MARKETING
SCIENCE
INSTITUTE

Electronic Shopping for Wine: How Search Costs Affect Consumer Price Sensitivity, Satisfaction with Merchandise, and Retention

John G. Lynch, Jr.
Dan Ariely

Working Paper
Report No. 99-104
1999

MSI

This research was supported by a grant from the Marketing Science Institute. The authors are grateful for comments on design issues from Joel Huber and for comments on earlier drafts of the paper by Jim Bettman, Bill Boulding, and seminar participants at Cornell University, Hong Kong University of Science and Technology, the University of Michigan, the INFORMS miniconference on Marketing Science and the Internet, and the Fuqua Marketing Workshop. We also thank Lisa Abendroth, Dave Arnold, Matt Bailey, and Chris Hollander for assistance in conducting the experiment, Craig Heffley for providing wine descriptions and ratings, and Susan Keane and Michele Rainforth for editorial and production assistance.

Electronic Shopping for Wine: How Search Costs Affect Consumer Price Sensitivity, Satisfaction with Merchandise, and Retention

John G. Lynch, Jr., and Dan Ariely

Although the potential for interactive home shopping (IHS) to reduce search costs is attractive to consumers, retailers fear that electronic retailing will intensify competition and lower margins. In this study, authors John Lynch and Dan Ariely carry out an experiment to see if this concern is warranted. They investigate whether a potential increase in price sensitivity accompanying easier access to price comparisons is offset by a decrease in price sensitivity brought about by easier access to nonprice information that differentiates products.

They hypothesize that maximally transparent IHS systems (those in which price and quality information are easily accessible and across-store price comparisons are made easy) increase consumers' welfare (i.e., make it easier for them to choose products that they will like), and thus increase customer retention.

The Experiment

In an experiment using on-line wine merchants, the authors varied the ease with which consumers were able to access price information and quality information, and the ease with which they were able to make cross-store comparisons. Consumers spent their own money to purchase wines from two competing electronic merchants selling some overlapping and some unique wines.

Results and Implications

The experiment revealed that making price information more accessible did not necessarily have ruinous effects on price sensitivity. Price sensitivity:

- decreased as the usability of quality information increased,
- was unaffected when price information was made more usable, and
- increased when store comparison was made easier, but only for overlapping wines, not for unique wines.

Increasing the transparency of the information environment also increased consumer welfare: the shopping task was more enjoyable and a subsequent taste test proved that consumers were able to purchase wines they liked. Transparency increased retention when consumers were given an opportunity to resubscribe to the same electronic wine-buying services.

In conclusion, it appears that the challenge facing electronic retailers is not how to prevent comparison but how to use emerging electronic venues to provide consumers with better information about product quality. When they cooperate with comparison-shopping engines, it is in retailers' interests to make the informational base include richer differentiating information.

John Lynch is Hanes Corporation Foundation Professor of Business Administration at the Fuqua School of Business at Duke University. Dan Ariely is Assistant Professor of Management Science at the Sloan School of Management at the Massachusetts Institute of Technology.

Contents

Introduction

Emerging electronic channels create a fundamental dilemma for retailers with stand-alone sites on the World Wide Web and for those attempting to build electronic malls for delivery via the Internet, on-line services, or interactive television. Alba et al. (1997) present the case that, for consumers, the main attraction of interactive electronic retailing is a reduction in search costs for products and product-related information. However, it is precisely this lowering of search costs that retailers fear most. Their concern is that electronic retailing will intensify competition and lower margins by expanding the scope of competition from local to national and international (Anders 1998; Gittins 1997; Kuttner 1998; Quelch and Klein 1996; Reeve 1998).

Retailers seem to view these emerging channels as inevitable but potentially lethal. They therefore configure their individual electronic stores such that it is difficult to compare their merchandise with that of other stores selling on the same channel. In addition, when third-party agents like Anderson Consulting's Bargain-Finder are created to facilitate cross-store electronic search, merchants attempt to block them from their sites (Bakos 1997; Pazgal and Vulcan 1998; Quick 1998b).

Our paper argues that retailers are overly fearful that electronic shopping will lower search costs for price information and heighten competition. Their fears drive them to create defensive, toe-in-the-water interactive offerings that offer consumers no real benefits. Defensive-minded retailers tend to underinvest; they create Web sites that are nothing more than inconvenient, slow catalogs with far fewer SKUs than their print catalogs and relatively little information about the few products offered. Retailers then point to the tepid sales over their interactive channels as evidence that the electronic shopping format itself has limited appeal for consumers, rather than seeing poor sales as a reflection of their defensive implementation of the electronic shopping concept.[1]

Alba et al. (1997) argue that consumers value highly the potential for electronic shopping to lower search costs by offering a third route. By supporting comparisons of merchandise sold by competing vendors, electronic shopping increases consumers' ability to choose merchandise that will maximize consumption utility (Haeubl and Trifts 1998). We therefore investigate the impact of interstore comparison on the benefits of interactive shopping systems and consumer price sensitivity. One might expect, a priori, that making interstore comparisons easier should increase price sensitivity for items carried by multiple comparable stores but not affect differentiated merchandise sold exclusively by one retailer. Moreover, if the stores differ in the information they offer, in their appearance, or in the benefits they provide, the additional information consumers will derive from interstore comparison should result in decreased price sensitivity.

Alba et al. (1997) maintained that efforts by electronic merchants to make cross-shopping difficult are doomed to failure. Such efforts remove much of what makes the electronic venue more appealing than other retail formats. Moreover, if consumers value the benefit of cross-shopping on line, someone will offer it (e.g., http://www.compare.net, http://www.personallogic.com, and http://www.zdnet.com) and consumers will demand it (Erlich and Fisher 1982). Consumers may choose to patronize those sellers who deliver this benefit, bypassing those who do not.

Our paper does not attempt to make any truly new conceptual points. Our purpose is to highlight the inconsistency between conventional wisdom in business and academic writings on electronic commerce and well-established principles drawn from consumer information-processing theory and from the economics of information. Our central argument is that retailers' fears may be misdirected (Alba et al. 1997), leading them to underestimate their incentives to invest in the kinds of electronic stores that might offer real benefits to consumers. Electronic shopping may indeed reduce the cost of search in ways that enlarge consumers' consideration sets and that make price comparisons easier. Ceteris paribus, if electronic shopping lowers the cost of acquiring price information, it should increase price sensitivity, just as price advertising does (Popkowski-Leszczyc and Rao 1990; cf. Boulding, Lee, and Staelin 1994).

At the same time, these media can, when used well, convey nonprice information related to quality more effectively than can conventional malls, catalogs, etc. The issues closely parallel those arising in Mitra and Lynch's (1995, 1996) work on the effects of advertising on brand price elasticity. Advertising can convey differentiating information that reduces consumers' price sensitivity (Kaul and Wittink 1995). So can interactive channels (Alba et al. 1997; Degeratu, Rangaswamy, and Wu 1998). If there are real differences among retailers in merchandising, assortment of complementary products, and service, interactive channels may be more effective than existing modes of retailing in conveying those points of differentiation. See Bakos (1997) for related arguments supporting this conjecture.

Our thesis is that the effect of electronic shopping on consumer price sensitivity is twofold. First, insofar as search costs for price information are lowered, consumers' price sensitivity will increase. But insofar as search costs for quality information that differentiates products are lowered, consumers' price sensitivity will decrease, and the latter effect may outweigh the former. Thus, if a site decreases search costs for price alone, price sensitivity will increase. However, if a site decreases search costs for both price and quality information, price sensitivity need not increase. Similarly, cross-store comparisons, argued by Alba et al. (1997) to be valuable to consumers, may not be a threat to retailers if their assortments are mostly nonoverlapping.

Second, by lowering search costs, full-featured electronic-shopping systems may help consumers better match heterogeneous brands to their personal tastes. If so, electronic shopping begins to provide informational benefits to consumers that might make them actually want to utilize it. Consumers may be more likely to reward full-featured, "transparent" electronic merchants with repeat business. Hoffman, Novak, and Chatterjee (1995) speculate that on-line merchants "will only be successful in the long run if they generate repeat traffic" and that "the repeat visit problem is partly a function of Web site design."

We report below an experiment in which consumers shop with their own money from two competing electronic wine merchants carrying partially overlapping inventories. We vary independently three components of search costs in electronic shopping: the ease with which consumers can access price information and quality information, and the ease with which they can make cross-store comparisons. We consider the current status quo for most Internet retailing sites to correspond to our condition in which it is difficult to access price and quality information and difficult to make store comparisons. Our condition in which search costs for price and quality are low and store comparisons are made easy approximates the kind of transparent electronic shopping system that consumers might see as having advantages over competing retail formats. Before conducting our experiment, our expectation was that despite the increased ease of cross-store comparison, price sensitivity would not be higher in the latter condition than the former. This, of course, is a matter of calibration. We expected that lowering search costs for price information would increase price sensitivity, and that lowering search costs for quality information would decrease price elasticity. We expected that making store comparisons easier would increase price sensitivity for wines carried by both stores but not for wines carried only by one of the stores.

After our consumers shopped electronically, they had the opportunity to taste a battery of wines, some of which they had purchased in the earlier shopping phase and some of which they had not. We expected that in the taste test, the chosen wines would do increasingly better than the unchosen wines as the transparency of the quality information available increased—i.e., as it became easier to access quality information and to make store comparisons. We expected also that consumers would enjoy their shopping experiences more in more transparent environments—i.e., when ease of accessing price, quality, and store-comparison information was greatest. Finally, we expected that more transparent environments would lead to greater retention for the electronic wine-shopping service over time.

Experiment

Method

Overview. Seventy-two M.B.A. and Ph.D. students and staff were recruited to participate in a test of an electronic shopping system described as being similar to Virtual Vineyards (www.virtualvin.com). Participants were told that wines would be sold at significant discounts relative to prices for the same wines from area merchants and that the researchers would contribute $5.00 to the M.B.A. Games charity fund-raiser for each participant who bought one or more bottles. Participants first shopped for wine from our two competing electronic wine merchants, Jubilee and Dionysus. A total of 100 wines were available. Each store sold 60 wines, 20 of which overlapped and 40 of which were unique to a store. Consumers went on a series of eight shopping trips. Across the eight trips, the prices of the different wines varied independently. This way we could measure price sensitivity at the individual subject level by measuring how the quantity of wine purchased depended on its price level.

We independently varied Price Usability (high or low) x Quality Usability (high or low) x Store Comparability (easy or hard) in a 2x2x2 between-subjects design. Participants were randomly assigned to one of the eight conditions. Primary dependent variables at this stage were price sensitivity for wines, computed separately for blocks of wines common to the two merchants and for wines unique to each (there was also an overall measure pooling all wines), and liking for the shopping experience. We also took measures of breadth and depth of search.

After performing the shopping part of the study, participants were asked to taste 10 of the wines available earlier to see which wines they actually preferred. We computed from each participant's ratings a measure of liking for wines chosen earlier and for wines not chosen in the shopping phase.

Two months later an e-mail announcement was sent to participants to ask if they would like to continue using the same electronic wine-shopping system for future purchases of wines from home. This measure of service subscription was taken as an indicator of retention of the service.

