Chapter 13

# FAILURE OF BAYESIAN UPDATING IN REPEATED BILATERAL BARGAINING

Ching Chyi Lee
*The Chinese University of Hong Kong*

Eythan Weg
*Purdue University*

Rami Zwick
*The Hong Kong University of Science and Technology*

**Abstract**

Ever since Camerer and Weigelt (1988) concluded in their important experimental work that "sequential equilibrium describes actual behavior well enough," we might be tempted to use this theory confidently in various domains. To assess the robustness of the above conclusion, the present study attempts to explore Bayesian updating in a bilateral negotiated sale setup injected with a whiff of an ultimatum aroma. We conclude that the ultimatum nature of the basic game tends to overwhelm rational behavior on the part of the sellers and that buyers are not cognizant of favorable prices occurring later in the game.

## 1. INTRODUCTION

Ever since Camerer and Weigelt (1988) concluded in their important experimental work that "sequential equilibrium describes actual behavior well enough," we might be tempted to use this theory confidently in various domains. The present study explores Bayesian updating in a bilateral negotiated sale setup injected with a whiff of an ultimatum aroma, in order to assess the robustness of the above conclusion.

It is quite natural for people or institutions to misrepresent their true nature in pursuit of gaining some benefits which otherwise could not be attained. To misrepresent one's true nature is to act as someone or something else – thereby creating confusion on the true identity of the actor.

Situations in which a party may have an incentive to misrepresent are prevalent: business-to-business suppliers supply in the present and expect to be paid in the future; credit card issuers rely on credit card holders to pay for their purchases.

These are just two examples. Whether a business or an individual is likely to misrepresent depends on the situation and, perhaps, on the law.

Although misrepresentation may touch on questions of the law (for example, the energy company Enron, that borrowed huge sums of money based on a complicated web of holding companies although the chances of the company paying back the loans were very small), there are situations in which misrepresentation may only be a matter of benign convenience and opportunity as the framework explored in the present paper will clearly show.

Bargaining appears to be a natural scenario for misrepresentation because *reputation* is important. Feigning toughness for a sufficient length of time may convince the opponent in the bargaining to relent because of the cost incurred. Raiffa (1982), in his classic *The Art and Science of Negotiation*, writes that although repeated bargaining is often done cooperatively, this is not always the case, especially when there is information disparity between the two sides. "With repetition, a negotiator might want to establish a reputation for toughness that is designed for the long-term rather than short-term rewards" (p. 13). To be successful in this attempt, toughness needs to be communicated in some way to sow the seed in the opponent's mind that the toughness is real.

## 2. THE MODEL

The basic setting for this study includes a buyer and a seller. The buyer is one of two types: low cost (L) and high cost (H). In the current version of the experiment, this is operationalized as the low or high costs related to the seeking of an alternative supplier for an identical product that the seller proposes to sell. Upon receipt of the proposal to sell the product at a specific price, the buyer may accept it and thus terminate the transaction, or opt to *search* at a cost, $c$, for a better price by another supplier. The search for another supplier is always successful; however, the price may be better or worse than the current one proposed by the present seller. If the buyer elects to search, she abandons the opportunity to purchase the unit at the original seller asking price and is committed to pay the "searched" price even if it is higher than the current asking price (i.e., this is a no recall environment). A specific example is a bakery that can not roll over baked goods from one day to the next and the described encounter takes place just prior to closing. A buyer comes in and the baker is making a take-it-or-leave-it offer of price p for a cake. If the buyer accepts the offer the sale is conducted, otherwise, if the buyer rejects the offer the baker disposes of the cake and closes the shop. The buyer in this case is searching for the lowest available price elsewhere, incurring the search cost, and purchase the cake at that price. In general, the no recall environment allows for the negotiation to re-open if the buyer returns from an unsuccessful search. Presumably, the buyer's bargaining position in this case is much weakened. The seller is characterized by uncertainty about the real nature of the buyer. Obviously, a buyer with a known low cost would extract a better price than one characterized by a high search cost. Hence, regardless of her type, a buyer always has an incentive to be known as a low-cost type. As a result, the buyer's "telling" of his type is useless because it cannot be *trusted*. The only way a low-cost search buyer can reveal her true identity is by behaving in a such a way that it is inconsistent with her being a high-cost search buyer. In our case, this can be achieved by the buyer's willingness to reject low price offers. In this way, the buyer behaves as if she has a low search cost and thus can eventually creates a *reputation* for being a low-cost buyer.

