The Intertemporal Dynamics of Consumer Lock-In

GAL ZAUBERMAN*

This article examines how dynamic changes in information cost structure and time preferences affect consumers' search and switching behavior over time and lead to lock-in. The information cost structure is conceptualized as a trade-off of initial setup costs and ongoing usage costs. Lock-in is defined as consumers' decreased propensity to search and switch after an initial investment, which is determined both by a preference to minimize immediate costs and by an inability to anticipate the impact of future switching costs. The results of three experiments support the proposed mechanism. Experiment 1 shows that a small initial investment is sufficient to produce lock-in. Experiment 2 shows that the results of a prior investment on lock-in are not due to psychological commitment but to a shift in relative costs of incumbent and new options. Experiment 3 shows that respondents fail to anticipate how their prior investment will lock them in.

The emergence of the Internet has generated increased interest in the effect of information acquisition and processing costs on search and choice behavior. Economic theory generally assumes that lower search and evaluation costs enable consumers to search more and find the best option. However, anecdotal evidence and initial research (e.g., Brynjolfsson and Smith 2000; Johnson et al., forthcoming) do not support this view. Rather, they suggest even greater consumer stickiness (i.e., reduced search and switching) on the Internet compared with the bricks-and-mortar retail environment. The effects of the Internet on consumers' search and switching provide a timely and important exemplar of the general effects of information cost structure on such behavior over time. The goal of this article is to gain a better understanding of the behavioral mechanisms underlying these effects.

The factors that affect consumer choices over time have been the subject of extensive research in multiple disciplines. This article specifically focuses on the joint effects of relative information costs and time preferences on consumers' reluctance to search extensively for better alternatives and, ultimately, switch. "Lock-in" is defined as consumers' decreased propensity to search and switch after an initial investment. Lock-in is driven by a preference to minimize immediate costs and an underweighing of the impact of future switching costs. Consumers tend to focus on short-term considerations and to select the more easily accessible option at the time of the transaction, even at the cost of forgoing future benefits. Consumers also fail to anticipate the impact of future switching costs, and when the future arrives, these switching costs dominate these later decisions in ways that consumers do not anticipate when making the initial decision.

*Gal Zauberman is assistant professor of marketing at the University of North Carolina, Chapel Hill, Chapel Hill, NC 27599-3490 (zauberman @unc.edu). This article is based on the author's dissertation research at Duke University. The author is grateful to his dissertation advisors Jim Bettman and John Lynch for their help and support throughout this process and for many hours of stimulating discussion. The author also thanks his dissertation committee members John Payne and Tom Wallsten as well as Kristin Diehl, Oded Koenigsberg, Laura Kornish, and Jonathan Levav. In addition, the author acknowledges the helpful comments of the editor, associate editor, and reviewers.

THEORETICAL BACKGROUND

Switching Behavior

It is useful to consider lock-in and the factors that affect the likelihood of switching within the broader context of customer loyalty. Customer loyalty is a complex aspect of consumer behavior that has been studied from multiple perspectives, ranging from customer satisfaction to emotional commitment to simple habit (e.g., Fournier and Yao 1997; Jacoby and Chestnut 1978; Oliver 1999). The literature on customer loyalty also distinguishes between true loyalty and simple repeat purchase behavior (Jacoby and Kyner 1973). The current article focuses strictly on the repeat choice aspect of behavior resulting from how information costs are traded off over time.

The current approach is more in line with the economic perspective on product replacement decisions and the adoption of new options (e.g., Cripps and Meyer 1994; Klemperer 1987), which centers on the relative costs and benefits associated with each option. For example, Klemperer (1987) argues that switching costs can provide a source of differ-

entiation for otherwise homogeneous products and can lead to lock-in. Shapiro and Varian (1999) further identify search costs as a type of switching cost and as a potential cause of lock-in.

The search for new alternatives is essential for switching to occur when alternatives are not known with certainty. The economics of an information model of search is a cost-benefit framework in which search continues as long as the marginal cost of search is lower than the expected marginal benefit (e.g., Stigler 1961; Stigler and Becker 1977). As search cost decreases, search increases, ceteris paribus. Motivated by these ideas, prior research in marketing examined the effects of the Internet and the reduction of search costs on search and choice behavior (e.g., Alba et al. 1997; Bakos 1997; Diehl, Kornish, and Lynch 2003; Lynch and Ariely 2000). In this article, however, the important comparison is not changes in absolute levels of search costs but, rather, the relative costs of different options at any two points in time. That is, even if the overall level of search costs is reduced, a previously used alternative still holds a relative cost advantage that determines switching rates (cf. Lal and Sarvary 1999).

Time Preferences and Choice over Time

A central feature of the current approach is that it explicitly incorporates time preferences. One of the most established research findings on intertemporal choice is that individuals behave as if they have high discount rates and, therefore, tend to avoid immediate investments that would yield future benefits (e.g., Frederick, Loewenstein, and O'Donoghue 2002; Hausman 1979). There is also evidence suggesting that people tend to consider only the immediate decision and ignore future consequences (e.g., Herrnstein et al. 1993; Kahneman and Lovallo 1993). For instance, in an extensive research program on melioration, Herrnstein et al. (1993) found that when people select among alternatives, they focus on immediate utility and are insensitive to the effects of these choices on the later relative utility of other alternatives in the set.

Consumers not only hold high implicit discount rates and focus on the present, but research has also shown that the level of discounting changes as a function of time delay, which potentially leads to a reversal of preference over time (e.g., Loewenstein and Prelec 1992). That is, when considering trade-offs between options at two future moments, people tend to give stronger relative weight to earlier moments as the moments come closer in time (Akerlof 1991; O'Donoghue and Rabin 1999; Thaler 1981). Consistent with this effect, Soman (1998) shows that consumers underweigh future effort when evaluating different effort and payoff transactions over time. The implication for lock-in is that the cost of switching seems smaller the farther away it is in time, potentially leading consumers to choose options that are attractive in the short-run because they do not fully anticipate how painful it will feel to switch later. This effect is a critical aspect of the proposed framework.

Consumers weigh costs more heavily the closer the costs are in time. Are they aware of this tendency and able to anticipate this preference shift? Existing research indicates that people are not particularly accurate at predicting their future choices and preferences (e.g., Ausubel 1991; Loewenstein and Schkade 1999; Nunes 2000). Consumers perceive near and distant events differently (Trope and Liberman 2000) and lack the ability to anticipate accurately the implications of their current choice on a future choice (Herrnstein et al. 1993). On the basis of the reviewed research, I propose that consumers unwittingly create switching costs with their initial investments that are more psychologically binding than they initially anticipated.

THE LOCK-IN MECHANISM

The proposed lock-in mechanism centers on the impact of information cost structure and consumers' time preferences on search and switching behavior over time. The information cost structure refers to any cost associated with acquiring and processing information about the environment, including "search costs" and "transaction costs." Specifically, "information costs" refer to the effort that it takes to initiate, process, evaluate, and finalize a "search" task. "Time preference" refers to a preference for the timing of costs and benefits and, in this article, refers more narrowly to the preference for immediate utility over delayed utility (and delayed cost over immediate cost).

To incorporate temporal concepts into the current framework, I conceptualize the information cost structure in terms of two temporally distinct information cost categories: (1) an initial one-time "setup cost" (S), including the cost of locating a provider (search) and initiating the service (entry and customization); and (2) ongoing "usage costs" (O) that apply each time the consumer uses the service, including the cost of processing, evaluating, and finalizing the search.

Choosing between alternatives involves two key trade-offs between these two types of costs. First, there is a trade-off between the two types of costs for each option. For example, customizing an Internet portal requires a significant initial investment; however, because it results in a site that is structured to the consumer's preferences, subsequent visits are more efficient. The alternative is to use the portal with only a short registration but without customization (lower S) and to forgo the customization benefits (higher O). Second, selecting an option at time t creates switching costs at $t + 1$ that result from incurring the setup for that selected option. In other words, investing in a setup cost creates an advantage for that selected option compared with other options. These shifts in information cost structure, combined with consumer preference to minimize immediate total costs and a failure to anticipate the impact of switching costs, are the main elements utilized in the current approach. An outline of this conceptual framework of lock-in is presented in figure 1.