Procedure for Computer Shopping Task. Respondents reported to the M.B.A. computing lab at prearranged times. They were told that they would go on a series of eight independent shopping trips with the prices of the wines varying from trip to trip. Respondents were asked to buy as much wine as they normally might consume in a month. On each trip, respondents searched through the wines in either or both stores and purchased as much or as little wine as they deemed appropriate. Once they had finished purchasing wine for that month's shopping trip, they indicated so and started a new month's shopping trip. This continued until all eight trips were completed.

Respondents expected to—and actually did—use their own money to pay for the wines they selected. We were concerned, however, that our price sensitivity measures (Quantity Difference and Price Elasticity) would be less stable if respondents with real budget constraints purchased only a few wines. Therefore, respondents were told that they would take eight shopping trips but that they would actually purchase only the wines they chose on one of the trips, to be randomly determined at the end of the experiment. Because respondents could not know which trip would be chosen, they were told to treat each one as if it were the one selected for the real transaction. As a consequence, we were able to get eight times as many purchases as we would have if respondents were paying for what they selected on all the shopping trips. (We also avoided inventory effects.) All the wines sold were on discount four of the eight trips and were sold at regular price on the remaining four trips, though the discounted wines were not specially noted in any way. Respondents were also told that they were not obliged to purchase any wines at all—after all, they were spending their own money. However, all participants bought at least three wines. The shopping task took between 30 and 75 minutes.

At the end of the computer-shopping task, participants were asked to rate how enjoyable the shopping experience was for them. This response was given on a scale from 0 (not enjoyable at all) to 100 (very enjoyable). In addition, participants answered a battery of questions that were aimed at assessing their knowledge of wine. There were two types of questions, one that related to experience with wine (amount generally purchased, frequency purchased, price, etc.) and one that asked respondents to identify different varieties of wine as being red or white. From their answers, we were able to construct a measure of wine expertise that was unidimensional and that exhibited marginally acceptable reliability ($\alpha = .65$). We had anticipated that expertise might moderate the effects of Quality Usability. However, expertise had no effects and will not be discussed further.

Independent Variables. As mentioned earlier, the independent variables were: Price Usability (high or low), Quality Usability (high or low), and Store Comparability (easy or hard). For analyses of price sensitivity, there was a fourth, repeated factor of Unique vs. Common Wine Blocks. This factor was not relevant to other dependent variables.

When Price Usability was high, price information was displayed in the first-level list of available wines, with a tool available to permit sorting by price. When Price Usability was low, the initial list of wine names did not show their prices nor was a tool available to sort wines by prices; respondents had to click on a wine name to bring up a screen with its price.

Quality Usability was varied by a parallel manipulation. When Quality Usability was high, the first-level list of wine names displayed descriptions of the wines using differentiating sensory attributes. Wines at Dionysus were described in terms of complexity, acidity, body, and sweetness/dryness, using bar graphs patterned after those used by Virtual Vineyards (www.virtualvin.com). Wines at Jubilee were described in terms of body, sweetness/dryness, intensity, and tannin, with numerical values of 1 to 7 for each dimension. This difference in format was intended to mirror the real world, in which competing vendors are unlikely to make the same information available or to use common display forms. In addition, when Quality Usability was high, respondents could sort the wine by varietal (e.g., Chardonnay, Merlot, etc.). Finally, respondents in this condition could click or "drill down" to see further differentiating comments (e.g., "Fun red wine? Here it is! A very pleasing bardolino with cherry and grape flavors and an easygoing demeanor. It's soft, juicy, and even sports hints of complex flavors such as vanilla and jam. But don't be fooled; it's down to earth and fun.")[2] When Quality Usability was low, the standardized descriptions on sensory dimensions did not appear on the first screen containing the list of wines. Instead, participants had to click on a wine's name on the first screen to see them and no tool was available to sort wines by varietal. Furthermore, there was no ability to drill down to see a further differentiating comment.

Store Comparability was varied by the nature of the display subjects saw on first-level viewing screens. When Store Comparability was easy, the screen was divided in half with Dionysus on the left and Jubilee on the right. The navigation tools mentioned earlier were provided at the bottom of each store's display and the respondent could independently view and navigate both stores. The wine list in each half was displayed initially in alphabetical order, but any sorting tool available in one store would simultaneously sort the wines from both stores on the same criterion. When Store Comparability was hard, only one store appeared on the screen at a time. If the respondent was shopping at Dionysus, the righthand side of the screen for Jubilee was blank. If respondents were shopping at one store and then wanted to visit the other, their shopping carts emptied; they had to start again from scratch at the first store if they returned after visiting the second. Moreover, sorting tools used at one store had no effect on the ordering of wines displayed at the other store. Such interfaces again mimic current reality: merchants do not let consumers hold on to their purchase indication while they leave the store. Note that this is also the case in brick-and-mortar stores in which a consumer has to consolidate and transact a purchase in order to take merchandise from the store.

Figure 1 shows a thumbnail of the first and second screens in two conditions: when Price Usability and Quality Usability were high and Store Comparability was easy (on the left) and when Price Usability and Quality Usability were low and Store Comparability was hard (on the right). Moving among the different screen levels was done by clicking with a mouse on the tools that appear at the bottom of the screen. Figures 2 shows the actual first and second screens seen by subjects in the low Price Usability, low Quality Usability, and easy Store Comparability condition, and Figure 3 shows the first and second screens seen by subjects in the high Price Usability, high Quality Usability, and easy Store Comparability condition.

Figure 1. A Schematic Representation of the Different Electronic Shopping Interfaces

The actual screens used were significantly more detailed and are available from the authors. The left panel shows what participants saw when Price Usability and Quality Usability were high, and Store Comparability was easy. The right panel shows the condition where Price Usability and Quality Usability were low, and Store Comparability was hard. Within these two panels, the top panel shows the information that was available at the highest level of the interface (without any search cost), and the bottom panel shows the information that was available at the second level of the interface (with search cost). In the full-featured electronic shopping (left panel), all the information was presented at the highest level, and additional descriptive information was presented at the second level. In the impoverished status-quo electronic shopping (right panel), none of the information was presented at the highest level, and participants had to drill down to get information about price, quality, or the other store.

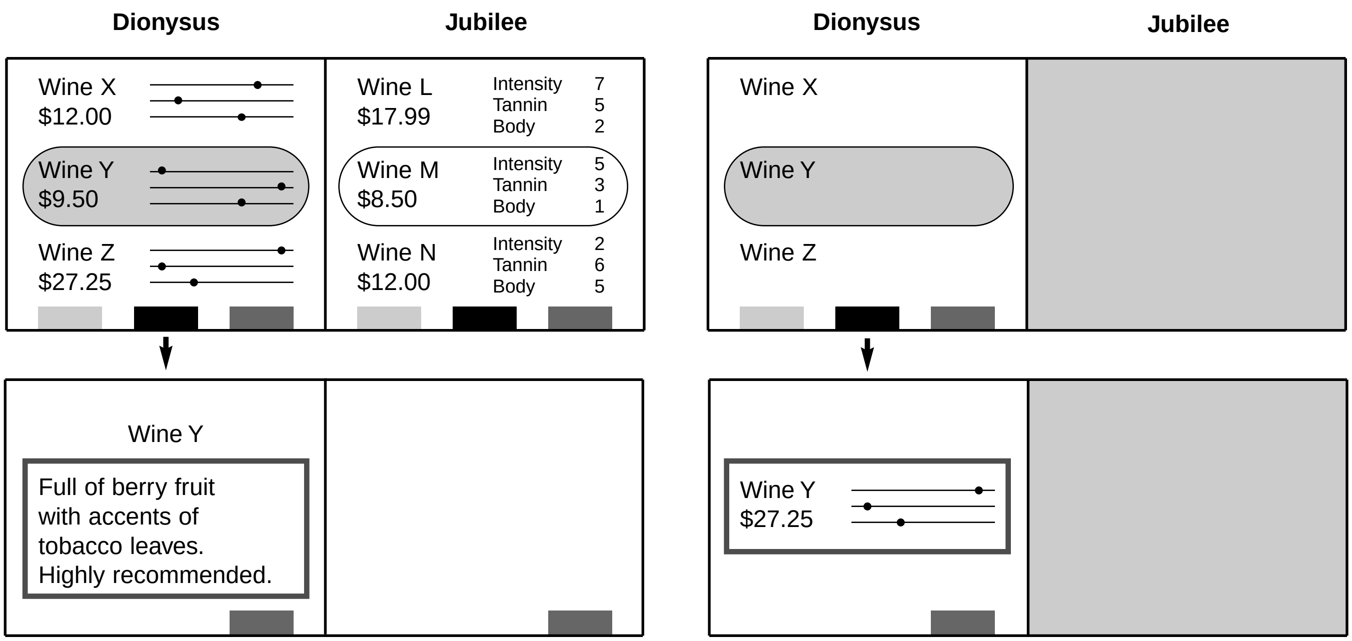

Figure 2. Screens 1 and 2 for Low Price Usability, Low Quality Usability, and Easy Store Comparability

