

Bargaining and Search: An Experimental Study*

RAMI ZWICK

Department of Marketing, The Hong Kong University of Science and Technology

CHING CHYI LEE

Department of Decision Sciences & Managerial Economics, The Chinese University of Hong Kong

Abstract

We study experimentally two versions of a model in which a buyer and a seller bargain over the price of a good; however, the buyer can choose to leave the negotiation table to search for other alternatives. Under one version, if the buyer chooses to search for a better price, the opportunity to purchase the good at the stated price is gone. Under the second version, the seller guarantees the same price if the buyer chooses to return immediately after a search (presumably because a better price could not be found). In both cases, the buyer has a fairly good idea about what to expect from the search, but because the search is costly, he has to weigh the potential benefits of the search against its cost. It turns out (theoretically) that adding search to a simple bargaining mechanism eliminates some unsatisfactory features of bargaining theory.

Our experiment reveals that the 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, we find that subjects follow simple rules of thumb and distributional norms in choosing strategies, which are reflected in the behavioral consistencies observed in this study.

Key words: bargaining, search, game theory, experimental economics

Introduction

A recent survey by The Dohring Co. (PR News Wire via DowVision, February 14, 1995) found that when the concept of a one-price selling format¹ was explained to the respondents, 15 percent saw no advantage to the plan, and 88.1 percent of the respondents said they would visit a one-price dealership, obtain the lowest price quote on a vehicle and visit other dealerships to try to find a better price or deal.

Clearly, the success of a no-haggling selling format is contingent on demand and buyers' outside opportunities (Wernerfelt 1994). For example, the success of the Saturn one-price strategy is attributed, among other factors, to the fact that Saturn dealers' geographic territories are so large that it is too expensive for consumers to search for a better deal, and their products are in high demand – so dealers usually get sticker prices even if negotiation is allowed (Advertising Age, March 22, 1993).

*This paper is part of a research project on negotiation that is financed by the Hong Kong Research Grants Council (Project No. HKUST6239/97H) and by the Chinese University of Hong Kong.

Here, we study experimentally a model motivated by the above illustration in which a buyer and a seller bargain over the price of a good; however, the buyer can choose to leave the negotiation table to search for other alternatives. This model corresponds stylistically to the following marketing scenario.

A buyer enters a shop in a market and indicates his interest in a particular item. After some haggling, the seller announces her “rock bottom price” and challenges the buyer to search for a better price. Under one version of the story (the No Recall version²), the seller puts extra pressure on the buyer by declaring (credibly) that the price is valid only for now. That is, if the buyer chooses to search for a better price some place else, the opportunity to purchase the good in the seller’s shop at the stated price will be gone.³ Under the second version of the story (the Recall version), the seller invites the buyer to search and guarantees the same price if the buyer chooses to return immediately to the shop after a search (presumably because a better price could not be found). In both cases, the buyer has a fairly good idea about what to expect from the search, but because the search is costly, he has to weigh the potential benefits of the search against its cost.

The above scenarios combine two important elements of marketing activities, namely, bargaining and search. Yet, despite the obvious connection, almost all previous empirical research has explored either bargaining or search independently, or bargaining with certain outside options⁴ (Binmore, Shaked and Sutton 1989; Weg, Zwick and Rapoport 1996). Only Brucks and Schurr (1990) have looked at the two processes simultaneously, though they investigated only the buyers’ behaviors because a computer program simulated sellers.

The following briefly summarizes the major theoretical and experimental evidence directly relevant to our paper. For a comprehensive review of the experimental bargaining literature, we refer the reader to Roth (1995) and, for search literature, see McMillan and Rothschild (1994).

Two key issues that characterize the goals of most game theoretical models of bargaining are (1) the allocation of bargaining outcome and (2) delay.

In general, delay in reaching a bargaining agreement can often be ascribed theoretically to the need for bargainers to communicate their private information. There are many incomplete information models of bargaining that address this issue.⁵ Without the need for such a communication, bargainers are generally expected to be able to reach an agreement immediately under complete information, and this point can be verified by the fact that most of the complete information models of bargaining predict immediate resolution of bargaining.⁶

Ever since Nash’s pioneering work in bargaining (Nash 1950), bargaining outcome has been found to be closely related to the bargainers’ reservation prices (disagreement outcome or status quo point). Nash modeled bargaining outcome explicitly as a function of players’ reservation prices. In general, bargainers with “better” reservation prices are expected to get a better outcome. Although Nash himself did not consider explicitly how the reservation price of a bargainer might be determined, recent attempts in relating sequential bargaining models to Nash’s static model have revealed some crucial factors that determine its value. The reservation price of a bargainer, for example, may depend on his time preference (Rubinstein 1982), his alternative options (Shaked and Sutton 1984; Binmore, Shaked and Sutton 1989; Shaked 1987), or his ability to “search” for alternative offers (Wolinsky 1987; Chikte and Deshmukh 1987; Muthoo 1995). In

particular, in the bargaining and search model of Chikte and Deshmukh that is closely related to our paper, the authors found that search costs are monotonically related to players' reservation prices – that is, a smaller search cost leads to a higher reservation price for the seller and a lower reservation price for the buyer. As a result, they conclude that a decrease in a player's search cost can increase his or her payoff without increasing the other player's payoff.

Although these general results are theoretically appealing, they do not, for the most part, characterize bargaining realities. Experimental evidence and field surveys suggest that disagreements and costly delays are pervasive even in situations that eliminate the most obvious potential sources of incomplete information (see, for example, Forsythe, Kennan and Sopher 1991). Second, bargainers are not always sensitive to the same intrinsic or extrinsic characteristics that theoretically determine their reservation prices and they demonstrate similar biases and heuristics discovered in individual decision making environments (Bazerman and Neale 1992; Zwick and Weg 1996).

Given the above evidence, it is natural to pursue a theoretical model that, although preserving the most desired properties of minimum rationality on the part of the bargainers, can, nevertheless, account for some of the above described deviations. It turns out that examining the bargaining foundations of search theory, or adding search to bargaining theory, eliminates some unsatisfactory features of both search theory and bargaining theory (McMillan and Rothschild 1994).

A simple bargaining and search model

The experiment reported in this paper is based on two simplified versions of Lee's bargaining and search model (Lee 1994). These simple versions were designed to (1) retain the main features of the original model and yield qualitatively similar outcomes, (2) be simple enough to allow experimental manipulations, and (3) have a simple structure to reduce the potential breach of the bargainers' rationality.

The model goes as follows: there are two bargainers – a seller and a buyer, who are both risk neutral. The seller owns a good that she values at 0.⁷ The buyer is willing to pay up to 1 for that good. The bargaining starts with the seller making an offer, p , to the buyer. The buyer can then either accept or search. Acceptance of p by the buyer ends the bargaining and yields a payoff of p to the seller and $1 - p$ to the buyer. If the buyer decides to search, he pays a cost c ⁸ and draws an outside offer x from a uniform distribution with support on $[0, 1]$. After x is observed, the buyer accepts an offer and the negotiation ends. Everything is assumed to be common knowledge.

There are two variations of this basic model relating to what happens if a search takes place.

No Recall: In this version, the buyer is not allowed to recall the seller's offer once it has been rejected.⁹ That is, after the search, the buyer can only accept the outside offer, x . In this case, the payoffs after realization of the outside offer are $(1 - x - c)$ for the buyer and 0 for the seller.

Recall: In this model, the buyer is allowed to recall the seller's offer after the search. Hence, after the search (if it takes place) the buyer can decide whether to accept the outside offer, x , or to recall the seller's offer, p . If the outside offer is accepted, the payoffs are $(1 - x - c)$ for the buyer and 0 for the seller. If the seller's offer is accepted, the payoffs are $(1 - p - c)$ for the buyer and p for the seller.

