Protocols for automated negotiations with buyer anonymity and seller reputations

Lorrie Faith Cranor [a] and Paul Resnick [b]

[a] *AT&T Labs-Research*
E-mail: lorrie@research.att.com
[b] *School of Information, The University of Michigan, Ann Arbor, MI 48109-1092, USA*
E-mail: presnick@umich.edu

Received 27 January 1999

In many Internet commerce applications buyers can easily achieve anonymity, limiting what a seller can learn about any buyer individually. However, because sellers need to keep a fixed web address, buyers can probe them repeatedly or pool their information about sellers with the information obtained by other buyers; hence, sellers' strategies become public knowledge. Under assumptions of buyer anonymity, publicly-known seller strategies, and no negotiation transaction costs for buyers, we find that take-it-or-leave-it offers will yield at least as much seller profit as any attempt at price discrimination could yield. As we relax those assumptions, however, we find that sellers, and in some cases buyers as well, may benefit from a more general bargaining protocol.

1. Introduction

Because it is so easy to gather personal information during electronic interactions, there has been increasing concern about the privacy implications of the World Wide Web and of electronic commerce conducted using the Web. For example, when we provide demographic information to a Web site, will it be sold to direct marketers? Or, perhaps more disturbingly, if our browsing behavior can be captured and analyzed, will health insurers raise premiums or refuse coverage for people who retrieve information about certain illnesses that are costly to treat? One principle of fair information practices is that consumers should receive notice about what information is being collected and how it will be used, and have a choice about whether to interact on those terms [7]. When browsing, however, people do not want constant interruptions that require them to read the fine print of privacy notices and click to indicate that they accept those terms and conditions.

The Platform for Privacy Preferences Project (P3P), under development by the World Wide Web Consortium (W3C, 1998), is designed to allow negotiations over privacy terms and conditions to take place in the background without constant user intervention. Users should be able to enter their privacy preferences once and have their software act accordingly. For example, some people may be willing to share their address and phone number with Web sites that promise to use it only for order fulfillment. At each Web site these people visit, their software would negotiate in the background with software running on the site, divulging the information only if the Web site agrees to this limitation on its use. Outside the scope of P3P, two automated agents – one representing the user and the other the Web site – might also negotiate prices or terms and conditions for a good or the use of a piece of intellectual property. For example, negotiations may lead to agreements in which a Web site allows a piece of music to be downloaded by a browser that promises to play the music only once and then discard it [11].

Throughout this paper we shall refer to Web sites as *sellers* and Web browsers (and the people who use them) as *buyers*. In our examples, the seller will generally offer some content or service in exchange for some combination of money and the right to use personal information for specified purposes. While it certainly may be viewed from the opposite perspective as well, we view the content or service as the good being sold and the money and personal information as the currency of the transaction.

What bargaining protocol should the agents use in their negotiations? A bargaining protocol determines when players can make offers and the structure of offers. In its most general form, negotiation involves one or more rounds of offers that are either accepted or rejected. The particular rules governing the sequence of offers, however, can vary. The parties may alternate offers and counter-offers, or one party may make all the offers, or anything in between. It may be required that an offer, once made, not later be withdrawn, or offers may be valid only if accepted within a specified time period or before the end of the round in which the offer was made.

The protocol for a particular negotiation or limitations on the validity of an offer may not be enforceable without external mechanisms such as legal statutes. Such rules may none-the-less be credible when announced by a negotiator who has established a reputation for following a particular set of rules. For example, customers of most car dealers are unlikely to believe an announcement that the first offer is the best and final offer, because car dealers are known to haggle with their customers. On the other hand, some car dealers have established a reputation for offering a fixed price.[1] At such dealerships customers understand that the dealer's first offer is his final offer; the dealer will not make a second offer for fear of tarnishing his reputation. Negotiation rules more complicated than those in the above example may prove even more difficult to enforce in the absence of external mechanisms.

In an automated negotiation, the software implementation of the automated protocol itself may enforce some restrictions on player actions. For example, the protocol

[1] The Saturn corporation has invested in advertising in order to develop a reputation that its dealers will not haggle over price.

can bar one or both players from making counter-offers in certain situations. We assume that any communication not specified within the bounds of the agreed-upon protocol will be completely ignored by the agents. Such restrictions would be implausible in live negotiations.[2]

While automated negotiations allow for credible restrictions that may shorten or simplify negotiation, they also allow for negotiations in which complicated strategies are used. In some negotiations, the optimal strategy for one or both of the players involves contingencies or randomization that would be difficult for most people to discover and execute. However, computerized agents might be programmed to execute the optimal strategy for every negotiation, regardless of the complexity of that strategy.

In addition, computerized agents can discover and execute strategies and communicate offers much more quickly than humans can. Thus the delay costs from multiple rounds of negotiation will normally be negligible. However, when delay costs place unequal burdens on the parties to a negotiation, the optimal strategy for one party may involve creating a deliberate delay that creates significant costs.

In this paper we consider questions of two kinds. First, given a negotiation protocol (the rules that govern offer sequences), what strategies are optimal for players to adopt? Second, given an analysis of players' optimal strategies under various rules, how should a social planner design a negotiation protocol?

The Web negotiation situation has three peculiar properties that we will explore:

- First, the browser (buyer) may have some degree of anonymity, while the Web site does not. The browser may be able to conceal its identity or use a pseudonym, while the Web site's address is public. Thus, the Web site is somewhat analogous to an identified store, which buyers can enter or leave at will. One implication of this difference in anonymity is that buyers may be able to pool information about a seller's negotiating strategy. Interestingly, a buyer who chooses to reveal its identity can also make its past negotiating strategies with other sellers visible to the current seller; a buyer who develops a reputation as a hard bargainer may benefit from revealing its identity.
- Second, only the browser can initiate an interaction. This is a basic feature of the HTTP communication protocol that governs interactions between browsers and Web sites. With the advent of "push" technology, this asymmetry between browser and web site is eroding, but many transactions of the type we envision will be based on the more conventional "pull" technology. One implication is that if a buyer ends a negotiation, the seller may have no way of contacting the buyer to make an additional offer. If the seller cuts off a negotiation, however, the buyer can still make additional offers, and it may be irrational for the seller to refuse to consider such offers.

[2] For example, in a live negotiation a player who makes an offer and says he will refuse even to listen to any counter-offer has a hard time making that threat believable.

- Third, a buyer may be able to restart a negotiation under a new pseudonym without the Web site being able to detect this. This gives a buyer another potential informational asymmetry, even if the buyers do not pool their information about a seller's negotiating strategy. A buyer who returns under a different pseudonym still retains its knowledge of the seller's behavior in the prior negotiation.

The ability of automated agents to recognize and execute complicated strategies will be limited if their underlying programs are unsophisticated. Indeed, programming an agent to behave optimally in a wide range of negotiation situations is a non-trivial problem. This problem may be simplified for negotiation protocols that bound the number of rounds of offers, especially if only one or two rounds are permitted. The optimal strategies will be easier for people to recognize – and, in turn, program automated agents to recognize and execute – and, hence, there will be less danger of players adopting irrational strategies. That, in turn, makes it safer for players to assume rational opponents when choosing their own strategies.