The Lock-in Model

The following model formalizes the conceptual framework. The modeling approach is structurally similar to several recent intertemporal choice models (e.g., Laibson, 1997; O'Donoghue and Rabin 1999) but was modified to capture dynamics in both the information cost structure and time preferences. In the current model, U_i is the total perceived utility in period i, where i is the point in time from which the consumer is evaluating the stream of utility in time periods t through T. I define $D = (t - i)$ as the temporal distance until the stream of utility begins. The perceived utility U_i is determined by the levels of S and O as well as the per-period benefits (Q):

$$U_i \equiv \beta(D)\delta^{t-i}[Q - (S_t + O_t)] + [1 - \beta(D)] \sum_{\tau=t+1}^{T} [\delta^{\tau-i}(Q - O_\tau)],$$

$$\text{where } 0 < \beta \leq 1,\ 0 \leq \delta \leq 1.$$

The utility U_i is a weighted sum of the discounted first-period transaction t and all discounted subsequent transactions $t + 1$ to T. Utility decreases in S and O and increases in Q (i.e., $\partial U/\partial S < 0$, $\partial U/\partial O < 0$ and $\partial U/\partial Q > 0$). In this model, U_i is also determined by two distinct components of time preferences: β represents the psychological bias toward the present, and δ represents standard economic discounting. Central to the model, β captures the relative weight given to the first period compared with the stream of discounted utilities for all other periods. Reflecting consumers' bias toward the present and time inconsistent preferences, β decreases with D (i.e., $d\beta/dD < 0$). The closer the transaction period t is to the current point of evaluation i (the smaller D), the larger is β and the greater is the relative weight given to the first period. For simplicity, the model assumes that Q is equal for all options and constant across time. Holding Q constant ensures that the insights from the model only result from the trade-off of costs over time.

The utilities of options A and B at the initial selection are thus given by:

$$U_i^A = \beta(D)\delta^{t-i}\left[Q - \left(S_t^A + O_t^A\right)\right] + [1 - \beta(D)] \sum_{\tau=t+1}^{T} \left[\delta^{\tau-i}\left(Q - O_\tau^A\right)\right];$$

$$U_i^B = \beta(D)\delta^{t-i}\left[Q - \left(S_t^B + O_t^B\right)\right] + [1 - \beta(D)] \sum_{\tau=t+1}^{T} \left[\delta^{\tau-i}\left(Q - O_\tau^B\right)\right].$$

Assume that, at the initial selection, $S \neq 0$, $S^A < S^B$, $O^A > O^B$, $(S^A + O^A) < (S^B + O^B)$. Note that A is preferred to B if there is only a single transaction. However, even with multiple transaction periods, people often prefer options with low immediate costs, as predicted by research on intertemporal choice. Consumers are more likely to select A than B when β is larger (greater bias toward the present) or when δ is smaller (the greater the long-term discounting).

FIGURE 1

A CONCEPTUAL FRAMEWORK OF LOCK-IN

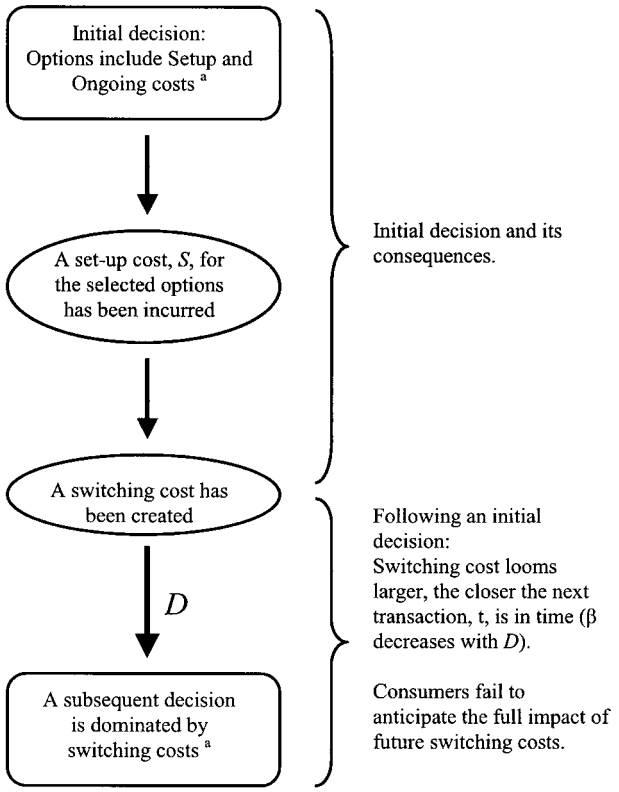

NOTE.—Subscript *a* = time preference (i.e., a preference for minimizing immediate total costs) affects the decision.

Consumers will be less likely to consider and switch to a (possibly superior) new alternative after the initial selection (associated utility is denoted as U_2) than before the initial selection (associated utility is denoted as U_1). This is true because reusing the previously selected option does not require any setup cost, whereas the total cost for switching to any other option includes both setup and usage costs. In other words, after initial selection of A, the setup cost has been incurred, so that the setup cost S^A at the next decision point U_2 is zero. When A is chosen over B in period 1, it therefore follows that $(U_2^A - U_2^B) > (U_1^A - U_1^B)$.

The gap in relative costs (i.e., switching costs) will be wider the lower the usage cost of the current option (O^A) is compared with the combined setup and usage costs of the alternative ($S^B + O^B$). For example, consider an initial choice between A and B in which option A has a first-time total cost of 35 (period 1: $S^A = 5$ and $O^A = 30$) and option B

has a first-time total cost of 55 (period 1: $S^B = 40$ and $O^B = 15$), with a difference of 20. Assume A was chosen in the first period. When consumers decide whether to switch in the second period, the cost structure will consist of a total cost of 30 (period 2: $S^A = 0$ and $O^A = 30$) for A and a total cost of 55 (period 2: $S^B = 40$ and $O^B = 15$) for B, with a difference of 25, which is an increase of 25% over the previous cost gap of 20. This simple example illustrates the shift in the cost structure following the first decision. Thus,

H1: Incurring a setup cost for a given option creates a future relative cost advantage for that option. Because of these switching costs and a preference for minimizing present costs, consumers are more likely to select the same option following previous selection than if that option was not previously selected.

Hypothesis 1 asserts that the effect of prior investment on lock-in is driven by its implication for subsequent switching costs. This perspective is distinct from the other well-documented effects of the extent of prior investment on future preference and choice (e.g., Arkes and Blumer 1985; Aronson and Mills 1959; Staw 1976). In general, these theories predict that the extent of commitment (i.e., lock-in) is greater if the choice has been made after extensive investment in search and deliberation. One category of such effects relates to the sunk-cost fallacy (e.g., Arkes and Blumer 1985) and research on escalation of commitment (e.g., Staw 1976) and focuses on people's inability to ignore unrecoverable prior costs. Another category of such effects relates to dissonance theory (Festinger 1957), in which greater effort results in increased liking (e.g., Aronson and Mills 1959). By contrast, in the current framework, lock-in is not a result of the extent of prior investment per se but, rather, is determined by the consequences of that investment for the subsequent usage occasion. Therefore,

H2: An initial investment will reduce consumers' propensity to search and switch because of its implications for future switching costs, not because of commitment.

Lock-in and Intertemporal Dynamics

To this point, the discussion of lock-in is consistent with the economic notion of switching costs. Next, I incorporate the implication of time preferences for switching following a selection. As a result of consumers' underweighing of future switching costs, these costs are perceived to be smaller the farther they are in the future. This temporal inconsistency is captured by a decrease in β the farther the next decision is in time (i.e., the greater D is). As D increases and β becomes smaller, the perceived switching costs (i.e., the perceived postinitial selection difference $U_2^A - U_2^B$) also become smaller. Moreover, to the extent that consumers do not anticipate and incorporate the temporal shift in importance weights into their decisions at period t, they will mispredict their likelihood of switching at $t + n$ after incurring a setup cost. Therefore, even if consumers select an alternative with only a single use in mind, they can become locked in if future usage occurs.

H3: The effects of an initial investment and the associated levels of switching costs on consumers' anticipated future switching intentions will be smaller than their effects on actual future switching behavior.

To test the proposed theory and the corresponding three hypotheses, I conducted three studies that employed different methodologies. Experiments 1 and 2 examined the effect of changes in the cost structure (S and O) on consumers' search and switching behavior (hypotheses 1–2). Experiment 3 examined directly the intertemporal aspects of the model (changes in β) and tested consumers' ability to anticipate the implication of initial investments on their future search and switching behavior (hypothesis 3).

