A Theory of Peer-induced Fairness in Games

Teck-Hua Ho and Xuanming Su[1]
Haas School of Business
University of California at Berkeley

January 27, 2008

[1]Authors are listed in alphabetical order. We thank George Akerlof, Eduardo Andrade, Colin Camerer, and Vince Crawford for helpful comments. Direct correspondence to any of the authors. Ho and Su: Haas School of Business, University of California at Berkeley, Berkeley, CA 94720-1900, Email: hoteck@haas.berkeley.edu; Su:Xuanming@haas.berkeley.edu.

A Theory of Peer-induced Fairness in Games

Abstract

A long-standing assumption in economics is that people are purely self-interested. This assumption has been challenged recently by accumulating experimental evidence based on the so-called ultimatum game. Behavioral economists propose several models of distributional fairness to relax the self-interest assumption. In this paper we introduce the concept of *peer-induced fairness* because people have a drive to make social comparison. That is, they look to similar others as a reference in order to form their opinions and evaluate their endowments. We investigate peer-induced fairness by considering two independent ultimatum games played in sequence by a leader and two followers. In the first ultimatum game, the leader makes a take-it-or-leave-it offer to the first follower. Before the next ultimatum game is played, the second follower obtains a public signal of this offer. Then, in the second game, the leader makes an offer to the second follower. The second follower infers what the first follower receives, uses this inference to form a reference point, and is averse to accepting offers that falls short of this benchmark. This generalized model nests the standard and several existing models of fairness.

This model makes two sharp predictions. First, the leader's offer to the second follower should be non-decreasing in the common belief of what the first offer is. Second, condition on an offer, the second follower's likelihood of acceptance is inversely proportional to the reference point derived from the signal. We test both predictions experimentally and find strong support for them. We structurally estimate the model and show that peer-induced fairness is 2.5 times larger than distributional fairness. We incorporate heterogeneity by allowing subjects to be either purely self-interested or fairness-minded. Our estimation results suggest that half of the subjects exhibit peer-induced fairness. We show how peer-induced fairness might influence the occurrence of labor strikes, explain low variability in CEO compensation, and limit the extent of price discrimination.

Keywords: Social Comparison, Peer-induced Fairness, Distributional Fairness, Behavioral and Experimental Economics

1 Introduction

Standard theories in economics generate predictions of market behavior by invoking two fundamental assumptions. First, agents are self-interested in that their utility function depends only on their own material payoffs. Second, market behavior is at equilibrium so that no individual can achieve a higher payoff by unilaterally deviating from the equilibrium. Recent advances in behavioral economics relax both assumptions by, for example, allowing agents to care about others' payoffs and to make mistakes (see Camerer, 2003, and Ho, Lim, and Camerer, 2006a for comprehensive reviews). This paper focuses on the self-interested assumption and provides a new perspective to incorporate concerns for fairness.

A simple and powerful way to demonstrate that people are not purely self-interested is to study their behavior in the so-called ultimatum game. In this game, a leader and a follower divide a fixed pie. The leader moves first and offers a division of the pie to the follower. The follower can accept or reject. If the follower accepts, the pie is distributed according to the proposal. If the follower rejects, both players earn nothing. The subgame perfect equilibrium predicts that the leader should offer a small amount (e.g., a dime) to the follower and the follower should accept (since a dime is strictly preferred to receiving nothing) if players care only about their own material payoffs. However, many experiments (where subjects are motivated by substantial financial incentives) contradict this sharp prediction. Typically, there are almost no offers below 20% of the pie. A majority of offers are between 30% to 40%. Low offers are frequently rejected and the probability of rejection decreases with the offer. These findings are robust to stake size (Slonim and Roth, 1998), persist with repeated trials (Roth et. al, 1991), and prevail across diverse cultures (Henrich, 2000; Camerer et. al, 2001).

Several solutions have been proposed to resolve this anomaly. These solutions modify a player's utility function by allowing it to depend on the payoffs of other players in the game (for a review see Fehr and Fischbacher, 2002). In the ultimatum example, each player's utility function now depends on what both players receive. Fehr and Schmidt

(1999) propose the so-called "inequity aversion" model in which each player has a disutility of receiving a payoff that is different from the other players. The extent of disutility depends on the player's relative payoff position; players exhibit a stronger disutility from "being behind" than from "being ahead." Charness and Rabin (2002) extend the inequity aversion model to incorporate reciprocity in the utility function (see also Rabin, 1993 and Fehr and Gachter, 2000). The generalized utility function allows players to reciprocate when others have been nice or misbehaved towards them. Bolton and Ockenfels (2000) propose the so-called Equity-Reciprocity-Competition (ERC) model in which each agent's utility depends on her absolute payoff as well as her relative share of the total payoff. Under ERC, given an absolute payoff, an agent's utility is maximized when her share is equal to the average share. Note that all three models assume that fairness concerns are *integral* in that agents' social preferences depend only on payoffs of other players in the game. We call this *distributional* fairness concerns.[2]

However, in many real-life situations, people are also driven by social comparison (Festinger, 1954). They have a drive to look to others who are in similar situations (i.e., their peers) to evaluate their endowments and judge whether they are treated fairly. We call this phenomenon *peer-induced* fairness concerns. We posit that the peer-induced fairness concerns can be more salient than distributional fairness concerns when agents engage in social comparison. This is so because social comparison creates a powerful reference point or benchmark for players to compare their well-being with other peer groups. Note that this social comparison process is "outside" the game in that the actions of peer groups do not directly affect the material payoffs of the players.

In this paper, we study peer-induced fairness in a social situation involving 3 economic agents. There is one leader and 2 followers. The followers have a similar endowment and the leader plays an ultimatum game with each follower in sequence. Each game involves the leader making a take-it-or-leave-it offer to a follower. The two games are identical

[2]Cui et. al (2007) provides an application of distributional fairness in a business-to-business channel setting. They show that the channel can be efficient even if linear pricing contract is used as long as channel partners are sufficiently fair-minded.

and independent in that each leader-follower pair plays the same game and actions of one game have no bearing on the material payoffs of the other game. However, in between the two games, the second follower obtains an informative but imperfect public signal of the first offer, and uses this signal to form a belief of the first follower's payoff. We analyze this social situation but allow all agents to have distributional fairness concerns and the second follower to have peer-induced fairness concerns by comparing against the first follower. The equilibrium prediction indicates that, other things being equal, the second follower's likelihood of accepting an offer is decreasing in the signal, suggesting a same offer can become less attractive as the second follower's belief of the first offer increases. In addition, the leader's offer to the second follower is contingent on the signal. The higher the signal the more attractive the offer will be.

Let's consider three classes of examples of the above game. First, consider a seller that interacts with multiple buyers (e.g., a manufacturer and multiple retailers, a firm and multiple customers). Each seller-buyer transaction is independent in that actions within a transaction do not influence material outcomes of other transactions. As distributional fairness would suggest, each individual buyer may care about the seller's payoff in their own respective transaction (in addition to their own material payoffs). On top of that, peer-induced fairness suggests that each individual buyer may care a lot about what other buyers receive in their interactions with the same seller. For example, a customer cares about what other customers pay for the same product. Similarly, a retailer cares about what terms other retailers receive from the same manufacturer. The buyer will treat any differences in price and terms to be entirely *unfair*. Second, consider a boss that hires multiple workers with the same skills. Clearly, workers care not only about their own wages but also about how much other peer workers receive. In fact, bosses often pay their workers similar wage despite wide differences in productivity in order to avoid demoralizing less productive workers.[3] Third, consider a family with multiple children. Sibling's rivalry is common and it suggests that children want their parents to

[3] Akerlof and Yellen (1990) shows that if workers proportionately withdraw their effort because of peer-induced fairness concerns, this behavioral tendency can cause unemployment. Similarly, Fehr et. al (1993) shows that sellers respond to higher prices from buyers by offering superior quality products.

treat them without favoritism. Clearly, this phenomenon implies that each child's utility function depends also on other children's payoffs.

We test our model's predictions experimentally by engaging subjects in two independent ultimatum games as described above. Using this set up, we find strong support for our model predictions. The responder in the second ultimatum game is more likely to reject an offer as the obtained signal increases. The leader's offer is strategic in that he exploits the second follower when the signal is low (even if he has made a good offer to the first follower) and concedes more when the signal is high. In addition, we estimate the model parameters using the data. The estimated peer-induced fairness parameter is 2.5 times stronger than the distributional fairness parameter suggesting that the former is more salient in such social settings. We also incorporate heterogeneity in subjects' taste for fairness by using a latent-class approach. We allow for two different segments, one that is purely self-interested and another that has distributional and peer-induced fairness concerns. Our estimation results suggest that half of the subjects are purely self-interested.

