# Agents that Buy and Sell: Transforming Commerce as we Know It

Pattie Maes, Robert H. Guttman, Alexandros G. Moukas Software Agents Group MIT Media Laboratory 20 Ames Street, E15-305 Cambridge, MA 02139 {pattie, guttman, moux}@media.mit.edu http://ecommerce.media.mit.edu

## Abstract

Software agents have become very popular in the last six or so years. They have been used successfully to filter information, match people with similar interests and automate repetitive behavior. More recently, the capabilities of agents have been applied to electronic commerce, promising a revolution in the way we conduct transactions, whether business-to-business, business-to-consumer, or consumer-to-consumer. This article surveys the technologies involved in buying and selling agents and discusses several agent-mediated electronic commerce systems in the context of a general model of the buying process.

## 1. Introduction

The Internet and World Wide Web are becoming an increasingly important channel for retail commerce as well as business to business transactions. Various recent studies By analysts from Nielsen, Forrester and IDC have shown that the numbers of web buyers, sellers and transactions are growing at a rapid pace.

However, the potential of the Internet for truly transforming commerce is largely unrealized to date. Electronic purchases are still largely non-automated. While information about different products and vendors is more easily accessible and orders and payments can be dealt with electronically, a human is still in the loop in all stages of the buying process, which adds to the transaction costs. A human buyer is still responsible for collecting and interpreting information on merchants and products, making decisions on merchants and products and finally entering purchase and payment information.

Software agent technologies can be used to automate several of the most time consuming stages of the buying process. Unlike "traditional" software, software agents are personalized, continuously running and semi-autonomous [1]. These qualities are conducive for optimizing the whole buying experience and revolutionizing commerce as we know it today [2]. For example, a company which needs to order additional paper supplies could have agents monitoring the quantity and usage patterns of paper within the

company, launching buying agents when supplies are becoming low. Those buying agents automatically collect information on vendors and products that may fit the needs of the company, evaluate the different offerings, make a decision on which merchants and products to pursue, negotiate the terms of transactions with these merchants and finally place orders and make automated payments.

## 2. Roles of Agents as Mediators in Electronic Commerce

It is useful to explore the roles of agents as mediators in electronic commerce in the context of a common framework. The model we present stems from consumer buying behavior (CBB) research and comprises the actions and decisions involved in buying and using goods and services. Although CBB research covers many areas, it is important to recognize its limitations up-front. For example, CBB research focuses primarily on retail markets (although most concepts pertain to business-to-business and consumer-to-consumer markets as well). Even within retail, not all shopping behaviors are captured (e.g., impulse purchasing). Also, electronic commerce covers a broad range of issues, some of which are beyond the scope of a consumer buying behavior model (e.g., back-office management, supply chain management and other merchant issues).

## 2.1. Buying Behavior Model

There are several descriptive theories and models that attempt to capture buying behavior -- e.g., the Nicosia model, the Howard-Sheth model, the Engel-Blackwell model, the Bettman information-processing model, and the Andreasen model. Although different, these models all share a similar list of six fundamental stages of the buying process. These six stages also elucidate where agent technologies apply to the shopping experience and allow us to more formally categorize existing agent-mediated electronic commerce systems [3]:

**1.** *Need Identification:* This stage characterizes the buyer becoming aware of some unmet need. Within this stage, the buyer can be stimulated through product information.

**2.** *Product Brokering:* This stage comprises the retrieval of information to help determine what to buy. This encompasses the evaluation of product alternatives based on buyer-provided criteria. The result of this stage is the "consideration set" of products.

**3.** *Merchant Brokering:* This stage combines the "consideration set" from the previous stage with merchant-specific information to help determine who to buy from. This includes the evaluation of merchant alternatives based on buyer-provided criteria (e.g., price, warranty, availability, delivery time, reputation, etc.).

**4.** *Negotiation:* This stage is about how to settle on the terms of the transaction. Negotiation varies in duration and complexity depending on the market. In traditional

retail markets, prices and other aspects of the transaction are often fixed leaving no room for negotiation. In other markets (e.g., stocks, automobile, fine art, local markets, etc.), the negotiation of price or other aspects of the deal are integral to the buying process.

**5.** *Purchase and Delivery:* The purchase and delivery of a product can either signal the termination of the negotiation stage or occur sometime afterwards (in either order). In some cases, the available payment (e.g., cash only) or delivery options can influence product and merchant brokering.

