to child	$ to self
1	4	40
2	10	40
3	28	40
4	34	40
5	4	46
6	28	46

Examining the effect of an increase in income is crucial for determining the comparative statistics of the motives for giving. While it is generally assumed that the impure altruism model predicts incomplete crowd out, this prediction only holds under normality. Thus prior to examining the comparative static predictions of the pure and impure altruism model we need to determine whether the two goods are normal. Comparing total contributions to the child between budget 1 and 5 and between 3 and 6 we determine whether the public and private good are normal by seeing if the $6 increase in income increases total giving and the amount of money the individual keeps for himself.

Balanced Budget Increases in Contributions: Comparing individual contributions between budget 2 vs. 5 and 4 vs. 6 determines the effect of balanced budget increases in contributions.

In moving from budget 5 to 2 we increase the initial contribution to the child by $6 while decreasing the participant's endowment by $6. Examining balanced budget increases in contributions we can examine both the pure and impure altruism model. According to the pure altruism model this balanced budget increase in giving should cause individual contributions to decrease by $6. The reason is that the concern for the well-being of the child is the sole motive for the pure altruist to give, thus the increase in initial contributions will be viewed as a perfect substitute for private giving. Similarly a pure altruist is predicted to decrease private giving by $6 when moving from budget 6 to 4.

Provided that the two goods are normal the impure altruism model instead predicts that contributions by others be an imperfect substitute for private giving. Thus an impure altruist will not view a forced contribution to the child as equivalent to a voluntary one, and the \$6 balanced budget increase in contributions will result in less than a \$6 decrease in private contributions. For an impure altruist total contributions to the public good increase as a result of budget balanced increases in giving, and the degree of crowd out will be incomplete.

Unfunded Increases in Initial Contributions: Comparing individual contributions between budget 1 vs. 2 and 3 vs. 4 determines the effect of unfunded increases in contributions.

Unfunded increases in initial giving enable us to test comparative statics of three different motives of giving. First, a pure egoist or a pure warm-glow giver is someone for whom the sole motive for giving is a concern for the amount he contributes (i.e., preferences are of the form $U_i(x_i, g_i)$). Thus changing the initial contribution to the child should have no effect on the individual's contribution, and contributions should be the same for budgets 1 through 4.

Second, for a pure altruist unfunded increases in initial contributions should trigger a response equivalent to that of an increase in income. Consider for example a pure altruist who views contributions to the public good as a luxury good, this person's sensitivity to increases in initial contributions will diminish as the initial contribution level increases. Figure 4.1 demonstrates such a scenario. The consumption of the public good is measured on the horizontal axis and the consumption of the private good on the vertical. Consider first the situation where the individual's endowment is w and the initial contribution to the public good is G_{-i}. Since the individual can only add to this initial contribution he is free to choose any budget along the kinked bold budget line in Figure 4.1. Given the budget line and the individual's income expansion path we note that his preferred level of the public good is at G^*. Thus his contribution to the public good equals $g_i = G^* - G_{-i}$. Now consider the response to an unfunded increase of Δ in the initial contribution to the public good. This increase expands the budget set as shown below, and as seen by the income expansion path the individual's preferred contribution level increases to G^{**}. Thus the individual wants the Δ increase in initial contributions to result in an increase in consumption of both the private and public good. To secure this outcome the individual reduces his contribution to offset the increase in initial giving such that the increase in provision is limited to that from G^* to G^{**}, as indicated by the dark line along the horizontal axis. Now consider instead the case where the same Δ increase in initial contributions occurred at a higher initial contribution level such as G'_{-i}. When the public good is viewed as a luxury, the income expansion path will be concave, and as seen in Figure 4.1, relative to our initial evaluation point the increase in initial giving results in a larger increase in the participant's desired total contribution and a small increase in the consumption of the private good. Thus the change in private contribution to an unfunded increase in giving will be smaller when the initial contribution level is large. For a pure altruist who views the good as a luxury we should therefore see diminished sensitivity to unfunded changes in initial contributions.

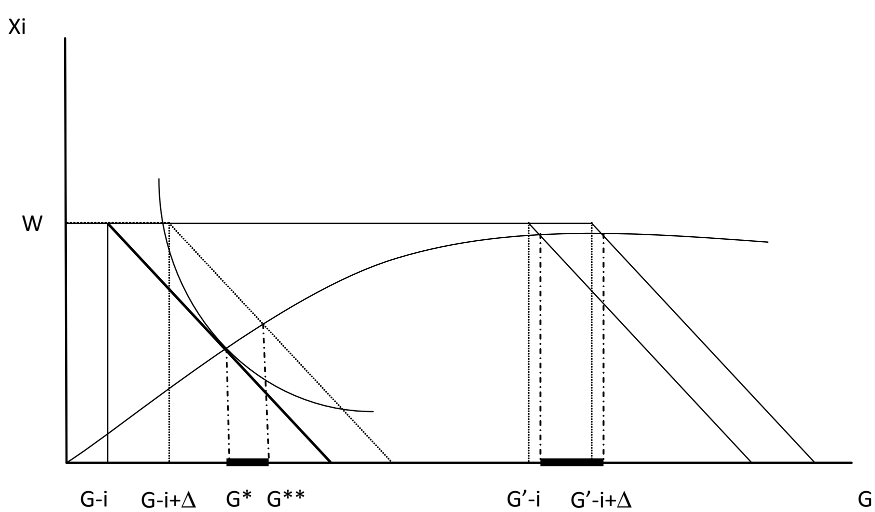

Figure 4.1: Response to Unfunded Increase in Initial Contributions

If we instead consider the case where the public good is viewed as a necessity, then the income expansion path will be convex and a pure altruist will exhibit increased sensitivity to unfunded increases in initial contributions. Finally, when a pure altruist has homothetic preferences the change in private giving in response to unfunded increases should be independent of the evaluation point.

Considering the response for a pure altruist we may then ask what the predictions are for an impure altruist's response to increases in initial contributions? In moving from budget 1 to 2 and from 3 to 4 the initial contribution to the child increases by $6, however given that the point for evaluating the change also increases (from $4 to $28) the marginal motive for giving may shift from a concern for the child's well being to a concern for the private benefit associated with giving. Thus as the initial contribution increases, donations by others become a worse and worse substitute for individual giving, and this shift will all else equal result in diminished sensitivity to changes in initial contributions, i.e., all else equal the response to the $6 change in initial giving is predicted to be larger when comparing budget 1 to 2, than when comparing budget 3 to 4.

Combined the predicted response to unfunded increases in initial contributions is as summarized in Table 4.2. While we are unable to predict the response for an impure altruist who views the good as a necessity, it is clear that constant or diminished sensitivity will be seen as strictly supportive of the impure altruism model. As for balanced budget increases in contributions we will examine the response to changes in income to determine what the relevant comparative statics should be.

Table 4.2: Predicted Response to Unfunded Increases in Initial Contributions

	Pure Warm-glow	Pure Altruism	Impure Altruism
Luxury	No response	Diminished sensitivity	Diminished sensitivity
Homothetic	No response	Constant sensitivity	Diminished sensitivity
Necessity	No response	Increased sensitivity	?

4.4 RESULTS

In characterizing the results we start by examining whether the cause we selected motivated participants to contribute. Across the six budgets there is substantial variation in giving and we are fortunate to have only a few participants who consistently appeared truncated in their choices. Only one of our participants stuck to a contribution of zero for every one of the six budgets, and an additional seven participants had average contributions at or below $3. At the other extreme we have five participants who gave their entire endowment away for everyone of the six budgets, and an additional two participants who made average contributions of $38 or more towards the child. Average individual giving across the six budgets was $20.82 (std.err. 1.16), however as seen by the contribution distribution in Figure 4.2 the individual variation was large.

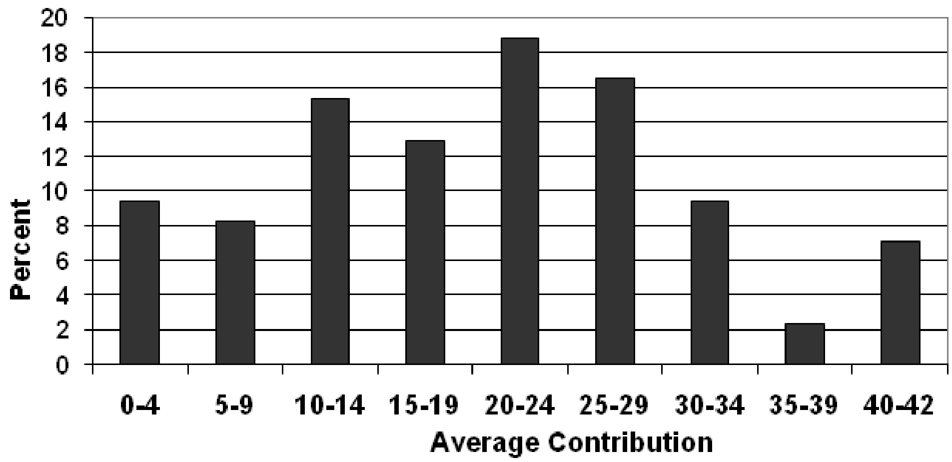

Figure 4.2: Distribution of Average Contributions across the Six Budgets

In our analysis of the data we first determine the response to changes in income. Next we examine the standard balanced budget crowd-out measure to determine the evidence in favor of the pure versus the impure altruism model. We then proceed to the response to unfunded increases in initial contributions to the child to determine how sensitive the response is to the point of evaluation. Finally, to evaluate the weight attached to the private benefit of giving we estimate a representative Cobb Douglas utility function.

4.4.1 Income Effect

To determine the predictions of the impure altruism model we need to examine how participants respond to changes in income. Figure 4.3 demonstrates how the contribution distribution shifts to the right as income increases from $40 to $46 while holding the initial contribution fixed at $4 (panel A) or $28 (panel B). Increasing income by $6 increases contributions by \$2.39 at low initial giving (\$4) and by \$2.46 at high initial giving (\$28). These increases are both significantly different from zero ($p < 0.001$) and smaller than \$6 ($p < 0.001$), thus participants on average respond as if the charity and their private consumption were normal goods.[5]

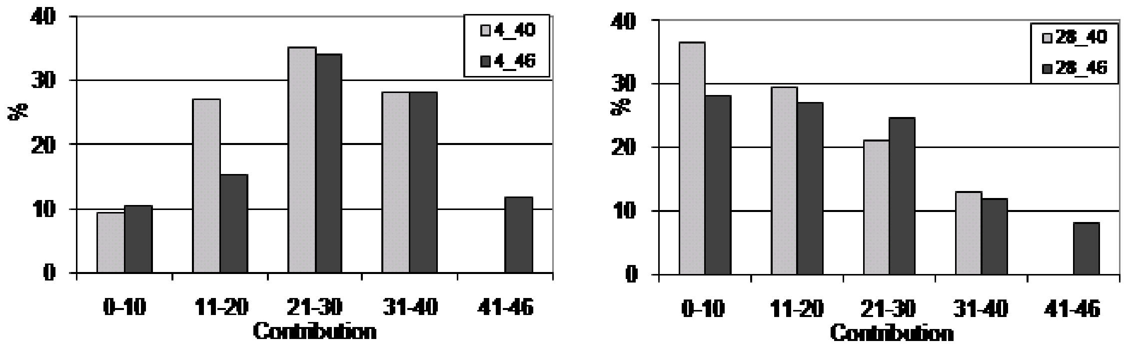

Panel A: Initial contribution \$4　　Panel B: Initial contribution \$28

Figure 4.3: Contribution Distributions Conditional on Participant Income

Individual responses also show evidence of the public good being normal. Of the 2x85 cases where we may observe changes in demand as a result of an increase in income, we only find a decrease in demand in 10 percent of the cases. Thus it is fair to argue that participants respond to the public good as if it is a normal good.

4.4.2 Balanced Budget Crowd-out

Given normality we can use our balanced budget crowd-out test to determine whether the contributions we observe are consistent with the impure versus impure altruism model. As the studies before us we test for pure altruism by determining whether a forced contribution completely crowds out a voluntary one. In contrast to earlier studies however we do not rely on just one measure of crowd-out. Rather our design makes it possible to measure the response to a forced \$6 contribution at two different contribution levels. In one case

[5] We are unable to reject the hypothesis that the response is the same at high and low initial contribution levels ($p = 0.87$).

initial giving increases from \$4 to \$10 and in the other from \$28 to \$34, with the participant's endowment decreasing from \$46 to \$40 in both instances. Table 4.3 reports average contributions for these four budgets as well as the associated degree of crowd-out.

