Cooperative vs. Competitive Multi-Agent Negotiations in Retail Electronic Commerce

Robert H. Guttman and **Pattie Maes**
MIT Media Laboratory
20 Ames Street, E15-301
Cambridge, MA 02139
{guttman,pattie}@media.mit.edu
http://ecommerce.media.mit.edu

Abstract

A key lesson learned from economic and game theory research is that negotiation protocols have substantial, rippling effects on the overall nature of the system. online auctions are increasingly popular negotiation protocols for software agents (and humans) to compete on the prices of goods and services. This paper takes a critical look at these competitive protocols in retail markets from economic, game theoretic, and business perspectives. Our analysis suggests that online auction protocols are, in fact, less efficient and more hostile than would be expected (or desired) in retail markets. Furthermore, we identify the importance of customer satisfaction and propose more cooperative multi-agent decision analysis tools (e..g, Multi-Attribute Utility Theory) and negotiation protocols (e.g., Distributed Constraint Satisfaction) as promising techniques to support it.

Keywords: agent-mediated retail, electronic commerce, cooperative negotiation

Cooperative vs. Competitive Multi-Agent Negotiations in Retail Electronic Commerce

Robert H. Guttman and **Pattie Maes**
MIT Media Laboratory
20 Ames Street, E15-301
Cambridge, MA 02139
{guttman,pattie}@media.mit.edu
http://ecommerce.media.mit.edu

Abstract

A key lesson learned from economic and game theory research is that negotiation protocols have substantial, rippling effects on the overall nature of the system. online auctions are increasingly popular negotiation protocols for software agents (and humans) to compete on the prices of goods and services. This paper takes a critical look at these competitive protocols in retail markets from economic, game theoretic, and business perspectives. Our analysis suggests that online auction protocols are, in fact, less efficient and more hostile than would be expected (or desired) in retail markets. Furthermore, we identify the importance of customer satisfaction and propose more cooperative multi-agent decision analysis tools (e..g, Multi-Attribute Utility Theory) and negotiation protocols (e.g., Distributed Constraint Satisfaction) as promising techniques to support it.

1. Retail Market Negotiations

Today's mass market retail is largely defined as *monopolistic competition* [1]. Whenever improved or unique products (e.g., faster computers or Tamagotchis) create a significant demand, similar products will eventually come to market that are very close (but not perfect) substitutes for the original. This new supply dismantles the monopoly and dissipates the demand. Thus, today's retail can be described as a competition amongst merchants for consumers' patronage.

The relationship that a merchant wishes to have with its customers, however, is not competitive. On the contrary, today's retail merchants desire highly cooperative, long-term relationships with their customers to maximize *customer satisfaction* [2]. This goal of maximizing customer satisfaction is to increase the probability of repeat purchases and new customers through positive reputation. The relationship among today's merchants and a given consumer is depicted in Figure 1.

But what will the landscape of retail look like in the future? Will tomorrow's online retail look anything like today's physical-world retail? We are already seeing online businesses that challenge the status quo (e.g., Amazon.com) and technologies that are dramatically changing the face of retail commerce – e.g., agent systems that reduce transaction costs for both merchants and consumers and create personalized and community-based experiences to help merchants increase sales [3].

However, certain fundamental economic characteristics of retail will likely *not* change. For example, future online retail markets will still have consumers who wish to satisfy their individual needs with well-matched merchant offerings. It is also likely that monopolistic competition will remain a defining characteristic of most online retail markets.[1] As such, merchants will still strive for highly cooperative, long-term relationships with their customers to maximize customer satisfaction.

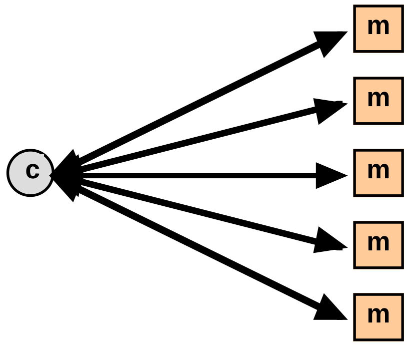

Figure 1 - Traditional retail markets. Merchants compete with one another for consumers' patronage. However, merchants desire a cooperative relationship with each of their customers.

As we have learned from economic and game theory research, a system's protocols have substantial, rippling effects on the overall nature of the system. Therefore, as designers of agent systems for mediating online transactions, we need to seriously consider which protocols we choose to employ. Although we have (and should exploit) the opportunity to prescribe new solutions to old problems, we may find that accurately modeling the competitive and cooperative levels among retailers and consumers will lead to more effective and efficient retail marketplaces as well as the long-term adoption and validation of our agent technologies for electronic commerce.

In this vain, this paper analyzes several electronic markets and their corresponding negotiation protocols. In particular, we take a critical look at competitive negotiation protocols for their appropriateness in online retail markets from economic, game theoretic, and business perspectives. Finally, we explore cooperative negotiation protocols that show more promise in accurately modeling the economic relationships in retail electronic commerce.