6. Product Service and Evaluation: This post-purchase stage involves product service, customer service, and an evaluation of the satisfaction of the overall buying experience and decision.

As with most models, these stages represent an approximation and simplification of complex behaviors. The stages listed above often overlap and migration from one to another can be non-linear and iterative.

Given the above set of stages, we can identify the roles of agents as mediators in electronic commerce. The personalized, continuously-running, autonomous nature of agents make them well-suited for mediating those consumer behaviors involving information filtering and retrieval, personalized evaluations, complex coordination, and time-based interactions. These roles correspond most notably to the need identification, product brokering, merchant brokering, and negotiation stages of the buying behavior model. Table 1 lists the six buying behavior stages and shows where several representative agent systems fall within this space.

	Persona Logic	Firefly	Bargain Finder	Jango	Kasbah	Auction Bot	T@T
1. need identification							
2. product brokering	П	П					П
3. merchant brokering			П	П			П
4. negotiation					П	П	П
5. payment & delivery							
6. service & evaluation							

Table 1: The online shopping framework with representative examples of agent mediation

## 2.2. Need Identification

To some extent, agent technology can be helpful in automating or assisting the buyer with the need identification stage. Specifically, agents can play an important role for those purchases that are repetitive (e.g. supplies) or predictable (e.g. habits). One of the oldest and simplest examples of software agents are so called "monitors": continuously running programs which monitor a set of sensors or data streams and take actions when certain pre-specified conditions apply. Examples of such monitoring agents are in abundant use in the stock market, as well as on e-commerce sites. For example, Amazon.com offers its customers a "notification agent" called "Eyes" which monitors the catalog of books for sale and notifies the customer when certain events occur that may be of interest to the customer (e.g., when a new book by author X becomes available, when a new book in category Y becomes available, etc).

## 2.3. Product Brokering

Once a buyer has identified a need to make a purchase (possibly with the assistance of a monitor agent), the buyer has to determine what to buy through a critical evaluation of retrieved product information. Table 1 shows several agent systems that lower consumers' search costs when deciding which products best meet their needs: PersonaLogic, Firefly, and Tete-a-Tete. PersonaLogic (http://www.personalogic.com) is a tool that enables consumers to narrow down the products that best meet their needs by guiding them through a large product feature space. The system filters out unwanted products within a given domain after a shopper specifies constraints on product features. A constraint satisfaction engine then returns a list of products that satisfy all of the shopper'Hs hard constraints ordered by how well they satisfy the shopper's soft constraints.

Tete-a-Tete uses comparable techniques to recommend complex products based on multiattribute utility theory. However, in contrast with PersonaLogic, Tete-a-Tete also assists buyers and sellers in the merchant brokering and negotiation stages and, hence, is discussed in section 2.5.

🔲 Netscape: Pa	assport 📃		
firefly guttman	guttman , here are so	me more recommended a	lternative artists.
directory	How recommendations w	gutunan , here are some mo How recommendations work	ore movies you should check out.
members	rale it!	*6: great stuff	Reservoir Dogs (1992) = P
venues	•		Directed by: Quentin Tarantino
filmfinder	-	*6: great stuff	Apocalypse Now (1979) 🔤 🛷
bignote	6: great stuff		Directed by: Francis Ford Coppola
My Yahoo!	•	*6: great stuff	The Godfather (1972) =
	-		Directed by: Francis Ford Coppola
talk help mess	sages exit	*6: great stuff	Citizen Kane (1941) 🗮
			Directed by: Orson Welles
		*6: great stuff	The Bicycle Thief (1948) 🗮
			Directed by: Vittorio de Sica

Figure 1: Firefly recommends simple products based on opinions of like-minded people

Like PersonaLogic, Firefly (http://www.firefly.com) and other systems based on collaborative filtering technology [4] help consumers find products (Figure 1). However, instead of filtering products based on features, Firefly recommends products via an automated "word of mouth" recommendation mechanism called collaborative filtering. The system first compares a shopper's product ratings with those of other shoppers. After identifying the shopper's "nearest neighbors" (i.e., users with similar taste), the system recommends products that neighbors had rated highly but which the shopper may not yet have rated, potentially resulting in serendipitous finds. Essentially, Firefly uses the opinions of like-minded people to offer recommendations. The system is used to recommend commodity products such as music and books, as well as harder to characterize products such as web pages and restaurants.