Table 4.3: Balanced Budget Crowd-Outs

Initial	Budget		Actual (n=85)		Non-truncated	
giving	Initial	Endowment	Contribution	Crowd-out	Contribution	Crowd-out
Low	4	46	27.2 (1.26)		26.7 (1.16)	
	10	40	21.6 (1.13)	-5.66 (0.51)	21.0 (0.99)	-5.70 (0.57)
High	28	46	19.5 (1.38)		19.8 (1.22)	
	34	40	14.8 (1.25)	-4.63 (0.47)	14.6 (1.06)	-5.21 (0.43)

The crowd-out at low initial giving is astoundingly high. Although the actual decrease of \$5.66 is slightly lower than the predicted \$6, we cannot reject that the decrease is consistent with complete crowd-out ($p = 0.51$). Had we followed the procedures of previous studies and only examined one crowd-out measure, this result would have led us to conclude that on average participants appear to be purely altruistic. However we would have reached the opposite conclusion if instead we had elicited the crowd-out at a higher initial contribution level. At a high initial contribution level the crowd-out is 77 rather than 94 percent, and we easily reject the hypothesis that the \$4.63 decrease in giving equals the predicted \$6 ($p = 0.004$). In testing if the crowd-out measures at high and low initial giving are the same we get a p-value of 0.11.

While the decrease in crowd-out is inconsistent with the pure altruism model it is entirely consistent with the impure altruism model, and it demonstrates that to identify the motives for giving one needs to elicit multiple crowd-out measures. We do however need to use some caution before readily accepting this decrease in crowd-out as evidence of impure altruism. The reason is that when participants' decisions are truncated the calculated crowd-out measure may lead us to make incorrect inferences regarding the underlying preferences.

First, at the higher initial contribution level we are more likely to observe participants contributing nothing, and this will bias our results in favor of the impure altruism model.

For example, a pure altruist may decrease giving by $6 between the two low-contribution budgets and opt for a zero contribution at both of the high-contribution budgets. Thus consistent with the impure altruism model we will see a decrease in crowd-out ($6 to $0) despite the participant being purely altruistic. Indeed contributions are more truncated at higher initial giving. While 3 people opted to give nothing at (10,40), a total of 9 people gave nothing at (34,40). If we calculate the degree of crowd-out for the participants who made positive contributions at (10,40) and (34,40), respectively, these increase to -5.7 (n=82) for low initial giving and -4.96 (n=76) for high initial giving. Second, we also need to account for the bias that may arise from participants being truncated at the other extreme, i.e., participants who contributed their entire endowment. While the direction of the bias is less obvious in this case, there is the possibility that the crowd-out will be exaggerated at low initial contribution levels. Consider for example a pure warm-glow giver who contributes his entire income at (4,46), while this individual is forced to decrease the contribution to $40 at (10,40) the unconstrained crowd-out is likely to have been smaller. To account for this truncation we eliminate participants who contributed their entire endowment at (10,40) and (34,40), respectively.[6] The last column of Table 4.3 reports crowd-outs that result when eliminating participants who are truncated either at the top or the bottom. Although the degree of crowd-out increases when eliminating truncated observations, increases in initial giving continues to decrease the degree of crowd-out (from 95% to 87%), and as before we can only reject that there is complete crowd-out in the latter case ($p = 0.6$ and $p = 0.07$). However, when we only include the 69 participants who were not truncated at (10,40) nor at (34,40) we cannot reject that the two crowd-out measures are the same ($p = 0.40$).

Further evidence for the decrease in crowd-out may be seen at the individual level. Table 4.4 reports the number of participants for whom the calculated crowd-out was positive, zero, incomplete, complete, or greater than the forced increase in giving. While for low initial giving the mode of participants responded to the forced $6 contribution by decreasing their contribution by precisely $6, at a higher initial contribution level the mode is for participants

[6] The reason for eliminating participants who are constrained at this level is that an impure altruist may have made an unconstrained contribution of say $45 at (4,46) however in not viewing the forced contribution as a perfect substitute for their own contribution he may wish to contribute in excess of $40 at (10,40) causing the resulting decision to be constrained and the measured degree of crowd-out smaller than it would have been given an unconstrained choice set.

to decrease the contribution by less than \$6. Thus as demonstrated by the average crowd-out measure, there is greater evidence of impure altruism at the larger initial contribution level.

Table 4.4: Participants by Degree of Crowd-out

		Number of participants (n=69)	
	Change in giving	Low initial giving (\$4)	High initial giving (\$28)
Crowd in	> 0	5	2
No change	$= 0$	4	4
Crowd-out	(0,-6)	18	**30**
	= -6	**24**	15
	< -6	18	18

Note however that a number of observations are inconsistent with the standard interpretation of the impure altruism model. Specifically a small number of participants respond to the forced increase in giving by increasing their own contribution, and about a fifth of the participants decrease their contribution by more than the increase in the initial contribution. While inconsistent with the standard interpretation of the impure altruism model, this behavior may arise if an impure altruist views the public good as an inferior good. Interestingly participants who exhibit excessive crowd out are much more likely to make contribution decisions in \$5 increments.[7]

4.4.3 Unfunded Increases in Initial Contributions

Although the evidence from the balanced budget increase in initial contributions demonstrates a strong altruistic motive, the decrease in crowd-out at the high initial contribution level suggests that impure altruism may be a better approximation of the participants' motives for giving. We therefore proceed by examining the response to unfunded increases in initial contributions. We start by examining the evidence in favor of participants being motivated solely by warm glow.

[7] Participants with excessive crowd out (greater than \$6) on average make decisions that are divisible by \$5 on 3.5-4 budgets out of the 6. By comparison those with complete or incomplete crowd out on average make decisions that are divisible by \$5 on 2-2.5 of the 6 budgets.

Pure Warm-Glow Giving

A pure warm-glow giver is predicted to make contribution decisions that only depend on his income and not on the initial contribution to the child. Holding the participant's income fixed at $40, Figure 4.4 demonstrates how the distribution of contributions changes as the initial amount given to the child takes on the values $4, $10, $28 and $34. In contrast to a pure warm-glow motive an increase in the initial amount given decreases that of the participant shifting the distribution to the left.

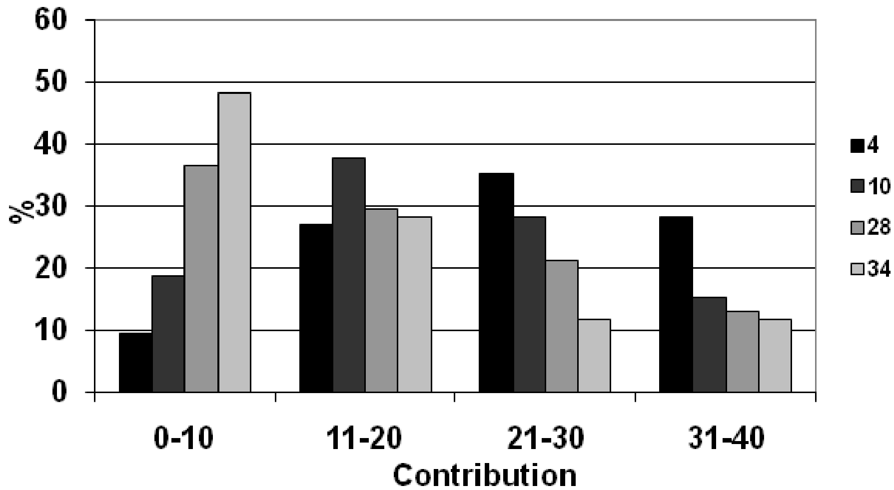

Figure 4.4: Contribution Distributions Conditional on Initial Contribution

As the initial contribution to the child increases from $4 to $10 and then again from $28 to $34, the average amount given by participants decreases from $24.8 to $21.6 and then from $17.0 to $14.8, with every one of these decreases being significant. Similarly when holding income fixed at $46 and increasing initial giving from $4 to $28 (budgets 5 and 6) the contribution to the child decreases significantly from $27.2 to $19.5. Thus we easily reject that on average participants are solely motivated by the private benefit from giving. Looking at the individual patterns of contributions a total of 14 participants have contributions that only depend on their endowment. As noted earlier one of these participants never gave anything and five consistently gave their entire endowment. At a minimum it therefore appears reasonable to argue that the remaining eight behaved as if they only were concerned about the amount they contributed.[8] The vast majority of participants can however not be characterized as pure warm-glow givers and we continue by examining the evidence in favor of pure versus impure altruism.

Pure versus Impure Altruism

As the initial contribution to the child increases an impure altruist's marginal motive for giving will shift from a concern for securing the public good to a concern for the warm-glow from giving, that is it shifts from a concern for the public good, G, to the private contribution, g_i. Thus the impure altruism model pulls behavior in the direction of diminished sensitivity to changes in initial contributions.

We study the sensitivity to unfunded changes in initial contributions by comparing the budgets for which income is held fixed at \$40 and the initial contribution to the child increases first from \$4 to \$10 and then from \$28 to \$34. While the change in initial contributions is \$6 in both comparisons, the marginal motive for contributing may be changing between these two levels. We first examine the response at the aggregate level. The participants' contributions and the resulting changes in giving are reported in Table 4.5.

Table 4.5: Sensitivity to Changes in Initial Contributions

Initial	Budget		Actual (n=85)		Non-truncated	
giving	Initial	Endowment	Contribution	Crowd-out	Contribution	Crowd-out
Low	4	40	24.8 (1.10)		24.5 (1.00)	
	10	40	21.6 (1.13)	-3.28 (0.44)	21.0 (0.99)	-3.59 (0.48)
High	28	40	17.0 (1.27)		17.1 (1.15)	
	34	40	14.8 (1.25)	-2.16 (0.39)	14.6 (1.06)	-2.54 (0.44)

We find that there is diminished sensitivity to initial giving, as the \$6 increase in initial contributions causes a greater decrease in private giving between \$4 and \$10 than between \$28 and \$34. The \$6 increase in initial giving results in a significantly ($p = 0.03$) greater decrease in contributions when the initial contribution is low (-3.28) rather than high (-2.16).

[8] The mean contribution for these eight participants was \$22.45 (std. err. 3.64).

As for our earlier analysis we eliminate the participants for whom the change in demand may be truncated. Eliminating participants who for the (10,40) budget gave nothing or gave their entire endowment reduces the sample to 76 participants and increases the change in giving slightly, similarly the change in giving is greater and the sample drops (n= 70) when removing truncated participants at (34,40). Looking at the 69 participants who were neither truncated at (10, 40) nor at (34, 40) we continue to see a significantly greater response in giving at low than at high initial contribution levels ($p = 0.04$).

Although diminished sensitivity to initial contributions is consistent with the impure altruism model, it need not be inconsistent with that of pure altruism. In particular a pure altruist, who views the public good as a luxury, would behave in much the same way, with total contributions increasing at an increasing rate with the initial contribution. To distinguish between the predictions of the pure and impure altruism model we therefore need to determine the income elasticity of demand. We calculate two measures of the income elasticity of demand for total giving to the charity by returning to the four budgets where income is increased from \$40 to \$46 and the initial contribution fixed at either \$4 or \$28.[9] Including all participants the average income elasticity of demand equals 0.65 (0.16) at initial giving of \$4 and 0.63 (0.10) at initial giving of \$28. Although we are unable to reject that these two measures are the same ($p = 0.89$), we can reject that they equal 1 ($p < 0.04$). Thus on average participants do not respond to the public good as if it were a luxury. The results are the same when we eliminate participants with truncated decisions. For the 69 non-truncated participants the income elasticity of demand is 0.53 (0.12) when the initial contribution is \$4 and 0.71 (0.12) when it is \$28.

To further examine the sensitivity to initial contributions, we classify participants according to their average income elasticity of demand. We see that the 18 participants who on average view the public good as a luxury demonstrate diminished sensitivity to changes in initial contributions. That is their response to changes in initial contributions is greater at low initial contribution level than at high initial contributions. When the initial contribution is \$4 these participants respond to a \$6 increase in initial contributions by reducing their individual contribution by \$1.89, however when the same \$6 increase in contributions occurs

[9] Income elasticity of demand = % change in total giving / % change in social income.

at an initial level of \$28, then the decrease in individual giving is only \$0.44. Thus when the good is viewed as a luxury, participants demonstrate significant diminished sensitivity to changes in the initial contributions ($p = 0.06$). As shown earlier this result is consistent with the predictions of both the pure and impure altruism model. In contrast the impure and pure altruism model may generate different results for the 46 participants who view the public good as a necessity. While the pure altruism model predicts that there be increased sensitivity to changes in initial contributions, the impure altruism model may result in increased, constant, or diminished sensitivity. As seen in Table 4.6, we do in fact find that consistent with the impure altruism model there is diminished sensitivity to initial contributions, and we reject the hypothesis that there consistent with pure altruism is increased sensitivity to changes in the initial contribution.

Table 4.6: Sensitivity to \$6 Change in Initial Contribution

	N	Income Elasticity	Change in initial contribution Low initial (\$4)	High Initial (\$28)	Change in CO	p-value
Luxury	18	1.49	-1.89	-0.44	1.45	0.06
Necessity	46	0.44	-3.76	-3.11	0.65	0.40

Our results suggest that the observed diminished sensitivity to initial contributions is evidence of impure altruism, and combined with the decrease in crowd-out on balanced budgets it seems reasonable to conclude that our participants are not only motivated by altruism. It is however important to note that our data also suggest that altruism plays a substantial role. To determine the relative weight attached to either motive we proceed by estimating a representative utility function for our sample.