The equilibrium

We now describe the *subgame perfect equilibrium* for each of the two versions of the basic model.

No Recall: Since the buyer cannot recall the seller's offer, the buyer's expected payoff if he rejects the seller's offer is $-c + (1 - E(x)) = \frac{1}{2} - c$, where $E(x)$ is the expected value of outside offer. Hence, the buyer should accept any offer from the seller when it is less than or equal to $\frac{1}{2} + c$ (thus getting $\frac{1}{2} - c$ or more) and reject an offer otherwise. To ensure acceptance, the seller's offer must be less than or equal to $\frac{1}{2} + c$ (otherwise, she gets nothing). Of course, knowing that a price equal to $\frac{1}{2} + c$ will be accepted for certain, the seller should not offer anything strictly less. Hence, in equilibrium, the seller offers $\frac{1}{2} + c$ and the buyer accepts. The equilibrium payoffs are $\frac{1}{2} + c$ for the seller and $\frac{1}{2} - c$ for the buyer. Notice that the equilibrium of this model verifies the two standard results of bargaining and search models with complete information. First, in equilibrium, agreement will be reached immediately without delay. Second, the players' payoffs are monotonically related to the buyer's search cost.

Recall: Since the buyer can recall the seller's offer after the search, his expected payoff if he searches is $-c + (1 - E(x \wedge p))$, where " $\wedge$ " is the minimum operator and $E(x \wedge p)$ is the (conditional) expected value of the "best" offer given p . Hence, the buyer will accept any seller's offer p satisfying $1 - p \geq -c + (1 - E(x \wedge p))$. It can be shown that if x is uniformly distributed on $[0,1]$, at $p = (2c)^{\frac{1}{2}}$, the above inequality becomes an equality. Hence, the buyer will accept all seller's offers less than or equal to $(2c)^{\frac{1}{2}}$ and search otherwise.¹⁰ The seller can offer $p = (2c)^{\frac{1}{2}}$ (she has no reason to offer anything less), in which case the offer gets accepted, and she gets $(2c)^{\frac{1}{2}}$ and the buyer gets $1 - (2c)^{\frac{1}{2}}$. The seller can also offer a price greater than $(2c)^{\frac{1}{2}}$, in which case the buyer will search. As a result, the seller's expected payoff for making such an offer is $p[1 - F(p)]$ (this is because, with probability $1 - F(p)$, the buyer will get an outside offer greater than p and hence accept p and the buyer's is $-c + (1 - E(x \wedge p))$). For the case of uniform distribution with support on $[0,1]$, the maximum of $p[1 - F(p)]$ occurs at $p = \frac{1}{2}$. Hence, the seller faces the decision regarding whether to offer $p = \frac{1}{2}$ and expect to get $\frac{1}{4}$ or to offer $p = (2c)^{\frac{1}{2}}$ and get $(2c)^{\frac{1}{2}}$ for certain. Simple calculation shows that the seller should offer $p = (2c)^{\frac{1}{2}}$ if $c \geq 1/32$ and offer $p = \frac{1}{2}$ otherwise. Hence, this model has two types of equilibria – search equilibrium and no search equilibrium. Search equilibrium arises when $c < 1/32$, whereas no search equilibrium arises when $c \geq 1/32$. Hence, we see here that delay in reaching

an agreement can occur even with complete information (when the search equilibrium arises). In addition, since there is a sudden jump in the buyer's expected payoff function (at $c = 1/32$), the buyer's expected payoff does not monotonically decrease in c .

Contrary to the standard results of previous complete information bargaining models, the above model predicts that (i) complete information renders no guarantee for immediate resolution of the bargaining and (ii) bargaining outcomes are not always monotonically related to the buyer's search cost. The basic intuition for these results is that when the buyer's search cost is very small, it may be too costly for the seller to discourage the search (by asking a low price). Hence, instead of making a "low" offer to deter the search, the seller chooses to charge the buyer a "high" price, knowing that the buyer will search and hoping that this "high" offer will be accepted later if the buyer's search turns out to be unfruitful.

Hence, theoretically, we know that allowing or disallowing recall of past offers, or changing the buyer's search cost, could lead to qualitatively quite different results. However, it is not clear if consumers can recognize the effect of these differences on their strategic power and hence act accordingly.

Hypotheses

We list below several testable hypotheses implied by the model – H1–H3 concern the sellers' behaviors; H4 concerns the buyers' behaviors; and H5 concerns both.

- H1:** In the No Recall mode, asking prices should increase with the buyers' search costs.
- H2:** In the Recall mode, the lowest asking prices are expected when $c = 1/32$. Asking prices should be higher when search costs are lower or higher than $1/32$.
- H3:** Asking prices should be higher in the No Recall compared to the Recall mode for all search cost values.¹¹
- H4:** Buyers should adopt a cut-off rule. That is, in the No Recall mode, buyers should accept any price that is lower than $\frac{1}{2} + c$, and, in the Recall mode, buyers should accept any price that is lower than $(2c)^{1/2}$.

In general, equilibrium solutions may be expected to account for interactive behavior if the solutions are transparent so that players may adopt them naturally, or if, as a result of much practice with the game (or similar games that allow positive transfer of experience), behavior converges to an equilibrium solution. The above model may be considered conceptually transparent but requiring some non-trivial computations. Consequently, we expect subjects to approximate the equilibrium rationality only after considerable experience in the negotiation environment.

- H5:** Buyers' and seller' behavior should converge to the equilibrium solution as they gain more experience in the negotiation environment.

In the next section we report on a study designed to test the above hypotheses in a marketing scenario where modes and search costs are orthogonal to each other, thus providing clear attribution of causes and effects.

Method

Subjects

Two hundred thirty-four male and female subjects, who were mainly undergraduate business students in groups of ten to twenty, participated in a role-playing session that lasted about 60 minutes.¹² 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.

Experimental design

Each of the bargaining games consisted of bargaining on a surplus of \$100¹³ with an uncertain outside option to the buyer uniformly distributed on the range [0, 100] using the trading rules described above.

We used a 3 (Search Cost) $\times$ 2 (Recall / No recall) $\times$ 15 (trials) design. The first two factors were between subjects and the last was within subjects. The search cost was \$2, \$6, or \$12. Note that \$2 is below the predicted threshold of $1/32$ ($100 \times 1/32 = 3.125$) required to generate search equilibrium in the Recall mode. Consequently, the predicted asking prices in the No Recall mode are \$52, \$56 and \$62 for search costs \$2, \$6, and \$12, respectively. The corresponding asking prices in the Recall mode are: \$50, \$35, and \$49 (rounded to the nearest dollar). In particular, note that prices increased monotonically in the No Recall mode, whereas the lowest price in the Recall mode is expected when search costs equal \$6 (see hypotheses H1 and H2).

Subjects assumed the same role (a buyer or a seller) in all 15 trials in a session and faced different anonymous opponents during each trial. Table 1 presents the number of subjects in each of the six experimental conditions.

During each trial, the seller is asked to sell an indivisible good, which has no value to her except its selling price. The value of the good to the buyer is \$100. Bargainers know both reservation prices. The game proceeds as follows: the seller announces a selling price for the good (the asking price). The buyer then has the following options:

Table 1. Number of subjects in each of the experimental conditions

	Search Cost			
	\$2	\$6	\$12	ALL
No Recall	42	32	44	118
Recall	42	34	40	116
All	84	66	84	234

1. Accept the asking price, thereby terminating the game.
2. Search for an alternative price. In this case, the buyer has to pay a search cost, and a price is randomly generated (the outside offer) from the range [0, 100].
3. After learning about the outside offer, the buyer must accept this offer in the No Recall mode or can accept either the asking price or the outside offer (but not both) in the Recall mode.