EXPERIMENT 1: THE EFFECT OF PAST PURCHASE IN ONLINE SHOPPING

The goal of this experiment was to study the effect of changes in information costs on lock-in and to illustrate the minimal investment in S needed for lock-in to take place in a complex environment. Because this experiment involved an actual Internet search, it was impossible to manipulate setup costs compared with usage costs directly. Therefore, experiment 1 used a manipulation of prior experience as a proxy measure of the information cost. It was assumed that with an initial (trial 1) retailer selection, participants would incur a setup cost. Manipulating whether participants could or could not use their prior investment in a second (trial 2) retailer selection would then determine whether all other options had a relative disadvantage because switching would require them to incur another setup cost. The key test in this experiment was whether investment at trial 1 would change subsequent propensity to search and adopt superior, new alternatives (hypothesis 1).

Method

Participants and Design. A total of 57 undergraduate students participated for class credit and were randomly assigned to one of two conditions. The experimental design was a 2 (search restriction) × 2 (time) × 4 (stimulus domain) three-factor mixed design.

The between-participants factor was search restriction with two conditions: restricted search and unrestricted search. The search-restriction factor was implemented by allowing participants in one condition to search and purchase freely from any retailer in their second trial (unrestricted), while participants in the other condition were not allowed to make a second purchase from the retailer used in trial 1 (restricted).

The first within-participant factor was time, with two decision blocks. The second within-participant factor was domain replication. Each decision block included four purchase decisions, one in each of four product categories. Categories were pretested to include different levels of familiarity on participants' parts. The four categories were watches, books, music compact disks, and airplane tickets. The specific item for each product category was different for each of the two blocks (e.g., for watches: Swatch Skin, SFB101, in trial 1 and Seiko Kinetic, SKH297, in trial 2), for a total of eight products in four product categories.

Procedure. This experiment used an Internet search as the main task. Participants were given an instruction booklet that asked them to conduct a series of Internet searches on their own time. Participants first completed a questionnaire about their familiarity with the Internet and e-commerce in general and about their prior experience purchasing each of the relevant categories used in this study. These measures, however, did not have a significant effect and will not be discussed further. In each trial, participants were instructed to search for and select a specific item online, using their browser of choice, and to do so in the way that they normally would if they were actually purchasing the item. Participants were instructed to search sequentially for specific products in all four different product categories. Then they engaged in a second search for four products, searching for a different item selection in each category. Participants completed a total of eight trials, two in each product category.

The search-restriction manipulation was implemented using an instruction page between the two blocks for the restricted search group. Participants in the restricted search group were instructed not to purchase from the retailer from whom they had purchased the previous time. In the unrestricted condition, participants purchased all products sequentially, with no interruption between blocks. The two experimental conditions were identical in all other aspects.

For each product, participants went through the purchase process up to the point of submitting the purchase order (participants printed the shopping basket). Process measures included participants' reported search time and the number of retailers searched. The outcome measures included the selected retailer, the price paid for the item, and satisfaction with the retailer and the shopping experience.

Results

The most important results were from the second block of search items. (As expected, both search-restriction conditions were identical on all measures for the first block.) The raw score dependent variables for each category were converted to *z*-scores so that all categories would have equal effects on the repeated measures analyses. The results of the second trial indicated the expected impact of prior investment on a subsequent search. Participants in the unrestricted search condition searched less time (−.23 vs. .27; $t(215) = 3.80$, $p < .001$) and visited fewer retailers before purchase (−.15 vs. .17; $t(208) = 2.38$, $p < .02$) than did participants in the restricted condition.

I examined lock-in more directly by assessing the degree to which participants in the unrestricted condition used the same sites on both trials. Overall, participants used the same site 54.9% of the time. When the results are examined by category, three of the four categories showed significant lock-in (books = 70.0%, music CDs = 66.7%, and travel = 70.4%). The watch category did not show lock-in: participants used a different site 92.3% of the time (apparently because many Internet sites did not carry both watch models). Overall, these findings lend support to hypothesis 1.

The results also demonstrated the potential economic consequences of the lock-in effect. Prices on the first block were identical for the two groups (−.004 vs. .004; $t(204) = .06$, $p = .95$); however, prices paid in the second block were lower for the restricted search group than for the unrestricted search group (−.29 vs. .25; $t(205) = 4.04$, $p < .0001$). Last, the results for satisfaction with both the experience and the retailers indicated that lock-in was not associated with diminished satisfaction. On the contrary, participants in the unrestricted search condition were more satisfied with the retailer they had selected in the second block (−.31 vs. .27; $t(215) = -4.47$, $p < .001$) and with the entire shopping experience (−.34 vs. .30; $t(214) = -4.93$, $p < .001$).

Discussion of Experiment 1

This Internet experiment demonstrates the role of the initial investment in *S* and the associated switching costs in determining lock-in. Participants who were allowed to use prior relevant experience with a specific product type based their subsequent decisions on that experience and tended to return to the previously used retailer (cf. Hoyer and Brown 1994). Prior relevant investment leads participants to search less and pay more than those who cannot use a previously selected site. The low absolute levels of search were consistent with behavior observed on the Internet (e.g., Johnson et al., forthcoming). Moreover, it seems that the prior investment needed to produce lock-in was minimal. To ensure that this result results from the experience in the first block and not from experience prior to the experiment, for each product participants were asked whether they had ever used that online retailer before. Previous usage for the relevant first-trial unrestricted condition was the following: airline tickets, 8%; books, 22%; music compact disks, 15%; and watches, 0% (results were identical for both search-restriction conditions). These results indicate that only a small minority of participants used these online retailers before the experiment (the experiment was conducted in January 1999), and indeed the main experimental results hold when those participants with prior experience are excluded from the analysis. These results therefore refute the argument that prior usage levels, rather than the experimental treatment, produced the observed lock-in.

Another explanation for the results is that this task did not involve actual purchase or real money; because participants value their time, an implicit goal might have been to minimize

effort. On close inspection, however, this argument is unsatisfactory. With the random assignment of participants, if one group's goal was to minimize effort, the other group's goal should have been the same. Yet the restricted group systematically attained lower prices (i.e., they did not select randomly). Nevertheless, to examine this issue in more detail, experiment 2 includes both a direct manipulation of the information costs and real monetary consequences.

The satisfaction results indicate that participants in the unrestricted search group (i.e., those who paid a higher price) were more satisfied with the retailer and the entire experience than were participants in the restricted search group. This suggests that the cost associated with the search process for a retailer was psychologically bundled with the benefits of using that retailer. That is, the impact on satisfaction of the search cost advantage of the unrestricted group in search for a retailer outweighed the price advantage of the restricted group. Another explanation could be that this effect results from negative affect associated with the search restriction itself. Experiment 2 overcomes this alternative explanation by manipulating setup costs without the use of explicit restriction.

EXPERIMENT 2: THE EFFECT OF COST STRUCTURE ON LOCK-IN

Experiment 1 shows that a small initial investment is sufficient to produce lock-in. Experiment 2 directly examines the effect of the information cost structure (*S* and *O*) on lock-in under controlled laboratory conditions. This experiment tests consumers' propensity to search and switch as a function of the initial selection and the time horizon (more details below). Participants in trial 1 were asked to choose between a high setup, low usage cost option and a low setup, high usage cost option, which makes it possible to examine the consequences of this choice for search and propensity to switch in trial 2, when superior options enter the set.

The time horizon was defined as the number of transactions participants anticipated that they needed to complete. Participants in the long-term condition are expected to be more likely to invest in a high setup cost option in trial 1 than those in the short-term time horizon condition. This is not surprising, but it allows for a comparison of the effects of the time horizon at time 1 with the effects of time horizon at trial 2, when both participants in the long-term condition and those initially expecting a short time horizon are told that the time horizon is long. The choice of the high versus low setup cost option at time 1 is expected to determine both subsequent search and switching (hypothesis 1). That is, although the time horizon condition had a significant effect on trial 1 decisions, it was predicted to have no effect on trial 2 decisions, when trial 1 decisions were controlled for. Trial 2 decisions are expected to be driven by the relative cost of the previously selected alternative compared with the other previously unused alternatives (hypothesis 2). Note that I make no prediction about the absolute level of lock-in; this is obviously a matter of calibrating the cost structure and the expected benefits. Lock-in is expected to be more severe for respondents who chose the high setup cost option in trial 1 because, for those respondents, the relative advantage of switching is lower.

