# The Enhancement Bias in Consumer Decisions to Adopt and Utilize Product Innovations

Shenghui Zhao
Robert J. Meyer
Jin K. Han*

*The authors are listed reverse-alphabetically. Shenghui Zhao is a doctoral candidate in Marketing at the Wharton School of Business, University of Pennsylvania, Robert Meyer is the Gayfryd Steinberg Professor of Marketing at the same institution, and Jin Han is an Associate Professor of Marketing at Singapore Management University. This work was supported by a 2003 research grant from the Wharton-Singapore Management University Research Center. Correspondence should be directed to the second author at meyerr@wharton.upenn.edu.

# The Enhancement Bias in Consumer Decisions to Adopt and Utilize Product Innovations

## Abstract

The ability of consumers to anticipate the value they will draw from new product generations that expand the capabilities of incumbent goods is explored. Drawing on previous research in affective forecasting, the work explores a hypothesis that consumers will frequently overestimate the benefits they envision drawing from new added product features and underestimate the learning costs required to realize those benefits. This hypothesis is tested using a computer simulation in which subjects are trained to play a *Pacman*-like arcade game where icons are moved over a screen by different forms of tactile controls. Respondents are then given the option to play a series of games for money with an incumbent game platform or pay to play with an alternative version that offer either expanded (Experiments 1 and 2) or simplified (Experiment 3) sets of controls. As hypothesized, subjects displayed an upwardly-biased valuation of the new sets of controls as measured by actual versus forecasted usage rates and performance gains. Yet, when given the opportunity to be paid to trade down to a more efficient device in exchange, few accepted. We thus observe a paradox where the presence of forecasting mistakes in product adoptions does little to induce regret in ownership.

According to a 2003 *Harris Poll,* 45% of American cell phone owners have never used voice mail, and 50% have never exercised the option to set their phones to silent or vibrate[1]. Yet, we suspect that few of these same decision makers would have considered buying a phone that did not have these functions. This apparent paradox illustrates what we call the *enhancement bias* in new product adoption decisions: given the opportunity to purchase new generations of an existing product with more bells and whistles, consumers may display an eagerness to adopt that goes beyond that which could be reasonably justified by their later utilization of these features and the real increase in happiness they yield.

Although consumer evaluation of product innovations and their adoption behavior has long been one of the focal topics in consumer and marketing research (see, e.g., Gatignon and Robertson 1991), to our knowledge the paradox of buying and usage described above has not been empirically examined. The purpose of this paper is to undertake such an investigation, exploring biases that arise when consumers who are trained to use one generation of products are given the opportunity to purchase and utilize an enhanced generation that enriches the way that the product can be utilized.

Our central thesis is that when exposed to new products that offer enhanced features consumers will over-project the degree to which their initial positive reactions will hold in the future, a bias similar to that which has been observed in other domains of affective forecasting (e.g., Loewenstein and Schkade 1999; Wilson and Gilbert 2003). This bias is fostered by a tendency to insufficiently project the contrary likelihood that future usage and pleasure will be suppressed by factors such as learning costs (e.g., Johnson, Bell, and Lhose 2003; Mukherjee and Hoyer 2001; Zauberman 2003). The consequence is a systematic bias in which consumers are

---

[1] http://biz.yahoo.com/prnews/030701/sftu019a_2.html

attracted to new product generations by the allure of innovative features, but then under-utilize them after adoption.

The remainder of the paper is organized as follows. We first provide a more formal description of the innovation adoption problem that forms the focus of our work. We then review prior literature on how it might be solved by consumers and draw hypotheses about how actual decisions may depart from the normative solution. We then report the results of three experiments that test these hypotheses. We conclude with a general discussion of the implications of this work for both basic research in consumer response to new technologies as well applied work in new-product design.

## Normative and Descriptive Elements of New-Product Adoption Decisions

In this work we consider how consumers solve a class of new-product adoption problems with the following structure. A consumer currently owns a durable good that conveys utility through the utilization of a set of features (such as options in software or capabilities of a home entertainment device). A manufacturer offers the consumer the opportunity to purchase an enhanced version that retains the features of the old but adds a new set of discrete attributes of uncertain value. The existence of these new attributes does not affect the functionality or utility derived of the older attributes, however, the consumer must choose to use either the new or the old features at any given point in time. Analogous cases include software packages that provide users with the option to utilize either older or newer interfaces (similar he to *Windows XP*), or digital cameras that give users the option to operate it with a basic set of automatic settings or a more advanced set of customized settings. In addition, we assume that the product will be privately consumed, excluding social-conformity considerations in the adoption decision.[2]

This problem and its normative solution can be more formally stated as follows. For simplicity assume that the utility from consuming an incumbent good with an attribute set $\alpha$ at any point in time $t$ is scaled to be $0$. Let $d_t \in \{0,1\}$ denote the consumer's decision whether or not to utilize some new feature set $\delta$ given its ownership at time $t$, let $x_t = u(\delta)_t - c(\delta)_t$ be the net utility that is realized from a decision to utilize $\delta$ at $t$ where $c(\delta)_t$ is a single-period usage or learning cost. Let $z_t$ denote the consumer's beliefs about the probability distribution associated with $x_t$[3]. In addition, let $\pi_T = d_0,...d_T$ be a sequence of attribute-usage decisions defined over a $T$-period ownership horizon, and let $V_0(\pi_t)$ be the total discounted expected utility implied by this sequence, defined as follows:

$$V_0(\pi_T) = E_0 \sum_{t=0}^{T} \beta^t v(x_t, z_t, d_t) \quad , \quad 0<\beta<1 \qquad (1)$$

where β is the discount factor. If a consumer were to optimally solve this problem, he or she would first try to find that sequential usage policy $\pi_T^*$ that maximizes expression (1), yielding an optimal ownership valuation ($V^* = V_0\ (\pi_T \mid \pi_T^*)$). The consumer would then buy the innovation if it were offered at a price less than $V^*$ (the expected lifetime value of the new features).

## The Nature of Intuitive Solutions

Consumers, of course, would be unlikely to make new-product adoption decisions as described above. To apply the normative solution the consumer would need to possess good

---

[2] In some cases consumers may buy enhanced versions of new products not because its features are assumed to convey benefits but because of desires to achieve social conformity. While real, the influence of such considerations is not taken up in this work.

[3] In a standard analysis these beliefs would be assumed to evolve as a Markov process given decision to utilize $\delta$; that is, associated with $z_t$ is a first-order cumulative conditional distribution function $G(z',z)$.

skills not only in intuitive dynamic programming (to derive the optimal ownership policy $\pi_T^*$), but also hedonic and utilitarian forecasting—accurately anticipating the various possible states of long-term pleasure one might come to associate with a new technology (the distribution over net asymptotic values of $u(\delta)-c(\delta)$) as well as the sequence of utilization decisions one will take in the future ($\pi_T$). While the literature is replete with examples of intuitive decisions that closely correspond with those prescribed by highly complex normative models (e.g. Hogarth 1981; Meyer and Hutchinson, 2001), this same robustness may *not* extend to intuitive forecasting. We will briefly review lines of evidence suggesting systematic biases that may arise when consumers attempt to develop two kinds of forecasts central to the above normative solution: 1) forecasts of the mean potential value of an innovative attribute (beliefs about $x$ and $z$); and 2) forecasts of the dynamic utilization of the new attribute (beliefs about the decision policy $\pi_T$).

## Intuitive forecasts of new attribute values: what's a deluxe widget worth?

Few would contest the suggestion that new products often have intrinsic appeal. New rides at amusement parks have the longest lines, movies have their greatest attendance on the weekend of first release, and we are lured by new flavors in ice cream shops. In some of these instances, of course, the allure of novelty has a rational basis: one *should* try new options because doing so provides information about product quality that can be exploited in future choices. Yet there is also evidence that the appeal of novelty often goes beyond what which can be explained by rational curiosity alone. Miller and Kahn (2003), for example, show that merely affixing novel names to the color or flavor of an otherwise familiar product can enhance its predicted quality among consumers. They explain this finding as arising from conversational norms (Grice 1975): consumers have simply come to associate novel names and appearances as the means by which firms communicate superior quality in markets.

In the current case positive reactions to a new-product generation would be further enhanced by the fact that there is a tangible basis for appeal aside from sheer novelty: that accruing to the expanded set of capabilities offered by the innovation. Even if the consumer were unsure whether these new product features enhanced value, there still might be rational reason for acquiring it. Specifically, expression (1) implies that the normative value of a novel attribute is not simply the consumer's best guess about future utility applied over a time horizon, but rather its *option value* over that horizon. That is, the expected long-term valuation drawn by a consumer who discovers a new attribute's value, and then uses this knowledge to make an optimal series of day-to-day decisions about whether or not to use it. Hence, a rational consumer who sees little merit in owning a new attribute in the short run might nevertheless be prescribed to acquire it simply to exploit the chance that it will prove valuable later.

There is considerable evidence suggesting that not only will consumers be attracted to alternatives that offer option values, but that this attraction might exceed that prescribed by rational analysis. To illustrate, Shin and Ariely (2003) report data from a sequential search task in which people are willing to pay to keep search options open even when the odds of their being utilized later is negligible—a result consistent with the strong preference that decision makers display for revocability observed by Gilbert and Ebert (2002). Likewise, Simonson (1990) and Loewenstein and Adler (1996) document a diversification bias whereby people choose a wider assortment of products when deciding now for the future compared to the assortment that they actually choose when decisions are made individually over time. A natural extension of this research is that consumers will be particularly attracted to new products that offer flexibility in the set of attributes that define their use—flexibility that provides an escape hatch in the event new attributes prove less valuable than was first envisioned.

Yet, while consumers may find themselves strongly attracted to the prospect of enhanced new product generations at the outset, it is by no means foregone that this visceral reaction will translate to a positive adoption decision. As suggested by Mukherjee and Hoyer (2001), this appeal can vanish if consumers contemplate the more negative realities of ownership, such as recalling that the joys of new ownership often wear off quickly, and that learning to utilize new features can be costly and frustrating.

Will these contrary considerations offset positive initial assessment biases? We suggest that in most cases it will not. The first basis is the large body evidence showing that while decision makers are generally correct about the valence of future affective states (i.e., whether one will feel happy or unhappy in the future), they tend to over-project the magnitude and duration of these feelings in forecasts (the *impact bias* described by Gilbert *et. al* (1998); Wilson and Gilbert (2003)). Applied to the context of product innovations, the implication is that consumers who have a positive initial reaction to a new product will overly over-project this sense of affect when forecasting the utility that the good will yield in the distant future.