### *2.1. The Bargaining Game*

There are two players: a buyer and a seller. The seller possesses five units of a product that he intends to sell to the buyer in five periods, one unit in each period. The buyer is known to the seller to be one of two *types*: $H$ or $L$. Initially, the degree to which the seller believes that the buyer is of type $L$ is $0 \leq \pi \leq 1$. The search cost in buying the product is 0 if the buyer buys the product from the current seller. Otherwise, if the buyer does not purchase the product from the current seller, she needs to search for it at a cost $c = h$ or $c = l$ ($h > l$) corresponding to her type. A search means a realization of a random variable $D$ having a uniform distribution on the interval [0, 100]. Thus, when the seller considers a price offer, $p$, he needs to assess the expected price and the cost that the buyer incurs in case the offer is rejected. The profit for the buyer is $100 - p$ when the price is accepted and $100 - D - c$ otherwise. Thus, the buyer has 500 pay units at his disposal and she wants to minimize the total price over the five periods of the game with the same seller. The seller's profit from each period is the price he can obtain for the unit. A failure to sell the product in a period results in zero profit for this period. The seller's goal is to maximize the total revenue over the five periods. All the information as described above is commonly known to the participants.

### *2.2. The Equilibrium*

The equilibrium of the game described above is very similar to that of the game described by Kreps and Wilson (1982) and also to that of the one described by Camerer and Weigelt (1988). We will only describe the equilibrium path here. The details and the derivation of the equilibrium can be found in Lee (2000).

Let $p_H$ and $p_L$ be the prices that make the high- and low-type buyers indifferent between searching and not searching. That is, $p_H = -h + [100 - E(D)]$ and $p_L = -l + [100 - E(D)]$, where $E(D)$ is the expected price if the buyer searches. Let $\pi_t$ be the seller's belief at Period $t$ about the probability that the buyer is a "low-cost buyer". Notice that we count the "periods" backwards. That is, in period $t$, there are $t$ periods left. Furthermore, let $\theta^t = [(p_H - p_L)/p_H]^t$. Then, in equilibrium, the seller offers $p_L$ until the first time that $\pi_t$ becomes lower than $\theta^t$, at which time the seller offers $p_H$ until the last period. The low-cost buyer accepts $p_t = p_L$ and rejects $p_t = p_H$. The high-cost buyer accepts $p_t = p_L$ with probability one and rejects $p_H$ with probability $\beta_t$ if she has not yet accepted any price greater than $p_L$. If the high-cost buyer has accepted a price greater than $p_L$ earlier, she accepts $p_H$ with probability one. Here, $\beta_t$ satisfies:

$$\beta_t = \frac{\pi_t(1 - \theta^{t-1})}{(1 - \pi_t)\theta^{t-1}}.$$

## 3. THE EXPERIMENT

### 3.1. Method

#### 3.1.1. Subjects

Two hundred and forty male and female subjects, who were mainly undergraduate business students at the Hong Kong University of Science and Technology, in groups of 24 students per session, participated in a session that lasted about 60 minutes. Subjects were recruited through advertisements placed on bulletin boards on campus and made during class announcements. The announcements promised monetary reward contingent on performance in a bargaining study.

#### 3.1.2. Experimental Design

Each of the bargaining games consisted of five periods in which the same players bargained on a surplus of HK$100[1] in each period with an uncertain outside option for the buyer uniformly distributed on the range [0, 100] using the trading rules described above.

We used a 2 (High search cost) × 5 (Degree of belief) × 8 (trials) design. The first two factors were between subjects and the last was within subjects. In all sessions the level of the low search cost was fixed at $5 (a commonly known fact). The high search cost was either $10 or $45. The prior belief about the buyer being a low-cost type was 0.00, 0.01, 0.05, 0.10, or 0.50.

Subjects assumed the same role (a buyer or a seller) in all eight trials in a session and faced different anonymous opponents on each trial. Obviously, the five periods of each trial were played by the same buyer and seller.