This experiment used a principal-agent task, in which participants chose hotels for a target consumer who was defined by a set of attribute importance weights. This enabled explicit monetary incentives to be included so that the impact of lock-in on consumer welfare could be studied. In particular, I predict that lock-in will result in the choice of lower utility options. Nevertheless, as in experiment 1, despite the selection of objectively inferior options, lock-in will not necessarily result in reduced satisfaction.

Method

Participants and Design. A total of 112 members of a university community were recruited through advertising. In addition to a $5 participation fee, participants could earn up to an additional $5 extra, contingent on their performance. Participants earned between $6.50 and $10 total for approximately 30 minutes of their time. The experimental design was a 2 (time horizon) × 2 (time) two-factor mixed design. The within-participant factor was time, with two decision blocks. The between-participant factor was time horizon, with two conditions: short-term + long-term and long-term. In the long-term condition, participants were told a priori that they would need to complete multiple purchase decisions. In the short-term + long-term condition, participants were initially instructed to select a provider for only a single use. After this first decision was completed, participants were informed that because of changed circumstances for their target consumer, they would need to engage in multiple additional purchase decisions (i.e., identical to the initial long-term manipulation). The rationale for this factor was to examine whether there would be any partial effect of the time horizon manipulation in future periods after the trial 1 choice was controlled for. In reality, all participants stopped after the second trial.

Stimulus Design. The stimulus included a travel portal that offered eight hotel providers, with each provider offering 15 hotel options. The set of hotel options was described by four attributes: accommodations, affordability, dining, and entertainment (see fig. 2). Calculated using an agent's attribute weights, each option was associated with a utility score that participants needed to infer. The providers were defined by a set of costs, operationalized by a trade-off between the setup effort required to use the provider (e.g., registration) and postselection usage costs (e.g., checkout). High setup costs required participants to select a login name and password and to provide personal and shipping information as well as to answer a short survey about their online shopping experiences. Participants were then able to start shopping by pressing a "shop" button. After this initial investment, participants logged in to start shopping for any subsequent purchase. In contrast, low setup costs required

FIGURE 2

EXPERIMENT 2: AN ILLUSTRATION OF THE HOTEL SELECTION TASK

A

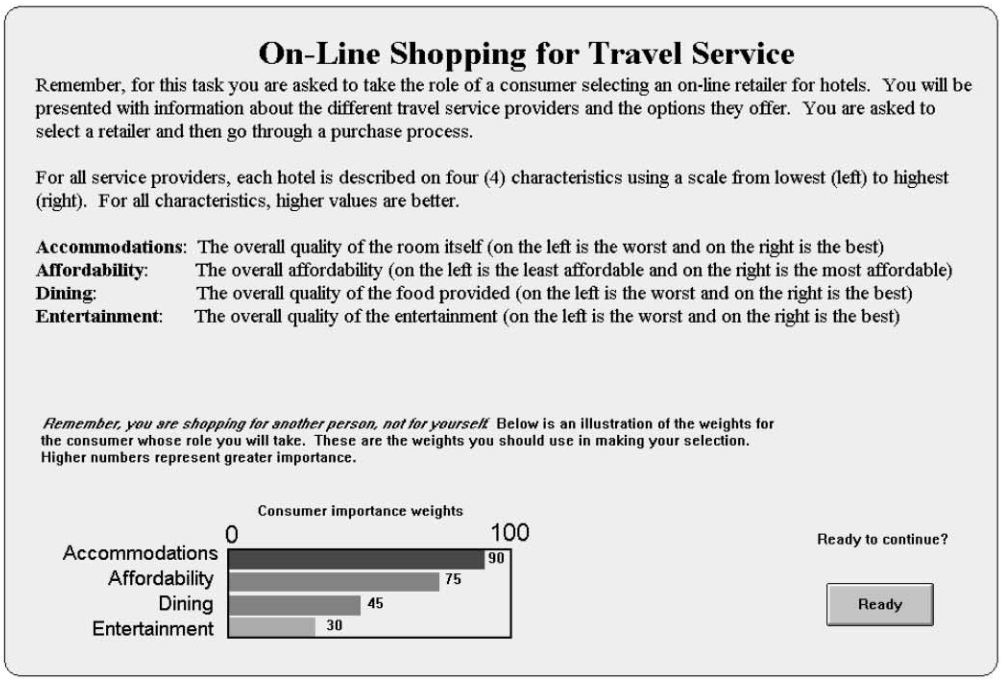

B

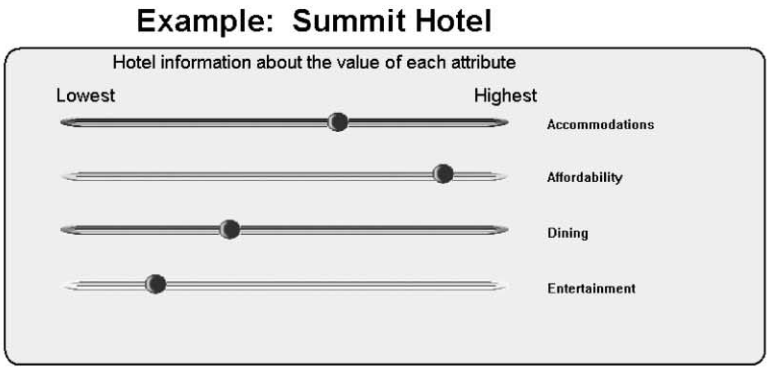

NOTE.—*A*, Principal-agent task illustration. *B*, Hotel option illustration.

participants to press only a "shop" button to get to the options and start shopping. As the ongoing usage cost manipulation, high usage costs required participants to submit their personal and shipping information after each purchase to complete the purchase. Low usage costs entailed only a one-click checkout process after selecting an option.

To test the idea that consumers forgo superior options as they get locked in, at trial 2 better options were introduced into the set of nonselected providers. On trial 1, four of the eight providers had a high setup and low usage cost structure, and the other four providers had a low setup and high usage cost structure. On trial 2, two new providers replaced two trial 1 providers: the chosen provider was maintained, but two of the remaining seven providers were randomly replaced. On trial 2, the provider chosen in trial 1 carried the same set of 15 options (low utility list). However, all

other, previously unused providers carried an improved set of options (high utility list). The improved list included six options that had a higher utility than the best available from the low utility list, and the best option had a utility that was 11% higher than the maximum utility available from the low utility list. The two new providers always had high setup and low usage costs.

Procedure. Participants were provided a CD-ROM with a computerized, repeated purchase shopping task that they completed on their own time at their preferred location. The data were written to a floppy disk that participants returned to receive payment. Each participant was required to act on behalf of a consumer who was shopping for hotels. Participants were asked to select a provider and purchase travel packages in a computerized, Internet-like environment. Participants first completed a computerized questionnaire about their familiarity with the Internet and e-commerce in general, and about their prior experience purchasing hotel services online. Again, however, these measures did not have a significant effect and will not be discussed further. After completing this initial part, participants were introduced to the main experimental task, which included the time horizon manipulation. Participants repeated the shopping task twice; the short-term + long-term group received the long-term manipulation instructions between the two trials. Between trial 1 and trial 2, participants in the short-term + long-term condition were told that the circumstances for the consumer for whom they were doing the purchasing had changed and that they would need to make several more purchases.

In trial 1, participants selected a hotel provider from the eight available providers and then selected a hotel option that best fit their target consumer. There were two elements to help participants assess the trade-off information when they selected a provider. A brief description of the provider was available without participants' needing to enter the site (a text box opened when the participant pointed to the provider's name with the mouse). An introduction to the provider (after participants entered that site) included a brief description of the cost structure. In trial 2, participants could choose either to see all available providers or to go directly to the provider they had used in the previous trial. After each trial, participants were asked several questions about their experience. At the completion of both trials, participants were asked questions about the entire experience, followed by questions about their age and gender.

Payoff. The contingent payoff is an important feature of the design because it provides both real incentives to locate the best alternative and a way to quantify the consequences of lock-in. Participants' pay was contingent on how well the selection of the hotels fit with the preferences of their target consumer from the principal-agent task. Participants were told that they would be paid $5 for their participation and could earn up to an additional $5, depending on how well their choices fit their target consumer. For each option, a utility score could be calculated as defined by the target consumer's attribute weights. The payoff for each of the two trials was between $0 and $2.50 and was calculated as follows:

$$\text{Pay} = (2.5 - \{2.5 \times [(\max \text{ utility} - \text{chosen utility})/(\max \text{ utility} - \min \text{ utility})]\}).$$

Dependent Measures. The dependent variables included both outcome and process measures. The outcome measures included the choice of a provider; quality of the chosen option; and satisfaction with the entire shopping experience, the provider, and the option selected. The other outcome measures related to the change in the provider selected and in satisfaction. The process measures included the number of providers considered (searched), total search time, and the number of options searched within each provider. Finally, a measure of participants' value for their time was collected at the end of the experiment. All participants were asked, "On average, how much time would you be willing to spend searching the Internet in order to save $25?" This was designed to examine whether a normative value of time argument can account for lock-in and its consequences.