Limited protocols such as a take-it-or-leave-it offer in the first round are especially easy to implement. Ease of implementation may, in turn, spur faster adoption. In addition, in some cases, restricting the number of rounds in the protocol may focus players on more socially efficient outcomes. Thus, we pay special attention to those situations where all parties benefit at least as much from a limited protocol as they do from a protocol with unbounded rounds. For protocols like P3P that are motivated in part by a desire to empower consumers, situations in which the buyer benefits at least as much from the limited protocol may be acceptable, even if the seller might benefit more from the unbounded protocol under some circumstances.

Besides ease of implementation, limited protocols may sometimes be desirable because they restrict the amount of time it may take to complete the negotiation. Although computerized agents allow each round of negotiation to be executed quickly, negotiations can take a significant amount of time if they persist for many rounds (and indeed they may sometimes proceed indefinitely unless one of the agents is programmed to end the negotiation after a fixed number of rounds or after a certain amount of time has passed without reaching an agreement). Time considerations are most significant when the negotiation occurs while users are waiting.

In the next section we describe three assumptions – buyer anonymity, publicly-known seller strategies, and no buyer transaction costs – and show that a restricted protocol will yield the same equilibrium outcomes as a more complicated protocol when these assumptions hold. Arguably, our three assumptions do not always hold in the bargaining situations of interest. If we could prove the same result when these assumptions were relaxed, then designers of P3P and similar systems would be well advised to adopt restricted protocols. However, as we relax these assumptions, we find that sellers, and in some cases buyers as well, may benefit from a more general bargaining protocol. In the sections that follow we consider these results and what they mean for protocol designers.

2. Modeling

A seller faces an infinite sequence of buyers $\{b_1, b_2, \ldots\}$. Each buyer b_i has a type defined by a reservation value v_i, and sometimes a delay or transaction cost k_i per round of negotiation. The buyer types are drawn from independent, identically distributed random variables having distribution B. The seller adopts a strategy $S = \{S_i\}$, where S_i defines the negotiation strategy employed with buyer b_i. S_i may consist of an initial offer, responses to potential counter-offers, planned further counter-offers, etc. Moreover, any of these plans may involve randomization. S_i may be contingent on the outcome of previous negotiation sessions. For example, the seller may want to experiment with various strategies in order to obtain information about the distribution of buyer types or, when individual buyers' histories are available, about a particular buyer's strategy. We will mostly be concerned, however, with the steady state strategy of the seller, after any initial learning period.

For simplicity, we initially analyze the protocols as if the negotiation concerned prices rather than personal data to be revealed or other transaction terms. One limitation of this modeling choice is that prices are normally continuous, while only a few discrete values may be possible for other terms and conditions. This difference between connected and unconnected agreement sets can be important (see, for example, Osborne and Rubenstein [8, section 3.10]), but none of our results depend on an assumption of connected agreement sets.

A more serious limitation of modeling trade in terms of money rather than data is that personal data may be more valuable to some buyers than others. To rectify this, we sometimes include in a buyer's type a positive linear utility function u_i defined over the unit of trade. Note that we introduce the utility function not for its usual purposes of modeling risk aversion or wealth effects but merely because our unit of trade may be data rather than money. Thus, positive linear utility functions are sufficient for our goal of allowing buyer types to have different utilities for personal data. If some buyer types are truly indifferent about release of their personal data, they will be willing to reveal data even if an agreement can be reached with less data revelation. This suggests that the P3P protocol should always include the option for users to reveal additional personal data,[3] but we consider that as an action outside of the negotiation process. For the remainder of this paper, we assume that all buyer types have some positive valuation for personal data, even if it is very small.

3. Cases

Our analysis begins with a simplified case, in which buyers have full anonymity, full knowledge of the seller's strategy, and no transaction costs. Not surprisingly, the

[3] This might be implemented in a protocol by allowing agreements to be reached in which transferring certain data elements is optional. Thus, a buyer's agent might reach the optimal agreement involving the minimum amount of data revelation acceptable to the Web site, but an indifferent buyer can send all or some of the optional data as well.

seller has little bargaining power in this situation, and can do no better than to make take-it-or-leave-it offers in the first round. Under these assumptions, a protocol that permits only one round of negotiation offers rational players the same outcome as a more general protocol, and eliminates the need to consider irrational strategies on the part of opponents. The remaining sub-sections relax each of the assumptions in turn. In some cases, the main result still holds, while in others the seller and/or some of the buyers may prefer a more complicated protocol.

3.1. Base case

We begin our discussion by exploring a simplified case based on three assumptions. We will show that these are reasonable assumptions in some situations, and that when these assumptions hold the optimal strategy for the seller is a take-it-or-leave-it offer.

Assumption 1 (The buyer is anonymous). Individuals who wish to browse the web anonymously can take advantage of a variety of tools and techniques that offer varying degrees of anonymity or pseudonymity. Pseudonymous relationships with Web sites can be established by offering a pseudonym rather than one's real name when prompted to register at a site [3]. By offering a different pseudonym at each site, individuals can thwart sites' attempts to compile profiles based on information collected by multiple sites. However, every time the individual returns to the same site, that site will recognize them by their pseudonym. Cookies and other persistent state technologies [5,6] can also allow sites to recognize repeat visitors, even those that have not registered with the site. To prevent sites from recognizing them as repeat visitors, individuals can install software on their computers that prevents the use of cookies and masks all identifying information except IP address. For even greater anonymity they can use an anonymizing proxy server [1,9] that hides all browser information. Thus, individuals who wish to browse anonymously have the opportunity to do so.

The ability of buyers to be anonymous has two important implications for online negotiations. *First, it prevents the seller from taking advantage of any historic information it may have on the buyer.* The seller is unable to determine that a particular buyer is one it does business with frequently, one that has a reputation for accepting high offers, or one known for driving hard bargains. *Second, it prevents the seller from distinguishing buyers who are initiating new negotiations from those restarting a recent negotiation.* Thus, the seller has no way of knowing if the present buyer is the one who walked away five minutes ago without reaching an agreement, or a new buyer with whom there has been no negotiation. As a result, buyers can come and go at will, returning repeatedly to probe the seller's strategy, without weakening their position in future negotiations.

Assumption 2 (The buyer has complete knowledge of the seller's strategy). Another implication of buyer anonymity is that buyers may have the opportunity to probe the seller until they can deduce the seller's strategy. Even if buyers are not completely anonymous, buyers may pool their information to determine a seller's strategy. Once a seller's strategy is determined, it may be posted to a server where it may be publicly accessible (or accessible to those participating in some sort of buyers' consortium) and available to future buyers. We can imagine systems that would allow buyers to automate the process of checking servers for sellers' strategies before engaging in a negotiation. Thus we assume that buyers have complete knowledge of the seller's strategy.

Assumption 3 (There are no transaction costs for buyer or seller). Because online negotiations may be completely automated and occur quickly with minimal consumption of resources (and any consumption that occurs is likely to involve resources that would otherwise be idle and for which there are no per-minute, or per-use fees), they may be said to be free of transaction costs.

The seller's goal is to maximize expected revenue per arriving buyer b_i. Intuitively, a single fixed price is the best a seller can do because, under any other strategy, the buyer may restart the negotiation as many times as necessary to elicit the seller's best offer. Because the seller's strategy is common knowledge, the buyer is aware of what the seller's lowest price is, and what strategy she must use to obtain it. Thus, for any seller strategy, the buyer can select the appropriate strategy so as to reach an agreement at the seller's lowest price. Hence, the seller can do no better than to offer that price right away as a take-it-or-leave-it offer.

