




2004

The Psychology of Dynamic Product Maintenance

Robert J. Meyer
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/marketing_papers

 Part of the [Advertising and Promotion Management Commons](#), [Applied Behavior Analysis Commons](#), [Behavioral Economics Commons](#), [Business Intelligence Commons](#), [Cognitive Psychology Commons](#), [Marketing Commons](#), and the [Sales and Merchandising Commons](#)

Recommended Citation

Meyer, R. J. (2004). The Psychology of Dynamic Product Maintenance. Retrieved from https://repository.upenn.edu/marketing_papers/327

This is an unpublished manuscript.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/marketing_papers/327
For more information, please contact repository@pobox.upenn.edu.

The Psychology of Dynamic Product Maintenance

Abstract

The processes that underlie consumer decisions to invest in the maintenance of a durable good over time are examined. The work centers on a hypotheses that consumers make decisions about whether to repair or replace a good that has suffered a decrease in performance through a process that assesses the value of repair actions relative to two points of reference: the normal rate at which the performance of goods declines as they age (age-indexing), and how the timing and cost of the repair compares to parallel norms for repair expenditures (expenditure indexing). We show how these heuristics can be represented by a cognitive algebra that models maintenance decisions as a series of myopic utility-maximization problems. This process yields outcomes that can approximate those that would emerge from an optimal dynamic maintenance policy in some cases, but significantly depart from optimality in others. The algebra is then used to generate a series of predictions about how maintenance decisions may depart from normative benchmarks that are tested in a dynamic computer-pet ownership simulation. Actual maintenance behavior is characterized by a number of biases that are consistent with theoretical predictions, including a seemingly contradictory tendency to undermaintain and prematurely replace goods of superior value when they were acquired, yet be overly reluctant to part with and *over*-maintain inferior goods. A discussion of the implications of the work for understanding real-world biases in product care and maintenance behavior is offered.

Disciplines

Advertising and Promotion Management | Applied Behavior Analysis | Behavioral Economics | Business | Business Intelligence | Cognitive Psychology | Marketing | Sales and Merchandising

Comments

This is an unpublished manuscript.

The Psychology of Dynamic Product Maintenance

Robert J. Meyer*

* Robert Meyer is the Gayfryd Steinberg Professor of Marketing at the Wharton School of Business, University of Pennsylvania., Philadelphia, PA 19104. This work was supported by Grants from the Huntsman Center for Global Competition of the University of Pennsylvania and Singapore Management University. The author thanks participants at colloquia at Stanford University, the University of California, Berkeley, the University of Miami, and the University of Pennsylvania for helpful suggestions on early versions of this work.

The Psychology of Dynamic Product Maintenance

Abstract

The processes that underlie consumer decisions to invest in the maintenance of a durable good over time are examined. The work centers on a hypotheses that consumers make decisions about whether to repair or replace a good that has suffered a decrease in performance through a process that assesses the value of repair actions relative to two points of reference: the normal rate at which the performance of goods declines as they age (age-indexing), and how the timing and cost of the repair compares to parallel norms for repair expenditures (expenditure indexing). We show how these heuristics can be represented by a cognitive algebra that models maintenance decisions as a series of myopic utility-maximization problems. This process yields outcomes that can approximate those that would emerge from an optimal dynamic maintenance policy in some cases, but significantly depart from optimality in others. The algebra is then used to generate a series of predictions about how maintenance decisions may depart from normative benchmarks that are tested in a dynamic computer-pet ownership simulation. Actual maintenance behavior is characterized by a number of biases that are consistent with theoretical predictions, including a seemingly contradictory tendency to under-maintain and prematurely replace goods of superior value when they were acquired, yet be overly reluctant to part with and *over*-maintain inferior goods. A discussion of the implications of the work for understanding real-world biases in product care and maintenance behavior is offered.

The idea there was that consumers would bring their broken electronic devices, such as television sets and VCR's, to the destruction centers, where trained personnel would whack them (the devices) with sledgehammers. With their devices thus permanently destroyed, consumers would then be free to go out and buy new devices, rather than have to fritter away years of their lives trying to have the old ones repaired at so-called "factory service centers," which in fact consist of two men named Lester poking at the insides of broken electronic devices with cheap cigars and going, "Lookit all them WIRES in there!"

-Dave Barry, from "Mister Mediocre' Restaurants"

Paying for the upkeep of durable goods is an expense in life that most of us would just as soon do without. While we all enjoy contemplating the prospect of acquiring new goods and technologies, we often forget that with their acquisition comes the burden of repair and maintenance—a stream of ongoing costs that can become quite sizable. In 1997, for example, consumers in the United States spent over sixty-two billion dollars keeping their cars running, four billion cleaning them at car washes, nine billion repairing office machines and computers, three-hundred and fifty million repairing watches and jewelry, and over seven-billion dollars on dry cleaning¹.

Yet, as universal as these expenditures may be, they are also investment decisions that we are often remiss at making. We drive cars with under-inflated tires, live in homes with clogged gutters and un-weeded gardens, and have teenagers whose rooms, well, may never see upkeep at all. On the other hand, one could also point to examples that seem to go the other way, instances where we seem to *over*-invest in maintenance. Many of us have imposed draconian “no food in the car” rules during the early days of ownership, only to become lax in care a shortly thereafter, or found ourselves paying for the repair of an old appliance that we might have been better off simply replacing.

What is the consumer decision process that leads to these varied—and possibly sub-optimal--behaviors? Although the study of how consumers decide to acquire or

replace durable goods has long formed a central part of the literature of consumer decision making (see, e.g., Bayus 1991; Cripps and Meyer 1994; Pickering 1984; Winer 1985), comparatively little attention has been given to a natural complement to this work, decisions about how much to invest in maintaining these goods once they have been acquired. As a consequence, the efficiency of maintenance investments has remained largely a matter of speculation, and little theoretical guidance exists that might help influence how these decisions are made.

The purpose of this research is to take a step toward closing this gap by reporting the results of an investigation into the processes that underlie consumer decisions to invest in the maintenance of a durable good. The approach is both theoretical and empirical. We begin by advancing a hypothesis that consumers often overcome the computational difficulties associated with assessing the normative long-term costs and benefits of maintenance strategies by utilizing two reference-point heuristics: comparing the how a good's condition compares to that which would be typical its age (*age indexing*) and how an expenditure compares to that which would be typical in amount and timing (*expenditure indexing*). We then describe how these heuristics can be represented by a mathematical theory of dynamic maintenance. The theory is one that yields outcomes that closely approximate normative theory in some contexts, yet predicts significant departures from optimality in others. We conclude by examining degree to which the predictions of the theory are observed in a realistic dynamic product-ownership simulation where respondents care for a computerized "pet" over a multi-period horizon.

The Psychology of Repair

The normative principles of repair and replacement

Consider the following problem:

It is raining and you discover that your 10-year-old roof has a small leak above the attic. You phone a contractor who informs you that you have three options: you could replace the roof for \$10,000, patch the roof for \$500, or put a bucket under the leak and defer the decision to later. You can afford any option and have no immediate plans to move. Which of these options should you pursue?

As routine as this problem might appear, it is not one that has a straightforward normative answer. While mathematical methods for solving maintenance problems such as this forms one of the largest—and oldest—literatures in operations research management (see, e.g., Barlow and Hunter 1960; Chao-Ton 2000; Hayre, 1983; Usher, Ahmed, and Syed 1998; Zhang and Jardine 1998), the advice it offers for how to derive optimal policies is rarely simple.

To illustrate, consider what optimal-maintenance theory would say about how a homeowner should rationally decide what roof-repair action to undertake. The homeowner would be presumed to seek a general policy for making maintenance decisions both now and in the future that maximizes expected net utility over a planning horizon—such as the duration of tenure in the home. The best current action is that returned by this policy when applied to the current circumstances—a small leak above the attic with \$0, \$500, or \$10,000 repair options.