During each trial, the seller was asked to sell five indivisible goods, one in each period. The goods had no value to the seller except their selling prices. The value of the good to the buyer was $100. The bargainers knew both reservation prices. Before the beginning of each trial the computer randomly selected the buyer's type based on the known probabilities of each type. Of course, the buyer was informed of her type, whereas the seller only knew the sampling probabilities. Once selected, the buyer's type was fixed for the five periods duration of the game. The game proceeded as follows: at the beginning of each period the seller announced a selling price for the good (the asking price). The buyer then had the following options:

1. Accept the asking price, thereby terminating the period.
2. Search for an alternative price. In this case, the buyer had to pay a search cost, and a price was randomly generated (the outside offer) from the range [0, 100]. The price generated through the search (if any) was known to both bargainers.

After learning about the outside offer, the buyer must accept this offer, thereby terminating the period.

At the end of each period, the buyer was informed of her profit as well as that of the seller. The seller was informed of his profit and that of the buyer's profit if the buyer accepted the seller's asking price, or the buyer's profit up to the uncertainty as to the buyer's search cost if the buyer decided to search (i.e., the buyer's profit was presented to the seller as "outside price – search cost." Whereas outside price was explicitly specified as a number, the search cost was understood to be the unknown buyer's type and as such was literally presented as "search cost"). After the profits were presented the game proceed to the next period, unless it was the last (fifth) period, whereas the game was terminated.

The subjects interacted in a computer laboratory arranged in such a way that it was impossible for the subjects to know with whom they were negotiating or to see each other's screens. Asking prices, acceptances, and searches (including their outcomes) were transmitted through computer terminals. No other communications were allowed.

Throughout the experiment, the subjects were informed of every *known* characteristic of the game being played. Moreover, the *known* dynamic aspects of the settled and searched prices (during the five periods of each game) were registered in a history log and were clearly visible on the screen for the duration of the game.

The subjects were informed that they would be paid their net payoff from one randomly selected game (cumulative over the five periods). In addition, each subject was paid $10 for participation. On the average, subjects earned $73.64 for a session[2].

## 4. RESULTS

### *4.1. Data Consistency*

The data includes 4,800 plays of the same basic game arranged in 10 sessions, covering 960 plays of the repeated game.

Subjects knew that searching will result in a randomly drawn price from the interval [0, 100] and that within this interval every price is "equally likely". That is, the expected search price was 50. If we add the commonly known minimal cost to pursue the search (5) or the maximum possible cost to search (45), we arrive at the fundamental conclusion, which requires no deep thinking or understanding, that asking prices (demands) should be no less than 55 and no more than 95.

Inspecting the data for coherence, we find two anomalies, one minor and one major:

- There are two unusual price demands of 553 and 47,055. These are beyond the range of the upper bound of the random search price, and are also beyond any reasonable demand. The seller's computer interface allows entering a price expressed in dollars and cents, but it was the seller's task to enter the decimal point. However, the interface failed to alert sellers to unusual demands such as we have here. We believe it reasonable to assume that subjects miss writing a decimal point and intended to demand 47.55, and 55.3, respectively.

- A total of 3006 demands (close to two-thirds of all the plays) were *below* the lowest reasonable price demand of 55, 1673 of them below 50.

Table 1 presents the number of price demands below several cutoff points:

*Table 1. Distribution of prices below the lower bound price demands*

| Cutoff | < 50 | < 45 | < 40 | < 35 | < 30 |
|---|---|---|---|---|---|
| Count | 1,673 | 882 | 448 | 212 | 75 |

Table 2 presents the frequencies of demands below the minimum expected demand of 55 in the ten experimental conditions.

The two distributions presented in Table 2 shed light on the problem. These demands, although irrational, are consistent (column-wise) with the fact that the frequency of price demands below 55 is lower when $h$ is 45 than when it is only 10. Thus, when the real higher bound is likely to be higher than 60, subjects tend to deviate from the lower bound of 55 less often. Or, in other words, when the real higher bound is likely to be higher, subjects tend to make higher price demands more often.

*Table 2. Frequency of demands below 55*

High Search Cost (H)

| $\pi$ | *45* | *10* |
|---|---|---|
| 0.00 | 185 | 433 |
| 0.01 | 143 | 320 |
| 0.05 | 273 | 434 |
| 0.10 | 168 | 399 |
| 0.50 | 228 | 423 |

### 4.2. Typical Demands

Given the high level of irrational behavior, which is further accentuated when one takes into consideration mild risk aversion, there is little surprise with the following results.