Results

Trial 1 Search and Choice. As expected, the anticipated time horizon had a significant effect on search in trial 1 ($F(1, 110) = 7.26$, $p = .008$). Participants searched 62% more providers in the long-term condition (4.95) than in the short-term condition (3.05). Participants in the short-term condition were less likely to select the high setup cost option than participants in the long-term condition ($\chi^2(1) = 22.91$, $p < .0001$). In the short-term condition, the majority of participants selected the low setup cost option (high setup = 35.7%, low setup = 64.3%), but in the long-term condition, the majority selected the high setup cost option (high-setup = 80.4%, low-setup = 19.6%).

Trial 2 Search. The search results indicated that there was no significant difference in search in trial 2 between the two time-horizon conditions ($F(1, 110) = .39$, $p = .54$); participants searched 1.61 providers in the short-term + long-term condition and 1.82 providers in the long-term condition. Overall, the search level was significantly reduced from trial 1 (4.00) to trial 2 (1.71) ($t(111) = 6.54$, $p < .0001$). Note that cumulative search over both trials was marginally less for those in the short-term + long-term condition ($M = 3.05 + 1.61 = 4.66$) than for those in the long-term condition ($M = 4.95 + 1.82 = 6.77$) ($t(1, 110) = 1.74$, $p = .06$). That is, short-term + long-term participants who saved effort in trial 1 did not compensate with additional search when faced with a changed situation in trial 2.

Next, I examine whether participants in trial 2 conducted a new search or went directly to their previously selected provider. A logistic regression with direct access as the dependent variable and the time horizon and trial 1 choice (whether high or low setup) as independent variables showed

that the coefficient associated with the time-horizon factor (the experimental manipulation) was not significant (.69; $\chi^2 = 2.04$, $p = .15$). However, the coefficient associated with the trial 1 choice was statistically significant (1.81; $\chi^2 = 12.84$, $p < .001$); the higher the setup cost, the more likely participants were to go directly to their previously used provider (without any additional search), which lends further support to hypothesis 1.

Trial 2 Choice and Switching. A logistic regression, with stay-switch as the dependent variable and the time horizon and trial 1 choice (whether high or low setup) as independent variables further, supported hypothesis 1. The regression coefficient associated with the time-horizon factor (the experimental manipulation) was not significant ($-.64$; $\chi^2 = 1.70$, $p = .19$), but the coefficient associated with trial 1 choice was highly statistically significant (-3.67; $\chi^2 = 13.66$, $p < .001$). Consistent with the search data, across both time-horizon conditions, participants who selected the low setup cost option in trial 1 (low investment) switched 85.1% of the time, but participants who selected the high setup cost option in trial 1 (high investment) switched only 50.8% of the time (see fig. 3).

Prior Investment and Choice over Time. The results reported above show that effects of time horizon at trial 1 on switching behavior at trial 2 are fully mediated by the choice of a high or low setup cost options at trial 1 (and the associated switching costs). Hypothesis 2 asserts that the effect of time horizon and investment in setup costs are not mediated by the effort expanded by respondents (both in terms of the number of providers they searched and the time they searched), as might be predicted by commitment driven accounts.

A logistic regression with stay-switch as the dependent variable and the number of providers searched and trial 1 choice (whether participants chose high or low setup) as independent variables showed that the number of providers searched was not significant ($-.07$; $\chi^2 = 1.40$, $p = .24$), but the regression coefficient associated with trial 1 choice was statistically significant (-1.82; $\chi^2 = 13.50$, $p < .001$). A similar set of analyses was conducted using time spent shopping at time 1 as a mediator.[1] The regression coefficient associated with the search time was not significant (-1.13; $\chi^2 = 1.27$, $p = .26$), but the coefficient associated with trial 1 choice was statistically significant (-3.65; $\chi^2 = 13.31$, $p < .001$).

Neither the time nor the effort that it takes to select a provider can explain switching behavior. However, it is also possible to examine whether the time participants spent in their chosen provider at trial 1 affects switching. The total time consumers spent at the selected provider had a significant effect on switching (-2.14; $\chi^2 = 4.56$, $p = .03$); switching was less likely the more time consumers spent at the provider. However, when prior selection of a provider (whether high or low setup) was included in the analysis, the coefficient associated with the time at the selected provider was not significant (-1.54; $\chi^2 = 2.37$, $p = .12$), but the coefficient associated with trial 1 choice was highly statistically significant (-3.33; $\chi^2 = 11.08$, $p < .001$). These results directly support hypothesis 2.

FIGURE 3

EXPERIMENT 2: SWITCHING AS A FUNCTION OF TRIAL 1 DECISIONS

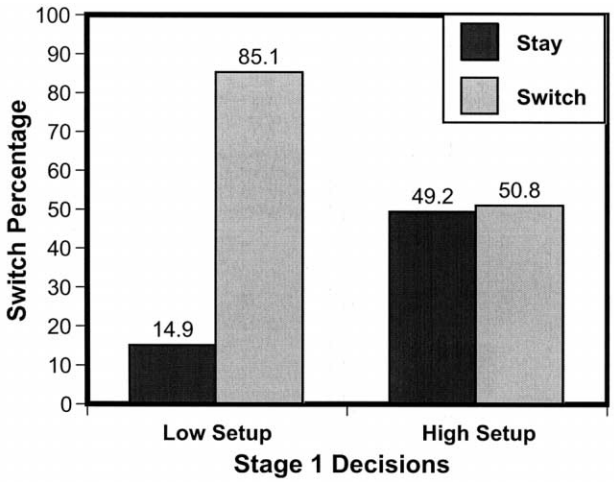

The next analyses examined in more detail the effect of the setup cost investment on subsequent choice by using timing data for different parts of the decision process. The total time that participants spent in a provider was partitioned into three parts: (1) time to set up, (2) time to process and select an alternative, and (3) time to check out. A logistic regression with stay-switch as the dependent variable and the three timing measures as independent variables further supported hypothesis 2. The regression coefficients associated with the alternative selection time ($-.003$; $\chi^2 = .00$, $p = .99$) and the checkout time ($-.48$; $\chi^2 = .23$, $p = .63$) were not significant, but the regression coefficient associated with setup time was statistically significant (-2.26; $\chi^2 = 5.10$, $p = .02$). The greater the setup time, the less the participants switched. Moreover, when prior selection of a provider was included in the analysis, the coefficients associated with these times were not significant (for all three components). Specifically, setup time was no longer significant (-1.34; $\chi^2 = 1.79$, $p = .18$), but the coefficient associated with trial 1 choice was highly statistically significant (-3.21; $\chi^2 = 10.31$, $p < .002$). Thus, this mediation analysis provides direct support that prior selection affected switching because of its implications for the later relative cost advantage of that option compared with other previously unused options.

Welfare Results. The final set of results focuses on the implications of lock-in for consumer welfare by examining

[1]Note that the arbitrary unit of time used (i.e., seconds, minutes, or hours) changes the magnitude of coefficients; therefore, the standard error–adjusted coefficients are reported for this and all other regression models that include time.

the monetary outcome of lock-in and reported satisfaction. The results show that participants who switched earned 36% more in trial 2 ($M = \$2.04$) than did those who were locked in ($M = \$1.50$) ($t(110) = 7.04$, $p < .0001$); note that these two groups did not differ in their trial 1 payoff ($t(110) = .74$, $p = .46$). In addition, at trial 2, participants who invested in the low setup cost option at trial 1 earned 11.5% more ($1.97) than those who invested in the high setup cost option ($1.77) ($t(110) = 2.33$, $p = .02$), because they were more likely to switch.