2. Competitive Negotiations

Negotiation is a form of decision-making where two or more parties jointly search a space of possible solutions with the goal of reaching a consensus [4]. Economics and Game Theory describe such an interaction in terms of protocols and strategies. The protocols of a negotiation comprise the rules (i.e., legitimate actions) of the game. An example of a simple negotiation protocol is the non-discriminatory English auction where (in one form) the only legal action is to (publicly) bid higher than the current highest bid by at least the minimum bid amount before the auction closes.

For a given protocol, a bidder uses a rational strategy (i.e., a plan of action) to maximize his or her *utility*. Decision analysis tools help identify optimal strategies given a bidder's preferences and knowledge (e.g., motivation, valuation, risk, information asymmetry, etc.) and is captured by a *utility function*. *Expected utilities* are utilities that

[1] There has been much speculation on the role of intermediaries in online markets. Many believe that intermediaries (e.g., retailers) will be decreasingly relevant as technologies begin to take over their roles and as manufacturers bypass their current distribution channels for direct sales. However, even if this will be our future, the resulting marketplaces will be more competitive (not less) forcing merchants to compete ever more fiercely for consumers' patronage.

consider probabilistic outcomes and events such as an opponent's future reaction to a player's action. Oftentimes, a utility function only reflects a player's self-interest. In other cases, it encompasses desirable social welfare or global performance goals such as system-wide equilibria [5, 6].

Competitive negotiations can be described as the decision-making process of resolving a conflict involving two or more parties over a single mutually exclusive goal. The Economics literature describes this more specifically as the effects on market price of a limited resource given its supply and demand among self-interested parties [1]. The Game Theory literature describes this situation as a zero-sum game where as the value along a single dimension shifts in either direction, one side is better off and the other is worse off [4].

The benefit of dynamically negotiating a price for a product instead of fixing it is that it relieves the seller from needing to determine the value of the good a priori. Rather, this burden is pushed into the marketplace itself. A resulting benefit of this is that limited resources are allocated fairly – i.e., to those buyers who value them most. As such, competitive negotiation mechanisms are common in a variety of markets including stock markets (e.g., NYSE and NASDAQ), fine art auction houses (e.g., Sotheby's and Christie's), flower auctions (e.g., Aalsmeer, Holland), and various ad hoc haggling (e.g., automobile dealerships and commission-based electronics stores). More recently, software agents have been taught competitive negotiation skills (e.g., auctioneering and auction bidding skills) to help automate consumer-to-consumer, business-to-business, and retail shopping over the Internet [3].

2.1. Classified Ad Negotiations

Kasbah [7, 8] is a Web-based multi-agent classified ad system where users create buying agents and selling agents to help transact goods. These agents automate much of the buyers' and sellers' brokering and negotiations. A user wanting to buy or sell a good creates an agent, gives it some strategic direction, and sends it off into a centralized agent marketplace. Kasbah agents proactively seek out potential buyers or sellers and negotiate with them on behalf of their owners. Each agent's goal is to complete an acceptable deal, subject to a set of user-specified constraints such as a desired price, a highest (or lowest) acceptable price, and a date by which to complete the transaction.

Negotiation between Kasbah buying and selling agents is bilateral, competitive, and straightforward as shown in Figure 2. After buying agents and selling agents are matched, the only valid action in the negotiation protocol is for buying agents to offer a bid to sellers. Selling agents respond with either a binding "yes" or "no".

Given this protocol, Kasbah provides buyers with one of three negotiation "strategies": anxious, cool-headed, and frugal

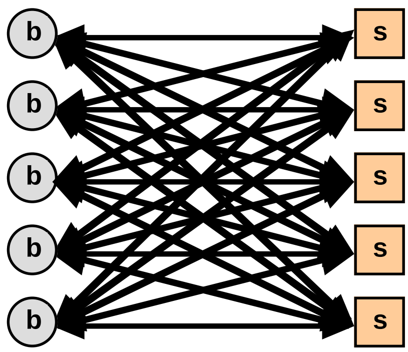

Figure 2 - Traditional classified ad markets. Buyers and sellers are in ad hoc, bilateral, competitive negotiations with one another respectively for unique (but potentially similar) limited resources.

– corresponding to a linear, quadratic, or exponential function respectively for increasing its bid for a product over time. The simplicity of these negotiation heuristics makes it intuitive for users to understand what their agents are doing in the marketplace.[2] This was important for user acceptance as observed in a recent Media Lab experiment [7]. A larger Kasbah experiment is now underway at MIT allowing students to transact books and music [8].