The seller cannot take advantage of any historic knowledge of the buyer, and it cannot make a credible threat of cutting off negotiation with a buyer if an offer is not accepted after a certain amount of time. In addition, the seller cannot take advantage of any desire on the part of the buyer to reach a quick agreement in order to avoid excessive transaction costs.

The seller may be able to price discriminate among buyer types by planning different strategies for different negotiation sessions. For example, suppose the seller adopts a strategy S_N that demands a price 2; S_{N+1} demands 3 with probability 0.8 and demands 1 with probability 0.2; finally, S_{N+2} and beyond all demand a price of 3. A buyer who has a reservation value of 2.5 will pay the price 2 immediately. A buyer who has a reservation value of 1.5 will refuse and restart the negotiation, hoping to get a chance to purchase at price 1 in session $N + 1$.

While such price discrimination strategies may work in the short run, they are not effective in steady state. In order for a buyer to accept a price right away when a lower price might be offered in a future round, not only must the lower price be unlikely, but it must be unlikely that even the current price will be available again. Thus, the seller's usual "best offer" must be continually rising in order for price discrimination to work. In steady state, the amount of profit that can be gained from this is vanishingly small.

Trading data rather than money. The result holds even if the unit of trade is data rather than money, as modeled by the introduction of buyer utility functions on the unit of trade. Even if buyers have a small valuation for their personal data they will always hold out for a better price (revelation of less data) if they can do so without cost to themselves.

3.1.1. Case 1a: No anonymous return

We now relax the assumption that buyers are anonymous and assume instead that buyers cannot change their identifiers in the short run, though they may be able to do so in the long run. Thus, a buyer cannot end a negotiation and immediately return under a new pseudonym, although the buyer may return under a new pseudonym after some amount of time has passed. While sellers may carry out threats to discontinue negotiation with a buyer, sellers are not able to establish long-term profiles of individual buyers, nor are they able to share profiles among sellers. In short, a seller can make inferences about the buyer's type or strategy only from the progress of the current negotiation session.

This situation may occur when sellers identify buyers by their IP addresses. Buyers who dial into an Internet Service Provider (ISP) to access the Internet may receive a different IP address every time they dial in. The IP address can be thought of as a pseudonym for the buyer that exists throughout a particular online session. However, when the buyer disconnects from the ISP and dials in again, a new session is established with a new pseudonym. This situation may also occur when a seller leaves a cookie on a buyer's machine that the buyer periodically deletes.

As with the base case, the optimal seller strategy in this situation is a fixed-price strategy. Riley and Zeckhauser [10] analyze this bargaining problem. In their model, sellers can adopt and commit to strategies that, with some probability, cut off negotiation with the buyer. In order to make such a threat practicable, the seller must be able to identify the buyer should the buyer try to return immediately. This is possible under our assumption of "no anonymous return". Thus, buyers who do not accept a seller's offer risk losing the ability to reach any agreement. The seller can parlay this threat into successful price discrimination. Buyers for whom the value of reaching an agreement is low will be more willing to risk not reaching an agreement than those for whom the value of reaching an agreement is high.

Even though price discrimination is possible in this model, Riley and Zeckhauser prove that it is never more profitable for the seller than all fixed-price strategies. The lost value from carrying out the threat of cutting off some buyers more than offsets the added value from charging higher prices to some of the buyers. While it is not easy to summarize their proof, we offer a proof for a more restricted case that provides some insights about why the result holds.

Proposition 1. Assume that buyers have full knowledge of the seller's strategy and no transaction costs, and sellers cannot take advantage of any historic information about a buyer obtained during previous online sessions, but can identify buyers who end a negotiation and immediately return. Suppose that a seller with reservation price s faces, with probability q, a buyer with reservation value v_1 and, with probability $1-q$, a buyer with reservation value v_2. No seller strategy is more profitable than every take-it-or-leave-it strategy.

Proof. Without loss of generality, assume $v_2 < v_1$. Here, we need concern ourselves only with a strategy for a single negotiating session since if a buyer returns, the seller can recognize the buyer and continue the previous negotiating session. Suppose there is an optimal strategy S that is more profitable than any take-it-or-leave-it strategy. If the first offer of S is randomized, pick one of the possible offers, call it x_1, that yields maximum expected profit to the seller. Without loss of generality, assume that at least one player type accepts the initial offer (else consider the equivalent strategy that begins with the second round after the offer of x_1 is rejected). If both player types accept x_1, the strategy is equivalent to one where the initial offer is made on a take-it-or-leave-it basis. Hence, x_1 must be accepted by one player type and not the other.

After the first round, S has some probability p of making a deal with players of the other type, and some average price x_2 when a deal is made. Thus, S has the same profit as a strategy where the seller in round 2 offers a price x_2 with probability p and, with probability $1-p$ cuts off negotiation with the buyer. We must have $x_2 < x_1$, else no player would reject x_1 in the first round in order to have the chance of paying x_2 in the second. Clearly, it is the higher-valued player type that accepts the first round-offer, since if the higher price is the best option for the low-valued player type it certainly will be for the higher-valued type. To sell to the lower valued buyer, the second round price x_2 must be no higher than v_2, and to maximize profits, the seller might as well set $x_2 = v_2$.

In summary, we restrict our attention to strategies that offer an initial price x_1, and a lower price $x_2 = v_2$, with probability p, in the second round. The high-valued buyers accept the higher price immediately while the low-valued buyers reject the initial price and accept the lower price if it is made in the second round.

Figure 1 shows the potential profit regions for both buyers and sellers. Low-valued buyers never gain any profit, since the price x_2 is set at their reservation value v_2. A high-valued buyer gains $v_1 - x_1$ with probability 1 from buying in the first round, so the profit is regions A and B. If she waits until the second round, then with probability p, the gain is $v_1 - x_2$ (regions A and C). The seller gains $x_1 - s$ from selling at price x_1 (regions C, D, E, and F). If the game goes to the second round, the seller gains $x_2 - s$ with probability p (region E). So long as the price $x_2 > s$, we can, without loss of generality, assume $s = 0$, which will simplify some of the equations.

Intuitively, the proof proceeds as follows. In order for price discrimination to be better than offering everyone the lower price v_2, the seller needs x_2 (and, hence, v_2) to be much lower than x_1, so that the extra possibility of getting regions C and D from the high-valued buyers outweighs the losses from not collecting region F from the low-valued buyers. But in order for price discrimination to be better than offering

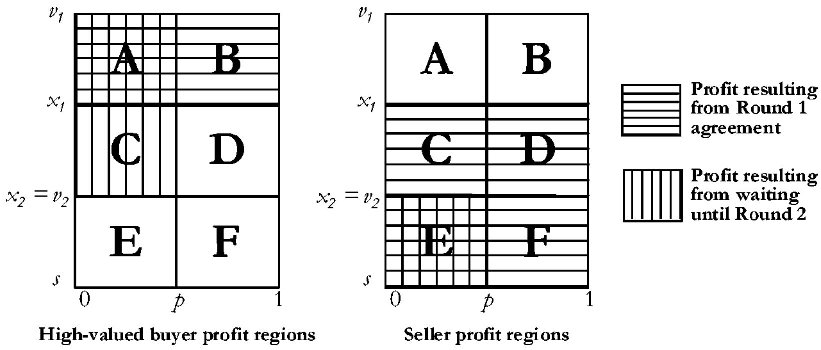

Figure 1. Buyer and seller profit regions when anonymous return is impossible.

only the higher price v_1, the seller needs x_2 to be close to x_1, so that the addition of region E from low-valued buyers more than makes up for the loss of A and B from high-valued buyers. No matter how the seller chooses x_1 and p, it is impossible to successfully price discriminate and make more money than is possible with the take-it-or-leave-it prices v_1 and v_2. The remainder of the proof fills in the details of this argument.