Although participants who were locked in selected an inferior option (based on the principal's utility function) and earned less money compared with those who switched, they were more satisfied with their choice of hotel, their choice of provider, and the entire shopping experience. On a 0–100 scale, participants who were locked in indicated significantly greater satisfaction with their chosen hotel (73.4) than did those who switched (65.8) ($t(110) = 2.11$, $p = .037$). The same pattern of results held for satisfaction with the provider (75.0 vs. 63.4; $t(110) = 3.15$, $p = .002$) and satisfaction with the entire shopping experience (71.2 vs. 63.4; $t(110) = 1.89$, $p = .062$). A logistic regression with stay-switch as the dependent variable and trial 1 satisfaction and trial 1 choice (whether high or low setup) as independent variables indicated that this was not simply a selection effect. Both the regression coefficient associated with trial 1 satisfaction ($-.026$, $p = .02$) and the regression coefficient associated with trial 1 choice (-1.85, $p = .0002$) were significant. Moreover, although trial 1 satisfaction had a significant effect on switching, this effect was much smaller than that of the trial 1 choice (whether high or low setup cost option) and the relative cost associated with it.

To complete the welfare argument, I tested whether participants' value for their time could explain lock-in. The reported value of time (the time they were willing to devote in order to save $25) did not differ between those who were locked in (53.5 minutes) and those who switched (41.9 minutes) ($U = 1{,}269$, $p = .40$).[2] Moreover, an analysis of covariance indicated that the value of time could not explain the differences in payoff between the two groups. The results indicated a significant main effect for stay-switch ($F(1, 107) = 28.27$, $p < .0001$) but no significant main effect of the value for time covariate ($F(1, 107) = 1.66$, $p = .20$) and no significant interaction ($F(1, 107) = .086$, $p = .77$).

Discussion of Experiment 2

The goal of this experiment was to examine the impact of changes in the information cost structure on initial choice, search, and switching over time. The key results supporting hypotheses 1 and 2 and the cost structure elements of the model are from trial 2. These results show how prior investment affects trial 2 decisions: the best predictor for trial 2 decisions is the type of provider selected in trial 1 (high or low setup), not the manipulation of the time horizon. As stated previously, the effect of time horizon on trial 1 decisions is not surprising, but it is important for understanding trial 2's results. The significant effect of time horizon in trial 1 indicates that participants were responding to the manipulation, yet trial 2's change in the time horizon alone cannot explain switching behavior. The switching data indicate that the time horizon manipulation had no partial effect on trial 2 switching when the choice of high versus low setup cost is controlled for. Moreover, when informed of a long time horizon that might justify further search (similar to trial 1 for long-term participants), the short-term + long-term participants, who searched less than long-term participants in trial 1 when they thought it did not matter, did not compensate in trial 2 for the search they had missed in trial 1.

As discussed previously, an alternative explanation to the effects of prior investment might be that the extent of commitment would have been greater if the choice had been made after extensive search and deliberation (e.g., Arkes and Blumer 1985; Aronson and Mills 1959; Staw 1976). However, that is not what the results show. Instead, in support of hypothesis 2, the results indicate that it is not the extent of search that matters, but whether respondents chose an option on the first occasion that had low or high setup costs (and, consequently, high or low switching costs). An examination of the time that participants spent in each of the three parts of the purchase process at their selected provider (time to set up, time to process and select an alternative, and time to check out) indicates that the only significant effect results from the setup time. When prior selection of a provider was included in the analysis, only the trial 1 choice (whether the high or low setup cost option) had any explanatory power. Thus, this internal analysis provides further direct evidence that prior selection affected switching because of its implications for the later relative cost advantage of that option compared with other previously unused options and not because of greater commitment.

This experiment also indicated that even in a controlled laboratory environment, lock-in has significant implications for consumer welfare like those reported in experiment 1. Locked-in consumers chose worse options and, hence, earned less money. They nonetheless expressed greater satisfaction with the hotel option that they selected, the provider, and the entire experience. A simple heterogeneity in value-of-time explanation cannot explain the difference between participants who switched and those who were locked in, nor can a self-selection explanation. Finally, it is noteworthy that trial 1 satisfaction was a significant predictor of future switching, though it had a smaller effect compared with the information cost structure. Nevertheless, these findings demonstrate that satisfaction is a dynamic process that should be studied over time (cf. Fournier and Mick 1999).

EXPERIMENT 3: PREDICTED SEARCH AND SWITCHING BEHAVIOR

The first two experiments focused on how initial investment in setup costs creates switching costs that then lead to

[2] Because of differences in the variance between the two groups, I report a Mann-Whitney nonparametric test. A standard contrast yielded the same results ($t(109) = 1.24$, $p = .22$).

lock-in (hypotheses 1–2). However, the first two experiments did not test the intertemporal dynamics of the model. Experiment 3 directly tests whether consumers anticipate the impact of switching costs generated by initial selection on subsequent choice. Hypothesis 3 asserts that consumers do not appreciate that switching costs from an initial investment will loom larger in the future than they appear at the time that setup investments are initially made. In terms of the model presented earlier, β increases as the decision to switch draws nearer (as D is getting smaller).

Prior research shows that future effort (feasibility) is discounted more than future benefits (desirability) (Soman 1998; Trope and Liberman 2000). If consumers do not understand this, they will overpredict their likelihood to search and switch in the future. Can this inability to anticipate correctly the impact of future switching costs help understand the results observed in experiment 2? Will participants be able to anticipate the effect of prior investment on search and switching and, more important, will the size of the initial investment moderate the degree of misprediction? It is hypothesized that regardless of their initial choice, participants will show similar intentions (prediction) to initiate new search and to switch, with little sensitivity to future switching costs. On the other hand, I also predict that actual behavior will be significantly affected by the switching cost created by the initial selection.

Method

Participants and Design. Seventy undergraduate students participated in the experiment as part of a research requirement. They were randomly assigned to two conditions in a 2 (prediction) × 2 (time) two-factor mixed design. The within-participant factor was time, with two decision blocks. The between-participants manipulated factor was prediction and included two conditions: prediction and no prediction. After trial 1, half of the participants were asked to predict their trial 2 behavior (whether they would engage in additional search, what the extent of the search would be, and whether they anticipated switching), and half proceeded directly to the second trial.

Procedure. The stimulus and procedure were similar to the short-term + long-term condition in experiment 2; therefore, only the main differences are presented here. As before, participants were asked to shop for a hotel for another consumer whose preferences were defined with specified attribute weights. On the first trial, participants selected a provider for a single use from eight available providers and then selected a hotel option that best fit their target consumer. Each provider offered 15 hotel options and was defined by a set of information costs. After trial 1 and before trial 2, all participants were told that the circumstances of the consumer for whom they were purchasing had changed and that they would need to make several more purchases (identical to the short-term + long-term condition in experiment 2). At this point, as the main manipulation, half of the participants predicted whether they would engage in additional search and whether they anticipated switching. The other half proceeded directly to the second trial, where they had the option of either to see a list of all available providers or to go directly to the provider they had used in the previous trial. The other elements of the task were identical in both trials.

Results and Discussion of Experiment 3

The results of experiment 3 demonstrate that participants overpredict their propensity to initiate search for new alternatives and their likelihood to switch at a later time, which is consistent with prior research. The results show that participants in the prediction condition indicated greater likelihood to search for available providers after trial 1 (82.9% new search) than the observed trial 2 search behavior (57.1% new search) ($\chi^2(1) = 5.51$, $p < .02$). A similar relationship was found between predicted and actual switching: participants predicted marginally greater likelihood of switching (74.3% switch) than the actual trial 2 switching behavior (54.3% switch) ($\chi^2(1) = 3.05$, $p = .08$).

The most important aspect of the results comes from a two-stage internal analysis comparing predicted and actual frequency of trial 2 search and switching as a function of trial 1 investment (whether consumers selected high or low setup cost options). Parallel two-stage analyses were conducted for predicted versus actual search and predicted versus actual switching. In the first stage of the analysis, for each dependent variable, two separate standard 2 × 2 contingency tables were computed to examine the likelihood of searching or switching as a function of whether participants selected a high or low setup cost option in trial 1 for each of the between-subjects conditions (predicted and actual). In the second stage of the analysis, the observed frequency for the prediction condition and the actual trial 2 behavior are compared with the pooled observations that are used as the expected frequency for a second 2 × 2 contingency table. This analysis assesses whether any differences between the effect of initial investment on the predicted behavior versus on actual behavior result from sampling error. This chi-square test is, in effect, a test of the two-way interaction between prediction and actual behavior and high versus low trial 1 investment. That is, by determining whether the predicted switching frequency differs from the actual switching frequency across conditions, the sensitivity of the prediction accuracy to initial investment can be observed.