Search cost as well as the price generated through the search (if any) is known to both bargainers.

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 terminals. No other communications were allowed.

To motivate the subjects, they were informed that they would be paid their net payoff from one randomly selected game. In addition, each subject was paid \$50 for participation. On the average, subjects earned \$89.77 for a session.¹⁴

Results

Sellers' behaviors

Figure 1 presents the frequency distribution of asking prices by Search Cost, Recall Condition (Recall and No Recall), and Trial. The bottom and top edges of the box on each trial are located at the sample 25th and 75th percentiles. The vertical lines extend from the box as far as the data extend, to a distance of at most 1.5 interquartile range. Any value more extreme than this is marked by a dot. A line joins the median points of the boxes. A reference line is drawn horizontally at the \$50 price.

Prices were subjected to a repeated measure ANOVA with Search Cost, Recall Condition and Trial as the independent variables. The results of the overall analysis as well as the analysis within Cost and Recall levels are presented in Table 2.

The following can be observed by skimming through the figures and is confirmed by the statistical analysis reported in Table 2.

- (1) Cost and Recall conditions affect prices in a significant way. However, these are not simple effects as is evidenced by the significant interaction between the Cost and Recall conditions (see the test of hypotheses for between subjects effects in Table 2).
- (2) Prices do not seem to follow a consistent pattern across trials (see Table 2 for an overall insignificant trial effect). The exception is the No Recall \$2 Cost level where prices decrease consistently with experience (see the significant interaction between the Trial and Recall conditions for Search Cost \$2). A simple repeated measure ANOVA within this cell reveals a significant trial effect ($F(14,7) = 3.55, p < 0.049$). H5 predicts that prices would converge on the equilibrium solution as players gained

more experience with the task. The fact that prices do not seem to follow a consistent pattern across trials does not immediately falsify H5 since subjects might have behaved according to the equilibrium prediction from the start. We will discuss H5 in greater detail in the sections on adaptive behavior below.

- (3) In the No Recall mode, prices increase with the Search Cost. They are highest when the search costs \$12 and lowest when the search costs \$2 (see Table 2 for a significant cost effect in the No Recall mode). A Scheffe test for multiple comparisons (on the overall mean prices across all 15 trials) revealed that all three cost levels are significantly different from each other at the 5% significant level.¹⁵ This supports H1.
- (4) In the Recall mode, prices are highest when the search costs \$2 and lowest when the search costs \$6 (see Table 2 for a significant cost effect in the Recall mode). A Scheffe test for multiple comparisons (on the overall mean prices across all 15 trials) revealed that prices in the \$6 cost level are significantly lower than in the




Figure 1. Asking prices

other two cost levels, but prices in cost levels \$2 and \$12 do not differ significantly from each other at the 5 percent level.¹⁶ This supports H2.

- (5) When search costs are \$2 and \$6, prices are higher in the Recall compared to the No Recall mode. Prices are about the same when the search cost is \$12 (see *p*-values in Table 2). This contradicts H3.

Table 2. MANOVA test criteria and F statistics for the hypothesis of no source effect on asking price

	Source	Wilks' Lambda	F	Num DF	Den DF	Pr > F
	TRIAL (T)	0.83	1.44	14	95	0.15
	T*RECALL ^b	0.92	0.61	14	95	ns ^c
	T*COST	0.79	0.85 ^a	28	190	ns
	T*RECALL*COST	0.76	0.99 ^a	28	190	ns
	By Cost					
\$2	TRIAL (T)	0.54	1.63	14	27	0.13
	T*RECALL	0.44	2.46	14	27	0.02
\$6	T	0.57	0.82	14	15	ns
	T*RECALL	0.54	0.91	14	15	ns
\$12	T	0.74	0.69	14	27	ns
	T*RECALL	0.79	0.50	14	27	ns
	By Recall					
No	TRIAL (T)	0.67	1.39	14	40	0.20
	T*COST	0.44	1.42	28	80	0.11
Yes	T	0.74	1.04	14	42	0.44
	T*COST	0.69	0.61	28	84	ns
Tests of Hypotheses for Between Subjects Effects						
	Source	DF	SS	F	Pr > F	
	RECALL	1	10889.08	13.89	0.000	
	COST	2	5001.00	3.19	0.045	
	RECALL*COST	2	10260.82	6.54	0.002	
	Error	108	84670.21			
	By Search Cost					
\$2	RECALL	1	18326.65	23.18	0.000	
	Error	40	31618.63			
\$6	RECALL	1	3034.13	3.81	0.061	
	Error	40	22322.16			
\$12	RECALL	1	62.98	0.08	ns	
	Error	40	30729.42			
	By Recall					
No Recall	COST	2	13349.97	8.34	0.001	
	Error	53	42413.28			
Recall	COST	2	4824.97	3.14	0.051	
	Error	55	42256.94			

^aApproximate *F* value. ^bRecall refers to the manipulated variable of Recall vs. No Recall option.

^cns = not significant.

Sellers' adaptive behaviors

Although there is no overall trial effect, a closer look at the individual seller's behavior reveals a consistent adaptation process that can be described as a search for the highest acceptable price via an anchoring and adjustment process.

Following Ochs and Roth (1989), Zwick et al. (1992) and Weg et al. (1996), we examined the current trial asking price (p_t) in relation to the previous trial asking price (p_{t-1}), the buyer's response made to the latter (accept or search), and the outside offer (x_{t-1}) if search occurred in trial $t-1$. This information, aggregated over all sellers, is presented in Table 3.

Six of the 117 sellers in the study (4 in the No Recall and 2 in the Recall mode) demanded the same price in all 15 trials (4 of the 6 demanded \$50 in each trial)¹⁷. Moreover, no other sellers demanded the same price in the last five trials.

We define an adaptive (monotonic) price-setting behavior to be a non-decrease of the current trial asking price (p_t) compared to the immediate predecessor trial asking price (p_{t-1}) when the latter is accepted and a non-increase when it is rejected (immediately or after a search). We regard the highest acceptable price as an individual buyer index, not unlike the level of aspiration that may fluctuate during the session due to the experience gained by the buyers.¹⁸ Additionally, the sellers' knowledge that their opponents may change from trial to trial may cause some sellers to violate one or more of the above conditions on some of the trials.

We tested the "adaptive behavior" hypothesis against an alternative "random" hypothesis stipulating that $p_t < p_{t-1}$, $p_t = p_{t-1}$, or $p_t > p_{t-1}$, with equal probabilities. Under the "random" hypothesis, the probability of less than 2 violations in 14 games is 0.027, and less than 3 violations is 0.105.

Table 3. Adaptive Behavior^a: Asking prices and Buyers' search rates^b

Buyer's Reply to p_{t-1}		No recall		Recall	
		p_t Rate	Search Rate ^b	p_t Rate	Search Rate ^b
Accepted	$p_{t-1} < p_t$	0.55	0.61	0.55	0.85
	$p_{t-1} = p_t$	0.33	0.58	0.30	0.48
	$p_t < p_{t-1}$	0.12	0.16	0.15	0.39
Search					
$p_{t-1} < sp_{t-1}$	$p_{t-1} < x_{t-1} < p_t$	0.02	0.67	0.04	1.00
	$p_{t-1} < p_t < x_{t-1}$	0.22	0.77	0.46	0.91
	$p_{t-1} = p_t < x_{t-1}$	0.27	0.65	0.21	0.54
	$p_t < p_{t-1} < x_{t-1}$	0.48	0.32	0.29	0.43
$sp_{t-1} < p_{t-1}$	$x_{t-1} < p_{t-1} < p_t$	0.08	0.87	0.20	1.00
	$x_{t-1} < p_{t-1} = p_t$	0.15	0.40	0.12	0.80
	$x_{t-1} < p_t < p_{t-1}$	0.71	0.32	0.57	0.33
	$p_t < x_{t-1} < p_{t-1}$	0.06	0.20	0.11	0.50

^aThe table presents data from all trials except the first one. ^bBuyers' search rates are computed for only the 41 buyers who did not exhibit a consistent threshold strategy.