Why does over-prediction occur? While several mechanisms have been proposed (e.g., the anchor-and-adjustment process hypothesized by Loewenstein, O'Donoghue, and Rabin 2003), most center on the idea strong affective reactions to a primary event inhibit our ability to process cues that are incongruent with that affect—a bias that Wilson and Gilbert's (2003) term *focalism.* In principle, of course, a consumer who is considering acquiring a new product should be just as inclined to contemplate the reasons why it *should not* be bought as the reasons why it *should.* Yet, focalism predicts that anticipatory reasoning will not be balanced; if a consumer's initial reaction to a new product is positive, it will be much easier to imagine future product use in a way that is consistent with this prior than inconsistent —hence an over-projection of value.[4]

In addition, even if consumers were to try to objectively consider the learning costs of adoption, the predicted size of these costs may be downwardly biased by recency effects in memory (e.g., Kahneman 1999). Recall that the greatest costs of learning to use a new tool or device are typically incurred during the early days of ownership, when the power-law of achievement will be its steepest (Johnson, Bell, and Lhose 2003). Once consumers become familiar with use, however, memories of these early costs will likely become less salient as perception become dominated by more recent, more pleasurable, usage experiences (e.g., Soman 2003). Hence, in the same way that it may be difficult to fully recall how difficult it was to learn to drive, consumers who try to attempt to anticipate the learning costs that will be associated with a new product generation may be prone to underestimate them.

Taken together, the converging evidence reviewed above leads to the following hypothesis:

> ***H1: The enhancement bias**. When given the opportunity to purchase a new product that possesses an expanded set of features relative to the incumbent, consumers will overvalue the new product, revealing rates of adoption and levels of willingness-to-pay in excess of those that would be justified by both actual and subsequent utilization patterns and a rational a priori options valuation.*

---

[4] The selective processing argument may also help resolve the apparent inconsistency between our argument and the well-known tendency for consumers to overweight out-of-pocket (purchase) costs relative to opportunity costs, since owning the new features is a form of opportunity (e.g., Thaler 1980). Our argument suggests that, though underweighting opportunity costs may be a general tendency, it may not hold true when other forms of cognitive biases are activated and point to a different direction. In our case, we argue that the strong affective component and the resultant selective processing may overcome the tendency to overweight out-of-pocket costs.

While **H1** predicts a global tendency to overvalue product enhancements, there will clearly be individual and contextual variations in the *degree* to which this will hold. To illustrate, Mukherjee and Hoyer (2001) offer data suggesting that if novel features are added to a product that is already perceived as complex and difficult to use, this enhancement bias may be muted by the greater salience of learning costs. Hence, the magnitude of the bias will likely differ depending on the quality of a consumer's experience with the adoption and use of the current generation of a product. Formally,

> ***H1a: The moderating effect of prior experiences****: the mean tendency of consumers to overvalue product enhancements will be moderated by the experienced utility and learning costs associated the incumbent product.*

## Now that they've bought it, will they use it?

After consumers purchase a new product normative theory prescribes that they should experiment with its new features to determine their true value as early as possible and make the appropriate switch-or-stick decision based on the outcome of their experimentation. However, there are reasons to suspect that this will not occur. After paying for new features consumers may be prone to prematurely abandon experimentation and subsequent utilization, limiting the potential utility they might otherwise draw from the acquisition.

One reason for this was suggested above: when making purchasing decisions consumers will tend to excessively focus on the imagined future benefits of an innovation's new features, and under-attend to learning costs that will later be required to realize them. When these costs are finally observed they may suppress usage for the rational reason that they are now accurately seen as being greater than the benefits that learning a new attribute could provide.

We suggest, however, that the actual degree of suppression will be even greater than that which would be predicted by such a rational re-assessment. The rationale is as follows. Unlike a one-time adoption decision, decisions about whether to utilize a new feature are frequent and recurrent; during the course of ownership one is making a continuous series of decisions whether to use familiar incumbent features with a known payoff versus a new features with both an uncertain payoff and, possibly, a trial cost. This dynamic nature of feature utilization brings to the table two new task elements that may degrade decision making: the need to see long-run benefits in incurring short-run costs, and the ability to defer experimentation.

A review of the substantial body of research on dynamic decision making strongly supports a conclusion that human decision makers are myopic, overweighing short-term consequences at the expense of long-term consequences (see, e.g. Hutchinson and Meyer 1994; Herenstein and Prelec 1991). Of particular relevance is work that has studied the ability of decision makers to solve "armed bandit" problems that involve a similar recurrent choice between an option that yields a reward with known probability versus another whose probability must be discovered through experimentation (e.g., Banks, Olson, and Porter 1997; Meyer and Shi 1995). A consistent finding of this work is that subjects abort experimentation too early when they encounter negative feedback, consistent with a tendency to view future consequences over too short a time horizon (e.g., Meyer and Shi 1995). Applied to the current task, this work would predict a similar inclination among consumers to prematurely abandon the learning of new product features when the benefits are not immediately realized, or if learning costs prove higher than expected.

This base aversion for experimentation may then be exacerbated by the omnipresent ability to defer the start of learning. Specifically, when a consumer acquires a good there is obviously no mandatory time at which they must start learning; a consumer may prefer to "ease in" to the new product generation by first using the more familiar features, then gradually learn to use the newer ones. The evidence from work that has examined how individuals make deferral decisions, however, suggests that the onset of learning may be delayed longer than normative theory would prescribe, and indeed may never begin at all (Dhar 1997; Tversky and Shafir 1992). Excessive deferral can be seen as a temporal extension of the well-known status-quo bias (Samuelson and Zeckhasuer 1988): given a choice between undertaking a risky, costly, action (learning a new feature) versus postponing the action for the moment and accepting a satisfactory status quo (using a familiar feature), there will be an inherent local preference for the latter. A myopic consumer might then be prone to repeatedly make this same decision, never realizing that doing so precludes them from realizing the higher long-run utility that motivated the purchase of the enhanced product generation in the first place.

Thus, the combination of myopia and propensity for deferral predicts that consumers will under-utilize the new features once they have them in their possession. We state the prediction as the following hypothesis:

> ***H2: The Under-Utilization Bias***. *Given a decision to acquire an innovation that possesses a mixture of innovative and familiar attributes, utilization of the new attributes will fall short of the levels implied by stated wiliness-to-pay for the good, direct forecasts of benefits, and objective benefits that would come from optimal usage.*

We might note that the processes that are theorized to lead to overvaluation of new-product generations and subsequent underutilization imply a possible paradox in how individual differences in post-purchase attribute utilization might relate to pre-purchase willingness to pay. In **H1a** we proposed that consumers who had more positive experiences when consuming past technologies would produce the most optimistic assessments of the prospective value of a new product that offered an attribute innovation. Yet, because much of this optimism comes from the under-forecasting of learning and switching costs (as above), it is these same consumers who would most likely experience the greatest disappointment when they come to utilize attribute innovation that they paid for. This disappointment, in turn, could lead to more rapid decisions to abandon the use of the new attributes relative to those who entered ownership with more modest expectations. We summarize this idea in the following hypothesis:

**H2a The Paradox of the Technological Optimist:** *Consumers who reveal the greatest optimism in their willingness to pay for a technological innovation will also be the most prone to abandon trial usage of attribute innovations given ownership.*

**Remark: why would usage predictions not be self-fulfilling?** Before we turn to the empirical evidence bearing on **H1** and **H2** a natural question might arise: why would consumer decisions about new attribute utilization not be subject to sunk-cost effects, hence at least be somewhat self-fulfilling? There is, of course, evidence that when consumers undertake monetary investments they often subsequently behave in a way that would serve to mentally justify that investment (e.g., Thaler 1980). In the current context one might hypothesize that adopters of new technologies will make a point of using its innovative features simply to justify the acquisition.

Yet, work that has examined the effect of investment commitments on behavior *over time* suggests that if adopters of technology are influenced by sunk-cost effects, the effect is likely to be short-lived. Of direct relevance are the works by DellaVigna and Malmendier (2002) and Gourville and Soman (1998) who have looked at over-time consumption behavior for subscription events such as health-club memberships. The consistent finding is that consumers systematically over-predict the degree to which they will utilize subscriptions as implied by the size of their initial investment; while usage is invariably initially high (perhaps fostered by sunk-cost effects), subsequent compliance often rapidly falls. What is particularly telling about these examples is that these are instances where consumers make heavy up-front investments in the hope that it will encourage compliance—yet we still see an over-forecast of usage. Once again, the most likely explanation is that as consumers we seem to have difficulty projecting how we will utilize products prior to actual ownership, with prior assessments biased upwards by the positive emotions that motivate the initial purchase (e.g., a sudden desire to become fit).

## Empirical Analysis

### Overview and Design Consideration

In this section we describe the results of three experiments designed to test the empirical validity of the research hypotheses summarized in **H1, H1a**, **H2**, and **H2a,** as well as provide descriptive insights into the process by which consumers make decisions to buy and then subsequently utilize product innovations. These issues were examined by observing how a sample of experimental subjects learned to play an original arcade-like computer game where performance was rewarded by a monetary incentive. After a period of training with one of several basic platform designs subjects were given the opportunity to purchase an enhanced platform that offered a combined set of features that were drawn from the basic platforms. In a third experiment we examine the reciprocal case: subjects trained on the enhanced platform are given the opportunity to exchange it with a reward for a simplified platform containing only one type of control.

The game was called "Catch'em" and bore similarities to the popular late 70's, early 80's arcade game *Pac Man.* In the game players viewed a square grid (Figure 1) on which was superimposed a number of stationary green dots called “cookies”. Also on the grid were two larger red and black dots that depicted the starting position of the player and his or her robotic opponent, termed the “Monster”. Upon triggering the start of the game both the Monster’s and player’s icons began moving over the grid. While the Monster moved at a random speed and direction, the player controlled the speed and direction of his or her icon. Each time either the player’s icon (or the Monster) moved over a cookie a point was scored for the player (or the Monster). If all of the cookies were consumed from the board by the player and/or the Monster, the play ended and the player received a point total equal to the number of cookies he or she had captured. If, however, at any point the Monster’s icon touched the player’s icon, the player’s icon was declared “caught” and play also ended, with all points having been earned from that play being forfeited.

We chose this—admittedly unusual—stimulus context because it was one that satisfied four ideal design criteria:

1. It provided us with experimental control over the design and familiarity subjects had with a basic generation of a technology;
2. It allowed experimental introduction over the value of enhanced features in a new technology;
3. It provided a natural objective for measuring performance that could be used for providing a monetary incentive to subjects; and
4. The task context—an arcade game—was one that was likely to be seen as highly involving and familiar to the subject pool, primarily undergraduate college students.

The technology in this case was the nature, complexity, and quality of the controls available to subjects for moving their icon. A basic technology was one where subjects had access to only one kind of control at one calibrated level of performance, while the enhanced technology was one where subjects had access to multiple controls—both those with which they were familiar and a "new" set that was derived from one of the other basic models (the existence of which was unknown to subjects).

