Examining Electronic Markets in Which Intelligent Agents Are Used for Comparison Shopping and Dynamic Pricing

by

Bryan M. Hertweck

Dissertation submitted to the faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in
Business, Management Science

APPROVED

____
Dr. Terry R. Rakes, Chairman

____
Dr. Loren P. Rees

____
Dr. Lance A. Matheson

____
Dr. Christopher W. Zobel

____
Dr. Evelyn C. Brown

September 2005
Blacksburg, VA

Keywords: Electronic Commerce, Intelligent Agents, Simulation, Neural Networks

Examining Electronic Markets in Which Intelligent Agents Are Used for Comparison Shopping and Dynamic Pricing

Bryan M. Hertweck

(ABSTRACT)

Electronic commerce markets are becoming increasingly popular forums for commerce. As those markets mature, buyers and sellers will both vigorously seek techniques to improve their performance. The Internet lends itself to the use of agents to work on behalf of buyers and sellers. Through simulation, this research examines different implementations of buyers' agents (shopbots) and sellers' agents (pricebots) so that buyers, sellers, and agent builders can capitalize on the evolution of e-commerce technologies.

Internet markets bring price visibility to a level beyond what is observed in traditional brick-and-mortar markets. Additionally, an online seller is able to update prices quickly and cheaply. Due to these facts, there are many pricing strategies that sellers can implement via pricebot to react to their environments. The best strategy for a particular seller is dependent on characteristics of its marketplace. This research shows that the extent to which buyers are using shopbots is a critical driver of the success of pricing strategies. When measuring profitability, the interaction between shopbot usage and seller strategy is very strong—what works well at low shopbot usage levels may perform poorly at high levels. If a seller is evaluating strategies based on sales volume, the choice may change. Additionally, as markets evolve and competitors change strategies, the choice of most effective counterstrategies may evolve as well. Sellers need

to clearly define their goals and thoroughly understand their marketplace before choosing a pricing strategy.

Just as sellers have choices to make in implementing pricebots, buyers have decisions to make with shopbots. In addition to the factors described above, the types of shopbots in use can actually affect the relative performance of pricing strategies. This research also shows that varying shopbot implementations (specifically involving the use of a price memory component) can affect the prices that buyers ultimately pay—an especially important consideration for high-volume buyers.

Modern technology permits software agents to employ artificial intelligence. This work demonstrates the potential of neural networks as a tool for pricebots. As discussed above, a seller's best strategy option can change as the behavior of the competition changes. Simulation can be used to evaluate a multitude of scenarios and determine what strategies work best under what conditions. This research shows that a neural network can be effectively implemented to classify the behavior of competitors and point to the best counterstrategy.

Acknowledgements

Thank you Lara Burkholder, my best friend and fiancée. Creating this document was difficult for many reasons, but meeting you allowed me to persevere and complete it. Your support was strong and unwavering and I am forever grateful for it. The years spent working toward this degree could have been the toughest of my life, but you made them the best of my life.

Thank you to my committee members for your guidance and support. Much of the progress required to complete this work was motivated by discussions with Dr. Terry Rakes and Dr. Loren Rees. Their leadership and insight is greatly appreciated. Thanks also to Drs. Christopher Zobel, Lance Matheson, and Evelyn Brown. Your service on my committee and the associated support was integral to my success.

Table of Contents

List of Tables and Figures

Chapter One: Introduction

The proliferation of e-commerce has already had a significant impact on how and where consumers shop. It also has the potential to make dramatic changes in the way goods are priced and how purchase decisions are made. The cost (in both time and money) of comparison shopping is drastically decreased with Internet shopping. Additionally, the Internet has opened the door for software agents to act on behalf or in support of buyers and sellers. Buyers' agents have become known as shopbots. Shopbots can collect and display market data—typically (but not limited to) prices. They can potentially make buying decisions and carry out sales transactions as well. Pricebots are similar to shopbots, but are employed by sellers. Pricebots adjust selling prices based on data obtained from the market and stored algorithms. Traditional static pricing remains the norm, but new markets based on dynamic pricing, auctions, reverse auctions, and ad hoc negotiations are gaining popularity. Software agents can assist buyers and sellers in all those markets. This paper examines marketplaces in which both buyers and sellers employ agent technology.

For those who work in information technology (IT) fields, the perception may be that e-commerce has been around for quite some time and that the field is relatively mature. While e-commerce has certainly had significant publicity and plays a major role in the shopping habits of many, it still represents a very small percentage of actual retail sales. According to U.S. Department of Commerce statistics (http://www.census.gov), retail e-commerce sales hit an all-time high in the second quarter of 2005 (at over $21 billion for the quarter). However, that figure is only 2.2% of total retail sales for the same period. That percentage may seem surprisingly small. The reasons are twofold:

(1) many use the Internet as part of the shopping process, but not to actually buy, and (2) many people still have limited access to computers and the Internet (i.e., they are on the other side of the "Digital Divide"). While e-commerce has created quite a buzz in the last decade, it still has massive potential for maturation with respect to evolution as an actual marketplace. This research aims to study what may happen as that maturation occurs.

Part of that maturation is likely to include increased use of agents. Websites such as mySimon.com and CNET.com provide agent-like services, such as listing and organizing product prices and specifications. Those services, however, are the tip of the iceberg for what true intelligent agents might be able to offer shoppers. A true intelligent agent can not only gather and present information, but can apply reasoning to data collected and make rational, unassisted decisions in the interest of its human employer. Sellers can also take advantage of the high visibility of Internet prices, using agents not only to gather information about competitors, but also to strategically set prices. This research examines scenarios where agents are employed in full-force within e-markets.

Finally, it is important to note that e-commerce is not just a retail phenomenon. Due to the high volume and repeat sales common the B2B arena, the availability of truly intelligent agents, which could save buyers money or maximize sellers' profits, would undoubtedly change the B2B landscape. Selling those agents, or their services, would likely become an industry of its own.

As e-commerce markets mature and become even more commonplace, both buyers and sellers will be confronted with new issues and questions to be addressed. Use of pricebots and shopbots will play a significant part in the evolution of e-commerce.

This research seeks to better understand how future electronic markets might work so that buyers, sellers, and agent builders can make better decisions.

Purpose and Plan

This examination of agent-based markets is geared toward generating advice for sellers, buyers, and agent builders. All of the planned work will give the above interested parties insight into the potential pricing patterns and attributes of agent-rich markets. Computer simulation of electronic markets will be the primary tool utilized in all phases of the research. Various statistical analyses of simulation output will generate results that can support quality decision-making.

For managers who oversee pricing within selling organizations, the research in Chapter Four will show how different pricing schemes perform versus others under varying market conditions. This research will demonstrate whether and how much the comparison shopping complexion of a market affects the choice of pricing strategy. Which strategies perform best in response to others will also be examined. The knowledge gained here will be very useful to managers and executives in markets where the Internet has made prices more visible than ever and raised the level of competition.

For buyers, the impact of varying parameterizations of buying agents on prices will be studied in Chapter Five. This work will provide insight into how valuable the service of a shopbot can be and will show how buyers' decisions about agent setup can directly impact their wallets. The interaction of parameters and their effects on the price a shopbot ultimately pays will be studied so that buyers can understand the potential trade-offs and compromises associated with their temporal requirements for purchase versus their need to save money. Sellers would also be wise to stay informed of the methods that buyers in their markets are using. Those who build shopping or buying agents also need to understand how agents can be built and what parameters can be varied with what effects.

Finally, the use of artificial intelligence in sellers' agents will be investigated in Chapter Six. This final portion of research will study agents built with artificial intelligence capabilities, specifically neural networks used to identify decision rules to follow, to determine whether such intelligent methods can enhance seller profitability. A pricebot system that can properly identify competitors' pricing strategies and react appropriately will be studied. Clearly, sellers will be interested in finding agent designs that are able to boost their profits. Those who design and build agents also need to understand what techniques work best and how their customers can benefit from those techniques. The level of potential profitability increase drives how much a builder is willing to spend on research, design, and development, and also drives pricing for agent products or services. The ultimate goal of this work is to show that agents can do more than just automate and facilitate what a human decision-maker could do—that an agent can be intelligent and have worth beyond merely saving time and effort.

Scope and Limitations

This research will utilize simulation techniques to examine price behavior in markets where agents are employed by both buyers and sellers. The simulations use five sellers and 100 buyers. Buyers in the simulations can have varying buying strategies. Some shop at all five sellers, mimicking the use of a shopbot. Others shop at one or two sellers, mimicking manual "browsing" or "surfing" on the Internet. Buyers also have an associated valuation, which is the price that represents the maximum the buyer is willing to pay for the good. Shoppers shop according to their assigned buying strategy and buy from the lowest-priced seller observed, assuming that price is below valuation. Buyers who enter the market have a need for only one unit of the available product and do not re-appear in the market.

Sellers change their prices periodically according to an assigned pricing strategy and are assumed to have full visibility of other sellers' prices. Sellers are also assumed to always have an adequate supply of goods available to meet demand. Both fixed and marginal costs are assumed to be equivalent across all sellers. The good offered in the market is assumed to be durable—i.e., it has no shelf-life or any other limitation on its availability. The good offered is also treated as a commodity—i.e., buyers make buying decisions according to price only.

As this research will rely heavily on simulation, any conclusions drawn will have inherent limitations. Law and Kelton (2000) assert that "one of the most difficult problems facing a simulation analyst is that of trying to determine whether a simulation model is an accurate representation of the actual system being studied." In this work, the "actual system being studied" is not a specific, existing market. Instead, this work seeks to model representative future markets where buyers and sellers capitalize on agent technology. Because there is no "actual system," validation of the simulation model is challenging. Therefore, the applicability of conclusions drawn from simulation results are dependent on how accurate the assumptions made about future markets turn out to be.

Generalizability is a concern with any scientific research. The results of this research will have certain limitations in their applicability. As mentioned above, the model used will focus on durable commodities. While the goods in the simulation model are true commodities, the results of this work may also be relevant to markets where the goods are similar to commodities—i.e., they are not identical across sellers, but the range of features and quality is relatively narrow. However, results will not be generalizable to perishable goods. In addition, many parameterization choices will be made in developing the simulation model for this research. These choices will be integral to what markets and scenarios can make use of the knowledge gained through this research. Sensitivity analysis will be performed on key parameters to gain insight into the extent to which results can be generalized to markets similar to the specific ones studied. Models will be built to maximize the range of applicability within the constraints of time and technology available.

Chapter Two: Background and Literature Review

Use of Agents in Markets

Kephart et al. (2000) discuss their vision of a coming "information economy" where *billions* of software agents exchange information goods and services with both humans and other agents. They predict use of agents by both buyers and sellers, as well as by intermediaries. They envision sophisticated economic agents that work independently and make use of algorithms designed to help them deliver optimal results for their human owners. As this information economy matures, they expect a new class of intermediary agents that serve the needs of other human-employed agents to appear. They also anticipate that virtually all agents will charge for their services. Kephart et al. note some key characteristics of agents that make studying the markets they will inhabit important. Economic software agents will be capable of making decisions and acting upon them more quickly than human decision makers. While these agents may become very sophisticated, they will likely be expert only in narrowly-defined tasks, lacking the more general capabilities and flexibility of humans. Therefore, agent-rich markets may have very different qualities from traditional markets.

Kanaan and Kopalle (2001) discuss the potential impacts of agent-enabled dynamic pricing on e-commerce markets and make associated propositions. They suggest that multiple aspects of consumer behavior will be changed by the emergence of pricebots. Their first proposition suggests that the use of shopbots will decrease with increasing dynamic pricing in a market. Kanaan and Kopalle argue that because many current bots rely on their own stored data to lessen search overhead and improve responsiveness they may be caught with out-of-date data which will lead to performance deterioration. This assertion assumes that while pricebot technology evolves, shopbot technology will stagnate, which seems an unlikely scenario in the current era of rapid technical innovation. Shopbot technology will likely evolve as well; in fact, it is not inconceivable that an arms race of sorts could evolve between buyers' agents and sellers' agents. They also assert that many sellers may block attempts of shopbots to learn price information. However, it seems unlikely that sellers would actively shut themselves off from a potentially large source of buyers. As Moukas et al. (2000) note, a seller may be able to effectively block requests from some sort of centralized source, but future agents may come directly from individual consumers, i.e., a non-agent request from a manual "surfer" may look exactly like a shopbot request. It would not be prudent to block or ignore that type of agent request. Kanaan and Kopalle also assert that trust in vendors will be negatively impacted by dynamic pricing. One segment they expect to be significantly impacted in this regard is that of less price sensitive, more loyal customer segments. If indeed that change occurs, those consumers will likely become less loyal and more likely to comparison shop. Because shopbots can help them shop, this pricebot-induced lack of trust may actually increase shopbot use in certain segments.

Smith (2002) surveys the shopbot literature, characterizes the impacts that buyers' agents might have on markets, and suggests avenues for future research. He discusses the notion that decreased search costs brought about by shopbots may drive retailer margins down. If that were to occur, clearly sellers would have a heightened interest in shopbot strategies that would give them a competitive edge. Additionally, Smith notes research that indicates that price dispersion can be increased in shopbot-enabled markets, due to strategies employed by sellers to take advantage of uninformed customers. Finally, Smith discusses other potential seller reactions to the introduction of shopbot technology to their markets. These seller reactions include "various obfuscation and bait-and-switch strategies" as well as paying shopbots for preferential listing of their products. Smith suggests that further research into when and where sellers should employ these strategies is needed.

Smith (2002) also discusses price sensitivity of Internet shoppers. Evidence suggests that while some Internet shoppers are very price sensitive, the value of factors such as product quality, branding, and delivery time still impact buyers' choices. He notes that while many analysts predicted that shopbot technology would give consumers more power in the marketplace, "research is far from unanimous" in supporting that viewpoint. Because of the many potential changes agent technology may introduce to e-markets, study of multiple scenarios is important.

Kauffman and Walden (2001) survey economic-oriented literature in the e-commerce arena and make suggestions about future paths of study. They offer four research directions:

1. Study the effect of "transparent" agents on Internet markets. Transparent, in this case, refers to the capability of an agent to reveal its source code, thus making its actions (for example, threats during negotiation) more credible.
2. Study the principal-and-agent nature of software agents—i.e., refine methods and models for how agents represent their principals' desires.
3. Develop knowledge about how to build agents that are robust information processors in markets with complex product characteristics.
4. Develop theories regarding the impact of new technologies on Internet markets—how different agent protocols affect markets.

The fourth direction is the aim of this dissertation. This research seeks to understand and describe market behaviors under varying conditions so that buyers, sellers, and agent builders can make smart decisions.

Dynamic Pricing Strategies

Kannan and Kopalle (2001) provide a taxonomy for e-commerce dynamic pricing strategies. They describe three high-level categories: posted price, auction pricing, and bundle pricing. Posted price includes both dynamic price updating and the use of e-coupons. Dynamic price updating simply involves the periodic update of a posted, fixed, and non-negotiable price. The Internet permits these updates to happen more frequently and at less cost than in traditional markets. E-coupons are electronic coupons that reduce posted prices and can vary over time and across customers. Auction pricing includes traditional auctions (e.g., eBay) as well as reverse auctions (e.g., Priceline) and exchange pricing (e.g., financial and commodity markets). Bundle pricing has two variations. The first is a volume discount, which could be earned by an individual consumer purchasing a large quantity or by a group of consumers bundling their demand and buying together. The second variation involves changes in price for an item dependent on what other items are purchased. All of these dynamic pricing schemes are important to the future of electronic markets, but this research is focused on dynamic posted price updating. As Pedersen (2000) notes, the Internet facilitates dynamic posted pricing by decreasing menu costs—the cost associated with changing prices. Virtual stores can change their price in one place for customers worldwide. Cortese (1998) makes the same assertion, noting that streamlined networks can reduce menu cost and time to nearly zero, even for companies with large product lines.

Varian (1980) develops a model for dynamic pricing designed to mimic price behavior in markets where items are periodically put “on sale.” He proposes a market model where there are both informed and uninformed consumers. Informed shoppers compare prices and buy at the lowest available price. Uninformed shoppers merely pick one retailer and purchase blindly. The pricing model he suggests has sellers drawing prices at random from a price density function. The seller who sets the lowest price in the market gets the business of all informed customers and the uninformed customers who choose him. Other sellers are limited to selling to only those uninformed customers who choose them. Varian makes a series of mathematical propositions about this model and concludes that sellers may find it in their interest to randomize prices in the interest of price discriminating between informed and uninformed customers. The Varian model is not a dynamic pricing “strategy,” but is useful for modeling the behavior of sellers who change their prices, but not necessarily according to a shopbot-implemented strategy.

Deck and Wilson (2003) simulate a market where there is a mix of informed and uninformed buyers. They, however, go a step beyond Varian (1980) by utilizing three types of buyers. One is uninformed, but there are two levels of informed—those who shop at two sellers and those who shop at all sellers. The level of comparison shopping is fixed throughout the study. They also utilize different pricing strategies from the randomized model proposed by Varian—they focus on undercutting, low-price matching, and trigger pricing. The undercut strategy simply involves setting a price for the next period that is lower than the lowest price observed in the prior period. Low-price matching involves querying currently available prices in the market and setting a price equal to the current minimum. A seller employs a trigger strategy by setting a current price along with a trigger value and associated potential new price. If the lowest price observed in the market is less than or equal to the trigger, a change to the new price is made.

Deck and Wilson (2003) carry out an experiment with computerized buyers and human subject sellers to investigate the strategies described above. A "manual" strategy where the subjects were free to set prices at their own discretion is also studied—a simulation where this was the only available strategy is used as the baseline treatment. Other treatments give subjects a choice between manual pricing and pricing utilizing one of the automated strategies. Throughout the experiment, sellers receive feedback in the form of competitors' prices. Four findings are described:

1. When sellers are allowed to set prices manually, the resulting distribution of prices is very similar to the theoretical Nash equilibrium, but with lower profits
2. Subject sellers use automated pricing strategies more often than they use manual pricing.
3. The undercut algorithm leads to median prices equivalent to the baseline. The low-price matching algorithm increases median prices above baseline. Trigger pricing leads to median market prices below those of the baseline.
4. Greater commitment to low-price matching shifts prices closer to the joint profit-maximizing outcome.

It is important to note that while Deck and Wilson (2003) make several conclusions in terms of the strategies studied, the data collected were not purely a result of the strategy employed. The data were also a result of decisions made by the human subjects—decisions of whether to adopt strategies and how to parameterize them. For each strategy treatment human decision makers had the option of employing the designated strategy *or* utilizing the baseline, non-algorithmic pricing method (i.e., manual pricing). The strategies were not universally adopted by subjects. In fact, the average adoption rate was as low as 62.5% (for low price matching) and as high as 81% (for trigger). Consequently, the competition for each strategy studied was variable (i.e., the low price match subjects competed with manual pricing more often than the trigger subjects did). Stricter control on seller behavior would allow for stronger conclusions about the relative performance of the pricing strategies.

Deck and Wilson (2002) also study varying implementations of a low-price match (LPM) strategy. Similar to the study above, they simulate a marketplace with computerized buyers and human subject sellers. Their study is a 2x2 factorial experiment where one factor is “LPM” vs. “no LPM” and the other factor is the level of customer search. For the “no LPM” condition, all four sellers submit a schedule of prices on which they prescribe prices depending on how many sellers a buyer had visited prior to inquiring about their price. The “LPM” treatment gives sellers the option of utilizing the “no LPM” price setting method or employing an LPM strategy. LPM consists of setting a price to be used if the seller is visited first. If they are not first, their price is set equal to the lowest of the prices that that buyer has already observed. The customer search factor is used to manipulate the level of informed vs. uninformed consumers. For the “low” condition, 61% of buyers visit a single seller, while visiting two, three, or four sellers are equally likely at 13% each. For the “high” condition, the probabilities for visiting one, two, three, or four sellers are equal at 25%.