In order to induce high-value players to pay the higher price x_1, sellers must select values for x_1 and x_2 that satisfy the following constraint:

$$v_1 - x_1 \geqslant (v_1 - x_2)p, \qquad \begin{array}{|c|c|}\hline \mathrm{A} & \mathrm{B} \\ \hline \end{array} \geqslant \begin{array}{|c|}\hline \mathrm{A} \\ \hline \mathrm{C} \\ \hline \end{array}. \tag{1}$$

High-valued buyer prefers high price in first round over waiting.

Simplifying, we get:

$$(v_1 - x_1)(1 - p) \geqslant (x_1 - v_2)p, \qquad \boxed{\mathrm{B}} \geqslant \boxed{\mathrm{C}}. \tag{1a}$$

It is clearly possible to select values for x_1 and x_2 in which this is true, indicating that the seller can price discriminate among the two buyer types, at least for some values of v_1 and v_2.

Now consider whether it is profitable for the seller to price discriminate in this fashion. In particular, when is the price discrimination strategy more profitable than the two single-price strategies of v_1, selling only to the high-valued buyers, and v_2, selling to all buyers at a somewhat lower price? For price discrimination to be profitable, the following constraints must hold:

$$qx_1 + (1 - q)v_2 p \geqslant qv_1, \qquad q\begin{array}{|c|c|}\hline \mathrm{C} & \mathrm{D} \\ \hline \mathrm{E} & \mathrm{F} \\ \hline \end{array} + (1 - q)\boxed{\mathrm{E}} \geqslant q\begin{array}{|c|c|}\hline \mathrm{A} & \mathrm{B} \\ \hline \mathrm{C} & \mathrm{D} \\ \hline \mathrm{E} & \mathrm{F} \\ \hline \end{array}. \tag{2}$$

Mixed strategy more profitable than single price v_1.

$$qx_1 + (1-q)v_2 p \geqslant v_2, \qquad q\begin{array}{|c|c|}\hline \mathrm{C} & \mathrm{D} \\ \hline \mathrm{E} & \mathrm{F} \\ \hline \end{array} + (1-q)\begin{array}{|c|}\hline \mathrm{E} \\ \hline \end{array} \geqslant \begin{array}{|c|c|}\hline \mathrm{E} & \mathrm{F} \\ \hline \end{array}. \tag{3}$$

Mixed strategy more profitable than single price v_2.

Simplifying, we get:

$$(1-q)v_2 p \geqslant q(v_1 - x_1), \qquad (1-q)\begin{array}{|c|}\hline \mathrm{E} \\ \hline \end{array} \geqslant q\begin{array}{|c|c|}\hline \mathrm{A} & \mathrm{B} \\ \hline \end{array}, \tag{2a}$$

$$q(x_1 - v_2) \geqslant (1-q)v_2(1-p), \qquad q\begin{array}{|c|c|}\hline \mathrm{C} & \mathrm{D} \\ \hline \end{array} \geqslant (1-q)\begin{array}{|c|}\hline \mathrm{F} \\ \hline \end{array}. \tag{3a}$$

Substituting constraint (1) into (2a), we get

$$(1-q)v_2 p \geqslant q(v_1 - x_1) \geqslant q(v_1 - v_2)p, \qquad (1-q)\begin{array}{|c|}\hline \mathrm{E} \\ \hline \end{array} \geqslant q\begin{array}{|c|c|}\hline \mathrm{A} & \mathrm{B} \\ \hline \end{array} \geqslant q\begin{array}{|c|}\hline \mathrm{A} \\ \hline \mathrm{C} \\ \hline \end{array}. \tag{2b}$$

Substituting constraint (1a) into (3a), we get:

$$\begin{aligned} (1-q)v_2(1-p) &\leqslant q(x_1 - v_2) & \qquad (1-q)\begin{array}{|c|}\hline \mathrm{F} \\ \hline \end{array} &\leqslant q\begin{array}{|c|c|}\hline \mathrm{C} & \mathrm{D} \\ \hline \end{array} \\ &= q\big[(x_1 - v_2)p + (x_1 - v_2)(1-p)\big] & &= q\left(\begin{array}{|c|}\hline \mathrm{C} \\ \hline \end{array} + \begin{array}{|c|}\hline \mathrm{D} \\ \hline \end{array}\right) \\ &\leqslant q\big[(v_1 - x_1)(1-p) + (x_1 - v_2)(1-p)\big] & &\leqslant q\left(\begin{array}{|c|}\hline \mathrm{B} \\ \hline \end{array} + \begin{array}{|c|}\hline \mathrm{D} \\ \hline \end{array}\right) \\ &= q(v_1 - v_2)(1-p), & &= q\begin{array}{|c|}\hline \mathrm{B} \\ \hline \mathrm{D} \\ \hline \end{array}. \end{aligned} \tag{3b}$$

In order to make (2b) and (3b) comparable, we divide (2b) through by p, and (3b) by $(1-p)$.

$$(1-q)v_2 \geqslant q(v_1 - v_2), \qquad (1-q)\begin{array}{|c|c|}\hline \mathrm{E} & \mathrm{F} \\ \hline \end{array} \geqslant q\begin{array}{|c|c|}\hline \mathrm{A} & \mathrm{B} \\ \hline \mathrm{C} & \mathrm{D} \\ \hline \end{array}, \tag{2c}$$

$$(1-q)v_2 \leqslant q(v_1 - v_2), \qquad (1-q)\begin{array}{|c|c|}\hline \mathrm{E} & \mathrm{F} \\ \hline \end{array} \leqslant q\begin{array}{|c|c|}\hline \mathrm{A} & \mathrm{B} \\ \hline \mathrm{C} & \mathrm{D} \\ \hline \end{array}. \tag{3c}$$

All of the inequalities must be binding in order to satisfy both (2c) and (3c). If price discrimination is at least as profitable as the take-it-or-leave-it offers of v_1 and v_2, then, in fact, all three strategies must yield exactly the same profit. Thus, the seller cannot improve on a single take-it-or-leave-it strategy. $\square$

Trading data rather than money. Proposition 1 holds even if the unit of trade is data rather than money, as modeled by the introduction of buyer utility functions on the unit of trade. Only equation (1) is different, becoming $u_1(v_1 - x_1) \geqslant u_1(v_1 - x_2)p$. But since the utility function is linear, this just reduces to the constraint from the original equation (1).

3.1.2. Case 1b: Seller has perfect knowledge of buyer type

Next, we consider a situation with even less buyer anonymity. Here, a player has one identity for all time, allowing sellers to accumulate long-term profiles and to share their information about buyer types with other sellers. Hence, the seller begins the negotiation with full information about the buyer's type.

This case may occur when buyers have fixed IP addresses, for example, buyers who have direct Internet connections (as opposed to dial-up connections). It may also occur when buyers routinely provide personally-identifying information to the Web sites they visit.