Our analysis focuses on the results of three experiments conducted within this paradigm. The purpose of Experiment 1 was to conduct a basic test of the four hypotheses in a setting where there was minimal measurement intervention; we observed learning paths, the dynamics of control utilization, and adoption decisions in the absence of direct elicitations of either forecasts of behavior or elicitations of reasons for decisions—interventions that might influence behavior. In Experiment 2 we attempt to more deeply probe the process that underlies the data uncovered in Experiment 1 by gathering process measures. In Experiment 3 we examine the degree to which decisions to upgrade were drawn more by global preferences for expanded controls versus mistaken forecasts about usage by studying preferences for simplified platforms among those trained on multiple controls.

## Experiment 1

### Design, Subjects, and Procedure

Subjects were 138 business-school undergraduates who volunteered to complete the task for a monetary incentive. Subjects performed the experiment seated in computer cubicles in the school's behavioral research lab. At the outset of the experiment subjects were told that the purpose of the experiment was to learn how consumers such as themselves learned to play gaming devices, and that they would be paid depending on their performance in the game. Subjects were told that there would be a show-up fee of $5 (US) per subject, and they could earn up to $10 more depending on how well they learned to play the game.

All subjects were told that they would be playing the "Catch'em" game a total of 30 times, with the first 15 being practice rounds that would not count toward their final earnings, and the second 15 being money rounds on which their pay would be based. After reading this basic instruction subjects were randomly assigned to either a control or treatment condition, with which they were also assigned to play one of three different basic game platforms (described below). Subjects in the control condition played with the same platform over all 30 rounds of the experiment. Subjects in treatment condition played the first 15 training rounds with one platform, but were then given the opportunity to pay to play the money rounds with a new platform that offered a broader range of controls. The opportunity to pay to switch to a new platform was offered only once; if a subject declined the purchase he or she played the 15 money rounds with the same game platform that they trained on, the same as those in the control condition.

**The game platforms.** The three basic game platforms on which subjects trained were defined by the physical form and reliability of the controls used to move the player's icon. There were three mechanisms:

1. *A Scroll Bar Control (henceforth scroll bars version; Figure 1a)*: Subjects continuously adjusted the speed and direction of movement of their icon by moving each of two horizontal scroll bars displayed on the computer screen. Use of the directional control was aided by a steering-wheel-like graphic that displayed the current directional heading of the icon.
2. *A Button Control with high reliability (henceforth good buttons version; Figure 1b).* Subjects adjusted speed and direction by repeatedly clicking two sets of button controls. One pair of buttons allowed subjects to reverse the current heading of their icon either horizontally or vertically, while the other set induced discrete increases or

decreases in speed. High reliability meant that the icon's movement responded 80% of the times to player actions in the intended manner given activation of any control.

3. *A Button Control with low reliability (henceforth bad buttons version; Figure 1b).* The appearance and function of this platform was identical to (2), except that random noise was added to the responsiveness of controls. Specifically, given activation of a given control there was a 60% chance that it would momentarily fail, resulting in no change in movement of the icon.

In Figure 2 we plot the average performance attained by subjects using each of these control formats during the training rounds. The figure yields an important feature of this training manipulation: in addition to varying the tactile experience with controls that subjects had entering the money rounds, the three control conditions also manipulated the qualitative nature of their learning experience. Specifically, subjects found the button controls to be a more natural way of moving the icon than the scroll bars, and when the buttons were reliable they realized high levels of performance after a short period of familiarization. For subjects given the scroll-bar control, however, their learning experience was quite different: while they ultimately developed the same level of skill as those displayed by subjects who trained on the reliable buttons (as measured by average realized scores) this achievement was achieved only after they incurred more substantial learning costs as evidenced by the low average scores realized at the outset of training. Finally, subjects who trained on the low-reliability buttons would have found the training rounds to be a far more frustrating experience; while there was tactile ease in using the buttons, they would have experienced little improvement in achievement over time movement was inherently difficult to control.

**The enhanced platform.** The central interest in the experiment was how subjects in the treatment groups responded to the opportunity to play their money rounds of the game with a new platform that offered an expanded set of controls. This version--called the *combo platform*—provided subjects with access to *both* sets of controls that appeared in the basic platforms: buttons as well as scroll bars (Figure 1c). Note that since subjects trained on only one kind of control and were unaware of the existence of the other, the added controls that appeared on the combo version represented an innovation: the scroll bars would have been novel to those who trained on buttons, and the buttons novel to those who trained on scroll bars.

To insure that the locus of perceived benefits of the combo platform would be isolated to the new control, the function and reliability of the more familiar controls was identical to that which subjects had experienced during the training rounds. Hence, the reliability of the button controls in the combo platform was low for those who trained on low-reliability buttons and high for those who trained on high-reliability buttons. For subjects who trained on the scroll bar, the new button controls were of moderate reliability. In addition, the physical appearance of the combo platform was identical to that of each of the basic platforms with the exception of the presence of a second set of controls (Figure 1c).

It should be observed that the design implied that the *objective* incremental value of the new combo platform thus varied depending on the platform on which subjects trained initially. For subjects who trained on the low-reliability buttons the combo platform subjects access to a more reliable control (the scroll bars) that could potentially allow them to realize significantly higher scores in the money rounds. For subjects who trained on the scroll bars or the high-reliability buttons the objective advantage of the combo version was simply tactile flexibility; since both controls yielded comparable asymptotic levels of achievement (see Figure 2), higher mean achievement mean could be expected only if subjects differed in their natural aptitude for each of the two controls, and made optimal self-selection decisions upon ownership. Of course, subjects could only discover these comparative benefits if they chose to purchase the combo platform and then experimented with the performance of the new control.

**The pricing and purchase mechanism.** After completing the training phase of the game subjects in the control groups moved on to the money rounds of the game, while those in the treatment condition read a mock news announcement that a new version had been developed which they had the opportunity to purchase for play during the money rounds rather than the platform they trained on. Subjects were given an illustration of what the new game platform looked like. It was emphasized that the more familiar controls would function just like the old ones did, and no statement was made about whether the new control would yield better or worse game results than the old one; subjects were told that the new controls simply gave them greater flexibility in how they controlled their icon.

After reading this announcement subjects were then told that they could acquire the new platform by paying a point handicap that would be applied to their realized score in the money round. Before being shown what this price would be, however, they would have to indicate the maximum price that they would be willing to pay for the game, and they will obtain it if the actual price turns out to be less than this value—an elicitation procedure akin to that suggested by Becker, de Groot, and Marschak (1964; BDM). To insure that subjects fully understood how the process would work subjects first participated in a practice round where they set a *WTP* price for a movie ticket and an illustrative actual price was drawn by lottery. Subjects were given the opportunity to repeat this exercise until they felt comfortable with the procedure.

The actual price of the combo game was held constant for all subjects at 120 points, a price at which subjects would break even if the new game allowed them to realize a modest (8 point-per-game) increase in performance over the incumbent platform. This price thus implied that subjects who saw the prospect of either only nominal or no improvements in performance with combo platform would play the money rounds with their existing game, whereas subjects with more optimistic estimates would play with the combo game. After subjects submitted *WTP*s, those who submitted valuations greater than 120 were informed that they would be playing with the combo platform, and the purchase price was immediately reflected as a negative number in the cumulative score box on their game screen (see Figure 1c).

**Results**

Among the 68 subjects in the treatment condition who were given the opportunity to purchase the new game platform, 57 (84%) provided willingness-to-pay levels that were sufficient to attain ownership of the combo platform (valuations greater than 120). Hence, on the whole subjects were quite optimistic about the score improvement they could potentially realize by playing the version. A subsequent analysis of the performance of the 11 non-adopters during the money rounds revealed a pattern of achievement similar to that observed among those in the control condition, hence these two groups were pooled in subsequent analyses.

**The efficiency of adoption decisions.** Subjects' stated willingness to pay for the new platform is, of course, an implicit forecast of how having the ability to use a second control will improve their score beyond that which could be realized by the basic platform. Since the raw measure of *WTP*s is highly skewed, we utilize and report log-transformed *WTP*s in all subsequent analyses unless otherwise noticed. In Figure 3 we plot the mean *WTP* for each initial platform (Figure 3a) and the mean *WTP* of subjects who adopted the innovation by training condition relative to two standards of achievement (Figure 3b): the improvement in scores they actually realized relative to that realized over the last 6 games of the training round, and the improvement relative to the scores realized by control subjects who did not upgrade. The figure yields two insights that suggest initial support for **H1** and **H1a**:

1. *Excessive mean optimism in the projected benefits a new control*. The mean stated *WTP* for the new platform across training conditions was 341 game points (SD=232, median = 300), equivalent to an expectation that having access to a second control would allow them to realize a 20% (SD=14%) improvement in score over retaining the basic platform. These implicit forecasts, however, turned out to be quite poor: on average subjects who bought the new platform (henceforth "adopters") did not perform better than those who never upgraded. In addition, the mean *WTP* was not lower than the average increase in raw score by control subjects (mean *WTP* = 401 [median 375] versus the mean performance increase of 443 [median 385], p>.6). Moreover, *WTP* was negatively correlated with raw increase in total score (r= -.36, p<.001).
2. *The optimism bias was conditioned by the training platform*. By visual inspection, Figure 3a offers some initial support to **H1a**. That is, those starting with high-reliability buttons which offered the least frustrating experience also tended to give higher *WTP*s for the new platform than those starting with the more difficult scroll-bar platform (mean=388 for high-reliability buttons and 298 for scroll-bar, median=400 and 250, respectively).

To more rigorously explore the effect of training experience on *WTP*s, we modeled individual estimates as a function of the initial platform and subjects' experience during the training period (we used the maximum score over the last six games during training rounds [*MAX6*] as the proxy for experienced ease of learning). The regression results are presented in Table 1. It is clear that, in support of **H1a**, both factors contributed significantly to stated *WTP*s. Specifically, subjects who trained on either button platform stated significantly higher *WTP*s for the new platform than did those who trained on the scroll bar platform, presumably as a result of their better experience with the game in training rounds. In addition, *WTP*s were positively related to the experienced ease of learning ($p<.02$). Notice that our proxy for experienced ease of learning (*MAX6*) incorporates the recency bias in retrospective evaluations.

To more directly examine the degree to which subjects were able to anticipate their actual performance using the combo platform we modeled each player's cumulative score during the money rounds as a function of their average score in the training rounds, their *WTP*s for the combo game, and gender (see Table 2). The data yield a surprising result: after controlling for training performance, the marginal effect of increasing statements of *WTP* was *negative* ($t(1,63)=-3.33$; $p=.0014$) among those who purchased the new platform[5]. In short, at the margin those with the most optimistic estimate of how well they would do in the money rounds tended to have the lowest actual achievements.