The concept of peer-induced fairness has wide economic implications. We briefly discuss three applications in this paper. First, we show how peer-induced fairness can constrain a monopoly's ability to price discriminate. Without peer comparisons, firms have the complete freedom to maximize profits in separate markets that have different economic characteristics. However, when consumers are averse to paying more than others, firms may have to charge the same price in different markets. Second, we show that peer-induced fairness can lead to wage compression. In particular, we show that the low variability in CEO compensation packages is necessary in order to prevent dissatisfaction resulting from peer comparisons (i.e., with other CEOs). Third, we show that peer-induced fairness can severely restrict the set of feasible negotiation outcomes. Specifically, under peer-induced fairness, any mutually agreeable outcome cannot deviate too much from the outcomes of comparable negotiations in the past. In fact, when negotiating parties select different comparison benchmarks, it may become impossible to reach an agreement. This may explain the widespread occurrence of labor strikes.

The rest of the paper is organized as follows. The next section formulates the model and presents the main equilibrium results. Section 3 describes the experimental design and procedure. Section 4 presents the main experimental results and calibrates the model using the data. Section 5 studies the model with heterogeneity. Section 6 describes three economic applications of peer-induced fairness. Section 7 concludes.

2 Basic Model

2.1 Model Setup

There are 3 players, one leader and 2 followers. The leader plays an identical ultimatum game with each of the followers in sequence. In each game, there is a fixed pie π to be divided between the leader and one of the followers. The leader moves first and offers s_1 to the first follower. The first follower's decision a_1 can either be accept ($a_1 = 1$) or reject ($a_1 = 0$). If $a_1 = 1$, the leader receives $\pi - s_1$ and the follower receives s_1. Otherwise, both receive 0. The second follower obtains a signal $z = s_1 + \epsilon$ where ϵ is a random noise term with any arbitrary distribution function $F(.)$ where $F(.)$ has a mean of zero. The second follower infers a reference point $\hat{s_1}$ of what the first offer is from signal z. This same signal is observed by the leader before the second game begins. Similarly, the leader makes a offer s_2 to the second follower. Again, the follower's decision a_2 can be accept ($a_2 = 1$) or reject ($a_2 = 0$). If $a_2 = 1$, the leader receives $\pi - s_2$ and the follower receives s_2. Otherwise, both receive nothing. Note that the leader receives material payoff in both games while each of the follower receives material payoff in their respective game.

Let us define the agents' utility functions. Consider the utility function of the first follower $U_{F1}(s_1, a_1)$. The follower's utility function has two components. The first component is the agent's material payoff from the game and the second component reflects the first follower's disutility from receiving a payoff that is behind that of the leader. Hence, the second component captures distributional fairness concerns.

$$U_{F1}(s_1, a_1) = \begin{cases} s_1 - \delta_B \cdot \max\{0, (\pi - s_1) - s_1\}, & \text{if } a_1 = 1, \\ 0, & \text{if } a_1 = 0. \end{cases} \tag{2.1}$$

Here δ_B is the parameter capturing the degree of aversion from being distributionally behind.[4]

The second follower's utility $U_{F2}(s_2, a_2)$ is defined similarly except that it also contains an additional component that reflects the disutility from being behind arising from the drive of comparing oneself to a similar peer (i.e., the first follower). It is given below:

$$U_{F2}(s_2, a_2|z) = \begin{cases} s_2 - \delta_B \cdot \max\{0, (\pi - s_2) - s_2\} & \\ \quad -\rho_B \cdot \max\{0, \hat{s_1}(z) - s_2\} & \text{if } a_2 = 1, \\ 0\ , & \text{if } a_2 = 0. \end{cases} \tag{2.2}$$

where ρ_B is the degree of aversion from being behind in a social comparison with a peer and $z = s_1 + \epsilon$ is the signal observed by the second follower and the leader. Note that the disutility associated with social comparison requires oneself to be in a similar social situation with a peer. In other words, the second follower assumes that the first follower receives $\hat{s}$ when considering acceptance and assumes that the first follower receives 0 when considering rejection.[5]

---

[4]Our model can be extended to include an additional disutility term resulting from being ahead. This is in the spirit of Charness and Rabin (2002) and Fehr and Schmidt (1999). For example, the first follower's utility function can additionally include the negative term $-\delta_A \cdot \max\{0, s_1 - (\pi - s_1)\}$. However, we shall show below that distributional fairness concerns associated with being ahead is absent in our experimental data. We choose to use the simplest possible model to demonstrate the effect of peer-induced fairness concerns because it allows us to generate some sharp predictions about subjects' behaviors (see Ho, Lim, and Camerer, 2006b for other rationales for using the simplest possible model).

[5]The aversion of being behind is similar to the notion of loss aversion (Kahnemann and Tversky, 1979; Camerer, 2001). People have a negative transaction utility when receiving a payoff that is below a well-defined reference point. This same idea has been applied to a business-to-business channel setting to show why nonlinear pricing contracts may not work as well as promised by the standard models because these pricing contracts yield negative transaction utility (Lim and Ho, 2007 and Ho and Zhang, in press).

The leader receives payoffs from both ultimatum games. In the second game, the leader receives the utility $U_{L,II}(s_2, a_2|z)$ given as:

$$U_{L,II}(s_2, a_2|z) = \begin{cases} \pi - s_2 - \delta_B \cdot \max\{0, s_2 - (\pi - s_2)\}, & \text{if } a_2 = 1, \\ 0, & \text{if } a_2 = 0. \end{cases} \tag{2.3}$$

Recall that the signal $z = s_1 + \epsilon$. In the first game, the leader receives the utility $U_{L,I}(s_1, a_1)$ as given below:

$$U_{L,I}(s_1, a_1) = \begin{cases} \pi - s_1 - \delta_B \cdot \max\{0, s_1 - (\pi - s_1)\}, & \text{if } a_1 = 1, \\ 0, & \text{if } a_1 = 0. \end{cases} \tag{2.4}$$

In the second game, the leader chooses s_2 to maximize $U_{L,II}(s_2, a_2|z)$. In the first game, the leader chooses s_1 to maximize $U_{L,I}(s_1, a_1)+U_{L,II}(s_2, a_2|z)$. This follows from standard backward induction.

2.2 Second Follower's Inference of the First Offer, $\hat{s_1}(z)$

The model assumes that the second follower has a prior belief about what the first follower receives and denotes the distribution of this prior by $G(.)$. The second follower has a noisy rational expectation in that $G(.)$ has a mean of s_1 and a standard deviation of σ_1. Given the signal $z = s_1 + \epsilon$, the second follower forms a posterior belief of the first offer, with distribution $H(.)$, given by:

$$h(x|z) = \frac{g(x) \cdot f(z - x)}{\int_{\infty}^{-\infty} g(x) \cdot f(z - x)dx}. \tag{2.5}$$

The second follower's inference of the first offer $\hat{s_1}(z)$ is the expected value of $H(.|z)$ and is given by[6]:

$$\hat{s_1}(z) = \int_{-\infty}^{\infty} x \cdot h(x|z)dx. \tag{2.6}$$

We add an information inference process by the second follower for 3 reasons:

[6]There are at least 2 ways to develop an inference for s_1. First, one can use the expected value of the posterior belief to infer s_1. Second, one can use the mode as a surrogate for s_1. Both are possible and we have chosen the former because it provides an unbiased estimate for s_1.

- This information inference process makes our model more realistic. In many real-life situations, the negotiation outcomes are often kept confidential so as to avoid social comparison (e.g., employees are told not reveal their raises to their peers). By allowing for this process, we make our model applicable to more social settings.

- By introducing imperfect information, we allow the leader to change his behavior as a result of the signal realization. Had the second follower perfectly known the offer to the first follower, the leader's offers in the two games will be the same in equilibrium. Hence, imperfect information provides an extra degree of freedom to test the model and to quantify the degree of peer-induced fairness due to social comparison.

- The information inference process also allows us to separate two fundamentally different kinds of peer-induced fairness from the leader's perspective. The leader may inherently want to treat both followers the same way (e.g., parents showing no favoritism among their children).[7] In contrast, the leader may care about treating the two followers the same way only to the extent that the second follower is averse to being behind. In the former, the leader will divide the pie the same way in the two games independent of the signal. If the latter is true, the leader will in fact condition the offer to the second follower on her belief of what the first follower received (the higher the belief the higher the offer).

As a simple example, consider the special case where the observation noise ϵ follows some uniform distribution over $[-k, k]$. That is, $f(u) = \frac{1}{2k}$ for $u \in [-k, k]$ and $f(u) = 0$ otherwise. This uniform special case is used in our experimental test below. The following lemma states the relationship between the signal and the inference.

[7] One can extend the basic model by allowing the leader to have an intrinsic preference for treating the two followers similarly (e.g., parents may not want to show favoritism among their children). For example, this can be accomplished by adding an extra term $-\beta \cdot (|s_1 - s_2|)$ to the leader's utility function. However, our experimental data indicates that the leader tends to condition the second offer on the second follower's inference of the first offer, which may be different from the actual first offer. Such behavior suggests that the leader does not have a strong intrinsic preference to treat the two followers the same way.