Deck and Wilson (2002) report the following three findings:

1. High search leads to lower prices than low search for one-seller buyers. Low price matching leads to higher prices for four-seller buyers (regardless of high vs. low search)
2. Price matching leads to less discriminatory pricing for high search. With or without matching, different buyer types receive considerably different prices in the low treatment
3. Price matching increases seller profits in the high treatment, but has no effect in the low case.

Because sellers are given visibility to past activities of their potential buyers, the LPM strategy can be used to price discriminate between informed and uniformed buyers.

Because the human subjects were given the option of choosing LPM, the results of the Deck and Wilson (2002) study have the same shortcomings as the 2003 study described above. Deck and Wilson (2002) do not report adoption rates as in the 2003 work. Additionally, the strategies used in the 2002 study make significant assumptions about the capability of sellers to have visibility into the history of prospective buyers. Both the LPM and "no LPM" treatments assume that the seller knows which competitors have already been visited by the buyer before making a pricing decision. While the Internet enables a wealth of information sharing, the extent to which current or future technology can support that assumption is debatable.

Dasgupta and Smith (2003) also examine a seller strategy that attempts to price discriminate between buyers. They simulate a repeat buying market where 25% of the buyers shop at only one seller while the other 75% shop at all sellers. Buyers are repeat buyers of some consumable commodity and can be uniquely identified by the seller. The sellers set prices using an algorithm that seeks to determine which kind of buyer is requesting a price based on the historical behavior of that buyer. Buyers determined to be comparison shoppers are offered a price determined by a model optimizer (MO) algorithm. The MO uses a set of weighted historical price and profit data points to estimate (via nonlinear regression) the price that would maximize profits. A higher price can be charged to those buyers who are assumed to not be comparison shoppers. That optimum price is determined by attempting to learn the distribution of buyer valuations. The sellers are all trying to gain the largest market share, but do not have information about competitors' prices or the extent to which buyers are informed. Dasgupta and Smith find that seller profits can be improved by 15-20% by utilizing this tiered pricing strategy.

DiMicco et al. (2003) introduce a Java-based simulation tool called "Learning Curve." It is designed to allow sellers to investigate agent-based marketplaces and evaluate the relative effectiveness of dynamic pricing strategies. Their study is different from those above in that it focuses on perishable good markets where supplies and their time horizon of availability are finite. They focus on two seller strategies—Goal-Directed (GD) and Derivative Follower (DF)—both of which require no assumptions about or knowledge of the buyer population. The GD strategy strives to sell all inventory by the last day of the market, but not before. The DF strategy experiments with incremental price changes by examining changes in profits after a change occurs. If an increase (decrease) in price results in greater profits, the price is increased (decreased) again. If profits decline, the next change is a decrease (increase) in price. DiMicco et al. studied the performance of these two strategies under both monopolistic and competitive conditions. As an overall conclusion, they state that both strategies performed quite well, despite their lack of knowledge of the buyer population. The GD strategy is generally successful in reaching its stated goal, but “at the expense of drastically over- and under-shooting the buyer valuation curve early and late in the market.” The DF strategy is successful in consistently selling at the highest price possible on a given day, but its performance is sensitive to the timing of peaks in demand. They also note that the complexion of the buyer population with respect to comparison shopping can impact price patterns—i.e., when the market is “extremely price sensitive (100% comparison shoppers), adaptive strategies can easily breakdown into price wars.” (p.)

Kephart et al. (2000) examine a broad range of topics, all of which are related to agents that dynamically price goods or services. They first investigate an online bookselling market, where five sellers compete for business. They employ various combinations of pricebots strategies and examine the resultant price patterns. The strategies examined are game-theoretic (GT), myoptimal (MY), and derivative-follower (DF). The GT strategy involves the determination of a price density function that can be computed when given the number of sellers, the valuation distribution of the buyers, and the comparison shopping complexion of the market (i.e., what portion of the market shops at 1, 2, … , n sellers). The MY strategy requires the same market knowledge as GT, with the additional need to know current prices of all other sellers. The MY strategy uses that knowledge to calculate the value that will maximize profits, up until the moment that another seller changes its price. The DF strategy, as discussed above, requires no knowledge of competitors or buyers in the market.

Kephart et al. (2000) simulate nine scenarios: three pure markets where all sellers utilize the same strategy, and six hybrid markets where four sellers utilize one strategy while the fifth utilizes one of the other strategies. Half of the buyers in their study shop at all five sellers, a quarter at two of the sellers, and the final quarter shop at only one seller. Of the pure markets, the DF group maintains the highest level of profitability, despite being the least informed. Their profitability is not far from the optimal profitability that could be attained by a cartel. The MY pure market has the second highest level of profitability, with the GT profits far behind. In the hybrid environments, both the GT and MY pricebots outperform the DF's by consistently undercutting prices. They each lead to profits that are more than double what the DF seller in the market earns. Thus, while a pure market of DFs leads to high profitability, that market is inherently unstable. Self-interest would motivate individual sellers to switch to other strategies. MY is the most effective strategy to switch to (given that the necessary market data is available to the pricebot). Consequently, Kephart et al. study the MY strategy in greater depth.

They first examine the effect of the re-pricing rate on sellers' profitability. Through simulation, they confirm that in a pure MY market, the pricebot that re-prices most frequently becomes the most profitable. Next, they add a learning element to the MY pricebots to enable anticipation of competitor behavior. They employ a reinforcement learning technique called Q-learning in two-agent competitive scenarios. A Q-learning seller learns a function that expresses the anticipated future-discounted profit for each possible price that it might charge, given current competitor prices. When only one agent learns, the learning agent becomes more profitable. When both agents are Q-learners, the outcome is unpredictable—the outcome depends on the value of the future discount parameter associated with the learning.

Kutschinski et al. (2003) perform a series of sophisticated simulations that involve various pricing strategies. In their simulations, sellers and buyers both have accounts that track their performance. Sellers have the capability to set not only price, but also production levels. Buyers have periodic, fixed income and seek the good available in the market according to a pre-set demand curve. Buyers have two associated accounts: one to track financial resources available and another to track rewards earned through buying goods. Mirroring the work of Varian (1980), Kutschinski et al. create a market with both informed and uninformed buyers. Informed shop the entire market (at a cost), while uninformed shop at only one seller.

The first simulation Kutschinski et al. (2003) describe is one where four sellers are matched with 20 buyers. The sellers' production levels are fixed, i.e., price is the only decision made by the sellers. All sellers make that price decision utilizing the same strategy, which implements price changes according to inventory levels. High inventory levels (compared to a specified threshold) lead to price decreases and low levels lead to price increases. Increases are performed in fixed increments and they find that the value of the inventory threshold is unimportant to the qualitative results of the simulations. They perform a series of simulations with varying proportions of informed buyers. Increasing competitiveness (higher proportions of informed buyers) leads to lower profitability for the sellers and greater rewards for the buyers. The range of prices charged by sellers decreases as competitiveness is increased. In all scenarios, the market prices move from initial conditions up to a theoretical market clearing price. In higher competition scenarios, prices drop from that clearing price and later stabilize at a lower price.

Next, Kutschinski et al. (2003) examine markets where production levels are allowed to vary. Sellers again control only prices; however, unlike the previous simulation, production levels are varied automatically to meet demands. They first examine two seller markets where the DF and MY strategies, as described above in the context of Kephart et al. (2000), are employed. Results are comparable to the prior Kephart et al. study. Next, Kutschinski et al. offer a strategy to bridge the gap between DF and MY. The drawback of DF is its reliance on only one historical data point at a time. The disadvantage of MY is its need for perfect information about both competitors and buyers. Their price-profit (PP) adaptation strategy does not require the same level of knowledge as MY, but uses a more robust approach to tracking profitability than DF. The PP strategy learns from observed profits by building a model of profit expectations for a given price through a machine learning mechanism. In examining heterogeneous seller markets (with asynchronous price changes across those sellers), they discover that all the PP sellers converge to the same profit expectation model.

Finally, Kutschinski et al. (2003) examine a learning pricing strategy that takes competitors prices into account. They pit a Q-learning seller against DF, MY, and PP sellers in simulations with asynchronous price changes as above. The Q-learning seller is able to learn the profit function and undercutting best-reply strategy against all competitor types. The asynchronous nature of the price changes leads to slower convergence than prior similar studies.

Kutschinski et al. (2003) study an interesting and broad range of pricing strategies with a robust simulation model. Their results and conclusions, however, seem rather limited in their ability to guide the decisions of managers who make pricing decisions. In one study they examine the impact of varying levels of comparison shopping behavior, but they do not compare performance of multiple strategies. In simulations involving multiple strategies, the focus of their conclusions is more on qualitative aspects of price dynamics than on relative quantitative performance of pricing strategies. The research of Chapter Four seeks to make pricing strategy comparisons across varying comparison shopping levels that will aid decision making for pricing managers.

Leloup (2003) examines a dynamic pricing strategy that involves a selling agent attempting to learn an optimal strategy through a statistical decision theory approach. In Leloup's work, the market studied is one with only one selling agent and multiple buying agents. The selling agent constructs choice probabilities associated with potential buyer actions and sets a price to maximize profit. Dynamic programming can be used to "learn" an optimal strategy over time. Because the computational aspects of such an approach can become overwhelming, Leloup utilizes a beta-logistic formulation as an approximation of the dynamic programming technique. Because no other pricing strategies are studied and the market studied included only one seller, no conclusions about the effectiveness of this pricing strategy can be drawn. Leloup also studies interactions between agents that allow the communication of prices—details of this aspect of his work are detailed later in this chapter.

Zacharia et al. (2000) studies an agent-based service marketplace where both price and seller reputation are factored into a buyer's utility (as determined by a shopbot). Buyers evaluate sellers not only by price but by reputation. Also, their needs have varying levels of importance. They incorporate a utility function that leads to a willingness to spend more for (relatively) important problems but allows the tradeoffs of quality for price in the case of (relatively) unimportant needs. Zacharia et al. run simulations to examine the effectiveness of varying shopbot strategies with respect to varying market conditions. In this case, the market condition studied is the relative levels of supply vs. demand. Markets with less demand than supply are characterized as "unemployed," while markets where demand exceeds supply are called "overemployed." The initial set of market simulations focus on three seller strategies: Derivative Follower (DF), Reputation Follower (RF), and random. The DF strategy here is the same as discussed in multiple studies above. The RF strategy is built upon the DF algorithm. To implement RF, a shadow price is maintained on which the DF algorithm is applied. The price advertised by an RF seller, however, is that shadow price multiplied by the seller's reputation. Zacharia et al. create this strategy in hopes of enabling sellers to react to changing reputations faster than a DF strategy would allow. In the simulation, reputations were updated after each transaction between seller and buyer.

Zacharia et al. (2000) pit these strategies against one another via simulated marketplaces. In an underemployed market, the sellers who utilized RF agents fare best. However, in an overemployed market, the DF sellers are most profitable. In both cases, the strategies that make use of some sort of intelligence outpace random pricing. Zacharia et al. also examine market equilibrium with respect to varying seller strategy. They find that when seller reputation is allowed to vary in homogenous markets, DF markets reach an equilibrium, while RF markets do not. As a result, they develop a new, theoretically "socially optimal" strategy called Profit-Maximizing Reputation Follower with memory (PMRF). The PMRF market leads to an equilibrium consistent with the theoretic optimal. The Zacharia et al. work suggests that varying market conditions and varying seller strategies both impact the nature of markets. Further study in those areas is needed to better understand the dynamics at work. The research of Chapter Four examines the relative performance of various pricing strategies.

Comparison Shopping

For many decades, marketing researchers have studied consumer shopping habits. The question of what drives the depth, breadth, and length of a shopper's search (i.e., "involvement") is an interesting one. Early research attempts to tie consumer involvement to characteristics of the product or service sought. Intuitively, it seems reasonable that the price of an item directly impacts the extent to which comparison shopping is done. Through a survey of housewives, Bucklin (1966) finds that low-priced goods are more likely to be "one-stop" purchases, while high-priced goods are more likely to be "two-plus stops." Udell (1966) surveys buyers of small appliances and also finds a direct relationship between price and involvement. For inexpensive purchases, only 28% of buyers visited more than one store. However, for expensive purchases, 68% of consumers shopped in multiple stores. Additionally, stark contrast in same-store visits is found between cheap and expensive appliances. For inexpensive purchases, only 3% of buyers made multiple trips to the store of purchase, versus 57% for buyers of expensive goods. Newman and Staelin (1972) examine "information seeking behavior" in appliance and car buyers. They find that seeking increased with cost for appliance buyers. For car buyers, they find the same direct relationship between cost and information seeking for those shoppers who initially considered more than one brand. However, for car buyers who initially considered only one brand, information seeking decreases with cost. The researchers propose that the relationship does not hold for expensive cars because of the limited number of dealers available. Kiel and Layton

(1981) also study new car buying. They find that net price and cost of purchase relative to income are significantly, positively correlated with search time. Finally, Pustis and Srinivasan (1994) study the decision-making duration for car buying. They find that the greater the level of "stock adjustment" (the net change in value of the purchasers stable of cars), the longer the deliberation of purchasers.

While literature exists to support the notion that price is a driver of search behavior in markets, involvement literature also suggests that many individual factors can intervene. One significant driver of search behavior is the market knowledge of the consumer. Bucklin (1966) finds that shoppers who start their search with brand knowledge and a store preference shop half as much as others. Newman and Staelin (1972) find that appliance and car buyers who had bought at least two of those items in the last 10 years had a substantial influence on search behavior. Anderson et al. (1979) propose a model that relates information search, product satisfaction, attitude toward business, and experience. Their hypothesis that "the greater the experience, the less the search activity" is supported by data collected in a survey of consumers. Urbany et al. (1989) survey large appliance buyers about their level of uncertainty and their shopping experiences. They study two dimensions of uncertainty, choice and knowledge, and capture three search measures: shopping time, number of brands considered, and number of stores shopped. They find that both types of uncertainty are positively correlated with all three search measures. However, when they employ multivariate regression, choice uncertainty shows a positive effect on search, while knowledge has a negative impact. Finally, Urbany et al. examine mean search measures with respect to the four permutations of high/low choice/knowledge uncertainty. Low knowledge and high choice uncertainty lead to the most search (average number of stores shopped: 3.51) while the opposite combination leads to the least search (1.81 stores shopped). Later research indicates that the relationship between product familiarity and search may be more complex. Punj and Staelin (1983) propose a model to describe the drivers of information search. Their model hypothesizes that greater usable prior knowledge about the product sought would lead to decreased information search. They survey new car buyers to examine search behavior. Through regression analysis, they find a significant negative relationship between prior car buying knowledge and the extent of search.

The research described above was all focused on the buying habits of individual consumers making personal buying decisions. Some research indicates that product or service characteristics, such as cost, are a predictor of search behavior. However, other literature indicates that individual attributes, such as level of familiarity with (or uncertainty about) a product may confound the effects associated with the product sought. In corporate purchasing situations, these individual factors may be removed, or at least mitigated. A common approach to inventory management is the ABC classification system. Under this system, inventoried items are split into high, medium, and low value categories, based on their portion of total annual inventory value. Managers are taught to focus their efforts on the high-value items (value may come from individual item cost, volume, or some combination) at the expense of low-value items. Meredith and Shafer (2002) describe a case from a division of Johnson & Johnson where ABC was implemented and significant cost savings were realized. The ABC implementation included ordering from the lowest cost supplier and solicitation of competitive bids for high-value items. Due to this common inventory management practice, it is likely that some highly-commercial markets see higher levels of comparison shopping due to the relative value of the goods involved.

Another factor in the extent to which shoppers search is the cost (in both time and money) associated with that search. In the Punj and Staelin (1983) information search model referred to above, search cost is also a proposed determinant of the extent of information search. As with prior knowledge, search cost is hypothesized to have a negative impact on information search. Punj and Staelin's automobile buying survey data support this hypothesis as well. Moorthy et al. (1997) also study consumer search behavior in the automobile buying context. They survey people who recently purchased cars as well as those who were in the process of buying. They also make a distinction in their data analysis between shoppers who performed a "directed" search versus a "random" search, based on brand preferences exhibited by survey respondents. For all regressions performed on the different data sets, search cost is found to be negatively related to amount of search.

The Internet and its burgeoning number of e-marketplaces have decreased search cost in relation to traditional brick-and-mortar markets. Bakos (1997) states that electronic marketplaces reduce the incremental cost of obtaining price information from additional sellers. Additionally, he notes that e-markets make it more difficult for sellers to obscure their prices (by including or excluding certain fees or promotions). He claims that electronic markets move commodity markets closer to the "classical ideal of a Walrasian auctioneer where buyers are costlessly and fully informed about seller prices." Use of agents within an Internet marketplace can further reduce search costs, especially in terms of time. One would expect such decreases to lead to increased use of shopbot services as e-markets mature.

Pedersen (2000) conducts a laboratory study to examine differences between conventional Internet shopping (i.e., “manual” browsing) and agent-enabled shopping. Subjects in the experiment are given a task that involved choosing a bank for multiple financial needs. A control group of participants is given access to a Web page that listed financial-service providers; the experimental group is shown the same page, but with an additional link to a shopbot service. A survey was used to measure characteristics of the subjects’ shopping experience. Consistent with Pedersen’s hypothesis, shopbot use leads to significant increases in the number of information sources used, the amount of information obtained, and the satisfaction of the users. These findings support the notion that shopbot use will increase as e-markets mature.

Kephart and Greenwald (2002) note that “shopbots outperform and out-inform humans by providing extensive product coverage in just a few seconds, far more than a patient, determined human shopper could achieve after hours of manual search.” It is important to note that Internet and agent technology, while dramatically reducing search time and opportunity cost may potentially introduce new costs to shopping. Advanced intelligent agents could demand a price for the service they provide. Kephart and Greenwald simulate agent-based markets and examine price behavior when various assumptions about those search service costs are made.

Because search cost is a driver of consumer involvement and the Internet tends to reduce that cost, it is reasonable to assume that as a market’s online presence grows, its complexion with respect to comparison shopping changes. Evolution of electronic marketplaces should lead to increased comparison shopping. In the simulations done by Kephart and Greenwald (2002), comparison shopping complexion of the market was varied. They represented comparison shopping complexion as a vector of three percentages representing the portion of the population that shopped at one, two, and all sellers of a market. They show that markets are very sensitive to that buyer strategy vector. For both the five and 20 seller cases, the way sellers choose prices for a market in equilibrium is dramatically dependent on the buyer strategy mix. Additionally, Kephart and Greenwald propose that there is a lower limit to the percentage of one-seller shoppers in a market—i.e., there will always be some non-comparison shoppers. They demonstrate that the value of that lower limit has an effect on how equilibrium is achieved in a market.

Clark (2000) notes the results of a Jupiter Communications survey about e-commerce buying behavior. Those results indicate that:

- 22% of shoppers survey only one website for a purchase
- 25% survey two sites
- 46% survey three to five sites, and
- 7% survey six or more sites

These figures are not specific to any particular e-market. Whatever the determinants of search behavior may be, it is likely that there is some variability in levels of comparison shopping across various markets (or within a market over time). The research of Chapter Four examines the impact of differing levels of comparison shopping on the choice of pricing strategy.