Additional insight into why subjects who acquired the new platform may have underperformed relative to their *WTP*s is contained in Figure 4, which plots performance over all 30 trials for treatment versus control subjects by training condition. The figure suggests one contributing explanation for the exaggerated *WTP* estimates: while subjects who bought the new platform seem to have correctly anticipated that their performance would improve on the money trials playing with the new platform, they failed to foresee two factors that would also naturally mitigate achievable relative performance:

---

[5] Excluding those who chose not to upgrade did not change the conclusion.

1. The fact that there would also be improvements in skill levels playing with the basic platform; and
2. Any potential incremental benefits of the combo version would not be immediately realized as control usage would likely alternate, at least initially, between the two options.

In short, it is as if the *WTP* estimates reflected a comparison of an envisioned *asymptotic* value of the combo platform to the *current* value of the basic platform—a comparison that naively overlooks the dynamics that would govern actual relative performance during the money period but is consistent with the projection and impact biases we discusses earlier.

**Feature utilization.** It should be emphasized, of course, that the conclusion of overstated *WTP*s exploits the hindsight knowledge that was not in evidence at the time subjects made these assessments: the objective incremental value of the added control option. Recalling the principles of rational product adoption we discussed at the start, the apparent overvaluation of the combo device might simply be seen as the case of rational investments in an experiment that did not pay off. That is, subjects who purchased the innovation ended up achieving levels of performance similar to those who did not simply because they discovered, after experimentation, that there was no added value.

In **H2** and **H2a**, however, we hypothesize that while subjects may well acquire the combo platform with well-meaning intentions to learn about its value, its new features will be underutilized, even in settings where there would be a real normative gain. In the current experiment such is the case of subjects who trained on the low-reliability button control. For these subjects the new availability of the scroll bar offered a very real opportunity to increase earnings, though it would require them to incur a period of learning with a control that they are likely initially to find unnatural.

In Figure 5 we plot the relative frequency with which subjects in each of the training conditions used their new control in combo platform over trials. For subjects who were trained on either reliable buttons or scrolls bars (the dashed lines in Figure 5) the data give strong apparent support for **H2**: although subjects paid a substantial amount—and were prepared to pay more—for the ability to at least experiment with the use of the new control, few made use of this opportunity. Specifically, during the initial three games (block 6 in Figure 5) of the money period, when utilization of the novel control should rationally have been quite high, subjects who had trained on the high-reliability buttons and the scroll bar utilized the new (reciprocal) control on average only 21% of the time, a level that diminished over time thereafter. In addition—and perhaps shockingly—the data revealed 8 subjects in these two conditions who *never* utilized the new controls at all over the entire 15 games, yet their *WTP* were no less than that of others (mean=396 and median=342, p>.9).

Perhaps the most compelling evidence favoring **H2** is the level of scroll-bar usage displayed by subjects who had trained on the low-reliability button (the solid line in Figure 5). On one hand, unlike those who had positive experiences in the training rounds, here we see subjects display a much higher rate of initial usage of the scroll bar, though its level (54%) is still below that which one would normatively prescribe if subjects were active experimenters. In addition, contrary to the normative recommendation, 3 subjects did not start experimenting with the new controls at least until after the first 3 games. On the other hand, the more disturbing feature of the data is that utilization never increased much in the task beyond this level—even though subject would have clearly benefited if it had. In essence, subjects seemed unable to abandon use of a familiar control in favor of a new one, despite the objective superiority of the latter.

As a final analysis we examined how individual differences in novel attribute utilization related to subjects' willingness-to-pay for the combo platform. This relationship is, of course, normatively positive; since *WTP* should reflect, in part, the value a subject sees in experimenting with the new control, the higher the *WTP*, the more a subject should invest in its usage, at least until its true value is established. **H2b**, however, predicts the opposite: because high *WTP* measures are theorized to be induced not by rational assessments of the value of information but rather by projected expectations of high immediate returns from the innovation (**H1a**), the more upwardly-biased this assessment, the more likely subjects will be to terminate usage after limited trials.

To test this hypothesis we estimated two models explaining the proportion of uses of the novel control for each subject over games: one that modeled usage as a function of stated *WTP* for the combo platform, indicator variables for a subject's training platform, and game trial (both linear and quadratic functions; Model 1 in Table 3), and another that that modeled usage as a function of game trial (linear and quadratic terms) and experienced ease of learning (MAX6; Model 2 in Table 3). The results of these analyses, reported in Table 3, support **H2b**: rather than serving to foster new attribute usage, in this case optimistic beliefs serve to suppress it. The greater the degree to which a subject had positive experiences in the training period (MAX6) and/or emerged from it with superior expectations about the value of the new attribute (*WTP*), the lower the mean utilization. Similar pattern of results still hold if we examine the usage behaviors in only the first game or first three games.

## Experiment 2

## Motivation and Description

While Experiment 1 offers apparent support for the hypothesis of over-estimation of the value of product innovations, it leaves uncertain the degree to which the revealed levels of *WTP* accrued directly to over-forecasts of new control attribute utilization and value versus other hedonic drivers, such as a simple attraction to novelty. To resolve this issue a second group of 27 subjects were recruited to replicate a version of the study reported in Experiment 1, but with one major design change. After subjects read the description of the new combo platform and before the elicitation of their willingness-to-pay they were posed with a series of questions designed to tap their beliefs of the value of the new platform. Subjects were asked to make four forecasts:

1. The likely percentage change over all 15 games if they continue to use their existing platform (positive or negative);
2. The likely percentage change over all 15 games if they switch to the new platform (positive or negative);
3. The likely percentage change in their score over the first three games compared to that which would be realized using the basic platform (positive or negative); and
4. The percentage of time during the first three games of the money rounds that they would likely utilize the new control offered by the new platform.

All forecasts were provided by checking a box on a discrete category scale that offered a range of possible percentage responses.

In addition, after *WTP* measures were elicited and the mock lottery was run, subjects who received the new combo platform were asked an additional series of questions designed to elicit their reasons for setting their *WTP* levels as they did. These took the form of a series of bipolarscaled question asking for the degree to which they agreed or disagreed with a set of possible motivations for wanting to acquire the combo platform, including expected performance increases, flexibility, aesthetics, and a desire for a change of pace.

Because of the limited size of the subject pool only one (rather than three) training-platform conditions was replicated: the high-reliability button—the condition where the optimism bias was most acute in Experiment 1.

**Results**

Mirroring the results of Experiment 1, subjects revealed a high willingness-to-pay (mean raw *WTP*s = 366, SD=248, median=385) for the new game platform, with 100% of subjects adopting the new platform given the mock-lottery price. The *WTP*s were not different from those in Experiment 1 (p>.6). In Table 4 we report a comparison of how their actual performance and new-control utilization compared to the forecasts. The data suggest that the over-forecasts of performance and utilization provide at best a partial explanation for their high levels of willingness-to-pay. Specifically,

1. Subjects did not show excessively optimistic forecasts of their performance with the new platform; in fact, the overall trend is that of an *under*-forecast relative to the benchmark of last six games during training rounds, especially for cumulative performance over the entire money rounds. On average, the forecast percentage improvement in performance with the existing platform roughly corresponded to the actual percentage change by control groups in Experiment 1 (p>.3). Their mean actual performance in the money rounds was 2241 (SD=713), comparable to that of the control group in Experiment 1 (mean = 2336, SD=862); but

2. Subjects severely over-forecasted the degree to which they would be utilizing the new control; while the average forecast rate during the first three games was 50%, actual usage was closer to 20% in the first three games. On average they used the new controls only about 10% over all 15 games.

Since subjects provided performance forecasts for both the incumbent and new platforms, these estimates can be used to derive normative *implied WTP*s: estimates of what these assessments *should* have been had they been based solely on their comparative performance forecasts; that is, the forecast performance for the new platform minus that for the incumbent. In Figure 6 we present a scatter plot of implied versus actual *WTP*s. The figure yields a striking result: the absence of a systematic positive relationship between the two constructs. Indeed, the relationship between the two measures was nominally negative, with nearly 80% of subjects (21 out of 27) stating a higher or equal score forecast for the incumbent platform compared to the new, resulting in an aggregate mean implied *WTP* of -154 (SD=360, median = -190). In contrast, recall that their average stated *WTP*s revealed only a moment later was significantly higher (Mean=366, SD=248, median=385; paired test t(26)=5.50, p<.0001). Hence, at the very least, the data reject the idea that subjects formed assessments of *WTP* by contrasting forecasts of the value of the new platform with what this *would* have been had they stayed with the incumbent.

What were subjects' own reasons for being attracted to the new game platform? In Table 5 we summarize the degree to which subjects, on average, agreed or disagreed with each of four possible reasons for preferring to acquire the combo platform. The data suggest that if hedonic or boredom-related factors were influencing this decision, subjects were not inclined to admit to them. Subjects were most inclined to agree that the decisions were motivated by the two factors that high *WTP* valuations *should* have been based on: expectations of higher performance and a desire for flexibility. In contrast, subjects were less inclined to agree that attraction was motivated by the aesthetic appeal of the new platform and nominally disagreed that it was motivated by task boredom.

One interpretation of the data is thus that approached their assessments of *WTP* with normatively-correct beliefs about the factors that should drive these evaluations. Where the *WTP* assessments went awry was a tendency to overestimate the degree to which they would utilize the new control feature (the value of flexibility) and an underestimate the degree to which their performance would likely improve to a similar degree using the older platform. What is curious about this latter result is that subjects *did* reveal such knowledge when directly asked what their performance would be using the older platform; it was simply not incorporated when *WTP* judgments were made.

## Experiment 3

### Motivation

While the first two experiments offer apparent evidence that respondents placed too high a value on the opportunity to adopt a platform that offered an enriched set of controls, they do not conclusively resolve the degree to which the enhancement bias was due to biases in forecasts or more subjective appeal (such as design sophistication and/or usage flexibility). For example, one might argue that we would not have seen this bias had subjects been given an ample opportunity to try out the enhanced platform prior to purchase, when their limited utilization of all controls would have been transparent. To help resolve this issue we undertook an investigation of choice behavior in the reciprocal case of Experiments 1 and 2: we examined the degree to which subjects who were initially trained on the combo machine would accept point rewards for trading *down* to a simpler platform that utilized only one set of controls.