Lemma 1 *The inference on the first offer, $\hat{s}_1(z)$, is increasing in the signal z.*

Proof: See Appendix.

2.3 Equilibrium Analysis

We work backward to derive the equilibrium predictions. In the second game, the leader makes an offer s_2 to the second follower, who then decides whether to accept or reject it (i.e., $a_2 = 1$ or 0). Recall that the signal of the first offer is z and the utility function of the second follower is:

$$U_{F2}(s_2, a_2|z) = \begin{cases} s_2 - \delta_B \cdot \max\{0, (\pi - s_2) - s_2\} & \\ \quad -\rho_B \cdot \max\{0, \hat{s}_1(z) - s_2\} & \text{if } a_2 = 1, \\ 0\ , & \text{if } a_2 = 0. \end{cases} \tag{2.7}$$

Thus, the second follower accepts the offer s_2 if and only if $U_{F2}(s_2, 1|z) \geq 0$. The leader's utility function is

$$U_{L,II}(s_2, a_2|z) = \begin{cases} \pi - s_2 - \delta_B \cdot \max\{0, s_2 - (\pi - s_2)\}, & \text{if } a_2 = 1, \\ 0, & \text{if } a_2 = 0. \end{cases} \tag{2.8}$$

The leader faces two alternatives. First, he may offer zero, which induces the follower to reject, and this leaves the leader with zero utility. Second, he may choose the optimal offer, among all the offers that are acceptable to the second follower. In other words, the leader solves the following problem:

$$\max_{s_2} \quad U_{L,II}(s_2, 1|z) \tag{2.9}$$

$$\text{s.t.} \quad U_{F2}(s_2, 1|z) \geq 0. \tag{2.10}$$

Note that this problem is equivalent to finding the smallest offer s_2 satisfying $U_{F2}(s_2, 1|z) \geq 0$, since the leader's utility $U_{L,II}(s_2, 1|z))$ always increases as s_2 decreases. Let us denote this $s_2^0 \equiv \min\{s_2 : U_{F2}(s_2, 1|z) \geq 0\}$. In general, if this optimal offer s_2^0 leaves the leader with non-negative utility, the leader will make this offer and the second follower will accept. Otherwise, the leader will offer zero and the second follower will reject. The following proposition characterizes the optimal offer s_2^*.

Proposition 1 *The leader's optimal offer to the second follower s_2^* as a function of the follower's inference $\hat{s_1}$, is*

$$s_2^*(\hat{s_1}) = \min\left\{\max\left\{\frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}, \frac{\pi\cdot\delta_B+\rho_B\cdot\hat{s_1}}{1+2\cdot\delta_B+\rho_B}, \frac{\rho_B\cdot\hat{s_1}}{1+\rho_B}\right\}, \frac{\pi(1+\delta_B)}{1+2\delta_B}\right\}. \tag{2.11}$$

Proof: See Appendix.

Note that the optimal offer is the minimum of two terms: 1) $\max\left\{\frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}, \frac{\pi\cdot\delta_B+\rho_B\cdot\hat{s_1}}{1+2\cdot\delta_B+\rho_B}, \frac{\rho_B\cdot\hat{s_1}}{1+\rho_B}\right\}$ and 2) $\frac{\pi(1+\delta_B)}{1+2\delta_B}$. The first term yields the leader's most preferred offer while satisfying the incentive compatibility constraint (i.e., it is the smallest offer that induces the second follower to accept). The second term provides an upper bound of the offer beyond which the leader will make a negative utility, even if accepted by the follower. The first term derives from taking the maximum of three fractions of which one of them is not a function of $\hat{s_1}$. The first/second/third fraction is dominant when $\hat{s_1}$ is small/moderate/large. Note that the follower cares only about distributional fairness when $\hat{s_1}$ is small (i.e., the first fraction is independent of ρ_B) and only about peer-induced fairness when $\hat{s_1}$ is large (i.e., the third fraction is independent of δ_B). The follower cares about both kinds of fairness when $\hat{s_1}$ is moderate (i.e., the second fraction depends on both δ_B and ρ_B).

As Proposition 1 demonstrates, the equilibrium offer s_2^* in the second game is non-decreasing in the second follower's inference $\hat{s_1}$. In fact, when $\hat{s_1}$ is sufficiently large, s_2^* is strictly increasing in a piecewise linear manner. This provides a sharp prediction on the leader's behavior. If the second follower has peer-induced fairness concerns (i.e., $\rho_B > 0$) and the leader strategically anticipates such preferences, the leader should contingent the offer on the inference $\hat{s_1}$. We provide an experimental test of this prediction below.

In the first game, the leader makes the offer s_1 to the first follower. Recall that the first follower's utility function is

$$U_{F1}(s_1, a_1) = \begin{cases} s_1 - \delta_B\cdot\max\{0, (\pi - s_1) - s_1\}, & \text{if } a_1 = 1, \\ 0, & \text{if } a_1 = 0. \end{cases} \tag{2.12}$$

Therefore, the first follower accepts the offer s_1 if and only if $U_{F1}(s_1, 1) \geq 0$, which can be shown to be equivalent to $s_1 \geq \frac{\pi \cdot \delta_B}{1+2\cdot\delta_B}$.

How much should the leader offer to the first follower? This decision influences the leader's payoffs in both the first and the second games. That is, the leader chooses s_1 to maximize $U_{L,I}(s_1, a_1) + U_{L,II}(s_2, a_2|z)$.

Condition on s_1 in the first game and along the equilibrium path in the second, the term $U_{L,II}(s_2, a_2|z)$ can be written in terms of the signal z as

$$U^*_{L,II}(z) = U_{L,II}(s^*_2(z), a^*_2(z)|z). \tag{2.13}$$

Since the signal $z = s_1 + \epsilon$, the expected value of the above utility given a first offer s_1 is

$$EU^*_{L,II}(s_1) = \int_{-\infty}^{\infty} U^*_{L,II}(s_1 + \epsilon)dF(\epsilon). \tag{2.14}$$

Therefore, the leader chooses the first offer s_1 to maximize $U_{L,I}(s_1, a_1) + EU^*_{L,II}(s_1)$. Recall that the leader may either offer zero (which induces the first follower to reject) or make the optimal offer that is acceptable to the follower. The following lemma states the relationship between the first offer s_1 and the leader's expected utility in the second game at equilibrium.

Lemma 2 *Condition on s_1 and along the equilibrium path, the leader's expected utility in the second game, $EU^*_{L,II}(s_1)$, is decreasing in s_1.*

Proof: See Appendix.

The lemma suggests that the leader incurs two costs of making a high offer s_1. First, a high s_1 will lower the leader's material payoffs in Game I. Second, the same high offer also leads to a lower expected utility for the leader in Game II. This is because a high s_1 sets a high reference point for social comparison by the second follower and this effect forces the leader to make a more generous offer s_2 to induce the follower to accept. Consequently, to mitigate peer-induced fairness concerns in Game II, the leader will make the smallest possible offer in Game I. This offer is however constrained by the follower's distributional fairness concerns.

The following proposition states that the optimal s_1^* is precisely the lower bound imposed by considering the first follower's distributional fairness concerns.

Proposition 2 *The leader's optimal offer to the first follower s_2^* is*

$$s_1^* = \frac{\pi \cdot \delta_B}{1 + 2 \cdot \delta_B}. \tag{2.15}$$

Proof: See Appendix.

Let's consider a numerical example. Suppose that $\delta_B = 0.5, \rho_B = 1.5, \pi = 100$. Suppose that the noise term ϵ has distribution F that is uniform over $\{-20, -10, 0, 10, 20\}$, and suppose that the second follower's prior belief of the first offer is normally distributed with mean s_1^* and variance 20. With these parameters, the equilibrium first offer is $s_1^* = 25$. Given the offer the first follower will accept (i.e., $a_1^* = 1$). Condition on the first offer and the distribution of the noise term (ϵ), the possible signal values are $\{5, 15, 25, 35, 45\}$. The equilibrium second offers condition on the signal are given in Table 1 below.

Signal (z)	Equilibrium Second Offer (s_2^*)
5	25
15	25
25	25
35	27.54
45	30.25

Table 1: Equilibrium Second Offers in a Numerical Example

The second follower always accepts the offer at equilibrium. Note the following about the relationship between the second offer and the signal:

1. The second offer s_2^* is non-decreasing in the signal. At the highest possible signal, the offer is 20% above the first offer s_1^*.

2. The second offer s_2^* is always greater than or equal to $s_1^* = 25$, a constraint imposed by distributional fairness concerns. This result implies that the leader is more generous to the second follower in general.