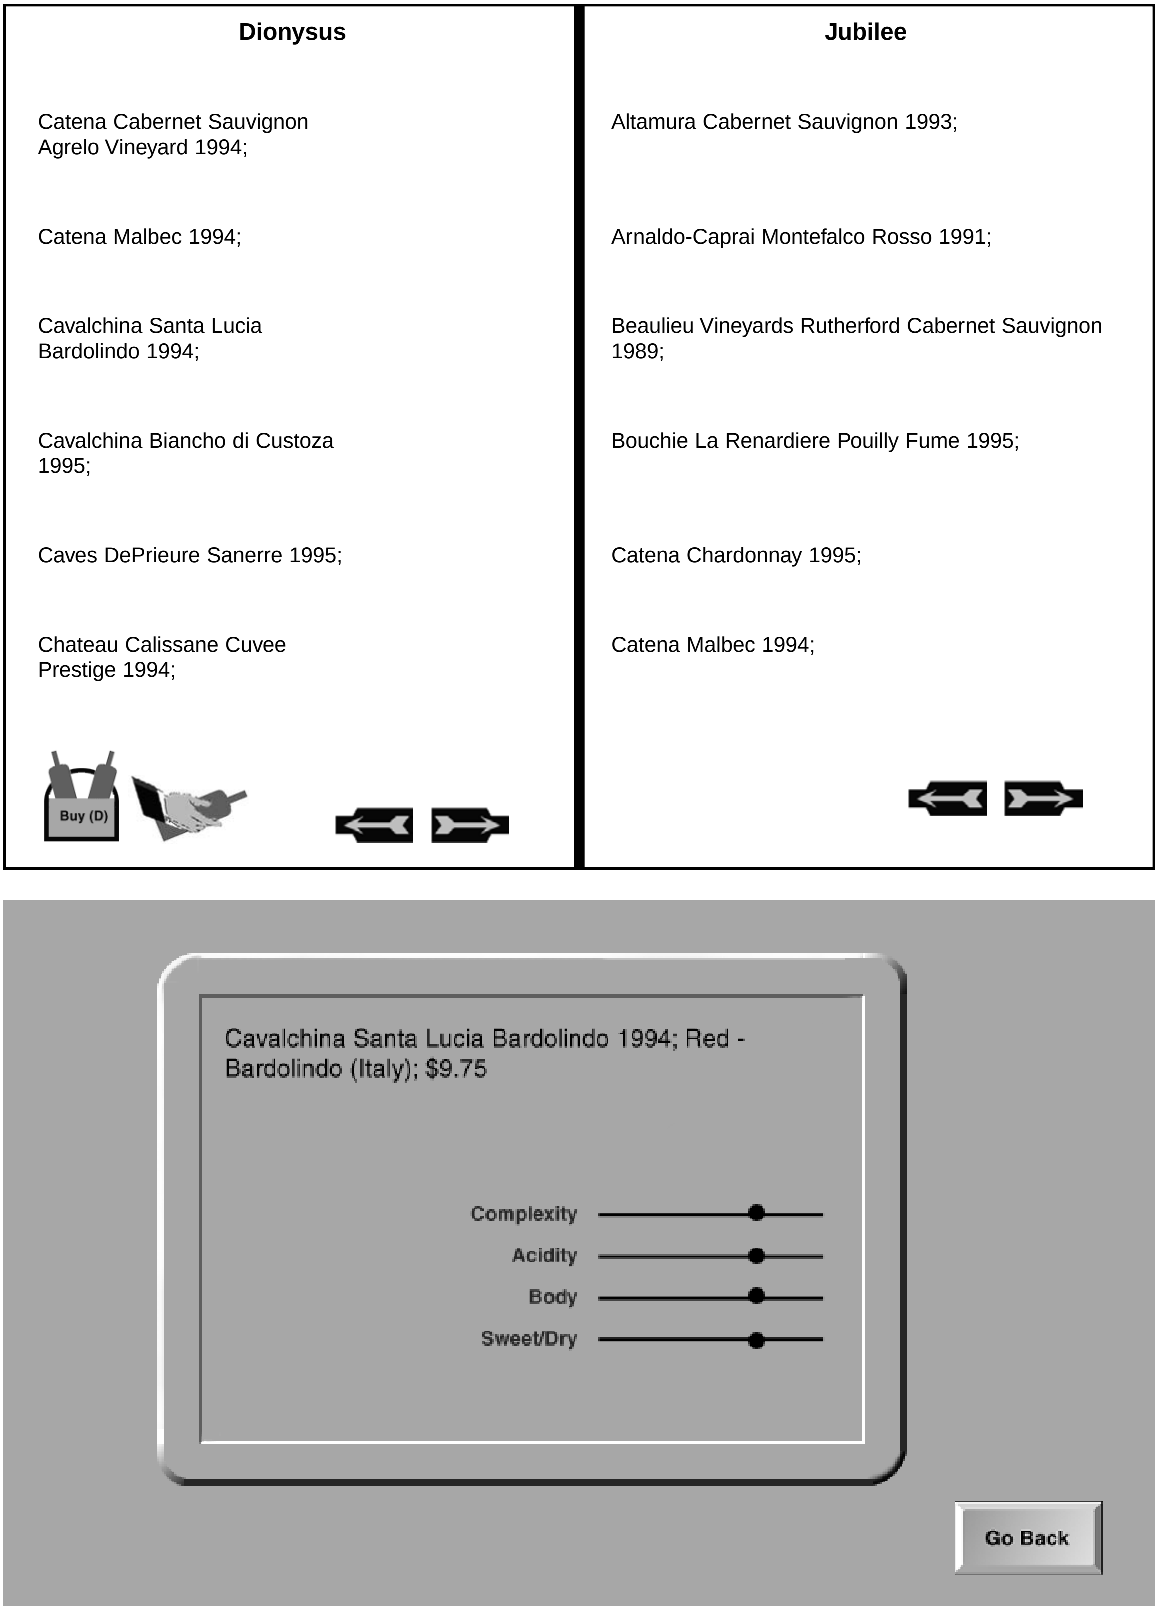

Figure 3. Screens 1 and 2 for High Price Usability, High Quality Usability, and Easy Store Comparability

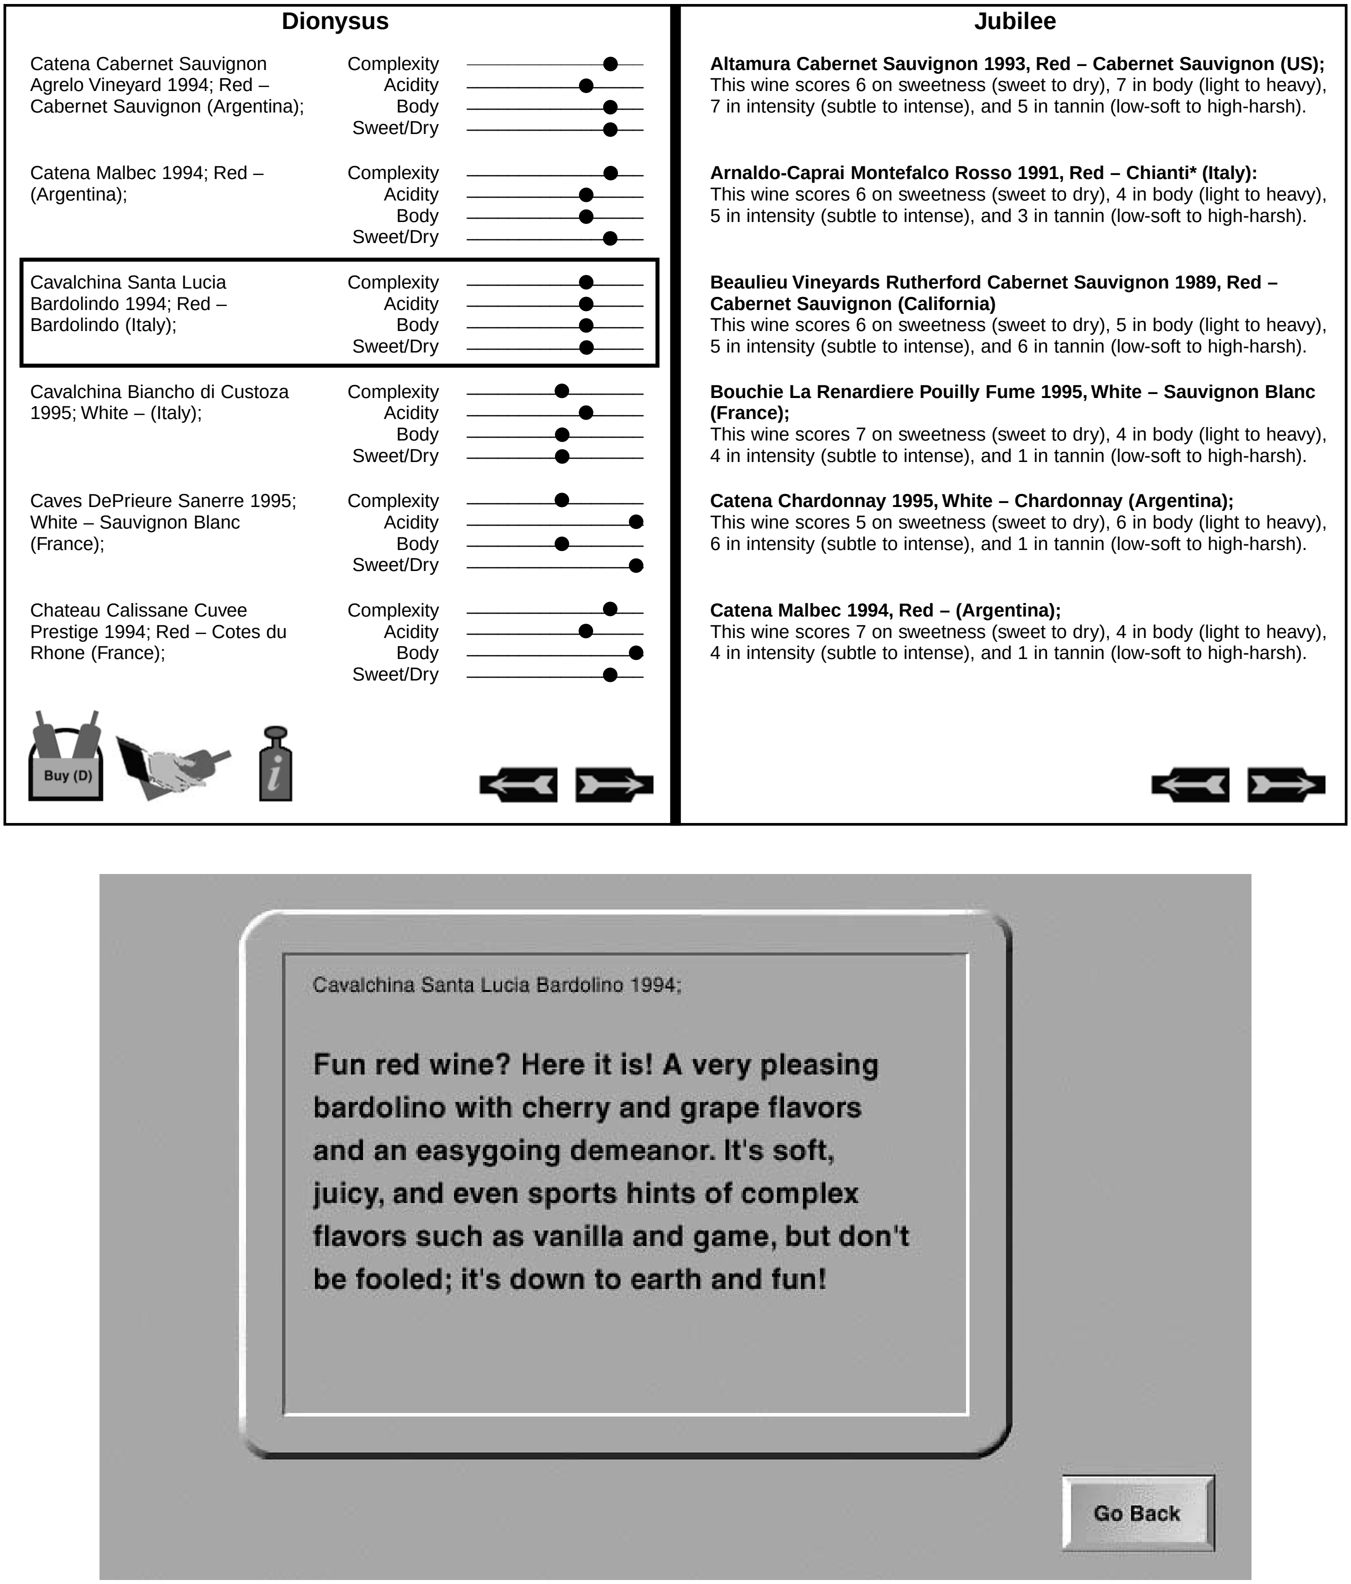

Dependent Measures: Price Sensitivity. The major dependent measures were based on the difference in wine purchasing when the wines were on discount and when they were not. Because we had only eight shopping trips, it was not possible to vary the price of each of the 100 wines independently. We therefore divided the wines up into six "wine blocks," and we varied the prices of these wine blocks independently across the eight shopping trips taken by each respondent. Each wine block was sold at its regular price on four of the eight trials and at a 15 percent discount on the other four trials. We used eight price combinations, orthogonally changing the prices of the six wine blocks. The order of exposure to the eight price combinations was counterbalanced. The order factor produced no reliable effects and will not be discussed further.