If players found little use in one of the controls during the training rounds, the opportunity to be compensated for playing a simpler platform *should* be quite attractive; they would be receiving a bonus to use a platform whose essential functionality was unchanged. Hence, if the interest in upgrading we observed in the first experiments accrued primarily to exaggerated beliefs about the benefits of the new control (from an inability to fully try them out), we would expect to see high levels of interest in the exchange when subjects had direct knowledge of benefits. In contrast, if the upgrades were motivated more by less rational sources of appeal of the combo platform—such as a preference for design sophistication--we should see a reluctance to enter into such an exchange. This same reluctance might also be exacerbated by endowment effects (Kahneman, Knetsch, and Thaler 1990); even if the added controls offered by the combo platform are found to be of only mildly positive value, decision makers may be reluctant to part with them by the mere fact that they are currently possessed.

**Design, Subjects, and Procedure**

205 undergraduate business majors participated for partial class credit and for a monetary incentive. The basic structure of the experiment mirrored that described in experiments 1 and 2 but in reverse; here all subjects played the training rounds with the combo platform, and were then given the opportunity to play the money rounds with a simplified platform that utilized only one set of controls. In addition, rather than eliciting *WTP* measures for the new platforms, in this case willingness-to-accept (*WTA*) measures were elicited: after the training rounds subjects were told that they would be given a point *bonus* if they were willing to play the money rounds with a simplified device. Subjects were randomly assigned to one of two trade-down conditions: one

where the simplified platform used only scroll-bar controls, and one offering only button controls. In addition, to provide a second measure of control use, after completing the training rounds subjects were asked to indicate which of the two sets of controls they found most useful in play—a measure that turned out to correlate quite highly with actual usage.

The *WTA* measures were elicited by a modified *BDM* procedure mirroring that used in Experiments 1 and 2. In this case the actual acceptance threshold was set at 300 points, slightly more than twice the threshold value used in the *WTP* lotteries in Experiments 1 and 2. This threshold was such as to compensate the average player for a 20% decrease in performance in the money rounds playing with a simplified platform---a highly generous level of compensation.

## Results

Is there a market for goods that remove unused features from familiar devices? In the current context the answer is a strong “no”. First, consistent with previous research on both endowment effects and preferences for flexibility (e.g., Shin and Ariely 2002), on average subjects were averse to entering into exchanges that would require them to give up one of the combo game controls. Subjects indicated a mean *WTA* of 513 (median 500), compared to the mean of 341 that they were willing to pay to acquire them in the first experiment—a result that produced 26% successful trade-downs across both platform conditions.

The most interesting aspect of the data, however, centers on those subjects for whom the new platform contained just the controls that they indicated that they found most useful in the first experiment. While the opportunity of an exchange should have been quite attractive to such subjects, most set *WTA* prices that precluded successful transactions, revealing a mean WTA of 455 (median 500), leading to 31% successful trade-downs. Even more dramatically, even the 38 subjects who *never* used the control that was being eliminated displayed an aversion for playing with a simpler platform, revealing a mean *WTA* of 409 (median 455), an assessment that led to 39% successful transactions.

The data thus reject the idea that the strong interest that subjects in the first two experiments displayed in acquiring the combo platform accrued entirely to mistaken forecasts of the net value of having access to a second control, or a desire for a change-of-pace in the game interface. Rather, much of the appeal likely rested in a root attraction to the more sophisticated design and choice of controls that the new platform offered—points of appeal that had little objective value, but for which subjects were reluctant to give up once they were acquired.

## Discussion

One often hears it said that consumers frequently buy more technology than they can realistically make use of. We compare products by counting the number of features they offer (often without knowing what they are used for), and feel bad when we discover that our incumbent devices are no longer state-of-the-art, regardless of how adequately they may be fulfilling our needs And this apparent bias is by no means limited to consumers; firms as well have recently been criticized as well for being prone to invest in more material and programmatic innovations that actually get implemented (Kim, Pae, Han, and Srivastava 2002). On the other hand, it also seems to be the case that firms who sell technology also worry about the possibility that consumers may walk away from technologies that are seen as *too* innovative. Hence, for example, when releasing Windows *XP* took pains to insure that its innovation would be seen by consumers as only modestly different from its old operating systems, to the point of allowing users the option eliminate new screen views if they wanted (through the "revert to classic view" command).

Yet, as pervasive as these observations may be, there has been little prior work that has formally studied the biases that characterize consumer new-technology adoption and subsequent usage decisions, and the psychology that may underlie these biases. The goal of this paper was to take a step toward gaining this knowledge by observing how a individuals made decisions whether or not to buy a new technology—an improved gaming device—in a laboratory setting where we could measure both the actual and perceived value of the technology as well as how it was utilized after purchase, and manipulate the kind of experiences with prior similar technologies that subjects had coming into the buying decision.

Central to the work was the idea that a general over-buying bias may, in fact, have a systematic cognitive basis. Drawing on prior work in affective forecasting, we hypothesized that when buying new technologies consumers will usually have a difficult time anticipating how they will utilize a product after it is purchased, and will be prone to believe that the benefits of attribute innovations that are perceived now will project in a simple fashion into the future. Implicit to this over-forecast is a tendency to underestimate the impact of factors that may likely serve to diminish usage in the future that are not in evidence now, such as frustration during learning and satiation. Consequently, there is a tendency for consumers to systematically evaluate product innovations through rose-colored glasses, imagining that they will have a larger and more positive impact on the future lives than they most often will likely end up having.

The experimental data reported here provide strong apparent support for this view of new product valuation. What is notable about the current demonstration is that the evidence for the optimism bias we report was derived from a context designed to facilitate rational assessments of innovation value. Specifically, subjects were given a clearly-stated metric by which the objective value of the innovation would be assessed, there was a direct monetary penalty for overstating value (the game innovation was paid for by a point deduction), and the innovation itself was a purely functional rather than aesthetic one (a new control added to the same graphic game platform). Yet, subjects still succumbed to the same biases that we suggest may be pervasive in real markets: a tendency to overvalue prospective innovations, and then under-utilize their features upon acquisition, even for the limited purposes of experimentation (e.g., in the current task 14 out of 84 subjects across the first two experiments who purchased the innovation never used its added control at all).

A further intriguing aspect of the results is that while these upgrade decisions proved to be ill-advised in terms of the benefits respondents actually realized from them, few seemed to regret it; there was no reciprocal rebate market for simplified platforms that allowed players access to just those controls they found most useful. Hence, respondents were apparently drawn to the enhanced platform not simply because of an (erroneous) calculation of expected benefits minus costs, but rather by a battery of affective forces that have a more limited rational basis, such as a desire to own top-of-the-line, and a pure preference for flexibility in control usage--even when never exploited.

An interesting question is why people do not anticipate the effects of cognitive lock-in at the time of making forecasts and adoption decisions. The core explanation we offered was cognitive availability: consumers who are wrapped in the optimistic euphoria of an adoption will simply find negative associations less available in memory, hence fail to fully incorporate them in forecasts (the focalism bias). The work of Liberman and Trope (1998), however, suggest a somewhat different way in which this availability bias may be working. Specifically, recall that adoption decisions primarily involve consideration of the long-run benefits of acquiring new features rather than the experiences on specific short-run usage occasions—a distant focus that often encourages a desirability bias (Liberman and Trope 1998). In contrast, the utilization decision is primarily concerned with the short-run decision of whether ones get greater pleasure from using familiar or unfamiliar attributes—a focus that may trigger a more myopic mindset. These different temporal frames, once activated, may lead to selective processing of information, filtering out the information that is inconsistent with the current frame.

Likewise, the observed reluctance to trade down to more efficient platforms might be driven by factors other than an inherent aversion for abandoning flexibility (e.g., Shin and Ariely 2003). A complementary mechanism may be that found by Dhar and Wertenbroch (2000), who studied how consumers make decisions to acquire goods that require them to give up on hedonic features but gain utilitarian ones. They show that when making such decisions consumers are more reluctant to give up attributes over utilitarian attributes because the former triggers greater degrees of cognitive elaboration about the consequences of loss. In the current context, owners of the enhanced platform were reluctant to trade down to a simpler device simply because it was easier to think of the possible downsides of loosing flexibility (a hedonic attribute) than the upsides of gaining scoring efficiency (a utilitarian attribute).

### Caveats

While the current findings offer support for the hypothesized effects of product enhancements, care must obviously be taken before presuming that the findings will hold in all product-adoption settings. First, a quite natural question—one that the current work only begins to address—is the effect of long-term learning on the enhancement bias. For example, it is natural to argue that once a consumer recognizes that they have overbought a technology they will be less inclined to do the same the next time around; i.e., they would become more astute forecasters of how they really make use of new technologies. The current data, however, offer a mixed view of the likelihood that such learning effects would be widely observed in natural settings. On one hand, consistent with this prediction, subjects who expressed the most conservative willingness-to-pay for the game innovation were those who experienced the highest learning costs with the game on which they first trained. On the other hand, Experiment 3 suggests that such effects may be offset by the simple fact that consumers may not see decisions to buy top-of-the-line goods that are under-utilized as “mistakes”; while there will always be a market for upgrades, establishing a market for downgrades (for current users of enhanced products) may be more difficult.

Another important question that might be raised is that the findings of this research seem, at first blush, to cut against the grain of those who have offered evidence of an *under*valuing of new technology options due to lock-in effects, such as those reported by Johnson, Bell, and Lhose (2003; in web-site visitation patterns) and Zauberman (2003; in valuations of new search engines). The current work differs from these efforts, however, in that here we consider the case of customer valuations of new products that offer *separable enhancements* to an existing platform; that is, by buying the new product one does not have to abandon what one has already learned (i.e., start a new learning curve). Congruent with their findings, we observed that subjects who developed a strong familiarity with one type of motion control indeed tended to stick with it even when they paid for the option to use an alternate, and even when the new control offered an objective normative benefit. The bias we observed is that subjects seemed unable to anticipate this underutilization when forming initial valuations.

In this same vein, an interesting challenge for future work would be to better identify situations where consumers systematically *undervalue* attribute innovations. Recent work by Mukherjee and Hoyer (2001) suggest such a boundary condition. They offer data showing that when an underlying product already has a complex structure adding more novel features can degrade attractiveness by increasing perceived learning and usage costs. Hence, there is almost certainly an upper limit to the enhancement effect documented here: added product features likely increase attractiveness only to a point. Likewise, the current data also show that the appeal of new attributes will also be conditioned by past experienced learning difficulty (Hypothesis H1a; Experiment 1).

Finally, an important goal of future work would be to better resolve the psychological mechanisms that underlie consumer assessments of product novelty. In this work we show that these assessments are biased in a way that is consistent with biases found in other domains of hedonic judgment, and for some of the same apparent reasons—for example, failing to anticipate future reluctance to experiment with the new control. But account is clearly a blunt one; the actual mechanism by which consumers develop visions of the future through analogical and other forms of structured reasoning (e.g., Moreau, Lehmann, and Markham 2001) is clearly a complex one, and more thoroughly understanding it may help better resolve the empirical boundaries of assessment biases and, possible, their correction.