3 Experimental Procedure

Fifty-seven undergraduate students at a western university participated in the experiment.[8] There were three experimental sessions. Each session had between 15 and 21 subjects and always consisted of 24 decision rounds. Each subject played the game 24 times. The matching protocol was such that subjects were randomly matched with others in each round and they never knew the identities of other players. Each session lasted for one and a half hours. Subjects earned an average payment of about $19. Before the experiment began, subjects were read the instructions aloud and were given a chance to ask question in private. A copy of the instructions is given in Appendix B. The entire experiment was computerized to facilitate information passing and random matching.

We simplified the decision task as much as possible. For example, in the instructions, a table was given to depict the possible first offers given a signal value. The anonymous matching of subjects was intended to avoid any communication between subjects. Since a random matching protocol is used in each round, we controlled for collusion, reciprocity, and reputation building behaviors. Therefore, each round could be framed as a one-shot game with new partners. In each round, subjects were randomly grouped in triples. In each triple, the three subjects were randomly assigned the roles of RED (leader) or BLUE1 (the first follower) or BLUE2 (the second follower). The three players played

[8]It is common to use undergraduates to test theories of industrial organization (see Holt, 1995). The results could in principle be replicated with managers. Several previous studies comparing professionals and students find little difference between the two groups (see Plott, 1987, and Ball and Cech 1996). Alternatively, one could use student subjects with different levels of experience with the task to assess whether experts behave differently from novices (e.g., Jung, Kagel, and Levin 1994).

two independent ultimatum games each with a pie size of 100 points in sequence.

RED and BLUE1 played Game I first. RED moved first and chose the first offer s_1 (an integer between 0 and 100) at which she wished to divide the pie between herself and the first follower. The computer routed the information on s_1 to BLUE1. BLUE1 must decide whether or not to accept the offer. If the BLUE1 chose to accept, RED and BLUE1 received the allocated amount accordingly. If BLUE1 rejected, both players earned 0 point.

To construct the signal z, we drew a number from a discrete uniform distribution over the set $\{-20, -10, 0, 10, 20\}$ and added it to the first offer. Consequently, given a signal z, the subjects could infer what the possible first offers are. To measure $\hat{s_1}$ directly, we asked BLUE2 to make a guess of what the first offer was and rewarded the player a modest sum of 10 points for making a correct guess. All outcomes, including whether BLUE2 guessed correctly, were revealed only at the end of each decision round comprising of both Games I and II.

Finally, RED and BLUE2 played Game II. RED moved first and made an offer s_2 to BLUE2. BLUE could either accept or reject. If BLUE2 chose to accept, both players received payoffs as allocated. Otherwise, both received nothing. The outcomes of both Games I and II were given to all players at the end of each round.

Each player's total point earnings for a decision round were recorded. Note that the leader received point earnings from both Games I and II. At the end of the session, point earnings for all rounds were summed up and redeemed for cash payment at the rate of \$0.01 per point (i.e., each ultimatum game involved dividing a pie of \$1).

4 Estimation

4.1 Basic Results

Table 2 shows the basic results. Note that few offers are above 50% of the pie. Across the two games, less than 4.2% of the offers are within this range. The modal offer is between 30% and 35% for both games. Few offers are below 15% of the pie. No more than 4.6% of the offers fall into this range across the games. Hence the subgame perfect equilibrium prediction of a very low offer is strongly rejected. There is a clear pattern of a higher rate of rejection as the offer decreases. For example, there is no single offer in the range of 45% to 50% that was rejected, while the rate of rejection ranges from 24.4% to 27.1% when the offers are within the range of 25% to 30%. The overall results suggest that subjects are not purely self-interested. In general our results are comparable to those of prior studies except that the offers are slightly lower and followers tend to reject less frequently.

	Game I		Game II	
Offer Range	Offers (%)	Rejected (%)	Offers (%)	Rejected (%)
> 50	19 (4.2)	0 (0)	18 (3.9)	0 (0)
50	23 (5.0)	0 (0)	25 (5.5)	2 (8.0)
45 – 49.5	10 (2.2)	0 (0)	15 (3.3)	0 (0)
40 – 44.5	73 (16.0)	1 (1.4)	79 (17.3)	1 (1.3)
35 – 39.5	66 (14.5)	4 (6.1)	53 (11.6)	3 (5.7)
30 – 34.5	113 (24.8)	12 (10.6)	119 (26.1)	9 (7.6)
25 – 29.5	45 (9.9)	11 (24.4)	48 (10.5)	13 (27.1)
20 – 24.5	69 (15.1)	13 (18.8)	68 (14.9)	8 (11.8)
15 – 19.5	17 (3.7)	6 (35.3)	12 (2.6)	5 (41.7)
10 – 14.5	14 (3.1)	7 (50.0)	16 (3.5)	10 (62.5)
< 10	7 (1.5)	6 (85.7)	3 (0.7)	2 (66.7)
All	456 (100.0)	60 (13.2)	456 (100.0)	53 (11.6)

Table 2: The distribution of offers and the rate of rejection

4.2 Does Peer-Induced Fairness Exist?

In our experiment, we induce the second follower to make social comparisons by asking her to guess the offer received by the first follower. To check whether this inference is accurate, we regress the second follower's guess ($\hat{s_1}$) of the first offer against the actual first offer. Formally, we have:

$$\hat{s_1} = \omega_0 + \omega_1 \times s_1 \tag{4.1}$$

The best fitted regression line yields $\hat{\omega_0} = 20.1\%$ (p-value $< 1 \times 10^{-16}$) and $\hat{\omega_1} = 0.40$ (p-value $= 3.93 \times 10^{-16}$). This suggests that the second follower is responsive to the first offer and exhibit some biases. They tend to over-estimate the first offer when it is less than 33% and over-estimate when it is above 33%.

The central hypothesis of this paper is that the second follower has peer-induced fairness concerns. The second follower's utility function (equation 2.2) implies that, other things being equal, the second follower is less likely to accept an offer if the difference between $\hat{s_1}$ and the offer s_2 is high. Table 3 below shows how the rate of rejection varies depending on whether the second follower believes she is ahead ($s_2 - \hat{s_1} > 0$), on par ($s_2 - \hat{s_1} = 0$), or behind ($s_2 - \hat{s_1} < 0$).

Being Ahead ($s_2 - \hat{s_1} > 0$)		On Par ($s_2 - \hat{s_1} = 0$)		Being Behind ($s_2 - \hat{s_1} < 0$)	
N	Number of Rejection	N	Number of Rejection	N	Number of Rejection
166	6 (3.6%)	110	5 (4.5%)	180	42 (23.3%)

Table 3: Different Rates of Rejection when Follower 2 is Ahead or Behind

The results are clear: The second follower rejects a lot more frequently when they are behind than otherwise (23.3% versus 4%). We test this formally by running a simple logistic regression with BLUE2's decision a_2 against the second offer s_2 and how much it differs from BLUE2's guess (which is an estimate for $\hat{s_1}$). Formally, we have:

$$P(a_2 = 1) = \frac{e^{\omega_0 + \omega_1 \cdot s_2 + \omega_2 \cdot (\hat{s_1} - s_2)}}{1 + e^{\omega_0 + \omega_1 \cdot s_2 + \omega_2 \cdot (\hat{s_1} - s_2)}} \tag{4.2}$$

If BLUE2 has peer-induced preferences, we would expect ω_2 to be negative. The estimation results show that $\hat{\omega_2} = -0.02463$ (p-value $= 0.0213$), suggesting that the second follower is indeed reluctant to accept an offer that is inferior to that of a peer. This finding rejects the self-interested assumption and theories that ignore peer-induced fairness concerns.

4.3 Did the Leader Respond to Peer-Induced Fairness?

Proposition 1 suggests that the leader's offer in Game II is non-decreasing in $\hat{s_1}$. Indeed, it is piecewise linear in $\hat{s_1}$ if the latter is sufficiently high. Figure 1 shows the observed frequencies of the difference between the second offer and the guess $(s_2 - \hat{s_1})$. Note that this difference centers around zero and drops quickly as the difference gets larger suggesting that the offer may be influenced by the guess.

A simple test for this prediction is to regress s_2 against $\hat{s_1}$. Formally, we have:

$$s_2 = \omega_0 + \omega_1 \cdot \hat{s_1} \tag{4.3}$$

If the prediction is right, we expect ω_1 to be positive. The regression results suggest that ω_1 is indeed positive and highly significantly ($\hat{\omega_1} = 0.197$ and p-value $= 4.75 \times 10^{-7}$). This result implies that the leader is strategic and aligns her second offer with the second follower's inference of what the first offer is.