Each store had three wine blocks of 20 wines each, with varying prices. Two wine blocks in each store included only wines that were unique to that store and one block included only wines sold in both stores. (Note that the prices in the common block of wines in Dionysus were varied independently of the prices in the same block of wines in Jubilee.) For each participant, we separately calculated price sensitivity collapsed across all six wine blocks pooled across stores.

For each respondent, we calculated two measures of how sensitive a wine block's sales were to changes in its own price: Quantity Difference and Price Elasticity. Quantity Difference is a measure of the slope of an individual demand curve. It equals the total number of bottles purchased from a wine block in four high-priced trials minus the total purchased in four low-priced trials. For each respondent, we calculated one such measure collapsing across all six blocks, as well as separate measures collapsing across the four unique blocks and across the two common blocks.[3] Each Quantity Difference is divided by the number of wine blocks included in the measure, indicating the per-block difference in total quantity of the block purchased at high vs. low prices. Similarly, we calculated Price Elasticity collapsing across all six wine blocks and separately for the (collapsed) four unique wine blocks and for the two common blocks. The Price Elasticity measure was based on the proportional change in this quantity relative to the proportional price change (more details on this later). Again, negative values correspond to downward-sloping demand.[4]

To calculate Price Elasticity, let Q(R) refer to the quantity of wines sold at regular price, and Q(D) refer to the quantity sold at the discounted price. Let $(R) refer to the regular price of the wines, and $(D) refer to their discounted price. By relating the proportional difference in quantity purchased under the two pricing conditions to the proportional change in price, we get the formula for Price Elasticity shown in Equation 1:

$$\frac{\{Q(D) - Q(R)\}}{\{Q(D) + Q(R)\}/2} \Bigg/ \frac{\{\$(D) - \$(R)\}}{\{\$(D) + \$(R)\}/2}$$

Note that in our case, since the discount was always fixed at 15 percent, the price part of the equation (the denominator) is simply a constant with a value of -.1621622, because (.85 - 1.0) / (.85 + 1.0)/2 = -.1621622.

Dependent Measures: Search during Shopping. We also collected various measures of the shopping process for each respondent relating to depth and breadth of search (cf. Novak and Hoffman 1997). We will explore their relationships to price sensitivity.

a. Scroll is the sum of times subjects hit the "Next" and "Previous" scroll buttons. This is a measure of the amount of shallow but broad search.

b. Sort by Price is the number of times the respondent sorted the wines by price. This was possible only in high Price Usability conditions.

c. Sort by Varietal is the number of times the respondent sorted the wines by varietal (Chardonnay, Merlot, etc.). This was possible only in high Quality Usability conditions.

d. Drill for Wine Comment is the number of times the respondent drilled down conditions to view the differentiating comment (e.g., "Fun red wine? Here it is . . ."). This was possible only in high Quality Usability conditions.

e. Drill for Missing is a measure of the number of times respondents drilled down to a second screen to access "missing" information about either price (in low Price Usability, high Quality Usability conditions) sensory quality ratings (in high Price Usability, low Quality Usability conditions), or both (in low Price Usability, low Quality Usability conditions). In high Price Usability, high Quality Usability conditions, this drilling was not possible or necessary.

Procedure and Dependent Variable for the Wine-Tasting Task. After completing the computer-shopping task, respondents continued on to the wine-tasting task. The purpose of the task was to get a measure of how successfully consumers chose their wines, as measured by the difference in their ratings of wines purchased and not purchased. The local wine expert mentioned earlier chose which wines would be included in the wine-tasting test. Our goal was to include the most popular wines in the set. We hoped that this approach would maximize the probability that each consumer would have purchased at least some of the wines included in the wine-tasting, thereby permitting the above measures to be calculated. Finally, before starting the taste test, respondents were given a choice of whether to taste 10 red or 10 white wines taken from the earlier wine-shopping task.[5] We assumed that, given individual preferences for red and white wines, allowing the choice would again increase the probability that a respondent would taste some of the wines that he or she had purchased. Three of the 72 respondents had not purchased any of the 10 wines they tasted, and so the results for this dependent measure are based on the data from the 69 respondents who had bought at least 1 of the wines tasted.

Within each set, the wines were tasted in an order from light to heavy, as is recommended for wine tasting. Respondents tasted .5 ounces of each wine from a plastic cup. Baguettes and water were available to cleanse their palates. Each wine was rated on a scale from 1 (poor) to 10 (excellent). Because each respondent had purchased some of the wines and had not purchased others, we could compute how well they liked both wines they had previously chosen and ones they had not.

Respondents were told by the experimenter pouring the wines not to discuss their perceptions with other participants. The wine tasting was not blind. We were trying to mimic the real world, in which if information from electronic shopping makes people think that they like a product better, this affects their experienced utility. Thus the labels of the wines—but not their prices—were visible during tasting.

After all 10 wines had been tasted, respondents were thanked and dismissed. They were told that they could pick up their ordered merchandise later that week. In order to guarantee that all participants benefited from the promised discounted prices, we chose to fulfill orders from a shopping trip when half the wines were on sale. Prior to picking up their wines, participants were notified by e-mail of which wines they had ordered on the selected trial and of the prices of those wines. Participants paid by cash or check.

Retention Measure. Two months after the completion of the study, participants were asked if they would like to subscribe to the same electronic-shopping wine service from their homes. They received the following message:

"During Term 1, you participated in our study on electronic home shopping for wines. We would like to get your feedback on your experience and to assess your interest in continuing on in another phase of the study in which you would be able to order wines during Terms 3 and 4. Would you be interested in participating in the next phase of the study? If you say yes, we will e-mail you our wine program to install on your computer to use at your convenience. You would pick up your wines and pay for them at the kiosk on the following day. You would have exactly the same interface and merchants as you had in the earlier stage.

______ Yes, please e-mail me the program

______ No, I am not interested in participating."

Results

Price Sensitivity. Because we had two measures of price sensitivity—Quantity Difference and Price Elasticity—we have two tests of every key ANOVA effect in a 2x2x2x2 mixed design with Price Usability x Quality Usability of Information x Store Comparability as between-subjects factors and Unique vs. Common Wine Blocks as a repeated factor. In the results presented below, we use MANOVA as protection against escalating type 1 errors that would be expected if each measure were analyzed separately. Measures significant by multivariate tests are then examined separately for each dependent measure.[6]

The MANOVA results that were critical to our theory were that there was a significant multivariate main effect of Quality Usability [$F(2,55) = 6.50$, $p < .003$] and a significant multivariate interaction of Store Comparability x Unique vs. Common Wine Blocks [$F(2,55) = 3.87$, $p < .03$.] We followed up with univariate tests for Quantity Difference and Price Elasticity below.

For both dependent variables, there was a significant univariate effect of Quality Usability. Quantity Difference showed a more negative slope in the demand curve for low Quality Usability (M = -.90) [7] than for high Quality Usability (M = +.13), [F(1,56) = 13.08, p < .01]. Similarly, respondents showed greater Price Elasticity for low Quality Usability (M = -2.47) than for high Quality Usability (M = -.52), [F(1,56) = 3.88, p < .054]. See Figure 4.

Figure 4. Price Sensitivity Measures as a Function of Quality Usability

Self-elasticity measures are on the right and Quantity Difference (Quantity at Low Price–Quantity at High Price) on the left. For both measures, more-negative numbers imply greater price sensitivity.

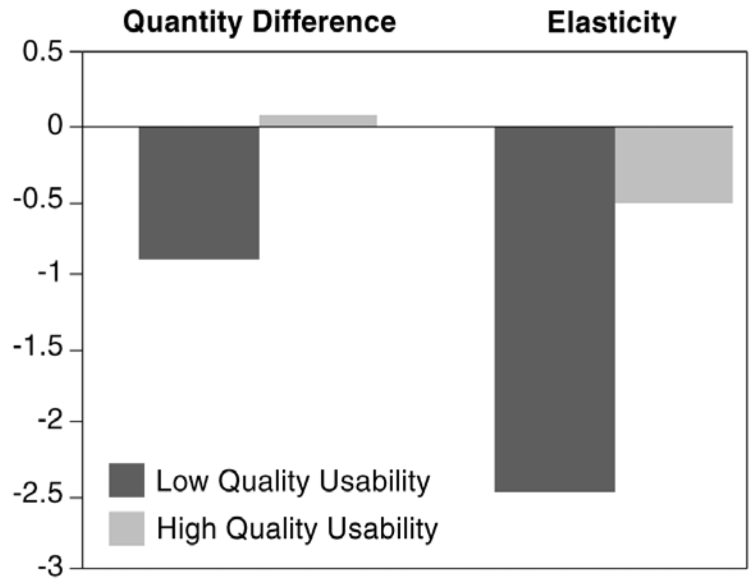

When the parallel analysis is done on all 72 respondents—by pooling across all six wine blocks to avoid missing values for eight participants—the effects described above are strengthened. The multivariate main effect of Quality Usability is significant (a measure of Equation 1 that pools across all six wine blocks), [F(2,63) = 6.83, p < .003], as are the univariate tests for Quantity Difference [F(1,64) = 13.67, p < .001] and Price Elasticity [F(1,64) = 7.89, p < .007]. The greater n and the balanced cell sizes from this approach increase statistical power.

There was a significant univariate interaction of Store Comparability x Unique vs. Common Wine Blocks for Quantity Difference [F(1,64) = 4.84, p < .04] and for Price Elasticity [F(1,56) = 6.89, p < .02]. The similar patterns for Quantity Difference and Price Elasticity can be seen in Figure 5 below.

Figure 5. Price Sensitivity Measures as a Function of Store Comparability and Merchandise Uniqueness

Left side shows the results for Quantity Difference and right side shows the results for Price Elasticity.

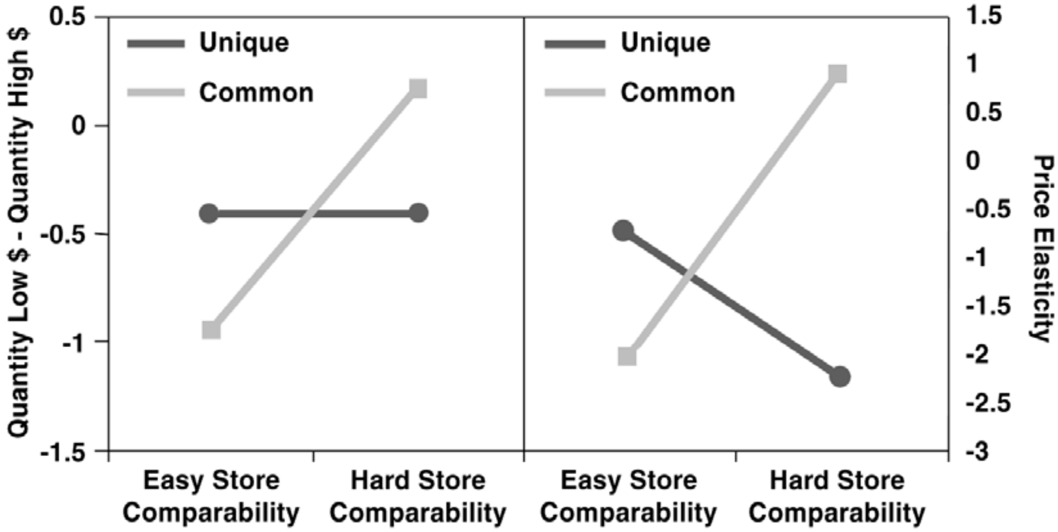

Simple-effects follow-up tests for Quantity Difference showed that for Unique wine blocks, increasing Store Comparability had no effect [F(1,64) = 0.00], with M= -.41 for easy Store Comparability and M = -.42 for hard Store Comparability. For Common wines, Quantity Difference was significantly higher for easy Store Comparability (M = -.96) than for hard Store Comparability (M = +.17), [F(1,64) = 10.45, p < .002].