Of the 58 sellers in the Recall mode, 20 never violated the adaptive hypothesis,¹⁹ 12 violated it once and 8 twice. For 15 of the 20 sellers who violated the hypothesis once or twice, the nature of the violation(s) was that they asked for a higher price after their offer was rejected in the previous trial and $x_{t-1} < p_{t-1}$. Such behavior can be described as a “gambler fallacy.” That is, the seller might believe that the next search, if taken by the buyer, would not yield as low a price as before.

Of the 59 sellers in the No Recall mode, 13 never violated the adaptive hypothesis,²⁰ 13 violated it once and 15 twice. For 19 of the 28 sellers who violated the hypothesis once or twice, the nature of the violation(s) was that they asked for a higher price after their offer was rejected in the previous trial and $x_{t-1} > p_{t-1}$. Such behavior violates the adaptive hypothesis but can be explained as anchoring on the outside offer rather than on the previous asking price for the downward adjustment.

Altogether, the price setting behavior of about 70 percent of the sellers in both Recall and No Recall modes can be well described as a search for the highest acceptable price and is in accordance with the group data reported in Table 3.²¹

Buyers' behaviors

Table 4 presents search rates in each of the experimental conditions. Clearly, buyers searched more often in the Recall compared with the No Recall cases (0.73 versus 0.43). This is intuitively obvious because, in the Recall mode, buyers can recall and accept a seller's offer when it is better (lower) than the outside offer, thus promoting more searches.

A perfect adherence to the model would imply no searching in all conditions except when search cost is \$2 and recall is possible. However, the adherence of the buyers to the theory can not be measured simply by the overall search rates because, as we have reported previously, the sellers did not follow the exact equilibrium path. Consequently, a better measure of buyers' rationality (H4) should take into account the actual asking prices to which the buyers have reacted.

H4 predicts that buyers will adopt a cut-off rule. That is, in the No Recall mode, buyers should accept any price that is lower than $50 + c$ (where c is the search cost), and in the Recall mode buyers should accept any price that is lower than \$20, \$34.6, and \$49 for search costs \$2, \$6, and \$12, respectively.

Table 4. Search rates

	Search Cost					
	2		6		12	
	Recall	No Recall	Recall	No Recall	Recall	No Recall
Total No. of Asking Prices	315	315	255	234	300	330
No. of Buyers' Searches	242	172	189	70	208	140
Search Rate	0.77	0.55	0.74	0.30	0.69	0.42
Search Rate: Recall Mode	0.73					
No Recall Mode		0.43				

The above predictions were tested in three stages. First, we identified all those buyers who used a cut-off rule. A buyer is said to use a cut-off rule if there exists a price (p^*) such that all asking prices below p^* were accepted and all asking prices above p^* were followed by a search. When ordering the asking prices encountered by a buyer from lowest to highest, the corresponding buyer's reply should depict a sequence with two runs – acceptances (that can be of the length zero) followed by searches. We also identified those buyers who used an almost perfect cut-off rule, defined as a pattern that can be transformed into a perfect cut-off pattern by ignoring one reply. Next, we investigated to what extent these cut-off prices agree with perfect rationality requirements (H4). Finally, we investigated the behavior of buyers who were not classified as adopting a perfect or almost perfect cut-off rule.

Table 5 presents the number of buyers who adopted a perfect or almost perfect cutoff rule in each condition.²² Although, by definition, buyers who accept all prices or always search follow a cut-off rule, they are presented as special groups in Table 5.²³

Cut-off rule

Out of 117 buyers in the study 76 (64.9%) used a consistent (or almost consistent) cut-off rule. However, buyers in the Recall mode were more likely to adopt this strategy (75.9% and 54.2% of the buyers in the Recall and No Recall modes, respectively). However, the difference is attributed to the number of subjects who always (or almost always) searched, with 13 in the Recall mode and none in the No Recall mode.

Figure 2 presents frequency distributions of cut-off prices.²⁴ The dashed lines connect the predicted perfect rationality thresholds (H4). Cut-off prices were subjected to ANOVA with Cost and Recall conditions as the independent variables. The interaction between the Cost and Recall conditions is significant as is evident from the figure ($F(2, 68) = 3.29, p < 0.05$). The interaction is evident in a significant cost effect within the No Recall mode ($F(2, 28) = 7.81, p < 0.002$) and a non-significant cost effect within the Recall mode ($F(2, 41) = 0.05, ns$). In the No Recall mode, the mean cut-off prices in the \$2 cost condition are significantly lower than in the other two cost conditions ($t(20) = 3.34, p < 0.001$; $t(19)$

Table 5. Number of buyers whose behavior can be described by a simple cut-off rule

Pattern	No Recall (N = 59)	Recall (N = 58)	ALL (N = 117)
Accepted all prices	2	2	4
Accepted all prices expect one ^a	2	1	3
Cut-off Rule	18	14	32
Imperfect Cut-off Rule	10	14	24
Always Searched	0	12	12
Always searched expect one ^b	0	1	1
All simple pattern	32	44	76

^aThe non-accepted price is not the highest price. ^bThe accepted price is not the lowest price.

= 1.99, $p < 0.051$, for comparing (\$2 and \$6) and (\$2 and \$12), respectively) but mean cut-off prices do not differ between the \$6 and \$12 cost levels. Similarly, the Recall condition is significant only within the \$12 cost level ($F(1, 22) = 7.51$, $p < 0.02$).

Contrary to H4 (see Figure 2) cut-off prices did not increase with the search cost level (except for the lower cut-off prices in the No Recall \$2 cost condition). Also, cut-off prices were not higher in the No Recall compared to the Recall mode (except for the \$12 cost level). In the No Recall mode, buyers adopted lower cut-off prices than predicted at all three cost levels. That is, buyers rejected prices that should have been accepted according to perfect rationality. In the Recall mode, almost all buyers in the \$2 and \$6 cost levels adopted cut-off prices that were higher than predicted.²⁵ That is, these buyers accepted prices (immediately) that should have been followed by a search. In the \$12 cost condition, most buyers (ten out of fifteen) used cut-off prices that were lower than predicted.

Non-cutoff rule

The behavior of 41 out of 117 buyers was not consistent with a stable cut-off rule. Their behavior is characterized by conditions where a price was accepted whereas a lower price was followed by a search (and this irregularity occurred more than once). Clearly, these




Figure 2. Distributions of cut-off prices.

buyers reacted to other signals in the decision environment in addition to the current asking price. Table 3 presents the empirical distributions of search rates on trial t (Search Rate) as a function of the ordinal relationship between the following three variables: (1) asking price in the current trial (p_t); (2) asking price in the previous trial (p_{t-1}); and (3) outside offer on trial $t-1$ (x_{t-1}) if such a search took place.²⁶

Findings:

- (1) In almost all combinations, the search rate is higher in the Recall compared with the No Recall mode.
- (2) The order relationship between the current (p_t) and the previous asking price (p_{t-1}) clearly affected the decision to search in trial t . If the previous trial asking price was accepted and the current asking price was higher ($p_{t-1} < p_t$), then a search followed in 61% (85%) of the time in the No Recall (Recall) mode. However, if the current asking price was lower ($p_t < p_{t-1}$), then a search followed in 16% (39%) of the times in the No Recall (Recall) mode. When the asking price was the same, a search occurred in about half the times. The corresponding rates when a search took place in the previous trial are 0.79 (0.94) and 0.32 (0.37) for $p_{t-1} < p_t$ and $p_t < p_{t-1}$ respectively, in the No Recall (Recall) mode.
- (3) When the buyer searched in trial $t-1$, the decision to search again in trial t was not only a function of the ordinal relationship between the current (p_t) and previous asking prices (p_{t-1}), as reported above in point 2, but was also affected by the ordinal relationship between the outside offer (x_{t-1}) and the asking price in trial $t-1$ (p_{t-1}). Everything else equal, the search rate in trial t is higher if the previous trial's search yielded a worse price than the seller's asking price ($x_{t-1} > p_{t-1}$) compared with when it yielded a better price ($x_{t-1} < p_{t-1}$). The comparable search rates are 0.53 and 0.37 in the No Recall mode, and 0.69 and 0.48 in the Recall mode.²⁷ These findings can be interpreted as an exhibition of the gambler fallacy. That is, after a bad draw, a better one is expected and vice versa.