4.4 Parameter Estimation

To formally estimate the relative importance of peer-induced and distributional fairness concerns, we structurally estimate the model parameters. The proposed model has 2 parameters δ_B and ρ_B. The model involves 4 decisions, s_1, s_2, a_1, and a_2. We assume normal error terms for the leader's decisions.

$$s_1 = s_1^* + \eta_1 \tag{4.4}$$

$$s_2 = s_2^* + \eta_2 \tag{4.5}$$

where η_1 and η_2 are normally distributed with mean of 0 and variances of σ_1^2 and σ_2^2 respectively. The followers' utilities have an extreme value error term so that their

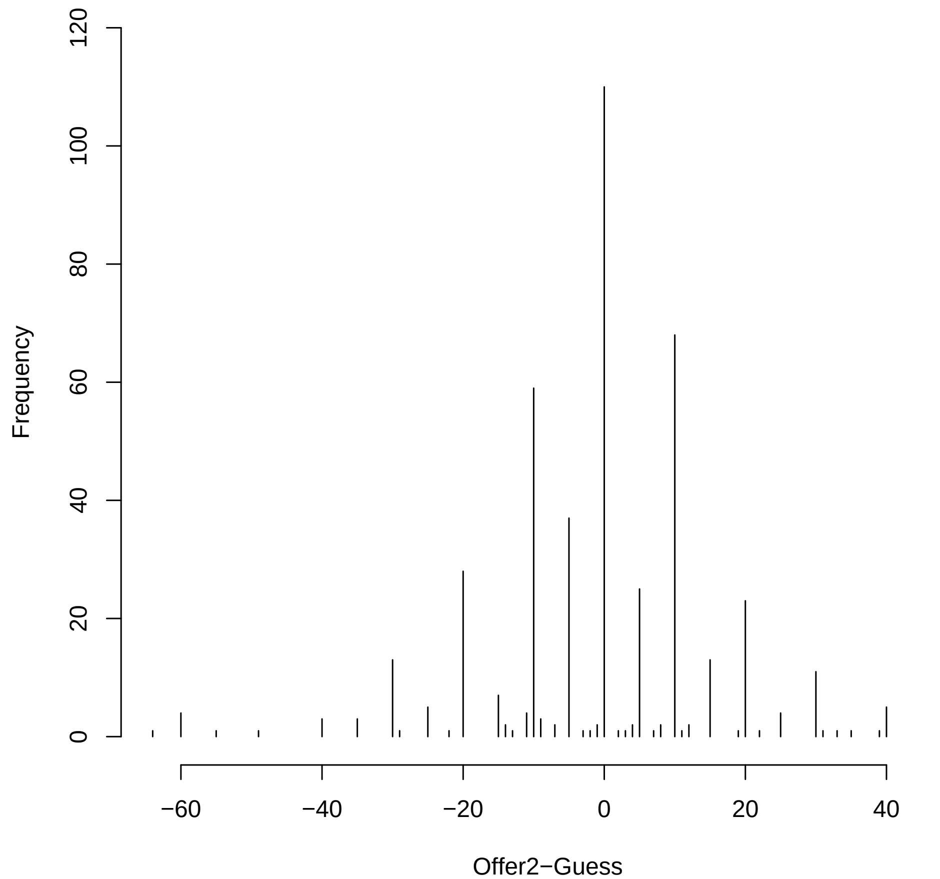

Figure 1: Observed Frequencies of $s_2 - \hat{s_1}$

acceptance probability has a logistic form with parameters λ_1 and λ_2 given below:

$$P_1(\delta_B, \lambda_1) = \frac{e^{U_{F1}(\delta_B)/\lambda_1}}{1 + e^{U_{F1}(\delta_B)/\lambda_1}} \tag{4.6}$$

$$P_2(\delta_B, \rho_B, \lambda_2) = \frac{e^{U_{F2}(\delta_B, \rho_B)/\lambda_2}}{1 + e^{U_{F2}(\delta_B, \rho_B)/\lambda_2}} \tag{4.7}$$

In summary, the likelihood function for a set of decisions s_1, s_2, a_1, and a_2 is:

$$\phi(s_1) \cdot \phi(s_2) \cdot (P_1)^{a_1} \cdot (1 - P_1)^{(1-a_1)} \cdot P_2^{a_2} \cdot (1 - P_2)^{(1-a_2)} \tag{4.8}$$

which we maximize over the parameters $\delta_B, \rho_B, \sigma_1, \sigma_2, \lambda_1, \lambda_2$.

Table 4 shows the estimation results. We estimate the full model and two nested models. The first column presents the nested model without any fairness concerns (i.e.,

$\delta_B = \rho_B = 0$ or agents are purely self-interested). The second column gives the results when players have only distributional fairness concerns (i.e., $\rho_B = 0$). The third column presents the full model. Both nested models are strongly rejected when compared to the full model indicating that subjects care about both distributional and peer-induced fairness. The self-interested hypothesis is clearly rejected ($\chi^2 = 1372.2$, p-value $< 1.0 \times 10^{-16}$). The nested model where the second follower has only distributional fairness is also strongly rejected ($\chi^2 = 64.2$, p-value $= 1.11 \times 10^{-15}$), suggesting that the second follower clearly has peer-induced fairness concerns. In the full model, the estimated peer-induced fairness parameter is $\hat{\rho_B} = 1.327$, which is two 2.5 times larger than the estimated distributional fairness parameter of $\hat{\delta_B} = 0.492$. This suggests that peer-induced fairness weighs more heavily than distributional fairness in the second follower's behavior.[9]

Parameter	No Fairness	Distributional Fairness Only	Full Model
δ_B	-	0.535	0.492
ρ_B	-	-	1.327
λ_1	13.811	15.692	13.586
λ_2	13.402	15.721	22.446
σ_1	34.292	14.132	14.869
σ_2	34.158	13.356	11.810
LL=	-4807.2	-4153.2	-4121.1

Table 4: Estimation Results

[9] We capture distributional fairness concerns by incorporating people's aversion to being behind others in material payoffs. When people are altruistic or have charitable intent (Charness and Rabin, 2002), they may also have an intrinsic aversion to being ahead of others in material payoffs, which can be captured by an additional disutility term $-\delta_A \cdot \max\{(\pi - s_i) - s_i, 0\}$ for the leader or $-\delta_A \cdot \max\{s_i - (\pi - s_i), 0\}$ for the follower. The parameter δ_A represents the degree of aversion to being ahead. However, in Table 4, for both the nested model with only distributional fairness (second column) and the full model (third column), additionally allowing for this parameter yields the estimate $\hat{\delta_A} = 0$. This suggests that subjects in our experiment are not altruistic and do not have distributional fairness concerns when they are ahead.

5 Incorporating Heterogeneity

Our basic model adopts a representative-agent approach and assumes that all players have identical fairness concerns. In this section, we incorporate heterogeneity by analyzing a 2-segment model in which one segment is purely self-interested and the other segment has both distributional and peer-induced fairness concerns. This extension is useful because a fraction of players is likely to be purely self-interested and we can then determine how self-interested players' behaviors are influenced by the existence of fairness-minded players. We structurally estimate this model using our experimental data. About one half of the subjects are estimated to be purely self-interested when the representative-agent assumption is relaxed.

In the two-segment model, let θ denote the fraction of the self-interested segment (i.e., the segment that has $\rho_B = \delta_B = 0$). The remaining segment has distributional and peer-induced fairness concerns, represented by the parameters δ_B and ρ_B as before. We shall derive the equilibrium using backward induction. The next proposition characterizes the leader's optimal offer s_2^* in the second game. The key observation is that the leader may either make the same offer characterized in Proposition 1 (which induces both types to accept) or simply offer zero (in which case only the purely self-interested followers will accept). The former is preferred when the fraction of fairness-minded players is sufficiently large (i.e., θ sufficiently small).

Proposition 3 *Suppose the follower's inference is $\hat{s_1}$. Denote*

$$k = \min\left\{\max\left\{\frac{\pi \cdot \delta_B}{1 + 2 \cdot \delta_B}, \frac{\pi \cdot \delta_B + \rho_B \cdot \hat{s_1}}{1 + 2 \cdot \delta_B + \rho_B}, \frac{\rho_B \cdot \hat{s_1}}{1 + \rho_B}\right\}, \frac{\pi(1 + \delta_B)}{1 + 2\delta_B}\right\}. \tag{5.1}$$

The leader's optimal offer to the second follower is

$$s_2^* = \begin{cases} k, & \text{if } \pi - k - \delta_B \cdot \max\{0, 2k - \pi\} \geq \theta\pi, \\ 0, & \text{if } \pi - k - \delta_B \cdot \max\{0, 2k - \pi\} < \theta\pi. \end{cases} \tag{5.2}$$

Proof: See Appendix.

Next, consider the first game when there are both self-interested and fair-minded types. Similarly as above, the leader faces a choice between making the minimum acceptable offer to induce the fair-minded types to accept, and offering zero (in which case only the self-interested types will accept). As the next proposition shows, the former is preferred when the fraction of self-interested types θ is sufficiently small. The cutoff value for θ can be calculated numerically.

Proposition 4 *There exists some cutoff $\tilde{\theta} \in [0, 1]$ such that the leader's optimal offer to the first follower is*

$$s_1^* = \begin{cases} \frac{\pi \cdot \delta_B}{1+2\cdot\delta_B}, & \textit{if } \theta \leq \tilde{\theta}, \\ 0, & \textit{if } \theta > \tilde{\theta}. \end{cases} \tag{5.3}$$

Proof: See Appendix.