Parallel simple-effects tests for Price Elasticity showed that, for Unique wine blocks, increasing Store Comparability had no effect [F(1,56) = 2.68, p < .11], with directionally lower Price Elasticity when Store Comparability was easy (M = -.80) as opposed to hard (M = -2.27). For Common wines, Price Elasticity was significantly higher for easy Store Comparability (M = -2.04) than for hard Store Comparability (M = +.91), [F(1,56) = 4.57, p < .04].

Process Measures and Price Sensitivity. Recall that experimental conditions varied in the opportunity or necessity for various processing operations during search. We correlated measures of Price Elasticity and Quantity Difference (pooling over all six wine blocks) with various measures of process. For all the correlations, bear in mind that a negative correlation implies that more of the processing activity is associated with more-negatively-sloped demand curves and more Price Elasticity. We make no causal interpretations.

Scroll, a measure of breadth of search, was associated with greater (negative) Quantity Difference (r = -.34, p < .01, df = 70) and more (negative) elasticity (r = -.24, p < .05, df = 70). This is as expected, because scrolling increases potential consideration-set size. Moreover, scrolling is associated with a pattern of shallow search focusing on price, not quality. Sort by Price (possible only for respondents in the high Price Usability conditions) also increases with price elasticity (r = -.35, p < .04, df = 34) and marginally increases with Quantity Difference (r = -.29, p < .09, df = 34). Sort by Varietal (possible only for respondents in the high Quality Usability condition) had no reliable effect. Drill for Wine Comment (drilling down for a differentiating wine comment in the high Quality Usability conditions) decreased price elasticity (r = +.45, p < .01, df = 34) and Quantity Difference (r = +.48, p < .01, df = 34) . As expected, Drill for Missing—drilling down for information that was missing from the first-level screen—is associated with more Price Elasticity (r = -.63, p < .01, df = 16) and Quantity Difference (r = -.60, p < .01, df = 16) when the information that was missing was only price information (in the high Quality Usability, low Price Usability condition). However, when drilling down revealed either quality information alone (high Price Usability, low Quality Usability) or both quality and price (low Price Usability, low Quality Usability), correlations with Price Elasticity and Quantity Difference did not differ significantly from zero.

Process Mediation of Effects of Search Costs on Price Sensitivity. In the foregoing zero-order correlation analyses, we treated a given process measure as missing whenever it was structurally zero—e.g., because a price-sorting tool was unavailable to subjects in low Price Usability conditions. In the analyses about to be reported, we examine whether variations in patterns of search mediate the two key effects of our design variables on price sensitivity: the Quality Usability main effect and the Store Comparability x Unique/Common Wine interaction. For these analyses, we coded the process variables as zero rather than missing when they were structurally unavailable.

As a preliminary step, we analyzed the effects of Price Usability, Quality Usability, and Store Comparability on Scroll, Sort by Price, Sort by Varietal, Drill for Wine Comment, and Drill for Missing. Table 1 shows cell means. Unsurprisingly, respondents made more use of informational tools that let them sort by price, sort by varietal, drill down for differentiating comments, and drill down for missing information when those measures were available than when they were not. Also, providing a tool that made one kind of processing operation easier decreased use of other tools.

Table 1. Analysis of Effects of Price Usability, Quality Usability, and Store Comparability

Price Usability	Quality Usability	Store Comparability	Price Elasticity	Quantity Difference	Scroll	Sort by Price	Sort by Varietal	Drill for Wine Comment	Drill Missing	N
High	High	Easy	-.66	.11	118.9	3.9	2.9	45.4	0	9
High	High	Hard	-.23	.42	119.8	5.3	5.1	41.7	0	9
High	Low	Easy	-3.99	-1.75	147.0	5.2	0	0	69.4	9
High	Low	Hard	-1.68	.85	203.2	10.1	0	0	58.4	9
Low	High	Easy	-.94	-.51	196.6	0	6.2	26.7	89.7	9
Low	High	Hard	-.26	.31	119.6	0	9.9	32.0	69.3	9
Low	Low	Easy	-2.69	-.58	207.0	0	0	0	83.9	9
Low	Low	Hard	-1.52	-.38	262.1	0	0	0	96.9	9

Next, we regressed Price Elasticity and Quantity Difference (calculated across all six wine blocks) on the number of times the participant performed Scroll, Sort by Price, Sort by Varietal, Drill for Missing (price or sensory ratings), and Drill for Wine Comment, in addition to the ANOVA terms for main and interaction effects of Price Usability, Quality Usability, and Store Comparability. The effect of Drill for Missing might be expected to differ as a function of whether the missing information was price information only (in low Price Usability, high Quality Usability), sensory information only (in high Price Usability, low Quality Usability), or both price and sensory ratings (in low Price Usability, low Quality Usability). Consequently, we added to the model the interactions of Price Usability x Drill for Missing and of Quality Usability x Drill for Missing.

Recall that the main between-subjects result was a main effect of Quality Usability, such that price elasticity and Quantity Difference were less negative when Quality Usability was high rather than low. A MANOVA with Price Elasticity and Quantity Difference calculated across all six wine blocks showed that Drill for Wine Comment had a significant multivariate main effect [$F(2,56) = 6.04$, $p < .005$], but that there was no remaining partial effect of Quality Usability with process variables in the model [$F(2,56) = .48$]. Separate univariate analyses of Price Elasticity and Quantity Difference revealed similar effects. Drill for Wine Comment significantly reduced Price Elasticity [$b = +.062$, $t(57) = 3.18$, $p < .003$], leaving the partial main effect of Quality Usability nonsignificant [$F(1,57) = .78$]. Drill for Wine Comment significantly reduced Quantity Difference [$b = .02$, $t(57) = 2.99$, $p < .005$], and the partial main effect of Quality Usability was nonsignificant [$F(1,57) = .73$].[8]

The other theoretically relevant effect on price sensitivity was the interaction of easy Store Comparability x Uniqueness of Wines. To test for mediation of this effect by process measures, we computed difference scores between the price sensitivity measures for (Unique/Common) wine blocks and analyzed the corresponding difference score measures as a function of the same model variables as above. Any significant effect on these difference scores reveals an interaction with Wine Uniqueness in the original scores. Recall that we dropped eight subjects who lacked elasticity measures for Common wines. None of the process variables were significant in this analysis, and the effect of Store Comparability (i.e., the Store Comparability x Wine Uniqueness interaction) remained highly significant. For Price Elasticity (Unique/Common) differences, Store Comparability [$F(1,50) = 8.25$, $p < .001$], and for Quantity Difference (Unique/Common) differences [$F(1,50) = 3.27$, $p < .08$]. It is unsurprising that process measures do not mediate the Store Comparability x Wine Uniqueness interaction, as the process measures do not distinguish between search for Unique and Common wines.

Shopping Evaluation. At the end of the shopping phase of the study, participants were asked to indicate how enjoyable the shopping experience had been for them on a 100-point scale. These responses were analyzed by a 2x2x2 Price Usability x Quality Usability x Store Comparability ANOVA. The results showed only three significant main effects. Each of the three components of increased transparency increased participants' enjoyment. Shopping was more enjoyable when Money Usability was high ($M = 68.3$) rather than low ($M=53.4$), [$F(1,64) = 10.14$, $p = .002$]; when Quality

Usability was high (M = 66.9) rather than low (M=54.8), [F(1,64) = 6.78, p =.011]; and when Store Comparability was easy (M = 66.9) rather than hard (M=54.7), [F(1,64) = 4.69, p = .034.] No interactions were significant.

We created a summary independent variable, Transparency of the interface, by summing the 1-0 dummy variables for Price Usability, Quality Usability, and Store Comparability. Regressing enjoyment on transparency is tantamount to a main-effects-only model with equality constraints on the weights of the 1-0 dummy variables for the three main effects. This is analogous to unit weighting schemes in decision making (Dawes and Corrigan 1974; Einhorn and Hogarth 1975). Transparency was highly significant [t(71) = 4.73, p < .0001]. Each incremental component of transparency added an average of 12.4 units of liking on a 100-point scale, as can be seen in Figure 6. Nested model comparisons showed that the model including only Transparency fit as well (R-squared = .242) as one estimating separate main effects for the three components of Transparency [F(2,68) = .27] and as well as a model estimating all main and interaction effects of the components [F(6,64) = .33]. R-squared values were .248 and .265 respectively, for the latter two models.

Figure 6. Shopping Enjoyment as a Function of Transparency of Shopping Environment

Transparency is the sum of dummy variables for Price Usability, Quality Usability, and Store Comparability.

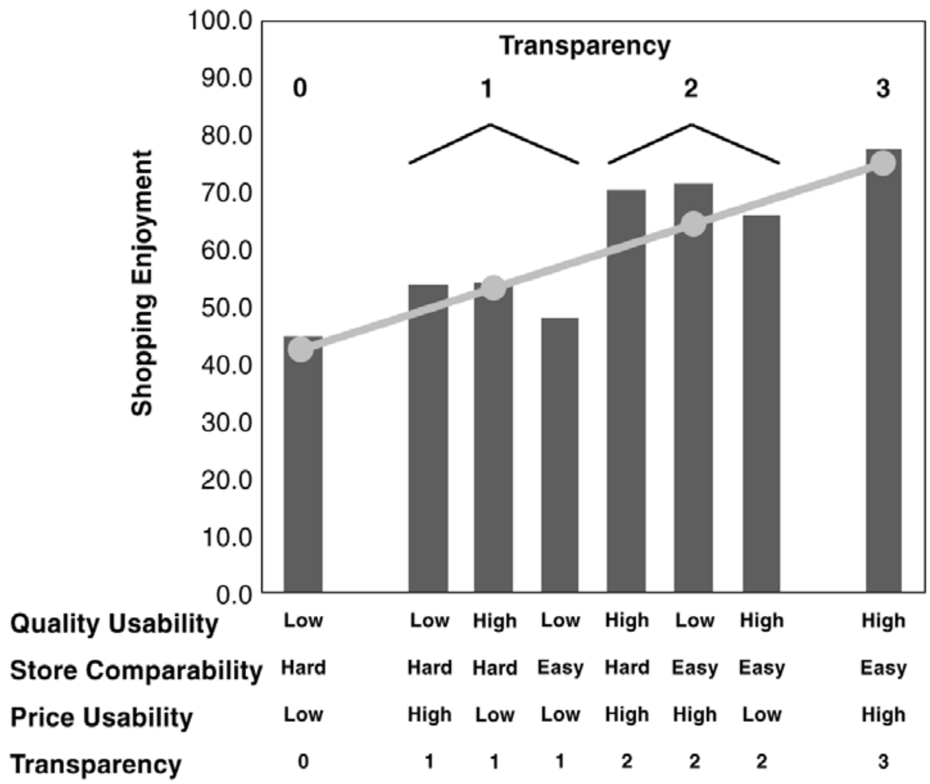

We can also examine correlations of shopping enjoyment with processing operations, using only cells in which a processing operation was not structurally precluded. Interestingly, though shoppers liked the task more in high Price Usability than in low, within the high Price Usability group, more use of the Sort by Price tool was associated with a less enjoyable shopping experience ($r = -.47$, $p < .005$, $df = 34$). Similarly, we found that when those in high Quality Usability conditions sorted by varietal more often, they rated the shopping task as less enjoyable ($r = -.51$, $p < .002$, $df = 34$). Finally, the shopping experience was less enjoyable for those who did more scrolling ($r = -.30$, $p < .02$, $df = 70$).