First, consider the data on actual versus predicted likelihood to initiate new search in the second trial. The results indicate that the predicted likelihood to initiate a new search is not significantly different for the two prior investment conditions ($\chi^2(1) = 1.7$, $p = .19$). Participants predicted similar likelihoods to search for available providers in the low setup cost group (91% new search) and the high setup cost group (75% new search). By contrast, the actual likelihoods for initiating a new search are significantly different for the two prior investment conditions ($\chi^2(1) = 4.4$, $p = .04$). Significantly more participants initiated a search for available providers in the low setup cost group (71% new

search) than in the high setup cost group (36% new search). In the second stage analysis, a comparison of the two groups showed a significantly smaller effect of prior investment on predicted than on actual new search ($\chi^2(1) = 7.6$, $p = .006$).

A similar pattern of results was found for the switching data (see fig. 4). Participants demonstrated statistically similar likelihoods to switch in the future in the low setup cost group (83% switch) and the high setup cost group (58% switch) ($\chi^2(1) = 2.4$, $p = .12$). By contrast, the actual rate of switching is significantly greater for participants choosing the low setup cost options (71% switch) compared with those choosing the high setup cost options (29% switch) ($\chi^2(1) = 6.2$, $p = .01$). Again, in the second stage analysis, a comparison of the two groups showed a significantly smaller effect of prior investment on predicted than on actual switching ($\chi^2(1) = 3.8$, $p = .05$).

In summary, experiment 3 demonstrates that participants overpredict how likely they are to initiate additional search and how likely they are to switch. More important, the results show that the initial selection (whether participants select a high or a low setup cost option) affects trial 2 search and switching behavior to a greater degree than expected by respondents. In support of hypothesis 3 and the intertemporal dynamics of the model (unanticipated increase of β as the decision to switch draws nearer), the results demonstrate how prior investment and the resulting switching costs moderate the relationship between predicted and actual behavior. Whether a high or low setup cost option was initially selected had a significantly smaller effect on prediction of future search and switching than on the actual behavior. Participants did not fully internalize the future implications of their decisions.

Finally, note that the actual temporal distance between trials was relatively short. Yet the results still demonstrate overprediction of investment (between-participants), which is consistent with prior research. Therefore, this experiment provides a conservative test of this effect. The magnitude of this inconsistency is expected only to increase with greater temporal distance.

GENERAL DISCUSSION

Examining the effects of the information cost structure on lock-in was motivated by the low switching rates observed in the low search cost environment of the Internet. The theory presented in this article, however, is more general. Conceptually, the focus is on the combined effect of the information environment and time preferences on consumer search patterns and choice behavior over time. A key element of the proposed model is that the information cost structure can be conceptualized in terms of two temporally distinct categories that are traded off: initial setup costs (S) and ongoing usage costs (O). The overarching hypothesis is that the relative difference between information costs over time and consumers' time preferences explains consumers' reluctance to search extensively for better alternatives and ultimately switch. Lock-in is induced by a preference to minimize immediate costs and a failure to anticipate the impact of future switching costs.

The proposed model captures this dynamic and includes two aspects of the lock-in process. First, consumers tend to minimize present costs, as if they weigh only (total) costs in the current period (i.e., β is relatively high). Second, consumers do not fully incorporate future switching costs into their decisions (i.e., β increases the closer the next decision is in time). When the future arrives, switching costs exert more powerful effects than consumers originally anticipate. The findings of the three experiments provide support for this model.

FIGURE 4

EXPERIMENT 3: PREDICTED AND ACTUAL PERCENTAGE OF PARTICIPANTS INITIATING NEW SEARCH AND SWITCHING AS A FUNCTION OF THE TRIAL 1 DECISIONS

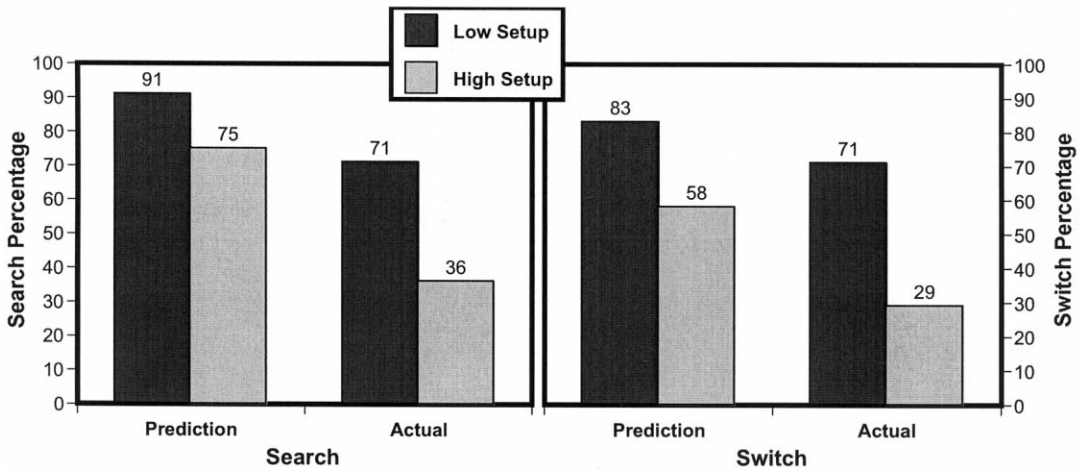

The results of experiments 1 and 2 demonstrate that the relative cost structure, whether manipulated or because of prior investment, leads to lock-in. Experiment 1 also demonstrates in an Internet setting how a minimal investment can lead to lock-in by changing the relative costs. Experiment 2 demonstrates the importance of the choice of high versus low setup cost in the initial period for subsequent search and switching behavior. A setup cost, once incurred, is associated with a lower ongoing usage cost, and consumers tend not to expend effort to look for and find superior new alternatives in future periods. Experiment 3 demonstrates that participants overestimate their propensity to search and switch at a later point. More important, the results demonstrate that prior investment and the resulting switching costs have a significantly smaller effect on consumers' prediction of future search and switching than on their actual behavior. Participants did not internalize the future implications of their prior decision and failed to incorporate the full impact of their decisions on later behavior. Switching seems easier and more likely in the future than it is at the time of the decision. This phenomenon could potentially lead people to delay switching perpetually—like a scene from the movie *Groundhog Day*, with a twist.

The results of experiments 1 and 2 also point to several consequences of lock-in for consumer welfare. Both experiments indicate that lock-in reduces the quality of the selected option, but the participants who were locked in also expressed greater satisfaction with their chosen provider and the entire shopping experience. Experiment 2 also showed that participants who were locked in were more satisfied with their chosen (yet inferior) option. These results are consistent with the view that customer satisfaction is not defined by the narrow economic value of the option, but the entire experience. That is, the search experience is psychologically bundled with other aspects of the shopping experience, including the utility of the selected option. The link between lock-in and satisfaction merits further research.

Taken together, the three experiments support the thesis that consumers' time preferences and relative costs are critical factors in determining the initial selection of providers and subsequent switching decisions. The current framework is distinct from other research examining the effect of prior investment on subsequent decisions. As discussed previously, other explanations for the effect of prior investment on subsequent choice are backward oriented and relate to the extent of past investments per se (e.g., Arkes and Blumer 1985; Aronson and Mills 1959; Staw 1976). By contrast, the theory presented here is a forward-oriented perspective, centered on the implications of past investment for the subsequent gap between the total cost of using the previously chosen provider and the cost of using other providers. The results of experiment 2 provide direct support for this forward-oriented effect and demonstrate that it is not how much is invested (time or number of providers searched), but how the investment (high or low setup) alters the cost structure. This is not to say that the pure extent of past investment does not play an important part in repeated choice decisions in the real world but, rather, that it is not the entire story. The current model provides one mechanism for why past investment matters.

The current findings relate to and are consistent with research on melioration reviewed earlier (e.g., Herrnstein et al. 1993). As in the melioration research, I also find that when consumers select among alternatives with costs and benefits over time, they focus on current higher utility and are insensitive to the effects of their own choices on the relative utility of options in the set. These general effects are well documented in psychology but have not received much attention in consumer research. The lock-in model presented here incorporates these ideas.