We next look at the mean behavior. Recall that in each experimental condition (i.e., crossing the seller's five levels of initial belief about the buyer's type with the two levels of the high search cost) consists of 12 pairs of subjects in fixed roles (buyer or seller) playing eight bargaining games. Each bargaining game consists of five periods in which a fixed pair plays the price-taking game as set by the seller. The buyer may rebuff the seller by irrevocably opting for an outside alternative through searching (at a cost).

In order to stabilize the data, we consider the median price demand of any given seller over the eight repetitions of the plays he participated in during the session and in each period (1–5). This eliminates the possible effects of extremely small or extremely high demands.

Table 3 presents the means of these medians over the 12 sellers in each session, by initial sellers' beliefs ($\pi$), high search cost ($H$), and period.

For comparison, Table 4 presents the equilibrium price demands, assuming risk neutrality.

By observing the equilibrium price demands in Table 4, we notice the following:

1. For $\pi = 0.00$, the price demand in each period should be higher when $H = 45$ than when $H = 10$.
2. For $\pi = 0.50$, the price demand in each period should be the same no matter whether $H = 45$ or $H = 10$.
3. For $\pi = 0.01$, 0.05, and 0.10, the price demands are generally lower in the earlier periods than in the later periods (they go up from 55 to either 60 or 95).

*Table 3. Mean medians price demands*

| $\pi\rightarrow$ | *0.00* | | *0.01* | | *0.05* | | *0.10* | | *0.50* | |
|---|---|---|---|---|---|---|---|---|---|---|
| $H\rightarrow$ *Period* | *45* | *10* | *45* | *10* | *45* | *10* | *45* | *10* | *45* | *10* |
| 1 | 56.8 | 46.8 | 61.1 | 52.6 | 51.2 | 43.2 | 56.0 | 46.6 | 53.6 | 49.1 |
| 2 | 56.5 | 45.8 | 59.9 | 49.9 | 51.8 | 43.2 | 56.6 | 46.0 | 54.7 | 49.8 |
| 3 | 56.2 | 45.9 | 59.6 | 49.5 | 53.2 | 43.8 | 56.5 | 45.2 | 55.4 | 49.5 |
| 4 | 55.8 | 47.2 | 60.7 | 51.2 | 53.4 | 42.7 | 56.7 | 46.5 | 55.3 | 49.6 |
| 5 | 56.0 | 47.3 | 60.0 | 50.8 | 53.5 | 42.2 | 56.5 | 45.6 | 55.0 | 49.4 |

*Table 4. Equilibrium price demands*

| $\pi\rightarrow$ | 0.00 | | 0.01 | | 0.05 | | 0.10 | | 0.50 | |
|---|---|---|---|---|---|---|---|---|---|---|
| $H\rightarrow$ Period | 45 | 10 | 45 | 10 | 45 | 10 | 45 | 10 | 45 | 10 |
| 1 | 95 | 60 | 95 | 55 | 55 | 55 | 55 | 55 | 55 | 55 |
| 2 | 95 | 60 | 95 | 55 | 55 | 55 | 55 | 55 | 55 | 55 |
| 3 | 95 | 60 | 95 | 55 | 95 | 55 | 55 | 55 | 55 | 55 |
| 4 | 95 | 60 | 95 | 55 | 95 | 55 | 95 | 55 | 55 | 55 |
| 5 | 95 | 60 | 95 | 60 | 95 | 60 | 95 | 55 | 55 | 55 |

The most striking observation from the data presented in Table 3 is the stability of the means over the five periods. The subjects do not learn! As mentioned in point 3 above, one would have expected some change as the seller learns about the buyer. We do not show it here, but this stability is also noted at the individual levels.

In Table 3, the first two columns on the left describe behavior under the well-known ultimatum conditions (Güth et al., 1982). The results deviate markedly both from rational expectation on the one hand and from traditional ultimatum results on the other hand, in which one expects about 2/3 of the range to be given to the seller. This translates to price of about 97 and 86.

The rest of the columns show that, except for one case (the Period 1 demand of 61.1 when $\pi = 0.01$ and $H = 45$), all means are outside the rational interval. This again shows that rational Bayesian behavior is far too much to expect. It seems that subjects frame the situation they are in quite differently.

The only understandable aspect of Table 3 is the relation between corresponding columns members under equal $\pi$. For a given $\pi$, the higher the search cost, the higher the mean price demand. However, this is a very weak prediction. It is also compatible with the corresponding number of deviations from the rational interval of prices, as we have seen above.