One of the central results of optimal maintenance theory is that if our homeowner can live with some simplifying assumptions about the dynamics of deterioration and repair—such as that the roof will deteriorate as a first-order Markov process and preferences are stationary over time--then the optimal policy can be described by a stable *control-limit rule*: it is optimal to repair, replace, or do nothing to the roof depending on how its current condition and/or age compares to a set of critical threshold values (e.g.,

Barlow and Hunter 1960; Assaf and Shanthikumar 1987). Armed with a bit of know-how in stochastic-dynamic programming, solving for these optimal thresholds then becomes a matter of computation. For example, a finite dynamic-programming algorithm would first derive the optimal threshold policy for some imagined terminal year (e.g., retirement) given all possible roof states, and then successively solve for the optimal thresholds for each early period assuming all subsequent decisions will be made optimally (e.g., Barlow and Hunter 1960).

Where intuitive and optimal worlds depart: Elements of a descriptive theory of maintenance

It is unlikely, of course, that consumers would make maintenance decisions in such a sophisticated manner. Yet, casual observation suggests that the intuitive rules that we actually use to make these decisions may not be that different from those prescribed by normative theory, at least on the surface. To illustrate, control-limit policies naturally arise in a wide range of actual maintenance tasks: we take our cars in for servicing when the mileage exceeds a manufacturer-suggested threshold, decide to paint the house if a certain number of years have elapsed since we last did so, and we have personal timing thresholds for knowing when it is time to update our wardrobes.² Likewise, intuitive maintenance decisions often respond to changes in task attributes in the direction predicted by optimal theory. Few people would see wisdom in taking a rental car to a carwash just before they are turned in (a rational response to a change in ownership horizons), and most would be more likely to replace rather than repair a good if the price difference is small (a rational response to long-term cost differentials).

On the other hand, while consumers may have an intuitive awareness of the some of the structural properties of normative theory, it less likely that these intuitions would extend to their actual numerical calculation. Central to this work is a hypothesis that consumers overcome the cognitive difficulties associated with computing the long-term costs and benefits of different maintenance strategies by utilizing heuristics that benchmark considered maintenance against norms for:

- 1) How condition of a goods typically degrade overtime in the course of ownership; and
- 2) How investments in maintenance are typically scaled and paced over time.

We will illustrate how these two classes of comparisons can guide—and possibly distort--maintenance decisions in turn.

Intuitive product life cycles. How bad we feel when witnessing damage to a possession is often not just a function of the severity of the damage but also *when* the damage has occurred in the lifespan of ownership. Most of us, for example, would likely feel a greater sense of trauma discovering a scratch on a brand-new car than one that had a few thousand miles on it (even when in the same starting condition), and greater distress getting a stain on a brand-new suit before we had a chance to wear it. Damage that occurs before a good has experienced a fair duration of use seems less forgivable than that which occurs later on, and—possibly--may yield a greater perceived benefit when repaired.

Would there be a rational basis for such feelings? In some cases, of course, there is: one *should* feel worse about damage to a brand-new possession because its diminishing effects on utility will be felt over a longer time horizon of ownership. But

similar product-age effects also arise in settings where ownership horizons would seem normatively irrelevant. To illustrate, we posed a sample of 87 business students with the following problem in furniture repair:

Imagine that you have a well-paying job in New York where you recently moved into in a new apartment on the upper-west side. Just yesterday a mint-new white couch was delivered to your apartment that you purchased for \$1500 from a furniture showroom in Soho. Unfortunately, while eating breakfast you accidentally spill a glass of red juice on the new couch that causes a large, quite visible, stain. The spill is particularly unfortunate since you are hosting a dinner party for some friends this evening who have never been to your apartment. You call the show room and they suggested two remedies for cleaning the stain:

- 1. They give you the name of a professional furniture restorer who, for \$195, will immediately come to your house and restore the couch back to its mint-new condition in time for the party; or*
- 2. For \$30 you can clean it yourself using a mix of commercial cleansers in time for the party, accepting that the repair might be less than perfect.*

You expect that you will own the couch for at least five more years, and you can afford the higher-cost repair if needed. Which option would you choose?

A separate group of 91 were given the same problem, but with the age described as being *five-years old* rather than brand new, but with all other features the same (it was in undamaged condition and would be owned for at least another five years)

A decision-theoretic analysis of this problem would prescribe that a rational decision maker who finds it worthwhile to spend \$195 for the more expensive repair in first scenario should find it no less worthwhile to do so in the second. The reason is that while the age of the couch is different in the two scenarios, the benefits of the expensive repair are the same: the couch would be restored to a like-new condition for a five-year horizon of ownership. Moreover, while it might well have imagined a greater level of liquidity in the second scenario, the effect would be to make the more expensive repair more rather than less attractive.

Did respondents behave in this way? In the aggregate they did not. In the case where the couch was brand-new, 62% of subjects preferred the expensive repair and 38% preferred the less-expensive option. When the couch was described as being five years old, however, respondents were more evenly divided: only 44% opted for the expensive repair while 56% voted for the less expensive ($\chi^2=5.85$; $p=.015$).

Two, closely related, explanations for the preference shift might be offered. One is that respondents who were contemplating the first scenario saw the spill as distressing by virtue of it being *premature* vis-à-vis the typical pace of wear and tear on a couch. Specifically, while couches are bound to suffer spills at some point, this one occurred before the owner had a chance to extract such benefits as the ability to show it off to friends for the first time, and/or enjoying the smell of its new fabric—considerations that enhanced the perceived severity of the loss. Alternatively, the locus of the effect may have been at the other end; respondents faced with the second scenario may have seen little benefit in trying to restore a 5-year old couch to a level *better* (mint condition) than it would typically be for that age, regardless of the fact that the benefits of the repair would be consumed over the same time horizon.

Intuitive expenditure patterns: the temporal disciplining of investments

In the same way that perceptions of the disutility of product damage may be indexed by norms about product aging and wear, decisions about whether and when to invest in maintenance may also be influenced by expectations about how such expenditures are typically paced over time. To illustrate, Cripps and Meyer (1994) report experimental evidence that individuals are reluctant incumbent goods with technologically superior new ones when the incumbent had just been purchased.

Likewise, Okada (2000) and Arkes (1996) report related survey evidence that when purchasing product replacements consumers are more responsive to promotions that allow them to trade in their existing good when it was acquired relatively recently (and is still functional), but monetarily-equivalent rebates given older goods whose functionality is perceived as having been exhausted.

The most common explanation that has been offered for these findings is that they are a consequence of a tendency for consumers to mentally amortize both the benefits and costs of new goods over the course of their ownership (see, e.g., Prelec and Loewenstein 1998; Gourville and Soman 1998; Okada 2000). By this account, consumers will be averse to replacing a durable good that had just been purchased because it would require them to abandon a stream of benefits that had been paid for, yet not yet fully extracted. In other words, the good would be seen as having a “residual book value” that is mentally added to the price of the prospective replacement as a transaction cost (Okada 2000).

An extension of these ideas to consumer decisions about the timing of investments in repair would seem natural. To illustrate, we posed a groups of 98 subjects the following variation of the couch problem described above:

Imagine that you have a well-paying job in New York that allows you to live in a nice apartment on the upper-west side. One of your favorite items of furniture is a designer couch that you purchased a couple of years ago that you keep in your living room. About a year ago a guest accidentally spilled a glass of wine on the couch that caused a rather ugly stain, and you paid \$150 to a professional furniture restoration company to restore it to like-new condition. Last night another accident occurred: this time you accidentally spilled a glass of tomato juice on the couch, leaving another stain. You consider three possible remedies:

1. *You could call the furniture restoration company back to restore the couch (for \$150);*
2. *You could try cleaning it yourself as best you can by experimenting with different commercial fabric cleansers for about \$30, accepting that the repair may be something less than perfect; or*

3. *Do nothing for the time being, and worry about cleaning it later. You have no plans to fully replace the couch and you can afford the highest-cost option if desired. Which option would you choose?*

A second group of 97 subjects was then given the same problem, but with an important variation: the \$150 repair was described as having been undertaken just one day—rather than one year—before the current spill.