# References

Banks, Jeffrey, David Porter, and Mark Olson (1997), “An Experimental Analysis of the Bandit Problem”, *Economic Theory*, 10 (1), 55-77.

Becker, Gordon. M., Morris H. DeGroot, and Jacob Marschak (1964), “Measuring Utility by a Single-Response Sequential Method,” *Behavioral Science*, 9 (July), 226-232.

DellaVigna, Stefano, and Ulrike Malmendier (2002), “Self Control in the Market: Evidence from the Health Club Industry,” *Working Paper*, Department of Economics, University of California, Berkeley.

Dhar, Ravi (1997), “Consumer Preference for a No-Choice Option,” *Journal of Consumer Research*, 24 (September), 215-231.

Dhar, Ravi, and Klaus Wertenbroch (2000), “Consumer Choice Between Hedonic and Utilitarian Goods”, *Journal of Marketing Research*, 37 (February), 60-71.

Gatigon, Hubert and Thomas S. Robertson (1991), “Innovative Decision Processes,” in eds. Thomas S. Robertson and Harold H. Kassarjian, *Handbook of Consumer Behavior*, Englewood Cliffs, NJ: Prentice-Hall, 316-348.

Gilbert, Daniel, and J. B. Ebert (2002), “Decisions and Revisions: The Affective Forecasting of Changeable Outcomes,” *Journal of Personality and Social Psychology,* 82, 502-514.

Gilbert, Daniel, Elizabeth C. Pinel, Timothy D. Wilson, Stephen J. Blumberg, and Thalia P. Wheatley (1998), “Immune Neglect: A Source of Durability Bias in Affective Forecasting,” *Journal of Personality and Social Psychology,* 75(3), 617-638.

Gourville, John, and Dilip Soman (1998), “Payment Depreciation: The Behavioral Effect of Temporally Separating Payments,” *Journal of Consumer Research,* 25(September), 160-174.

Grice, H. Paul (1975), “Logic and Conversation,” in eds. J. Cole and P. Morgan, *Syntax and Semantics*, *Vol. 3, Speech Acts*, New York: Academic Press, 43-58.

Herenstein, Roichard J., and Drazen Prelec (1991), “Melioration: A Theory of Distributed Choice,” *Journal of Economic Perspectives*, 5 (Summer), 137-156

Hogarth, Robin M. (1981). "Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgment Heuristics." *Psychological Bulletin,* 90, 197-217.

Hutchinson, J. Wesley, and Robert J. Meyer (1994), “Dynamic Decision Making: Optimal Policies and Actual Behavior in Sequential Choice Problems,” *Marketing Letters*, 5(4), 369-382.

Johnson, Eric J., Steven Bell, and Gerald L. Lhose(2003), "Cognitive Lock-in and the Power Law of Practice," *Journal of Marketing*, 67 (April), 62-

Kahneman, Daniel (1999), "Objective Happiness," in eds. Daniel Kahneman, Ed Diener, and Norbert Schwarz, *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation, 3-25.

Kahneman, Daniel, Jack Knetsch, and Richard Thaler (1990), "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy,* 98, 1325-1348.

Kim, Namwoon, Jae H. Pae, Jin K. Han, and Rajendra K. Srivastava (2002), "Institutionalization of Technologies for Technology Markets: A Buyer-Seller Approach," *Working Paper*, Hong Kong Polytechnic University.

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

Liberman, Nira, and Yaacov Trope (1998), "The Role of Feasibility and Desirability Considerations in Near and Distant Future Decisions: A Test of Temporal Construal Theory," *Journal of Personality and Social Psychology*, 75 (1), 5-18.

Loewenstein, George, and Daniel Adler (1996), "A Bias in the Prediction of Tastes," *Economic Journal*, 105, 929-937

Loewenstein, George, and David Schkade (1999), "Wouldn't it be Nice? Predicting Future Feelings," in eds. Daniel Kahneman, Edward Diener, and Norbert Schwartz, *Well Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation, 85-105.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin (2003), "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics,* 118.

Meyer, Robert J., and J. Wesley Hutchinson (2001), "Bubbling Geniuses: The Power of Everyday Reasoning in Mulitstage Decision Making," in eds. S. Hoch and H. Kunreuther , *Wharton on Making Decisions*, New York: John Wiley, 37-62.

Meyer, Robert J., and Yong Shi (1995), "Sequential Choice under Ambiguity: Intuitive Solution to the Armed Bandit Problem," *Management Science*, 41 (May), 817-834.

Moreau, C. Page, Donald R. Lehmann, and Arthur B. Markman (2001), "Entrenched Knowledge Structures and Consumer Response to Really New Products." *Journal of Marketing Research,* 38 (February), 14-29.

Miller, Elizabeth, and Barbara E. Kahn (2003), "The Effect of Color and Flavor Names on Consumer Choices," *Working Paper*, Department of Marketing, the Wharton School of Business, University of Pennsylvania.

Mukherjee, Ashesh, and Wayne D. Hoyer (2001), “The Effect of Novel Attributes on Product Evaluation,” *Journal of Consumer Research*, 28 (December), 462-472.

Samuelson, William, and Richard Zeckhauser (1988), “Status-Quo Bias in Decision Making”, Journal of Risky and Uncertainty, 1(March), 7-59.

Shin, Jiwoong, and Dan Ariely (2003), “Keeping Doors Open: The Effect of Unavailability on Incentives to Keep Options Viable,” *Management Science,* in press.

Simonson, Itamar (1990), “The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior,” *Journal of Marketing Research,* 32, 150-162.

Soman, Dilip (2003), “Prospective and Retrospective Evaluations of Experiences: How you Evaluate and Experience Depends on When you Evaluate it,” *Journal of Behavioral Decision Making*, 16(1), 35-52.

Thaler, Richard (1980), “Toward a Positive Theory of Consumer Choice,” *Journal of Economic Behavior and Organization,* 1(1), 36-90.

Tversky, Amos, and Eldar Shafir (1992) “Choice under conflict: The Dynamics of Deferred Decision,” *Psychological Science*, 3 (6), 358-361.

Wilson, Timothy D. and Daniel Gilbert (2003), “Active Forecasting,” in ed. Mark Zanna, *Advances in Experimental Social Psychology,* 35, New York: Elsevier, 345-411.

Zauberman, Gal (forthcoming), “Lock-In from an Inter-temporal Choice Perspective: The Combined Effect of the Information Cost Structure and Time Preferences,” *Journal of Consumer Research*, 30 (December).

## Figure 1: The Three Game Platforms

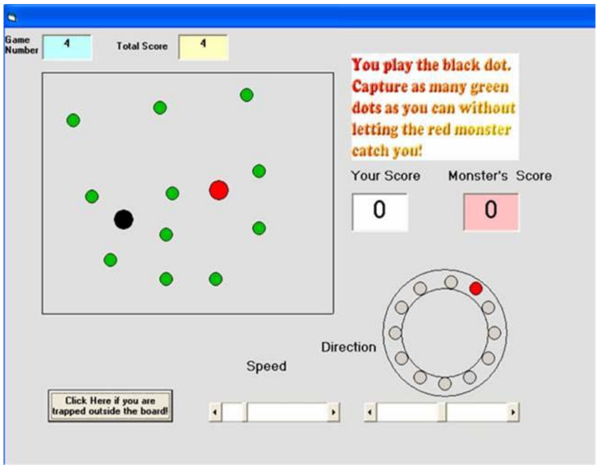

Game Number 4

Total Score 4

You play the black dot. Capture as many green dots as you can without letting the red monster catch you!

Your Score 0

Monster's Score 0

1a: Scroll-Bar Control

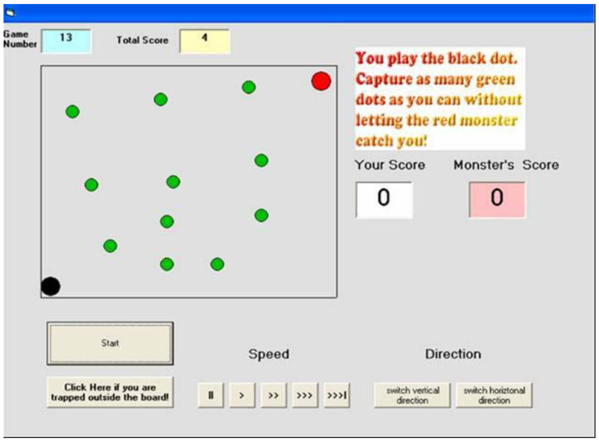

Game Number 13

Total Score 4

You play the black dot. Capture as many green dots as you can without letting the red monster catch you!

Your Score 0

Monster's Score 0

1b: Button Control

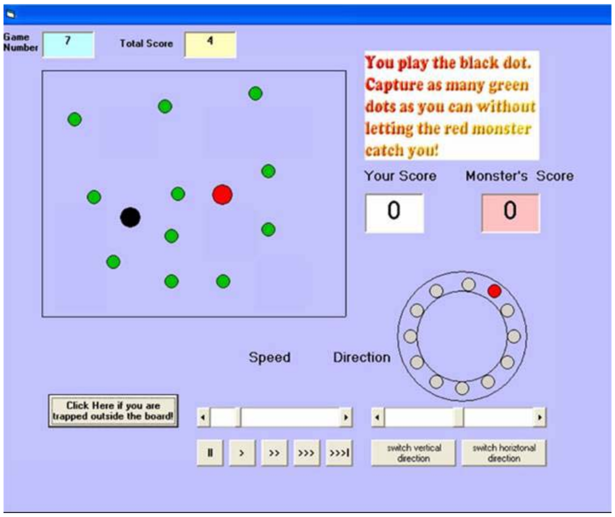

Game Number 7

Total Score 4

You play the black dot. Capture as many green dots as you can without letting the red monster catch you!