2.2. Stock Market Negotiations (CDAs)

AuctionBot [10, 11] is a general purpose Internet auction server at the University of Michigan. AuctionBot users create new auctions to buy or sell products by choosing from a selection of auction types and specifying its parameters (e.g., clearing times, method for resolving bidding ties, the number of sellers permitted, etc.). Buyers and sellers can then bid according to the multilateral, competitive negotiation protocols of the created auction.

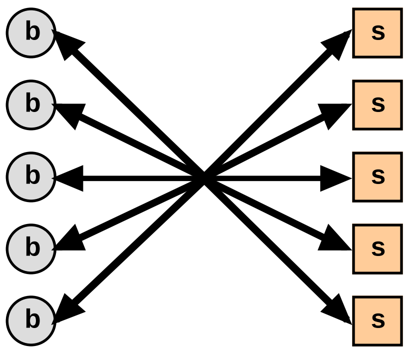

Figure 3 - Traditional stock markets. Within an unbiased, centralized marketplace, buyers and sellers are in multilateral, competitive negotiations with one another for limited resources.

One auction type that AuctionBot offers is the Continuous Double Auction (CDA). An example CDA market is the NASDAQ stock market. Kasbah's current protocols and negotiation strategies also resemble a CDA[3]; however, whereas Kasbah agents negotiate bilaterally, AuctionBot agents participating in a true and more efficient CDA negotiate multilaterally as shown in Figure 3.

Of the last two negotiation models discussed, the classified ad model (Figure 2) more closely resembles the retail market model (Figure 1). However, an important distinction is that there is no competition among consumers (i.e., buyers) in retail markets. This is because retailers typically sell production goods, not limited resources.[4] This lack of consumer competition in retail markets means that the actions of other consumers have negligible impact on retailers' current prices of the goods in question.

2.3. Retail Auction Negotiations

Two of the original (non-academic) auction Web sites are OnSale [12] and eBay's AuctionWeb [13] and are still very popular. Likely reasons for their popularly include their novelty and entertainment value in negotiating the price of everyday goods, as well as the potential of getting a great deal on a wanted product. In any case, the popularity of OnSale and eBay's AuctionWeb has quickly spawned an already competitive and growing

[2] Unlike other multi-agent marketplaces [9], Kasbah does not concern itself with optimal strategies or convergence properties. Rather, Kasbah provides more descriptive strategies that model typical haggling behavior found in classified ad markets.

[3] For more on this perspective, see [7].

[4] The pricing of production goods is based on marginal costs – a very different economic model than auctioning limited resources. (More on this later.)

industry. Yahoo! lists more than 90 active online auctions today [14]. Forrester Research reports that auctions will be core to making business-to-business transactions more dynamic, open and efficient [15]. online auctions like FastParts [16] and FairMarket [17] are already making this happen in the semiconductor and computer industries.

What's most relevant here is that many online auctions are augmentations to *retail* sites with retailers playing the roles of both auctioneer and seller (i.e., a sales agent). For example, First Auction [18] is a service of Internet Shopping Network, one of the first online retailers. Cendant's membership-driven retail site, netMarket [19], has also recently added auctions to its repertoire of online services. New auction intermediaries such as Z Auction [20] offer their auction services to multiple manufactures and resellers as a new sales channel.

With this much "auction fever," you would think that auctions are a panacea for retail shopping and selling. On the contrary, upon closer look we see that auctions have rather hostile characteristics. For example, although the protocols for the two most prevalent types of online auctions, first-price open-cry English and Yankee [21], are simple to understand and bid, determining the optimal bidding *strategy* is non-trivial[5] and, more importantly, can be financially adverse. In fact, in first-price open-cry auctions (i.e., highest bid wins the good for that price), the winning bid is always greater than the product's market valuation. This is commonly known as "winner's curse" as depicted in Figure 4. This problem is exacerbated in retail auctions where buyers' valuations are largely private[6]. Buyers with private valuations tend to (irrationally) skew bids even further above the product's true value.

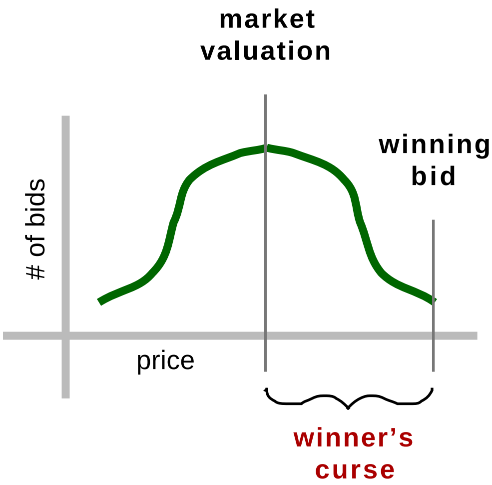

Figure 4 - "Winner's curse" is the paradox that the winning bid in an auction is greater than the product's market valuation. This occurs in all first-price open-cry auctions, the most prevalent type on the Internet.