We structurally estimate this two-segment model using the experimental data. This task helps to determine the fraction of the purely self-interested segment. Table 5 shows the estimation results. The first column, for convenience, replicates the estimation results from our basic model, while the second column adds one additional parameter that represents the size of the purely self-interested segment (θ).

Parameter	Basic Model	Heterogeneity
δ_B	0.492	0.787
ρ_B	1.327	2.573
λ_1	13.586	10.973
λ_2	22.446	10.392
σ_1	14.869	12.979
σ_2	11.810	10.974
θ	-	0.491
LL=	-4121.1	-3823.7

Table 5: Estimation Results for Model Extensions

These results suggest that there is substantial heterogeneity in subjects' preferences for fairness. Forty-nine percent of the subjects is estimated to be purely self-interested.

Consequently, the representative-agent assumption is strongly rejected ($\chi^2 = 594.8$, p-value $< 1.0 \times 10^{-16}$). Furthermore, observe that the model estimates for the fair-minded segment are $\hat{\delta_B} = 0.787$ and $\hat{\rho_B} = 2.573$. These estimates are higher than the estimates from our basic representative-agent model. Nevertheless, as noted previously, the degree of aversion to being behind a peer (ρ_B) is still much higher than the degree of aversion to being distributionally behind (δ_B). This again suggests that peer-induced fairness concerns are more significant than distributional fairness concerns.

6 Economic Applications of Peer-Induced Fairness

Many economic models can be substantially enriched by incorporating peer-induced fairness. In this section, we sketch three simple applications in which peer-induced fairness plays an important role. Specifically, we show how peer-induced fairness can limit the degree of price discrimination, account for low variability in CEO compensation and lead to the occurrence of labor strikes.

6.1 Price Discrimination

Many firms charge the same price in different markets even though the opportunity for price discrimination exists. Peer-induced fairness provides a plausible rational explanation for this phenomenon. Consider a monopoly selling in two separate markets. The marginal production cost in the markets are denoted c_1 and c_2, where $c_1 < c_2$, so we can think of Market 1 as the low-cost market and Market 2 as the high-cost market. For simplicity, assume that the demand function for each market is linear, $D_i(p_i) = 1 - p_i$ for $i = 1, 2$. Equivalently, we can think of each market as a unit mass of consumers, whose valuations are uniformly distributed between 0 and 1. By the standard textbook analysis, we can calculate the monopoly's profit-maximizing price in each market to be $p_i^* = (1 + c_i)/2$. Under this result, the monopolist charges a higher price in the high-cost market than in the low-cost market.

Now, suppose that consumers have peer-induced fairness concerns as described in our model. In this case, consumers in Market 2 will be averse to paying a higher price compared to customers in Market 1 where the price is lower. When a consumer in Market 2 with valuation v pays price p_2 for the product and the price in Market 1 is p_1, the consumer receives utility $v - p_2 - \rho_B \cdot (p_2 - p_1)$. Consequently, only consumers with valuations at least $p_2 + \rho_B(p_2 - p_1)$ are willing to buy. This implies that for each unit by which p_2 increases above p_1, the marginal decrease in demand is now larger. In other words, the impact of peer-induced fairness is to make it more costly for the monopolist to raise price in the high-cost market beyond that in the low-cost market. Following this reasoning, it can then be shown that under peer-induced fairness concerns, the optimal prices satisfy $(1 + c_1)/2 < p_1^* \leq p_2^* < (1 + c_2)/2$. Specifically, whenever it is optimal to charge a positive price differential (i.e., $p_1^* < p_2^*$), we have

$$p_1^* = \frac{1 + c_1 + 2 \cdot \frac{\rho_B}{1+\rho_B} \cdot \left(\frac{1}{1+\rho_B} - c_2\right)}{2 \cdot \left(1 - \left(\frac{\rho_B}{1+\rho_B}\right)^2\right)} > \frac{1 + c_1}{2}, \tag{6.1}$$

$$p_2^* = \frac{\left(\frac{1+\rho_B \cdot p_1^*}{1+\rho_B}\right) + c_2}{2} < \frac{1 + c_2}{2}. \tag{6.2}$$

However, when ρ_B is sufficiently large, the monopolist prefers to eliminate price discrimination completely by charging the same price $p_1^* = p_2^* = (1+\bar{c})/2$, where $\bar{c} = (c_1+c_2)/2$ is the average cost across both markets. This brief analysis clearly indicates that the price differential over the two markets $p_2^* - p_1^*$ is smaller when there is peer-induced fairness.

6.2 Executive Compensation

Why are CEO salaries so high? With the attractive executive remuneration packages in practice, the marginal utility gained from the last dollar in a CEO's pay is likely to be very small. That is, when the CEO's utility function $u(x)$ exhibits diminishing marginal utility, the marginal value of the x-th dollar $u'(x)$ is very small when x is very large. Since the CEO is not much worse off without that last dollar, the size of the compensation package does not seem to serve any major economic role. Why, then, are CEO salaries so high?

Peer-induced fairness concerns provide a possible explanation. Suppose that CEOs engage in social comparison with their peers, i.e. other CEOs. In this case, their utility function can be modeled as $v(x) = u(x) - \rho_B \max\{0, \hat{x} - x\}$, where $u(x)$ is the utility for money as above and $\hat{x}$ is the average compensation received by the focal set of CEOs. Since individuals are likely to engage in upward social comparison by selecting individuals who are better as comparison benchmarks, we expect $\hat{x} > x$. Then, the marginal value of the x-th dollar (where x is large and so $u'(x)$ is small) is $v'(x) \approx \rho_B$, which may be much higher than 0. This discussion suggests that CEO remuneration packages are high not because of their material value, but because of the need to avoid discomforting social comparison.

Peer-induced fairness also suggests that the reference or focal CEO set ($\hat{x}$) can significantly influence this social comparison process. For instance, O'Reilly, Main and Crystal (1988) shows that there is a strong association between CEO compensation and the compensation level of outside directors who serve on the compensation committee. This finding can be interpreted by our model if $\hat{x}$ is the average salary of the members of the compensation committee.

6.3 Labor Strikes

Many labor contract negotiations end up in a strike. In most cases, there had been ample time and opportunities for interaction between negotiating parties. This suggests that agreement is not feasible in the first place. The concept of peer-induced fairness provides a potential explanation.

Suppose two parties A and B are negotiating over a pie, the size of which is normalized to one unit. Both A and B will receive the outside option of zero if they do not come to an agreement. If they do, let x and $1 - x$ be the shares of A and B respectively. Then, by standard analysis, it follows that for any $x \in (0, 1)$, both parties will strictly prefer an agreement. In this case, we call $(0, 1)$ the feasible set.

Now, suppose that both parties exhibit peer-induced fairness concerns. A plausible benchmark might be the outcome of a similar negotiation in the past between another pair of players (A' and B') and suppose this negotiation resulted in A' receiving x' and B' receiving $1 - x'$. Then, in the current negotiation, A's utility from receiving x will be $x - \rho_B \max\{0, x' - x\}$ and B's utility from receiving $(1 - x)$ will be $(1 - x) - \rho_B \max\{0, x - x'\}$. It is easy to see that the feasible set of this game is now smaller, consisting only of allocations $x \in (\frac{\rho_B x'}{1+\rho_B}, \frac{1+\rho_B x'}{1+\rho_B}) \subset (0, 1)$.

Finally, each party in the negotiation may choose a different comparison benchmark. That is, A may be inclined to be compared with A' who received x' whereas B may prefer to be compared with B" who received $1 - x''$. We would expect $x' > x''$ since each party's "comparable" outcome is likely to be biased in favor of that party. For instance, Knez and Camerer (1995) experimentally show that people apply different benchmarks for comparison when they have different outside options. Babcock, Wang, and Loewenstein (1996) provides empirical evidence for such a self-serving bias in teacher contract negotiations. In this case, the feasible set of the game becomes $x \in (\frac{\rho_B x'}{1+\rho_B}, \frac{1+\rho_B x''}{1+\rho_B})$. In fact, when $x' - x'' > 1/\rho_B$, the feasible set is empty. This may occur when the two reference points diverge too widely (i.e., the gap $x' - x''$ is too large), or when the degree of peer-induced fairness ρ_B is too large.

7 Conclusions

In this paper, we propose a model of peer-induced preferences in games. This model is significant in two ways. First, any casual introspection would suggest that peer comparison is a pervasive phenomenon and yet it has been largely ignored in economic modeling and analysis. Second, there is a large literature in social psychology on the topic of social comparison (e.g., see Suls and Wheeler, 2000 for a comprehensive review). This stream of literature suggests that people have a drive to make social comparisons with their peers. Hence we believe it is important to recognize this well-documented behavioral phenomenon in the analysis of strategic games. This paper provides the first attempt in capturing this behavioral regularity mathematically.