Buying Strategies

In the "Seller Strategies" section above, many dynamic pricing methods were described. In the majority of those studies, buyer behavior is very simple and is held constant. For most of the studies the buyers behavior is comparable to that described by Varian (1980): buyers have a valuation that indicates their threshold price for buying and, with respect to market-wide prices, some buyers are informed while others are uninformed. A few studies vary the relative proportions of informed/uninformed buyers, but their results and conclusions are limited, as argued above. In most of the studies described, buyers do not persist in the market—they appear, make a buy/no buy decision and depart. Only in the work of Kutschinski et al. (2003) do the buyers persist in market. While the characteristics of those buyers vary over time, they have no means by which to attempt to learn about or adapt to their environment. In their study of pricing strategies, Dimicco et al. (2003) note that investigation of buyer strategy is a logical future direction for e-commerce research. They argue that buyers are likely to develop shopbots that can decipher the dynamics at work in a market in order to realize savings for their users. They suggest that a simulation approach is appropriate for examination of the potential impact of introducing more intelligent buyer agents into dynamic pricing markets. The research of Chapter Five examines methods by which buyers' agents might be able to employ some sort of intelligence in order to better perform for their employers.

Tassier et al. (2002) examine the effect of buyer behavior on pricing patterns in a durable goods market. They note that, traditionally, in studies of markets and economic models, consumers are assumed to behave statically. That is, a consumer's reaction to a price observed in the market at one time will be identical to the reaction at any other time. They argue, however, that this assumption may not be representative of the real world.

Therefore, the focus of the Tassier et al. study is on the impact of introducing memory for buyers modeled as rational agents. They suggest that a buyer in a durable goods market might delay a purchase if the currently observed price is greater than all of the prior prices in the buyer's memory. Similarly if the currently observed price is less than all of prices in memory, the buyer might choose to buy earlier than demand would otherwise dictate. Tassier et al. simulate an automobile market populated by a single monopolistic producer/seller and 500 buyers. The market contains both new and used cars (when a buyer buys a car, the previously-owned car enters the used car market). The monopolistic seller sets prices and production levels to maximize profits, given the market characteristics. Buyers only own one vehicle at a time and make decisions according to a utility maximization function. Tassier et al. run three sets of simulations: one in which buyers had no memory, one in which buyers can use a postpone buying strategy, and one in which buyers can use an early buy strategy. A market of buyers without memory leads to stable pricing over time. The addition of memory introduces price oscillations—periods of high prices and low sales followed by lower prices with higher sales and vice-versa. Postpone and early buy methods lead to qualitatively identical price and sales behavior in the market.

Hertweck et al. (2003) study a simulated market where buyers are memory-enabled as in the Tassier (2002) study. They examine a market where five sellers all employ a DF pricing strategy as described earlier in this chapter. Buyers in this study persist in the market until a purchase is made. Buyer valuations are drawn from a uniform distribution and increased over time according to a discount rate. Buyer memory is implemented such that upon entering the market, a buyer sets a threshold value equal to the minimum price observed in the buyer's memory window. That threshold is then aged in the same manner as the valuation. In order to make a purchase, the lowest price observed in the market has to be lower than both the current valuation and threshold values. They examine varying implementations of buyer memory—a base case with no memory and four memory-enabled implementations with varying discount rate values as seen in Table 2-1:

		Valuation	
		Low	**High**
Threshold	**Low**	*Buyers who seek bargains and are able to wait to purchase* ***(Bargain Hunter)***	*Buyer who wants to find a great deal but has an immediate need* ***(Conflicted Buyer)***
	High	*Buyer unconcerned about finding the great deal and can wait to purchase* ***(Casual Buyer)***	*Buyer with an immediate need and little concern for finding a bargain* ***(Impulse Buyer)***

Table 2-1: Hertweck et al. (2003) Buyer Categories.

Hertweck et al. conclude that the manner in which shopbots are implemented can have an impact on the way pricing in a market evolves and how profitable the sellers in that market are. Between the most extreme scenarios simulated, they note a 2.2% change in seller profit margin. They conclude that this area of research merits further study and should be of interest to sellers, buyers, and pricebot makers.

Another manner by which buyers (or their agents) could capitalize on market information to make smarter decisions is by communication with other buyers (or their agents). A shopbot may enter a market without knowledge of or visibility to historical prices. The shopbot could conceivably communicate with other shopbots to learn what prices have been recently paid for the good in question. Leloup (2003) discusses and studies neighborhood communication of agents. He suggests that, in a real, physical world scenario, buyers communicate their purchase experience with neighbors. If one shopper knows that a neighbor received a price that is lower than what is currently available, that shopper may delay his purchase, even if the current price is lower than her valuation. It is possible that this could occur in an online, agent-populated world as well. Leloup uses a two-dimensional matrix representation of the world of agents and shows various ways to represent neighborhoods within that world. He chooses to study a Moore neighborhood, which consists of one cell and the eight that immediately surround it. Buying agents strictly reject all prices higher than a price that has been charged in the past to a member of their neighborhood. As described earlier in this chapter, the market in question has only one seller who attempts to learn optimal pricing policy. He finds that the neighborhood communication model leads to prices that are below the initial optimal price. Chapter Five will address more sophisticated models of markets where shopbots have knowledge of past prices, seeking to quantify the value of that knowledge to the buying community.

Neural Networks

Most of the shopbot and pricebot techniques described above are relatively simple. They involve algorithms that could be performed relatively easily by humans—software agents merely lessen the human time and effort required. The buyer strategies described above are attempts to mimic human behavior—specifically memory and communication. Chapter Six of this work attempts to examine a more "intelligent" implementation of agent technology—an attempt to duplicate the way the human brain works. Specifically, the potential of a neural network-enabled shopbot is examined.

The greater the extent to which an agent can display intelligent behavior, the greater value it has for its employer. A dictionary (in this case, www.dictionary.com) can provide an indication of the generic, traditional meaning and usage for a term. The dictionary definition for "intelligent" includes "showing sound judgment and rationality." For "intelligence," the definition includes "the capacity to acquire and apply knowledge" and "the faculty of thought and reason." Combining those ideas, it would seem that anything deemed "intelligent" must be capable of taking in information and processing it in such a way that some sort of reasoned choice is made.

Elofson et al. (1997) assert that intelligent agents can be described by three characteristics: agency, intelligence, and mobility. Agency refers to the degree of authority and autonomy of the agent. Mobility refers to the extent to which an agent traverses a network. Intelligence is the degree to which an agent learns and reasons. They describe intelligence as a property existing on a continuum. At a minimum, an agent must be able to act on stated preferences through an inference engine or other reasoning mechanism. Higher intelligence is embodied in the use of user models or other more complex representations of what a user wants done. Further along an intelligence continuum is the capability to learn and adapt based on the user's objectives and the resources available to the agent.

Neural networks (NNs) are a biological approach to artificial intelligence. NNs are designed to mimic the way neurons of the human brain work. Typically, neurons of a NN combine a set of weighted inputs and apply a transfer function to create an output. Outputs from one neuron can become inputs of another, creating the network. The network learns through a training process. Input sets with known correct outputs are fed into the network and weights are adjusted according to how the network's output compares to the correct values. After the network is properly trained, it is able to generalize what it has "learned" to data beyond the training set. NNs have powerful capabilities to recognize patterns that other methods can not. Therefore, classification problems such as character recognition, pattern recognition in images, text-to-speech conversion, and natural language processing are popular choices for NN application. NNs can also be used to solve combinatorial problems and perform signal processing. NNs have also been employed as financial and economic models for prediction. The application of greatest interest for this research is classification, specifically economic or financial classification problems. Neural network classification is a very mature research area, therefore only a few of the most relevant works are described below.

Zhang et al. (1998) list bankruptcies, business failures, and bond ratings among the classification-type business problems that can be tackled with NNs. Piramuthu et al. (1994) add credit/loan evaluation, consumer choice, and tax planning to the list of business classification problems that might be addressed with NNs. Piramuthu et al. outline the advantages that NNs have over other classification problem approaches. While parametric statistical methods require assumptions about probability distributions, NNs do not. Statistical models also become problematic when multi-collinearity and autocorrelation are present in data; NNs are more robust with respect to that type of data. Statistical models often require the specification of the nature of relationships between variables (linear, quadratic, etc.), where NNs do not.

Piramuthu et al. (1994) go on to implement a NN to classify creditworthiness of potential borrowers. They utilize a three-layer NN, with 14 input nodes, 10 hidden layers nodes, and five output nodes (corresponding to five potential classes). The focus of their study was on the relative performance of various training algorithms with respect to convergence during training of the NN. They compared a Newton-Raphson training algorithm, a traditional backpropagation, steepest gradient algorithm, and a hybrid of those two. Piramuthu et al. concluded that the two non-traditional training approaches improved both convergence during training and classification accuracy. The difference in speed of convergence was dramatic. The difference in classification accuracy was much more subtle. Additionally, all three approaches led to classification accuracies that were as low as 44% and only as high as 68%. They also compared NN classification performance (across all three training methods) to three inductive learning techniques. They concluded that the NN techniques were better classifiers of creditworthiness.

Agarwal et al. (2001) studied a NN implemented in a decision support system context. Their system was used to classify firms according to their financial health. They note that many classification problems studied in the financial arena have been binary, while they tackle the task of four-way classification. Additionally, the classes studied were ordinal, where classes are merely nominal in many other similar classification studies. They compare a NN approach to three statistically-based approaches: multi-discriminant analysis (MDA), ordinal logistic regression (OLGR), and a “naïve” approach, which merely classifies all firms as the class that was the highest proportion of the training class. They also compared training with balanced samples to training with unbalanced samples. Balance here refers to the relative proportions of each class observed in the training data set.

When measuring classification accuracy, Agarwal et al. (2001) used four different measures: simple classification accuracy (SCA), distance-weighted classification accuracy (DWCA), expected misclassification cost (ECM), and ranked probability score (RPS). SCA simply scores each classification as right or wrong. Since the classification scheme is ordinal, DWCA can assess different penalties to varying levels of misclassification (distance between classes was assumed equal). ECM also assesses varying penalties for misclassification, but the direction of the misclassification matters and the distances are not equal as they are with DWCA. RPS evaluates the aggregate accuracy of state-probabilities. When examining performance in classifying the holdout (non-training) data set, Agarwal et al. conclude that the NN model clearly outperforms all others using the ECM criteria, for both balanced and unbalanced training sets. When balanced training was done, the NN outperformed both MDA and OLGR using the SCA measure. For all other scenarios, the performance of the NN was not significantly better than the other techniques, but it was also never outperformed. Neural network SCA percentages were 71% for unbalanced training and 63.9% for balanced training.

Partovi and Anandarajan (2002) studied a slightly different kind of classification—business-oriented, but not directly financial. They studied the problem of ABC inventory classification. ABC is a system for partitioning inventory into different levels of importance so that planning and control efforts are focused on the more important stocked items. Partovi and Anandarajan studied the performance of two different NNs against the performance of MDA. The two NNs differed in their training methods—one used traditional backpropagation (BP) while the other utilized a genetic algorithm (GA). They implemented a three-layer NN, with four input nodes, 16 hidden nodes, and three output nodes (for A, B, and C inventory classes). Partovi and Anandarajan used SCA to compare performance of the two NN models to MDA. They found that both the BP and GA networks significantly outperformed MDA for A items and B items. For C items, MDA actually had a higher SCA than both NNs, but the difference was not significant. The GA-trained NN outperformed the BP-trained NN in all categories, but no significant differences were reported. The GA and BP networks had overall SCA levels of 80.3% and 75.0% respectively.

Nath et al. (1997) used NNs to address another business classification problem—discriminating between firms that were acquired and those that were liquidated. The focus of their study was to evaluate a technique for evaluating the saliency of input variables for a NN. They implemented a backpropagation NN with 12 input nodes, four levels (9, 12, 18, and 25) of hidden nodes, and one output node. They used a saliency measure to evaluate the 12 input variables and came up with a reduced model consisting of the five most salient variables. Classification rates for full and reduced NN models of all hidden layer sizes were compared along with Fisher's linear discriminant analysis (FDLA). FDLA was also performed with both the full and reduced input variable sets. All classification percentages were between 60.0% and 65.7%. No significant differences were found in classification percentages, thus, Nath et al. concluded that their saliency measure was an effective way to simplify classification NNs without sacrificing accuracy.

Saad et al. (1998) compare the performance of time delay, recurrent, and probabilistic neural networks for classification of stock trends. Saad et al. do not focus strictly on predicting next values in a time series, but rather making predictions that would be helpful in implementing a trading strategy. They design NNs to predict short-term upward stock trends (2% increase over 22 days), which represent buying opportunities. They measure NN performance by counting false alarms—i.e., inaccurate prediction of short-term increases. Saad et al. make a number of comparisons among the NN implementations, using multiple measures. They conclude that all of the methods are capable performers, but selection of the "right" one involves tradeoffs. Probabilistic NNs have the advantage of implementation simplicity, but have disadvantages in higher requirements for storage requirements and testing time. Recurrent NNs have the advantage of incorporating past experience and requiring relatively little memory storage, but do have significant implementation complexities. TDNN is a compromise strategy, described as moderate in implementation complexity and memory requirements.

Berardi et al. (2004) propose a framework for building and evaluating NN classification models. They highlight the potential for NNs to be effective within the e-commerce domain. As other authors described above have done, Berardi et al. point out several advantages of NNs that make them especially appropriate for e-commerce applications. They note that NNs make no a priori assumptions about the data—the breadth of data available in e-commerce arenas limits the applicability of fixed-form models. They also discuss the fact that NNs are consistent estimators, which is a valuable property in any application. Finally, they cite the strength of neural networks' ability to deal with "noisy" data that might challenge traditional approaches. The beginning portion of the conclusions of Berardi et al. highlight the relevance and importance of the NN work undertaken in the final portion of this dissertation:

"The neural network is a promising modeling tool for e-commerce applications. Electronic commerce data is typified by its complexity and high level of noise. Neural networks are ideally suited for these problem characteristics and can be incorporated into intelligent agents and other decision support system components." (p. 245)

Chapter Three: Methodology

Chapter Four Methodology

The Chapter Four research consists of two rounds of simulation. All simulations were done with a custom-designed Excel/VBA simulation tool. The first round is designed to determine what dynamic selling strategy performs best when a seller is the first in a market to employ a pricebot approach. The goal of the first round is twofold: (1) to determine whether the level of comparison shopping (specifically the extent to which shopbots are used by buyers) matters in selection of pricebot strategy, and (2) to determine the best initial pricebot strategies to use when entering a market where no dynamic pricing strategies are in use. The second round focuses on those best strategies by examining how markets would evolve if all competitors convert to the best strategy. It also examines what the best reaction to the best strategy should be and includes some investigation into the effects of varying parameterization of given strategies.

The simulations make use of the following buyer types:

- Buyers who shop at only 1 seller and buy if price beats valuation (i.e., the price at or below which a buyer will purchase a good). This type represents buyers who have favorite sellers and blindly buy there, without benefit of agent services.
- Buyers who shop at two sellers and buy if lower price beats valuation. This type represents buyers who manually comparison shop between two favorite or well-known sellers, without benefit of agent services.
- Buyers who shop at all 5 sellers and buy if lowest price beats valuation. This type represents buyers who avail themselves of shopbot services to shop the entire market.

Unlike most of the shopbot simulations in the literature, buyers persist in this simulated market. They behave as if they have decided definitively to purchase and have to do so by some deadline. Accordingly, they arrive with a valuation and a time horizon, shop the market, and make a buy/no buy decision. Agent-enabled buyers shop every time period. For non-agent buyers, the decision on whether to shop in a given time period is a parameterized probabilistic decision. If buyers do not buy, their valuation is increased according to a discount rate, and the same shopping/buying process occurs in the next period. This cycle repeats until their time horizon is reached, at which point they buy from the lowest-priced seller, regardless of valuation.

Round one considers two factors: shopbot usage and pricebot strategy. The shopbot usage factor is represented by a percentage of buyers who use agents to shop the entire market. The remaining percentage of buyers is split equally between one and two seller shoppers. Four levels of shopbot usage are examined in this first round of simulation, constructed as follows:

- 20% agent use (40% one seller, 40% two seller)
- 40% agent use (30% one seller, 30% two seller)
- 60% agent use (20% one seller, 20% two seller)
- 80% agent use (10% one seller, 10% two seller)

Using a mix of 1, 2, and 5 seller shoppers is in line with the studies of both Kephart (2000) and Deck and Wilson (2003). Kephart examined a five seller market where 50% of buyers shopped all five sellers, 25% shopped two, and the remaining 25% shopped only one. Deck and Wilson examined a four-seller market where only 20% of buyers shopped all four sellers, 20% shopped two, and 60% shopped only one. However, both studies held buyer mix constant throughout experimentation.

The pricebot dynamic pricing strategy factor consists of five treatments. The first four are techniques that have been studied in previous research, while the final is a variation on a previously studied strategy. See Chapter Two for descriptions of prior work. The five treatments are:

- *Varian:* It would be worthwhile to have a baseline study for each market where all five sellers utilize the randomized Varian (1980) strategy. This treatment is designed to model a market where no strategic pricebot use is occurring.
- *Derivative Follower (DF):* A strategy where incremental increases/decreases are made; if an increase/decrease leads to increased profits over the prior period, another change in the same direction is made. If profits decrease, the direction of change is reversed. This strategy is easily automated by an agent, but requires only information about the seller performing the pricing—no competitor or buyer information is necessary.
- *Low-Price Match (LPM):* A strategy where the next period's price is set equal to the low observed in the prior period. This technique assumes that competitor prices are readily available to a pricebot.
- *Trigger:* A hybrid strategy where one strategy is followed until a trigger event is observed—then a secondary strategy is employed. For this study, derivative follower with a trigger to low-price match when 2 competitors are observed to have lower prices will be used. This technique also assumes perfect knowledge of competitor prices.
- *Beat Half the Market (BHM):* A strategy where the price is set such that it beats two of the four competitors, trading off sales to those who shop the whole market in exchange for extracting a higher price from the one-seller shoppers and some of the two-seller shoppers. BHM also requires perfect knowledge of competitor prices.

For each strategy treatment, the pertinent strategy is utilized by one seller (referred to below as the "fifth" seller) while the other four have prices set according to Varian's (1980) model of sales. Varian sellers re-price asynchronously, with a probabilistic decision made each period based on a re-pricing parameter. For the other strategies, assumed to be implemented by a pricebot, re-pricing occurs each period.

From an ANOVA perspective, the study is a 4x5 full factorial where each cell represents the combination of a shopbot usage level and a pricing strategy. Table 3-1 is a depiction of that research design. The issue of the appropriate number of iterations per cell will be addressed with the results in Chapter Four.