Although winner's curse is a short-term financial benefit to retailers, it can be a long-term detriment due to the eventual customer dissatisfaction of paying more than the value of a product. Two universal auction rules that compound this problem are: (1) bids are non-retractable and, worse yet, (2) products are non-returnable. This means that customers could get stuck with products that they're unhappy with and paid too much for. In short, online auctions are less lucrative and far less forgiving than would be expected in retail shopping.

5 Factors to be considered include information asymmetry, risk aversion, motivation and valuation.

6 The motivation of a buyer with private-valuation is to acquire goods for personal consumption (or for gifts). This is in contrast to a buyer with common-valuation (e.g., in stock) where the motivation is to make money through the buying (and later reselling) of goods which have no other intrinsic value to the buyer.

Another customer dissatisfaction problem owing to online auctions is the long delay between starting negotiations and purchasing the product. For example, due to communication latency issues and wanting a critical mass of bidders, the English and Yankee auction protocols as implemented over the Internet extend over several days. This means that after a customer starts bidding on a product, she/he must continuously bid for the product (or have a shopping agent do it as provided by AuctionWeb) up until the auction closes several days later. This does not cater to impatient or time-constrained consumers.[7] To make matters worse, only the highest bidder(s) of an auction can purchase the auctioned good meaning that the other customers need to wait until the good is auctioned again and then restart negotiations.[8] Additionally, since bids are non-retractable and binding, consumers are unable to reconsider earlier buying decisions during this delayed negotiation stage.

There are other buyer concerns with English and Yankee style auctions such as *shills*. Shills are bidders who are planted by sellers to unfairly manipulate the market valuation of the auctioned good by raising the bid to stimulate the market. Although deemed illegal in all auctions, shills can be hard to detect especially in the virtual world where it is relatively inexpensive to create virtual identities (and thus virtual shills). Also, there is usually no negative consequence to the seller if one of his/her shills (accidentally) wins the auction.

Competitive negotiation auctions in retail markets also pose problems for *merchants*. Although auctions can relieve merchants of the burden of establishing prices for limited resources (e.g., fine art and stocks), this benefit is less realizable for *production goods* as in retail markets. Unlike fine art, for example, it is relatively easy to determine the marginal costs of production goods.[9] If auctioning these goods, however, it is non-trivial for the merchant to determine the optimal size of the auctioned lots and the frequency of their auction [22]. Such a determination requires an understanding of the demand for the good since it directly affects inventory management and indirectly effects production schedule.[10] Therefore, retailers are still burdened with determining the value of their goods a priori.

In addition, where sellers may have shills, buyers may collude by forming *coalitions*. A buyer coalition is a group of buyers who agree not to outbid one another. In a discriminatory (i.e., multi-good) auction, the result of this is that the coalition can buy goods for less than if they competed against one another thus unfairly cheating the seller. The coalition can then distribute the spoils amongst themselves (e.g., evenly, by holding a second private auction, etc.). As with shills, collusion through buyer coalitions is also considered illegal. However, as with shills, it can be hard to detect buyer collusion, especially in online markets where bidders are virtual. In fact, Multi-Agent Systems research has developed technologies that can efficiently form coalitions even among previously unknown parties [24] — posing an additional threat to online retail auctions.

As explained, online auctions are unnecessarily hostile to customers and offer no

[7] In fact, such delays are the antithesis to *impulse buying*.

[8] Even in traditional static catalog retail (as well as CDAs), consumers can purchase products immediately.

[9] Granted, the pricing of retail products can get involved. This is where marketing tactics come into play such as branding, market segmentation, price discrimination, etc.

[10] This relates directly to the just-in-time (JIT) concept for manufacturing, inventory, and retailing [23]. However, it is not yet clear how best to gauge demand in JIT inventory and retailing (e.g., via negotiation or sales).

long-term benefits to merchants. Essentially, they pit merchant against customer in price tug-of-wars. This is not the type of relationship merchants prefer to have with their customers [2]. Unlike most classified ad and commodity markets (e.g., stock markets), merchants often care less about profit on any given transaction and care more about long-term profitability. This ties directly to customer satisfaction and long-term customer relationships. The more satisfied the customer and intimate the customer-merchant relationship, the greater the opportunity for repeat customer purchases and additional purchases through direct referrals and indirectly through positive reputation.

Furthermore, consumers are much more in the driver's seat in online markets than in physical-world markets largely due to the dramatic reduction in search costs [25]. This increases the competition among retailers and forces them to positively differentiate themselves online in value dimensions other than price.

Looking at Figure 5, we see more clearly why first-price, open-cry auctions are particularly poor negotiation protocols for retail markets. Comparing traditional retail markets (Figure 1) with online retail auctions (Figure 5), we see that these online auctions invert the relationships among retailers and consumers. Instead of merchants competing for consumer patronage, online retail auctions force consumers to compete with one another for a specific merchant offering. Also, rather than supporting cooperative negotiations between merchants and customers, the relationships in retail auctions are highly competitive. Furthermore, since online retail auctions are time-consuming and bids are non-retractable and binding, bidding in one auction (generally) precludes the consumer from bidding for the same (or similar) product in another merchant's auction. This means that other merchants' prices have no direct impact on consumers' bids during the extended auction bidding period (which contrasts with stock markets as shown in Figure 4).