If one can assume that the likelihood of future spills is the same in both scenarios (the overall frequency of historical spills is the same), optimal maintenance theory would predict that subjects respondents should reveal similar choices in both versions of the problem. The initial \$150 spent on repairs is a sunk cost that is normatively irrelevant to decisions about the second repair, regardless of whether it occurred a year later or a day later. Yet, subjects were quite sensitive to this timing change. In the condition where the first of the two spills occurred a year ago, subjects were equally divided between choosing the more or less expensive repair (43% preferred each), with 13% opting for deferral. When the accidents arose a day apart, however, there was a shift toward preferring the less-expensive repair for the second spill: 56% preferred the less expensive repair compared, 36% preferred the more expensive repair, and 5% deferred (overall $\chi^2(2)=5.54$; $p=.06$). In essence, subjects acted as if there was an allowable budget for repairs that was renewed only when a sufficient period of time had elapsed since the first repair.

To test for the possibility that this result might have accrued to an imagined liquidity or wealth constraint, we posed a third group of 81 subjects with a variation of the 2-day scenario in which there was a much larger expenditure the day before—the original purchase of the couch for \$1500 (paid by check). In contrast to the above result,

here we see a return to indifference between the two repair investments: 50% chose the more expensive repair, 48% the cheaper repair, and 1% deferred. As such, reduced liquidity does not seem to explain the reduced investment levels when the two spills occurred in close temporal proximity.

Why did subjects respond so differently to newly-*repaired* goods versus newly-*purchased* goods? One possible explanation is that the two kinds of expenditures were seen as being associated with quite separate mental amortization schedules. In this case the fact that that \$1500 had just been spent to *purchase* a couch appeared to be seen as irrelevant to the question of whether \$150 or \$30 should be spent to *repair* it. In contrast, when a second spill occurred shortly after \$150 had been spent on the same kind of expense (a repair), respondents acted as if the previous expenditure was now quite relevant.

A Cognitive Algebra of Dynamic Maintenance

Preliminaries

In this section we describe formal descriptive model of consumer maintenance that attempts to offer a unified account of a range of normative and non-normative features of intuitive maintenance decisions such as those described above. Within this framework both optimal and heuristic policies are shown to emerge as special cases of a more general dynamic decision process. After developing the model's structure we use it to derive a series of testable hypotheses about how actual maintenance behavior may depart from that prescribed by normative theory.

To lend tractability to the exercise we focus on modeling how consumers resolve a particular class of maintenance problems that we will be studying empirically in the

next section. Consider a risk-neutral consumer who wishes to consume the benefits of a class of durable goods that have a finite life expectancy E over a total time horizon T , $E < T$. At an initial time point $t=0$ the consumer purchases one of these goods and observes its initial performance level, s_0 . This level is a random draw from a known distribution $G(s)$ of new-product quality levels (reflecting, for example, chance variations in the returns from new-product searches). At each subsequent point in time the owner faces a constant hazard rate p that the good will suffer an accident that causes its performance level to decrease by some amount δ , where δ is a non-negative random draw from a known distribution $F(\delta)$. Finally, p and $F(\delta)$ are chosen such that $s_t > 0$ for all t ; i.e., the good is never subject to complete failure prior to its age limit⁴

Upon observing the state of the good (s_t) at each point in time the consumer is given the option to undertake one of three courses of action:

1. Consume the good in its current condition and receive the utility s_t ,
2. Pay a repair fee $c^R = r(s_0 - s_t)$ to restore the good to its initial condition, where r is a constant marginal cost of repair; or
3. Pay a replacement fee $c^B = B$, $B > c^R$ to acquire a new good that would allow the consumer to receive the performance level s_0^* where s_0^* is a new random draw from $G(s)$. If a replacement is purchased the forgone good is assumed to have no salvage value; that is, it is not sold or used in conjunction with the replacement.

The consumer's objective is to make a sequence of these decisions so as to maximize the total utility of ownership over the time horizon while minimizing total repair and replacement costs.

While this ownership problem is a potentially complex one, in Appendix One we show that it is associated with a comparatively straightforward—though not necessarily intuitive--optimal maintenance policy:

1. *On the initial trail, $t=0$, observe the quality of the initially-drawn good, s_0 . If s_0 is less than a threshold acquisition quality Q^* , immediately pay c^B for a replacement, and continue doing so until a good of quality Q^* or higher is obtained;*
2. *Once an acceptable good is acquired, proceed with its ownership for the duration of the life expectancy E , paying for the repair of all damages until a termination age of K^D and none thereafter.*

But while the general form of the policy might be straightforward, the task of computing optimal control parameters Q^* and K^D is a difficult one that few decision makers could be expected to intuitively undertake. Our central hypothesis is that when faced with such a decision task individuals will not make maintenance decisions as above, but rather by a heuristic process that mimics some of its prescriptive properties. Specifically, decision makers are posited to approach maintenance problems as a series of short-term utility-maximization decisions that differentially encourage or dissuade maintenance actions based on how similar they are to those that optimal policies would typically prescribe, or that have been found to work in similar settings in the past. The utility functions that describe this process exhibit two key properties:

- 1) *Age indexing of utility*: the utility of a good in each period is assessed by a function that differentially motivates repair versus replacement depending on

how a good's observed condition compares to that *expected* for a good of that age; and

- 2) *Temporal indexing of costs*: decisions about whether to invest in a maintenance action that would increase the utility of a good are based on a psychological cost function that compares the timing and magnitude of an expenditure to that typically expected over the life of a good.

We develop these ideas more formally in two stages. We first describe these properties in greater detail in the form of three behavioral axioms, and summarize the overall choice function that we hypothesize drives maintenance decisions. We then explore the implications of this model for likely empirical patterns of consumer maintenance behavior.

The Behavioral Axioms

1. Age Indexing

A common normative prescription of optimal maintenance models is that as goods age they become more worthwhile to replace it than repair. The reason is that goods are often made of wearable parts that have an inherent higher likelihood of failure with age (e.g., cars and washing machines), or are subject to obsolescence by the emergence of superior replacement goods (e.g., computers and fashion). Yet, knowing the exact threshold point at which it becomes more worthwhile to replace than repair—if one exists at all—is not an easy computational problem to solve.⁵

We hypothesize that consumers overcome the complexity of such calculations by employing a general heuristic we term **age indexing** that is applied to all product ownership problems, even those where there is no mechanical aging of parts (such as the

current one). Under age indexing consumers are assumed to assess the marginal disutility of a given observed unit of damage by comparing how the level of performance of results from that damage contrasts with that which would be typically expected from a good that has undergone a normal rate of wear. We formalize this idea in terms of the following axiom:

Axiom 1: the Age-Indexing of utility. Let s_0 be the objective performance state of a good when it is first acquired, s_t be its observed state at time t , and A_t be the age of the good at t . We hypothesize when assessing the attractiveness alternative maintenance actions at a given time t consumers assess the utility of consuming a good of age A_t in state s_t , $u(s_t)$, by the reference-dependent process:

$$u(s_t) = k_0 + k_1 u_o(s_0 H(A)) + k_2 u_r(s_t - s_0 H(A)) \quad (1)$$

where $H(A)$ is a monotonically decreasing aging function describing the consumer's beliefs about the typical pace at which the utility conveyed by a good declines as it ages, where $H(0)=1$, $\lim_{A \rightarrow \infty} H(A) = 0$ (e.g., a negative exponential), and k_0, k_1, k_2 are nonnegative scaling constants such that $k_0 + k_1 + k_2 = 1$. We assume that the marginal utility function for normal wear ($u_o(s_0 H(A))$) is convex over $s_0 H(A)$, and that for departures from the age-indexed state ($u_r(s_t - s_0 H(A))$) is asymmetric about its origin, being steeper in losses than gains⁶. Hence, initial decreases in utility from a good's initial state are hypothesized to be more salient than later ones, and the marginal displeasure of seeing a good age prematurely is assumed to loom larger to consumers than the pleasure of seeing it maintained in superior condition for its age.