In addition to constraint-based and collaborative filtering techniques, two other techniques are widely used to implement product brokering or product recommendation agents. A large set of sites uses simple rule-based techniques such as those provided by Broadvision Inc to personalize product offerings to individual customers. Finally, a few sites experiment with data-mining techniques to discover patterns in customer purchasing behavior and exploit these to help customers find other products that meet their needs.

#### 2.4. Merchant Brokering

Whereas the product brokering stage compares product alternatives, the Merchant Brokering stage compares merchant alternatives. Andersen Consulting's BargainFinder (http://bf.cstar.ac.com/bf) is the first shopping agent for online price comparisons. Given

a specific product, BargainFinder requests its price from each of nine different merchant web sites looks up its price from about nine different merchant Web sites using the same requests as from a web browser. Although a limited proof-of-concept system, BargainFinder offers valuable insights into the issues involved in price comparisons in the online world. For example, a third of the on-line CD merchants accessed by BargainFinder blocked all of its price requests. One reason for this was because merchants inherently do not want to compete on price alone. Value-added services that merchants offered on their web sites were being bypassed by BargainFinder and therefore not likely considered in the consumer's buying decision. However, it was also the case that Andersen Consulting received requests from an equal number of smaller merchants who wanted to be included in BargainFinder's price comparison. In short, companies competing on price and/or desiring more exposure wanted to be included, the others didn't.

	Shopping for Laptops & Notebooks	
	● Find Products & Prices  ○ Find Reviews & Info	
	Manufacturer: No preference  Model: Price Range: Up to \$2000 Processor: Pentium MMX Processor Speed: 200MHz Find Laptops Find Laptops	
Shoppin	ng Results	
Hitachi	<u>VISIONBOOK PLUS</u> <u>PENT-200 MMX 2.1GB</u> <u>32MB 12.1 DSCAN 20X</u> <u>56K W95</u>	<sup>00</sup> Buy!
Toshiba	SATELLITE 310CDS BUYCOMP.COM \$1829. PENT-200 MMX 2.0GB 32MB 12.1 DSCAN 16X	95 Buy!
Compaq	Presario 1220 Portable <u>ComputerESP</u> \$1950. Media GX 200MMX 2.1GB/ 32MB Sync 12.1HPA 20X	<sup>00</sup> Buy!

Figure 2: Jango captures shoppers' preferences for price and a limited set of product features and returns a list of product offerings differentiated by price

Jango (http://jango.excite.com), seen in Figure 2, can be viewed as an advanced BargainFinder. The original Jango version "solved" the merchant blocking issue by

having the product requests originate from each consumer's Web browser instead of from a central site as in Bargain Finder. This way, requests to merchants from a Jangoaugmented Web browser appeared as requests from "real" customers. This kind of "aggressive interoperability" makes it convenient for consumers to compare prices from multiple merchants' online catalogs whether merchants want this or not. While virtual database technology (such as offered by Junglee, Inc) and learning techniques for semiautomatically composing "wrappers" for websites [5] are helpful when building comparison shopping agents, the process is still largely done by hand and is extremely tedious. In the near future, XML (see "The XML Revolution" article, this issue) as well as mobile agents technology (see "The Promise of Java-Based Mobile Agents", this issue) may make comparison shopping agents a lot more flexible, more open-ended and easier to implement.

## 2.5. Negotiation

During the negotiation stage, price or other terms of the transaction are settled on. The majority of business-to-business transactions involve negotiation. In retail, we are mostly familiar with fixed prices, even though those were only invented 100 years ago. The benefit of dynamically negotiating a price for a product instead of fixing it is that it relieves the merchant from needing to determine the value of the good a priori. Rather, this burden is pushed into the marketplace. A result of this is that limited resources are allocated fairly -- i.e., to those who value them most. However, there are impediments to using negotiation. In the physical world, certain types of auctions require that all parties be geographically co-located, for example, in auction houses. Also, negotiating may be too complicated or frustrating for the average consumer. Finally, some negotiation protocols occur over an extended period of time which does not cater to impatient or time-constrained consumers. In general, real-world negotiations accrue transaction costs that may be too high for either consumers or merchants.

Fortunately, many of these impediments disappear in the digital world. For example, OnSale (http://www.onsale.com) and eBay's AuctionWeb (http://www.ebay.com/aw) are two popular web sites that sell refurbished and second-hand products using a choice of auction protocols. Unlike auction houses, these sites do not require that participants be geographically co-located. However, these sites still require that consumers manage their own negotiation strategies over an extended period of time. This is where agent technologies come in. Table 1 shows several agent systems that assist the customer in negotiating the terms of a transaction: AuctionBot, Kasbah, and Tete-a-Tete.