4.4.4 A Representative Cobb-Douglas Utility Function

To get a quantitative indication of the relative strengths of altruistic and warm glow motives in the aggregate, we use the data to estimate the parameters of a Cobb-Douglas specification of impure altruism:

$$U(x, G, g_i) = (1-\alpha-\beta)lnx + \alpha lnG + \beta lng_i$$

where x is the consumption of the private good, G the total contribution to the public good, g_i the individuals contribution to the public good, and α and β are the parameter weights on the altruistic and warm glow motives. The optimal contribution g_i is given by:

$$g_i = -G_{-i} + 1/2[(1-\beta)G_{-i} + (\alpha+\beta)Z_i + \sqrt{[(1-\beta)G_{-i} + (\alpha+\beta)Z_i]^2 - 4\alpha G_{-i}Z_i}]$$

where G_{-i} is the initial contribution to the child, $Z_i \equiv w_i + G_{-i}$ is social income. Throughout the paper we have shown that there is substantial heterogeneity in the individual contribution decisions. To account for the possibility that random departures of giving from the Cobb-Douglas specification are likely correlated within-subject we estimate the model using the following random effect specification.[10]

$$g_{ib} = -G_{-ib} + 1/2[(1-\beta)G_{-ib} + (\alpha+\beta)Z_b + \sqrt{[(1-\beta)G_{-ib} + (\alpha+\beta)Z_b]^2 - 4\alpha G_{-ib}Z_b}] + e_i + u_{ib}$$

Where the e_i part of the additive error term is the individual random effect, $i = 1, ..., 85$ indexes the individual subjects and $b = 1, ..., 6$ indexes the six budgets of own income and initial contribution that each individual faces. To account for the participants who at one budget or another hit a corner solution at $g_{ib} = 0$ or $g_{ib} = w_b$, we estimate the model using non-linear, random effects Tobit.[11]

Consistent with our earlier findings we see that the coefficient on the warm-glow component is significantly greater than zero, however the coefficient is relatively small. One way of assessing the weight on warm glow is to examine the predicted individual contribution at a particular budget and determine what fraction of total giving would be accounted for by the warm-glow component. Consider for example budget (4,46), based on the 85 participants

[10] The first-order condition is quadratic in g. This implies two solutions for the optimal g_{ib}, one solution as shown in (2) and the other as in (2) but with a "$-$" replacing the "+" before the square root term. Estimation based on the solution with the negative square root term moves toward nonsensical vales of $\widehat{\alpha}$ and $\widehat{\beta}$ and does not converge.

[11] We calculate the estimates of α and β using maximum likelihood, assuming that u_i and e_{ib} are normally distributed. To calculate the multivariate normal probabilities when $g_{ib} = 0$ we use STATA's maximum simulated likelihood routines (Cappellari and Jenkins 2006), adapting Barslund's (2007) multivariate Tobit program.

Table 4.7: NLLS Random Effects Tobit Estimate of Representative Utility Function

	n=85 participants		
	Coefficient	Std. err	p-value
Alpha	0.595	0.024	0.00
Beta	0.021	0.009	0.02

the predicted contribution at this budget is \$26.82. Actual contributions at this budget were an average of \$27.22 (std.err. 1.26). Note however that if at this budget we had instead set $\alpha = 0$ then a pure warm-glow giver would only contribute \$0.95. Thus, using the Cobb-Douglas specification the predicted net effect of warm-glow at this budget is less than 5 percent of the individual's total contribution.

4.5 CONCLUSION

Economists have long been skeptical of the extreme predictions of the pure altruism model, and we have enthusiastically embraced the impure altruism model. However to date we have failed to carefully subject this model to a direct empirical test. Using experimental methods we have developed a design that directly tests the comparative statics of the impure altruism model and simultaneously improves upon previous tests of the pure altruism model. When examining both the standard balanced budget crowd-out and the sensitivity to initial contribution we find behavior that is very much in line with that predicted by impure altruism. In particular we find diminished balanced budget crowd-out and diminished sensitivity to increases in initial contributions to the charity. That being said the data also demonstrate a substantial altruistic component. When estimating the weights attached to the altruistic and warm-glow components of a representative utility function, we find that quantitatively the weight on warm-glow is very small and if one was forced to characterize the participants as either altruistic or pure warm-glow givers, then the first would be the better choice.

As in any study one needs to use caution when trying to extend the results to other domains. For example, there is no reason to believe that the preferences estimated here extend to all types of charity. It may very well be that preferences for humanitarian causes are more altruistic than for non-humanitarian ones. However in contrast to earlier studies our results do demonstrate that there are causes for which economic man sacrifices his hard-earned money - not because he wants to feel good about his deeds - but rather because he is concerned for the welfare of others.

APPENDIX A

DISCUSSION OF CRITERION D_1

In this appendix we first describe two equilibria when the seller has a pessimistic ex ante belief ($\alpha < \widetilde{\alpha}$) and show that only Equilibrium A can survive if we impose more restrictions on the seller's belief off the equilibrium path. Then we give the formal definition of criterion D_1 and present the effect of applying this criterion.

Equilibrium A: The seller always offers $p_1 = p_2 = l$; All buyer types always accept p_t if $p_t \leq v_t$ and reject p_t if $p_t > v_t$; The seller always forms a belief less than the cutoff belief γ^*.

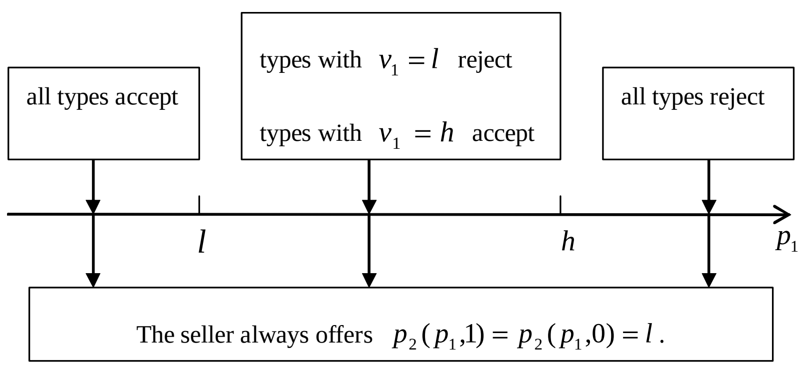

Figure A.1: Equilibrium A

Equilibrium B: The seller offers $p_1 = l + \delta q^F(h-l)$, $p_2 = h$ if $p_1 \leq l + \delta q^F(h-l)$ is rejected, $p_2 = l$ otherwise; Buyer types with $v_1 = l$ accept $p_1 \leq l + \delta q^F(h-l)$ and reject $p_1 > l + \delta q^F(h-l)$; Buyer types with $v_1 = h$ accept $p_1 \leq h$ and reject $p_1 > h$; The seller forms a belief greater than the cutoff belief γ^* after $p_1 \leq l + \delta q^F(h-l)$ is rejected and a belief less than the cutoff belief γ^* otherwise.

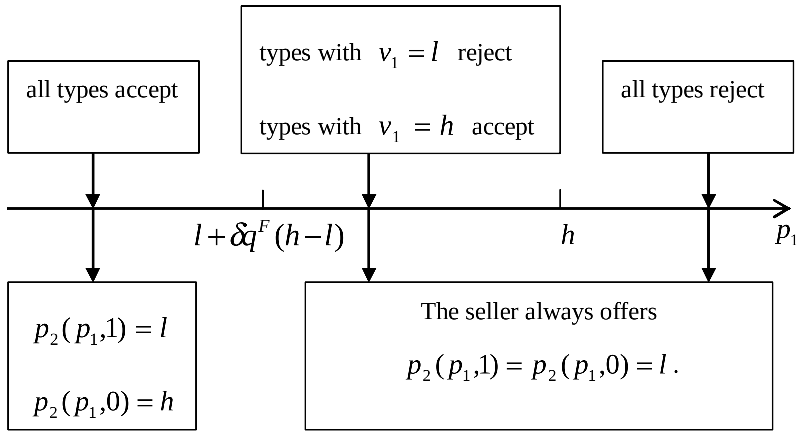

Figure A.2: Equilibrium B

We have discussed Equilibrium A in Proposition 1.

In Equilibrium B, the buyer types with $v_1 = l$ are willing to accept $p_1 \in (l, l+\delta q^F(h-l)]$ and incur a loss in the first period since the seller threatens to offer $p_2 = h$ if p_1 is rejected. In equilibrium, no type rejects p_1 in this range, so the seller can assign any belief after p_1 is rejected, so $p_2 = h$ after rejection of p_1 can be sustained as part of the equilibrium strategy.

Equilibrium B, however, is not very reasonable if we impose more restrictions on the seller's belief off the equilibrium path. Suppose the seller unexpectedly observes rejection of $p_1 \in (l, l+\delta q^F(h-l)]$ and tries to figure out which type is most likely to reject the offer. First note that the buyer types with $v_1 = h$ gain a positive payoff in the first period by accepting $p_1 \in (l, l + \delta q^F(h-l)]$ and the lowest offer $p_2 = l$ after acceptance of p_1, so the types with $v_1 = h$ are strictly worse off by rejecting p_1. Second, the seller finds that if her response to rejection makes type (G, l) better off or indifferent from the equilibrium by deviation, then type (F, l) is definitely better off by deviation. Thus, type (F, l) is most likely to reject p_1 among all the buyer types. If this is true, the seller should offer $p_2 = l$ instead of $p_2 = h$ when p_1 is rejected, and Equilibrium B is no longer sequentially rational.

Equilibrium A does not have this problem. In Equilibrium A, since the seller always offers $p_2 = l$ and the buyer always accepts $p_1 \leq v_1$, no buyer type can be better off by deviation.

Thus, the criterion applied on Equilibrium B does not have any effect on Equilibrium A.

Criterion D_1 formalizes the idea discussed above. The intuition conveyed by criterion D_1 is as follows. Consider a fixed equilibrium and the out-of-equilibrium action a_1 in the continuation game following any p_1. If whenever buyer type θ_1 has an incentive to defect from the equilibrium and send the out-of-equilibrium message a_1, or is indifferent between the equilibrium and defection, some other buyer type $\widetilde{\theta}_1$ is strictly better off by defection, then the seller's beliefs should assign zero probability to type θ_1 when the seller unexpectedly observes the out-of-equilibrium action a_1. Definition 3 formally defines criterion D_1 in the continuation game.

Definition 3 (Criterion D_1 in the Continuation Game Following p_1). *Consider a fixed equilibrium on the continuation of p_1, with action $a_1 \in \{0, 1\}$ reached with zero probability. Suppose $x(p_1, 1)$ and $x(p_1, 0)$ is the seller's equilibrium strategy.*

Step 1: *Find the sets of all (mixed) responses ϕ by the seller that would cause type $\theta_1 = (d, v_1)$ to defect from the equilibrium and to be indifferent. If $a_1 = 0$ is the out-of-equilibrium action, form the sets*

$$D_{\theta_1} \equiv \{\phi : (v_1 - p_1) + \delta q^d x(p_1, 1)(h - l) < \delta q^d \phi(h - l), \phi \in [0, 1]\},$$

$$D^0_{\theta_1} \equiv \{\phi : (v_1 - p_1) + \delta q^d x(p_1, 1)(h - l) = \delta q^d \phi(h - l), \phi \in [0, 1]\}.$$

If $a_1 = 1$ is the out-of-equilibrium action, form the sets

$$D_{\theta_1} \equiv \{\phi : (v_1 - p_1) + \delta q^d \phi(h - l) > \delta q^d x(p_1, 0)(h - l), \phi \in [0, 1]\},$$

$$D^0_{\theta_1} \equiv \{\phi : (v_1 - p_1) + \delta q^d \phi(h - l) = \delta q^d x(p_1, 0)(h - l), \phi \in [0, 1]\}.$$

Step 2: *For a given out-of-equilibrium action a_1, if for some type θ_1 there exists a second type $\widetilde{\theta}_1$ with $D_{\theta_1} \cup D^0_{\theta_1} \subsetneq D_{\widetilde{\theta}_1}$, then the combination (θ_1, a_1) may be pruned from the continuation game following p_1.*

Step 3: *Check whether the fixed equilibrium is still sequentially rational given that the seller's belief is restricted to the buyer types who survive from Step 2. If not, then the equilibrium does not survive from D_1.*

Given a PBE, if the corresponding equilibrium in all the continuation games survives from D_1, then we say that the PBE survives from D_1.

Definition 4 (D_1 Equilibrium). *A PBE of the game survives from D_1 if and only if the equilibrium on the continuation of p_1 survives from D_1 for all $p_1 \in \mathbb{R}$.*

The effect of applying criterion D_1 in our model is summarized in the following lemma.

Lemma 15. *The following results hold for a PBE in the continuation game following p_1:*

(i) If $p_1 > l$ and all buyer types accept p_1, the equilibrium in the continuation game can not pass criterion D_1;

(ii) If $p_1 \leq l$ and all buyer types accept p_1, the equilibrium in the continuation game passes criterion D_1;

(iii) If $p_1 < h$ and all buyer types reject p_1, the equilibrium in the continuation game can not pass criterion D_1;

(iv) If $p_1 \geq h$ and all buyer types reject p_1, the equilibrium in the continuation game passes criterion D_1.