Consumer Welfare: Liking for Purchased Brands. In addition to the purchase-related measures, one of the main goals of our experiment was to test implications of lowered search costs on consumers' welfare. Consumers tested 10 wines after shopping. For the set of 10, we computed the mean rated liking of the subset that the consumer had earlier purchased on at least one of the eight shopping trials, dropping from the analysis three participants who did not purchase any of the 10 wines that they later tasted. We regressed this measure of liking for the purchased brands on the main effects of Price Usability, Quality Usability, and Store Comparability, and the mean liking of wines not purchased. The latter is a covariate that controls for individual differences in scale usage, liking for wine in general, etc.[9] We expected that consumers would make better choices—liking the purchased wines more—when Quality Transparency was high—that is, when Quality Usability and ease of Store Comparability were high.

Results showed a significant effect of the covariate rating of unpurchased wines [$F(1,64) = 47.7$], a main effect for Quality Usability [$F(1,64) = 6.74$, $p = .012$], and a marginal main effect for Store Comparability [$F(1,64) = 2.98$, $p = .089$]. As expected, consumers were better able to choose wines they liked when Quality Usability was high ($M= 5.97$) rather than low ($M=5.21$) and when Store Comparability was easy ($M=5.85$) rather than hard ($M= 5.33$). Both these effects supported the hypothesis that information systems that are more complete and informative will elicit higher ratings. These results support the idea that increasing quality information makes consumers better able to choose merchandise that matches their personal tastes. As expected, there was no effect for Price Usability [$F(1,64) = 0.00$].

We replaced Quality Usability and Store Comparability dummies in the model with a Quality Transparency index equal to the sum of their dummies—in essence constraining their weights to be equal in magnitude and direction. Model fit is not reduced significantly when Quality Transparency replaces the two separate dummies. Model $R^2 = .502$ for the constrained model and .504 for the unconstrained one reported above. This analysis leads to the conclusion that liking for the chosen wines increases by .64 of a scale point on a 10-point scale for each improvement in Quality Transparency [$t(1,65) = 3.05$, $p < .003$], with Price Usability still showing zero effect [$t(1,65) = .02$]. Figure 7 plots the results.

Figure 7. Taste-Test Liking for Wine as a Function of Quality Transparency (Quality Usability + Ease of Store Comparability Dummies) and Price Usability

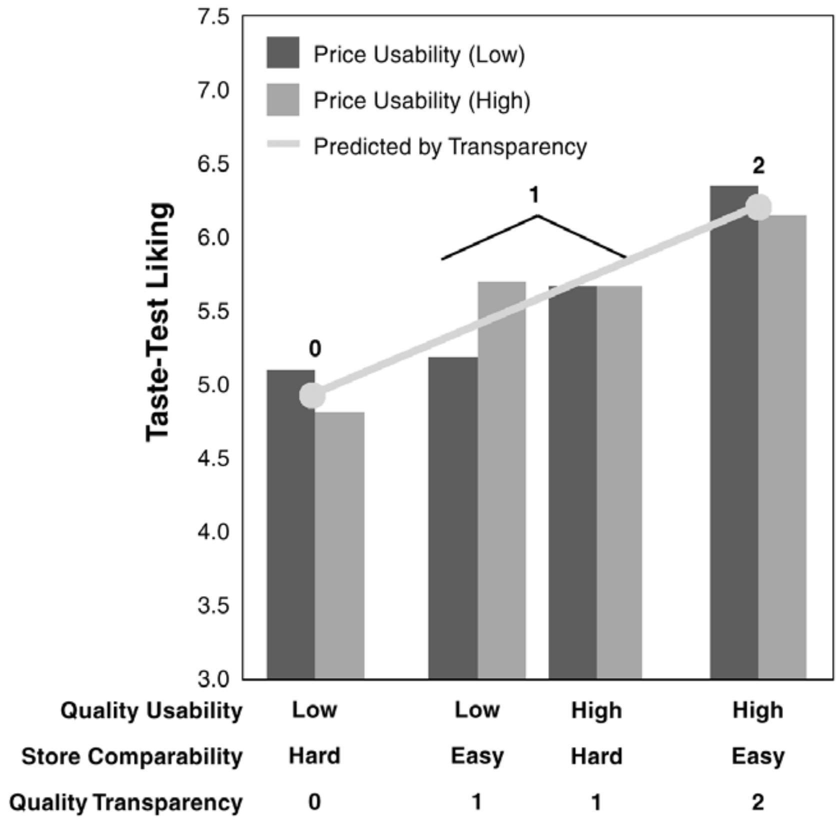

Two other findings for this dependent variable bear mention. First, full-featured systems allow consumers to choose wines that match their personal tastes. Put differently, less powerful systems may drive consumers toward random choices, or toward more stereotypical choices. We analyzed the number of bottles from our taste-test sample that participants purchased to see if inexperienced consumers purchased most frequently from our full inventory, as our wine expert expected would happen. We found a main effect for Store Comparability [$F(1,61) = 6.54$, $p = .013$]. Participants in the hard Store Comparability condition purchased more of our sample than participants in the easy Store Comparability condition. We speculate that in the easy Store Comparability condition consumers were better able to choose wines that deviated from the group norms.

Second, we correlated our search variables with the difference in rated liking of chosen and unchosen wines, using only cells for which a search variable was not structurally zero. In the high Price Usability conditions, sorting by price was associated with lower liking for chosen wines relative to unchosen ones ($r = -.34$, $df = 34$, $p < .05$).

Retention. Two months after the first phase of the study, participants were invited to take part in a second phase in which they would be e-mailed software from the experiment that would allow them to shop from home. Their agreement was our measure of retention. Three participants did not respond. We analyzed the data treating those nonresponses both as missing and as a failure to retain. We present the former results, which prove to yield slightly more conservative conclusions.

Yes/No responses were analyzed by a 2x2x2 Price Usability x Quality Usability x Store Comparability ANOVA. The results showed a marginally significant main effect for Quality Usability [$F(1,61) = 2.91$, $p = .093$] and a marginally significant main effect for Store Comparability [$F(1,61) = 3.77$, $p = .057$], both two tailed. Logit and probit regressions including the main effects of Price Usability, Quality Usability, and Store Comparability all lead to the same conclusions, albeit with slightly more power. When Quality Usability was low, 29 percent requested software to use the service further, compared with 49 percent when Quality Usability was high. When Store Comparability was hard, 27 percent were retained, whereas 50 percent were retained when Store Comparability was easy. The Store Comparability effect becomes significant at conventional levels if missing responses are treated as "No" responses.

Our expectation was that between-cell differences in retention probability would be completely accounted for by a 1 df contrast for Transparency, defined as the sum of the 1-0 dummy variables for Price Usability, Quality Usability, and Store Comparability. This expectation was supported. Retention significantly increased with Transparency [$F(1,61) = 7.11$, $p < .01$], and there was no significant residual between-cell variation after subtracting out variance due to Transparency [$F(6,61) = .24$]. If Retention is regressed on Transparency alone, Retention probability increases 17.4 percent for every added element of Transparency, as can be seen in Figure 8.

Figure 8. Retention as a Function of Transparency

(Sum of dummies for Price Usability, Quality Usability, and ease of Store Comparability). Predicted transparency is shown by the sloping line.

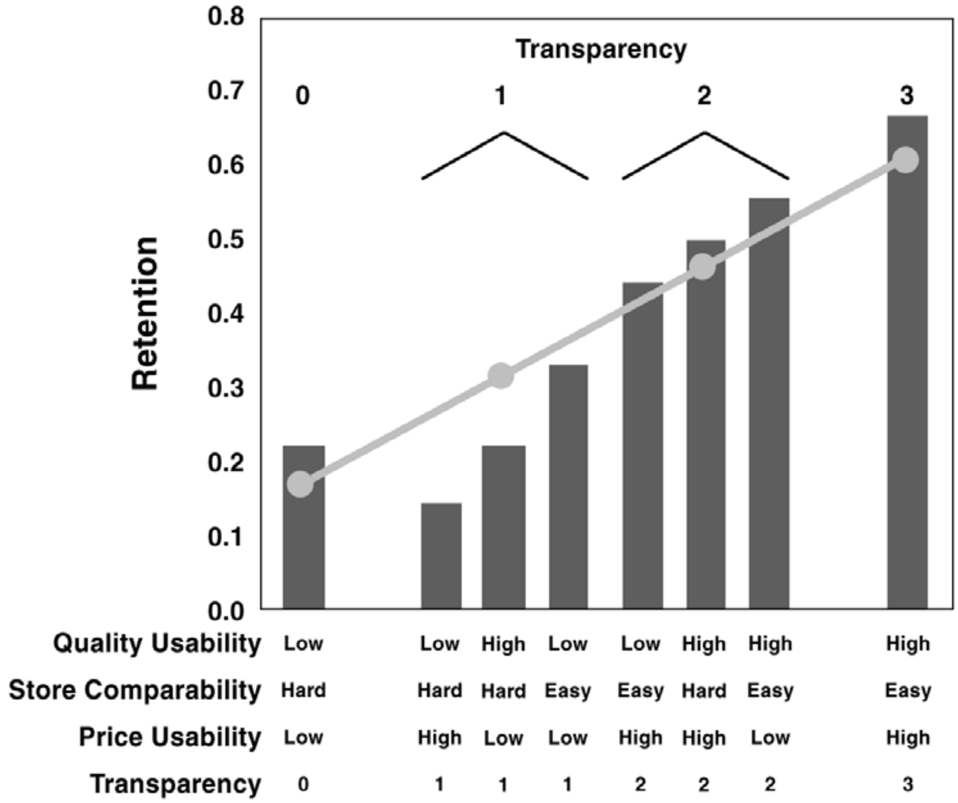

Discussion

Implications for Retailers

Both academics and popular business-press writers have stressed the potential for electronic retailing to increase competition, owing largely to easier price search. Following Alba et al. (1997), we have argued that incumbent retailers' fears of increased price competition have led them to underinvest and to attempt to create electronic venues that minimize exposure to competition. The result, ironically, is that their defensive offerings give customers little reason to shop electronically, leading to poor sales that are then misinterpreted as an indication that the electronic venue has low potential.