Expression (1) describes an assessment process in which the utility that consumer associates with owning a good reflects a balancing of two age-dependent forces:

- 1) A tendency to hold decreasing utility for goods as they age regardless of condition (the marginal utility function $u_o(s_0H(A))$); and
- 2) A tendency for this base aging effect to be conditioned by how the actual condition of the good contrasts with that which would be typical for its age (the marginal utility function $u_r(s_r - s_0H(A))$).

The model provides a characterization of the age effects on judgment illustrated earlier: a constant unit of damage will be seen as having a greater disutility (hence be more likely to be repaired) if it is the first incident of damage to the good (by the convexity of $u_o(s_0H(A))$) and/or when it is perceived as premature; i.e., results in a negative value of $u_r(s_r - s_0H(A))$ (by the asymmetry in slopes about the origin).

While a key implication of expression (1) is that consumers will be increasingly indifferent to the loss of goods as they age, the expression also allows for the reverse to be true in special cases. Specifically, if a good is maintained in mint condition as it ages—such as an antique—expression (1) implies that its utility will be marked increasingly positive values of the contrast function $u_r(s_r - s_0H(A))$. This, in turn, would imply that the sudden loss of such a good could invoke a sense of loss that is comparable or greater than that had it been brand new.⁷

2. Temporal indexing of expenditures

While age indexing describes the process by consumers assess the disutility of product damage, it does not characterize how consumer will decide to respond to these

assessments. We hypothesize that in the same way that consumers may use *age indexing* to assess the marginal disutility of damage to a good during the course of ownership, *expenditure indexing* is used to assess of the wisdom of a prospective investment in repair or replacement by contrasting its timing with that of the typical pace and magnitude of expenditures.

The hypothesized process is as follows. When a consumer initially decides to invest in a good—being it to pay for a repair or purchase a replacement—this investment is accompanied by an implicit expectation about its duration; specifically, the length of time that will elapse until another such investment is required. For example, when one pays for a plumbing repair there is an implicit expectation that it will hold up for some length of time, and new car purchases are usually accompanied by a belief about how long the car will be held before a new one will be bought. Building on the mental amortization models of Okada (2000) and Gourville and Soman (1998), we hypothesize that these expectations serve to discipline the pattern of consumer spending over time by imposing psychic transaction costs that discourage investments that would be premature relative to timing norms.

We describe this idea through the following axiom:

Axiom 2: The temporal indexing of costs. Let D_t^k be the length of time that has elapsed since the consumer last undertook maintenance action k for a currently-owned good, M^k be the consumer's expectation of the *normal* length of time between successive such investments. For repairs, the duration D_t^R would be either as the length of time since the last repair or, if a good has never suffered damage, its current age. Although we do not explicitly model the evolution of

beliefs about typical durations, M^k is assumed to reflect either externally-established norms about durations (e.g., what *Consumer Reports* recommends for how often one should paint a house) or past experiences (the last interval between house paintings). We posit that when considering whether to undertake a maintenance action the consumer's perception of its costliness is not just a function of the objective size of the expenditure, but also its temporal proximity to previous expenditures of the same type. Formally, the *psychological cost* of action k , at time t , $\varphi(c_t^k)$, is hypothesized to be given by the value function

$$\varphi(c_t^k) = k_0 + k_1 v_1(c_t^k) + k_2 v_2(M^k - D_t^k) + k_3 v_3(c_{t-D}^k / D_t^k) \quad (2)$$

where c_t^k is the objective cost of action k at time t , c_{t-D}^k is the size of the last investment in action k , and $v_1()$, $v_2()$, and $v_3()$ are positive monotone scaling functions with associated constants k_i . It is assumed that the difference function $v_2(M^k - D_t^k)$ is asymmetric about its origin, positive and steeply sloping for $M^k > D_t^k$ (the case of accelerated investments) and slightly sloping for all $M^k < D_t^k$ (decelerated investments).

Expression (2) describes a heuristic mechanism by which consumers discipline the temporal pacing of expenditures without engaging in explicit forward planning or dynamic reasoning. The timing function $v_2(M^k - D_t^k)$ serves to make expenditures for a particular category seem psychologically more costly if the considered duration (D_t^k) is much shorter than would the norm (M^k), and the historical expense function $v_3(c_{t-D}^k / D_t^k)$ serves to either temper or inflate this effect based on the size of the last investment. That is, while accelerated expenditures (positive values of $(M^k - D_t^k)$) will

inflate psychic costs, this effect will be diminished if the last expense was a trivial one, but amplified if it was large. Expression (2) implies that consumers will be least prone to undertake those maintenance actions that were immediately preceded by a large investment of the same type (i.e., $M^k \gg D_t^k; c_{t-D}^k / D_{t-D}^k \gg 0$).

Expression (2) exhibits two important properties. First, while the core implication is a psychic penalty for accelerated maintenance expenditures, it also implies that consumers may be excessively *prone* to undertaking expenditures when the time since the last investment increases beyond what is normal (i.e., cases where $M^k < D_t^k$). Hence, a consumer who has been fortunate to own an appliance that has remained trouble-free for a much longer period than would be predicted by (2) to be more prone than usual to pay for an expensive repair when the need finally arises. In essence, this delay acts as a kind of mental “savings” that is deducted from the perceived cost of the next similar maintenance (or replacement) action⁸.

Second, note that the expression also allows a description of the temporal mental budgeting process that we illustrated earlier in which the psychological treatment of repair expenditures are treated quite differently from that of replacement expenditures. Specifically, in (4) the transaction penalty associated with a repair is a function solely of the timing and cost of the last repair—not the timing and cost of the last replacement. Hence, (2) predicts, *ceteris paribus*, that that a consumer would have the same adverse reaction to having to undertake a premature repair of a good that was purchased a year earlier compared to one that was purchased a few days earlier.

3. The Maintenance Choice Process

Given assessments of the marginal benefits and psychic costs of different maintenance actions (Axioms (1) and (2)), consumers are hypothesized to choose maintenance actions over time through a sequence of static (non-strategic) utility-maximization decisions. Formally, at each point in time the consumer is hypothesized to observe the state of the good, and decide whether to repair, replace, or accept it in that condition (defer) by choosing the option that offers the best prospective stream of utility as follows:

Axiom 3: The choice axiom. Let V_t^D, V_t^R, V_t^B be the prospective values of deferral, repair, and replacement, respectively at time t , $u(s_t^i)$ be marginal utility of consuming a good in state i at time t as in expression (1) and $\varphi(c_t^k)$ be the psychic cost of undertaking maintenance activity k at time t as in expression (2). In addition, let $f(\beta|q)$ be a $(0,1)$ bounded function that describes the subjective rate of discounting for events q periods in the future, M^k be the expected duration of activity k (as above), E be the finite life expectancy of the good, T the length of the total planning horizon, and s_0^* be the expected value of the performance a product replacement (as described above).

We posit that at each point in time the consumer chooses the maintenance action that satisfies the general maximization rule:

$$\text{MAX} \begin{cases} V_t^D = k_D + k_{1D}u(s_t) \\ V_t^R = k_R + k_{1R}v_R \left(\sum_{q=0}^{\text{Min}(E-t, M^R)} f(\beta|q)u(s_0 H(A)) - \varphi(c_t^R) \right) \\ V_t^B = k_B + k_{1B}v_B \left(\sum_{q=0}^{\text{Min}(T-t, M^B)} f(\beta|q)u(s_0^* H(A)) - \varphi(c_t^B) \right) \end{cases} \quad (3)$$

In the appendix we show that expression (3) can be derived as a reduced-form approximation to a normative maintenance policy that assumes that goods decay deterministically (by $H(A)$) for finite durations (M^k) under a high rate of mental discounting. Behaviorally, it describes a choice process where consumers solve dynamic maintenance problems by comparing at each point three envisioned future utility streams:

- 1) That restoring the good to its original performance level s_0 for the cost c^R_t and the consuming it under a normal aging rate ($H(A)$) for either the balance of the planning horizon or the expected duration of the repair;
- 2) That of replacing the good with a new one that provides an expected performance level s_0^* for the cost c^B_t and consuming it under a normal aging rate for the expected time until the next replacement; or
- 3) That of accepting current performance s_t for the current time period.