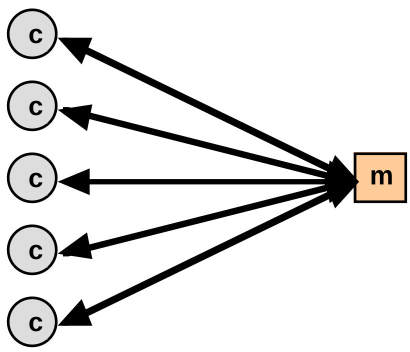

Figure 5 - online retail auctions. Consumers compete with one another for merchant offerings. Also, the relationship between the merchant and each of its customers is highly competitive.

3. Cooperative Negotiations

The degree of cooperation among negotiators falls within a continuum. After all, even in competitive negotiations, all parties need to cooperate sufficiently to engage in negotiation as well as agree on the semantics of the negotiation protocols.

However, one clear distinction that can be made between competitive and cooperative negotiations concerns the number of dimensions that can be negotiated across. For example, all of the competitive negotiation protocols discussed in the previous section allow for negotiation only within the price dimension. The cooperative negotiation protocols that we discuss in this section, on the other hand, allow agents (and humans) to negotiate over multiple dimensions.

Therefore, cooperative negotiations can be described as the decision-making process of resolving a conflict involving two or more parties over multiple interdependent, but *non*-mutually exclusive goals [26]. The study of how to analyze multi-objective decisions comes from economics research and is called multi-attribute utility theory (MAUT) [27]. The game theory literature describes cooperative negotiation as a non-zero-sum game where as the values along multiple dimensions shift in different directions, it is possible for *all* parties to be better off [4].

In essence, cooperative negotiation is a win-win type of negotiation. This is in stark contrast to competitive negotiation which is a win-lose type of negotiation.

Desired retail merchant-customer relationships and interactions can be described in terms of cooperative negotiation — the cooperative process of resolving multiple interdependent, but non-mutually exclusive goals. A merchant's primary goals are long-term profitability through selling as many products as possible to as many customers as possible for as much money as possible with as low transaction costs as possible. A customer's primary goals are to have their personal needs satisfied through the purchase of well-suited products from appropriate merchants for as little money and hassle (i.e., transaction costs) as possible. A cooperative negotiation through the space of merchant offerings can help maximize both of these sets of goals.

From a merchant's perspective, cooperative negotiation is about tailoring its offerings to each customer's individual needs resulting in greater customer satisfaction. From a customer's perspective, cooperative negotiation is about conversing with retailers to help compare their offerings across their full range of value resulting in mutually rewarding and hassle-free shopping experiences.

3.1. Multi-Attribute Utility Theory

Multi-objective decision analysis prescribes theories for quantitatively analyzing important decisions involving multiple, interdependent objectives from the perspective of a single decision-maker [27]. This analysis involves two distinctive features: an uncertainty analysis and a utility (i.e., preference) analysis. Techniques such as bayesian network modeling aid uncertainty analysis. Multi-attribute utility theory (MAUT) analyzes preferences with multiple attributes.

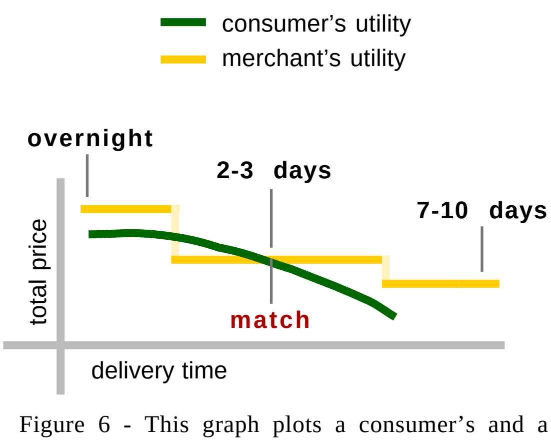

Figure 6 - This graph plots a consumer's and a merchant's multi-attribute utilities for a product's total price vs. delivery time (in days). In this example, the merchant offers three delivery options at different price points of which the "2-3 days" option best matches the consumer's utility profile.

Examples of uncertainty in retail shopping are "will she like this product as a gift?" and "how much do I trust this merchant?" Such uncertainties weighed against other factors play a part in consumers' buying decisions. From a merchant's perspective, analyzing an uncertainty like "what will be the demand for this product?" is vital for pricing products and managing inventory.