Round One		**Pricing Strategy (vs. 4 Varians)**				
		Varian	**DF**	**LPM**	**Trigger**	**BHM**
Shopbot Usage	**20%**	*n replications*	*n replications*	*n replications*	*n replications*	*n replications*
	40%	*n replications*	*n replications*	*n replications*	*n replications*	*n replications*
	60%	*n replications*	*n replications*	*n replications*	*n replications*	*n replications*
	80%	*n replications*	*n replications*	*n replications*	*n replications*	*n replications*

Table 3-1: Research design for Chapter Four

ANOVA will be performed on both the profit and sales of the fifth seller. The multi-factor, multi-level research design lends itself to this type of test. ANOVA is a more powerful test than similar non-parametric tests when the assumptions of normally distributed, uncorrelated errors with heterogeneity of variance are met. The model assumed is:

$$Y_{ijk} = \mu_{...} + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \varepsilon_{ijk}$$

where:

- Y_{ijk} is the profit or sales (over a length of time to be determined) of the fifth seller when simulated with the ith random number seed, the jth pricing strategy, and within the kth shopbot usage level
- $\mu_{...}$ is the grand mean
- α_j is the effect of the jth pricing strategy, $\Sigma\alpha_j = 0$
- β_k is the effect of the kth shopbot usage level, $\Sigma\beta_k = 0$
- $(\alpha\beta)_{jk}$ is the interaction effect of the jth pricing strategy and kth shopbot usage level, $\Sigma_k(\alpha\beta)_{jk} = 0$ for all k and $\Sigma_j(\alpha\beta)_{jk}$ for all j
- ε_{ijk} is random variation, assumed to be $N(0, \sigma^2)$
- $i = 1 \ldots n$; $j = 1 \ldots 5$; $k = 1 \ldots 4$

Three hypotheses are of concern:

- H_{10}: all $\alpha_j = 0$
  H_{1a}: not all α_j equal zero
- H_{20}: all $\beta_k = 0$
  H_{2a}: not all β_k equal zero
- H_{30}: all $(\alpha\beta)_{jk} = 0$
  H_{3a}: not all $(\alpha\beta)_{jk}$ equal zero

The first hypothesis tests whether there is a pricing strategy effect. The second tests whether there is a shopbot usage effect. The third tests whether there is interaction between the strategy and shopbot usage factors.

Assuming the interaction null hypothesis is rejected, mean values from each row will be examined to determine if the most lucrative strategy changes depending on shopbot usage. An examination of performance on a per-replication basis will also be done. Within each shopbot usage level, a check will be performed to determine which strategy produced the greatest revenue for each replication in order to verify that the best strategy is consistent across replications.

If the interaction is not significant, post hoc multiple comparison tests will be performed to determine where the differences in treatment means occur. Assuming null hypotheses one or two are rejected, Fisher's protected Least Significant Difference (LSD) will be used to examine all pairwise comparisons and identify differences. If either of those hypotheses is not rejected, Tukey's Honestly Significant Difference (HSD) can be utilized for pairwise comparisons. Fisher's LSD is more powerful than Tukey's HSD and provided the ANOVA results are significant, Fisher's LSD controls experiment-wide error properly. The conclusions from the pricing strategy comparisons should help to answer the following question: What pricing strategy should I implement if I do not know what to assume about shopbot usage in my market? The conclusions from the shopbot usage comparisons should help to answer the following question: At what level of agent use does the nature of the market change with respect to pricebot strategy decisions?

The second round of simulations examines what happens when a market shifts to the seller strategy that was most successful in the first round. For the low price match strategy, that combination does not make sense, so other parameters could be studied, such as an undercut percentage/amount or varying the frequency of re-pricing across sellers (which would be an interesting question to study for any of the strategies). Another part of the follow-on to round one would be examining which strategy is the best response when the entire market moves to the best as defined by the first round of simulations. This second round of simulations could also be an opportunity to study the effect of having a market dominator who is more likely to be the preferred seller. Whether the best strategy for the dominator is the same as for the other sellers is an interesting question to study.

Chapter Five Methodology

The research of Chapter Five will make use of the same basic simulation tool as was used in Chapter Four. However, the entity of interest shifts from seller to buyer. The primary goal of the Chapter Five work is to understand how various properties and parameterizations of the buyer's shopbot affect its performance. The buyers in Chapter Four were not sensitive to the market beyond polling prices and comparing them to an independently chosen valuation that increased over time. In Chapter Five, variable discount rates for the valuation as well as the addition of a memory-like component will be introduced to the shopbot buyers. Various shopbot configurations will be studied to determine their relative worth. To accomplish this goal, the primary measure of interest for Chapter Five will be price paid by the buyer. A secondary goal will be to examine the effects of those buyer changes on the sellers and the overall market pricing dynamics—portions of Chapter Four will be revisited to determine whether changing buyer behavior affected sellers' results.

To begin the Chapter Five work, the first round study from Chapter Four will be repeated using a mix of buyer types instead of the homogeneous, simple buyer market from Chapter Four. This heterogeneous market will include a mix of buyer types representing different levels of valuation discount rate as well as differing capabilities with respect to the ability to see and remember past prices in the market. The Chapter Four, Round One ANOVA will be re-done with the new, heterogeneous buyer market to determine any effects that the buyer complexion might have on seller pricing strategy decisions.

To examine the relative performance of differing shopbot implementations, select markets from Chapters Four and Five will be studied further. Here, a "market" is defined by a combination of shopbot usage level and fifth seller strategy—i.e., one cell in the Table 3-1. At each shopbot usage level, the market with the best-peforming fifth seller strategy will be studied—if Chapter Four and Five studies do not agree on the best-performing strategy, both markets will be studied.

For each market selected, buyers of each shopbot implementation type studied will be selected at random. The average price paid will be used to compare performance across shopbot types. As described above for the number of simulation iterations needed for the Chapter Four work, the number of buyers to be examined needs to be determined for Chapter Five. For data analysis, one-way ANOVA will be performed with buyer behavior as the treatment of interest. The treatments will mirror the work done by Hertweck et al. (2003). The baseline case will be taken from Chapter Four, while the other four treatment levels will be as described in Table 2-1: Bargain Hunter, Casual Buyer, Conflicted Buyer, and Impulse Buyer. The model assumed is:

$$Y_{ij} = \mu_{..} + \rho_i + \alpha_j$$

where:

- Y_{ij} is the price paid by the ith buyer using the jth shopbot implementation
- $\mu_{..}$ is the grand mean
- ρ_i represents the variability across buyers (due to different conditions at market entry), ρ_i are independent $N(0, \sigma_\rho^2)$
- α_j is the effect of the jth buyer behavior implementation, $\Sigma\alpha_j = 0$
- $i = 1 \ldots n$; $j = 1 \ldots 5$

The hypothesis of primary concern is:

H_0: all $\alpha_j = 0$
H_a: not all α_j equal zero

This hypothesis tests whether there is a buyer behavior implementation effect. Similar to the approach for the Chapter Four data analysis, Fisher's protected LSD will be used to examine all pairwise comparisons and identify differences if the null hypothesis is rejected. Relative differences between treatments (for the regret measures) can give an indication of the worth of intelligent agent technology to buyers. That knowledge is important to potential shopbot consumers as well as providers.

Chapter Six Methodology

The focus of the Chapter Six research returns to the seller. From the perspective of seller pricing strategies, the work of Chapter Four is focused on a very narrow range of markets. The focus of Chapter Four is on how the level of shopbot use in a market interacts with seller strategy choices. For Chapter Six, the shopbot usage level becomes fixed (at 50%) and the focus shifts to sellers' pricing strategy choices across a much broader range (as defined by the mix of pricing strategies in use) of markets. Chapter Four assumed all competitors of the seller of interest were not yet using pricebots—i.e., their pricing was performed using the Varian (1980) model. Chapter Six allows for a broad range of competitor selling strategies—competitors can use any of the strategies studied in Chapter Four. Chapter Six utilizes a Neural Network (NN) to determine the pricing strategies being used by competitors in a simulated market. The competitor strategy mix determined by the neural network is then used as an input to a decision matrix that can advise a seller what the best response pricing strategy is for the observed competition. The knowledge stored within the decision matrix is an output of simulations and data analysis comparable to what was done in Chapter Four (without the complications of varying shopbot usage but performed across a broader range of seller strategy mixtures). Once the best response strategy (in terms of average profit level against a given mix of pricing strategies) is determined, the strategy of the seller of interest is changed to that best response (competitors are left unchanged) and the simulation is continued. Profit levels after the strategy change can then be compared to those before the change to determine the effectiveness of the NN.

The decision matrix consists of a set of rules for choosing the best performing pricing strategy based on the strategies that competitors are using. Simulations (using the same tool as was used in Chapter Four) will be used to determine the best response pricing strategy for each of the competitor strategy mixes studied. Details of this analysis and the rules generated can be found in Chapter Six.

After the NN and decision matrix are constructed, the final step will be to generate markets and compare their profitability levels for the time periods before and after the recommended strategy change. To accomplish this, seller strategy mixes (for one seller of interest and four competitors) will be randomly generated. Next, the market will be simulated for 30 periods. End conditions for those simulations will be recorded in order to serve as initial conditions for subsequent simulations. Also, the historical price data will be fed through the NN system in order to classify the pricing strategy of each of the four competitors. That classification will then be fed into the decision matrix in order to recommend the most profitable pricing strategy for the market as described by the NN output. Additionally, the known, correct set of competitor strategies will serve as in input to the decision matrix to determine the recommended strategy assuming a perfectly accurate NN. Next, three parallel 30-period simulations will be run, using the previously-recorded end conditions as initial conditions: (1) the seller of interest will continue to use that same strategy as was used in the initial 30 period run, (2) the seller of interest switches to the strategy recommended by the NN/decision matrix, and (3) the seller of interest switches to the strategy that would be recommended by a theoretically perfect NN that feeds the decision matrix. The profitability levels for each of the three conditions will be recorded. Paired t-tests will be used to compare seller profitability realized using the actual and perfect NN systems to profitability assuming the seller does not change pricing strategy.

Chapter Four: Choosing a Dynamic Pricing Strategy

Chapter Two describes a multitude of studies that involve the use of pricebots and shopbots in e-markets. However, the practical question of what strategy a manager responsible for pricing in a selling organization should employ has received little attention. Two Deck and Wilson studies (2002 and 2003) examined the relative financial performance of multiple pricing strategies. However, in both cases, the results are confounded by the fact that human subjects were making decisions in their simulations. The variability in the human behavior meant that strategies were not necessarily implemented equivalently and that the competition for each strategy may not have been comparable. DiMicco et al. (2003) introduced a computer simulation tool for the study of pricing strategies. They study the relative performance of pricing strategies, but in a perishable goods market. This chapter focuses on durable goods. They also only studied two strategies, one of which was very specific to the perishable good scenario. Kephart et al. (2000) studied multiple strategies and evaluated relative performance, but two of the three strategies studied required perfect knowledge of the buyer population, which is not a practical assumption. Kutschinski et al. (2003) performed a series of robust simulations with multiple strategies, but did not make a quantitative assessment of the relative merit of those strategies for sellers. Leloup (2003) and Dasgupta and Smith (2003) studied only one strategy. Therefore, the question of what strategy is best remains largely unanswered.

Two of the studies mentioned in Chapter Two allowed for variability in comparison shopping complexion. DiMicco et al. (2003) studied comparison shopping levels of 0% and 100%, both of which seem too extreme to be practical. In the portion of the Kutschinksi et al. (2003) study in which comparison shopping levels were allowed to vary, only one pricing strategy was examined. In another study, Kephart and Greenwald (2002) study the effects of varying comparison shopping behavior on equilibria attained in e-markets. However, they do not utilize any sort of active pricing strategy in their simulations; they merely allow the sellers to act such that equilibrium is achieved. The research of this chapter seeks to better understand how varying the comparison shopping complexion of a market (in this case, represented by the extent of shopbot use by buyers) impacts the selection of dynamic pricing strategies for sellers' pricebots.

A two factor, full factorial study was done using an Excel/VBA simulation as described in Chapter Three. Due to the ease of generating data and the difficulty inherent in accounting for interaction effects when planning sample size a priori, a cumulative sampling approach was used. As additional simulation runs were performed, ANOVA was performed iteratively on cumulative data. Changes in p-values were tracked to determine the existence of stable sampling state and effects that are significant and of practical importance.

Round One Results

In the first round of simulations, the base set of parameters described in Table 4-1 was used.

Parameter description	Value/Distribution
Top end of price range	2100
Bottom end of price range	1900
Number of sellers	5
Number of buyers	1000
Number of time periods per iteration	365 (days/year)
Percentage of buyers who shop at all 5 sellers. Remaining percentage is split equally between 1- and 2-seller buyers.	Set by iteration to 0.2, 0.4, 0.6, or 0.8
Valuation discount rate	0.01
Rate of repricing for Varian sellers.	0.0667 (re-pricing an average of twice per month)
Rate of shopping for non-agent buyers.	0.5 (shopping an average of once every two days)
Unit cost to seller. Set equal to bottom end of price range.	1900
Initial buyer valuation	Normal with ±2σ in price range
Buyer horizon	Normal centered on 7.5, with ±2σ going from 0 to 15 and truncated at 1.

Table 4-1: Base parameterization for round one simulations

One hundred iterations per treatment (two thousand total iterations) were performed. For the base case, ANOVA and other analyses were performed on per iteration total profits for 100 iterations per treatment. Both primary factors and the interaction were clearly significant (all p-values were less than 10^{-37}). The significance of the interaction is the most compelling of the ANOVA results. The presence of an interaction means that the relative performance of the pricebot strategies is dependent on the extent to which shopbots are employed in the market. That is, pricing managers should be aware of the comparison shopping complexion of their market when making pricebot decisions. The multiple intersections of lines in Figure 4-1 clearly illustrate the interaction between pricebot strategy and shopbot usage. The Trigger and LPM strategies have very similar lines, with rapidly increasing profits as the level of shopbot usage increases. Since LPM strives to beat all competitors, it is logical that it performs better at higher comparison shopping levels. Trigger makes use of an LPM element, as it uses a DF approach as its primary technique, but switches to LPM when two competitors are observed to have lower prices. It is interesting to note that this approach generally outperforms the pure LPM approach. The DF strategy shows a very slight increase as shopbot usage increases. Beat Half the Market shows a dramatic decrease in profitability from the lowest shopbot usage level to the highest. The Varian model becomes slightly less profitable as shopbot usage increases.

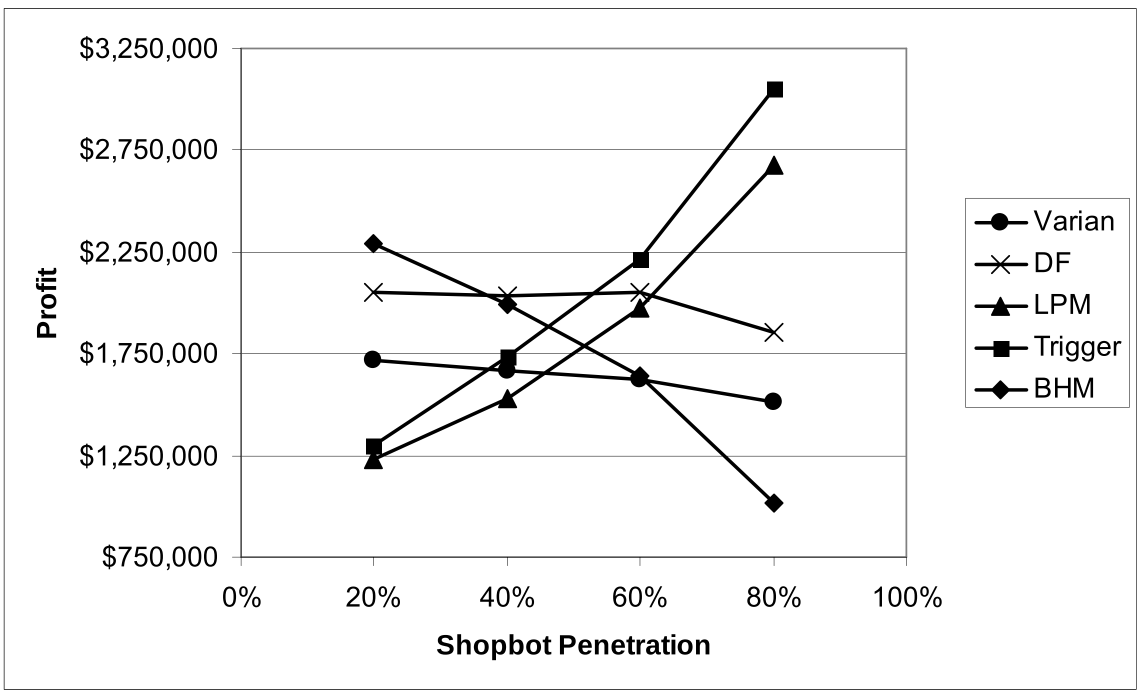

Figure 4-1: Profitability interaction between shopbot usage and pricebot strategy (100 iterations per treatment)

Figure 4-1 clearly shows why pricing managers need to understand the comparison shopping complexion of their markets. For instance, consider a manager who chooses to implement a BHM strategy. It might turn out to be the most profitable option in a market where 20% of buyers use shopbots, but the least profitable in a market with 80% shopbot usage. Of course, situations may exist where it is difficult to reliably estimate the behavior of consumers in a particular market. For that situation, multiple comparisons between strategies were made across all levels of shopbot usage. Fisher's Least Significant Difference (LSD) procedure was used to examine strategy differences across all shopbot usage levels. Table 4-2 shows average profitability levels along with group designations. The LSD for 100 iterations is $59,640, which is 3.2% of the grand mean profitability—a variation in profits that most managers would likely see as practically important in addition to statistically significant (at the $\alpha = 5\%$ level). The 100 iteration level was chosen as the appropriate level of analysis for the remainder of scenarios studied in this chapter. In reporting groups, different letters indicate significant differences. Any time two strategies (or other objects of study) are part of the same-lettered group, the Fisher's test indicated no significant difference between them.

Strategy	Average Profit (Group)
Trigger	$2,076,657 (A)
Derivative Follower	$1,997,368 (B)
Low Price Match	$1,850,476 (C)
Beat Half the Market	$1,733,119 (D)
Varian	$1,630,044 (E)

Table 4-2: Strategy profit levels and groupings for the base parameterization (100 iterations/treatment analyzed, $\alpha = 0.05$)

The Trigger and DF are the most successful strategies for generating profit when shopbot usage is not considered. This result is not surprising due to the fact that those two strategies make moves based on profit results. The next two strategies, LPM and BHM, react to the market, but do not inherently consider profit levels. Those two strategies merely make decisions in order to put set prices at a certain level relative to the competition. The presence of the Varian model at the bottom of the profitability levels makes sense since it is not really a "strategy," it is merely a mathematical model of how retail pricing (without agents) occurs.

The shopbot usage effect was also significant and Fisher's LSD procedure was used to examine differences among the four levels. Table 4-3 shows the average profits for the fifth seller for each shopbot usage level, ignoring the pricing strategy used. Each level is significantly different from all others (at the $\alpha = 5\%$ level) and profitability increases as shopbot usage increases. This increase comes from the fact that with greater comparison shopping, purchases are made more quickly and more often. That is, buyers generally have visibility to more prices and therefore are more likely to quickly find a price that is lower than their valuation. Because this simulation holds the size of the buyer population constant and immediately replaces shoppers who have purchased with new shoppers, the overall level of sales increases with shopbot usage. This result is interesting to note, but taken by itself is of little importance to managers because it involves nothing they have control over.

Shopbot Usage	Average Profit (Group)
80%	$2,022,577 (A)
60%	$1,898,874 (B)
40%	$1,792,508 (C)
20%	$1,716,172 (D)

Table 4-3: Shopbot usage profit levels and groupings for the base parameterization (100 iterations/treatment analyzed, $\alpha = 0.05$)

The presence of an interaction between the two primary effects described above makes it logical and important to examine the relative performance of the pricing strategies at each shopbot usage level. Fisher's LSD procedure was used to make all pairwise comparisons at each shopbot usage level. Table 4-4 shows the relative performance of the pricebot strategies for each shopbot usage level. The table lists the strategies in order of decreasing profits and shows grouping designations from the application of Fisher's LSD procedure ($\alpha = 5\%$). As Figure 4-1 showed, Table 4-4 also demonstrates the interaction effect present. Three of the four non-Varian strategies (DF, Trigger, and BHM) are in the highest Fisher's grouping for at least one shopbot usage level. The fourth (LPM) is in the second highest group twice. Other positions in the rank order are highly variable depending on shopbot usage as well. It is clear from these results that there is an advantage to be gained by being the first in a market to adopt a strategic pricebot approach. For each shopbot usage level, the Varian model was outperformed by at least two pricebot strategies. At its best, Varian placed in the middle-performing Fisher's LSD group. Table 4-5 shows the percentage gains/losses in average profitability achieved by the fifth seller by choosing a non-Varian approach. For the highest and lowest shopbot usage levels, the performance range of the shopbot strategies is very wide—in fact, at 80% shopbot usage, the difference between the best-performing Trigger and worst-performing BHM is actually greater than the Varian profit level. The range of profit levels at the middle levels of shopbot usage, while smaller, are still clearly sizable and important. The highest and lowest shopbot usage levels also show greater risk of decreased performance resulting from a switch from the Varian model to a shopbot strategy.