APPENDIX B

PROOFS FOR CHAPTER 2

Proof of Lemma 3. Suppose $\Psi(p_1, 1) \in (0, 1)$ and $x(p_1, 0) < x(p_1, 1)$ in a PBE. Then $l < C(F, l) < C(G, l) < h < C(F, h) < C(G, h)$; $\Psi(p_1, 1) \in (0, 1)$ only if $p_1 \in [C(F, l), C(G, h)]$. We will show that it reaches a contradiction for any $p_1 \in [C(F, l), C(G, h)]$.

If $p_1 \in [C(F, l), C(G, l))$, then only type (F, l) rejects p_1 and $x(p_1, 0) = 1 \geq x(p_1, 1)$.

If $p_1 \in (C(F, h), C(G, h)]$, then only type (G, h) accepts p_1 and $x(p_1, 1) = 0 \leq x(p_1, 0)$.

If $p_1 \in (C(G, l), C(F, h))$, then $\gamma(p_1, 0) = \underline{\gamma}(\alpha) < \alpha < \overline{\gamma}(\alpha) = \gamma(p_1, 1)$ and $x(p_1, 0) \geq x(p_1, 1)$.

Denote X' as the probability for type (G, l) to reject $p_1 = C(G, l)$ and Y' as the probability for type (F, h) to reject $p_1 = C(F, h)$.

If $p_1 = C(G, l)$,

$$\gamma(p_1, 0) = \frac{\alpha X'(1 - q^G)}{\alpha X'(1 - q^G) + (1 - \alpha)(1 - q^F)} < \underline{\gamma}(\alpha)$$

and

$$\gamma(p_1, 1) = \frac{\alpha q^G + \alpha(1 - X')(1 - q^G)}{\alpha q^G + (1 - \alpha)q^F + \alpha(1 - X')(1 - q^G)} > \overline{\gamma}(\alpha),$$

so $x(p_1, 0) \geq x(p_1, 1)$.

If $p_1 = C(F, h)$,

$$\gamma(p_1, 0) = \frac{\alpha(1 - q^G)}{\alpha(1 - q^G) + (1 - \alpha)(1 - q^F) + (1 - \alpha)Y'q^F} < \underline{\gamma}(\alpha)$$

and

$$\gamma(p_1, 1) = \frac{\alpha q^G}{\alpha q^G + (1-\alpha)(1-Y')q^F} > \overline{\gamma}(\alpha),$$

so $x(p_1, 0) \geq x(p_1, 1)$.

Therefore, every case is contradictory to $x(p_1, 0) < x(p_1, 1)$, and the seller offers $x(p_1, 0) \geq x(p_1, 1)$ in a PBE if $\Psi(p_1, 1) \in (0, 1)$. □

Proof of Lemma 4. (*i*) $\Psi(p_1, 1) \in (0, 1)$ and $x(p_1, 0) > x(p_1, 1) \Rightarrow \alpha \in [\widetilde{\alpha}, \widehat{\alpha}]$.

Suppose $\Psi(p_1, 1) \in (0, 1)$, $x(p_1, 0) > x(p_1, 1)$, and $\alpha \in (0, \widetilde{\alpha}) \cup (\widehat{\alpha}, 1)$. Then $C(G, l) < C(F, l) < l < C(G, h) < C(F, h) < h$, and $\Psi(p_1, 1) \in (0, 1)$ only if $p_1 \in [C(G, l), C(F, h)]$. We will show that it reaches a contradiction for any $p_1 \in [C(G, l), C(F, h)]$.

If $p_1 \in [C(G, l), C(F, l))$, then only type (G, l) rejects p_1 and $x(p_1, 0) = 0 \leq x(p_1, 1)$.

If $p_1 \in (C(G, h), C(F, h)]$, then only type (F, h) accepts p_1 and $x(p_1, 1) = 1 \geq x(p_1, 0)$.

If $p_1 \in (C(F, l), C(G, h))$ and $\alpha < \widetilde{\alpha}$, then $\gamma(p_1, 1) = \overline{\gamma}(\alpha) < \gamma^*$ and $x(p_1, 1) = 1 \geq x(p_1, 0)$.

If $p_1 \in (C(F, l), C(G, h))$ and $\alpha > \widehat{\alpha}$, then $\gamma(p_1, 0) = \underline{\gamma}(\alpha) > \gamma^*$ and $x(p_1, 0) = 0 \leq x(p_1, 1)$.

Denote X as the probability for type (F, l) to reject $p_1 = C(F, l)$ and Y as the probability for type (G, h) to reject $p_1 = C(G, h)$.

If $p_1 = C(F, l)$,

$$\gamma(p_1, 0) = \frac{\alpha(1-q^G)}{\alpha(1-q^G) + (1-\alpha)X(1-q^F)} > \underline{\gamma}(\alpha)$$

and

$$\gamma(p_1, 1) = \frac{\alpha q^G}{\alpha q^G + (1-\alpha)q^F + (1-\alpha)(1-X)(1-q^F)} < \overline{\gamma}(\alpha).$$

So $\gamma(p_1, 1) < \overline{\gamma}(\alpha) < \gamma^*$ when $\alpha < \widetilde{\alpha}$ and $x(p_1, 1) = 1 \geq x(p_1, 0)$. $\gamma(p_1, 0) > \underline{\gamma}(\alpha) > \gamma^*$ when $\alpha > \widehat{\alpha}$ and $x(p_1, 0) = 0 \leq x(p_1, 1)$.

If $p_1 = C(G, h)$,

$$\gamma(p_1, 0) = \frac{\alpha Y q^G + \alpha(1-q^G)}{\alpha Y q^G + \alpha(1-q^G) + (1-\alpha)(1-q^F)} > \underline{\gamma}(\alpha)$$

and

$$\gamma(p_1, 1) = \frac{\alpha(1-Y)q^G}{\alpha(1-Y)q^G + (1-\alpha)q^F} < \overline{\gamma}(\alpha).$$

So $\gamma(p_1, 1) < \overline{\gamma}(\alpha) < \gamma^*$ when $\alpha < \widetilde{\alpha}$ and $x(p_1, 1) = 1 \geq x(p_1, 0)$. $\gamma(p_1, 0) > \underline{\gamma}(\alpha) > \gamma^*$ when $\alpha > \widehat{\alpha}$ and $x(p_1, 0) = 0 \leq x(p_1, 1)$.

Therefore, every case is contradictory to $x(p_1, 0) > x(p_1, 1)$.

(ii) It is directly derived from Lemma 3 and (i) of Lemma 4. □

Proof of Proposition 5. **Step 1:** It is the unique D_1 equilibrium strategy for the buyer to accept p_t if and only if $p_t \leq v_t$.

Suppose $x(p_1, 0) > x(p_1, 1)$. Then $\Psi(p_1, 0) = 1$ by Lemma 4 given $\alpha < \widetilde{\alpha}$. Since $x(p_1, 0) > x(p_1, 1)$, $C(F, h) = \max_{\theta_1}\{C(\theta_1)\} < h$. By Lemma 15, a PBE cannot pass criterion D_1 if $\Psi(p_1, 0) = 1$ and $p_1 < h$. Therefore, $p_1 \geq h$ and $\Psi(p_1, 0) = 1$, i.e., all types reject $p_1 \geq h$.

Suppose $x(p_1, 0) < x(p_1, 1)$. Then $\Psi(p_1, 1) = 1$ by Lemma 4 given $\alpha < \widetilde{\alpha}$. Since $x(p_1, 0) < x(p_1, 1)$, $C(F, l) = \min_{\theta_1}\{C(\theta_1)\} > l$. By Lemma 15, a PBE cannot pass criterion D_1 if $\Psi(p_1, 1) = 1$ and $p_1 > l$. Therefore, $p_1 \leq l$ and $\Psi(p_1, 1) = 1$, i.e., all types accept $p_1 \leq l$.

Suppose $x(p_1, 0) = x(p_1, 1)$. Then all buyer types accept p_1 if and only if $p_1 \leq v_1$, otherwise the seller offers $p_1 = v_1 - \epsilon$ and no equilibrium exists.

Therefore, combining three cases above, it is the unique equilibrium strategy for the buyer to accept p_1 if and only if $p_1 \leq v_1$.

Step 2: Given the buyer's strategy, the seller offers $p_1 = l$ or $p_1 = h$, and always offers $p_2 = l$ on the equilibrium path. The respective payoffs for the seller is:

$$\begin{cases} \pi(l) = l + \delta l; \\ \pi(h) = \alpha h + \delta l. \end{cases}$$

Since $\alpha < \widetilde{\alpha} < \gamma^*$, it is optimal to offer $p_1 = l$. □

Proof of Proposition 6. **Step 1:** It is the unique D_1 equilibrium strategy for the buyer to accept p_t if and only if $p_t \leq v_t$.

Suppose $x(p_1, 0) > x(p_1, 1)$. Then $\Psi(p_1, 1) = 1$ by Lemma 4 given $\alpha > \widehat{\alpha}$. Since $x(p_1, 0) > x(p_1, 1)$, $C(G, l) = \min_{\theta_1}\{C(\theta_1)\} < l$. By Lemma 15, a PBE can pass criterion D_1 if $\Psi(p_1, 1) = 1$ and $p_1 \leq l$. Therefore, $p_1 < C(G, l)$ and $\Psi(p_1, 1) = 1$, i.e., all types accept $p_1 < C(G, l)$.

Suppose $x(p_1, 0) < x(p_1, 1)$. Then $\Psi(p_1, 0) = 1$ by Lemma 4 given $\alpha > \widehat{\alpha}$. Since $x(p_1, 0) < x(p_1, 1)$, $C(G, h) = \max_{\theta_1}\{C(\theta_1)\} > h$. By Lemma 15, a PBE can pass criterion D_1 if $\Psi(p_1, 0) = 1$ and $p_1 \geq h$. Therefore, $p_1 > C(G, h)$ and $\Psi(p_1, 0) = 1$, i.e., all types reject $p_1 > C(G, h)$.

Suppose $x(p_1, 0) = x(p_1, 1)$. Then all buyer types accept p_1 if and only if $p_1 \leq v_1$, otherwise the seller offers $p_1 = v_1 - \epsilon$ and no equilibrium exists.

Therefore, combining three cases above, it is the unique equilibrium strategy for the buyer to accept p_1 if and only if $p_1 \leq v_1$.

Step 2: Given the buyer's strategy, the seller offers $p_1 = l$ or $p_1 = h$, and always offers $p_2 = h$ on the equilibrium path. The respective payoffs for the seller is:

$$\begin{cases} \pi(l) = l + \delta\alpha h; \\ \pi(h) = \alpha h + \delta\alpha h. \end{cases}$$

Since $\alpha > \widehat{\alpha} > \gamma^*$, it is optimal to offer $p_1 = h$. □

Proof of Lemma 7. **Step 1:** Suppose $x(p_1, 0) < x(p_1, 1)$. Then $l < C(F, l) < C(G, l) < h < C(F, h) < C(G, h)$. $\Psi(p_1, 1) = 0$ or 1 given $\widetilde{\alpha} < \alpha < \widehat{\alpha}$ by Lemma 4. Therefore, all buyer types accept $p_1 \leq C(F, l)$ when $\widetilde{\alpha} < \alpha < \gamma^*$, and all types reject $p_1 > C(G, h)$ when $\gamma^* < \alpha < \widehat{\alpha}$. By Lemma 15, if all types accept $p_1 \in (l, C(F, l)]$, the equilibrium cannot pass criterion D_1. So all types accept $p_1 \leq l$ when $\widetilde{\alpha} < \alpha < \gamma^*$ and reject $p_1 > C(G, h)$ when $\gamma^* < \alpha < \widehat{\alpha}$, under the assumption of $x(p_1, 0) < x(p_1, 1)$.

Step 2: Suppose $x(p_1, 0) = x(p_1, 1)$. Then $C(G, l) = C(F, l) = l < C(G, h) = C(F, h) = h$. Therefore, all buyer types accept $p_1 < l$ and all types reject $p_1 > h$. At $p_1 = l$ or h, see footnote 8. For $p_1 \in (l, h)$, types (F, l) and (G, l) reject p_1 and types (F, h) and (G, h) accept p_1, so $x(p_1, 0) = 1$ and $x(p_1, 1) = 0$ for $\widetilde{\alpha} < \alpha < \widehat{\alpha}$, which leads to a contradiction.

Step 3: Suppose $x(p_1, 0) > x(p_1, 1)$. Then $C(G, l) < C(F, l) < l < C(G, h) < C(F, h) < h$.

If $\Psi(p_1, 1) \in \{0, 1\}$, then all buyer types reject $p_1 > C(F, h)$ when $\widetilde{\alpha} < \alpha < \gamma^*$ and accept $p_1 < C(G, l)$ when $\gamma^* < \alpha < \widehat{\alpha}$. By Lemma 15, if all types reject $p_1 \in (C(F, h), h)$, the equilibrium cannot pass criterion D_1. So all buyer types reject $p_1 > h$ when $\widetilde{\alpha} < \alpha < \gamma^*$ and accept $p_1 < C(G, l)$ when $\gamma^* < \alpha < \widehat{\alpha}$.