Often, decisions have multiple attributes that need to be considered. For example, in retail shopping, the price of a product could be important, but so could its delivery time. What is the

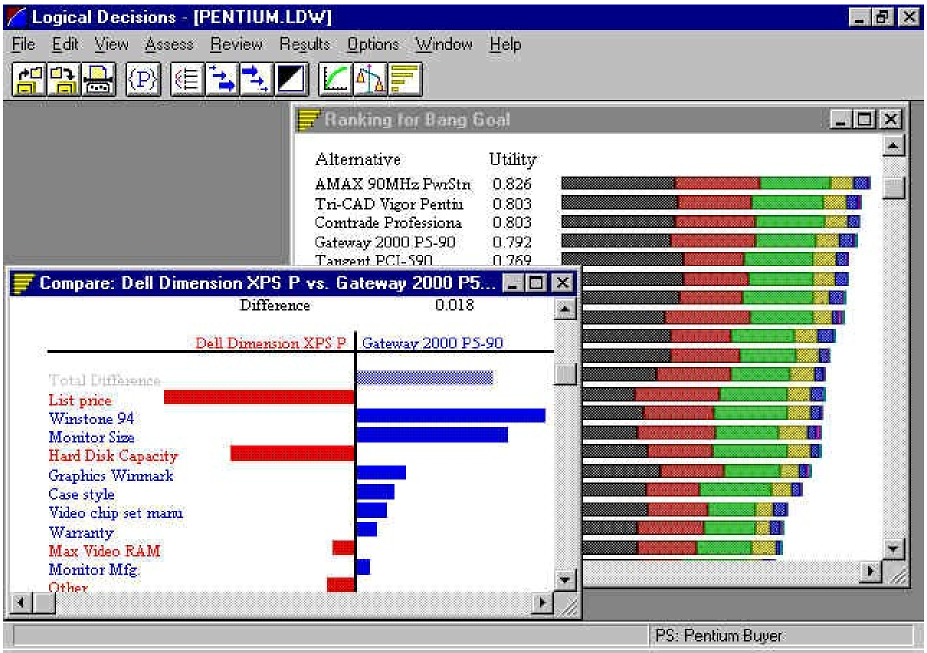

Figure 7 - A screenshot of Logical Decisions for Windows (LDW). This screenshot shows the results of a computer buying decision after LDW captured the decision-maker's utilities across multiple product attributes and value add. One results window shows the product rankings and the other a side-by-side comparison of two product

relationship and tradeoff between these two? Figure 6 gives a simple example of this.

Multi-objective decision analysis and MAUT can (and have) been used to tackle many different types of decision problems including electrical power vs. air quality, airport location, heroin addiction treatment, medical diagnostic and treatment, business problems, political problems, etc. These theories have also been instantiated in computer systems. The PERSUADER system at Carnegie Mellon University, for example, integrates Case-Based Reasoning and MAUT to resolve conflicts through negotiation in group problem solving settings [28]. Logical Decisions for Windows (LDW) by Logical Decisions, Inc. [29] is a general-purpose decision analysis tool for helping people think about and analyze their problems. Figure 7 shows LDW at work on a computer buying decision problem.

MAUT can help consumers make retail buying decisions. Likewise, MAUT techniques can help retailers create pricing policies for their added value – e.g., delivery options, extended warranty options, loan options, payment options, restocking fees, gift services, etc. However, MAUT alone does not constitute a cooperative negotiation protocol. Rather, MAUT is a decision analysis tool that can help capture and execute a strategy for an existing protocol. One such promising protocol is distributed constraint satisfaction.

3.2. Distributed Constraint Satisfaction

MAUT analyzes decision problems *quantitatively* through utilities. Constraint Satisfaction Problems (CSPs) analyze decision problems more *qualitatively* through constraints. A CSP is formulated in terms of variables, domains, and constraints. Once a decision problem is formulated in this way, a number of general purpose (and powerful) CSP techniques can analyze the problem and find a solution [30].

Finite-domain CSPs are one type of CSP and are composed of three main parts: a finite set of *variables*, each of which is associated with a finite *domain*, and a set of *constraints* that define relationships among variables and restricts the values that the variables can simultaneously take. The task of a CSP engine is to assign a value to each variable while satisfying all of the constraints. A variation of these "hard" constraints is the ability to also define "soft" constraints (of varied importance) which need not be satisfied. The number, scope, and nature of the CSP's variables, domains, and constraints will determine how constrained the problem is and, for a given CSP engine, how quickly a solution (if any) will be found.

Many problems can be formulated as a CSP such as scheduling, planning, configuration, and machine vision problems. In retail markets, CSP techniques can be used to encode hard constraints such as "I'm not willing to spend more than $2,000 for this product," and soft constraints such as "availability is more important to me than price." Even constraints such as "I prefer the Gateway 2000 P5-90 over the Dell Dimension XPS P (but I don't know why)" are legitimate. PersonaLogic (see Figure 8) uses CSP techniques to help shoppers evaluate product alternatives. Given a set of constraints on product features, PersonaLogic filters products that don't meet the given hard constraints and prioritizes the remaining products using the given soft constraints. This approach can also be applied to sales configuration systems such as Dell's "Build Your Own System" [31] and Trilogy's Selling Chain™ [32].