Figure 3: AuctionBot offers many auction protocol permutations

AuctionBot (http://auction.eecs.umich.edu) is a general purpose Internet auction server at the University of Michigan. AuctionBot users create new auctions to sell products by choosing from a selection of auction types and then specifying its parameters (e.g., clearing times, method for resolving bidding ties, the number of sellers permitted, etc.) (Figure 3). Buyers and sellers can then bid according to the multilateral distributive negotiation protocols of the created auction. In a typical scenario, a seller would bid a reservation price after creating the auction and let AuctionBot manage and enforce buyer bidding according to the auction protocols and parameters. What makes AuctionBot different from most other auction sites, however, is that it provides an application programmable interface (API) for users to create their own software agents to autonomously compete in the AuctionBot marketplace. However, as with the Fishmarket Project [6], it is left to the users to encode their own bidding strategies. Fishmarket is not currently being used as a real-world system, but it has hosted tournaments to compare opponents' hand-crafted bidding strategies along the lines of Axelrod's prisoner's dilemma tournaments [7].



Figure 4: Kasbah is one of the first online agent systems for negotiating consumer products

MIT Media Lab's Kasbah (http://kasbah.media.mit.edu) [8] is an on-line, multi-agent consumer-to-consumer transactions system. 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 (Figure 4). Kasbah agents pro-actively 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 on behalf of its user subject to a set of user-specified constraints, such as a initial asking (or bidding) price, a lowest (or highest) acceptable price, a date by which to complete the transaction and restrictions on what parties to negotiate with and how to change the price over time. Kasbah's agents automate much of the Merchant Brokering and Negotiation stages for both buyers and sellers.

Negotiation in Kasbah is straightforward. 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 selling agents with no restrictions on time or price. 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 -- corresponding to a linear, quadratic, or exponential function respectively for increasing its bid for a product over time (similar functions exist for selling agents). The simplicity of these negotiation heuristics makes it intuitive for users to understand what their agents are doing in the marketplace. This was important for user acceptance as observed in a recent Media Lab experiment [8]. A larger Kasbah experiment is now underway at MIT allowing students to transact books and music.

The Kasbah system incorporates a trust and reputation mechanism called the "Better Business Bureau". Upon the completion of a transaction, both parties may rate how well the other party managed his/her half of the deal (e.g., accuracy of product condition, completion of transaction, etc.). Kasbah agents use accumulated ratings to determine if they should negotiate with agents whose owners fall below a user-specified reputation threshold. More details on the specific algorithms used in Kasbah's better business bureau can be found in [9].

Tete-a-Tete (http://ecommerce.media.mit.edu/tete-a-tete/) [10] provides a unique negotiation approach to retail sales. Unlike most other online negotiation systems which competitively negotiate over price, Tete-a-Tete's consumer-owned shopping agents and merchant-owned sales agents cooperatively negotiate across multiple terms of a transaction including warranties, delivery times, service contracts, return policies, loan options, gift services, and other merchant value-added services. Based on bilateral argumentation [11], Tete-a-Tete's integrative negotiations comprise an exchange of XML-based proposals, critiques, and counter-proposals. For example, a shopping agent may receive proposals from multiple sales agents. Each proposal defines a complete product offering including a product configuration, price, and the merchant's value-added services. The shopping agent evaluates and orders these proposals based on how well they satisfy its owner's preferences (expressed as multi-attribute utilities) (Figure 5). If the shopper is unsatisfied with the presented proposals, she can critique them along one or more dimensions. Her shopping agent broadcasts these preference changes to the sales agents which, in turn, use them to counter-propose better product offerings.



Figure 5: Tête-à-Tête's shopping interface

Tete-a-Tete considers product features and merchant features equally throughout negotiations to help the shopper simultaneously determine what to buy and who to buy from. This integration of Product and Merchant Brokering through integrative negotiations has the unique benefit that constraints on product features can affect the decision of who to buy from. For example, only a certain merchant may be able to support a particular product configuration. Likewise, constraints on merchant features can affect the decision of what to buy. For example, if no merchant can accommodate the overnight delivery of a specific product, an alternate product which can be delivered overnight may be determined to have a better overall value.