Proposition 2. When buyers have full knowledge of the seller's strategy and no transaction costs, and sellers have perfect knowledge about each buyer's type, the optimal strategy for the seller is a take-it-or-leave-it offer.

Because the seller knows the buyer's type, the seller knows the set of offers that the buyer will accept. Thus, the seller can engage in perfect price discrimination, offering a different take-it-or-leave-it offer to each buyer, where the offer is the most favorable to the seller among all those that the buyer would accept. It will then be an equilibrium for the buyer to accept the offer. Of course, since the buyer also has perfect knowledge of the seller, there are many other equilibria, including one where the buyer refuses to pay more than the seller's reservation price.

If the seller-optimal equilibrium is the one that is normally arrived at, buyers may adjust their behavior in anticipation that any information they reveal could be used against them in later negotiations. For example, a buyer who might prefer to accept an offer in the current game if it stood on its own may reject the offer because to accept it would cause future offers from other sellers to be higher. Whatever type the buyer chooses to portray herself as, however, the buyer will have an incentive to continue to behave as if she has that type. From the sellers' perspective, then, the buyer really is of the type that she pretends to be, and the seller's optimal strategy is still to offer the reservation price of that buyer type as a take-it-or-leave-it offer.

3.2. Case 2: Limited knowledge of seller strategy

Proposition 3. When buyers have full anonymity and no transaction costs (but only limited knowledge of seller strategy), the optimal strategy for the seller is a take-it-or-leave-it offer.

In this case we relax the assumption that the buyer begins with complete knowledge of the seller's strategy. This would occur when there are no buyers' clubs or servers where sellers' strategies are posted. However, even if we relax this assumption, the other two assumptions imply that the buyer can acquire such knowledge at no cost. That is, the buyer can probe the seller by conducting a negotiation but then breaking it off before reaching an agreement. The buyer can then return, under a different pseudonym, and conduct another negotiation to further probe the seller's strategy. Thus, the ability to be anonymous and interact without transaction costs are sufficient to give the buyer as good a knowledge of the seller's strategy as is desired. The argument from the base case, then, implies that the seller can do no better than to use a fixed-price strategy, making a take-it-or-leave-it offer in the first round.

3.3. Case 3: Transaction costs

We return now to the assumptions of anonymous buyers who have perfect knowledge of the sellers' strategies. We consider, however, a situation where there are some transaction costs, which we can model either as a fixed cost for each round of bargaining, or a time discount factor on the value of the agreement. This corresponds to a situation of the Web browser negotiating with a Web site while the human buyer waits. A fixed cost per round of bargaining seems somewhat more realistic in this case: the total delay will usually not be significant enough to change the value of the interaction between buyer and Web site, but there is an opportunity cost to the lost time the person spends waiting. If the seller has a way to exploit the delay costs for the purposes of price discrimination, the seller can insert artificial delays to drive up those costs, even when they would otherwise be negligible. We consider two possibilities, one in which all people have the same opportunity cost for time, so that transaction costs are uniform, and another in which people exhibit varying levels of patience.

3.3.1. Case 3a: Uniform transaction costs

Suppose that delay costs are uniform among all buyers. Analogous to the "no anonymous return" scenario, some price discrimination is possible. Here, the seller cannot use the threat of cutting a buyer off completely (an action that would have variable cost to buyers with different values), because a buyer can return under a different pseudonym. Instead, the seller can threaten to delay, imposing a cost k per round of delay on the buyer. Faced with an offer from the seller, a buyer can accept it, wait for the next offer, drop out of the game, or return anonymously to restart the negotiation. There is no variability in the actual cost of delay among buyers, which on first analysis suggests that the seller cannot use delay as a means of price discrimination. There may, however, be variability in the lost revenue to a buyer from not making a deal in the first round: a high value buyer may lose a sure deal by waiting, while a low value buyer only forgoes a deal that he would not have accepted in any case. This opens an opportunity for price discrimination, if the seller's strategy is to offer a lower price in a later round with probability less than 1. As in the earlier scenario, however, we provide a proof that with two buyer types there is no way to make a simple price discrimination strategy more profitable than all the take-it-or-leave-it strategies. We conjecture but have not yet been able to prove that this result holds for any number of buyer types, and more complex price discrimination strategies involving more than two rounds.

Proposition 4. Assume that buyers have full anonymity, full knowledge of the seller's strategy, and uniform transaction costs. Suppose that a seller with reservation price s faces, with probability q, a buyer with reservation value v_1 and, with probability $1-q$, a buyer with reservation value v_2. Consider strategies where the seller offers a price x_1 in the first round, with probability p offers x_2 in the second round and otherwise always offers x_1. Moreover, assume that the strategy is restarted for each negotiating session, so that a buyer can end a negotiating session and return anonymously to face the same negotiating strategy. No such strategy is more profitable than every take-it-or-leave-it offer.

Proof. Assume that there exists a price discrimination strategy more profitable than any take-it-or-leave-it strategy. As in the proof of proposition 2, we can assume without loss of generality that this strategy induces high-valued players to pay x_1 immediately but lower valued players stay in the game until the next round, accepting the lower price x_2 if it is offered. There are two possibilities to consider if the lower price is not offered in the second round:

A. The low-valued player returns anonymously and waits through the first round, in the hopes of obtaining a lower price. If this is the case, the seller will eventually sell to all buyers and all players of type 2 will pay x_2, although some will wait longer than others before getting that offer. As we shall see below, however, the only prices that could induce this behavior have $v_2 > x_1$, which means that the seller could sell more profitably to all player types at the single take-it-or-leave-it price v_2.

B. The low-valued player drops out if the lower price is not offered. If this is the case, the proof proceeds analogously to that for proposition 2 and a contradiction is derived.

Proof for case A: Anonymous return for low-valued players if x_2 not offered in round 2. Let z be the expected value to players of type 2 from participating. There is always a delay cost of k since the player rejects the initial offer of x_1 and continues to the second round. Then, profit $v_2 - x_2$ is available with probability p. With probability $1-p$, the player will have to incur another delay cost k to return anonymously and restart the negotiation, which yields an expected profit z. Thus, z satisfies the following recurrence equation:

$$z = -k + p(v_2 - x_2) + (1-p)(-k + z) \tag{4a}$$

which simplifies to

$$z = v_2 - x_2 - k\left(\frac{2-p}{p}\right). \tag{4b}$$

In other words, players of type 2 will always pay x_2 if they wait long enough, and the expected delay cost from waiting is $k((2-p)/p)$. The difference in value between waiting for the price x_2 and taking the offer x_1 immediately is

$$v_2 - x_2 - k\left(\frac{2-p}{p}\right) - (v_2 - x_1) = x_1 - x_2 - k\left(\frac{2-p}{p}\right). \tag{4c}$$

Similarly, let y be the value that high-valued buyers could get from adopting the same strategy as low-valued buyers, continually returning anonymously and waiting for an offer of x_2. Following similar reasoning, we get

$$y = v_1 - x_2 - k\left(\frac{2-p}{p}\right). \tag{4d}$$

For high-valued players, the difference in value between waiting for the price x_2 and taking the offer x_1 immediately is

$$v_1 - x_2 - k\left(\frac{2-p}{p}\right) - (v_1 - x_1) = x_1 - x_2 - k\left(\frac{2-p}{p}\right). \tag{4e}$$

The difference in value between accepting x_1 and waiting for x_2 is the same for both player types, yet the two types choose different actions, so the difference in value must be 0, which makes both player types indifferent:

$$x_1 = x_2 + \frac{k(2-p)}{p}. \tag{4f}$$

Substituting (4f) into (4b), we find that:

$$z = v_2 - x_1. \tag{4g}$$

But we must have $z > k$ if low-valued buyers incur the delay cost k to return, rather than dropping out, and hence:

$$v_2 \geqslant x_1 + k. \tag{4h}$$

So, in this case, the seller could more profitably employ a take-it-or-leave-it offer of v_2, getting all buyer types to accept immediately.