We investigate peer-induced fairness in a sequence of two payoff independent ultimatum games played by a leader and 2 followers. The leader and the first follower play an ultimatum game, and then the same leader and the second follower play the same ultimatum game. The games are payoff independent in that each follower receives material payoff *only* in their respective game. Between the two games, there is an information collection stage. That is, after the first ultimatum game, the second follower observes an imperfect signal of the first offer before playing the second ultimatum game. Existing models of distributional fairness predict that the leader's offer in the second game and the second follower's decision should not be influenced by the signal. In contrast, our basic model of peer-induced fairness predicts that the second follower's behavior will be influenced by her inference of the first offer based on the signal and that the leader will align the second offer with the second follower's inference.

We test our model predictions experimentally. Subjects are randomly assigned the roles of leader and followers and are motivated by substantial financial incentives. We find strong support for the predictions. Specifically, the second follower's rate of rejection increases with the difference between the second offer and her inference of the first offer. Also, the leader aligns her offer close to the inference of the first offer in order to avoid rejection by the second follower. In combination, these results strongly suggest the existence of peer-induced fairness. We also structurally estimate our model using the experimental data. Our estimation results show that peer-induced fairness is distinct from distributional fairness and the former is crucial in explaining subjects' behaviors. The parameter estimates suggest that the second follower has a preference for peer-induced fairness that is 2.5 times as strong as her preference for distributional fairness (i.e., the former weighs more heavily in follower's decision).

We extend the basic model by allowing a fraction of the subjects to be purely self-interested. Our structural estimation results indicate that about half of the subjects are purely self-interested while the other half exhibit fairness concerns. This result suggests that it is important to incorporate heterogeneity in the strategic analysis of bargaining.

Finally, we show how peer-induced fairness plays a key role in several economic applications. For example, peer-induced fairness can restrict a monopoly's ability to price discrimination, account for the low variability in CEO compensation, and lead to the occurrence of labor strikes.

Appendix A: Proofs

Proof of Lemma 1 The posterior distribution of the first offer s_1, conditional on the signal z, is simply the distribution G over the truncated support $[z-k, z+k]$. As z increases, the truncated support shifts to the right, so the posterior expectation $\hat{s_1}(z)$ also increases.

Proof of Proposition 1 Introducing the variables $w_1 = \max\{\pi - 2s_2, 0\}$ and $w_2 = \max\{\hat{s_1} - s_2, 0\}$, we obtain the following problem (the solution of which yields s_2^0):

$$\begin{aligned}
\min_{s_2,w_1,w_2} \quad & s_2 && (7.1)\\
\text{s.t.} \quad & s_2 - \delta_B \cdot w_1 - \rho_B \cdot w_2 \geq 0 && (7.2)\\
& w_1 \geq \pi - 2s_2 && (7.3)\\
& w_2 \geq \hat{s_1} - s_2 && (7.4)\\
& w_1, w_2 \geq 0 && (7.5)
\end{aligned}$$

Notice that the above feasible region can be expressed in terms of only s_2 to yield:

$$\begin{aligned}
\min_{s_2} \quad & s_2 && (7.6)\\
\text{s.t.} \quad & s_2 - \delta_B(\pi - 2s_2) - \rho_B(\hat{s_1} - s_2) \geq 0 \qquad \Leftrightarrow s_2 \geq \frac{\pi\delta_B + \rho_B\hat{s_1}}{1 + 2\delta_B + \rho_B} && (7.7)\\
& s_2 - \delta_B(\pi - 2s_2) \geq 0 \qquad \Leftrightarrow s_2 \geq \frac{\pi\delta_B}{1 + 2\delta_B} && (7.8)\\
& s_2 - \rho_B(\hat{s_1} - s_2) \geq 0 \qquad \Leftrightarrow s_2 \geq \frac{\rho_B\hat{s_1}}{1 + \rho_B} && (7.9)\\
& s_2 \geq 0 && (7.10)
\end{aligned}$$

Therefore, among all the offers that are acceptable to the second follower, the offer that maximizes the leader's utility $U_{L,II}(s_2, a_2|z)$ is

$$s_2^0 = \max\left\{\frac{\pi\delta_B + \beta\hat{s_1}}{1 + 2\delta_B + \rho_B}, \frac{\pi\delta_B}{1 + 2\delta_B}, \frac{\rho_B\hat{s_1}}{1 + \rho_B}, 0\right\}. \tag{7.11}$$

Next, notice that the offer s_2^1 that leaves the leader with zero utility is

$$s_2^1 = \pi - \frac{\pi\delta_B}{1 + 2\delta_B} = \frac{\pi(1 + \delta_B)}{1 + 2\delta_B}. \tag{7.12}$$

Finally, we see that the leader's equilibrium offer in the second game must be $\min\{s_2^0, s_2^1\}$, as given in the proposition.

Proof of Lemma 2 Consider two possible offers s_1 and $s_1' = s_1 + c$ with $c > 0$. It is clear that the set of possible signal realizations z' corresponding to s_1' is a translation of the set of possible signal realizations z corresponding to s_1. It then follows that the distribution of inferences $\hat{s_1}(z')$ is a translation (to the right) of the distribution of inferences $\hat{s_1}(z)$. Since $s_2^*(\hat{s_1})$ increases as the estimate $\hat{s_1}$ increases, along the equilibrium path, the leader's expected utility in Game II decreases as s_1 increases.

Proof of Proposition 2 By Lemma 2, we know that $EU_{L,II}^*(s_1)$ is decreasing in s_1. Also, note that $U_{L,II}^*(z) \leq U_{L,I}(\frac{\pi\delta_B}{1+2\delta_B}, 1)$. This holds because for any z, $U_{L,II}^*(z) = U_{L,II}(s_2^*(z), a_2^*(z)|z) \leq U_{L,II}(\frac{\pi\delta_B}{1+2\delta_B}, 1|z) \leq U_{L,I}(\frac{\pi\delta_B}{1+2\delta_B}, 1)$.

Now, we evaluate the two alternatives facing the leader: offer zero (and the first follower rejects) or offer the optimal acceptable offer (and the follower accepts). Recall that the leader wishes to maximize $U_{L,I}(s_1, a_1) + EU_{L,II}^*(s_1)$. When the leader offers zero to the first follower, the first term is zero and the second term is at most $U_{L,I}(\frac{\pi\delta_B}{1+2\delta_B}, 1)$. Alternatively, the leader may make an offer that is acceptable to the first follower. Recall that only offers $s_1 \geq \frac{\pi\delta_B}{1+2\delta_B}$ are acceptable. Since both $U_{L,I}(s_1, a_1)$ and $EU_{L,II}^*(s_1)$ are decreasing in s_1, the leader's optimal offer that is acceptable to the follower is $s_1 = \frac{\pi\delta_B}{1+2\delta_B}$. In this case, the first term is $U_{L,I}(\frac{\pi\delta_B}{1+2\delta_B}, 1)$ and the second term is non-negative. The proposition thus follows.

Proof of Proposition 3 It is clear that k is the minimum offer that is acceptable to the type with fairness concerns. The leader may either: (i) offer k and receive $U_{L,II}(k,1)$, or (ii) offer 0 and receive π with probability θ and 0 with probability $1-\theta$ (i.e., the expected utility is $\theta\pi$). The leader thus chooses the better alternative, as characterized in the proposition.

Proof of Proposition 4 Recall that the leader wishes to maximize $U_{L,I}(s_1,a_1) + U_{L,II}(s_2,a_2)$, where a_1 and a_2 now refers to the acceptance decisions of the fair-minded types. Note from Proposition 3 that along the equilibrium path, we have $U^*_{L,II}(s_2,a_2) = \max\{U_{L,II}(k,1), U_{L,II}(0,0)\} = \max\{U_{L,II}(k,1), \theta\pi\}$. Thus the reasoning in the proof of Lemma 2 continues to apply to the first term and thus Lemma 2 holds. Therefore, the only candidates for the first offer s_1 are 0 and $\frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}$. We will analyze the increase in the leader's utility when he offers $s_1 = 0$, compared to when he offers $s_1 = \frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}$; in this proof, we term this his incremental utility.

In Game I, the leader's utility from offering $\frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}$ does not depend on θ; however, the leader's utility from offering 0, which is $\theta\pi$, has derivative π with respect to θ.

Next, consider the leader's incremental utility from Game II along the equilibrium path. When $s_1 = 0$, the leader's utility is $\max\{U_{L,II}(k,1|s_1=0), \theta\pi\}$. When $s_1 = \frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}$, the leader's utility is $\max\{U_{L,II}(k,1|s_1=\frac{\pi\cdot\delta_B}{1+2\cdot\delta_B}), \theta\pi\}$. In both cases, the first term does not depend on θ and the second term has derivative π with respect to θ. Therefore, the derivative of the incremental utility (i.e. the difference) with respect to θ must be at least $-\pi$.