Now consider the case $\Psi(p_1, 1) \in (0, 1)$. First consider pure strategy, i.e., suppose $x(p_1, 0) = 1$ and $x(p_1, 1) = 0$. By the proof of Lemma 16, it is not possible for any p_1 to separate a single type from other types. So pure strategy is possible to be adopted for $p_1 \in (C(F, l), C(G, h)] = (\underline{p}, \widetilde{p}]$. Types (F, l) and (G, l) reject p_1, and types (F, h) and (G, h) accept p_1.

Next consider mixed strategy, i.e., $0 < x(p_1, 0) - x(p_1, 1) < 1$. Then either $x(p_1, 0) = 1$ and $x(p_1, 1) \in (0, 1)$ or $x(p_1, 0) \in (0, 1)$ and $x(p_1, 1) = 0$, since the knife-edge condition $\alpha = \gamma^*$ is omitted. The former implies $\gamma(p_1, 0) < \gamma^*$ and $\gamma(p_1, 1) = \gamma^*$, and the latter implies $\gamma(p_1, 0) = \gamma^*$ and $\gamma(p_1, 1) > \gamma^*$. Therefore, $\gamma(p_1, 1) = \gamma^*$ when $\widetilde{\alpha} < \alpha < \gamma^*$, and $\gamma(p_1, 0) = \gamma^*$ when $\gamma^* < \alpha < \widehat{\alpha}$. From Lemma 16, it is not possible for type (G, l) or (F, h) to randomize, otherwise the seller can at least sometimes separate type (G, l) or (F, h) from other types. So only type (F, l) and (G, h) will play mixed strategy.

Case 1: When $\widetilde{\alpha} < \alpha < \gamma^*$, type (F, l) randomizes to reject p_1 with probability X^*, (G, l) rejects p_1, and (G, h) and (F, h) accept p_1. Then

$$\gamma(p_1, 1) = \frac{\alpha q^G}{\alpha q^G + (1 - \alpha) q^F + (1 - \alpha)(1 - X^*)(1 - q^F)} = \gamma^*.$$

Type (F, l) is indifferent from accepting and rejecting p_1, then

$$l - p_1 + \delta q^F x(p_1, 1)(h - l) = \delta q^F (h - l).$$

So type (F, l) rejects $p_1 \in (\underline{p}, l]$ with probability $X^* = 1 + \frac{q^F}{1 - q^F} - \frac{\alpha q^G (1 - \gamma^*)}{(1 - \alpha)(1 - q^F)\gamma^*}$, and the seller offers $x(p_1, 1) = 1 - \frac{l - p_1}{\delta q^F (h - l)}$, $x(p_1, 0) = 1$.

Case 2: When $\gamma^* < \alpha < \widehat{\alpha}$, type (F, l) randomizes to reject p_1 with probability X^{**}, (G, l) rejects p_1, and (G, h) and (F, h) accept p_1. Then

$$\gamma(p_1, 0) = \frac{\alpha(1 - q^G)}{\alpha(1 - q^G) + (1 - \alpha)X^{**}(1 - q^F)} = \gamma^*.$$

Type (F, l) is indifferent from accepting and rejecting p_1, then

$$l - p_1 = \delta q^F x(p_1, 0)(h - l).$$

So type (F, l) rejects $p_1 \in (\underline{p}, l]$ with probability $X^{**} = \frac{\alpha(1-q^G)(1-\gamma^*)}{(1-\alpha)(1-q^F)\gamma^*}$, and the seller offers $x(p_1, 0) = \frac{l-p_1}{\delta q^F(h-l)}$ and $x(p_1, 1) = 0$.

Case 3: When $\widetilde{\alpha} < \alpha < \gamma^*$, type (G, h) randomizes to reject p_1 with probability Y^*, (F, l) and (G, l) reject p_1, and (F, h) accepts p_1. Then

$$\gamma(p_1, 1) = \frac{\alpha(1 - Y^*)q^G}{\alpha(1 - Y^*)q^G + (1 - \alpha)q^F} = \gamma^*.$$

Type (G, h) is indifferent from accepting and rejecting p_1, then

$$h - p_1 + \delta q^G x(p_1, 1)(h - l) = \delta q^G (h - l).$$

So type (G, h) rejects $p_1 \in (\widetilde{p}, h]$ with probability $Y^* = 1 - \frac{(1-\alpha)q^F\gamma^*}{\alpha q^G(1-\gamma^*)}$, and the seller offers $x(p_1, 1) = 1 - \frac{h-p_1}{\delta q^G(h-l)}$ and $x(p_1, 0) = 1$.

Case 4: When $\gamma^* < \alpha < \widehat{\alpha}$, type (G, h) randomizes to reject p_1 with probability Y^{**}, (F, l) and (G, l) reject p_1, and (F, h) accepts p_1. Then

$$\gamma(p_1, 0) = \frac{\alpha Y^{**} q^G + \alpha(1 - q^G)}{\alpha Y^{**} q^G + \alpha(1 - q^G) + (1 - \alpha)(1 - q^F)} = \gamma^*.$$

Type (G, h) is indifferent from accepting and rejecting p_1, then

$$h - p_1 = \delta q^G x(p_1, 0)(h - l).$$

So type (G, h) rejects $p_1 \in (\widetilde{p}, h]$ with probability $Y^{**} = \frac{(1-\alpha)(1-q^F)\gamma^*}{\alpha q^G(1-\gamma^*)} - \frac{1-q^G}{q^G}$ and the seller offers $x(p_1, 0) = \frac{h-p_1}{\delta q^G(h-l)}$ and $x(p_1, 1) = 0$.

Lemma 7 comes from the combination of three steps. □

Proof of Proposition 8. **Step 1:** $U_2 > \max\{U_4, U_5\}$ for $\widetilde{\alpha} < \alpha < \gamma^*$.

$$
\begin{aligned}
& U_4 - U_2 \\
= \; & \delta(1-\alpha)q^F l(q^G - 1) + \delta\alpha q^G l(q^G - 1) \\
& +\delta(1-\alpha)q^F h(q^F - q^G) + [\alpha q^G h + (1-\alpha)q^F h - l] \\
< \; & 0
\end{aligned}
$$

Each item on the right hand side of the equation is negative for $\widetilde{\alpha} < \alpha < \gamma^*$.

By plugging Y^* into the definition of U_5, $U_5 = \frac{1-\alpha}{1-\gamma^*} q^F h + \delta l$, which is decreasing in α. So

$$
\begin{aligned}
& \max_{\alpha \in (\widetilde{\alpha}, \gamma^*)} (U_5 - U_2) \\
< \; & \frac{1-\widetilde{\alpha}}{1-\gamma^*} q^F h - l \\
= \; & \frac{h}{q^G + q^F - l/h} (l/h - q^G)(l/h - q^F) < 0.
\end{aligned}
$$

Step 2: Assume all buyer types accept $p_1 \in [\underline{p}, l]$, then $p_1 = l$ is the equilibrium price given that $U_2 > \max\{U_4, U_5\}$.

(i) For an arbitrary $p_1^* \in [\underline{p}, l]$, assume all buyer types accept $p_1 \in [\underline{p}, p_1^*]$, type (F, l) and (G, l) reject $p_1 \in (p_1^*, l]$, and type (F, h) and (G, h) accept $p_1 \in (p_1^*, l]$. Then p_1^* is the optimal p_1 given $U_1 > \max\{U_4, U_5\}$.

(ii) Since $U_1 = \underline{p} + \delta l < \max\{U_4, U_5\} < U_2 = l + \delta l$, there exists $p' \in (\underline{p}, l)$ such that $u(p') = p' + \delta l = \max\{U_4, U_5\}$. For an arbitrary $p_1^* \in [p', l]$, assume all buyer types accept $p_1 \in [\underline{p}, p_1^*]$, type (F, l) and (G, l) reject $p_1 \in (p_1^*, l]$, and type (F, h) and (G, h) accept $p_1 \in (p_1^*, l]$. Then p_1^* is the optimal p_1 given $u(p') = \max\{U_4, U_5\}$. □

Proof of Proposition 9. (i) By the definition of U_3, U_4, and U_5, they are the possibly highest payoffs when the buyer adopts any semi-separating equilibrium strategy. Since the lowest payoff from a pooling offer, U_1, is greater than $\max\{U_3, U_4, U_5\}$, there is no semi-separating equilibrium.

(ii) For all $p_1 \in (\underline{p}, l]$, the buyer can adopt two semi-separating equilibrium strategies: 1) types with $v_1 = l$ reject p_1 and types with $v_1 = h$ accept p_1, or 2) types with $v_1 = h$ accept p_1, type (G, l) rejects p_1 and type (F, l) randomizes. If the first strategy is adopted at $p_1 \in (\underline{p}, l]$, the seller's payoff by offering p_1 is dominated by U_4. If the second strategy is adopted, the payoff is dominated by $U_3 < \max\{U_4, U_5\}$. Suppose the buyer adopts either strategy, given $U_1 < \max\{U_4, U_5\}$, a semi-separating equilibrium exists and the path is unique, with $p_1 = \widetilde{p}$ or $p_1 = h$, depending on whether U_4 or U_5 is larger.

(iii) Define $U(p_1, X^*) = [\alpha q^G + (1-\alpha)q^F + (1-\alpha)(1-q^F)(1-X^*)]p_1 + \delta l$, which is increasing in p_1. First suppose $\max\{U_4, U_5\} < U_1 < U_3$. By definition $U(\underline{p}, X^*) < U_1 < U_3$. Therefore, there exists $p'' \in (\underline{p}, l)$ such that $U(p'', X^*) = U_1$. For any arbitrary $p_1^* \in [p'', l]$, assume the buyer uses the second strategy for $p_1^* \leq p''$ and uses the first strategy described in part (ii) for $p_1^* > p''$, then $p_1^* \in [p'', l]$ is the optimal p_1.

Then suppose $U_1 < \max\{U_4, U_5\} < U_3$. Since $U(\underline{p}, X^*) < U_1 < U_3$, $U(\underline{p}, X^*) < \max\{U_4, U_5\} < U_3$. Define $p'' \in (\underline{p}, l)$ such that $U(p'', X^*) = \max\{U_4, U_5\}$. If for any arbitrary $p_1^* \in [p'', l]$, the buyer uses the second strategy described in part (ii) for $p_1^* \leq p''$ and uses the first strategy for $p_1^* > p''$, then $p_1^* \in [p'', l]$ is the optimal p_1. If for any $p_1 \in (\underline{p}, l]$, the buyer uses the first strategy, then $p_1 = \widetilde{p}$ or $p_1 = h$ is optimal, depending on whether U_4 or U_5 is larger. □

Proof of Proposition 12. **Step 1:** The equilibrium in the two-period version of Hart and Tirole's (1988) rental model is as follows. In period 2, both types accept p_2 if and only if $p_2 \leq v_2$ and reject p_2 otherwise. In the first period, the l-type buyer accepts p_1 if and only if $p_1 \leq l$ and reject p_1 otherwise. If $\mu < l/h$, the h-type buyer accepts $p_1 \leq h - \delta(h-l)$ and reject $p_1 > h - \delta(h-l)$ in the first period. If $\mu > l/h$, the h-type buyer accepts $p_1 \leq h - \delta(h-l)$, randomizes to accept $p_1 \in (h - \delta(h-l), h]$ with probability $y^* = \frac{\mu h - l}{\mu(h-l)}$, and reject $p_1 > h$ in the first period. Therefore, if $\mu < l/h$, the seller offers $p_1 = p_2 = l$; if $l/h < \mu < \frac{hl + \delta l(h-l)}{hl + \delta h(h-l)}$, the seller offers $p_1 = h - \delta(h-l)$, $p_2 = h$ if p_1 is accepted, and $p_2 = l$ if p_1 is rejected; if $\mu > \frac{hl + \delta l(h-l)}{hl + \delta h(h-l)}$, the seller offers $p_1 = p_2 = h$. The corresponding revenues are as follows.

$$\pi = \begin{cases} l + \delta l, & \text{if } \mu < l/h; \\ \mu[h - \delta(h-l)] + \delta\mu h + \delta(1-\mu)l = \mu h + \delta l, & \text{if } l/h < \mu < \frac{hl+\delta l(h-l)}{hl+\delta h(h-l)}; \\ \mu y^* h + \delta\mu h = \frac{\mu h^2 - hl + \delta\mu h^2 - \delta\mu hl}{h-l} & \text{if } \mu > \frac{hl+\delta l(h-l)}{hl+\delta h(h-l)}. \end{cases}$$

This part of proof is available upon request.

Step 2: Compare the revenue in our model with that in Step 1, assuming that $\mu = \alpha q^G + (1-\alpha)q^F$.