Proof for case B: Low-valued players drop out if x_2 not offered in round 2. Now consider the case where low-valued players, instead of returning anonymously, drop out if the lower price is not offered. Figure 2 shows the potential profit regions for both buyers and sellers in this situation. Low-valued buyers may gain some profit if $v_2 - x_2 - k > 0$. A high valued buyer gains $v_1 - x_1$ with probability 1 from buying in the first round, so the profit is regions A and B. If she waits until the second round, at a cost of k, then with probability p, the gain is $v_1 - x_2$ (regions A, C and E). The seller gains $x_1 - s$ from selling at price x_1 (regions C–H). If the game goes to the second round, the seller gains $x_2 - s$ with probability p (region E). So long as the price

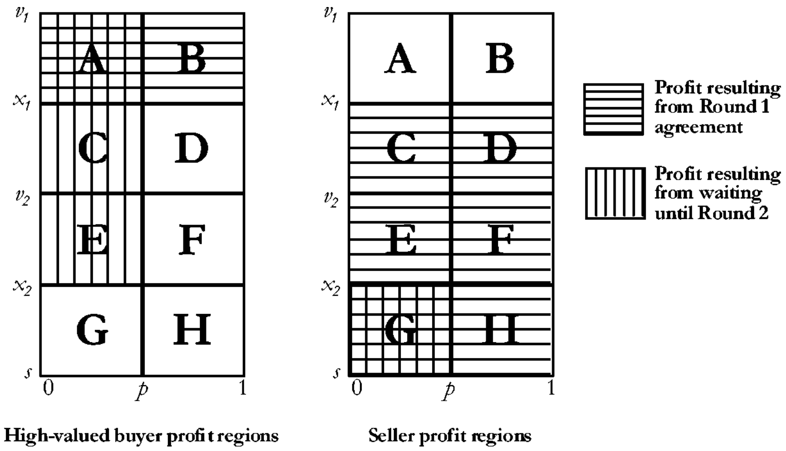

Figure 2. Buyer and seller profit regions with uniform buyer delay costs.

$x_2 > s$, we can, without loss of generality, assume $s = 0$, which will simplify some of the equations. We will use the labeled regions in figure 2 to illustrate our example below.

In order to induce the high-valued buyer to pay price x_1 immediately, sellers must select values for x_1 and x_2 such that

$$\mathrm{A} + \mathrm{B} > \mathrm{A} + \mathrm{C} + \mathrm{E} - k. \tag{5a}$$

In order to induce the low-valued buyer to stay in the game for the second round, we must have

$$\mathrm{E} - k > 0. \tag{5b}$$

Combining (5a) and (5b), price discrimination will be possible if and only if

$$\mathrm{A} + \mathrm{B} > \mathrm{A} + \mathrm{C}, \quad \text{or equivalently}, \quad \mathrm{B} > \mathrm{C}. \tag{5c}$$

If price discrimination is more profitable than charging a fixed price v_1, we must have

$$q(\mathrm{C} + \mathrm{D} + \mathrm{E} + \mathrm{F} + \mathrm{G} + \mathrm{H}) + (1 - q)\mathrm{G} \geqslant q(\mathrm{A} + \mathrm{B} + \mathrm{C} + \mathrm{D} + \mathrm{E} + \mathrm{F} + \mathrm{G} + \mathrm{H}). \tag{6a}$$

Simplifying,

$$(1 - q)\mathrm{G} \geqslant q(\mathrm{A} + \mathrm{B}). \tag{6b}$$

Substituting (5c),

$$(1 - q)\mathrm{G} \geqslant q(\mathrm{A} + \mathrm{C}). \tag{6c}$$

Since $\mathrm{E} \geqslant 0$,

$$(1 - q)(\mathrm{E} + \mathrm{G}) \geqslant q(\mathrm{A} + \mathrm{C}). \tag{6d}$$

Finally, dividing through by p,

$$(1-q)(\mathrm{E}+\mathrm{F}+\mathrm{G}+\mathrm{H}) \geqslant q(\mathrm{A}+\mathrm{B}+\mathrm{C}+\mathrm{D}). \tag{6e}$$

If price discrimination is more profitable than charging a fixed price v_2, we must have

$$q(\mathrm{C}+\mathrm{D}+\mathrm{E}+\mathrm{F}+\mathrm{G}+\mathrm{H})+(1-q)\mathrm{G} \geqslant \mathrm{E}+\mathrm{F}+\mathrm{G}+\mathrm{H}. \tag{7a}$$

Simplifying,

$$q(\mathrm{C}+\mathrm{D}) \geqslant (1-q)(\mathrm{E}+\mathrm{F}+\mathrm{H}). \tag{7b}$$

Substituting (5c),

$$q(\mathrm{B}+\mathrm{D}) \geqslant (1-q)(\mathrm{E}+\mathrm{F}+\mathrm{H}). \tag{7c}$$

Since $\mathrm{E} \geqslant 0$,

$$q(\mathrm{B}+\mathrm{D}) \geqslant (1-q)(\mathrm{F}+\mathrm{H}). \tag{7d}$$

Finally, dividing through by $1-p$,

$$q(\mathrm{A}+\mathrm{B}+\mathrm{C}+\mathrm{D}) \geqslant (1-q)(\mathrm{E}+\mathrm{F}+\mathrm{G}+\mathrm{H}). \tag{7e}$$

But (6e) and (7e) together imply that all the inequalities are binding, and thus the profits from price discrimination at best equal the largest profits from a take-it-or-leave-it strategy. Intuitively, v_2 must be large in order to make price discrimination better than ignoring the low-valued buyers, but it must be small to make price discrimination better than selling to everyone at v_2. These constraints are in conflict.

Trading data rather than money. Proposition 5 no longer holds if the unit of trade is data rather than money, as modeled by the introduction of buyer utility functions on the unit of trade. The natural interpretation of the uniform delay cost is still in terms of money or utility, and not in terms of data, which is the unit of trade. This effectively turns the case of uniform transaction costs into that of variable transaction costs. The cost of delay, when expressed in the same units as those used for trades, varies between buyers. As we shall see in the next section, the seller can sometimes profitably price discriminate when transaction costs are variable. For example, suppose that one player type cares very little about releasing personal data but that the other type cares somewhat more. If the seller offers to complete the transaction immediately with a large transfer of data, or in the next round for a smaller transfer, low-privacy buyers will accept the immediate offer while high-privacy buyers will wait. Even though the transaction cost for both player types is the same, the amount of value gained by transferring less data is variable and that is sufficient to enable profitable price discrimination.

3.3.2. Case 3b: Variable transaction costs

If transaction costs vary among buyers, sellers again have an opportunity to price discriminate, offering a higher price in early rounds and a lower price in later rounds to induce buyers with higher transaction costs to pay a higher price. Here, however, the seller has more latitude in its price discrimination than before, and can, in some cases, do so profitably.