Combining the two games, the derivative of the incremental utility with respect to θ must be non-negative. In other words, as θ increases, offering $s_1 = 0$ always becomes more attractive. The proposition thus follows.

Appendix B: Instructions

This is an experiment in economic decision making. The instructions are simple and if you follow them carefully and make good decisions, you could earn a considerable amount of money which will be paid to you in cash before you leave today. Different subjects may earn different amounts of cash. What you earn today partly depends on your decisions, partly on the decisions of others, and partly on chance.

The experiment consists of 24 decision making rounds. There are 3 subjects in this room. In each round, we will randomly group you into 1 triplets. In each round and in each triplet, one subject will be RED player and two subjects will be BLUE players (BLUE1 and BLUE2). You have an equal chance of playing the role of RED, BLUE1 or BLUE2 in each round. The decision making task of each player will be explained below.

It is important that you do not look at the decisions of others, and that you do not talk, laugh or exclaim aloud during the experiment. You will be warned if you violate this rule the first time. If you violate this rule a second time, you will be asked to leave and you will not be paid. That is, your total earnings will be zero.

Experimental procedure

In each round, the decision making task occurs in 3 stages, namely, I, II, and III. Each RED player and the 2 matched BLUE players (BLUE1 and BLUE2) undertake the task as follows. Again the assignment of your role is determined randomly so that each person in the triplet has an equal chance of playing RED, BLUE1 or BLUE2.

In Stage I, RED and BLUE1 will have a pot of 100 points to divide between them (BLUE2 will sit still in this stage). RED will make an offer of OFFER1 (ranging from 0 to 100 points) to give it to BLUE1. After receiving the offer OFFER1, BLUE1 must decide whether or not to accept it. If BLUE1 accepts the offer, RED will receive 100 - OFFER1 points and BLUE1 will receive OFFER1 points. However, if BLUE1 rejects the offer, both RED and BLUE1 will receive nothing in that decision making round. Note that the outcome of Stage I (i.e. whether BLUE1 accepts the offer) will only be revealed to RED at the end of Stage III.

In Stage II, we randomly draw a number from a set of 5 numbers: -20, -10, 0, 10, 20.

That is, each number has an equal chance of being drawn. We call the drawn number X. We generate a signal called SIGNAL1 by adding X to OFFER1. We will use the number SIGNAL1 in Stage III. Note that each triplet involves a different independent draw in each decision round. However, each draw is always from the same set consisting of the same 5 numbers.

Let's consider two examples to see how this signal generation process works. If SIGNAL1=30, then there are five possible scenarios:

Offer 1	X	Signal1
50	-20	30
40	-10	30
30	0	30
20	10	30
10	20	30

Note that if SIGNAL1=30, OFFER1 can range from 10 to 50 depending on the value of the random number X.

Similarly, if SIGNAL1=70, we have the following five possible scenarios:

Offer 1	X	Signal1
90	-20	70
80	-10	70
70	0	70
60	10	70
50	20	70

That is, OFFER1 can range from 50 to 90. Note that the above two examples are chosen purely for illustration purposes. In no way, the shown values are indicative of the optimal choices.

BLUE2 will guess what OFFER1 is. If BLUE2 guess correctly, he or she will receive a total of 10 points. If BLUE2 guess wrongly, he or she will receive nothing. Note that BLUE2's guess, and whether it is correct, will be revealed to RED and BLUE2 only at the end of Stage III.

In Stage III, RED and BLUE2 will have a pot of 100 points to divide between them (i.e., BLUE1 will sit still). Before RED makes her offer, both RED and BLUE2 will be informed of the value of SIGNAL1. Note that SIGNAL1 is generated by adding the random draw X described in Stage II to the OFFER1 made by RED to BLUE1 in Stage I. Then, RED will make an offer OFFER2 (ranging from 0 to 100 points) to give it to BLUE2. After receiving the offer OFFER2, BLUE2 must decide whether or not to accept the offer. If BLUE2 accepts the offer, RED will receive 100 - OFFER2 points and BLUE2 will receive OFFER2 points. However, if BLUE2 rejects the offer, both RED and BLUE2 will receive nothing in that decision round.

At the end of Stage III, the RED and both BLUE subjects will be informed of their respective decision outcomes and point earnings. The above decision task is repeated for 24 times. In each round, 1 triplets will be formed. Each player in the triplet will have an equal chance playing RED, BLUE1 or BLUE2.

Payoffs

Your dollar earnings for the experiments are determined as follows. First, we will sum up your total point earnings from all 24 rounds. Then we will multiply your point earnings by 0.01. This is the amount you will be paid when you leave the experiment. Note that the more points you earn, the more money you will receive.

References

[1] Akerlof G. and Yellen, J. (1990) 'The Fair Wage Effort Hypothesis and Unemployment,' Quarterly Journal of Economics, 255-283.

[2] Babcock, L., X. Wang, and G. Loewenstein (1996), 'Choosing the Wrong Pond: Social Comparisons that Reflect a Self-Serving Bias', Quarterly Journal of Economics, 111(1) 1-19.

[3] Bolton, G. and A. Ockenfels (2000), 'ERC: A Theory of Equity, Reciprocity and Competition', American Economic Review, 90, 166-193.

[4] Camerer, C. (2001), 'Prospect Theory in the Wild: Evidence from the Field', in Kahneman, D. and A. Tversky (eds.), Choices, Values and Frames, Cambridge: Cambridge University Press.

[5] Camerer, Colin. F., Joseph Henrich, Robert Boyd, Samuel Bowles, Ernst Fehr, Herbert Gintis, and Richard McElreath. May 2001. 'Cooperation, Reciprocity and Punishment in Fifteen Small-Scale Societies.' American Economic Review, 91, 73-78.

[6] Camerer, C. (2003), Behavioral Game Theory: Experiments in Strategic Interaction, Princeton: Princeton University Press.

[7] Charness, G. and M. Rabin (2002) 'Understanding Social Preferences with Simple Tests', Quarterly Journal of Economics, 117, 817-869.

[8] Cui, H., Raju, J. and J. Zhang (2007), 'Fairness and Channel Coordination,' Management Science, August 2007, 53(8) 1303-1314.

[9] Fehr, E. Fischbacher U. (2002), "Why Social Preference Matter - The Impact of Non-selfish motives on competition, cooperation and incentives," The Economic Journal 112 (478), C1C33.

[10] Fehr E. and Gachter S. (2000) "Fairness and Retaliation: The Economics of Reciprocity" The Journal of Economic Perspectives, Vol. 14, No. 3., pp. 159-181.

[11] Fehr, E. and K. Schmidt (1999), 'A Theory of Fairness, Competition and Cooperation', Quarterly Journal of Economics, 114, 817-868.

[12] Fehr, E., G. Kirchsteiger and A. Riedl (1993), "Does Fairness prevent Market Clearing?" Quarterly Journal of Economics 108, 437-460.

[13] Festinger, L. (1954). A theory of social comparison processes. Human Relations, 7(2) 117-140.

[14] Henrich, J. (2000) 'Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining among the Machiguenga of the Peruvian Amazon,' American Economic Review, 90: 973979.

[15] Ho, T. and J. Zhang (2007), 'Designing Pricing Contracts for Boundedly Rational Customers: Does the Framing of the Fixed Fee Matter?' Management Science, Forthcoming.

[16] Ho, T-H., Lim, N and Camerer, C. (2006a) 'Modeling the Psychology of Consumer and Firm Behavior Using Behavioral Economics,' Journal of Marketing Research, Vol. XL III, 307-331.

[17] Ho, T-H., Lim, N and Camerer, C. (2006b) 'How Psychological Should Economic and Marketing Models Be?' Journal of Marketing Research, Vol. XLIII, 341-344.

[18] Kahneman, K. and Tversky A. (1979) 'Prospect Theory: An Analysis of Decision under Risk,' Econometrica, Vol. 47, No. 2, 263-292

[19] Knez, M. and Camerer, C. (1995) 'Outside Options and Social Comparison in Three-Player Ultimatum Game Experiments,' Games and Economic Behavior, 10, 65-94.

[20] Lim, N. and Ho T-H 'Designing Price Contracts for Boundedly Rational Customers: Does the Number of Blocks Matter?' Marketing Science, 26 (3), 312-326, 2007.

[21] O'Reilly, C.A., B. Main and G. Crystal (1988), 'CEO Compensation as Tournament and Social Comparison: A Tale of Two Theories,' Administrative Science Quarterly, 33(2) 257-274.

[22] Rabin, M. (1993) "Incorporating Fairness into Game Theory and Economics" The American Economic Review, Vol. 83, No. 5., pp. 1281-1302.

[23] Roth, A.E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991), 'Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,' American Economic Review, 81, 1068-1095.

[24] Slonim, R. and A.E. Roth, (1998) 'Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic,' Econometrica, 66,3, 569-596.

[25] Suls, J. and Wheeler, L. (2000) Handbook of Social Comparison: Theory and Research, New York: Kluwer Academic Publishers.