Implications

The normative policy for this problem characterizes maintenance as an all-or-none-affair. If one feels that one could obtain a better-performing good by purchasing a replacement, the optimal time to do so is at the immediate outset of ownership, when its benefits can be realized over the longest horizon. Likewise, repairs should never be intermittent; if one is considering repairing a good one should do so either right away or never. Would a consumer who makes decisions by the process summarized in expression (4) follow these principles? The answer is both yes and no. On one hand, consistent with normative theory, (3) yields such straightforward predictions as:

R1: Consumers will be more likely to buy a replacement when initially endowed with of a low-valued good ($s_0 > E[s_0^*]$); and

R2: Be more reluctant to invest in both repairs and replacements given decreases in the anticipated horizon of amortization (the summation limit) and increases in the perceived cost of repair.

On the other hand, the expression (3) also implies a number of likely systematic departures from normative theory. We illustrate these in terms of four empirical hypotheses.

Under-maintenance when vigilant care is optimal. A key feature of the normative policy is if it is optimal for a consumer to invest in the maintenance of a good, the commitment to maintenance should be complete. Hence, a rational consumer would never let a good deteriorate through several accidents before undertaking repair, or gradually withhold investments as it ages. Yet, application of expression (3) to a setting where repair is always optimal would not likely yield such vigilance. By *R1* above, as a consumer's subjective rate of discounting increases relative to that assumed by a normative analysis (the consumer becomes increasingly myopic), the out-of-pocket costs of *all* repairs will loom increasingly large relative to doing nothing, discouraging investments in maintenance. Moreover, if the consumer holds a subjective aging function $H(A)$ that is decreasing in A and/or approaches maintenance with a belief that replacement intervals are finite ($M^B < T$), this under-maintenance bias will be exacerbated over time. Specifically, increases in the slope of $H(A)$ will decrease the marginal disutility associated with damage for older goods (by expression (1), Axiom 1) and

decreases in M^B will shorten the imagined duration over which the benefits of repairs will be realized (by expression (2), Axiom 2). Hence,

H1: Given ownership of goods for which normative theory prescribes vigilant maintenance, consumers will exhibit a global under-maintenance bias that is exacerbated as the good ages.

Excessive ownership of and investment in low-valued goods. Consider the converse case where a consumer is initially endowed with a low-valued good, formally, one whose quality is lower than that likely to be provided by a random replacement ($s_0 < E[s_0^*]$). As noted above, the optimal ownership strategy in this case is straightforward: one should immediately replace it with a new one, prior to any investments in maintenance. There would, therefore, be no normative basis for an ownership strategy that holds the good for a period of time, invests in limited repairs, and then pays to acquire a replacement.

Axioms 1 and 2, however, imply that such ownership behaviors, may, in fact, be commonplace. First, as above, Axiom 2 predicts that consumers will be reluctant to pay for purchases of replacement goods until the purchase cost of the incumbent has been at least somewhat mentally amortized. In addition, Axiom 1 predicts a bias toward seeing an enhanced disutility for damage to *all* newly-acquired goods, something that would encourage investments in repair even for goods that will likely soon be replaced. In sum,

H2: When endowed with an inferior good for which that normative theory prescribes immediate replacement, consumers will retain ownership longer than they should before undertaking replacement, and undertake positive investments in their maintenance..

Exacerbation of the under-maintenance bias for the old and diminutive. A consequence of the linearity in the cost and reward structure of the current problem is that the optimal maintenance policy is indifferent to both the quality (s_0) of the currently-

owned good as well as its age—up to optimal care-termination period. The hypothesized decision process, however, implies that this principle may be frequently violated empirically. First, it can be verified by inspection of expression (3) that the relative utility of a repair option V_t^R is predicted to be strictly decreasing in a good's initial quality, and by the concavity of the normal aging function $H(A)$ (expression 1) the relative utility of repair versus deferral ($V_t^R - V_t^D$) is also decreasing in s_0 . Second, expression (3) also characterizes the increased aversion for investing in repairs as a good ages as arising gradually (through decreasing summation limits on the returns to repairs) rather than as a discrete step. Hence, maintenance of goods will likely be gradually withdrawn earlier than would be prescribed by optimal theory, which, in turn, would encourage an earlier-than optimal replacement. In summary,

H3: The global tendency to under invest in maintenance when it is optimal will be increasing in the age of a good and decreasing in the initial quality of the good.

Temporal restrictions on spending: premature, delayed, and multiple damage effects. Above we offered tentative survey support for two explicit predictions that emerge from Axioms 1 and 2: a diminished marginal likelihood of investing in repairs that are temporally proximate (by expenditure indexing), and a heightened likelihood in investing in the repair of damage to just-acquired (by age indexing). We also noted, however, that the cost-indexing model (Axiom 2), also predicts a heightened enhanced likelihood of investing in the repair of goods for which the first incidents of damage are *delayed*, when the psychological cost of a repair is diminished by the absence of previous expenditures. More formally,

H4. The likelihood of investing repairs will be enhanced when damage is perceived as either premature or highly delayed relative to normal wear rates, and diminished when multiple damage events requiring independent investments occur in close temporal proximity.

Caveats: learning and the mechanism of damage

As currently formulated the theory says little about two aspects of maintenance behavior that may nevertheless prove empirically important: temporal changes in the form of assessment functions and the mechanism by which damage occurs. As formulated the model is structurally static, something that is unlikely to hold in practice over successive generations of consumer ownership of goods. The most likely way that decision rules will evolve through experience, however, is far from clear. On one hand, heuristics may evolve toward optimality over time as decision makers become more experienced in making maintenance decisions. On the other hand, one might argue just the opposite: biases may be exacerbated by a tendency for beliefs about normal durations of ownership (M^B) and aging rates ($H(A)$) to be self-fulfilling. Specifically, a consumer who begins ownership of a good with overly pessimistic beliefs about M^B and $H(A)$ will be discouraged from investing in maintenance to the level that would be prescribed by optimal theory. This under-maintenance bias, in turn, would translate to reduced expectations about ownership durations, which would, in turn, further reduce lower levels maintenance for the next good that is owned.

Likewise, the nature of biases may also be sensitive to the mechanism by which decreases in performance through wear occur. To illustrate, above we modeled product decay as a series of probabilistic decreases in the observed quality of the product that

occur with a constant hazard rate (p). We might just as easily, however, modeled decay in as a series of increases in the hazard rate itself; a difference that may be critically psychologically if not mathematically.

Because of the theoretical indeterminacy of these issues, we leave them as empirical issues to be explored in the experimental work below. Possible generalizations of the theory to account for their observed effects will be raised in the concluding discussion

Empirical Analysis

Overview and general procedure

In this section we describe a program of experimental work that examines the ability of the proposed theoretical structure to explain maintenance behavior in a realistic controlled setting. The experiment posed a sample of subjects with the task of purchasing, owning, and maintaining a series of computerized electronic pets over a lengthy time horizon.⁹ This context and its implementation was the outgrowth of several generations of game design that utilized different approaches to manipulating product utility, the nature of experienced damage, and the mechanics of repair and replacement. Our central goal was to create a stimulus environment that would be seen by subjects as offering a reasonably realistic and involving portrait of how these variables manifest themselves in real-world settings, yet would be sufficiently abstract as to minimize the confounding effects of beliefs about product care and replacement policies that might be unique to the real world.