(i) For an optimistic seller ($\alpha > \widehat{\alpha}$), there is a unique equilibrium outcome as shown in Proposition 2, and the seller's revenue in our model is

$$\begin{aligned} & (\alpha q^G + (1-\alpha)q^F)h + \delta(\alpha q^G + (1-\alpha)q^F)h \\ = \; & \mu h + \delta\mu h > \mu y^* h + \delta\mu h. \end{aligned}$$

So the seller's revenue in our model is higher than in Hart and Tirole's (1988).

(ii) For a moderately optimistic seller ($\gamma^* < \alpha < \widehat{\alpha}$), the corresponding μ is greater than l/h. Denote $W_1 = \mu h + \delta l$ and $W_2 = \frac{\mu h^2 - hl + \delta\mu h^2 - \delta\mu hl}{h-l}$. Then it suffices to compare the potential optimal revenues W_1 and W_2 in step 1 with the optimal revenue in our model. Our proof consists of the following results.

Result 1: $W_1 > V_2$ for $\gamma^* < \alpha < \widehat{\alpha}$.

$$\begin{aligned} W_1 - V_2 &= (\mu h + \delta l) - \{l + \delta[\alpha q^G + (1-\alpha)q^F]h\} \\ &= (\mu h + \delta l) - (l + \delta\mu h) = (1-\delta)(\mu h - l) > 0 \end{aligned}$$

The inequality holds since $\alpha > \gamma^*$ and $\mu > l/h$.

Result 2: $W_1 > V_4$.

$$\begin{aligned} W_1 - V_4 &= \{\mu[h - \delta(h-l)] + \delta\mu h + \delta(1-\mu)l\} \\ &\quad - \{\mu[h - \delta q^G(h-l)] + \delta[\alpha(q^G)^2 + (1-\alpha)(q^F)^2]h + \delta(1-\mu)l\} \\ &= \delta\{\mu - [\alpha(q^G)^2 + (1-\alpha)(q^F)^2]\}h - \delta\mu(1-q^G)(h-l) \\ &= \delta[\alpha q^G(1-q^G) + (1-\alpha)q^F(1-q^F)]h - \delta\mu(1-q^G)(h-l) \\ &= \delta\mu(1-q^G)h + \delta(1-\alpha)q^F(q^G-q^F)h - \delta\mu(1-q^G)(h-l) \\ &= \delta(1-\alpha)q^F(q^G-q^F)h + \delta\mu(1-q^G)l > 0. \end{aligned}$$

Result 3: $V_5 > W_2$ if $(1-q^F)(1-l/h) < q^G - l/h$.

$$\begin{aligned} V_5 - W_2 &= [\alpha q^G(1-Y^{**}) + (1-\alpha)q^F]h - \mu y^* h \\ &= \alpha(q^G - q^F)(\frac{1-q^F}{q^G - l/h} - \frac{1}{1-l/h})h + [q^F - \frac{(1-q^F)(l/h - q^F)}{q^G - l/h} + \frac{l/h - q^F}{1-l/h}]h \end{aligned}$$

If $(1-q^F)(1-l/h) < q^G - l/h$, then $V_5 > W_2$ when

$$\alpha < \frac{l/h - q^F}{q^G - q^F} - \frac{q^F}{q^G - q^F}\frac{1}{\frac{1-q^F}{q^G-l/h} - \frac{1}{1-l/h}}.$$

The RHS of the inequality is decreasing in q^G and converges to 1 when $q^G \to 1$, therefore the RHS of the inequality is greater than 1, so the inequality always holds when $(1-q^F)(1-l/h) < q^G - l/h$.

Result 4: There exists $\overline{\alpha} \in (\gamma^*, \widehat{\alpha})$ such that for $\alpha \in (\overline{\alpha}, \widehat{\alpha})$ $W_2 > W_1$, if $q^G > (1-q^F)(1-l/h) + l/h$ and $q^F < \frac{\delta(l/h)(1-l/h)}{l/h+\delta(1-l/h)}$.

$$\begin{aligned} W_2 &> W_1 \Rightarrow \mu > \frac{l/h + \delta(l/h)(1-l/h)}{l/h + \delta(1-l/h)} \\ &\Rightarrow \alpha > \frac{1}{q^G - q^F}[\frac{l/h + \delta(l/h)(1-l/h)}{l/h + \delta(1-l/h)} - q^F] \equiv \overline{\alpha} \end{aligned}$$

It is easy to show that $\overline{\alpha} > \gamma^* = \frac{l/h - q^F}{q^G - q^F}$. Next we need to show the conditions under which $\overline{\alpha} < \widehat{\alpha} = \frac{1-q^F}{(1-q^F)-(q^G-l/h)} * \frac{l/h-q^F}{q^G-q^F}$.

$$\begin{aligned} q^G - l/h &> (1-q^F)(1-l/h) \\ &\Leftrightarrow (1-q^F) - (q^G - l/h) < (1-q^F) - (1-q^F)(1-l/h) \\ &\Leftrightarrow (1-q^F)(l/h) > (1-q^F) - (q^G - l/h) \\ &\Leftrightarrow \frac{l/h - q^F}{l/h} < \frac{(1-q^F)(l/h - q^F)}{(1-q^F) - (q^G - l/h)} \end{aligned}$$

To show $\overline{\alpha} < \widehat{\alpha}$, it is sufficient to show that

$$\frac{l/h + \delta(l/h)(1-l/h)}{l/h + \delta(1-l/h)} - q^F < \frac{l/h - q^F}{l/h},$$

which is satisfied when $q^F < \frac{\delta(l/h)(1-l/h)}{l/h+\delta(1-l/h)}$.

Combining Result 1, 2, 3, and 4, we have shown that, if $q^G > (1-q^F)(1-l/h) + l/h$ and $q^F < \frac{\delta(l/h)(1-l/h)}{l/h+\delta(1-l/h)}$, there exists $\overline{\alpha} \in (\gamma^*, \widehat{\alpha})$ such that for $\alpha \in (\overline{\alpha}, \widehat{\alpha})$ $V_5 > W_2 > W_1 >$

$\max\{V_2, V_4\}$. Therefore, V_5 is the optimal revenue in our model and it is higher than the optimal revenue in the two-period version of Hart and Tirole (1988).

(iii) For a pessimistic seller or moderately pessimistic seller ($\alpha < \gamma^*$), there always exists a pooling equilibrium in which the seller offers $p_1 = p_2 = l$ and all buyer types accept the offer as shown in Proposition 1 and 3. This equilibrium yields revenue $l + \delta l$, which is the same as in Hart and Tirole's (1988). □

Proof of Lemma 15. **Part 1:** Suppose $\Psi(p_1, 1) = 1$ in the continuation game following $p_1 > l$. Then $x(p_1,1) > x(p_1, 0)$ and $x(p_1,1) = 1$ without considering the knife-edge case that $\alpha = \gamma^*$. Since $\max\{x(p_1,1) - x(p_1, 0)\} = 1$ and all types accept p_1, $p_1 \leq \min\limits_{(d,v_1)}\{v_1 + \delta q^d(h-l)\} = l + \delta q^F(h-l)$ by the definition of cutoff value.

Apply Definition 1 in the case that $a_1 = 0$ is the out-of-equilibrium message and form the sets D_{θ_1} and $D^0_{\theta_1}$ for each buyer type θ_1. So $D_{\theta_1} = \{\phi : \phi > x(p_1, 1) + \frac{v_1 - p_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$ and $D^0_{\theta_1} = \{\phi : \phi = x(p_1, 1) + \frac{v_1 - p_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$. Therefore, for $x(p_1,1) = 1$ and $p_1 \in (l, l + \delta q^F(h-l)]$, $D_{\theta_1} \cup D^0_{\theta_1} \subsetneq D_{(F,l)}$ for all $\theta_1 \neq (F, l)$. All the combinations $(\theta_1, a_1 = 0)$ with $\theta_1 \neq (F, l)$ are pruned from the game. Given the seller's belief is restricted on type (F, l) after rejection, $x(p_1, 0) = 1$ and it is contradictory to $x(p_1,1) > x(p_1, 0)$. So the equilibrium fails to pass criterion D_1.

Part 2: Suppose $\Psi(p_1, 1) = 1$ in the continuation game following $p_1 \leq l$. From Part 1, $D_{\theta_1} = \{\phi : \phi > x(p_1, 1) + \frac{v_1 - p_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$ and $D^0_{\theta_1} = \{\phi : \phi = x(p_1, 1) + \frac{v_1 - p_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$.

If $p_1 = l$ and $\alpha < \gamma^*$, $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for $\theta_1 \in \{(F, h), (G, h)\}$ and $D_{\theta_1} \cup D^0_{\theta_1} = \{1\}$ for $\theta_1 \in \{(F, l), (G, l)\}$.

If $p_1 = l$ and $\alpha > \gamma^*$, $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for $\theta_1 \in \{(F, h), (G, h)\}$ and $D_{\theta_1} \cup D^0_{\theta_1} = [0, 1]$ for $\theta_1 \in \{(F, l), (G, l)\}$.

If $p_1 < l$ and $\alpha < \gamma^*$, then $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for all buyer types θ_1.

If $p_1 < l$ and $\alpha > \gamma^*$, then either $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for all buyer types θ_1 or $D_{\theta_1} \cup D^0_{\theta_1} \subsetneq D_{(G,l)}$. If the latter happens, the seller's belief is restricted on type (G, l) after rejection and she offers $x(p_1, 0) = 0$. It is still sequential rational for all buyer types θ_1 to accept $p_1 < l$ given $x(p_1, 1) = x(p_1, 0) = 0$.

In all the cases above, the equilibrium passes criterion D_1.

Part 3: Suppose $\Psi(p_1, 0) = 1$ in the continuation game following $p_1 < h$. Then $x(p_{1,}0) > x(p_1, 1)$ and $x(p_{1,}0) = 1$ without considering the knife-edge case that $\alpha = \gamma^*$. Since $\max\{x(p_{1,}0) - x(p_1, 1)\} = 1$ and all types reject p_1, $p_1 \geq \max_{(d,v_1)}\{v_1 - \delta q^d(h-l)\} = h - \delta q^F(h-l)$ by the definition of cutoff value.

Apply Definition 1 in the case that $a_1 = 1$ is the out-of-equilibrium message. So $D_{\theta_1} = \{\phi : \phi > x(p_1, 0) + \frac{p_1 - v_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$ and $D^0_{\theta_1} = \{\phi : \phi = x(p_1, 0) + \frac{p_1 - v_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$. Then for $x(p_{1,}0) = 1$ and $p_1 \in [h - \delta q^F(h - l), h)$, $D_{\theta_1} \cup D^0_{\theta_1} \subsetneq D_{(F,h)}$ for all $\theta_1 \neq (F, h)$. All the combinations $(\theta_1, a_1 = 1)$ with $\theta_1 \neq (F, h)$ are pruned from the game. Given the seller's belief is restricted on type (F, h) after acceptance, $x(p_{1,}1) = 1$ and it is contradictory to $x(p_{1,}0) > x(p_1, 1)$. So the equilibrium fails to pass Criterion D_1.

Part 4: Suppose $\Psi(p_1, 0) = 1$ in the continuation game following $p_1 \geq h$. From Part 3, $D_{\theta_1} = \{\phi : \phi > x(p_1, 0) + \frac{p_1 - v_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$ and $D^0_{\theta_1} = \{\phi : \phi = x(p_1, 0) + \frac{p_1 - v_1}{\delta q^d(h-l)}, \phi \in [0, 1]\}$.

If $p_1 = h$ and $\alpha < \gamma^*$, $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for $\theta_1 \in \{(F, l), (G, l)\}$ and $D_{\theta_1} \cup D^0_{\theta_1} = \{1\}$ for $\theta_1 \in \{(F, h), (G, h)\}$.

If $p_1 = h$ and $\alpha > \gamma^*$, $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for $\theta_1 \in \{(F, l), (G, l)\}$ and $D_{\theta_1} \cup D^0_{\theta_1} = [0, 1]$ for $\theta_1 \in \{(F, h), (G, h)\}$.

If $p_1 > h$ and $\alpha < \gamma^*$, then $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for all buyer types θ_1.

If $p_1 > h$ and $\alpha > \gamma^*$, then either $D_{\theta_1} \cup D^0_{\theta_1} = \emptyset$ for all buyer types θ_1 or $D_{\theta_1} \cup D^0_{\theta_1} \subsetneq D_{(G,h)}$. If the latter case happens, the seller's belief is restricted on type (G, h) after acceptance and $x(p_1, 1) = 0$. Then it is still sequential rational for all buyer types θ_1 to reject $p_1 > h$ given $x(p_1, 1) = x(p_1, 0) = 0$.

In all the cases above, the equilibrium passes criterion D_1. □

Lemma 16. *If $x(p_1, 0) \neq x(p_1, 1)$ and $\Psi(p_1, a_1) > 0$ for a given p_1 in a PBE, the seller's posterior belief $\gamma(p_1, a_1) \neq 0$ or 1.*

Proof of Lemma 16. **Step 1:** Suppose $x(p_1, 0) > x(p_1, 1)$. Then $C(G, l) < C(F, l) < l < C(G, h) < C(F, h) < h$.