### *4.3. Behavior Change*

We next look at the relations between price demands and acceptance behavior. Consider the correlation between the search behavior and the price demand. We code the search as 1 if the buyer searches and 0 otherwise. Table 5 presents the correlations between search behavior and price demand by the conditions of the experimental design, except that we collapse the data from periods 2 to 5. The goal

*Table 5. Correlations between search behavior and price demands*

| $\pi\rightarrow$ | 0.00 | | 0.01 | | 0.05 | | 0.10 | | 0.50 | |
|---|---|---|---|---|---|---|---|---|---|---|
| $H\rightarrow$ Period | 45 | 10 | 45 | 10 | 45 | 10 | 45 | 10 | 45 | 10 |
| 2–5 | 0.300 | 0.297 | 0.286 | 0.187 | 0.271 | 0.235 | 0.156 | 0.364 | 0.203 | 0.087 |
| 1 | 0.442 | 0.413 | 0.504 | 0.510 | 0.330 | 0.457 | 0.440 | 0.563 | 0.474 | 0.364 |

*Table 6. Search counts*

| $\pi\rightarrow$ | 0.00 | | 0.01 | | 0.05 | | 0.10 | | 0.50 | |
|---|---|---|---|---|---|---|---|---|---|---|
| $H\rightarrow$ Period | 45 | 10 | 45 | 10 | 45 | 10 | 45 | 10 | 45 | 10 |
| 1 | 35 | 34 | 38 | 47 | 28 | 45 | 26 | 41 | 39 | 18 |
| 2 | 26 | 26 | 27 | 37 | 24 | 40 | 18 | 30 | 34 | 23 |
| 3 | 25 | 31 | 25 | 35 | 27 | 35 | 22 | 37 | 38 | 24 |
| 4 | 25 | 31 | 24 | 39 | 26 | 34 | 20 | 30 | 33 | 28 |
| 5 | 11 | 30 | 25 | 37 | 15 | 30 | 13 | 41 | 34 | 24 |

is to observe if there is any change in the response behavior between the beginning and the rest of the game.

We do not test any formal significance. It is quite noticeable that (a relatively) high price demand is met with a tendency for the buyer to decide to search since all correlations are positive. This is, of course, not an indication of formal updating, since, as we have seen, most demands should have been accepted anyway.

Table 6 shows the total number of searches in each period by $\pi$ and $H$. The base is always 96 games played in each cell. We see that buyers initially tend to reject and search. This seems less natural when the buyer is equally likely to be weak or strong (the last two columns when $\pi = 0.50$).

## 5. DISCUSSION

We divide our discussion into two parts: the sources of difficulties in playing the game "correctly" and some modeling issues. The latter arises from the observation that sellers violated the most basic prediction of the rational model, even when there was no uncertainty as to the buyer's type ($\pi = 0$). Since sellers' asking prices were so low vis-à-vis the rational prices, buyers were "deprived" of the opportunity to strategically reject prices in early rounds in order to manage the seller's beliefs about the their type. The low asking prices were akin to the seller stating that he believes that the buyer's search cost is low, even though the objective probability of a low-cost type was much less than 50% in most cases.

### *5.1. The sources of difficulties in playing the game "correctly"*

We have seen that even superficial analysis of the bargaining game allows us to conclude that there is only meager evidence to indicate some support for rational behavior in this setting (see the correlation data above). Of course, what is rational is not absolutely clear. Let us check our assumptions again. The bargaining interval, which is the interval in which price demands must fall in order to satisfy the minimum price that the seller may demand and the maximum price beyond which the buyer will not settle, was derived from the distribution of the search price. The derivation assumes that it is common knowledge that all participants are risk neutral. The existence of a commonly known bargaining interval assumes that the participants would rather gain more than less out of playing the game.

Now we expressly admit that there is no reason in terms of game theory to require the players to figure out the bargaining interval in such a complex manner. We know that an average person's understanding of probability has its limitations, which may have nothing to do with game behavior. In our setting, in addition to the strategic uncertainty and the uncertainty (from the seller's point of view) as to the buyer's type, a third layer of uncertainty existed as to the realization of the actual outside price. We now believe that this added complexity might have unnecessarily clouded the strategic nature of the game. In fact, a next iteration of this experimental program intends to simplify the bargaining game by eliminating the probabilistic nature of the outside option. Consider, for example, a game similar to the one we have implemented except that the unknown (to the seller) buyer's type is characterized by her outside option. Now the bargaining interval is easy to figure out! Any solution of the original game is a solution of this proposed one. Eliminating one source of uncertainty might re-focus subjects' attention on the reputation nature of the scenario.