The goal of the empirical work was twofold. The first was to examine the degree to which subjects' decisions about how much to invest in the maintenance of their

computer pets displayed departures from optimality consistent with hypotheses **H1-H4** above. As such, we created experimental settings where the normative policy was to pursue either vigilantly maintenance of a good (for example, a case where repair costs are low relative to replacement costs and the performance of the current good is higher than that which could be expected from a replacements) or never pay to maintain it (the case where repair costs are high and/or the value of the current good is much lower than the expected value of replacements). The second goal was to establish a body of empirical evidence on the two exploratory issues raised above: the dynamics of biases and variations in the mechanism of product decay.

Subjects and Procedure

One hundred-twenty-four graduate and undergraduate students volunteered to participate in return both for course credit and a monetary incentive. Subjects were run in groups of five to twelve in a University computer laboratory over a two-week period, with the average session lasting forty minutes. The monetary incentive took the form of four fifty-dollar cash prizes to the top four performing players.

Upon entering the lab subjects viewed a computer screen that provided the following overview of the task:

“This game tests your skills in product ownership, where the stakes are real cash. For the next 30 minutes or so we want you to imagine that you live a world where people enjoy having as pets exotic electronic animals. They are expensive toys—they cost \$500 each—but ownership brings lots of pleasure that make the purchase worthwhile. But there are few drawbacks with these pets. First, they do not live forever; after they reach the age of 75(weeks) their batteries wear out and you have to go back to the store and buy a replacement. Second, the amount of pleasure they provide varies from pet to pet and you only find out how much pleasure it provides after your purchase it. Finally, they are rather fragile. When you play with them they occasionally suffer accidents that lower the pleasure they deliver, and occasionally suffer illnesses that increases the frequency with which these accidents occur. Both can be remedied, but these incur repair costs.”

The ownership game involved two recurrent phases that are illustrated in the sequence of screen captures in Appendix 3: a purchasing phase and an ownership phase. In the purchasing phase subjects were taken to a hypothetical “pet store” (Appendix 3, step 1) where they were asked to make a choice among three different cartoon images that depicted the appearance of the pets. After making a choice subjects were taken to a “training room” (Appendix 3, steps 2 and 3) where the quality of the pet they just purchased was randomly determined by having the subject spin a series of computer wheels. Subjects were told that the level of pleasure they derived from owning the pet (on a 100-point scale) was a function of the number of “tricks” it could be taught to perform, (e.g., playing cards), and the purpose of the spins was to discover this ability for their pet. Each pet came with a starting pleasure level of 30, and for each successful training exercise their pet’s pleasure was enhanced by 10, to a maximum of 100. The purpose of this rather involved process for determining initial pleasure levels was to dramatize to subjects the chance nature of performance levels, emphasizing the difficulty of replacing pets with high levels of performance (cases where $s_0 \gg E(s_0^*)$), and the ease of those with low levels (cases where $s_0 \ll E(s_0^*)$).¹⁰

After acquiring a pet subjects began the main phase of the experiment (Appendix 3, step 4), where, on each trial, subjects could accumulate pleasure points by clicking on a “play” button. At the start of each round subjects were also informed if the pet had suffered one of two forms of maladies (Appendix 3, Step 5):

1. An accident that dirties the pet, lowering its pleasure level from one to twelve points (a uniform random draw);
2. and/or an illness that causes odds of such an accident to increase by 25%.

Upon initially acquiring each malady had a 12% of occurring on a given trial. This 12% rate remained constant through the task for incidents of illness, but could increase for rates of accidents if illnesses were left untreated, to a maximum of 50%. Subjects were informed of the accident generating process at the start of the experiment, and were reminded throughout by being provided with information about the expected number of accidents they might expect for a pet of its age assuming that it was not ill, and, if the pet was sick, the updated probability of an accident.

To remedy accidents and illnesses subjects could take their pet either to a “grooming salon” that would restore a pet’s original quality level or a “hospital” that would reduce the damage probability to the original .12 level. The cost of the restoring a pet depended on how severely its condition had deteriorated from its new state. For the case of cleaning, subjects were assigned to one of two marginal cost conditions designed to induce optimal policies that either favored or discouraged paying for such repairs. In the low-cost condition subjects were charged \$5 for every pleasure unit a pet had deteriorated from its new condition, and in the high-cost condition subjects were charged \$40 per unit. The cost for curing illnesses or restoring accident rates was constant in the experiment, and was defined by a two-tiered structure mimicking that often associated with doctor visits. Respondents were charged \$25 to restore a pet that was mildly sick—defined as having accident rates that had not grown to more than 50% of the initial (.12) level--and \$100 if the pet was more severely ill. Finally, at any point in the task respondents had the option of returning to the pet store and replacing their current pet with a new one, for a cost of \$450. There was no salvage value of a discarded pet.

The total length of the game was 300 decision periods, with each pet having a maximum lifespan of 75 periods. A finite life for each pet was introduced to allow investigation of whether maintenance policies evolved over successive generations of ownership of pets.

Optimal play

Above we noted that the optimal maintenance policy for this task will be that of a two-stage acquisition-threshold and repair-duration rule. In the experiment respondents were faced with two kinds of repair decisions in the task—those restoring increases in accident rates and those restoring quality losses—hence the optimal maintenance policy requires us to solve for two care-termination ages—one for accident rates (K^S) and one for quality losses (K^D). We explored the solution surface for the control parameters Q^* , K^S , K^D by numerically simulating 5000 plays of the of the game at each point over an exhaustive gradient of parameter values. In Figure 1 we illustrate partial profiles of these analyses plotting measured earnings over a range of values of K^S and K^D for the high- and low-cost conditions for when the optimal acquisition strategy is followed (Fig. 1a and 1b), and a range of acquisition thresholds for when the optimal subsequent maintenance policy is followed (Fig. 1c).

This analysis suggested the following point optimum ownership policy:

1. *On the initial trail, $t=0$, observe the quality of the initially-drawn good, s_0 . If s_0 is less than 80, immediately pay for a replacement, and continue doing so until a good of quality 80 or higher is obtained.*
2. *Once an acceptable good is acquired, proceed with one of two repair strategies depending on the marginal cost of cleaning. If the cost is low (\$5) pay for the*

repair of all increases in accident rates until a termination age of 70 and all decreases in pleasure points until a termination age of 35. If the cost is high (\$40) pay for all increases accident rates until the age of 60, but never pay for cleaning..

We should emphasize that this policy is quite robust to deviations from optimal control parameters, but only *if* a player intuitively grasps three core elements of the optimal solution: one should avoid setting too high a hurdle for initially accepting a pet, never pay for cleaning when the costs are high, and be vigilant about repairing increases in accident rates. Once grasped, the optimal policy can be quite forgiving. For example, as shown in Figure 1c, while expected earnings are maximized when one initially rejects all pets with qualities below 80, a decision maker who plays a more modest acceptance rule of $Q^*=60$ can do almost as well realizing expected earnings that are within 6% of the optimum. Likewise, when faced with low cleaning costs (Figure 1a) a player could err in setting the termination period for cleaning by as many as 20 periods and still realized earnings that are within 10% of the optimum. On the other hand, the policy is much less forgiving of players who fail to intuit the above core principles. Players who retain a policy of vigilantly paying for cleaning in the high-cost condition after a pet reaches the age of 20, for example, will realize negative expected earnings, even given vigilant maintenance of accident rates (Figure 1b).