Payoffs

Table 6 presents buyers' and sellers' payoffs by Search Cost, Recall condition and buyers' search strategy.

Sellers

Sellers' payoffs were subjected to a repeated measure ANOVA with Search Cost, Recall and Trial as the independent variables.²⁸ The following can be observed by scanning Table 6. The sellers' mean profits in almost all conditions are within a narrow range of \$27.34 to \$30.77. The exception is the No Recall, \$2 Search Cost condition where the mean profit is much lower (\$18.31). The Cost and Recall conditions interacted to affect the sellers' payoffs ($F(2,108) = 6.04, p < 0.003$). Whereas the Search Cost did not affect

Table 6. Payoffs

			Buyer Type			Seller
Cost			Cut-off Rule	Non Cut-off Rule	ALL	
No Recall	\$2	Mean	53.30	54.43	53.84	18.31
		Std	20.92	22.41	21.62	20.90
	\$6	Mean	53.21	56.11	54.09	30.77
		Std	15.40	22.01	17.67	21.37
	\$12	Mean	43.81	46.65	45.49	28.86
		Std	19.56	21.42	20.70	25.86
Recall	\$2	Mean	59.58	54.34	58.58	27.34
		Std	18.23	15.15	17.78	26.02
	\$6	Mean	55.81	55.27	55.65	27.69
		Std	15.80	11.22	14.58	24.12
	\$12	Mean	49.25	49.15	49.23	29.73
		Std	17.73	14.40	16.94	25.07

the sellers' profits in the Recall mode ($F(2,55) = 0.53$, ns), it did so in the No Recall mode ($F(2,53) = 11.02$, $p < 0.001$). The latter is due to a significant lower mean profit in the \$2 Search Cost level compared to the other two levels, but profits in the \$6 and \$12 cost levels do not differ significantly from each other. The Trial variable and the interactions of trial with all other independent variables were not significant.

In Figure 1 we have shown that asking prices were lower than predicted in the No Recall mode and were higher in the Recall mode. Given the higher than predicted search rates by buyers (reported in Table 4), was it advantageous for sellers to take the chance and ask for higher prices, knowing that lower prices do not deter a search as much as expected? Figure 3 depicts regression analysis plots with 95% confidence limits for the relationship between a seller's mean payoff (vertical axis) and his/her mean asking price (horizontal axis). Overall, there is a significant positive linear relationship between asking prices and payoffs ($R^2 = 0.11$, $p < 0.01$). However, the relationship is only significant in the No Recall mode (see Figure 3). Clearly, given the lower than predicted asking prices in the No Recall mode and the higher than predicted search rates, quoting a low price to deter buyers' searches was not the most effective strategy. There is no such relationship in the Recall mode where prices were already higher than expected.

Buyers

Buyers' payoffs were subjected to a repeated measure ANOVA with Search Cost, Recall, buyer's type and Trial as the independent variables. The following can be observed by looking at Table 6 and is confirmed by the statistical analysis reported below.

- (1) The effects of buyer's type and Trial as well as the interactions of these factors with all other independent variables were not significant.




Figure 3. Sellers' mean payoffs as a function of mean asking prices.

- (2) Both Search Cost and Recall main effects are significant ($F(2,102) = 17.29$, $p < 0.001$; $F(1, 102) = 3.84$, $p < 0.053$, for the Search Cost and Recall conditions, respectively). The interaction between the two is not significant ($F(2,102) = 0.39$, ns).

- (3) Payoffs are higher in the Recall compared with to the No Recall mode. Analyzing the effect of Recall within cost level, however, revealed that the above reported effect is significant at cost level \$2 ($F(1,40) = 5.68, p < 0.02$), marginally significant at cost level \$12 ($F(1,40) = 3.21, p < 0.08$), but is not significant at cost level \$6 ($F(1,28) = 0.62, ns$).
- (4) Analyzing the effect of the cost within the Recall condition revealed that in both modes the overall effect (reported in point 2 above) is due to a significant lower mean profit in the \$12 cost level compared to the other two cost levels, but profits in the \$2 and \$6 cost levels do not differ significantly from each other.

Discussion

Our experiment suggests that the equilibrium predictions capture only some of the qualitative features of the data. In line with recent developments in behavioral decision theory and game theory, which assume bounded rationality, we find that subjects follow simple rules of thumb in choosing strategies, as are reflected in the behavioral consistencies revealed in the study. These are summarized and explained below.

Summary of sellers' behaviors

While the model fails to predict exact asking prices, it accurately reflects the relationship between the Search Cost and asking prices within both the Recall and No Recall modes (hypotheses H1 and H2). Special attention should be given to the findings that asking prices do not increase monotonically with buyers' Search Costs in the Recall mode (H2). One's initial intuition is that an increase in the Search Cost reduces consumers' incentives to search for alternative sellers and hence would allow sellers to charge higher prices. However, our model predicts, and the experimental results show, that sellers are sensitive to the buyer's ability to recall prices. If recall is possible, sellers find it too costly to discourage the search (by asking low prices). Hence, instead of making a "low" offer to deter the search, they choose to ask for a "high" price, knowing that the buyer will search and hoping that this "high" price will be accepted later if the buyer's search turns out to be unfruitful.

H3 is clearly false. Intuitively (and in accordance with the model), one might expect that since buyers have more options in the Recall compared with the No Recall mode, they would be more powerful in the first condition and command lower prices. However, we find that in general, prices are higher in the Recall compared with the No Recall mode. This indicates that the external constraint imposed on the buyers' options in the No Recall mode in fact strengthens their positions, presumably because the moment a price is rejected in the No Recall mode, the seller is assured a zero payoff, and there is no way for the buyers to reverse their decisions even if they so desired. In the Recall mode, on the other hand, a price has always some chance of being accepted after the initial rejection (although the probability depends on the price), if the outside offer is

high enough. The No Recall constraint has strategic value. It changes the sellers' expectations about the buyers' future responses. It fits well with Schelling's illustration of the commitment tactics as an irreversible sacrifice of freedom of choice. The tactics rest on the paradox that "the power to constrain an adversary may depend on the power to bind oneself; that, in bargaining, weakness is often strength, freedom may be freedom to capitulate, and to burn bridges behind one may suffice to undo an opponent" (Schelling 1956).

Approximately 70% of the sellers exhibited pricing behavior that may be roughly characterized as a systematic search for the highest acceptable price. This search took place despite the fact that the identity of the buyers could have changed from trial to trial. The percentage reported above is higher than those reported by Ochs and Roth (1989) and Zwick et al. (1992) who found that 55.2 and 44% (respectively) of their subjects were involved in a similar adaptive behavior over trials. Our figure is much closer to the 80% mark reported by Weg, Zwick and Rapoport (1996), presumably because this study and ours have enriched the bargaining environment with outside options that facilitate adaptation (see the rich environment hypotheses tested by Weg et al. 1996).