Relative Rank of Pricebot Strategies (Group)			
20%	**40%**	**60%**	**80%**
BHM (A)	DF (A)	Trigger (A)	Trigger (A)
DF (B)	BHM (A)	DF (B)	LPM (B)
Varian (C)	Trigger (B)	LPM (B)	DF (C)
Trigger (D)	Varian (B)	BHM (C)	Varian (D)
LPM (E)	LPM (C)	Varian (C)	BHM (E)

Table 4-4: Relative profitability rankings (and groupings) of pricebot strategies for the base parameterization (100 iterations/treatment analyzed, α = 0.05)

Gain (Loss) in Profits Compared to Varian Model							
20%		**40%**		**60%**		**80%**	
BHM	33.1%	DF	22.1%	Trigger	36.9%	Trigger	101.4%
DF	19.2%	BHM	19.6%	DF	26.5%	LPM	76.5%
Varian	0.0%	Trigger	4.1%	LPM	21.7%	DF	22.6%
Trigger	(24.3%)	Varian	0.0%	BHM	1.1%	Varian	0.0%
LPM	(28.7)%	LPM	(8.2)%	Varian	0.0%	BHM	(33.1%)

Table 4-5: Changes in average profitability realized by switching from Varian model to strategic pricebot (100 iterations/treatment analyzed)

Profit is an obvious choice as a measure of interest for comparisons of pricebot strategies, but sellers might also be interested in generating sales in order to increase market share. The same ANOVA analysis performed above on profitability data was performed on sales data. As with the profitability data, both primary factors and the interaction were significant ($p < 10^{-307}$ for all) when 100 iterations per ANOVA cell were considered. Figure 4-2 illustrates the relative sales performance of the strategies studied. Unlike the interaction found in the profit data, the sales interaction does not involve a multitude of intersecting lines as occurs in Figure 4-1, which illustrates the profit interaction. However, the gaps in performance between certain strategies change dramatically as shopbot usage is varied. The Trigger and LPM strategies perform comparably across all the shopbot usage levels, but their advantage over the other strategies increases as shopbot usage increases. The Varian, DF, and BHM strategies are grouped tightly together at 20% shopbot usage, but disperse as shopbot usage increases. Varian sales increase, BHM sales decrease, and DF sales remain relatively constant as shopbot usage is increased. For pricing managers who seek a pricebot strategy for gaining market share, Trigger and LPM are both clearly good choices. When examining data across all shopbot usage levels, Trigger is statistically the best when using Fisher's LSD procedure (α = 5%). When examining each shopbot usage level individually, Trigger is alone in the top Fisher's LSD group at every level except 20%, where LPM is also in the top group.

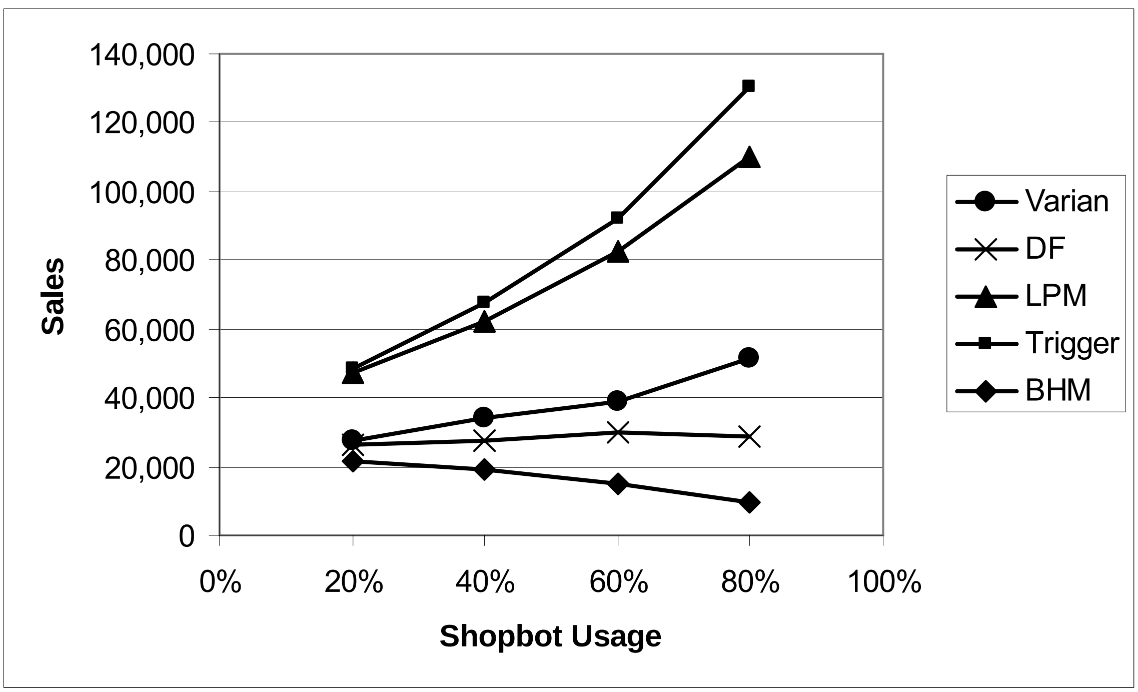

Figure 4-2: Sales interaction between shopbot usage and pricebot strategy (100 iterations per treatment)

To examine the sensitivity of the results above to the parameters described in Table 4-1, multiple alternate parameterizations were run and studied. Table 4-6 shows the combinations examined. Parameter values were chosen such that the ratio between alternate and base was at least on the order of two-to-one. For each alternate parameterization (AP-X), a single deviation was introduced from the base parameterization (BP) described in Table 4-1. All other parameters were left equal to their base values.

Alternate Parameterization	Change from Base
AP-1	Valuation discount rate increased from 0.01 to 0.05
AP-2	Valuation discount rate decreased from 0.01 to 0.0025
AP-3	Shopping rate for non-agent buyers decreased from 0.5 (every other day) to 0.25 (once every four days)
AP-4	Repricing rate for Varian sellers increased from 0.0667 (twice a month) to 0.1429 (once a week)
AP-5	Initial buyer valuation distribution changed from normal to uniform from 1900 to 2100

Table 4-6: Alternate parameterizations

In examining profit data for the alternate parameterizations, very few differences from the base case were found. For AP-1, overall profit levels were increased due to more frequent buying, but the relative performance of the strategies was very similar to the BP. The only notable difference was at the 20% shopbot usage level, where DF joins BHM in the top performing group (using Fisher's LSD procedure with $\alpha = 0.05$). For AP-2, overall profit levels were lower than the BP, but the relative performance of strategies in AP-2 was nearly identical the BP results. For AP-3, overall profit levels were suppressed due to less frequent shopping, but differences in the relative performance of strategies were very subtle. For AP-4, profit levels were comparable to the BP and relative performance of strategies was virtually identical to the BP. Finally, AP-5 results were virtually indiscernible from those of the BP. Sales data was also analyzed for all of the alternate parameterizations. Differences in sales results were also very minor, with no differences worthy of reporting. Because the altering of parameters led to very limited deviations in the results, the remainder of the analysis focused on the base parameterization.

Round Two Results

While Round One of simulations was designed to examine scenarios where a seller is the first to utilize a pricebot in its market, Round Two is geared toward a best response to a successful Round One strategy. Because having multiple LPM, Trigger, or BHM sellers in this simulation environment creates uninteresting markets where prices decrease almost exclusively, a treatment where DF was successful was chosen for study. The specific treatment chosen for Round Two study was the 40% shopbot usage level for the base parameterization, a treatment for which DF was the top performer. The scenario considered is that all sellers in the market have migrated to the DF strategy. The problem analyzed is: "if one seller wishes to switch from DF to gain an advantage, what is that seller's best strategy choice?"

One hundred simulation iterations were run for each of the five strategies versus four DF competitors, all with the base parameterization and shopbot usage at 40%. One-way ANOVA was performed on the strategy factor and it was significant ($p < 10^{-204}$). Table 4-7 shows average profit levels for each strategy and grouping indicators (using Fisher's LSD procedure at $\alpha = 0.05$). Trigger is clearly the best response to DF in this market—adding the low price match component to the derivative follower technique makes for an effective counter-strategy. If a market (with 40% shopbot usage) evolved to a point where all sellers migrated to DF pricebots, one of the five could switch to Trigger and realize a 28% increase in profits.

Strategy	Average Profit (Group)
Trigger	$2,526,657 (A)
Derivative Follower	$1,973,989 (B)
Varian	$1,697,156 (C)
Low Price Match	$1,434,573 (D)
Beat Half the Market	$855,814 (E)

Table 4-7: Strategy profit levels and groupings for Round Two simulations (100 iterations/treatment analyzed , $\alpha = 0.05$)

Sales were also examined for this Round Two scenario. One-way ANOVA was again performed and the strategy factor was significant ($p < 10^{-209}$). Table 4-8 shows average sales levels for each strategy and grouping indicators (using Fisher's LSD procedure at $\alpha = 0.05$). For market share seekers, Low Price Match is the best response to a market of DF sellers in this market, offering an improvement in sales of 87% over DF sellers. Trigger is second best choice. With Trigger proving to be the most profitable in Round Two, its strong sales performance makes it a good choice for sellers regardless of what their primary goal is.

Strategy	Average Sales (Group)
Low Price Match	74,362 (A)
Trigger	68,684 (B)
Varian	43,028 (C)
Derivative Follower	39,746 (D)
Beat Half the Market	29,617 (E)

Table 4-8: Strategy sales levels (per seller) and groupings for Round Two simulations (100 iterations/treatment analyzed, $\alpha = 0.05$)

Dominant Seller Effects

All of the work above assumes that no sellers were favored over any others—for buyers who shopped at one or two sellers (those not using shopbots), the probability of being selected by a buyer was equal across all sellers. Two final scenarios were chosen to be studied in Chapter Four, with each involving dominant sellers. Each of these scenarios was constructed such that one of the sellers was twice as likely as the others to be selected by the non-shopbot buyers. Both utilize the base parameterization described in Table 4-1. The first scenario examines the performance of the dominant seller, while the second examines one of the four non-dominant sellers.

Scenario 1

When examining the profitability of the dominant seller, two-way ANOVA was performed as above in the Round One studies. In this scenario, only the strategy primary factor and the interaction were significant ($p < 10^{-67}$). The shopbot usage factor had a p-value of 0.67. Figure 4-3 is a graphical depiction of the interaction present. There are some similarities to what was observed in Figure 4-3. Trigger and LPM are once again roughly parallel and increasing with shopbot usage. Beat Half the Market once again is the most profitable at the lowest shopbot usage level and the least profitable at the highest. Derivative Follower was constant or slightly increasing in the base case, but for the dominant seller, DF becomes much less profitable as shopbot usage increases. Also, the overall difference in performance across shopbot usage levels is diminished in the dominant seller case (as evidenced by the high p-value for the shopbot usage factor).

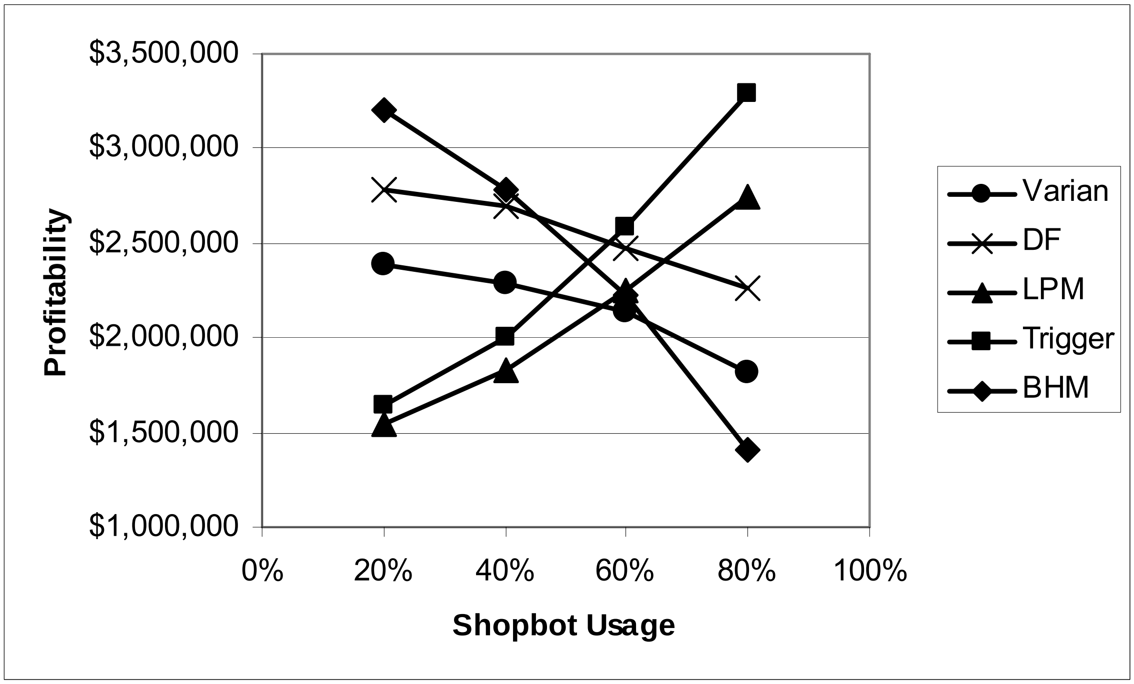

Figure 4-3: Profit interaction between shopbot usage and pricebot strategy for a dominant seller (100 iterations per treatment)

Table 4-9 shows average profitability (across all shopbot usage levels) for each strategy employed by the dominant seller. When sellers had equal probability of being chosen by buyers, Trigger was the best choice, while DF and LPM were both in the next highest performing group. For a dominant seller, however, DF becomes the best performer while Trigger drops into the next highest performing group with BHM. BHM had been the next-to-worst performer in the base case. LPM proves to be the worst choice when shopbot usage is removed from the analysis—it was a mid-pack performer in the base case.

Shopbot Usage	Average :Profit (Group)
Derivative Follower	$2,555,017 (A)
Beat Half the Market	$2,406,216 (B)
Trigger	$2,381,677 (B)
Varian	$2,157,001 (C)
Low Price Match	$2,090,458 (D)

Table 4-9: Profit levels and groupings for dominant seller simulations (100 iterations/treatment analyzed, $\alpha = 0.05$)

When examining the relative performance of the strategies within each shopbot usage levels, the rank orders are very similar to what was reported in Table 4-4. The only important deviation in the dominant seller scenario is the placement of DF in the top performing group (with Trigger) at 60% shopbot usage—Trigger had been alone in the top performing group for that market in the base case. Table 4-10 shows the percentage gains/losses in average profitability achieved by the fifth seller by choosing a non-Varian approach. When comparing Table 4-10 to Table 4-5, it can be seen that the range of performance for the lower shopbot usage levels is larger for the dominant seller case than for the base. However, that same range is lower for the dominant seller scenario when the higher shopbot usage levels are compared. While shopbot strategy selection is clearly important for any seller in any market, it is of heightened importance for a dominant seller in a low shopbot usage market. Sales data for the dominant seller were analyzed, but the results are not reported as they were very similar to those from the base case.

Gain (Loss) in Profits Compared to Varian Model							
20%		40%		60%		80%	
BHM	34.5%	DF	21.5%	Trigger	20.6%	Trigger	81.7%
DF	16.9%	BHM	17.6%	DF	15.6%	LPM	51.6%
Varian	0.0%	Varian	0.0%	LPM	4.9%	DF	25.0%
Trigger	-31.0%	Trigger	-12.4%	BHM	3.8%	Varian	0.0%
LPM	-35.2%	LPM	-20.4%	Varian	0.0%	BHM	-21.9%

Table 4-10: Changes in dominant seller average profitability realized by switching from Varian model to strategic pricebot (100 iterations/treatment analyzed)

Scenario 2

Finally, the same dominant seller market was studied from the perspective of one of the *dominated* sellers. In the case of the dominated seller, relative performance of the strategies for both profit and sales were nearly identical to those from the base case where all sellers were equivalent.

In summary, it is clear that decisions regarding strategic pricebots be made with an appreciation of the type of market in which the pricebot will be working. The Round One simulation data represents a case where a seller is the first in her market to make use of a pricebot. That data shows a strong interaction between strategy and shopbot usage for both profit and sales. For the profit case, the interaction is such that a strategy may be the best performer in one market but the worst performer in another. It is also clear that there are strong motives for implementing pricebots, as double digit percentage gains (in both profits and sales) were observed for multiple market scenarios. However, pricing managers must proceed with caution, as double digit percentage losses were also observed. Chapter Five shifts the focus from seller to buyer. It examines the relative merits of adding a memory component to the simple shopbot implementation modeled here in Chapter Four.

Chapter Five: Enabling Shopbots with Memory

Hertweck et al. (2003) studied the effects of varying implementations of shopbots on the qualitative nature of a specific marketplace. They simulated a market similar to the one studied in Chapter Four. The Hertweck et al. market had five Derivative Follower (DF) sellers. Fifty percent of the buyers in that market shopped at all five sellers, while 25% shopped at two sellers, and 25% shopped at only one. The focus of the Hertweck et al. study was the impact of differing buyer behavior on market prices. Hertweck et al. studied the buyer types represented in Table 5-1.

		Valuation	
		Low	**High**
Threshold	**Low**	*Buyers who seek bargains and are able to wait to purchase* ***(Bargain Hunter)***	*Buyer who wants to find a great deal but has an immediate need* ***(Conflicted Buyer)***
	High	*Buyer unconcerned about finding the great deal and can wait to purchase* ***(Casual Buyer)***	*Buyer with an immediate need and little concern for finding a bargain* ***(Impulse Buyer)***

Table 5-1: Buyer categories studied in Hertweck et al. (2003)

These buyer profiles were achieved by adding a memory component to the buyer type simulated in Chapter Four. In Chapter Four, buyers were assigned an initial valuation that was subsequently increased each period according to a fixed discount rate. The Chapter Four buyer buys when a price that is lower than his valuation is observed (or when the end of the assigned time horizon is reached). When the memory component is added, the buyer is assigned a threshold value in addition to a valuation. The threshold is based on the best price observed in a period just prior to the buyer's entrance into the market (a memory length parameter controls the size of that window). The threshold is also increased over time according to another discount rate. The memory-enabled shopbot buys only when it observes a price that is lower than both its valuation and threshold (or, again, when the end of the time horizon is reached).

For this chapter, the work of Hertweck et al. (2003) is continued. The Hertweck et al. (2003) work was an introductory and highly qualitative study. The research of this chapter builds on the prior work by broadening the scenarios covered and adding statistical rigor. The Chapter Four simulation tool was utilized, with enhancements to model the memory behavior described above. First, the Chapter Four study, which incorporated a homogeneous buyer market, was re-visited with a new heterogeneous mix of buyers to study the impact on seller strategy selection. Next, specific seller combinations were chosen to examine relative performance of various shopbot implementations. The parameters shown in Table 5-2 were used for the Chapter Five simulations.