Analysis and Results

Overall performance. Over the course of 300 total periods of ownership the 124 subjects in the low cleaning cost condition owned an average of 6.9 pets from which they realizing an average net performance score of 14167, while those in the high cleaning-

cost condition owned an average of 8.4 pets from which they realized a net performance of 11964. As a point of reference, a subject who followed the optimal policy would have been expected to have purchased an average of 6.72 pets over this same period in both conditions¹¹, and realized a mean net performance of 16596 points in the low-cost condition and 13960 in the high cost. Hence, the decision processes used by respondents yielded decisions that were not unskilled, revealing a level of achievement similar to that which would be attained by a range of policies that capture the gist—though not the precise details—of the optimum policy (Figure 1). As might be expected, there was considerable individual variation about these central values, with number of owned pets ranging from as few as four to as many as twenty, and realized earnings ranging from as little as 4789 points (in a high-cost condition, 34% of the optimum) to 21518 (in a low-cost condition, 129% of the optimum). The percentage of subjects who realized earnings that were below the expected optimum was 77% in the low cost condition and 74% in the high cost.

The discussion of the behaviors that formed the basis of this general result is organized in two phases: the character of decisions whether to pay for repairs given decreases in performance, and then the character of replacement decisions.

Repair decisions

The normative policy prescribes that once a decision is made to own rather than immediately replace a pet, owners should repair virtually all illnesses that cause increase in its accident rate, but pay to repair decreases in quality: more sparingly: never when the cost is high, and during the first half of the pet's life when the cost is low. The hypothesized evaluation process, however, predicts a more complex maintenance pattern:

a global tendency to under-maintain when vigilance is called for (H1), over-maintenance of inferior goods that would be better replaced (H2), and a tendency for maintenance to gradually decline with a pet's age and its initial quality (H3). In addition, the proposed algebra also makes other, more detailed, predictions about maintenance behavior, such as tendency to be averse to paying for temporally-proximate repairs (H4).

In Figure 2 we plot the observed relative frequency with which respondents paid for the repair of a good given either an accident or illness by its proportional age, defined as its current chronological age relative to its realized maximum. We plot proportional age to control for individual differences in the length of time a given pet was owned—a duration, as noted above, that was almost always shorter than the theoretical optimum. Figure 2a plots conditional maintenance rates over time by repair type (averaging over costs), and 2b plots cleaning rates over time for the two cost levels.

The data provide a mixed view of the intuitive rationality of respondents in the task. On one hand, congruent with normative theory, subjects were more consistently vigilant in paying to repair increases in accident rates (cures) than decreases in quality (cleans; Figure 2a), were less likely to pay for cleaning when the marginal costs were high (Figure 2b), and were less likely to paying for a repair as the terminal period of ownership approached. Subjects' implicit recognition of the greater importance of repairing increases in accident rates versus quality decreases is also notable in dispelling the suggestion above that subjects might be prone to overlooking repairs of damage that does not directly affect the utility drawn from a good, at least in this case.

On the other hand, the data also showed a number of systematic departures from optimality that was more supportive of the view of decision making provided by the

proposed algebra. Strongly supporting H1, in cases where repair was optimal, conditional maintenance rates were far less the vigilance prescribed by the optimal policy. Specifically, the mean conditional likelihood of repairing increases in the accident rate over the first 90% of each pet's lifespan was 41%, while the mean likelihood of repairing a loss in quality when it was optimal (\$5 repair cost, the first half of each lifespan) was 31.1%¹². In contrast, while it was *not* optimal to pay for cleaning in the high-cost condition, subjects *over* invested in maintenance. In the \$40 cost condition when no repairs should have been undertaken, there was nevertheless a 16.6% repair rate over all trials. Finally, consistent with H3, subjects' tendency to withdraw care as a pet aged was quite gradual, showing no clear evidence of the homogeneous termination point prescribed by optimal theory.

To explore the effect of pet quality on repair rates, in Figure 3 we plot histograms of the conditional probability of undertaking repairs given ownership of pets of each possible starting quality level. Although normative theory prescribes that decision makers should be indifferent to starting quality levels in their maintenance decisions once a decision to assume ownership has been made, the data reject this, particularly for decisions to cure increases in accident rates (Figure 3a), and clean decreases in the quality given low cleaning costs (Figure 3b). The starting-quality effect for cleaning decisions (3b), however, was not strictly monotonic: while subjects were more likely to invest in the maintenance of pets of quality 90 or 100 than those of lower quality, over the range of lower qualities the effect of starting quality vanishes, with the lowest-valued pets being given the same average level of care as those of intermediate quality.

A final focus of interest is whether subjects' maintenance behaviors followed the predictions of **H4**, that the under-maintenance bias will be exacerbated given multiple, temporally proximate repairs, and mollified when initial damage occurs either prematurely or is highly delayed. To test the first of these propositions we computed the mean number of repairs undertaken for each kind of damage when they occurred in isolation versus jointly (that is, an illness and a dirt accident occurred on the same trial). Supporting **H4**, repair frequencies given joint events were sub-additive; the marginal likelihood of repairing an illness in isolation was .43 and that for a dirt accident in isolation was .27, but the mean number of repairs when both occurred at the same time was .56, implying that the joint incidence prompted an increased likelihood of repair, but often of just one rather than both maladies ($F(\text{illness} \times \text{accident interaction}; 1, 3128 \text{ df}) = 89.02, p < .001$).

To test the effect of timing of the first accident, in designing the experiment we randomly assigned one of the pets (never the first) acquired by subjects to a condition in which the first damage event (always a dirt accident) did not occur randomly but by a controlled schedule unknown to the respondent. There were three such timing levels: one where the first damage event occurred immediately on the 2nd ownership trial (slightly premature), one where it was delayed until the 10th trial (moderately delayed), and one where it was delayed until the 20th trial (highly delayed). Because of the loss of 18 pets that were replaced in the first period of ownership, 106 pets were available for analysis, 29 in the age 2 condition, 38 in the age 10, and 29 in the age 20. It should be emphasized that while this analysis provides a strong test of the effect of highly delayed damage, it provides a weaker test of the predicted effects of pre-maturity. The odds that a pet would

normally suffer some kind of damage on the first feasible trial was .226, a likelihood that may not have been sufficient to induce perceptions of rarity among subjects. In contrast, the odds of not seeing damage until the 20th trial was .002, a level much more likely to induce perceptions of rarity.

Perhaps reflecting these different strengths of manipulation, the data gave partial support for the theoretical hypotheses. Consistent with predictions, subjects were more likely to pay for the repair of damage when it was highly delayed given the base rate. The rate of cleaning when the first accident was delayed until the 20th period was .59 compared to a mean rate of first repairs of .40 ($\chi^2=2.99$; $p=.08$). In contrast, the data did not support a significant increase in repair rates when the first accident was slightly premature .41 v. .4 or moderately delayed .47 v. .4.

To provide a more systematic statistical investigation of the drivers of repair decisions we subjected subjects' decisions about whether to pay for each form of product repair conditional on a given trial to binary logit analyses that modeled these decisions as a function of eight sets of explanatory variables:

- 1) The original quality of the pet;
- 2) The marginal cost of cleaning;
- 3) Whether the pet became "dirty" on the trial (external wear);
- 4) Whether the pet became "sick" on the trial (internal wear);
- 5) The proportional age of the pet and its square;
- 6) Proportional experience with pet ownership, defined as the pet number divided by the number of pets ultimately owned by the respondent; and
- 7) Selected two-way interactions among variables (1) through (6).

Experience with ownership was included in the analysis to allow us to detect possible learning effects in the task. For example, one might conjecture that as experience with the task grew subjects may have developed an intuitive awareness of the optimality of cleaning in the low-cost condition but its sub-optimality in the high-cost condition. Such

an effect would be manifested in a significant positive two-way interaction between experience and sickness incidents on cure rates and a significant negative three-way interaction among experience, sick rates, and repair costs for dirt accidents. In this analysis we measure ownership experience using a proportional measure to control for individual differences in place-replacement frequencies.¹³

The results of the analysis are summarized in Table 1. Supporting the visual findings reported above, the model reveals significant negative interactions between wear incidents and proportional age, implying a decreased conditional likelihood of maintenance as the age of a pet increased. In addition, the data offer a mixed verdict for what subjects managed to learn about pet maintenance over the course of successive generations of ownership. Supporting learning is the finding of a significant positive interaction between ownership experience (proportional pet number) and illness incidents, implying that subjects became more likely to invest in increased risks of probabilistic failures with ownership experience. Likewise, we also observe this same positive interaction for responses to dirt incidents, but do not observe a significant negative conditioning three-way interaction with cleaning costs. This implies that while subjects in the low-cost condition were appropriately learning to invest more in cleaning with increased ownership, subjects in the high-cost condition were *inappropriately* learning the same thing. Finally, the analyses support the visual effect of starting quality on repair rates noted above (Figure 3): there is a significant positive interaction between incidents of illness and the linear trend in quality on cure rates (3a), but that between dirt incidents and linear quality on cleaning rates – while positive in sign—is insignificant (3b).