Limitations, Extensions, and Boundary Conditions

The model proposed here highlights the dynamics of the cost structure over time and consumers' time preferences as important determinants of lock-in. This model, however, is not intended as a comprehensive account of consumer loyalty; rather, it characterizes an important element in a multidimensional phenomenon. In addition, the model can be extended to address limitations and boundary conditions. First, the discussion thus far has referred to actual search and transaction costs. However, the current model can be easily generalized to incorporate other more subjective, affective, and cognitive costs. For example, initial setup costs (S) can include psychological aversion to new alternatives and the comfort of the familiar. Setup costs can also include the memory advantage of past selections. Research is needed to investigate whether time preferences for different types of resources obey the same rules demonstrated here.

A second important extension of the model would be to relax the assumption that the benefit from the different alternatives (Q) is fixed and that the ongoing costs (O) are constant. For example, dynamics could be incorporated, with Q increasing with time because of loyalty programs and customization. With respect to the assumption that O is constant over time, Johnson, Bellman, and Lohse (2003) find that learning can facilitate lock-in and that the power-law shape of the learning curve can lead to a relative usage cost advantage for previously used sites compared with new sites. In the context of the present framework, the learning process would lead to ongoing usage costs that were reduced over time. The assumption of constant ongoing usage costs over time (i.e., no learning) is therefore a conservative assumption. Learning would only lead to an even stronger lock-in effect because it increases the total usage cost gap between the currently used option and other alternatives. Personalization and smart agents are other mechanisms that can affect this cost gap because of changes in ongoing usage costs. Research is needed to examine the degree to which consumers understand how cooperation with such marketer-provided decision aids can lock them in to specific providers.

In conclusion, the theoretical model developed in this article combines concepts from intertemporal choice and switching and search behavior to provide an understanding of lock-in. An important general message of this work is that dynamic intertemporal preferences are very powerful

determinants of behavior. Consumers do not appreciate how powerfully they will be affected in the future by current investments in a specific alternative. This has parallels to work by Loewenstein (1996), who has drawn attention to the fact that visceral factors like hunger, sexual desire, and sleep deprivation cause stimuli in close proximity to be far more motivating than anticipated when they are at greater (temporal or physical) distance; consequently, people overconsume these stimuli relative to what they want when reflecting coolly. Just as in the work I have just reported, consumers do not learn from experience about the dynamic changes in their preferences as a function of proximity. It is our challenge to incorporate these robust psychological tendencies further into our behavioral models.

[*David Glen Mick served as editor and William O. Bearden served as associate editor for this article.*]

REFERENCES

Akerlof, George A. (1991), "Procrastination and Obedience," *American Economic Review*, 81 (May), 1–19.

Alba, Joseph, John J. Lynch, Barton Weitz, Chris Janiszewski, Richard Lutz, Alan Sawyer, and Stacy Wood (1997), "Interactive Home Shopping: Consumer, Retailer, and Manufacturer Incentives to Participate in Electronic Marketplaces," *Journal of Marketing*, 61 (July), 38–53.

Arkes, Hal R. and Catherine Blumer (1985), "The Psychology of Sunk Cost," *Organizational Behavior and Human Decision Processes*, 35 (February), 124–140.

Aronson, Elliot and Judson Mills (1959), "The Effect of Severity of Initiation on Liking for a Group," *Journal of Abnormal and Social Psychology*, 59 (September), 177–181.

Ausubel, Lawrence M. (1991), "The Failure of Competition in the Credit Card Market," *American Economic Review*, 81 (March), 50–81.

Bakos, J. Yannis (1997), "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, 43 (December), 1676–1692.

Brynjolfsson, Erik and Michael D. Smith (2000), "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science*, 46 (April), 563–585.

Cripps, John D. and Robert J. Meyer (1994), "Heuristics and Biases in Timing the Replacement of Durable Products," *Journal of Consumer Research*, 21 (September), 304–318.

Diehl, Kristin, Laura J. Kornish, and John G. Lynch (2003), "Smart Agents: When Lower Search Costs for Quality Information Increase Price Sensitivity," *Journal of Consumer Research*, 30 (June), 56–71.

Festinger, Leon (1957), *A Theory of Cognitive Dissonance*, Stanford, Calif.: Stanford University Press.

Fournier, Susan and David Glen Mick (1999), "Rediscovering Satisfaction," *Journal of Marketing*, 63 (October), 5–23.

Fournier, Susan and Julie L. Yao (1997), "Reviving Brand Loyalty: A Reconceptualization within the Framework of Consumer-Brand Relationships," *International Journal of Research in Marketing*, 14 (December), 451–472.

Frederick, Shane, George F. Loewenstein, and Ted O'Donoghue (2002), "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40 (June), 351–401.

Hausman, Jerry A. (1979), "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," *Bell Journal of Economics*, 10 (Spring), 33–54.

Herrnstein, R. J., George Loewenstein, Drazen Prelec, and William Vaughan (1993), "Utility Maximization and Melioration: Internalities in Individual Choice," *Journal of Behavioral Decision Making*, 6 (July), 149–185.

Hoyer, Wayne D. and Steven P. Brown (1994), "Effects of Brand Awareness on Choice for a Common, Repeat-Purchase Product," *Journal of Consumer Research*, 17 (September), 141–148.

Jacoby, Jacob and Robert W. Chestnut (1978), *Brand Loyalty: Measurement and Management*, New York: Wiley.

Jacoby, Jacob and David B. Kyner (1973), "Brand Loyalty vs. Repeat Purchasing Behavior," *Journal of Marketing Research*, 10 (February), 1–9.

Johnson, Eric J., Steven Bellman, and Gerald L. Lohse (2003), "Cognitive Lock In and the Power Law of Practice," *Journal of Marketing*, 67 (April), 62–75.

Johnson, Eric J., Wendy Moe, Peter Fader, Steven Bellman, and Gerald L. Lohse (forthcoming), "On Depth and Dynamics of Online Search Behavior," *Management Science*.

Kahneman, Daniel and Dan Lovallo (1993), "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking," *Management Science*, 39 (January), 17–31.

Klemperer, Paul (1987), "Markets with Consumer Switching Costs," *Quarterly Journal of Economics*, 102 (May), 375–394.

Laibson, David (1997), "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112 (May), 443–478.

Lal, Rajiv and Miklos Sarvary (1999). "When and How Is the Internet Likely to Decrease Price Competition," *Marketing Science*, 18 (4), 485–503.

Loewenstein, George F. (1996), "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Processes*, 65 (February), 272–292.

Loewenstein, George F. and Drazen Prelec (1992), "Anomalies in Intertemporal Choice: Evidence and Interpretation," in *Choice over Time*, ed. George F. Loewenstein and Jon Elster, New York: Russell Sage, 3–34.

Loewenstein, George F. and David Schkade (1999), "Wouldn't It Be Nice? Predicting Future Feeling," in *Well-Being: The Foundation of Hedonic Psychology*, ed. Daniel D. Kahneman and Norbert Schwarz, New York: Russell Sage, 85–105.

Lynch, John G. and Dan Ariely (2000), "Wine Online: Search Costs and Competition on Price, Quality, and Distribution," *Marketing Science*, 19 (Winter), 83–103.

Nunes, Joseph C. (2000), "A Cognitive Model of People's Usage Estimations," *Journal of Marketing Research*, 37 (November), 397–409.

O'Donoghue, Ted and Matthew Rabin (1999), "Doing It Now or Later," *American Economic Review*, 89 (March), 103–24.

Oliver, Richard L. (1999), "Whence Consumer Loyalty?" *Journal of Marketing* ("Whence Consumer Loyalty," ed. George S. Day and David B. Montgomery), 63 (4), 33–44.

Shapiro, Carl and Hal R. Varian (1999), *Information Rules: A Strategic Guide to the New Economy*, Boston: Harvard Business School Press.

Soman, Dilip (1998), "The Illusion of Delayed Incentives: Evaluating Future Effort-Money Transactions," *Journal of Marketing Research*, 35 (November), 427–437.

Staw, Barry M. (1976), "Knee-Deep in Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action," *Organizational Behavior and Human Decision Processes*, 16 (June), 27–44.

Stigler, George J. (1961), "The Economics of Information," *Journal of Political Economy*, 69 (June), 213–225.

Stigler, George and Gary S. Becker (1977), "De Gustibus Non Est Disputandum," *American Economic Review*, 67 (March), 76–90.

Thaler, Richard H. (1981), "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, 8 (3), 201–207.

Trope, Yaacov and Nira Liberman (2000), "Temporal Construal and Time-Dependent Changes in Preferences," *Journal of Personality and Social Psychology*, 79 (December), 876–889.