We have argued that the conventional analysis overlooked two important countervailing principles that have received ample documentation in marketing research on the economics of information. First, electronic shopping can also reduce search costs for differentiating quality information (Alba et al. 1997; Bakos 1997). It has been well established that differentiating information can lower price sensitivity (Kaul and Wittink 1995; cf. Mitra and Lynch 1995). Second, making it easy for consumers to compare across stores need not intensify price competition—at least not if competing stores are selling exclusive (nonoverlapping) merchandise. Consequently, if a retailer cooperates in efforts to lower search costs for price, for quality information, and for comparison across stores, it may well be that consumer price sensitivity will be no greater than it is currently at electronic retailing sites. It is a matter of the relative strength of the three effects.

Our empirical results strongly support this conjecture. We examined the effects of Price Usability, Quality Usability, and ease of Store Comparability on price sensitivity. We found that increasing Price Usability had no effect. Increasing Quality Usability decreased price sensitivity. Increasing ease of Store Comparability increased price sensitivity for Common wines sold by both competing merchants but not for wines unique to one merchant. On net, even for wines sold by both, price sensitivity was no higher in the most transparent informational environments (high Price Usability, high Quality Usability, easy Store Comparability) than in the least transparent. Thus, in this market, retailers could offer consumers the benefits of lowered search costs without exposing themselves to increased price sensitivity and price competition that might erode margins.

The second half of our thesis is that offering consumers the benefits of transparency improves consumer welfare. We have shown that consumers were better able to choose wines that they ended up liking when they used electronic interfaces that maximized the transparency of quality information.

If information has the potential to reduce the full price of a good (purchase price + cost of search + costs from disappointing purchases), consumers will demand it (Erlich and Fisher 1982). We argued that by maximizing the transparency of the information environment for consumers, retailers would earn their repeat business—or, alternatively, that consumers would not be trapped in impoverished, defensive sites of low transparency. Our findings show exactly that. Retention was significantly higher for shoppers in more transparent informational environments when we recontacted them two months after the experiment with the offer to continue the same service.

None of these findings is surprising once one points out the pertinence of the relevant literatures. But the implications for manufacturers are quite different from the conventional analysis.

Implications for Manufacturers: Effects of Distribution on Market Share

Thus far we have taken a retailer's perspective on the problem of whether or not to participate in transparent electronic markets that permit easy comparison. It is worth considering the implications of our data for manufacturers. Alba et al. (1997) discuss the potential for electronic retailing to threaten (inter)national brands by shifting the formula for retail success from stocking branded goods that draw customers into the stores to stocking exclusive merchandise. If a customer can buy the same pair of Levis over the Internet from multiple vendors, price competition may erode dealer margins. Retailers may respond by replacing the branded item with a slightly less popular exclusive or private-label offering (Faust 1997; King and Bounds 1997; White 1998).

Manufacturers could respond by striking deals with retailers for exclusive distribution, but this would entail a potential penalty of lost sales through stores no longer used as distributors. The interesting conceptual question is whether lowered search costs through electronic commerce reduce returns to distribution—in essence, lowering the sales penalty for granting exclusivity.

We took a very preliminary look at this issue by examining the effects of Price Usability, Quality Usability, and Store Comparability on the per-customer market share of the common wines in our experiment. For each customer, we calculated (Quantity of Common Wines) / (Quantity of Unique Wines + Quantity of Common Wines). Results showed a main effect for ease of Store Comparability, [$F(1,64) = 4.99$, $p < .05$]. The common wines had an average 35.9 percent market share when Store Comparability was hard, and an average 26.2 percent market share when it was easy. The implication is that manufacturer returns to distribution are less in electronic environments that make cross-shopping easy rather than hard.

It is instructive to compare each of these market shares to two benchmark null hypotheses. Consider the null hypothesis that the Common wines' market share is proportional to their share of distribution—i.e., one-third of the merchandise at each of the two competing stores. Given the 35.9 percent share in the hard Store Comparability condition, one cannot reject that null hypothesis [$t(35) = .83$, $p > .4$], but the 26.2 percent Common wine share in the easy Store Comparability condition is significantly lower than 33.3 percent [$t(35) = -2.36$, $p < .05$].

We noted that even when Store Comparability was easy, wines sold in both stores got disproportionately high market share. We formed Common and Unique wine blocks by dividing our 100 wines into five stratified random samples of 20 wines each, then designating one of the five blocks as the one to be sold in both stores. If the customer realized the identity of the wines sold in both stores, then the market share of those wines would be one-fifth, not one-third.[10] Common wines had significantly more than 20 percent market share for both easy Store Comparability [$t(35) = 2.08$, $p < .05$)] and hard Store Comparability [$t(35) = 5.33$, $p < .0001$]. Thus, there are positive returns to distribution for electronic merchants even in our transparent markets. A manufacturer could not costlessly grant exclusive distribution to a single retailer.[11]

Conclusions and Future Research

The message of our research for retailers is that they should not be so defensive. Retailers should not fear electronic shopping—only poor electronic-shopping systems. The challenge to on-line merchants is not to fortify their defensive positions in an environment that is arguably doomed to fail. Rather, their task is to learn how to provide consumers with useful product-related information that will increase consumer retention by increasing consumers' satisfaction with the merchandise they purchase. When they cooperate with comparison-shopping engines, it is in their interests to make the informational base include richer differentiating information. Current comparison agents have relatively impoverished criteria, effectively increasing Price Usability and Store Comparability without increasing Quality Usability (The Economist 1997; Quick 1998a). See, e.g., http://www.zdnet.com/computershopper/.

The net effect that lowering search costs has on price sensitivity and consumer welfare is a matter of calibration of the strength of effects of our three search costs (compare results for Common and Unique wines). We chose wines quite deliberately in an attempt to demonstrate our rhetorical points. Thus, what we expect to generalize is our conceptual point about the tradeoff of these search costs rather than our specific finding that the differentiating effects dominate over the effects of lowering search costs for price information. We intend to provide an appropriate conceptual framework that can be adapted to any specific case.

In the same way, we would not attempt to predict from our results whether price sensitivity should be higher or lower in specific electronic markets than in their conventional brick-and-mortar or catalog counterparts. Though channels compete on multiple dimensions, one can argue that much of the competition between retail formats is on various dimensions of search costs (Alba et al. 1997). Consequently, the same conceptual dimensions that vary among our alternative electronic interfaces can be used to distinguish electronic and conventional shopping formats. Whether a particular electronic market yields more or less price sensitivity than brick-and-mortar or catalog formats would be a matter of the balance of these competing dimensions of search costs and the differing utility for customers choosing those formats.

Beyond these matters of calibration, there are interesting questions of the external validity of our findings that turn on the interaction of our conceptual search-cost variables with background factors held constant in our experiment (Lynch 1982). We found no effects of lowering search costs for price information. In our research, there were only two sellers. Also, differentiation was high in the sense that wines differ much more than alternatives in many other product classes in lack-of-fit penalties for making choices that deviate from "best."

Bakos (1997) argues that both these dimensions of an electronic market should moderate the effects of making price information more transparent. First, lowering search costs for price information should matter more when there are more sellers. Perhaps if we had replicated our results with, say, six sellers instead of two, we would have found that the net effect of lowering all three component search costs would have been to increase price sensitivity. Second, Bakos argues that lowering search costs for price should matter less when lack-of-fit costs are higher. These should occur in product classes characterized by a high degree of real differentiation—such as wines. In such a market there is great potential for well-done electronic shopping to increase consumers' ability to choose goods that they like better than those they would have chosen in another shopping medium. In commodity-type markets, perhaps the effect of making prices more transparent would prevail. Further empirical and analytical work on seller incentives to disclose parity vs. differentiation is needed (cf. Bakos 1997; Shaffer and Zettelmeyer 1998).

Second, prices in our study were exogenous. It is an open question how our results would change if pricing were endogenous. Electronic retailing makes it easier to monitor competitors and to respond in a rivalrous manner (Cortese 1998). Gatignon's (1984) research on advertising suggests that such circumstances might tip the balance of the effects of transparency toward greater price sensitivity.

Third, Mitra and Lynch (1995) argue that the path from advertising to differentiated preferences would be strongest when consumers begin with little knowledge. The same point should hold for differentiation via electronic shopping. If a consumer learns about alternatives off-line and only comes on-line to complete the transaction, there would be little differentiation effect to offset the lowered costs of search for price and for alternative sellers of the preselected good.

This same conceptual point suggests interesting avenues for future research on dynamic changes in price sensitivity over time after a cohort of consumers has adopted a full-featured electronic shopping system. Consider highly differentiated markets in which there is little new entry over time, but where prices remain volatile over time due to promotion, etc. A full-featured electronic shopping system, when first introduced, might produce results such as we observed, with the effects of increased quality transparency offsetting those of increased price transparency. So, for example, one might find that shoppers are less price-sensitive when they first shop electronically than when they shop in brick-and-mortar stores. But after some time, consumers will have learned about whatever product differences exist, while still relying on the electronic medium to inform them of (unstable) prices. Once real product differences are learned, the potential of electronic shopping to achieve further differentiation would diminish while consumers' ability to track and compare volatile prices would remain. Thus, one might expect that with the introduction of full-featured electronic shopping, price sensitivity would first decrease and then increase over time compared with prices in brick-and-mortar retail environments. In contrast, in a category such as wine in which there is constant turnover in the set of competing alternatives, the role of an electronic shopping system in explaining differences among products would remain.

All may not be lost even when consumers can learn. Electronic sellers can learn about customers' tastes; they then can use this information to provide better- and better-tailored advice about which alternatives would maximize customer utility.

Retailers can offer customers ancillary services such as the "Personal Lists" that Peapod offers. These reduce price sensitivity (Degeratu et al., 1998). They can use customer data to anticipate utility better, using smart agents to build trust (Urban 1998). Research is needed to learn how retailers can use electronic commerce not to compete on price but to capture the value of differentiation for their customers—and for themselves.

Notes

1. See, for example, Krantz's discussion of Sears's "cautious" approach to on-line retailing (1998, p. 40), or the discussion presented by Downes and Mui of an executive survey pointing to the prevalence of cautious strategies with only low-risk offerings (1998, p. 88): "There is considerable wishful thinking on the part of those, like the CEO of a large retail chain in our survey, who don't think customers are ready to do business with us in cyberspace." See also Steinhauer (1998) on the same topic, and Wigand and Benjamin (1995) for a similar discussion of incumbent disincentives to participate in electronic markets for travel services.

2. The ratings of each wine and the differentiating comments were provided by the head wine buyer at the top wine store in Durham, North Carolina. Some of the comments came from Wine Spectator magazine.

3. Because each store had two unique wine blocks and one common wine block, the overall measure of Quantity Difference is not the unweighted average of the measures for unique and for common blocks. The same is true for Price Elasticity measures.

4. Price Elasticity and Quantity Difference each have advantages and disadvantages. Price Elasticity is the normatively relevant measure, but individual-level Price Elasticity is not preserved by aggregation. That is, obtaining Price Elasticity by aggregating quantities across all respondents and then calculating Equation 1 is not identical to calculating Equation 1 separately for each individual and then averaging the individual values. Respondents buying few bottles of wine are weighted equally to those buying many in calculating average individual-level Price Elasticity, but those buying many bottles contribute more to aggregate Price Elasticity. Quantity Difference, on the other hand, is preserved by aggregation; the value of Quantity Difference obtained by averaging quantities purchased across individuals is equal to the average of the individual Quantity Differences.