Summary of buyers' behaviors

Most buyers (64.9%) used a consistent (or almost consistent) cut-off rule in accordance with H4. For these buyers, cut-off prices varied only slightly with the Recall condition and Search Costs. In the No Recall mode, cut-off prices were lower than predicted at all cost levels, and in the Recall mode, cut-off prices were higher than predicted at the \$2 and \$6 cost levels, and lower at the \$12 cost level.

Buyers who did not use a consistent cut-off rule were sensitive to the ordinal relationship between: (1) the asking price in the current trial (p_t), (2) the asking price in the previous trial (p_{t-1}), and (3) the outside offer in trial $t-1$ (x_{t-1}) if a search took place. In general, these buyers were more likely to search if the current asking price was higher than the previous price and were affected by the success of their previous search. After a successful search, the chance for searching again in the next trial was reduced, presumably because these buyers were susceptible to the gambler fallacy.²⁹

Why did cut-off prices not differ much between the Recall and No Recall modes? Computing the rational cut-off price in each condition differs in a subtle way. In the No Recall mode, a buyer should search if the asking price is higher than the expected value of the search minus the search cost. In the Recall mode, on the other hand, a buyer should search if the asking price is higher than the expected value of the search given that the asking price is already at hand, minus the search cost. The conditional expectation, in effect, truncates (from below) the range along which expectation should be computed. Computing conditional expectation is not an intuitive task. Consequently, we propose that buyers simplify the computation by segregating the search option into two distinct prospects (Thaler 1985). A buyer in the Recall mode might reason as follows: "I can get $100 - p$ for sure if I accept the asking price immediately. If I search, on the other hand, I will face a choice between a sure prospect of $100 - p - c$ and a prospect that *a priori* has an expected monetary value of $50 - c$." Clearly, for such a buyer to prefer to search, $50 -$

c must be better than $100 - p$, otherwise both prospects in the second choice are inferior to accepting immediately. Thus, we are back to the same reasoning as in the No Recall mode! Note that the deviation from rationality is due to segregating the two outcomes of the search rather than to integrating them.³⁰

The above hypothesis can explain the small difference in cut-off prices between the Recall conditions. However, it suggests that buyers should have accepted all prices below the $\$50 + c$ mark, whereas in reality many $\$50$ and even lower prices were rejected in favor of the search (see Figure 2 where almost all cut-off prices are below $\$50$).

Note that the game we play is an ultimatum game with a costly, uncertain outside option to the second mover.³¹ We know from ultimatum studies that first movers tend to offer amounts significantly higher than the minimum possible (the subgame perfect solution) and that a substantial percentage of games end in disagreement (that is, a substantial number of offers possessing positive value are turned down).³² On average, first movers can expect to get slightly more than 50% of the pie but hardly ever more than 65%.³³

Bolton and Zwick (1995) showed that in the context of the ultimatum game, strategic considerations can explain observed behavior once agent preferences over the relative division of the pie are properly taken into account. Specifically, they asserted that a player's preference for more money is modified to include a preference for disagreement over amounts he/she perceives as very small relative to his/her playing partner's share. Consequently, if the first mover proposes a significant unequal allocation, the second mover rejects this allocation to punish the first mover for unfair treatment. First movers, aware of the second movers' preferences, propose close to equal division.

It is intuitively clear that buyers in our game are strategically more powerful than second movers in the ultimatum game. Whereas rejection in the ultimatum game results in zero payoff to both players, rejection (and search) in our game can result in zero payoff to the seller but more than likely in a positive payoff to the buyer.³⁴ Consequently, the bargainers in our study considered the buyers' fair share to be more than what the buyers could have attained without the search option (the ultimatum game). Any adjustments due to search costs and recall conditions are made from this mark.³⁵

Buyers who did not use a consistent cut-off rule were sensitive in a predictable way to the fluctuation of asking prices from trial to trial as well as to the level of the outside offers when they were available.

Payoffs

Buyers fared much better than sellers in the current environment although in theory they should have done so only in the Recall mode. Apparently, the sellers' bargaining power derived from their structural position of moving first and the cost (to the buyers) to search was not enough to offset the buyers' power derived from their ability to take their business elsewhere. The prospect of being left out was the dominant factor in determining the power balance between buyers and sellers. Consequently, buyers profited about twice as much as sellers in almost all conditions.

It is interesting to relate asking prices to buyers' profits. Whereas asking prices were lower in the No Recall compared with the Recall mode, buyers' mean profits were higher in the Recall mode. Thus, the effect of the Recall condition on buyers' mean profits conforms to the theoretical model, yet the underlying reason for the effect is different. Whereas, theoretically, higher buyers' profits in the Recall mode should have been related to lower prices compared with the No Recall mode, in practice, the higher profits in the Recall mode were due to the much lower search rate and the fallback option of accepting the asking price after an unsuccessful search. Another interesting observation is that although a high search cost (\$12) in both Recall conditions significantly lowered buyers' profits, it did not correspond to a similar increase in the sellers' profits. That is, buyers' lower profits are due to unnecessarily high search rates rather than to higher prices.

Conclusion

Naturally, the model we have tested gives away much of the complexity of the bargaining between buyers and sellers by assuming that only the buyers can search and that the "bargaining" consists of the seller making a take-it-or-leave-it price offer to any buyer who has found him/her. The implicit assumption, therefore, is that sellers have the ability to make commitments. In some applications this is appropriate, such as in a store employing a sales clerk who is instructed to sell the good at a fixed price. The posted price offer then becomes credible because this agent is not authorized to negotiate price reductions. However, in many other cases, this is not a good assumption. Nevertheless, if the predictions of the theory are not met in the simplified experimental market, this raises a presumption that this theory will not be adequate for the study of a more complex market; but this is, of course, an open research question. We believe, however, that both laboratory play and real world bargaining are governed by the same general principles, and, hence, the behavioral regularities we have reported should facilitate prediction of outcomes in buyer-seller negotiation in simple bargaining contexts, where buyers have uncertain costly outside options, and where bargainers do not expect repeated interactions. As such, they can be appreciated by managers and consumers alike.

The theoretical model tested here assumes that bargainers are risk neutral. Since buyers rejected prices that were better than the expected monetary gain from a search (in the No Recall mode, for example), it can be argued that buyers are not risk neutral but rather exhibit risk seeking behavior. Sellers, on the other hand, asked for prices that were lower than expected, demonstrating risk aversion. Since subjects from the same population were assigned to play the roles of buyers and sellers, it is hard to explain why subjects' attitudes toward risk would vary with their role in the game. We suspect that the risk posture cannot account for the behavioral regularities found in our study. They must be due rather to perceived commitment power, task simplifications, fairness considerations (as proposed in the section above on buyers' behaviors) and simple adaptation processes.

On the more practical side, whether or not the asking price will be available after a search is usually under the sellers' control. From a planning point of view, when search cost is very low, sellers are much better off inviting the buyers to search and guaranteeing the

same asking price if the search does not yield a better price. Under all other conditions, asking price levels and search rates offset each other to produce remarkably stable seller's profits (see Rapoport, Erev and Zwick (1995) for a similar finding). Buyers, on the other hand, are better off in the Recall-mode environment at all cost levels. We found that a buyer who asks the seller to quote a price but adds, "if I am not satisfied with this price I will immediately leave the store and will not come back even if I do not find a better price elsewhere," will succeed in lowering the asking price. The tendency, however, to do too much searching under this condition more than wipes out the potential gains from lower prices.