An important side-feature of CSPs is that they can clearly explain why they made certain decisions such as removing a product from the results list (e.g., "Product X is not

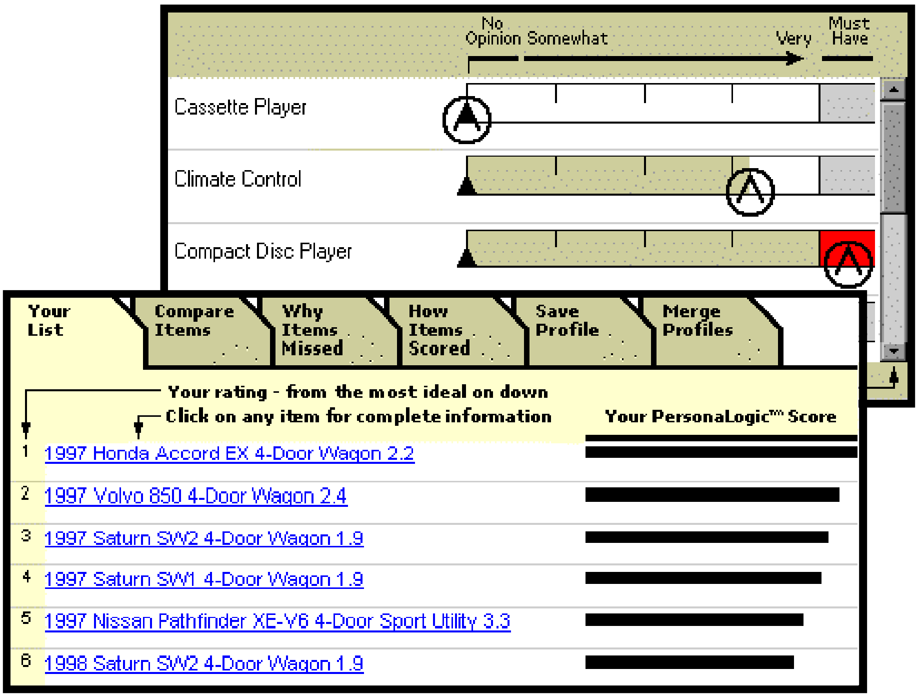

Figure 8 - Screenshots of PersonaLogic assisting a customer select automobile features with results.

an option because it has only 16 MB of RAM and you specified that the product should have at least 32 MB of RAM."). This feature is important because it relates to consumer trust (and ultimately, customer satisfaction). Trust is partially achieved by the shopping agent exhibiting somewhat predictable behavior and being able to explain its actions and decisions.

Most relevant to this paper, it should be possible to extend PersonaLogic to perform bilateral, cooperative negotiations between a merchant and a customer as shown in Figure 1 by using Distributed Constraint Satisfaction Problem (DCSP) protocols. DCSPs are similar to CSPs except that variables and constraints are distributed among two or more loosely-coupled agents [33]. This maps well to retail markets where consumers and merchants each have their respective set of constraints on merchant offerings. Furthermore, the loosely-coupled agents in a DCSP communicate to jointly solve a problem through established cooperative protocols [33]. This supports the desired cooperative relationship between a merchant and each of its customers.

However, DCSPs have been designed for fully cooperative group problem solving. Do DCSP techniques require more cooperation than is appropriate for merchant-customer interactions? For example, a customer may not be willing to divulge her reservation value (e.g., a willingness to pay up to $2,000 for a computer) to a merchant for fear of first-degree price discrimination with the merchant (unfairly) capturing all of the surplus in the market. However, first-degree price discrimination is tenuous in markets with monopolistic competition [1]. This suggests that DCSP techniques are not overly cooperative for bilateral, cooperative negotiations in retail markets.

4. Conclusion and Future Work

This paper analyzed several electronic markets and their corresponding negotiation protocols from economic, game theoretic, and business perspectives. We discussed how competitive negotiation protocols, and online auctions in particular, are inappropriate for online retail markets. Fundamentally, merchants strive for highly cooperative, long-term relationships with their customers to maximize customer satisfaction. This helps increase the probability of repeat purchases and new customers through positive reputation. Not surprisingly, none of the competitive negotiation protocols we discussed satisfied this need. Rather, they pitted merchant against customer in price tug-of-wars.

Cooperative multi-agent decision analysis tools and negotiation protocols, on the hand, appear to map much better to the retail market model depicted in Figure 1. For example, multi-attribute utility theory (MAUT) can help consumers make complex buying decisions (see Figure 7) taking into account multiple factors including merchants' unique added value (e.g., extended warranty options, delivery options, etc.). Constraint satisfaction techniques can also help consumers make complex buying decisions (see Figure 8) and this paper explored using cooperative distributed constraint satisfaction problem (DCSP) protocols to best support today's (and likely tomorrow's) retail market model (see Figure 1).