If $p_1 < C(G, l)$, all types accept p_1, and $\gamma(p_1, 1) = \alpha$.

If $p_1 > C(F,h)$, all types reject p_1, and $\gamma(p_1,0) = \alpha$.

If $p_1 \in (C(G,l), C(F,l))$, only type (G,l) rejects p_1 and $x(p_1,0) = 0$, so it contradicts with $x(p_1,0) > x(p_1,1)$.

If $p_1 = C(G,l)$, either all types accept p_1 or $x(p_1,0) = 0$ and it contradicts with $x(p_1,0) > x(p_1,1)$.

If $p_1 \in (C(G,h), C(F,h))$, only type (F,h) accepts p_1 and $x(p_1,1) = 1$, so it contradicts with $x(p_1,0) > x(p_1,1)$.

If $p_1 = C(F,h)$, either all types reject p_1 or $x(p_1,1) = 1$ and it contradicts with $x(p_1,0) > x(p_1,1)$.

If $p_1 \in (C(F,l), C(G,h))$, both type (F,l) and (G,l) reject p_1 and both type (F,h) and (G,h) accept p_1, and $\gamma(p_1,0) = \underline{\gamma}(\alpha)$ and $\gamma(p_1,1) = \overline{\gamma}(\alpha)$.

If $p_1 = C(F,l)$ and $\sigma^B((F,l),p_1) = 1$, then $x(p_1,0) = 0$ and it contradicts with $x(p_1,0) > x(p_1,1)$. If $p_1 = C(G,h)$ and $\sigma^B((G,h),p_1) = 0$, then $x(p_1,1) = 1$ and it contradicts with $x(p_1,0) > x(p_1,1)$. So more than one buyer types who are from different distributions accept or reject $p_1 = C(F,l)$ or $p_1 = C(G,h)$.

Therefore, in any case when $x(p_1,0) > x(p_1,1)$ in a PBE, the seller's posterior belief $\gamma(p_1,a_1) \neq 0$ or 1, if history (p_1,a_1) is reached with a positive probability in the continuation game following p_1.

Step 2: Suppose $x(p_1,0) < x(p_1,1)$. Then $l < C(F,l) < C(G,l) < h < C(F,h) < C(G,h)$.

If $p_1 < C(F,l)$, all types accept p_1, and $\gamma(p_1,1) = \alpha$.

If $p_1 > C(G,h)$, all types reject p_1, and $\gamma(p_1,0) = \alpha$.

If $p_1 \in (C(F,l), C(G,l))$, only type (F,l) rejects p_1 and $x(p_1,0) = 1$, so it contradicts with $x(p_1,0) < x(p_1,1)$.

If $p_1 = C(F,l)$, either all types accept p_1 or $x(p_1,0) = 1$ and it contradicts with $x(p_1,0) < x(p_1,1)$.

If $p_1 \in (C(F,h), C(G,h))$, only type (G,h) accepts p_1 and $x(p_1,1) = 0$, so it contradicts with $x(p_1,0) < x(p_1,1)$.

If $p_1 = C(G,h)$, either all types reject p_1 or $x(p_1,1) = 0$ and it contradicts with $x(p_1,0) < x(p_1,1)$.

If $p_1 \in (C(G,l), C(F,h))$, both type (F,l) and (G,l) reject p_1 and both type (F,h) and (G,h) accept p_1, and $\gamma(p_1,0) = \underline{\gamma}(\alpha)$ and $\gamma(p_1,1) = \overline{\gamma}(\alpha)$.

If $p_1 = C(G,l)$ and $\sigma^B((G,l),p_1) = 1$, then $x(p_1,0) = 1$ and it contradicts with $x(p_1,0) < x(p_1,1)$. If $p_1 = C(F,h)$ and $\sigma^B((F,h),p_1) = 0$, then $x(p_1,1) = 0$ and it contradicts with $x(p_1,0) < x(p_1,1)$. So more than one buyer types who are from different distributions accept or reject $p_1 = C(G,l)$ or $p_1 = C(F,h)$.

Therefore, in any case when $x(p_1,0) < x(p_1,1)$ in a PBE, the seller's posterior belief $\gamma(p_1,a_1) \neq 0$ or 1, if history (p_1,a_1) is reached with a positive probability in the continuation game following p_1. □

Lemma 17. *If $x(p_1,0) = x(p_1,1)$ and $\Psi(p_1,a_1) > 0$ for a given $p_1 \notin \{l,h\}$ in a PBE, the seller's posterior belief $\gamma(p_1,a_1) \neq 0$ or 1.*

Proof of Lemma 17. Suppose $x(p_1,0) = x(p_1,1)$. Then $C(F,l) = C(G,l) = l < C(F,h) = C(G,h) = h$.

If $p_1 < l$, all types accept p_1, and $\gamma(p_1,1) = \alpha$.

If $p_1 > h$, all types reject p_1, and $\gamma(p_1,0) = \alpha$.

If $l < p_1 < h$, both type (F,l) and (G,l) reject p_1 and both type (F,h) and (G,h) accept p_1, and $\gamma(p_1,0) = \underline{\gamma}(\alpha)$ and $\gamma(p_1,1) = \overline{\gamma}(\alpha)$.

Therefore, in any case when $x(p_1,0) = x(p_1,1)$ for $p_1 \notin \{l,h\}$ in a PBE, the seller's posterior belief $\gamma(p_1,a_1) \neq 0$ or 1, if history (p_1,a_1) is reached with a positive probability in the continuation game following p_1. □

APPENDIX C

INSTRUCTIONS FOR CHAPTER 3 (MIN_NO SESSIONS)

Overview

This is an experiment in decision-making. The department of economics has provided funds for this research. During the course of this experiment, you will be called upon to make a series of decisions. If you follow these instructions carefully and make good decisions, you can earn a considerable amount of money which will be paid to you in cash at the end of the experiment. We ask that you not talk with one another for the duration of the experiment.

Specifics

The experiment is divided into a series of sequences. A sequence will consist of an indefinite number of rounds. At the beginning of each sequence, you will be randomly assigned as a First Mover or a Second Mover. Your role will appear on your computer screen and will not change during the sequence. At the beginning of each round you will be randomly paired with another person who is assigned to the other role from your own. That is, if you are a First Mover (Second Mover), in each period you will be randomly paired with a Second Mover (First Mover) with all possible pairings being equally likely.

In each round, you and your paired player will play the game described in the following graph. First, the First Mover chooses between A and B. If the First Mover chooses A, the round is over. The First Mover receives 35 points and the Second Mover receives 0 points. If the First Mover chooses B, then the Second Mover must make a choice between C and D. If the Second Mover chooses C, then the First Mover receives 0 points and the Second Mover receives 100 points. If the Second Mover chooses D, then the First Mover receives 45 points and the Second Mover receives 55 points. Following the first round of a sequence, the First Mover will be told the decision that his/her paired Second Mover has made in the last round, if that player was able to choose between C and D. The First Mover never knows the identity of his/her paired Second Mover.

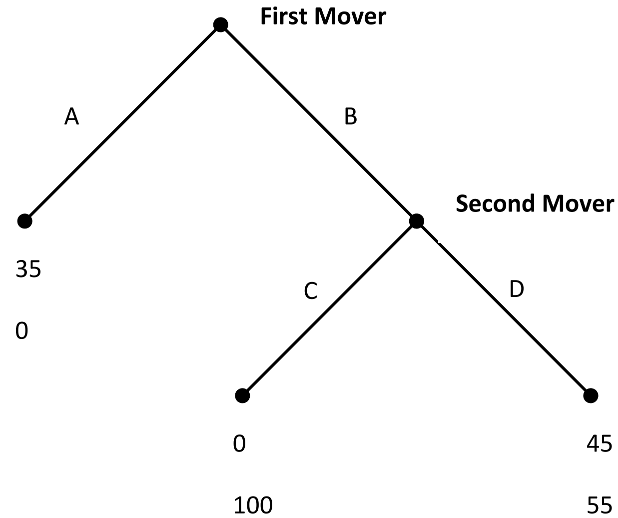

Figure C.1: Decisions and Earnings (in Points)

To complete your choice in each round, click on the decision button and then the OK button. The Second Movers need to wait for the First Movers to make a choice between A and B before making their own choice. Then the Second Mover will be told that the round is over (if the First Mover chooses A), or will be asked to make a choice between C and D (if the First Mover chooses B).

The computer program will record your choice and the choice made by the player paired with you in this round. After all players have made their choices, the results of the round will appear on your screen. You will be reminded of your own choice and will be shown the choice of your match, as well as the payoff you have earned for the round. Record the results of the round on your RECORD SHEET under the appropriate headings.

Immediately after you have received information on your choice and the choice of the player with whom you are randomly paired for the round, a ten-sided die with numbers from 0 to 9 on the sides will be thrown by each of you, one by one, to determine whether the sequence will continue or not. The experimenter will announce loudly the result of each die roll. If a number from 0 to 7 appears, the sequence will continue into next round. If an 8 or 9 appears, the sequence ends. Therefore, after each round there is 80% chance that you will play another round and 20% chance that the sequence will end.

Suppose that a number less than 8 has appeared. Then you will play the same game as in the previous round, but with an individual selected at random whose role is different from yours. You record the outcome and your earnings for the round. Then another throw is made with the same die to decide whether the sequence continues for another round.

If an 8 or 9 appears, the sequence ends. The experimenter will announce whether or not a new sequence will be played. If a new sequence is to be played then you will be randomly reassigned as a First Mover or a Second Mover. The new sequence will then be played as described above.

If the experiment does not end within 2 hours, you will be invited to continue the experiment in the next several days.

Earnings

Each point you earn is worth 0.5 cent ($0.005). Therefore, the more points you earn the more money you earn. You will be paid your earnings from all the rounds in cash, and in private, at the end of today's experiment as well as $5 show-up fee.

Final Comments

First, do not discuss your decisions or your results with anyone at any time during the experiment.

Second, your ID# is private. Do not reveal it to anyone.

Third, since there is 80% chance that at the end of a round the sequence will continue, you can expect, on average, to play 5 rounds in a given sequence. However, since the stopping decision is made randomly, some sequences may be much longer than 5 rounds and others may be much shorter.

Fourth, your role as a First Mover or a Second Mover will be randomly assigned when a new sequence begins. Your role will not change for the duration of that sequence.

Finally, remember that after each round of a sequence you will be matched randomly with a player whose role is different from yours. Therefore, the probability of you being matched with the same individual in two consecutive rounds of a game is 1/3 since there are 3 First Movers and 3 Second Movers in the room.

Questions?

Now is the time for questions. Does anyone have any questions?

Quiz

If there are no more questions, please finish the quiz. Your answers to this quiz will not affect your earnings. The purpose of the quiz is to help you understand the instruction better. After everyone has completed the quiz the answers will be reviewed.

APPENDIX D

FREQUENCIES EXCLUDING FIRST ROUNDS

Table D.1: Frequency of Invest

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	0.500	0.867	0.978
No_Info_Cost2	0.639	0.839	0.782
No_Info_Cost3	0.642	0.833	0.841
Info_No_Cost1	0.863	0.627	0.907
Info_No_Cost2	0.790	0.094	0.296
Info_No_Cost3	0.778	0.417	0.432

Table D.2: Frequency of Return-given-Invest

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	0.815	0.974	0.989
No_Info_Cost2	0.696	0.877	0.794
No_Info_Cost3	0.654	0.933	0.879
Info_No_Cost1	0.980	0.734	1.000
Info_No_Cost2	0.906	0.455	0.750
Info_No_Cost3	0.889	0.433	0.600

Table D.3: Frequency of Invest&Return

	1st treatment	2nd treatment	3rd treatment
No_Info_Cost1	0.407	0.844	0.968
No_Info_Cost2	0.444	0.736	0.621
No_Info_Cost3	0.420	0.778	0.739
Info_No_Cost1	0.846	0.461	0.907
Info_No_Cost2	0.716	0.043	0.222
Info_No_Cost3	0.691	0.181	0.259

APPENDIX E

INSTRUCTIONS FOR CHAPTER 4

Claim Check_____________

Welcome

Thank you for agreeing to participate in our study on decision making. There are two parts of the study today. In the first part you are asked to make six decisions and in the second part you are asked to fill out a survey. When you have completed your decisions we will randomly select one of your six decisions for payment. Your total payment from the study will be the sum of the payment that results from your decision and $5 for showing up to the study. The entire study should take about an hour, and at the end you will be paid privately and in cash. A research foundation has provided the funds for this study.

We ask that you do not speak to each other or make comments, except to ask questions about the procedures of the study. We also ask that you do not discuss the procedures of the study with others outside this room.

Your Identity

Your identity is secret. You will never be asked to reveal it to anyone during the course of the study. Your name will never be associated with your decisions or with your answers on the survey. Neither the assistants nor the other participants will be able to link you to any of the decisions you make. In order to keep your decisions private, please do not reveal your choices to any other participant.

Claim Check

Attached to the top of this page is a yellow piece of paper with a number on it. This is your Claim Check. Each participant has a different number. We use claim checks to maintain secrecy about your decisions, payment, and identity. You will present your Claim Check to an assistant at the end of the study to receive your cash payment.