Our findings suggest several possible directions for future research. First, for enhanced generalizability, the experimental work can be extended to include many of the variables that are assumed away in the present study. For example, bargaining can be extended to include more than one period, thereby weakening the ultimatum nature of the environment (Lee 1994). Sellers can be allowed to search, thereby rendering the search option a symmetric status (Wolinsky 1987), or exogenizing the seller's choice to allow or prevent recall. Second, the compression of prices at around the \$50 mark is due, in our opinion, to the buyers' and sellers' tendencies to compare and resist significant deviations from equal profits from agreements.³⁶ The comparison is only feasible in the full information environment provided by the present study. Relaxing the information requirement is simply a natural progression to follow. Third, in our study, different search costs reflect search efficiencies. However, the model does not capture the full strategic role of the search activity because it does not treat the search intensity as a decision variable that can be controlled by buyers and sellers. Wolinsky (1987) has shown that search intensity can play an important role in determining the negotiated outcomes. Finally, our study treats each trial as an independent encounter. However, in reality, reputation (e.g., for not negotiating) is an important consideration. To what extent the results are affected by an on-going relationship is an open question.

Notes

1. For example, the set prices for the Saturn and the 1993 subcompact Escort.
2. Note that "Recall" in this context does not refer to the buyer's ability to recall the offer from his/her memory but to the availability of the offer in the future once it is rejected at the present time.
3. A specific example is a bakery that can not roll over baked goods from one day to the next and the described encounter has occurred just prior to closing. In general, the No Recall environment exemplifies situations where negotiation re-open if the buyer returns from an unsuccessful search, but presumably the buyers bargaining position is weakened.
4. Certain here refers to the magnitude of the outside option and not to its availability.
5. See, for example, Fudenberg and Tirole (1983), Sobel and Takahashi (1983), Chatterjee and Samuelson (1987, 1988), and Grossman and Perry (1986). A comprehensive review of bargaining models with private information can be found in Kennan and Wilson (1993).
6. See, for example, the bargaining model of Rubinstein (1982), the bargaining with outside option models of Shaked and Sutton (1984), Binmore, Shaked and Sutton (1989), Shaked (1987), and the bargaining and search models of Wolinsky (1987), Chikte and Deshmukh (1987), and Muthoo (1995). Exceptions, however, are found in Rubinstein (1982), Chatterjee et al. (1993) and Lee (1994) for different reasons. In Rubinstein and Chatterjee et al., delay arises in only one of the many equilibria, whereas in Lee, delay arises in the unique equilibrium of the game.

7. This can be generalized, but for simplicity it is assumed that the good has already been produced and the seller cannot bargain with any other buyer.
8. To guarantee that the expected payoff from search is always positive, we assume that c is not very high compared to the potential value of the good. This implies in the uniform distribution case that $c < 0.5$.
9. See Footnote 2 above for our particular definition of recall.
10. To be more precise, in order for this statement to be true, we need to show that $(1 - p) - [-c + 1 - E(x \wedge p)]$ is monotonically decreasing in p , which is true. The proof is straightforward. See Lee (1994).
11. This is true if outside offers follow a uniform distribution (except when $c = 0$, in which case, the prices equal). For other distributions, the equilibrium offer in the Recall case can be higher than that in the No Recall case. For example, as a referee pointed out, when the outside offer is equally likely to be 10 or 90, the equilibrium price in the No Recall condition is $45 + c$, whereas in the No Recall case it is 90 (for $c < 35/2$).
12. A few subjects were from engineering, arts and the MBA program.
13. \$100 Hong Kong dollars. The exchange rate between Hong Kong and US dollars is approximately 7.8 to 1.
14. The amount of the cash prize was very attractive to students considering that the hourly wage for an on campus job was about \$35 during the first few sessions. The hourly wage was later raised to \$50.
15. The minimum significant difference at 0.05 is 2.2.
16. The minimum significant difference at 0.05 is 2.0.
17. Their behavior can not be characterized as non-adaptive (as defined next in the text) since almost all prices were accepted on each trial.
18. The existence of such cut-off prices will be discussed in the section on buyers' behaviors.
19. Of whom two always demanded the same price.
20. Of whom four always demanded the same price.
21. For reasons of brevity, the exact correspondence to the group data is not presented in the text. The reader can extract the information from Table 3.
22. In the No Recall mode the number of buyers so identified within Search Cost level are 11, 11, and 10 for Search Cost \$2, \$6, and \$12, respectively. In the Recall mode, the corresponding number of buyers are 17, 12, and 15.
23. Accepting all prices is equivalent to a low cut-off price and always searching is equivalent to a high cutoff price.
24. The box plot interpretation is as in Figure 1. The cut-off price for buyers who exhibited a perfect cut-off rule was computed as the average price between the highest acceptable price and the lowest price that was followed by a search. The cut-off price for buyers who exhibited an almost perfect cut-off rule was computed as above by ignoring the one reply needed to transform the pattern from an almost perfect to a perfect cut-off rule. If the identity of the ignored reply is not unique (as in the pattern "AAAASASSSSSSSS", where a buyer's replies (Accept or Search) to asking prices that are ranked from lowest to highest are presented), then we computed the cut-off price as above ignoring both replies. In the above example, the cut-off price was computed as the average of the 4th and the 7th prices. If all prices were accepted, then the cut-off price was computed as the highest acceptable price plus \$0.5. Similarly, if all prices were followed by a search, the cut-off price was computed as the lowest price that was followed by a search minus \$0.5.
25. Out of the 17 (12) buyers in the \$2 (\$6) cost condition, 5 (6) searched after any price. These subjects' cut-off prices were defined to equal \$0.5 less than the lowest price they encountered. This, of course, is not a definite estimate because it depends on the prices these buyers have actually encountered. We do not know what would have been their reaction to even lower prices. Consequently, the higher than predicted cut-off prices should only be treated as a proposition that needs further verification. However, the other buyers in these conditions whose cut-off prices were estimated without ambiguity exhibited cut-off prices higher than predicted.
26. A complete description should also take $x_{t-1} + c$ into account. In what follows, we consider a successful search (ex-ante) to be one that yields a price lower than the asking one. However, a buyer who considers the success of his/her previous search would probably compare p_t with $x_{t-1} + c$. Nevertheless, in only 5.5% of the trials with a search, adding the search cost to the outside offer reversed the ordinal relationship between p_t and x_{t-1} . For simplicity, we present the data ignoring the "location" of $x_{t-1} + c$. We have performed the same analysis with $x_{t-1} + c$ (in addition to the other 3 variables) and found virtually the same pattern.