5. The red wines were Elsa 1994 Malbec San Rafael, Masciarelli 1995 Montepulciano d'Abruzzo Italian Dry Red Wine, Dievole 1995 Chianti Classico, Moris 1995 Morellino di Scansano Dry Red Wine, Domaine du Vieux Chene 1995 Cotes du Rhone Villages, Mas Donis 1996 Garnatxa Syrah, Mestre Pere et Fils 1990 Santenay Gravieres Burgundy Red Wine, Kingston Estate 1993 Riverland Shiraz, Konrad 1992 Melange A Trois (Cabernet Sauvignon, Cabernet Franc, Merlot Blend), and Lava Cap 1995 Estate Bottled El Dorado Cabernet Sauvignon. The white wines were Francois Montand Blanc de Blanc Brut Premium Sparkling Wine, Rustico Nino Franco Sparkling White Wine, Vestini 1996 Trebbiano d'Abruzzo, Chartron La Fleur 1996 White Bordeaux, E. Guigal 1996 Cotes Du Rhone Blanc, Domaine Val Saint Jean 1996 Chardonnay, Christophe 1994 Napa County Chardonnay, Konrad

1993 Mendocino Chardonnay, Trumpeter 1996 Tupengato Chardonnay, and Millstream 1996 California White Zinfandel.

6. There are different patterns of missing data for the analyses reported below. Eight participants purchased no wines in the two blocks of common wines. This creates division by zero in Equation 1; consequently, we had missing values for Common wine Price Elasticity for those eight participants. MANOVAs dropped those eight participants. There were no missing values for Unique wine Price Elasticity or for Unique or Common wine Quantity Difference. When follow-up tests could be conducted either on the 64 participants with complete data or on all 72 participants, qualitative results were identical in terms of patterns and statistical significance. We report follow-up tests on the full data set whenever possible.

7. This means that, for each block of 20 wines, the average respondent in Low Quality-Usability conditions bought .90 less bottles in total during the four trips when those wines were at regular price than on the four trips when the same wines were discounted by 15 percent.

8. For Price Elasticity, there was an additional significant tendency for Sorting by Price to increase price sensitivity [$b = -.276$, $t(57) = -2.14$, $p < .05$]. For Quantity Difference, there was an additional marginal tendency for scrolling to increase price sensitivity [$b = -.002$, $t(57) = -1.81$, $p < .08$], and a marginal interaction of Quality Usability with Drill for Missing Information [$b = -.02$, $t(1,57) = -1.93$, $p < .06$]. As expected, drilling increased price sensitivity in high Quality Usability conditions when drilling could reveal price but not quality information.

9. Cook and Campbell (1979, pp. 153-9) explain why this ANCOVA approach is more powerful than the alternative approach of analyzing the difference in mean ratings of purchased vs. unpurchased wines. In essence, the analysis of difference scores is like ANCOVA if the slope of the covariate effect is exactly 1.0. If it is less, as is the case here, the difference score analysis overadjusts for any differences between conditions in mean ratings of unpurchased wines. In the present case, however, analysis of difference scores leads to identical conclusions.

10. This is a classic similarity effect, as in a choice between traveling by car vs. by red bus vs. by blue bus (Tversky 1972).

11. This simple analysis ignores effects of Common vs. Unique distribution on retailer pricing.

References

Alba, Joseph, John Lynch, Barton Weitz, Chris Janiszewski, Richard Lutz, Alan Sawyer, and Stacy Wood (1997), "Interactive Home Shopping: Incentives for Consumers, Retailers, and Manufacturers to Participate in Electronic Marketplaces." *Journal of Marketing* 61 (July), 38-53.

Anders, George (1998), "Cybersqueeze: Comparison Shopping Is the Web's Virtue—Unless You're a Seller." *Wall Street Journal* (July 23), A1-8.

Bakos, J. Yannis (1997), "Reducing Buyer Search Costs: Implications for Electronic Marketplaces." *Management Science* 43 (December), 1676-1708.

Boulding, William, Enkyu Lee, and Richard Staelin (1994), "Mastering the Mix: Do Advertising, Promotion, and Sales Force Activities Lead to Differentiation?" *Journal of Marketing Research* 31 (May) 159-72.

Cook, Thomas, and Donald Campbell (1979), *Quasi-Experimentation: Design and Analysis Issues in a Field Setting.* Chicago: Rand McNally.

Cortese, Amy E. (1998), "Good-Bye to Fixed Pricing: How Electronic Commerce Could Create the Most Efficient Market of Them All." *Business Week* (May 4), 70.

Dawes, Robin M., and Bernard Corrigan (1974), "Linear Models in Decision Making." *Psychological Bulletin* 81 (March), 95-106.

Degeratu, Alexandru, Arvind Rangaswamy, and Jianan Wu (1998), "Consumer Choice Behavior in Online and Regular Stores: The Effects of Brand Name, Price, and Other Search Attributes." In *Collected Working Papers: Marketing Science and the Internet*, eds. Donna Hoffman and John Little. Cambridge, MA: MIT Sloan School, INFORMS Mini-Conference.

Downes, Larry, and Chunka Mui (1998), *Unleashing the Killer App: Digital Strategies for Market Dominance.* Boston: Harvard University Press.

The Economist (1997) "Business: Roboshop: Intelligent Agents" (June 14).

Einhorn, Hillel, and Robin M. Hogarth (1975), "Unit Weighting Schemes for Decision Making." *Organizational Behavior and Performance* 13, 171-92.

Erlich, Isaac, and Lawrence Fisher (1982), "The Derived Demand for Advertising: A Theoretical and Empirical Investigation." *American Economic Review* 72 (June), 366-88.

Faust, Fred (1997), "Venture Likes Fit of Private-Label Apparel Lines: Turnaround Depends on Success of Brands Like Stone Mesa Jeans." *St. Louis Post-Dispatch* (November 9), E1.

Gatignon, Hubert (1984), "Competition as a Moderator of the Effect of Advertising on Sales." *Journal of Marketing Research* 21 (November), 387-98.

Gittins, Ross (1997), "Price Is Right on the Business Net." *Sydney Morning Herald* (April 30), 17.

Haeubl, Gerald, and Valerie Trifts (1998), "Consumer Decision Making in Interactive Online Shopping Environments: The Role of Information Agents and Tools for Organizing Information." In *Collected Working Papers: Marketing Science and the Internet*, eds. Donna Hoffman and John Little. Cambridge, MA: MIT Sloan School, INFORMS Mini-Conference.

Hoffman, Donna, and Thomas P. Novak (1996), "Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations." *Journal of Marketing* 60 (July), 50-68.

______ , ______, and Patrali Chatterjee (1995), "Commercial Scenarios for the Web: Opportunities and Challenges." *Journal of Computer-Mediated Communication, Special Issue on Electronic Commerce* 1 (3).

Kaul, Anil, and Dick R. Wittink (1995), "Empirical Generalizations about the Impact of Advertising on Price Sensitivity and Price." *Marketing Science* 14 (3), G151-60.

King, Ralph T., and Wendy Bounds (1997), "Clothing: Its Share Shrinking, Levi Strauss Lays Off 6,395." *Wall Street Journal* (November 4), B-1.

Krantz, Michael (1998), "Cover: The Cyberspace Marketplace." *Time* (July 20), 34-41.

Kuttner, Robert (1998), "The Net: A Market Too Perfect for Profits." *Business Week* (May 11), 20.

Lynch, John G., Jr. (1982), "On the External Validity of Experiments in Consumer Research." *Journal of Consumer Research* 9 (December), 225-39.

Mitra, Anusree, and John G. Lynch, Jr. (1995), "Toward a Reconciliation of Market Power and Information Theories of Advertising Effects on Price Elasticity." *Journal of Consumer Research* 21 (March), 644-59.

______ , and ______ (1996), "Advertising Effects on Consumer Welfare: Prices Paid and Liking for Brands Selected." *Marketing Letters* 7 (1), 644-59.

Novak, T. P., and D. L. Hoffman (1997), "New Metrics for New Media: Toward the Development of Web Measurement Standards." *World Wide Web Journal* 2 (1), 213-46.

Pazgal, Amit, and Nir Vulcan (1998), "Have Your Agent Call Mine: Software Agents, the Internet and Marketing." In *Collected Working Papers: Marketing Science and the Internet*, eds. Donna Hoffman and John Little. Cambridge, MA: MIT Sloan School, INFORMS Mini-Conference.

Popkowski-Leszczyc, Peter. T. L., and Ram C. Rao (1990), "An Empirical Analysis of National and Local Advertising Effect on Price Elasticity." *Marketing Letters* 1 (2), 149-60.

Quelch, John, and Lisa Klein (1996), "The Internet and International Marketing." *Sloan Management Review* (Spring), 60-75.

Quick, Rebecca (1998a), "How Comparison-Shopping Sites Work—and Sometimes Don't: Robots Can Be Useful but May Miss Bargains." *Wall Street Journal* (August 20), B-8.

Quick, Rebecca (1998b), "Web's Robot Shoppers Don't Roam Free." *Wall Street Journal* (September 3), B-1, B-8.

Reeve, Simon (1998), "Business: Net Savings When You Buy US Goods: Electronic Commerce." *The European* (April 27), 24

Reeves, Byron, and Clifford Nass (1996), *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places.* Stanford, CA: Cambridge University Press.

Rosen, Sherwin (1978), "Advertising, Information, and Product Differentiation." In *Issues in Advertising,* ed. David Tuerck, 161-9. Washington, D.C.: American Enterprise Institute.

Shaffer, Greg, and Florian Zettelmeyer (1998), "Channel Conflict Over Information Provision and the Role of the Internet." In *Collected Working Papers: Marketing Science and the Internet*, eds. Donna Hoffman and John Little. Cambridge, MA: MIT Sloan School, INFORMS Mini-Conference.

Shankar, Venkatesh, and Arvind Rangaswamy (1998), "The Impact of Internet Marketing on Price Sensitivity and Price Competition." In *Collected Working Papers: Marketing Science and the Internet*, eds. Donna Hoffman and John Little. Cambridge, MA: MIT Sloan School, INFORMS Mini-Conference.

Steinhauer, Jennifer (1998), "Old-Line Retailers Resist On-Line Life." *New York Times* (April 20), D-1.

Tversky, Amos (1972), "Choice by Elimination." *Journal of Mathematical Psychology* 9, 341-67.

Urban, Glen (1998), "Trust Based Marketing on the Web." In *Collected Working Papers: Marketing Science and the Internet*, eds. Donna Hoffman and John Little. Cambridge, MA: MIT Sloan School, INFORMS Mini-Conference.

Wigand, Rolf T., and Robert I. Benjamin (1995), "Electronic Commerce: Effects on Electronic Markets." *Journal of Computer-Mediated Communication, Special Issue on Electronic Commerce* 1 (3).

White, George (1998), "Wall Street, California Fashion Pushes Sales of Store Brands Forward.' *Los Angeles Times* (September 8), D-1, Home Edition.