Your Score 0

Monster's Score 0

1c: The Enhanced Platform: Combined Controls

**Figure 2: Performance over time during training rounds by initial platform (Experiment 1)**

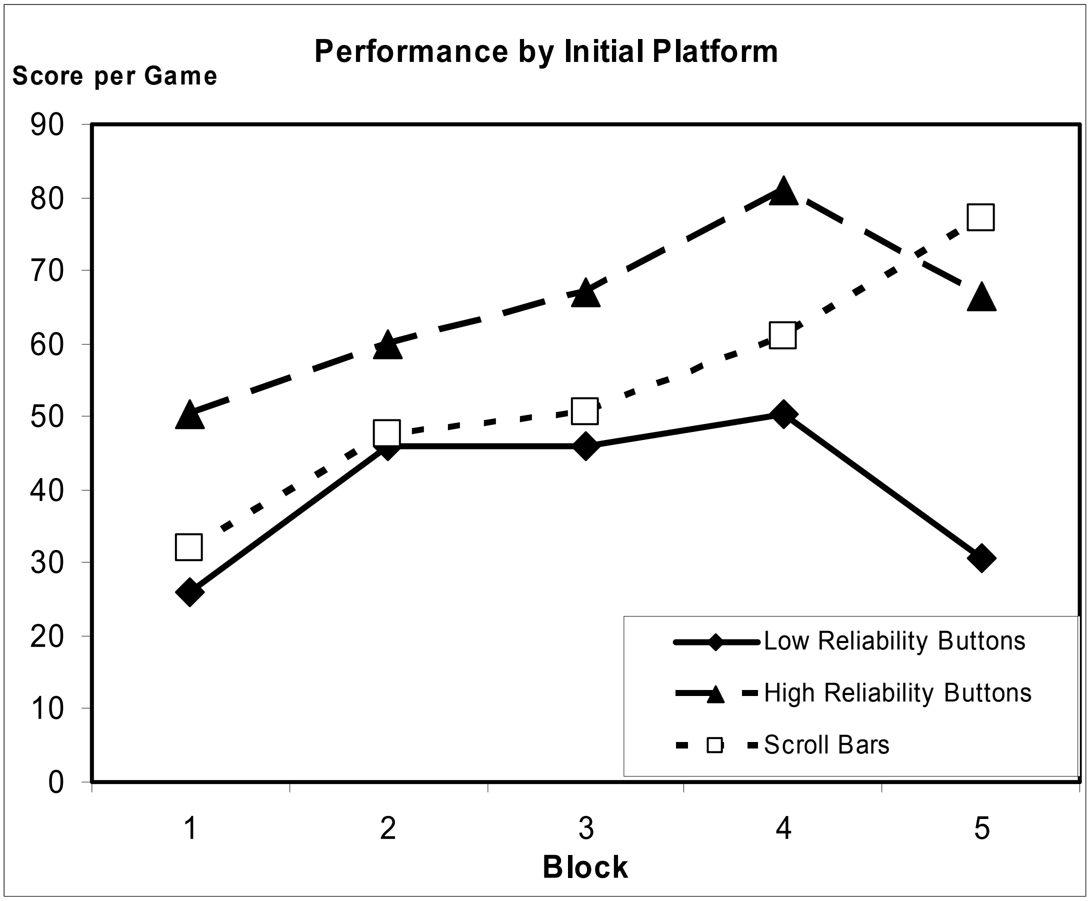

Note:
Game trials are reported as blocks of three games.

**Figure 3: Stated willingness-to-pay for the new combo platform by training platform (Experiment 1)**

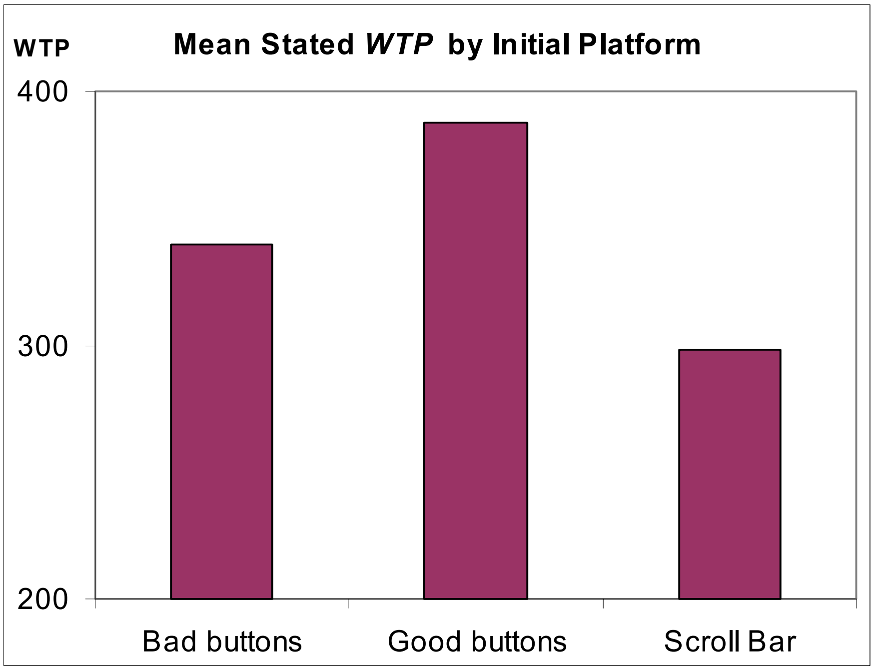

Figure 3a: Mean stated willingness-to-pay by type of training platform for all participants in the treatment condition

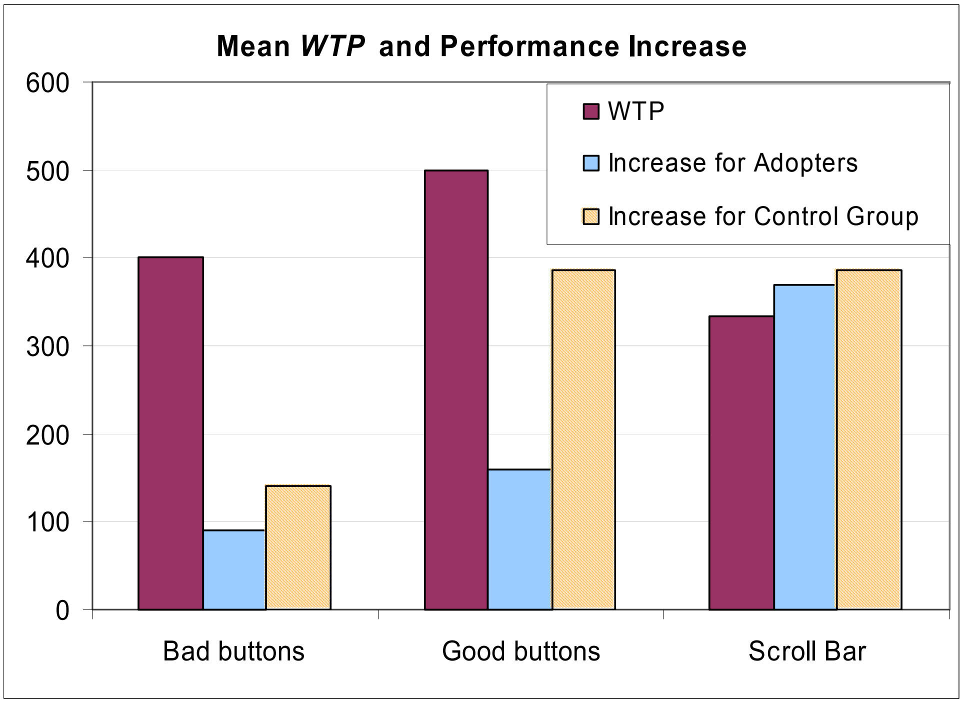

Figure 3b: Mean stated willingness-to-pay relative to performance increases by adopters and control group

**Figure 4: Performance over time by initial platform and upgrade decision**

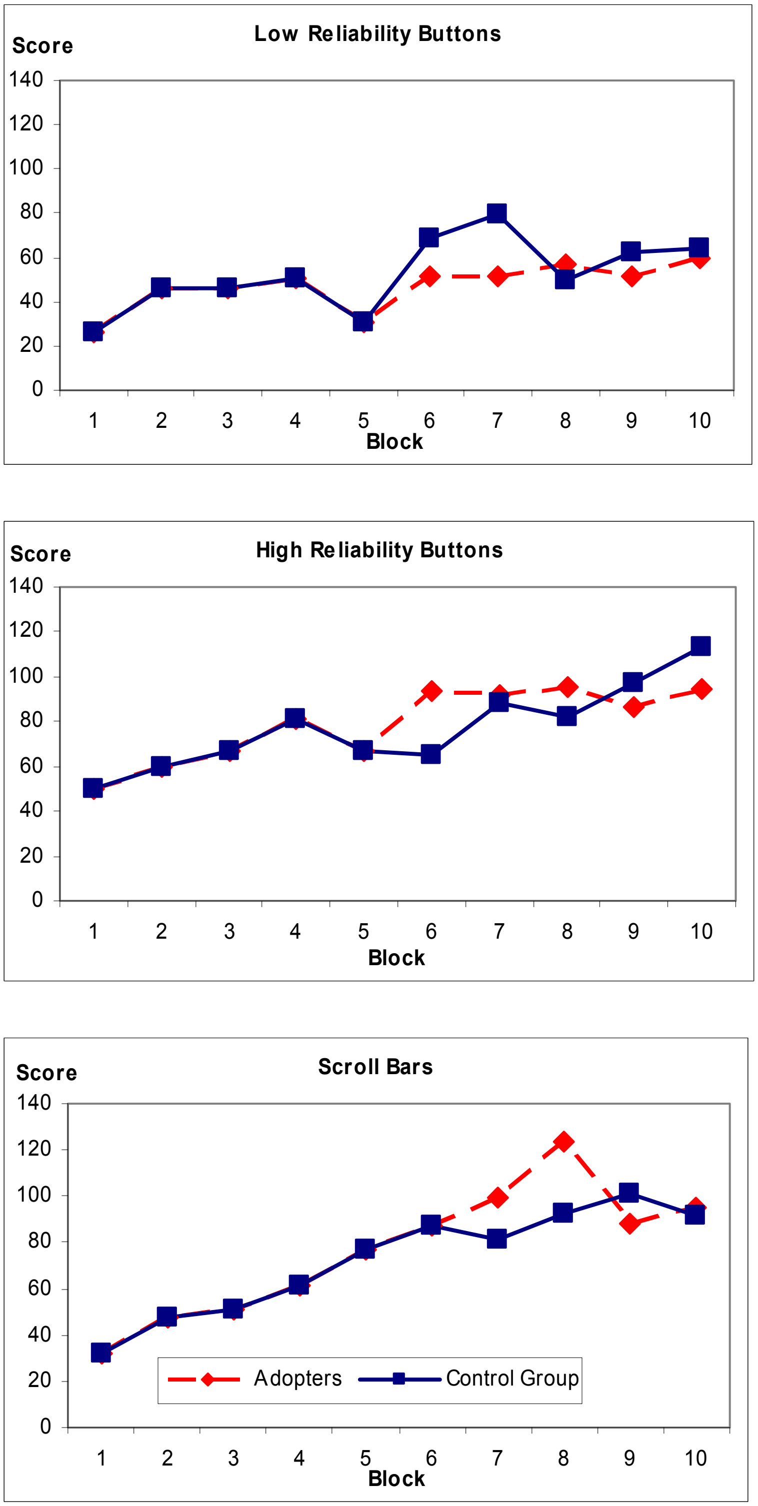

Note:
Game trials are reported as blocks of three games.