Parameter description	Value/Distribution
Top end of price range	2100
Bottom end of price range	1900
Number of sellers	5
Number of buyers	1000
Number of time periods per iteration	365 (days/year)
Rate of repricing for Varian sellers.	0.0667 (re-pricing an average of twice per month)
Rate of shopping for non-agent buyers.	0.5 (shopping an average of once every two days)
Unit cost to seller. Set equal to bottom end of price range.	1900
Initial buyer valuation	Normal with ±2σ in price range
Buyer horizon	Normal centered on 7.5, with ±2σ going from 0 to 15 and truncated at 1.
Memory length	7 (days)

Table 5-2: Base parameterization for Chapter Five simulations

The buyer markets studied in Chapter Four were homogeneous. Buyers had no memory component and all shared the same valuation discount rate. For this chapter, six buyer shopbot implementations were studied. Two of these six were non-memory buyers with two different valuation discount rates: 1% representing “high” and 0.25% representing “low.” The other four pricebot types were modeled after the categories of Table 5-1. For these four, 1% and 0.25% were again used to represent “high” and “low” discount rates—for both the valuation and the memory-based threshold. The markets studied in this chapter are heterogeneous—the shopbot portion (those who shop all 5 sellers) of the buyer population consists of an even mix of the six buyer shopbot implementations described above. The non-shopbot portion of the buyer population was also made heterogeneous, with an even split between the two valuation discount rates.

The first portion of this experiment involves replicating what was done in the previous chapter, but with the new, heterogeneous buyer population. Again, 100 replications per 20 treatments (4 shopbot usage levels by 5 pricebot strategies) were performed. As was done in Chapter Four, the scenario of one seller switching from a Varian model of pricing was studied. When ANOVA was performed on seller profitability data from the 100 replications, the strategy effect and the interaction were clearly significant ($p < 10^{-97}$). The shopbot usage factor was marginally significant, with a p-value of 0.064. Figure 5-1 clearly shows the interaction between strategy selection and shopbot usage.

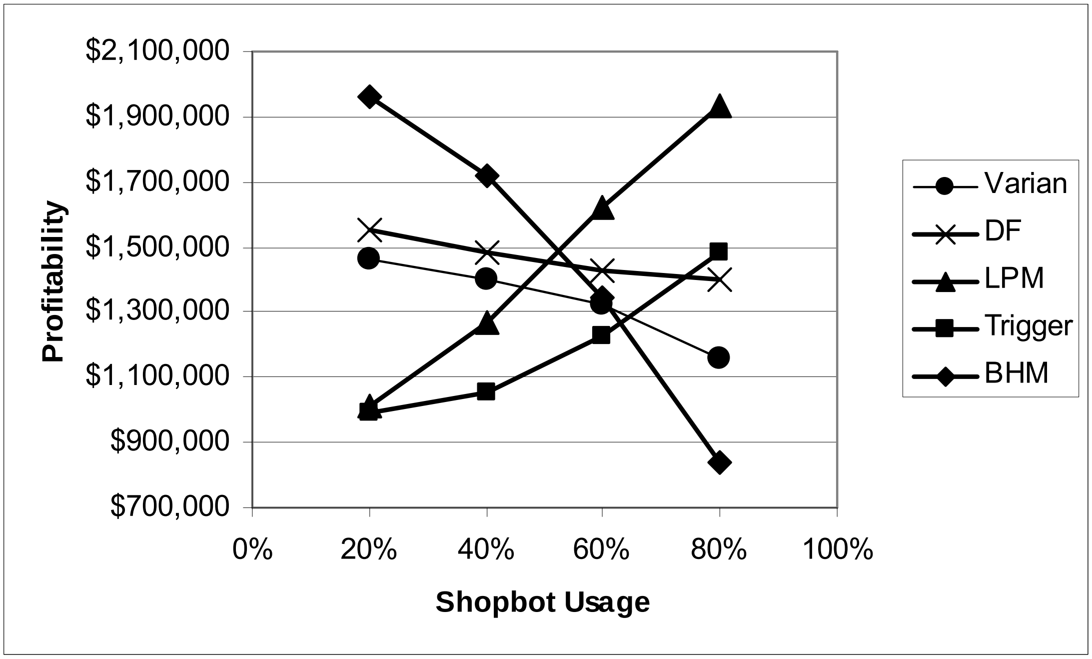

Figure 5-1: Profitability interaction between shopbot usage and pricebot strategy (100 iterations per treatment)

As in Chapter Four, the interaction present makes it clear that pricing managers need to understand the nature of the market in which they compete when making strategy decisions. For that situation, multiple comparisons between strategies were made across all levels of shopbot usage. Fisher's Least Significant Difference (LSD) procedure was used to examine strategy differences across all shopbot usage levels. Table 5-3 shows average profitability levels along with group designations. Profit levels are suppressed from the homogeneous buyer market with the new heterogeneous shopbot market. In this new market, many buyers wait longer to buy than they would have in the homogeneous market, therefore, the buyer population does not refresh as often. Additionally, interesting comparisons between the rank orders of strategies in Tables 4-2 and 5-3 can be made. In Chapter Four, Trigger was in the highest performing group, while here it is the worst performing strategy, ranking below even the non-strategic Varian model. Low Price Match and Beat Half the Market ascend to the highest performing group with the switch to a heterogeneous buyer market, having been the third- and fourth-highest previously. Derivative Follower, however, demonstrates relatively consistent performance. It is in the most profitable strategy group in this chapter and was the second-best performer with the homogeneous market.

Strategy	Average Profit (Group)
Beat Half the Market	$1,465,493 (A)
Derivative Follower	$1,465,473 (A)
Low Price Match	$1,459,192 (A)
Varian	$1,335,370 (B)
Trigger	$1,190,142 (C)

Table 5-3: Strategy profit levels and groupings for the base parameterization (100 iterations/treatment analyzed, $\alpha = 0.05$)

The presence of an interaction between the two primary effects described above makes it logical and important to examine the relative performance of the pricing strategies at each shopbot usage level. Fisher's LSD procedure was used to make all pairwise comparisons at each shopbot usage level. Table 5-4 shows the relative performance of the pricebot strategies for each shopbot usage level. The table lists the strategies in order of decreasing profits and shows grouping designations from the application of Fisher's LSD procedure ($\alpha = 5\%$). The interaction depicted in Figure 5-1 is also evident in Table 5-4. The heterogeneous market interaction, however, is not as strong as the one observed with the homogeneous market. In Table 4-4, three different strategies appeared at least once in the top performing group. In Table 5-4, only BHM and LPM earn top rankings, with BHM succeeding at low shopbot usage levels and LPM at high. Table 5-5 shows the percentage gains/losses in average profitability achieved by the fifth seller by choosing a non-Varian approach. The gains and losses shown in Table 5-5 for switching from the Varian pricing model are generally comparable to what was observed in Chapter Four, with a few exceptions. The potential loss in profits for choosing the wrong strategy at the 40% shopbot usage level is much higher for the heterogeneous shopbot environment of this chapter. Additionally, the market studied here offers lower potential profitability increases for the two highest shopbot usage levels than those that were observed for the homogeneous market. Sales data was also collected and analyzed. Results were very similar to those of Chapter Four and, consequently, are not reported here.

Relative Rank of Pricebot Strategies (Group)			
20%	**40%**	**60%**	**80%**
BHM (A)	BHM (A)	LPM (A)	LPM (A)
DF (B)	DF (B)	DF (B)	Trigger (B)
Varian (C)	Varian (C)	BHM (C)	DF (B)
LPM (D)	LPM (D)	Varian (C)	Varian (C)
Trigger (D)	Trigger (E)	Trigger (D)	BHM (D)

Table 5-4: Relative profitability rankings (and groupings) of pricebot strategies at varying shopbot usage levels (100 iterations/treatment analyzed, α = 0.05)

Gain (Loss) in Profits Compared to Varian Model							
20%		**40%**		**60%**		**80%**	
BHM	33.9%	BHM	22.9%	LPM	22.4%	LPM	67.5%
DF	6.1%	DF	6.1%	DF	7.7%	Trigger	28.7%
Varian	0.0%	Varian	0.0%	BHM	1.6%	DF	21.2%
LPM	-30.8%	LPM	-9.3%	Varian	0.0%	Varian	0.0%
Trigger	-32.3%	Trigger	-24.6%	Trigger	-7.2%	BHM	-27.5%

Table 5-5: Changes in average profitability realized by switching from Varian model to strategic pricebot (100 iterations/treatment analyzed)

The data discussed above indicates that the makeup of the shopbot population has important implications for the performance of pricebot strategies. Compared to the effects of the alternate parameterizations detailed in Table 4-6, the effects of changing the makeup of the shopbot population were considerable. Chapter Four demonstrated that it is important for pricing managers to understand the extent to which their customers are using shopbots. This chapter underscores that fact and adds another important dimension for sellers to consider—the types of shopbots that their customers use is an important consideration in pricebot strategy selection as well. While the effects of varying shopbot implementations on sellers are interesting, buyers likely have a greater interest in how these varying implementations affect them.

Comparing Shopbot Implementations

To examine the relative performance of the six shopbot implementations present in the heterogeneous buyer markets, shopbot purchase price data was sampled from the 100 iterations per treatment that were described above. Details of those six implementations are shown in Table 5-6. In order to isolate the differences between the six, only a subset of the shopbots was sampled. Any buyers who would have bought during their first period in the market regardless of their implementation type were excluded.

Shopbot Implementation	Valuation Discount Rate	Threshold Discount Rate
Base Low	0.25%	N/A
Base High	1.00%	N/A
Bargain	0.25%	0.25%
Casual	0.25%	1.00%
Conflicted	1.00%	0.25%
Impulse	1.00%	1.00%

Table 5-6: Shopbot implementation designations and parameters

The markets shown in Table 5-7 were chosen for study based on seller profitability results from above. Markets M-1 through M-7 represent the scenarios (for each level of shopbot penetration) in which an initial seller decides to utilize a pricebot and chooses the top performing strategy from the heterogeneous buyer market simulations described above. Markets M-2, M-4, and M-6 utilize the best performing strategies from Chapter Five. Markets M-3, M-5, and M-7 utilize top performers from Chapter Four. The homogeneous and heterogeneous market results were in agreement for the 20% shopbot usage market represented by M-1.

Designation	Seller Mix	Shopbot Penetration
M-1	4 Varian, 1 Beat Half the Market	20%
M-2	4 Varian, 1 Beat Half the Market	40%
M-3	4 Varian, 1 Derivative Follower	40%
M-4	4 Varian, 1 Low Price Match	60%
M-5	4 Varian, 1 Trigger	60%
M-6	4 Varian, 1 Low Price Match	80%
M-7	4 Varian, 1 Trigger	80%

Table 5-7: Market designations and descriptions

Price paid was chosen as the measure of interest for comparing the shopbot implementation types. ANOVA was again performed on various numbers of samples to determine the best level at which to conduct tests. At three hundred samples, the effect of the shopbot implementation was significant ($\alpha < 0.05$) for five of the seven markets studied. Markets M-3 and M-6 did not have a significant shopbot implementation effect. The average prices for each of the five markets with significant effects are listed (in increasing price order) in Table 5-8. Fisher's LSD procedure ($\alpha = 0.05$) was performed to determine groupings of implementations—those designations are also listed in Table 5-8. The differences among the buyer treatments, while statistically significant, were generally not dramatic. The differences between top and bottom performing treatment ranged from \$2.13 to \$12.44, or 0.11% to 0.64% of the average price paid for each respective market. While that level of savings may draw limited interest from individual consumers making individual purchases, those making high volume individual purchases or repeated purchases may see significant potential value in choosing the right type of shopbot.

Average Price Paid (Group)					
M-1		M-2		M-3	
Conflicted	$1,937.15 (A)	Conflicted	$1,940.15 (A)	Bargain	$1,927.79 (A)
Impulse	$1,938.86 (A)	Bargain	$1,940.66 (A)	Conflicted	$1,929.09 (AB)
Bargain	$1,940.62 (A)	Impulse	$1,941.43 (A)	Base Low	$1,930.16 (AB)
Casual	$1,941.88 (A)	Casual	$1,943.97 (AB)	Casual	$1,930.42 (AB)
Base Low	$1,943.06 (A)	Base Low	$1,945.45 (AB)	Base High	$1,930.93 (AB)
Base High	$1,949.58 (B)	Base High	$1,948.00 (B)	Impulse	$1,931.56 (B)
M-4		M-5		M-6	
Conflicted	$1,933.76 (A)	Bargain	$1,913.15 (A)	Bargain	$1,933.92 (A)
Bargain	$1,935.99 (AB)	Conflicted	$1,913.57 (AB)	Conflicted	$1,935.72 (A)
Base Low	$1,937.55 (ABC)	Impulse	$1,914.72 (BC)	Casual	$1,936.84 (A)
Casual	$1,937.83 (ABC)	Base Low	$1,914.96 (BC)	Base Low	$1,936.87 (A)
Impulse	$1,941.01 (BC)	Casual	$1,915.04 (C)	Base High	$1,938.07 (A)
Base High	$1,943.14 (C)	Base High	$1,915.29 (C)	Impulse	$1,938.76 (A)
M-7					
Conflicted	$1,912.56 (A)				
Bargain	$1,913.13 (AB)				
Impulse	$1,913.54 (AB)				
Base Low	$1,913.92 (BC)				
Casual	$1,914.37 (BC)				
Base High	$1,914.95 (C)				

Table 5-8: Average prices paid and groupings for all markets studied (100 samples/treatment, α = 0.05, groups excluded for non-significant ANOVAs)

There were no markets in which one particular shopbot implementation could be deemed the best with confidence—the top performing group (according to Fisher's LSD) had anywhere from two to five members. There are, however, some interesting consistencies across the eight markets studied. The two non-memory implementations, Base High and Base Low, were in the bottom half of the rank order for every market except M-4, where Base Low ranked third. Both were in the worst performing Fisher's group in every market except M-1, where Base Low was in the top group. Following that observation, the average memory-enabled implementation performance was tested against the average non-memory shopbot, using contrasts on the ANOVA data. The difference between average memory-enabled shopbot price paid and average non-memory shopbot average price paid was significant for the same five markets that had significant ANOVA results (all p values < 0.04), with the lower price being paid by the memory-enabled shopbots in all the markets. Table 5-9 shows the averages for each market.

Average Price Paid				
	M-1	**M-2**	**M-3**	**M-4**
Memory	\$1,939.63	\$1,941.55	\$1,929.71	\$1,937.15
Non-memory	\$1,946.32	\$1,946.73	\$1,930.55	\$1,940.34
	M-5	**M-6**	**M-7**	
Memory	\$1,914.12	\$1,936.31	\$1,913.40	
Non-memory	\$1,915.12	\$1,937.47	\$1,914.43	

Table 5-9: Average prices paid for memory vs. non-memory shopbots for all markets studied

Because memory was found to be an important discriminating characteristic for shopbot performance, an additional comparison was made among only the memory-enabled shopbot treatments. The threshold discount rate is the memory-related parameter that differentiates the Bargain Hunter and Conflicted buyers from the Casual and Impulse buyers. The contrast associated with comparing the averages of those two pairs of treatments was analyzed for all seven markets. For this contrast, only four of the seven markets showed a significant difference (those four p values < 0.02) between high and low discount threshold shopbots, with the low threshold discount rate treatments (bargain hunter and conflicted) paying the lower average price. The markets that did not show significance were M-1, M-2, and M-6. Table 5-10 shows the averages for each market.

Average Price Paid				
Threshold Discount Rate	**M-1**	**M-2**	**M-3**	**M-4**
Low (0.25%)	$1,938.88	$1,940.41	$1,928.44	$1,934.87
High (1.00%)	$1,940.37	$1,942.70	$1,930.99	$1,939.42
	M-5	**M-6**	**M-7**	
Low (0.25%)	$1,913.36	$1,934.82	$1,912.85	
High (1.00%)	$1,914.88	$1,937.80	$1,913.96	

Table 5-10: Average prices paid for low vs. high threshold discount rate shopbots for all markets studied

Alternate Parameterizations

Table 5-2 shows the base parameters used for all of the analyses above. As was done in Chapter Four, alternate parameterizations were studied to assess the sensitivity of the results to key simulation parameters. Table 5-10 shows the alternate parameterizations studied. AP-1 and AP-2 were not carried over from Chapter Four because the shopbot implementation treatments cover those variations. AP-6 and AP-7 were added so that the impact of varying memory length could be analyzed.

Alternate Parameterization	Change from Base
AP-3	Shopping rate for non-agent buyers decreased from 0.5 (every other day) to 0.25 (once every four days)
AP-4	Repricing rate for Varian sellers increased from 0.0667 (twice a month) to 0.1429 (once a week)
AP-5	Initial buyer valuation distribution changed from normal to uniform from 1900 to 2100
AP-6	Memory length decreased from 7 to 3 periods
AP-7	Memory length increased from 7 to 14 periods

Table 5-11: Alternate parameterizations

When seller profit levels were examined in the alternate parameterization cases, there were a few interesting differences. For the base Chapter Five case, Beat Half the Market, Derivative Follower, and Low Price Match were all part of the top performing Fisher's group, when profitability across all shopbot usage levels was examined. However, for each of the alternate parameterizations, only one of those three place in the top-performing group. For AP-3 and AP-4, DF was the top performer. For AP-5 and AP-7, BHM performed best. For AP-6, LPM was alone at the top. While the results across shopbot usage levels showed some interesting deviations for the alternate parameterizations, the results when profitability was examined within each usage level showed great consistency across all five AP's. Sales data was also examined and the alternate parameter results were very consistent with what was observed in the base case.

When shopbot implementations were analyzed, the alternate parameterizations did introduce significant changes in the results. The five alternate parameterizations of Table 5-11 were applied to Markets M-1 through M-7 (described in Table 5-7). While the shopbot implementation effect was significant (at $\alpha = 0.05$) for five of those seven markets when the base parameterization was used, eleven of thirty-five alternate parameterization scenarios resulted in no significant shopbot implementation effect. Three of those eleven were marginally significant, with p-values between 0.05 and 0.1. Those combinations are listed in Table 5-12

Alternate Parameterization	Markets With No Effect
Base	M-3, M-6
AP-3	M-2, M-4*, M-6
AP-4	—
AP-5	M-6
AP-6	M-2, M-4, M-5, M-6*, M-7
AP-7	M-1, M-6*

Table 5-12: Scenarios where no shopbot implementation effect was present ($\alpha = 0.05$, * indicates $0.05 < p < 0.1$)

The results summarized in Table 5-12 lead to some interesting observations. All alternate parameterizations except AP-4 have some markets where there is no shopbot implementation effect. In AP-4, the Varian sellers re-price more frequently than in the base case. For AP-6, only two of seven markets exhibit a significant implementation effect. In AP-6 the shopbot memory length is reduced from the base case. For AP-7, as in the base case, five of seven markets showed a significant implementation effect, with one non-significant market having a p-value of 0.096. These observations lead to the conclusion that memory length is likely an important choice when adding a memory component to the shopbot. Further, it could be that there is some sort of interaction between memory length and seller re-pricing that is important to understand when designing memory-enabled shopbots. Further study would be needed to determine the best approach for choosing memory length. For the combinations that had a significant shopbot implementation effect, the relative rankings of the implementation types were similar to the base case. The two base types were again typically in the lower half of the rankings and in the lowest group as determined by Fisher's Least Significant Difference

(LSD) procedure ($\alpha = 0.05$). The Bargain Hunter and Conflicted implementation types were typically among the best performers.