This analysis has guided the design of our new multi-agent system called Tete-a-Tete (T@T). T@T employs a combination of MAUT and DCSP techniques to mediate negotiations among consumer-owned shopping agents and retailer-owned sales agents. Once completed, we hope to show in our subsequent analysis of T@T that a bilateral, cooperative negotiation approach to retail electronic commerce allows merchants to tailor their offerings to each customer's individual needs resulting in more efficient markets and greater customer satisfaction than possible with competitive online auctions.

5. Acknowledgements

We would like to thank Peter Wurman and Michael Wellman for their help in refining our understanding of their AuctionBot system.

6. References

[1] R.H. Frank. *Microeconomics and Behavior, 3rd ed.* McGraw-Hill, Inc., 1996.

[2] Forrester Research Report. "Affordable Intimacy Strengthens online Stores." September, 1997.

[3] R. Guttman, A. Moukas, and P. Maes. "Agent-mediated Electronic Commerce: A Survey." To appear, *Knowledge Engineering Review*, June 1998.

[4] J. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* MIT Press, 1994.

[5] J. Cheng and M. Wellman. "The WALRAS algorithm: A Convergent Distributed Implementation of General Equilibrium Outcomes." To appear, *Computational Economics*.

[6] T. Sandholm and V. Lesser. "Equilibrium Analysis of the Possibilities of Unenforced Exchange in Multiagent Systems." *14th International Joint Conference on Artificial Intelligence (IJCAI'95)*. Montreal, Canada, 1995.

[7] A. Chavez, D. Dreilinger, R. Guttman, and P. Maes. "A Real-Life Experiment in Creating an Agent Marketplace." *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97).* London, UK, April 1997.

[8] Kasbah URL: <http://kasbah.media.mit.edu/>

[9] C. Sierra, P. Faratin, and N. Jennings. "A Service-Oriented Negotiation Model Between Autonomous Agents." *Proceedings of the Eighth European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW'97)*. Ronneby, Sweden, May 1997.

[10] P. Wurman, M. Wellman, and W. Walsh. "The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents." In the *Proceedings of the Second International Conference on Autonomous Agents*. May, 1998.

[11] AuctionBot URL: <http://auction.eecs.umich.edu/>

[12] OnSale URL: <http://www.onsale.com/>

[13] eBay's AuctionWeb URL: <http://www.ebay.com/aw>

[14] Yahoo! online Auction URL: <http://www.yahoo.com/Business_and_Economy/Companies/Auctions/online_Auctions/>

[15] Forrester Research Report, *Business Trade and Technology Strategies: Sizing Intercompany Commerce, vol. 1, no. 1.* July 1997.

[16] FastParts URL: <http://www.fastparts.com/>

[17] FairMarket URL: <http://www.fairmarket.com/>

[18] First Auction URL: <http://www.firstauction.com/>

[19] netMarket URL: <http://www.netmarket.com/>

[20] Z Auction URL: <http://www.zauction.com/>

[21] P. Milgrom. "Auctions and Bidding: A Primer." *Journal of Economic Perspectives*, pp. 3-22. Summer 1989.

[22] C. Beam, A. Segev, and J. G. Shanthikumar. "Electronic Negotiation through Internet-based Auctions." CITM Working Paper 96-WP-1019, December 1996.

[23] G. Morgenson. "The Fall of the Mall." Forbes, May 24, 1993.

[24] T. Sandholm and V. Lesser. "Coalition Formation among Bounded Rational Agents." *14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, Montreal, Canada, 1995.

[25] E. Schwartz. *Webonomics : Nine Essential Principles for Growing Your Business on the World Wide Web*. Broadway Books, 1997.

[26] R. Lewicki, D. Saunders, and J. Minton. *Essentials of Negotiation.* Irwin, 1997.

[27] R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* John Wiley & Sons, 1976.

[28] K. Sycara. "The PERSUADER." In *The Encyclopedia of Artificial Intelligence*. D. Shapiro (ed.), John Wiley and Sons, January, 1992.

[29] Logical Decisions URL: <http://www.logicaldecisions.com/>

[30] E. Tsang. *Foundations of Constraint Satisfaction*. Academic Press, 1993.

[31] Dell "Build Your Own System" URL: <http://www.dell.com/store/index.htm>

[32] Trilogy's Selling Chain™ URL: <http://www.trilogy.com/prodserv/products/main.html>

[33] M. Yokoo and E. Durfee. "Distributed Constraint Satisfaction for Formalizing Distributed Problem Solving." *Proceedings of the 12th IEEE International Conference on Distributed Computing Systems*, 1992.