27. The main difference between buyers who used a cut-off rule and those who did not can be demonstrated here. Buyers do not exhibit a stable cut-off rule exactly because they are affected by the relationship between p_{t-1} and x_{t-1} , whereas buyers who do just compare p_t to their own internal cut-off price, ignoring, for the most part, x_{t-1} . Computing search rates for the above two cases for buyers who used the cut-off rule verify this assertion. Search rates in trial t are about the same if the previous trial's search yielded a worse price than the seller's asking price ($x_{t-1} > p_{t-1}$) compared to when it was better ($x_{t-1} < p_{t-1}$). The comparable search rates are 0.56 and 0.53 in the No Recall, and 0.82 and 0.79 in the Recall modes.
28. Sellers were also classified into three groups based on how many times they encountered each type of a buyer. One segment contains all sellers who have been matched in 10 or more trials with buyers who have adopted a cut-off rule. The second segment contains all those sellers who have been matched in 10 or more trials with buyers who have not adopted a cut-off rule. The third segment contains all the rest of the sellers. The above ANOVA was repeated including seller's type as an independent variable. The effects of the seller's type as well as the interactions with all other independent variables were not significant. Thus we do not include this factor in the results reported next.
29. Virtually the same pattern exists if a "successful search" is defined as $x_{t-1} + c < p_{t-1}$. See footnote 25.
30. The difference is between computing $E[\min(x, p)]$ and $\min[p, E(x)]$.
31. In an ultimatum game, two players, a first mover and a second mover, must come to a mutual agreement on how to share a pie of k dollars. The first mover proposes a division to the second mover. If the second mover "accepts", the money is divided accordingly. If the second mover "rejects", both players receive nothing. For a review of experimental results of ultimatum games, see Roth (1995).
32. This characterization of the data has proved robust to both cultural (Roth, Prasnikar, Okuno-Fujiwara and Zamir (1991), Carter and Irons (1991)) and framing perturbations (Hoffman, McCabe, Shachat and Smith (1994)), as well as other minor variations (Forsythe, Horowitz, Savin and Sefton (1994); Weg and Smith (1993)).
33. Interestingly, Hoffman et al. (1994) found that naming the first mover in an ultimatum game a "seller" generates higher demands compared to a condition in which subjects were named Player 1 and Player 2. When named "seller", Player 1 was more likely to take advantage of his/her superior strategic position.
34. Given the relatively low search costs in our study, there is only a slight chance that a buyer will end up with a negative payoff if, in the No Recall mode, the outside offer is sufficiently high so that the buyer's surplus is less than the search cost. In the Recall mode, Search Cost has to be higher than the buyer's surplus derived from accepting whichever is lower, the outside offer or the asking price.
35. Buyers in our study can also punish what they may consider an unfair price demand by the sellers in the Recall condition by accepting the outside offer even though it is higher than the seller's price. This is similar to a disadvantageous counter offer in an alternating sequential bargaining game. In our study, a higher outside price was accepted in 17 out of 290 trials (5.9%). The same buyer made 6 out of the 17 acceptances. Only once a higher asking price was accepted in 349 trials.
36. Note that the unequal profits reported in Table 6 are due to disagreements or delayed agreements.

References

- Bazerman, M.H. and M.A. Neale. (1992). *Negotiating Rationally*. New York: The Free Press.
- Binmore, K., A. Shaked, and J. Sutton. (1989). "An Outside Option Experiment," *Quarterly Journal of Economics* 104, 753–770.
- Bolton, G.E., and R. Zwick. (1995). "Anonymity Versus Punishment in Ultimatum Bargaining," *Games and Economic Behavior* 10, 5–121.
- Brucks, M., and P.H. Schurr (1990). "The effect of Bargainable Attributes and Attribute Range Knowledge on Consumer Choice Processes," *Journal of Consumer Research* 16, 409–419.
- Camerer, C., G. Loewenstein, and M. Weber (1989). "The Curse of Knowledge in Economic Settings : An Experimental Analysis," *Journal of Political Economy*, 97, 1232–1254.

- Carter, J.R., and M.D. Irons. (1991). "Are Economists Different, and If So, Why?" *Journal of Economic Perspectives* 5, 171–177.
- Chattadee, K., and L. Samuelson. (1987) "Bargaining with Two-sided Incomplete Information: An Infinite Horizon Model with Alternating Offers," *Review of Economic Studies* 54, 175–192.
- Chattadee, K., and L. Samuelson (1988). "Bargaining under Two-Sided Incomplete Information: The Unrestricted Offers Case," *Operations Research* 6, 605–618.
- Chattadee, K., B. Dutta, D. Ray, and K. Sengupta (1993). A Non-Cooperative Theory of Coalitional Bargaining," *Review of Economic Studies* 60(2), 463–477.
- Chikte, S.D., and S.D. Deshmukh. (1987). "The Role of External Search in Bilateral Bargaining," *Operations Research* 35, 198–205.
- Forsythe, R., J. Kennan, and B. Sopher (1991). "An Experimental Analysis of Strikes in Bargaining Games With One-Sided Private Information," *American Economic Review* 81, 253–278.
- Forsythe, R., J. Horowitz, N.E. Savin, and M. Sefton. (1994), "Fairness in Simple Bargaining Experiments," *Games and Economic Behavior* 6, 347–369.
- Fudenberg, D., and J. Tirole. (1983). "Sequential Bargaining with Incomplete Information," *Review of Economic Studies* 50, 221–247.
- Fudenberg, D., D.K. Levine, and J. Tirole. (1985). "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in Roth (ed.), *Game Theoretic Models of Bargaining*. Cambridge: Cambridge University Press, pp. 73–98.
- Grossman, S.J., and M. Perry. (1986). "Sequential Bargaining under Asymmetric Information," *Journal of Economic Theory* 39, 120–154.
- Hoffman, E., K. McCabe, K. Shachat, and V. Smith. (1994). "Preferences, Property Rights and Anonymity in Bargaining Game," *Games and Economic Behavior* 7, 346–380.
- Kennan, J., and R. Wilson. (1993). "Bargaining with Private Information," *Journal of Economic Literature* 31 45–104.
- Lee, C.C. (1994). "Bargaining and Search with Recall: a Two-Period Model with Complete Information," *Operations Research* 42, 1100–1109.
- McMillan, J., and M. Rothschild. (1994). "Search," in R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory*, Vol. 2, Amsterdam: Elsevier Science, pp. 905–927.
- Muthoo, A. (1995). "On the Strategic Role of Outside Options in Bilateral Bargaining," *Operations Research* 43, 292–297.
- Nash, J.F. (1950). "The Bargaining Problem," *Econometrica* 18, 155–162.
- Ochs, J., and A.E. Roth. (1989). "An Experimental Study of Sequential Bargaining," *American Economic Review* 79, 355–384.
- Rapoport, A., I. Erev, and R. Zwick. (1995). "Buyer-Seller Negotiation With One-Sided Incomplete Information," *Management Science* 41(3), 377–394.
- Roth, A.E. (1995). "Bargaining Experiment," in J. Kagel and A.E. Roth (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press, pp. 253–348.
- Roth, A.E., V. Prasnikar, M. Okuno-Fujiwara, and S. Zamick. (1991). "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review* 81, 1068–1095.
- Rubinstein, A. (1982). "Perfect Equilibrium in a Bargaining Model," *Econometrica* 50, 97–109.
- Schelling, Thomas C. (1956). "An Essay on Bargaining," *American Economic Review* 46, 281–306.
- Shaked, A. (1987) "Opting Out: Bazaars versus 'High Tech' Markets," Discussion Paper 87/159 (Theoretical Economics), Suntory Toyota International Centre for Economics and Related Disciplines, London School of Economics.
- Shaked, A., and J. Sutton. (1984) "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model," *Econometrica* 52, 1351–1364.
- Sobel, J., and I. Takahashi. (1983). "A Multi-stage Model of Bargaining" *Review of Economic Studies* 50, 411–426.
- Thaler, R.H. (1985). "Mental Accounting and Consumer Choice," *Marketing Science* 4, 199–214.
- Weg, E., R. Zwick, and A. Rapoport. (1996). "Bargaining in uncertain environments: a systematic distortion of perfect equilibrium demands," *Games and Economic Behavior* 14, 260–286.

- Wernerfelt, B. (1994). "Selling Formats for Search Goods," *Marketing Science* 13, 298–309.
- Wolinsky, A. (1987). "Matching, Search, and Bargaining," *Journal of Economic Theory* 42, 311–333.
- Zwick, R., and E. Weg. (1996). "An Experimental Study of Buyer-Seller Negotiation: Self-Interest Versus Other-Regarding Behavior," Working Paper 96.071, The Hong Kong University of Science and Technology.
- Zwick, R., A. Rapoport, and J.C. Howard. (1992). "Two-Person Sequential Bargaining Behavior With Exogenous Breakdown," *Theory and Decision* 32, 241–268.