Please remove your claim check now, and put it in a safe place.

Decision Tasks

For the decision tasks you will be paired with a child in Southwestern Pennsylvania (Allegheny, Washington, Greene, and Fayette Counties). The child is between 1 and 12 years old, and the child's family home has suffered extensive fire damage. Most or all of the family's possessions have been lost. For each of your decisions you will be given an amount of money which you will be asked to allocate between the child and yourself. The money allocated towards the child will be spent on children's books. These books will be distributed to the child by the American Red Cross of Southwestern Pennsylvania, immediately after the child has been affected by a severe fire.

As soon as a fire is reported in Southwestern Pennsylvania, the American Red Cross is contacted and volunteers are dispatched to the site. They help the affected families find temporary shelter, provide them with clothing, a meal, and give them a comfort bag with essential toiletries. Each day an average of one family in Southwestern Pennsylvania experiences a severe fire. These families depend on the American Red Cross for emergency help to cope with the sudden loss of their home and belongings. Unfortunately the American Red Cross only has funds to provide these families with the bare essentials, and they do not provide any "comfort" items for the children of the affected families.

For the study today we have joined the American Red Cross of Southwestern PA to collect funds to buy books for the affected children. In each of the six decisions you will be given an amount of money which you are asked to allocate between the child you are paired with and yourself. In addition the foundation has agreed to donate a fixed amount of money towards the child independent of your allocation. Thus the total amount to be spent on the child is the sum of the foundation's fixed donation and the allocation you make to the child. The amount of money that you can allocate between the child and you, as well as the foundation's fixed donation to the child, will vary across the six decisions.

The American Red Cross will use the funds collected from your allocation and that of the foundation to purchase the child books. Each participant in this study is paired with a different child. If you choose not to allocate any funds to the child, then the money to be spent on the child will be limited to the research foundation's fixed donation. Only you have the opportunity to allocate additional funds to the child. Neither the American Red Cross nor any other donors provide books to the child. Your decision alone determines how much will be spent on the child.

In explaining why the American Red Cross is seeking funds for books, their Emergency Preparedness Coordinator Sandi Wraith states "Children's needs are often overlooked in the immediate aftermath of a disaster because everyone is concerned primarily with putting the fire out, reaching safety, and finding shelter, food and clothing...just the basics of life. So many times, I've seen children just sitting on the curb with no one to talk to about what's happening...for this reason I've found trauma recovery experts in the community to work with us to train our volunteer responders in how to address children's needs at the scene of a disaster.......being able to give the children fun and distracting books will provide a great bridge for our volunteers to connect with kids and get them talking about what they've experienced."

Once we are ready to proceed to the decisions, you will be given a decision folder and a calculator. The folder contains a decision task with six decisions on it, and an envelope. For each decision you will have to enter your preferred allocation. If you wish to receive a receipt from the American Red Cross for your allocation to the child, you will need to fill out the acknowledgment form. Note however that by doing so you will relinquish your anonymity. If you wish to remain anonymous, leave the acknowledgment form blank. When you have completed the decision form please place it in the envelope along with the acknowledgment form, instructions and the calculator.

When we have collected all the envelopes we will draw a number between 1 and 6 to determine which one of the decisions counts for payment. Since one decision is randomly selected for payment, you should be making your decision as if every decision counts.

Sample Decisions

Here is an example of the type of decision you will have to make. This is just an example to demonstrate how everything is calculated. The example is not meant to guide your decision in any way. On the actual decision sheets we want you to select the allocation that you like best.

Example: You have been given $20 to allocate between the child and yourself. The research foundation's fixed donation towards the child is $5. You must choose how much money to allocate towards the child and yourself.

You may choose to allocate nothing towards the child's books and $20 to yourself. If this decision is selected for payment the foundation's fixed donation of $5 is spent on the child and your payment from the decision will be $20.

Alternatively you may choose to allocate $20 towards the child and nothing to yourself. The money to be spent on the child's books will be $20+$5 = $25, and your payment from the decision is $0.

Finally, you may choose to allocate any amount between $0 and $20 to the child and the remainder to yourself. Suppose you choose to allocate $8 towards the child and $12 to yourself. If selected for payment the American Red Cross will receive $8+$5 = $13 to spend on the child's books and your payment for the decision will be $12.

Monitor Role

To verify that all the procedures of this study are followed we will select a participant to be the monitor of the study. If your Claim Check number is 8 you will be the monitor. The monitor will follow the assistants around to see that everything takes place as explained in these instructions. The monitor will receive a fixed payment for his or her time.

Once all decision forms have been collected all participants will be given a survey. While you are completing the survey the monitor will walk with two assistants to a separate room to oversee that the calculation of the funds for the child and you are performed as described in the instructions. Your payment will be placed along with a receipt in an envelope that has your claim check number on the face of it. The assistant will make out a check to the American Red Cross of Southwestern PA for the amount corresponding to the funds for the child determined by your allocation. One check will be made out for each child. This check as well as any relevant acknowledgment form will be placed in an addressed and stamped envelope to the American Red Cross. Once all the calculations have been completed an assistant will walk the monitor back to this room. A box of envelopes with your payments will be given to an assistant who has not seen your decision sheets. The monitor will then make a statement to you on the extent to which the instructions were followed as described in the instructions. Once you have completed your survey you may come to the front to collect your payment by showing your claim check. An assistant who has not seen your decision form will hand you the sealed envelope with your payment.

After the study is completed the monitor and an assistant will walk to the nearest mailbox (on Forbes next to the Hillman Library) where the monitor will drop the envelope in the mailbox. To prove that all procedures are followed the monitor will be asked to sign a certificate to that effect. This certificate will be posted outside 4916 Posvar Hall.

Upon receipt of the check and acknowledgment form the American Red Cross will send a letter affirming that the check has been used to buy books for the child according to the description above. This letter will be posted outside 4916 Posvar Hall.

If you are the monitor of this study please identify yourself by coming to the front of the room now.

If you have any questions about the procedures, please raise you hand now and one of us will come to your seat to answer your question.

Before we proceed to the decision task we want you to complete a brief quiz, to make sure you know how everything will be calculated.

BIBLIOGRAPHY

[1] Andreoni, J. (1988): "Privately Provided Public Goods in a Large Economy: The Limits of Altruism," Journal of Public Economics, 35, 57-73.

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

[3] Andreoni, J. (1990): "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," Economic Journal, 100, 464-477.

[4] Andreoni, J. (1993): "An Experimental Test of the Public-Goods Crowding-Out Hypothesis," American Economic Review, 83(5), 1317-1327.

[5] Ausubel, L. and Deneckere, R. (1989), "Reputation in Bargaining and Durable Goods Monopoly", *Econometrica*, 57, 511-531.

[6] Banks, J. and Sobel, J. (1987), "Equilibrium Selection in Signaling Games", *Econometrica*, 55, 647-661.

[7] Battaglini, M. (2005), "Long-Term Contracting with Markovian Consumers", *American Economic Review*, 95, 637-658.

[8] Berg, J., Dickhaut, J. and McCabe, K. (1995). "Trust, Reciprocity, and Social History," Games and Economic Behavior, 10 (1), 122-42.

[9] Bergstrom, T., Blume, L. and Varian, H. (1986): "On the Private Provision of Public Goods," Journal of Public Economics, 29, 25-49.

[10] Biehl, A. (2001), "Durable-Goods Monopoly with Stochastic Values", *Rand Journal of Economics*, 32, 565-577.

[11] Blume, A. (1990), "Bargaining with Randomly Changing Valuations" (mimeo., University of Iowa).

[12] Blume, A. (1998), "Contract Renegotiation with Time-Varying Valuations", *Journal of Economics and Management Strategy*, 7, 397-433.

[13] Bolton, G. and Katok, E. (1998): "An Experimental Test of the Crowding-Out Hypothesis: The Nature of Beneficent Behavior," Journal of Economic Behavior and Organization, 37, 315-331.

[14] Bolton, G., Katok, E. and Ockenfels, A. (2004). "How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation," Management Science, 50 (11), 1587-1602.

[15] Bulow, J. (1982), "Durable-Goods Monopolists", *Journal of Political Economy*, 90, 314-332.

[16] Camera, G. and Casari, M. (2008). "Cooperation among Strangers under the Shadow of the Future," Working paper, Purdue University.

[17] Camerer, C. (2003). Behavioral Game Theory: Experiments in Strategic Interaction, Russell Sage Foundation and Princeton University Press.

[18] Cho, I. K., 1987, "A Refinement of Sequential Equilibrium", *Econometrica*, 55, 1367-1389.

[19] Cho, I. K. and Kreps, D. (1987), "Signaling Games and Stable Equilibria", *Quarterly Journal of Economics*, 102, 179-222.

[20] Cho, I. K. and Sobel, J. (1990), "Strategic Stability and Uniqueness in Signaling Games", *Journal of Economic Theory*, 50, 381-413.

[21] Coase, R. H., 1972, "Durability and Monopoly", *Journal of Law and Economics*, 15, 143-149.

[22] Duffy, J. and Ochs, J. (2007). "Cooperative Behavior and the Frequency of Social Interaction," Working paper, University of Pittsburgh.

[23] Eckel, C., Grossman, P. and Johnston, M. (2005): "An Experimental Test of the Crowding Out Hypothesis," Journal of Public Economics, 89 (8), 1543-1560.

[24] Ellison, G. (1994). "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," Review of Economic Studies, 61, 567-588.

[25] Engle-Warnick, J. and Slonim, R. (2006). "Learning to Trust in Indefinitely Repeated Games," Games and Economic Behavior, 54, 95-114.

[26] Fehr, E. and Schmidt, K. (1999). "A Theory of Fairness, Competition and Cooperation," Quarterly Journal of Economics, 114, 817-68.

[27] Fischbacher, U. (2007). "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," Experimental Economics, 10(2), 171-178.

[28] Fudenberg, D., Levine, D. and Tirole, J. (1985), "Infinite-horizon models of bargaining with one-sided incomplete information", In Roth, A. (ed.) *Game-Theoretic Models of Bargaining* (Cambridge University Press).

[29] Fudenberg, D. and Tirole, J. (1991), *Game Theory* (MIT Press).

[30] Fudenberg, D. and Tirole, J. (1983), "Sequential Bargaining with Incomplete Information", *Review of Economic Studies*, 50, 221-247.

[31] Greif, A. (1989). "Reputation and Coalitions in Medieval Trade," Journal of Economic History, 49, 857-882.

[32] Greif, A. (1993). "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition," American Economic Review, 83, 525-548.

[33] Gul, F., Sonnenschein H. and Wilson, R. (1986), "Foundations of Dynamic Monopoly and the Coase Conjecture", *Journal of Economic Theory*, 39, 155-190.

[34] Hart, O. and Tirole, J. (1988), "Contract Renegotiation and Coasian Dynamics", *Review of Economic Studies*, 55, 509-540.

[35] Kandori, M. (1992). "Social Norm and Community Enforcement," Review of Economic Studies, 59, 63-80.

[36] Kennan, J. (2001), "Repeated Bargaining with Persistent Private information", *Review of Economic Studies*, 68, 719-755.

[37] Kennan, J. (1998), "Informational Rents in Bargaining with Serially Correlated Valuations" Working Paper, University of Wisconsin.

[38] Ledyard, J. (1995): "Public Goods: A Survey of Experimental Research," in Kagel, J. and Roth, A. ed., Handbook of Experimental Economics, 111-194, Princeton University Press.

[39] Lee, Y.-J. and Xie H. (2006). "Social Norms and Trust among Strangers," Working Paper, University of Pittsburgh.

[40] Loginova, O. and Taylor, C. (2007), "Price Experimentation with Strategic Buyers", Working Paper, University of Missouri and Duke University.

[41] Milgrom, P., North, D. and Weingast, B. (1990). "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs," Economics and Politics, 2, 1-23.

[42] Okuno-Fujiwara, M. and Postlewaite, A. (1995). "Social Norms and Random Matching Games," Games and Economic Behavior, 9, 79-109.

[43] Ribar, D. and Wilhelm, M. (2002): "Altruistic and Joy-of Giving Motivations in Charitable Behavior," Journal of Political Economy, 110(2), 425-457.

[44] Robert, R. (1984): "A Positive Model of Private Charity and Public Transfers," Journal of Political Economy, 92(1), 136-

[45] Sobel, J. (1991), "Durable Goods Monopoly with Entry of New Consumers", *Econometrica*, 59, 1455-1485.148.

[46] Sugden, R. (1982): "On the Economics of Philanthropy," Economic Journal, 92, 341-350

[47] Vesterlund, L. (2006): "Why do People Give?" in Steinberg, R. and Powell, W. ed., The Nonprofit Sector, 2nd edition, Yale Press.

[48] Warr, P. (1982): "Pareto Optimal Redistribution and Private Charity," Journal of Public Economics, 19, 131-138.

[49] Warr, P. (1983): "The Private Provision of A Public Good Is Independent of the Distribution of Income," Economics Letters, 13, 207-211.