Another problem is hinted at by the extreme case when $\pi$ is 0. In this case, the game is simply an ultimatum game with an uncertain outside option to the responder. This game was investigated by Zwick and Lee (1999) in their No Recall condition. They report that the classical game-theoretical model can account for some (but not all) of the behavioral regularities. In line with recent developments in behavioral decision theory and game theory, which assume bounded rationality and preferences over the relative division of a surplus, Zwick and Lee (1999) found that subjects followed simple rules of thumb and distributional norms in choosing strategies, which are reflected in the behavioral patterns they observed. We know that ultimatum games are notoriously difficult to play rationally. Some of the other $\pi$ values in our design may not be sufficiently different from 0. So, the framing of the bargaining situation with an ultimatum skeleton obviously, in hindsight, makes our data difficult to interpret.

It is not clear how to re-frame the game to avoid the ultimatum aspect. One way is to change the number of periods played and the information given to the seller.

### *5.2. The Modeling Issues*

Needless to say, the rational reasoning discussed in the introduction is not sufficiently powerful or intuitive to derive behavior in practice. In fact, the game we described requires some elaboration and additional scaffolding to obtain more natural predictions. To see why the data are not sufficient, we note that if we assume that the buyer may only choose to accept or search by some deterministic rule, we know that this rule must obey the property that if she accepts any price, then she must accept any lower price. And since the buyer is one of two types, once she accepts a price not on the boundary of the bargaining interval, she reliably signals her identity as a high-cost buyer. Since a deterministic rule is not sufficient to uncover what seems to be natural behavior, a probabilistic rule is one approach to enrich buyer's behavior.

One such operationalization is given by Lee (2000), except that in his approach both the buyer and the seller are equipped with probabilistic strategies. In any approach that relies on such rules, both players need to be commonly cognizant of the extension to the given game. So, aside from the computational complexity, the players need some common understanding of the extension. But are these "computational complexities"? In the 1960s, in a prelude to Kahneman and Tversky's work that initiated the heuristics and biases approach to decision theory, it was discovered that quite consistently humans are notoriously conservative in updating their degrees of belief (Edwards, 1968). Now, the work of Lee, as well as others, that relies on mixed extensions of sequential games also relies explicitly on Bayesian updating. While the experimental work discussed by Edwards (1968) explicitly described the probabilistic scenario in which people fail, the modeling in this case is implicit and the players need to commonly agree upon their existence and then update.

We conjecture that, in a plainer environment, one can at best expect that people are quite a bit slow in achieving the reputation they desire. But at least this may be experimentally attainable.

## ACKNOWLEDGMENT

The author acknowledges the financial support received from the Hong Kong Research Grants Council for this research (Project Number: CUHK4076/98H).

## NOTES

[1] $100 Hong Kong dollars. The exchange rate between Hong Kong and US dollars is approximately 7.8 to 1.

[2] The amount of the cash prize was very attractive to students considering that the hourly wage for an on-campus job was about $50.

## REFERENCES


Camerer, Colin and K. Weigelt, (1988). "Experimental Tests of a Sequential Equilibrium Reputation Model." *Econometrica*, 56(1), 1–36.

Edwards, W., (1968). "Conservatism in Human Information Processing," in *Formal Representation of Human Judgment*, B. Kleinmuntz ed., John Wiley, New York.

Güth, W., R. Schmittberger, and B. Schwarze, (1982). "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior & Organization*, 3(4), 3867–388.

Kleinmuntz, B., (1968). *Formal Representation of Human Judgment*, John Wiley, New York.

Kreps, D., and R. Wilson (1982). "Reputation and imperfect information." *Journal of Economic Theory*, 27, 253–279.

Lee, Ching Chyi, (2000). "Reputation in Repeated Bilateral Trading with Outside Alternatives." *Working Paper*, The Chinese University of HK.

Raiffa, H., (1982). *The Art and Science of Negotiation.* Harvard University Press, Cambridge MA.

Zwick, Rami and Ching Chyi Lee, (1999). "Bargaining and search: An experimental study." *Group Decision and Negotiation*, 8(6), 463–487.