## 3. Conclusion and Future Directions

This paper explored how Software Agents are helping buyers and sellers combat information overload and expedite specific stages of the online buying process. Today's first-generation agent-mediated electronic commerce systems are already creating new markets (e.g., low-cost consumer-to-consumer and refurbished goods) and beginning to reduce transaction costs in a variety of business processes. The industries affected the earliest are those dealing with perishables (tickets, bandwidth availability, etc), surplus inventory and commodities (gas, electricity, pencils, music, books, etc).

However, we still have a way to go before software agents transform how businesses conduct business. This change will occur as software agent technologies mature to better manage ambiguous content, personalized preferences, complex goals, changing environments, and disconnected parties. The greatest changes may occur, however, once standards are adopted and evolved to unambiguously and universally define goods and services [10], consumer and merchant profiles, value-added services, secure payment mechanisms, inter-business electronic forms, etc.

Looking even further into the future, agents will explore new types of transactions in the form of dynamic relationships among previously unknown parties (see "MultiAgent Systems for Electronic Commerce", this issue.) At the speed of bits, agents will strategically form and reform coalitions to bid on contracts and leverage economies of scale -- in essence, creating dynamic business partnerships that exist only as long as necessary. It is in this third-generation of agent-mediated electronic commerce where companies will be at their most agile and marketplaces will approach perfect efficiency.

## 4. References

- [1] P. Maes. "Agents that Reduce Work and Information Overload." Communications of the ACM, Vol. 37, No.7, pp. 31-40, 146, ACM Press, July 1994.
- [2] A. Moukas, R. Guttman, P. Maes, "Agent-mediated Electronic Commerce: An MIT Media Laboratory Perspective", to appear, Proceedings of the International Conference on Electronic Commerce
- [3] I. Terpsidis, A. Moukas, B. Pergioudakis, G. Doukidis, and P. Maes. "The Potential of Electronic Commerce in Re-Engineering Consumer-Retail Relationships through Intelligent Agents." J.-Y. Roger, B. Stanford-Smith, and P. Kidd. (eds.), Advances in Information Technologies: The Business Challenge, IOS Press, 1997.
- [4] U. Shardanand and P. Maes. "Social Information Filtering: Algorithms for Automating 'Word of Mouth'." Proceedings of the Computer-Human Interaction Conference (CHI'95), Denver, CO, May 1995.
- [5] R. Doorenbos, O. Etzioni, and D. Weld. "A Scalable Comparison-Shopping Agent for the World Wide web." Proceedings of the First International Conference on Autonomous Agents (Agents'97). Marina del Rey, CA, February 1997.
- [6] J. Rodriquez, P. Noriega, C. Sierra, and J. Padget. "FM96.5: A Java-based Electronic Auction House." Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97). London, UK, April 1997.
- [7] R. Axelrod. The Evolution of Cooperation. Harper Collins, 1984.
- [8] 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.

- [9] G. Zacharia, A. Moukas and P. Maes. "Collaborative Reputation Mechanisms in Electronic Marketplaces". Proceedings of the HICSS-99 Conference, Electronic Commerce Minitrack, Hawaii, 1999.
- [10] R. Guttman and P. Maes. "Agent-mediated Integrative Negotiation for Retail Electronic Commerce." Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET'98), Minneapolis, Minnesota, May 1998.
- [11] S. Parsons, C. Sierra, and N. Jennings. "Agents that Reason and Negotiate by Arguing." To appear, Journal of Logic and Computation.

#### Bio Pattie Maes

Pattie Maes is an Associate Professor at MIT's Media Laboratory, where she founded and directs the Software Agents Group. She holds a Bachelor's degree and Ph.D. degree in Computer Science from the VUB. She was a founder and director of Firefly Network, Inc (acquired by Microsoft in 1998), one of the first companies to commercialize software agent technology.

#### Bio Robert H. Guttman

Robert H. Guttman is a co-founder of Frictionless Commerce, Inc., a start-up building the next generation of shopping advisory solutions for retail electronic commerce. Robert received an MS from the MIT Media Laboratory's Software Agents Group in 1998 and, as a Telecom Italia Fellow, he developed agent-mediated integrative negotiation protocols and decision support tools for retail markets.

#### Bio Alexandros G. Moukas

Alexandros G. Moukas is a co-founder of Frictionless Commerce, Inc. and an AT&T Fellow and doctoral candidate at MIT Media Laboratory's Software Agents Group (on leave.) He holds a BS degree in Business Administration from the American College of Greece, an MS degree in Artificial Intelligence from Edinburgh University, and an MS degree from MIT.