**Figure 5: Utilization of new features over time by initial platform**

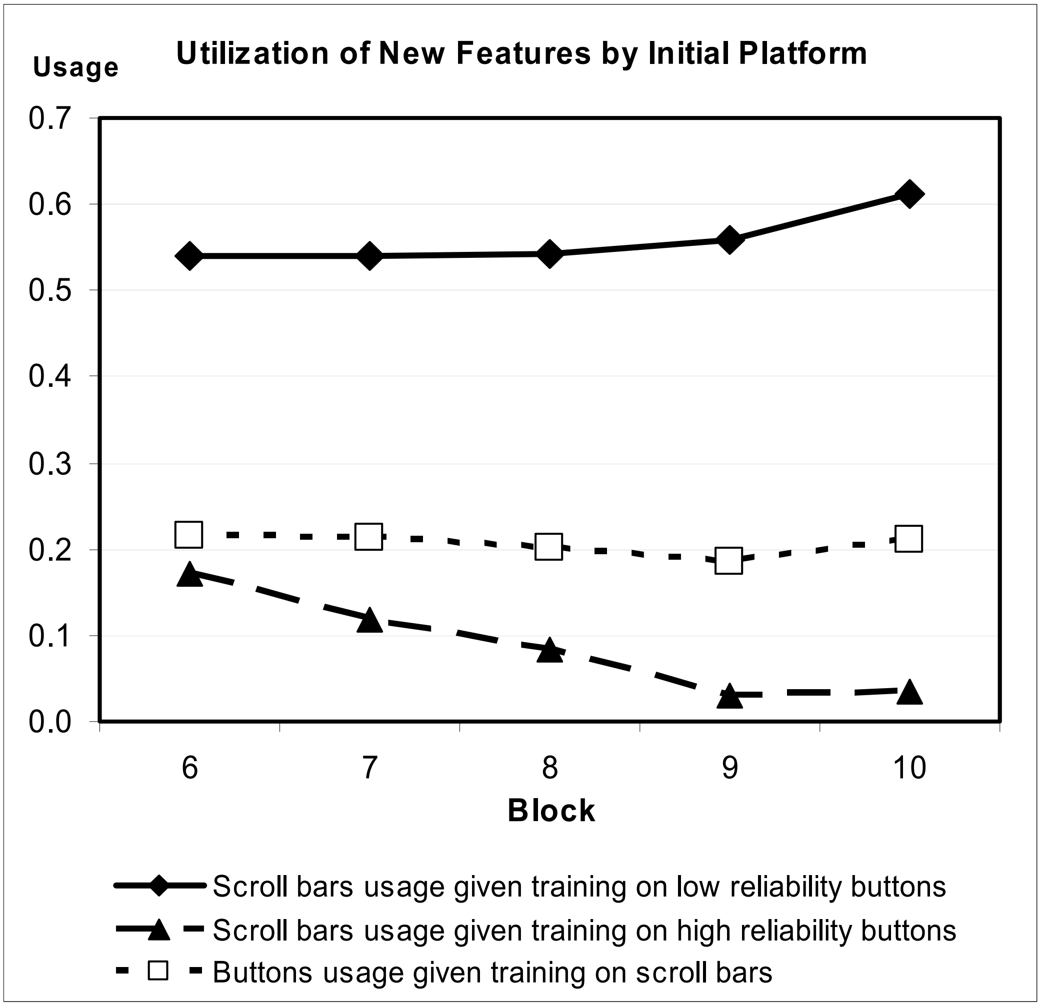

Notes:
Usage refers to the proportion of time the new features were used. Game trials are reported as blocks of three games.

**Figure 6: Scatter Plot of Implied and Stated *WTP*s**

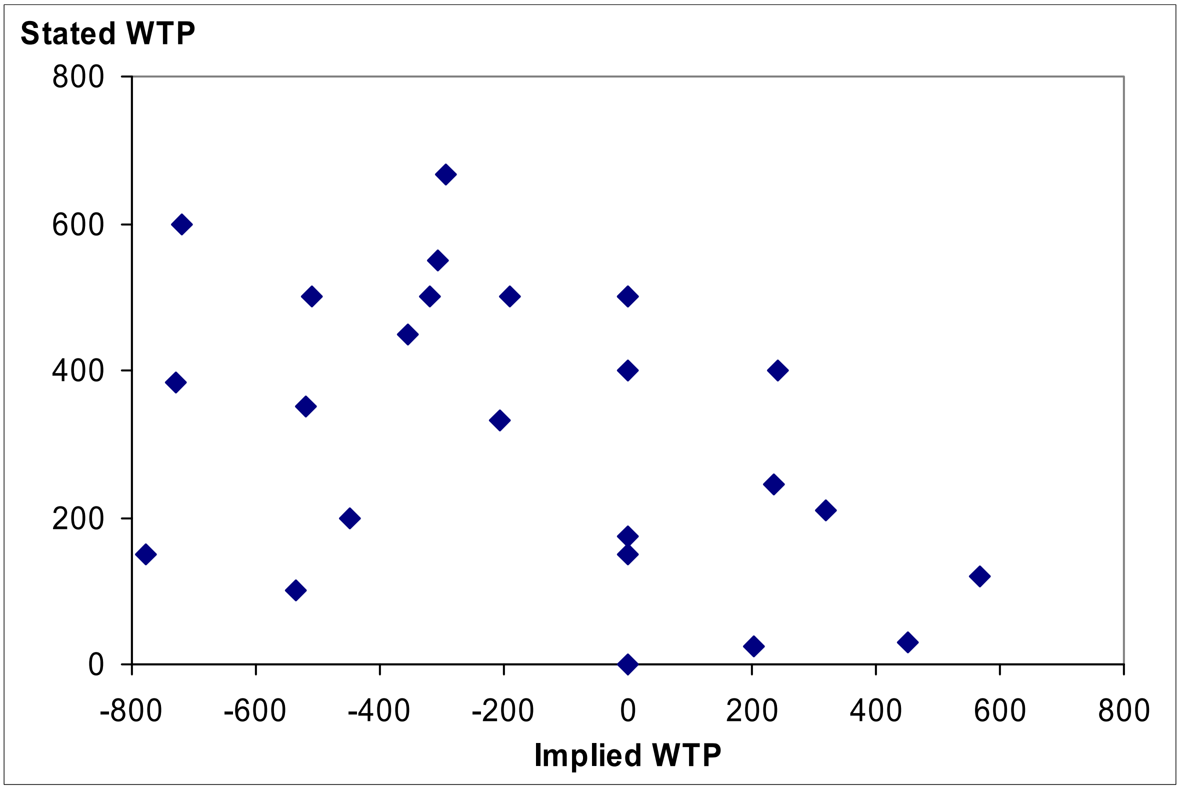

Note:
Implied WTPs were calculated by subtracting the predicted performance with the new combo platform from the predicted performance for using incumbent platform.

**Table 1: Determinants of stated *WTP*s**

| Dependent Variable: Log(*WTP*) | | | | |
|---|---|---|---|---|
| Variable | Parameter Estimate | SE | t-value | Pr > \|t\| |
| Intercept | 3.960 | 0.494 | 8.01 | <.0001 |
| Initial platform | | | | |
| Bad buttons | 0.718 | 0.379 | 1.89 | 0.0629 |
| Good buttons | 0.779 | 0.409 | 1.91 | 0.0611 |
| MAX6 | 0.005 | 0.002 | 2.14 | 0.0364 |

F(3,63) = 3.82, $p<.02$
$R^2 = 0.15$

Note:
MAX6 = best score over the last six games during training rounds

**Table 2: Effect of *WTP* on Subsequent Performance**

| Dependent variable: Cumulative performance during money rounds | | | | |
|---|---|---|---|---|
| Variable | Parameter Estimate | SE | t value | Pr > \|t\| |
| Intercept | 1102.58 | 350.57 | 3.15 | 0.0025 |
| Cumulative performance during training rounds | 0.97 | 0.14 | 6.85 | <.0001 |
| Gender[a] | -392.00 | 138.00 | -2.84 | 0.0061 |
| Log(*WTP*) | -169.98 | 51.00 | -3.33 | 0.0014 |

F(3,63)=21.57, p<.0001
Adj. $R^2$ = 0.51

Note:
[a] 1=Female and 0=Male

**Table 3: Determinants of New Control Utilization**

Dependent Variable: New control usage

| | Model 1 | | Model 2 | |
|---|---|---|---|---|
| Variable | Parameter Estimate | | Parameter Estimate | |
| Intercept | 0.472 | ** | 0.751 | ** |
| Initial platform | | | | |
| Bad buttons | 0.349 | ** | | |
| Good buttons | -0.113 | ** | | |
| Log(*WTP*) | -0.038 | * | | |
| MAX6 | | | -0.002 | ** |
| Game trial-linear | -0.003 | | -0.0026 | |
| Game trial-quadratic | 0.001 | | 0.0008 | |

| Model 1 | Model 2 |
|---|---|
| F(5, 849) =56.90, p<.0001 | F(3,851)=41.67, p<.0001 |
| $R^2 = 0.25$ | $R^2 = 0.13$ |

Note:
MAX6 = best score over the last six games during training rounds
* p<.05  ** p<.001

**Table 4: Forecasts of Performance with New Platform and New Control Utilization [a]**

| | Predicted | Actual |
|---|---|---|
| *WTP* | 366<br>(248) | |
| Performance change over existing platform implied by *WTP* | 22%<br>(15.6%) | |
| Performance change over 15 games for existing platform [b] | 17%<br>(20.1%) | 25% [c]<br>(35.1%) |
| Performance change over 15 games for new platform [b] | 9%<br>(18.7%) | 32%*<br>(42.7%) |
| Performance change over first 3 games with new platform [b] | 6%<br>(20.2%) | 21%<br>(101.5%) |
| New control utilization over first three games | 50%<br>(9.4%) | 24%*<br>(31.0%) |

Notes:

[a] The number reported are means with standard deviations in parenthesis (N=27).

[b] Percentage change, positive or negative, relative to last six games during training rounds.

[c] Mean for the control group in Experiment 1.

* Predicted is significantly different from actual (p<.05).

**Table 5: Self-reported reasons for upgrade**

| Reason for upgrade | Mean | | *SD* |
|---|---|---|---|
| New controls are useful | 0.76 | * | 1.48 |
| New controls offer flexibility | 0.48 | * | 1.05 |
| New platform is aesthetically more pleasing | 0.28 | | 1.24 |
| Desire for a change of pace | -0.12 | | 1.59 |

Note:
Items use a 7-point scale, anchored by -3 “Disagree a lot” to 3 “Agree a lot.”
* Mean is significantly different from 0, p<.05