The same constrasts that were studied in the base case (for memory vs. non-memory and high vs. low threshold rate) were again investigated for the alternate parameterizations. Table 5-13 lists the markets that showed no significance for these tests for all of the parameterizations studied.

Alternate Parameterization	Markets With No Memory Effect	Markets With No Threshold Rate Effect
Base	M-3, M-6	M-1, M-2, M-6*
AP-3	M-3, M-6	M-2, M-4, M-6
AP-4	—	M-1*, M-2
AP-5	M-4*, M-6	M-4*, M-6
AP-6	M-2, M-4, M-5, M-7	M-2, M-4, M-5, M-6, M-7*
AP-7	M-1*, M-2*, M-4	M-1*, M-2, M-6

Table 5-13: Scenarios where no memory or threshold rate differences were found (contrasts tested at $\alpha = 0.05$, * indicates $0.05 < p < 0.1$)

When examining the contrast associated with comparing memory to non-memory shopbots, 24 out of the 35 alternate parameter scenarios exhibited a significant difference between memory and non-memory implementations at the 5% level. Three scenarios showed marginal significance and eight show no significant difference between memory and non-memory shopbots. Half of those eight non-significant scenarios occurred under AP-6, where memory length is decreased from seven to three periods. With memory length seven (base case), two scenarios exhibited no significant memory/non-memory difference. With memory length fourteen, one scenario showed no significance and two were marginal. None of the scenarios with no difference between memory and non-memory bots occurred under AP-4, where Varian sellers re-priced at an increased rate. These results underscore the need to further investigate the potential interplay between seller re-pricing rates and buyer memory lengths. When examining the contrast associated with threshold discount rate, the pattern of significance in the scenarios studied is similar, but with fewer scenarios that are significant. Only 20 of 35 scenarios exhibited a significant (at the 5% level) difference between the high and low threshold discount rate levels. Three scenarios were marginally significant and the remaining twelve showed no threshold discount rate effect.

The simulations examined in this chapter show that the manner in which a shopbot is implemented can impact the price that the buyer ultimately pays. A shopbot implementation effect was identified in five of seven markets studied under the base parameterization. It was also shown that the implementations that included a memory component outperformed those that did not in five of the seven base case scenarios. Additionally, in a majority of the seven markets analyzed, memory-enabled shopbots with a low threshold discount rate outperformed those with a high threshold discount rate. The observed differences in price paid are generally small compared to the total average price paid, but there are certainly high volume and repeat buying scenarios encountered in business where such seemingly small savings could prove valuable in practice. Also, the differences observed indicate that further study into methods of implementing shopbots, specifically with a memory component, is warranted. When alternate simulation parameterizations were studied, again, not all markets showed significant shopbot implementation effect. However, a clear majority did. A majority of those scenarios also showed memory-enabled bots to have an advantage over non-memory bots. Also, as in the base case, low threshold discount rate shopbots outperformed high rate bots more often than not. Further study is necessary to better understand and explain the interplay between the simulation parameters and the shopbot implementations—most notably seller re-pricing rates and shopbot memory length. Finally, it was found that the heterogeneous (with respect to buyer types) market studied in this chapter led to some important differences in the analysis of seller pricebot strategy performance—differences greater than those observed when simulation parameters were varied as a part of the Chapter Four work above. Also, the alternate parameterizations led to differences in what strategy performed best across all shopbot usage levels. These results underscore the conclusion from Chapter Four that it is critical for pricing managers to understand their markets when choosing pricebot strategies. While this chapter focused on varied shopbot implementations, the next will shift back to the seller perspective. It will examine how sellers can utilize neural networks to evaluate and respond to competitors in their markets.

Chapter Six: Using Intelligent Agents to Classify and React to Competitors

Chapters Four and Five both examined the use of agents in e-commerce markets. The agents examined, however, utilized a very limited amount of intelligence. Buyers' agents (shopbots) surveyed market prices, maintained valuation and threshold values (updated over time) and made buying decisions based on comparisons among those values. Sellers' agents (pricebots) also had the capability to survey market prices and make pricing decisions based on observed prices and predetermined strategies. The work of Chapter Six seeks to show the potential power of agents that have capabilities that exhibit greater intelligence that the relatively rudimentary agents of Chapters Four and Five. Specifically, Chapter Six will examine the potential of pricebots that are able to classify competitor pricing behavior and implement the most effective counterstrategy for their market.

Chapter Four showed that the extent to which buyers use shopbots to comparison shop can impact the relative success of various pricing strategies—the interaction between seller strategy and shopbot usage was very strong. Additionally, the work of Chapter Four also demonstrated that competitors' strategies affect what selling strategy is most effective—at the 40% shopbot usage level, Derivative Follower was most effective against four Varian competitors, while Trigger was most effective against four Derivative Follower competitors. Chapter Six examines a broader range of competitor strategy mixes, but fixes the shopbot usage level of the market at 50%.

This chapter examines the use of neural networks (NNs) to identify what strategies are being used by competitors. After that determination has been made, the pricebot can consult a table of decision rules to decide what counterstrategy should be most effective for the NN-classified competitor behavior. The pricebot then adopts the strategy indicated by the appropriate rule. Figure 6-1 shows the scenario under study in this chapter.

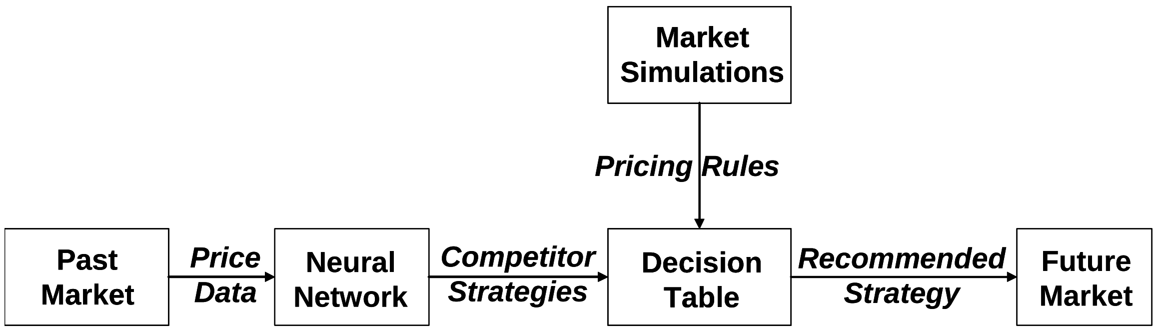

Figure 6-1: Chapter Six System of Study

The following steps were undertaken to examine the potential value of an agent with the capabilities described above.

1. Devised a NN classifier that can accurately determine a competitor's pricing strategy based on historical parameters of the observed market. The same simulation used in previous chapters was used again in this chapter to generate data for NN training and evaluation.
2. Constructed the decision table that shows the best counterstrategy given the strategies in use by four competitors. This table was built by first using the simulation tool to systematically generate data for each possible combination of seller strategies and then using ANOVA to determine the best performing strategy for a seller of interest (the "fifth" seller, as was done in Chapter Four).
3. Randomly generated seller strategy combinations and run the simulation tool for 30 periods for each combination. Measure the profit of the fifth seller. Capture data to be used as NN inputs.
4. Applied the neural network developed in #1 to the captured historical price data in order to classify each of the four competitor sellers by strategy.
5. For the strategy mix identified by the NN, identified the matching rule in the decision table and read the associated best counterstrategy.

6. Changed (if necessary) the strategy of the fifth seller to the best counterstrategy identified in #5.
7. Ran the simulation for 30 more periods. Measured the profit of the fifth seller for the second set of 30 periods.
8. Went back to the ending conditions from the first 30 periods and re-ran the final 30 periods with no strategy change from
9. Compared profit levels for the last 30 periods with and without the strategy change.

As mentioned above, the range of strategy combinations examined in this chapter is broader than what was studied above. However, because certain combinations of strategies lead to uninteresting scenarios where market prices are quickly driven to the market's price floor, certain restrictions had to be put on the markets. The markets examined in this chapter are those in which the total of Low Price Match (LPM), Trigger, and Beat Half the Market (BHM) sellers is no more than three. The basic simulation tool used is the same as Chapter Four and Five.

Tiberius for Excel was chosen as the NN tool for competitor strategy classification. Tiberius for Excel is written by Phil Brierly (www.philbrierley.com) and is described by the author as:

> *"…a feedforward multilayer perceptron trained with the backpropagation algorithm. It is written completely in VBA, the macro language of Microsoft Excel and is self contained in an Excel workbook"*

Several NN designs were considered, implemented, and evaluated. Tiberius is limited to one NN output, which made attacking a classification problem challenging. A multi-output NN may be better suited to strategy classification, but due to budget and time constraints, Tiberius was chosen for this chapter's work. The final decision on design was to train five different NNs to detect each of the five strategies being studied. Each of those five NNs had the same design. After examining the efficacy of a multitude of input combinations, eleven inputs were chosen for the final design and are described below. The aggregate measures were calculated over a 15-period window of prices. A seven-period window was also examined, but the longer window led to more accurate results. Where "rank" appears, it is the rank of a seller's *current* price relative to its competitors' *prior-period* prices

- *Bias:* As recommended by Mr. Brierly, a constant bias input was used.
- *Average of Ranks:* The mean rank of a seller's price.
- *Standard Deviation of Ranks:* The standard deviation of the ranks of a seller's price.
- *Rank Changes:* The number of times that a seller's price rank changed from one period to the subsequent period.
- *Price Standard Deviation:* The standard deviation of the seller's actual prices.
- *Difference from Average:* The difference between a seller's average price and the average price of the entire market.
- *Rank in Averages:* The ranking of a seller's average price against his competitors' average prices.
- *Difference from High:* The difference between a seller's highest observed price and the highest observed price in the market.
- *Difference from Low:* The difference between a seller's lowest observed price and the lowest observed price in the market.
- *Number of Price Changes:* The number of times a seller's price changed from one period to the subsequent period.
- *Average Price Change:* The sum of the absolute value of a seller's price changes divided by the number of price changes.

The NNs used each had 11 hidden nodes and one output node. A set of 1000 training scenarios was generated with the simulation tool. Each of five NNs was trained to recognize one of the strategies studied. The desired strategy to be identified was coded as a desired output of zero and all other strategies were coded as one. Training was done over 400 epochs with a learning rate of 0.8. Four hundred epochs offered an improvement in accuracy over the 200 epoch level, but increasing the number of epochs beyond 400 offered no further gain. The learning rate was left unchanged from the Tiberius default value in order to streamline the design process for this initial study. After each strategy's NN was trained, each could be applied to any price scenario data. The strategy of the NN with the lowest output was used to classify the seller being evaluated.

The trained NNs were then applied to a separate set of 250 evaluation scenarios. For those 250 scenarios, the NN system accurately predicted a seller's strategy 87.8% of the time. Table 6-1 shows the relative accuracy for each strategy. The range of classification accuracies was rather broad, spanning from 71.2% for DF sellers to 95.5% for BHM sellers. The inconsistency of the NN system across strategies is likely due to scenarios for which pricing output generated by one strategy may mirror the typical pattern of another. These confusing scenarios tend to arise from (1) the fact that Trigger is a hybrid strategy, and (2) the fact that LPM is exactly matching other sellers' prices. Trigger is not a unique strategy; it prices according to DF or LPM rules depending on its relative rank in the market. Therefore, over certain periods of time, it is easy for its price patterns to be confused with those that come from pure DF or LPM sellers. If the low-price seller in a market is unchanged over some period of time, a stream of prices from a LPM seller in that market are difficult to discern from whatever strategy is being matched—the price histories are identical, except for being out of step with one another by one period. Further work on NN inputs could lead to identification of data that would allow the NN to recognize when this situation is happening and classify the lagging strategy as LPM.

Strategy	NN Accuracy
Varian	92.2%
DF	71.2%
Trigger	89.7%
LPM	90.5%
BHM	95.5%

Table 6-1: Accuracy of NN classification for each seller strategy

Additionally, overall accuracy may be enhanced by using a NN tool that has broader capabilities than Tiberius. As mentioned above, Tiberius is restricted to a single output, but a multi-output design may prove more effective for this classification problem. Tiberius is also restricted to a single hidden layer. Modifications to allow hidden layers were not performed in order to keep this initial study of NNs as strategy classifiers simple.

After NN construction, the simulation tool was used to generate data to be analyzed for population of the decision table. Limiting the seller strategy mixes to avoid uninteresting markets resulted in 35 four-seller combinations of interest. For each of these 35 strategy combinations, the simulation was used to generate 500 iterations of data for each of the five possible counterstrategies. The base parameterization (with the exception of iteration length) from Chapter Five was used along with the same homogeneous mix of buy types. Table 6-2 shows the parameterization used for all of Chapter Six.

Parameter description	Value/Distribution
Top end of price range	2100
Bottom end of price range	1900
Number of sellers	5
Number of buyers	1000
Number of time periods per iteration	30 (days/month)
Rate of repricing for Varian sellers.	0.0667 (re-pricing an average of twice per month)
Rate of shopping for non-agent buyers.	0.5 (shopping an average of once every two days)
Unit cost to seller. Set equal to bottom end of price range.	1900
Initial buyer valuation	Normal with ±2σ in price range
Buyer horizon	Normal centered on 7.5, with ±2σ going from 0 to 15 and truncated at 1.
Memory length	7 (days)

Table 6-2: Base parameterization for Chapter Six simulations

Profit levels for each of the counterstrategies were measured and compared. For each market scenario, the counterstrategy with the highest average profitability across the 500 iterations was selected to populate the recommended strategy for that scenario in the decision table. Thus, the final decision table consisted of 35 rows and six columns. The 35 rows represent each of the competitor strategy mixes under consideration. Five of the six columns are the numbers of each of potential competitor strategy observed in the market. The sixth column is the recommended counterstrategy based on the simulations described above. Table 6-3 shows some example rows of the decision table. Appendix A shows the full table. Table 6-4 shows the total numbers of recommendations earned by each of the pricing strategies. The Varian pricing model was found to be most profitable for one of the market scenarios. In that situation, however, the second-place finisher (DF), was chosen as the recommended strategy. This choice was made in consideration of how a pricebot would be viewed by a potential customer if it recommended a "strategy" of randomly choosing prices. Additionally, the margin between Varian and DF was very small (<0.5%) and the two were in the same Fisher's LSD group using $\alpha = 0.05$.

# of Varians	# of DFs	# of LPMs	# of Triggers	# of BHMs	Recommend
2	0	1	0	1	DF
1	1	2	0	0	LPM
0	3	0	1	0	Trigger
2	0	0	1	1	BHM

Table 6-3: Sample rows taken from the seller strategy recommendation table

Strategy	# of Recommendations
DF	14
Trigger	14
LPM	5
BHM	2

Table 6-4: Frequency of recommendation for each pricing strategy studied

With the NNs constructed and the decision table populated, the final portion of the Chapter Six study involved an examination of the efficacy of the NN/decision table combination. The simulation tool was used to generate 1000 more 30-period iterations, with the mix of sellers randomized, restricted to the sets of strategy combinations as defined above. The profitability of the "fifth" seller was recorded for each iteration. End conditions (seller strategies, seller prices, buyer valuations, buyer thresholds, etc.) were also captured to serve as starting conditions for later simulations. Price data was also captured for all sellers over the last 15 periods of each iteration. The NN system was then applied to each of sellers 1-4 for each of the 1000 iterations to classify pricing strategies. The NNs categorized all four of the sellers of an iteration correctly 43.0% of the time. Exactly three sellers were properly classified 42.9% of the time, meaning that three or more sellers were classified correctly 85.9% of the time. Exactly two and exactly one correct classifications occurred 12.9% and 1.2% of the time, respectively. Therefore, at least half of the sellers were properly classified 98.8% of the time. There were no occurrences of zero correct classifications. While 43.0% total accuracy seems low, it is not a surprising result considering the classification scenario being studied. Whenever n independent classifications are done, the probability of all n classifications being correct is p^n, where p is the probability of a single correct classification. Thus, classification errors compound, much as type I error compounds when multiple hypotheses are tested.

The strategy mixes that resulted from NN classification were then cross-referenced with the decision table of Appendix A to obtain a recommended counter-strategy. Because the NN system was less than 100% accurate, situations occurred where the NN classification of competitor strategies did not match any of the rows of the decision table. For these situations, the rules were examined for any patterns exhibited. Because DF and Trigger were easily the most frequently recommended strategies, it seemed reasonable to assign one of those for each of the unmatched NN classifications. Analysis of the decision rules led to the adoption of the following rules for scenarios not listed in the table:

1. Count the total number of LPM and Trigger sellers in the NN classification.
2. If the count from #1 is zero, the recommended strategy is Trigger. If it is 2 or higher, DF is recommended.
3. If the count from #1 is 1, count only the number of LPMs.
4. If the count from #3 is zero, DF is recommended. If it is one or higher, Trigger is recommended.

After complete determination of the 1000 recommended counterstrategies, the process of determining whether the recommended strategy outperforms a seller's current strategy choice was undertaken. The captured end conditions from the 1000 previous iterations were used as initial conditions for two sets of 1000 new iterations to embody two treatments: one representing the fifth seller continuing on with its previous strategy and one representing the fifth seller switching to the strategy recommended by the NN/decision table. For each of these 2 sets of runs, sellers one through four retained the same strategies from the previous runs. These runs used the same base parameterization described in Table 6-2, including the iteration length of 30. Profit levels for these 30 periods were captured for each set of runs. The differences between the old strategy and recommended strategy could then be analyzed by using a paired t-test (one-tailed, assuming recommended strategy performance exceeds continued use of old strategy). Results were examined at various cumulative iteration levels and are summarized in Table 6-5. Finally, a third set of 1000 runs was performed to gauge the ultimate potential of a pricebot enabled with the NN and decision capabilities outlined in this chapter. This third treatment was the same as the previous two, with the exception that the fifth seller strategy was set to whatever the recommendation of the NN/decision table system would have been with a perfectly accurate NN.

	Actual NN		Perfect NN	
Cumulative Iterations	Profitability Difference	p-value	Profitability Difference	p-value
200	$438.96	0.442	$4368.98	0.099
400	$2040.37	0.155	$4569.34	0.026
600	$1657.10	0.158	$3827.13	0.020
800	$1928.97	0.087	$5248.15	0.001
1000	$1952.86	0.068	$4547.26	0.001

Table 6-5: Results from paired t-tests comparing profitability levels with original strategy vs. strategy recommended by the NN actually implemented and a theoretically perfect NN

Table 6-5 illustrates the potential of a pricebot that is able to classify its competitors' pricing behavior and react appropriately. At 1000 iterations, even a NN that accurately classifies strategies at less than 90% shows a marginally significant improvement due to the change to recommended counterstrategy. While 1000 iterations were required to obtain a p-value that is indicative of significance, that significance is not merely an artificial result achieved through large sample size. The price difference observed is 2.4% of the total profitability level for the 30 periods observed. Given the very thin profit margins observed in many industries, most managers would be very interested in a tool that would permit them to increase profitability by 2.4%. The profitability difference data also indicates a trend toward stabilization. The relative changes in profitability difference from one iteration level to the next consistently decrease as iteration level increases.