Proposition 5. When buyers have full anonymity, full knowledge of the seller's strategy, and variable transaction costs, the seller can sometimes price discriminate successfully.

Expanding on the simple example above, suppose the low-valued player has delay costs k_2, while the high-valued player has delay costs k_1, and $k_2 < k_1$. The seller can price discriminate with a pure strategy. The seller offers a price x_1, with $v_2 < x_1 < v_1$ in the first round and commits to offering the price x_2 in the second round, with $x_2 + k_2 < v_2$. In figure 2, this would be like setting $p = 1$, so that only the left column of the profit regions are relevant (the areas of regions B, D, F, and H are all 0). Like the scenario with uniform delay costs, the high-valued buyer can gain at least as much as the low-valued buyer by waiting until the second round. But here the high-valued buyer loses more in transaction costs than the lower-valued buyer does.

To induce the high-valued buyer to pay the high price, we must have

$$v_1 - x_1 \geqslant v_1 - x_2 - k_1, \quad \text{or, equivalently,} \quad x_1 \leqslant x_2 + k_1. \tag{8}$$

To induce the low-valued buyer to stay in the game, we must have

$$v_2 - x_2 \geqslant k_2, \quad \text{or, equivalently,} \quad x_2 \leqslant v_2 - k_2. \tag{9}$$

The seller's best strategy will be to choose x_1 and x_2 as large as possible, subject to these constraints, that is,

$$x_2 = v_2 - k_2, \quad \text{and} \tag{8a}$$

$$x_1 = x_2 + k_1 = v_2 - k_2 + k_1. \tag{9a}$$

For the price discrimination strategy to be more profitable for the seller than the fixed price v_1, we must have

$$qx_1 + (1 - q)x_2 \geqslant qv_1. \tag{10a}$$

Simplifying,

$$(1 - q)x_2 \geqslant q(v_1 - x_1). \tag{10b}$$

Substituting (8a) and (9a),

$$(1 - q)(v_2 - k_2) \geqslant q(v_1 - v_2 + k_2 - k_1). \tag{10c}$$

Simplifying,

$$v_2 - k_2 \geqslant q(v_1 - k_1). \tag{10d}$$

In other words, the second buyer type's value must be enough larger than its delay cost, so that the second round price can be high enough to make up for the opportunity cost of a first-round price x_1 that is less than v_1. The higher the probability of a high-valued buyer, and the greater the spread between the high-valued buyer's value and delay cost, the greater must be the spread between the low-valued buyer's value and and delay cost. If the delay costs are negligible in comparison to the buyers' values, (10d) simplifies to $v_2 \geqslant qv_1$; v_2 must be closer to v_1 the greater the probability of a high-valued buyer.

For the price discrimination strategy to be more profitable for the seller than the fixed price v_2, we must have

$$qx_1 + (1-q)x_2 \geqslant v_2. \tag{11a}$$

Simplifying,

$$q(x_1 - x_2) \geqslant v_2 - x_2. \tag{11b}$$

Substituting (8a) and (9a),

$$qk_1 \geqslant k_2. \tag{11c}$$

In other words, k_1 must be enough larger than k_2, and the smaller the probability of a high-valued buyer, the larger k_1 must be. While (10d) and (11c) will not always hold, when they do, the seller can derive greater profit from the price-discrimination strategy than from any single-price strategy.

Perhaps surprisingly, it is possible for the buyers as well as the sellers to benefit from the seller's price discrimination, because it may allow the seller to "open new markets" [4]. When the seller employs only single-price strategies, the seller chooses the price as a monopolist would, missing the opportunity for trade with some low-valued buyers in order to maintain a higher price for the high-valued buyers. A price discrimination strategy can reduce that deadweight loss. Usually, a monopolist who can price discriminate captures all the additional surplus and indeed takes away whatever consumer surplus would otherwise exist, so that buyers, both individually and as a group, fare worse under price discrimination. In this case, however, in order to practice price discrimination, the monopolist may need to offer some of the deadweight savings to some of the buyers. While this does not happen when there are only two buyer types, the following example with three buyer types illustrates the possibility.

Suppose there are three buyer types. Table 1 shows the reservation values and transaction costs for each buyer type. The last column shows the probability that the seller will face each type of buyer.

If the seller can only consider single-price strategies, the optimal price is 100. The lower-valued buyer types are so infrequent that the lost revenue from offering a lower price to attract them will outweigh any additional profits from selling to them.

Consider, however, a seller strategy of offering to sell at a price 100 in the first round, 58 in the second round, or 28 in the third round. The high-valued buyers will still buy at the high price, the middle-valued buyers at the middle price, and the low-

Table 1
Reservation values and transaction costs.

Reservation value	Transaction cost per round of delay	Probability
100	100	0.9
90	30	0.05
30	1	0.05

valued buyers at the low price. The seller, however, cannot capture the entire surplus from the middle-valued buyer. The third round price cannot be larger than 28 while still retaining participation of the lowest-valued buyers, since such buyers have to pay transaction costs of 2 to wait until the third round. The difference between the second and third round prices cannot be more than 30, else the middle-valued buyers would choose to wait until the third round; hence the second round price cannot be more than 58. A middle-valued buyer has a value of 90, and loses 30 by waiting one round, leaving a surplus of 2 for such a buyer. Compared to the single-price strategy of 100 for all buyers, no buyers are worse off, and middle-valued buyers are better off.

4. Related work

Osborne and Rubinstein [8] and Wang [13] consider related negotiation scenarios that offer useful analogies and insights but are not directly applicable because they do not consider the effects of buyer anonymity or buyers sharing with each other knowledge of seller strategies. Osborne and Rubenstein introduce seller delay costs in addition to the buyer delay costs. We can interpret these delay costs as computational and operational costs of running the Web server. In their models, buyers have identical valuations, so that the negotiation can be viewed simply as bargaining over the division of the known surplus from trade.

Given buyer and seller delay costs, some strategies are not credible (e.g., it may not be credible for a seller to threaten to reject certain offers, when it is clear that the seller would lose money by waiting until the next round to continue negotiation). Osborne and Rubinstein restrict their attention to subgame perfect equilibria, those where both seller and buyer strategies are credible in all the possible subgames, even those subgames that are not reached in the equilibrium negotiation. They find that if the buyer and seller know each other's delay costs, and the set of possible offers is connected (that is, if prices of $2 and $4 are possible, then so is $3), the only subgame perfect equilibrium has the buyer accepting the seller's first offer (section 3.8).

In some other circumstances, however, there are subgame perfect equilibria that involve agreement later than the first round. For example, if the set of allowable offers is disconnected, there can be equilibria where the early actions effectively signal a player's bargaining toughness and a player who prematurely suggests a compromise in the first round will be forced to accept an extreme outcome (section 3.10). In those cases, the absence of available intermediate offers makes it possible to enforce the extreme outcome. They also consider situations where the seller has imperfect knowledge of the buyer's delay costs (chapter 5) and find that, analogous to our scenario 3b, the seller can price discriminate by selling at different prices in different rounds of the bargaining.