Assuming some improvement in NN accuracy is possible with further work on design, that profitability increase may be even higher with an enhanced NN system. To get a sense for the potential improvement available, data assuming a 100% accurate NN was examined. At 1000 iterations, the pricebot with a theoretically perfect NN earns a clearly significant increase in average profitability of 5.6%—again, a profitability gain that is undoubtedly practically important to managers in addition to being statistically significant. The perfect NN treatment also leads to p-values below 0.05 at all iteration levels above 200. It is logical to assume that further development of a strategy classifying NN would push the performance of the actual NN closer to the theoretically perfect level. For these 1000 iterations studied, the recommendation of the actual NN/decision table combination made the "right" recommendation (as defined by the perfect NN) about 67.7% of the time. Given that performance gap, potential for performance improvement is promising.

In conclusion, the research of this chapter clearly indicates that a system combining neural network classification and decision rules based on previously-run simulations can be an effective tool for pricebot implementation. The system was able to achieve profitability improvements of a magnitude that was marginally significant statistically and practically important to managers. It was also shown that potential for enhanced performance exists through improvement of NN classification accuracy. While the perfect NN analyzed above will not be achieved, further study could uncover ways to close the gap between actual and theoretical maximum performance. Enhancements to the work done here could include both the use of a more sophisticated NN tool and incorporation of new NN inputs. A more sophisticated NN tool could allow multiple outputs as well as multiple hidden layers, increasing the range of design options available for experimentation. Incorporation of new NN inputs could be helpful if inputs can be identified that are effective at distinguishing different strategies when their price patterns are very similar to another—such as when a LPM pricebot is matching some other strategy over period of time and ends up looking like and being classified as the strategy it is matching rather than as LPM. For example, a count of how many times a seller's price is equal to a competitor's prior period price could be indicative of LPM and enhance NN accuracy. Both pricing managers and potential sellers of pricebots or pricebot services should be very interested in further examining the type of intelligent agent system studied in this chapter.

Chapter Seven: Summary

E-commerce markets have begun to transition from curiosity to ubiquity. Online shopping and buying is rapidly gaining popularity, but still represents a very small portion of total commerce. As the magnitude of Internet-based commerce increases, the tools and techniques utilized by players in e-commerce markets will mature. Intelligent agents working on behalf of buyers and sellers will likely be a part of that evolution. Shopbots work on behalf of buyers to find and purchase the right goods at the best prices. Pricebots assist sellers by monitoring the market and implementing pricing schemes that are geared toward maximizing profitability. This work has simulated commodity-type e-commerce markets where sellers post prices that can be updated periodically. The data collected and analyzed has led to conclusions about what techniques work best and what factors should be considered by buyers, sellers, and agent builders in e-commerce environments. From this study, there are lessons to be learned for e-commerce buyers and sellers as well as third parties who might provide agent products or services.

Chapter Four provides insight for sellers who are considering introducing strategic pricebot pricing into their e-commerce markets. The conclusions from Chapter Four are multi-faceted. Firstly, more often than not, benefits were realized from switching to a number of different strategies from a non-strategic pricing model, where prices were merely selected from a probability distribution to model non-pricebot pricing behavior. However, there were occurrences where pricing “strategically” was outperformed by continuing to price with the non-strategic pricing policy. For a manager trying to choose the best pricing strategy, it is important to understand that there is likely not a single best strategy. The simulations in Chapter Four showed that the best strategy has a strong dependence on the nature of the marketplace—specifically, the extent to which the buying population comparison shops with shopbots. The most profitable strategy at one shopbot usage level may turn out to be the worst at another level. Chapter Four also examined a scenario where all competitors in what was previously a non-pricebot market switch to what should be the most profitable strategy. For the scenario studied, it was shown that a more profitable counter-move was then available to a seller willing to switch pricing policies again. Finally, it was shown that the relative standing of a seller in the market plays a role in determining the relative performance of pricing strategies. When there is a preferred seller in a market, there are subtle changes in the way profits from varying strategies compare. Overall, Chapter Four shows that pricing strategy selection is a complicated venture that requires that managers understand the complexion of the market in which they sell.

Chapter Five studies choices available for the implementation of buyers' shopbots. While the focus of Chapter Five was intended to be on shopbots, there were some interesting findings with respect to pricebots as well. Chapter Five underscores the importance of considering the makeup of a market before choosing a pricebot strategy. Changing the buyer complexion from a homogeneous group of simple buyers to a heterogeneous mix of shopbot implementations introduced changes in the relative performance of pricing strategies when the Chapter Four study was repeated. With respect to shopbot performance, Chapter Five shows that certain implementation choices can make a difference in the price a buyer pays in certain markets. Adding a memory component to a simple comparison shopping agent generally decreased the price that buyers ultimately paid. The discount rate used to control how a remembered price is treated was also shown to have an impact on price paid in some markets. The differences in price paid observed in Chapter Five were generally small—which for a one-time retail consumer might seem insignificant. However, in business-to-business scenarios where bulk purchases are made repeatedly, the potential important of subtle price differences is heightened.

While Chapters Four and Five highlight important factors to be considered for agent implementation, the techniques employed were simple from an intelligence standpoint. Chapter Six introduces a more sophisticated tool with the application of neural networks (NNs). Chapters Four and Five showed that different strategies can perform well under different conditions. Therefore, it is likely that a wide variety of price strategy complexions could evolve in a market based on the assumptions made by sellers' pricing managers. Chapter Six studies a broad range of those markets and shows that a NN-enabled system can be an effective tool for sellers to employ. A pricebot can use a NN to classify the behavior of competitors in a market. Based on that knowledge, a counterstrategy can be selected from a decision matrix that was previously populated with results from simulations similar to those performed in Chapter Four. Chapter Six shows that a NN-system can lead to improved profitability by prompting sellers to change their strategies. Examination of a hypothetical, perfect NN classification environment shows that there is significant potential for improvement beyond the actual NN employed in Chapter Six.

This work has led to many important conclusions described above. Potential also exists to study many more aspects of this same general area. All of the following would be valuable extensions to the work already completed herein:

- Other pricing strategies could be examined. This work focused on a few previously studied strategies as well as a strategy not present in the literature—there are likely other approaches that are worthy of study.
- Non-commodity markets could be simulated. Introducing factors other than price for buyers to consider would certainly lead to multiple interesting research paths.
- Further research into the application of artificial intelligence is definitely warranted. The neural network approach detailed in Chapter Six could be refined. Other applications of neural networks such as price prediction could be investigated. Other artificial intelligence techniques could be employed to work on behalf of buyers and sellers as well.
- Specific methods of agent implementation need to be studied. This work assumes that the technology exists to do the work described, but does not attempt to actually implement the techniques used.
- Other types of markets could be investigated. Auctions, reverse auctions, and markets where agents negotiate with each other are all worthy endeavors.

Continued work along these paths will help buyers, sellers, and agent builders be successful in their respective roles within e-commerce markets as they continue to gain popularity and evolve with respect to their use of tools and technology.

Appendix A: Seller Strategy Recommendation Table

# of Varians	# of DFs	# of LPMs	# of Triggers	# of BHMs	Recommend
0	2	0	0	2	Trigger
0	2	0	1	1	Trigger
0	2	0	2	0	DF
0	2	1	0	1	DF
0	2	1	1	0	DF
0	2	2	0	0	LPM
0	3	0	0	1	Trigger
0	3	0	1	0	Trigger
0	3	1	0	0	Trigger
0	4	0	0	0	Trigger
1	1	0	0	2	Trigger
1	1	0	1	1	Trigger
1	1	0	2	0	DF
1	1	1	0	1	DF
1	1	1	1	0	DF
1	1	2	0	0	LPM
1	2	0	0	1	Trigger
1	2	0	1	0	Trigger
1	2	1	0	0	DF
1	3	0	0	0	Trigger
2	0	0	0	2	LPM
2	0	0	2	0	DF

# of Varians	# of DFs	# of LPMs	# of Triggers	# of BHMs	Recommend
2	0	0	1	1	BHM
2	0	1	0	1	DF
2	0	1	1	0	DF
2	0	2	0	0	LPM
2	1	0	0	1	Trigger
2	1	0	1	0	DF
2	1	1	0	0	DF
2	2	0	0	0	Trigger
3	0	0	0	1	LPM
3	0	0	1	0	BHM
3	0	1	0	0	DF
3	1	0	0	0	Trigger
4	0	0	0	0	DF

Sources

Agarwal, A., Davis, J.T., and Ward, T., "Supporting Ordinal Four-State Classification Decisions Using Neural Networks," *Information Technology and Management,* Vol. 2, No. 1, Jan 2001, pp. 5-26.

Anderson, R.D., Engledow, J.L., and Becker, H., "Evaluating the Relationships Among Attitude Toward Business, Product Satisfaction, Experience, and Search Effort," *Journal of Marketing Research,* Vol. 16, Aug 1979, pp. 394-400.

Bakos, J.Y., "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science,* Vol. 43, No. 12, Dec 1997, pp. 1676-1692.

Berardi, V.L., Patuwo, B.E., and Hu, M.Y., "A Principled Approach for Building and Evaluating Neural Network Classification Models," *Decision Support Systems,* Vol 38, 2004, pp. 233-246.

Bucklin, L.P., "Testing Propensities to Shop," *Journal of Marketing,* Vol. 30, Jan 1966, pp. 22-27.

Clark, D., "Shopbots Become Agents for Business Change," *(IEEE) Computer,* Vol. 33, No. 2, Feb. 2000, pp. 18-21.

Cortese, A.E., "Good-bye to Fixed Pricing?" *Business Week,* No. 3576, May 4, 1998, p. 70.

Dasgupta, P. and Melliar-Smith, P., "Dynamic Consumer Profiling and Tiered Pricing Using Software Agents," *Electronic Commerce Research.* Vol. 3, 2003, pp. 277-296.

Deck, C.A. and Wilson, B.J., "The Effectiveness of Low Price Matching in Mitigating the Competitive Pressure of Low Friction Electronic Markets," *Electronic Commerce Research,* Vol. 2, No. 4, Nov 2002, pp. 385-398.

Deck, C.A., and Wilson, B.J., "Automated Pricing Rules in Electronic Posted Offer Markets," *Economic Inquiry,* Vol. 41, No. 2, April 2003, pp. 208-223.

DiMicco, J., Maes, P., and Greenwald, A. "Learning Curve: A Simulation-Based Approach to Dynamic Pricing," *Electronic Commerce Research.* Vol. 3, 2003, pp. 245-276.

Hertweck, B.M., Rakes, T.R., Rees, L.P., "Dynamic Pricing in e-Commerce Markets: Strategic Pricebots and Memory-Enabled Shopbots," 2003 Decision Sciences Annual Meeting Proceedings.

Kannan, P. and Kopalle, P., "Dynamic Pricing on the Internet: Importance and Implications for Consumer Behavior," *International Journal of Electronic Commerce.* Vol. 5, No. 3, Spring 2001, pp. 63-83.

Kauffman, R. and Walden, E., "Economics and Electronic Commerce: Survey and Directions for Research, *International Journal of Electronic Commerce.* Vol. 5, No. 4, Summer 2001, pp. 5-116.

Kephart, J.O. and Greenwald, A.R., "Shopbot Economics," *Autonomous Agents and Multi-Agent Systems,* Vol. 5, 2002, pp. 255-287.

Kephart, J.O., Hanson, J.E., & Greenwald, A.R. (2000). Dynamic Pricing by Software Agents. *Computer Networks,* 32, 731-752.

Kiel, G.C. and Layton, R.A., "Dimensions of Consumer Information Seeking Behavior," Journal of Marketing Research, Vol. 18, May 1981, pp. 233-239.

Kutschinski, E., Uthmann, T., and Polani, D., "Learning Competitive Pricing Strategies by Multi-agent Reinforcement Learning," *Journal of Economic Dynamics and Control,* Vol. 27, No. 11-12, 2003, pp. 2207-2218.

Law, A.M. and Kelton, W.D. (2000), *Simulation Modeling and Analysis*, 3rd Ed., New York: McGraw-Hill.

Leloup, B., "Pricing with Local Interactions on Agent-based Electronic Marketplaces," *Electronic Commerce Research and Applications,* Vol. 2, No. 2, Summer 2003, pp. 187-198.

Meredith, J.R. and Shafer, S.M. (2002), *Operations Management for MBAs,* John Wiley and Sons: New York, pp. 296-298.

Moorthy, S., Ratchford, B.T., and Talukdar, D., "Consumer Information Search Revisited: Theory and Empirical Analysis," *The Journal of Consumer Research,* Vol. 23, No. 4, Mar 1997, pp. 263-277.

Moukas, A., Zacharia, G., Guttman, R., and Maes, P., "Agent-Mediated Electronic Commerce: An MIT Media Laboratory Perspective," *International Journal of Electronic Commerce.* Vol. 4, No. 3, Spring 2000, p. 5-21.

Nath, R., Rajagopalan, B., and Ryker, R., "Determining the Saliency of Input Variables in Neural Network Classifiers," *Computers and Operations Research,* Vol. 24, No. 8, 1997, pp. 767-773.

Newman, J.W. and Staelin, R., "Prepurchase Information Seeking for New Cars and Major Household Appliances," *Journal of Marketing Research,* Vol. 9, Aug 1972, pp. 249-257.

Partovi, F.Y. and Anandarajan, M., "Classifying Inventory Using an Artificial Neural Network Approach," *Computers and Industrial Engineering,* Vol. 41, 2002, pp. 389-404.

Pedersen, P.E., "Behavioral Effects of Using Software Agents for Product and Merchant Brokering," *International Journal of Electronic Commerce,* Vol. 5, No. 1, Fall 2000, pp. 125-141.

Piramuthu, S., Shaw, M.J., and Gentry, J.A., "A Classification Approach Using Multi-layered Neural Networks," *Decision Support Systems,* Vol. 11, No. 5, pp. 509-525.

Punj, G.N. and Staelin, R., "A Model of Consumer Information Search Behavior for New Automobiles," *Journal of Consumer Research,* Vol. 9, Mar 1983, pp. 366-380.

Saad, E.W., Prokhorov, D.V, and Wunsch, D.C. (1998), "Comparative Study of Stock Trend Prediction Using Time Delay, Recurrent and Probabilistic Neural Networks," *IEEE Transactoins on Neural Networks,* Vol. 9, No. 6, pp. 1456-1470.

Smith, M.D., "The Impact of Shopbots on Electronic Markets," *Academy of Marketing Science Journal,* Vol. 30, Iss. 4, Fall 2002, pp. 446-454.

Tassier, T., Everson, M.P, & Ostrowski, D. (2002). Agent-Based Models as a Complement to Economic Theory: A Durable Goods Example. *Proceedings of the 2002 Congress on Evolutionary Computation*, May 2002, Vol. 1, 729-734.

Udell, J.G., "Prepurchase Behavior of Buyers of Small Electrical Appliances," *Journal of Marketing,* Vol. 30, Oct 1966, pp. 50-52.

Urbany, J.E., Dickson, P.R., and Wilkie, W.L., "Buyer Uncertainty and Information Search," *Journal of Consumer Research,* Vol. 16, Sep 1989, pp. 208-215.

Varian, H., "A Model of Sales," *The American Economic Review,* Vol. 70, No. 4, Sep 1980, pp. 651-659

Zacharia, G, Evgeniou, T, and Maes, P., "Dynamic Pricing in a Reputation-Brokered Agent-Mediated Marketplace," *International Journal of Intelligent Systems in Accounting, Finance, and Management.* Vol. 9, No. 4, Dec 2000, pp. 271-286.

Zhang, G., Patuwo, B.E., and Hu, M.Y. (1998), "Forecasting with Artificial Neural Networks: The State of the Art," *International Journal of Forecasting,* Vol. 14, pp. 35-62.

http://www.census.gov/mrts/www/data/html/05Q2.html, accessed 9/2005

http://dictionary.reference.com/search?q=intelligent, accessed 9/2005

www.philbrierley.com, accessed 10/2004

BRYAN M. HERTWECK

Assistant Professor
Department of Information Systems
Earl N. Phillips School of Business
High Point University

833 Montlieu Ave.	Office: 336-841-4603
High Point University	Cell: 336-906-7272
High Point, NC 27262	bhertwec@highpoint.edu

EDUCATION

Ph.D	2005	Business, Management Science, Virginia Polytechnic Institute and State University (Virginia Tech) ***Dissertation Title:*** Examining Electronic Markets in which Intelligent Agents are Used for Comparison Shopping and Dynamic Pricing
M.S.	1995	Industrial and Systems Engineering, Virginia Tech ***Concentration:*** Management Systems Engineering ***Thesis Title:*** The Effects of Authority on Group Decision-Making Performance
B.S.	1993	Aerospace Engineering, Virginia Tech ***Minor:*** Mathematics

PUBLICATIONS

Kleiner, B.M. and **Hertweck, B.,** "By Which Method?: Total Quality Management, Reengineering, or Deengineering," *Engineering Management Journal*, Vol. 8, No. 2, June 1996, pp. 13-18.

REFEREED CONFERENCE PROCEEDINGS

Hertweck, B.M., Rakes, T.R., Rees, L.P., "Dynamic Pricing in e-Commerce Markets: Strategic Pricebots and Memory-Enabled Shopbots," 2003 Decision Sciences Annual Meeting Proceedings, Washington, DC.

CONFERENCE PRESENTATIONS

Hertweck, B.M., Rakes, T.R., Rees, L.P., "The Impact of Comparison Shopping Behavior on Intelligent Agent Pricing Strategy Selection," to be presented at the 2004 Decision Sciences Annual Meeting, Boston, MA.

TEACHING EXPERIENCE

Fall 2005	Professor, MIS 231, Database I; CIS 203; E-Commerce Development; and CIS 430: Database III
Summer 2005	Instructor, BIT 2405: Quantitative Methods I
Summer 2003, 2004, and Spring 2005	Instructor, BIT 2406: Quantitative Methods II
Spring 1995	Co-instructor, EF 1004: Technology and Productivity

WORK EXPERIENCE

2001-Present
Graduate Teaching Assistant, Virginia Tech, Blacksburg, VA
- Supported POM, DSS, Networking, and Client/Server Development courses

Summer 2002
Research Assistant, Virginia Tech, Blacksburg, VA
- Supported research into use of optimization techniques for commodity and future trading for Dominion Power, Richmond, VA by writing and executing VBA simulations

2000-2001
Business Analyst, Sentraliant, Richmond, VA
- Supported development, integration, and testing of customer care and billing system for Sirius Satellite Radio

2000-2000
SAP Consultant, Management Resources, Inc., Richmond, VA
- Managed SAP implementation for the Oregon Liquor Control Commission
- Managed post-implementation SAP maintenance for Wyoming Liquor Commission

1999-2000
SAP Consultant, Clarkston Consulting Group, Durham, NC
- Supported SAP implementations at multiple clients, including Lockheed-Martin Enterprise Information Systems
- Designed and tested SAP interface for eCredit.com credit and collections software

1996-1999
Team Leader, ERP Services, RWD Technologies, Columbia, MD
- Provided post-implementation SAP configuration and training at Boeing North American Services
- Developed and delivered SAP training materials for multiple clients, including Microsoft, Bristol-Myers Squibb, and Proctor & Gamble
- Wrote and presented multiple sales proposals for SAP implementation support

CERTIFICATIONS

- Certified SAP R/3 Financial Accounting (FI) and Controlling (CO) Consultant
- Certified SAP R/3 Public Sector/Funds Management (FM) Consultant
- Certified AcceleratedSAP (ASAP) Consultant
- American Production and Inventory Control Society (APICS) Certification in Production and Inventory Management (CPIM)

SPECIAL AWARDS and HONORS

R.B. Pamplin Doctoral Fellowship
Gilbert and Lucille Seay Undergraduate Engineering Scholorship
National Merit Scholarship