Negotiation protocol design is a close relative of mechanism design. Any equilibrium of bargaining under a particular negotiation protocol is also the outcome of some incentive-compatible mechanism.[4] The Revelation Principle[5] states that any outcome of an incentive-compatible mechanism can be achieved as the outcome of a direct mechanism, in which players do best by honestly reporting their types and an outcome is computed as a function of the set of types. Here, however, we are interested not merely in the set of equilibrium outcomes, but in features of the bargaining process, specifically whether multiple rounds are useful.

5. Discussion

We have identified three assumptions – anonymous buyers, perfect knowledge of seller strategies, and no transaction costs – that may sometimes hold in automated negotiations between Web sites and user agents, and analyzed several cases in which these assumptions hold in various combinations. While we have certainly not done an exhaustive analysis of every combination of assumptions, our limited analyses suggest some guidelines for the development of automated negotiation protocols.

For protocols designed to be used in situations where all three of our assumptions are likely to hold or where only the anonymity or knowledge assumptions are individually relaxed, designers should consider using restricted protocols that require sellers to make take-it-or-leave-it offers in the first round. In these cases, it is optimal for rational sellers who encounter rational buyers to make take-it-or-leave-it offers. By employing a restricted protocol, the optimal outcome may be reached, even when one of the players does not behave rationally. Thus it is not necessary to develop strategies that attempt to account for irrational players (an inherently difficult problem). Moreover, the overall implementation of the protocol is simplified.

One possible drawback of using a restricted protocol in the cases where take-it-or-leave-it offers are optimal is that sellers may not be able to learn as much from each prospective buyer during the learning phase (where sellers determine the distribution of buyer types in the population and decide where to set their prices). For some applications, building in a facility for prospective buyers to signal a seller about their type after they have turned down a seller's offer can allay this concern. For example, in a privacy negotiation protocol such as P3P, a seller might advertise that they require buyers to give them several pieces of personal information, which they

[4] If the corresponding mechanism were not incentive compatible, then, in the bargaining game, there is a player type that would pretend to be another type, thus ruining the equilibrium (cf. Osborne and Rubenstein [8, section 5.6]).

[5] See explanation in [12].

will use for a stated purpose. A buyer who refuses this offer might send a message to the seller explaining which piece of information or which part of the information use is objectionable. Further work should be done to model the benefits of such a signal to each party.

For protocols designed to be used in situations where there are transaction costs or where more than one of our other assumptions do not hold, it is not always optimal for the seller to make a take-it-or-leave-it offer. In these cases the protocol should probably allow for multiple rounds of negotiation. However, to gain some of the simplicity advantages of the restricted protocols, multi-round negotiation protocols should include a facility for sellers to make binding take-it-or-leave-it offers that are enforced by the protocol. Note that the ability to commit to take-it-or-leave-it offers does not confer any advantage on sellers that they did not already have, since they could pursue any strategy consistently and buyers would eventually learn of this consistency.[6] The main advantage of the ability to commit is that it will make it unnecessary for buyers to repeatedly probe sellers or consult with a buyers' club to determine that a seller will not change its offer.

Further work should be done to examine the cases where price discrimination through multi-round negotiation is an optimal seller strategy. If such cases are rare or difficult to identify, designers might consider using a restricted protocol; the lost opportunities for sellers caused by this restriction are likely to be outweighed by losses caused by irrational behavior that might otherwise result.

In cases where price discrimination may be an optimal strategy for sellers, it is also important to examine the consequences for buyers. For protocols such as P3P that are being designed with consumer protection in mind, if price discrimination does not produce an overall benefit for buyers, it might be desirable to force sellers to make a take-it-or-leave-it offer. Somewhat surprisingly, however, as we showed in case 3b, there may be some situations where price discrimination can benefit the buyers.

In cases where a fixed-price strategy is optimal for sellers, giving buyers the option of engaging in a multi-round protocol cannot benefit them. If a seller selects a fixed-price strategy, it would be irrational for them to accept any other offer from a buyer, even if such an offer was permitted, assuming that sellers do not pay any delay costs.

Another avenue for future research is to analyze when it may be in the interests of certain buyers to reveal their identities, despite the fact that anonymous interaction is possible. For example, buyers of certain types may wish to aid sellers in price discrimination if those buyers would not otherwise fare well in the seller's optimal one-price strategy.

[6] Unlike the single negotiation session model of Osborne and Rubenstein, we do not preclude strategies that are irrational in subgames of the current negotiation session, because pursuing such a strategy in that subgame may have reputational effects among buyers for future negotiation sessions.

6. Conclusion

Our analysis introduces interesting new twists to the study of negotiation protocols, based on the negligible transaction costs in automated negotiations and the asymmetry between buyers and sellers in Web-based negotiations. Sellers have fixed identities, permitting buyers to share knowledge of seller strategies; and buyers can be anonymous, foiling seller attempts to learn about particular buyers' strategies. These conditions give buyers significant bargaining power, so that, even if sellers have no transaction costs, they cannot profitably price discriminate. We expect that future research on negotiation and other mechanism design problems under varying conditions of anonymity and reputation sharing will prove fruitful.

References

[1] Anonymizer web site (1999). http://www.anonymizer.com/faq.html.

[2] L. Cranor and J. Reagle, Designing a social protocol: Lessons learned from the platform for privacy preferences, in: *Proc. of the Telecommunications Policy Research Conf.*, Washington, DC (1997). http://www.research.att.com/~lorrie/pubs/dsp/.

[3] E. Gabber, P. Gibbons, Y. Matias and A. Mayer, How to make personalized web browsing simple, secure, and anonymous, in: *Proc. of the Financial Cryptography 97*, Lecture Notes in Computer Sciences 1318 (Springer, New York, 1997).

[4] J. Hausman and J.K. MacKie-Mason, Price discrimination and patent policy, RAND J. Economics 19(2) (1988) 253–265.

[5] D. Kristol and L. Montulli, HTTP state management mechanism, IETF RFC 2109 (February 1997). http://www.ietf.org/rfc/rfc2109.txt.

[6] T. Negrino, So what are browser cookies, anyway?, Macworld Online (1996). http://www/macworld.com/netsmark/cookiestory.html.

[7] Organization for economic co-operation and development, OECD guidelines on the protection of privacy and transborder flows of personal data (1980). http://www.oecd.org/dsti/sti/it/secur/prod/PRIV-EN.HTM.

[8] M.J. Osborne and A. Rubinstein, *Bargaining and Markets* (Academic Press, San Diego, 1990).

[9] M.K. Reiter and A.D. Rubin, Crowds: Anonymity for Web transactions, Technical Report. http://www.cs.nyu.edu/~rubin/crowds.ps; DIMACS (August, 1997).

[10] J. Riley and R. Zeckhauser, Optimal selling strategies: When to haggle, when to hold firm, Quarterly J. Economics (May 1983) 267–289.

[11] M. Stefik, Trusted systems, Scientific American (March 1997) 78–81.

[12] H. Varian, Mechanism design for computerized agents, in: *Proc. of the Usenix Workshop on Electronic Commerce*, New York (1995). ftp://alfred.sims.berkeley.edu/pub/Papers/mechanism-design.pdf.

[13] R. Wang, Bargaining versus posted-price selling, European Economic Rev. 39 (1995) 1747–1764.

[14] World Wide Web Consortium, Platform for Privacy Preferences Project (P3P) (Web site containing overview and status of P3P. W3C recommendations with regards to P3P will be posted here when they become available publicly.) (1999). http://www.w3.org/P3/.