Replacement decisions

Under the optimal policy the best strategy for maintaining lower-quality pets is never to assume ownership at all. Specifically, if a pet is considered to of sufficiently low quality that one feels that investments in its maintenance are not worthwhile, or that one would be better off with a replacement, there is no normative advantage to delaying the replacement decision. We hypothesized (**H2**), however, that subjects would be reluctant to immediately replace pets immediately after they have been acquired, regardless of how low their quality might be.

In Figure 4 we plot the frequency with which pets acquired in the first 225 weeks of the task immediately replaced, never replaced, or held for an intermediate duration pets of varying qualities. In cases where the ownership duration was between 1 and 75 periods, we also report the median trial one which a voluntary replacement was made. On one hand, the figure offers evidence that subjects held at least a limited instinctual grasp of some aspects of the normative policy. As would be optimal, the most common replacement action for pets of the lowest quality was to replace them immediately, while that for pets of quality 100 was to retain them for their maximum life (75 periods). The larger sense of the data, however, is a strong rejection of the normative theory as a descriptive account of subjects' replacement decisions. While, indeed, 24% of all acquired pets of quality 30 were immediately replaced, 71% were not, and were held for a median duration of ownership of 19 periods. Even more disturbing was the displayed willingness to replace pets of superior quality. In the 66 instances where subjects were fortunate enough to draw pets of quality 100, 43% were voluntarily replaced before they had lived their full lives, with a median ownership duration of 40 periods. Respondents

undertook these acts even though the odds that the new pet would be of the same quality was quite small (.16).

What drove the replacements of high-quality pets? Within the proposed algebra it arises through a recursive process where norms of short ownership horizons discourage investments in maintenance, which, in turn, lead to degraded goods that enhance the appeal of replacements. Consistent with this account, in Figure 5 we plot the average degree to which the original quality of pets had degraded at the time of replacement by starting quality level. The figure reveals a remarkable empirical regularity: regardless of the initial quality of a pet, replacements arose when its quality was allowed to deteriorate to roughly 75% of its initial level (ranging from 71% to 78%). Hence, as pets aged subjects appeared to view them as if they were undergoing unavoidable wear-out, even though in this task a repair investment would restore them to a like-new condition.

Finally, the data offered no evidence that the general bias of replacing pets prematurely diminished as subjects became more familiar with the process of pet ownership in simulation. Among pets that were owned for at least 2 periods (i.e., eliminating immediate replacements), the data show no systematic lengthening of ownership over pets; the mean ownership durations over the first four such pets owned by respondents (excluding those held at the terminus of the simulation) were 32, 34, 33, and 30 periods, respectively. Because these means pool over pets of various qualities--some for which it was optimal to replace earlier—as a supplemental analysis we computed the percentage of pets of quality 70 or higher that were retained for a full 75 periods by ownership order. While the percentage long-duration ownerships increased over the first three pets owned (from 12% to 21%), there was no continuation to the fourth, where

experience (and, for some, an approaching game terminus) would most favor long ownership: in this case the percentage dropped to 16%.

Discussion

Knowing how much to invest in the upkeep of durable goods is not an easy task. While there is a large literature in operations management that describes how one might derive analytic solutions to some of these problems, in practice we are more likely to rely on our intuitions or advice from others, such as that recently offered by a consumer web columnist:

“While you can make decisions based on broad statistics such as the life expectancy of appliances and frequency of repair records, still, part of your decisions should be based on your personal preferences and “gut feeling”. Just because an appliance is nearing its average life doesn’t mean there isn’t a lot of spunk left in those coils and wires! But don’t get suck with a mean machine that eats paychecks for breakfast, either!” (Pat Veretto, About Frugal Living, frugalliving.about.com.

What constitutes this “gut feeling”, and what is the quality of the decisions it yields? The purpose of this research was to explore this issue by investigating the process by which consumers make dynamic decisions about the maintenance of a durable good. The research centered on a hypothesis that consumers choose among maintenance actions by an assessment process that indexes the perceived benefits and costs of different maintenance actions by two dynamic benchmarks: how goods typically wear over time and how expenditures for repair and replacement are typically paced. The process allows consumers to make maintenance decision that may often be not that different from those prescribed by optimal dynamic models, but without requiring the recursive reasoning and extended foresight required by such analyses.

This hypothesis was represented within a formal algebraic theory of dynamic product care. The purpose of this exercise was to show how a wide range of seemingly disparate consumer maintenance behaviors—including normative ones--can be explained within a common mathematical framework. To illustrate, at the outset we noted that while the stereotypic account of consumers is that we are lax in the care of many of our possessions, there are also instances where bias seems to go the other way—cases of seeming excessive attachment to older, deteriorated goods. The proposed framework provides a means for explaining both tendencies by proposing that the utility that is drawn from goods in various conditions is assessed relative to beliefs about the expected condition of a good given its age. In the early stages of ownership utility is monotonically declining as a good ages through normal wear-and-tear, but if it survives past the age of typical abandonment increasing age enhances utility. As a result, consumers would be predicted to lament the loss of the very old and very new, but be comparatively indifferent to the loss of those in between.

We began the process of assessing the empirical validity of the theory by reporting the results of an experiment that examined the ability of a sample of 124 subjects to learn to maintain a series of computer “pets” over a 300-period horizon. The goal of the simulation was to investigate whether the major predicted departures from optimality would be evident in a reasonably realistic decision environment—one where respondents were faced with many of the same challenges that mark real product ownership, and where incentives were offered for performance. For the most part the data supported the predicted departures from optimality: subjects under-invested in maintenance when it was optimal (H1), retained and invested in low-quality pets that

should have been immediately replaced (H2), displayed repair rates that gradually declined as a pet aged and were lower for pets of inferior quality (H3), were further deflated in the margin when multiple damage events occurred on the same trial (H4), but were more inclined to pay for the repair of damage that was highly delayed in its arrival (H4).

Data from other task settings

One obvious limitation of the empirical work described here is that it provides a portrait of decision achievement in only one task setting. For example, the pets owned by subjects had a deterministic lifespan (75 periods) and there was never a chance that it might fail completely due to poor maintenance. In the real world, of course, neither circumstance usually holds: the lifespan of a possession is something that is usually under our control, and, for some goods (e.g., cars and computers) there *is* an ongoing risk of complete failure.

To examine the degree to which the findings reported here would generalize to such a setting, 189 new subjects participated in a new version of the simulation that imposed no limit on ownership duration (respondents could own a single pet for up to 300 trials if desired), and, most significantly, the accident rate (p) in the task reported above was replaced by a *failure probability*.¹⁴ Specifically, all new and fully-repaired pets began life with a failure probability of zero, but this became increasingly positive when a pet suffered illnesses. All other features of the task were the same as the current: the pet suffered probabilistic decreases in quality that could be repaired by cleaning (though here there was a constant accident rate), and the quality of a pet replacement was probabilistic. The normative policy here was simply a more extreme version of the

current one: the rational player should begin the simulation by immediately replacing pets until one with a quality of 90 or 100 is drawn, then vigilantly maintain it for the full 300 periods of possible ownership (cleaning and repair costs were held constant).

The essential finding was that the biases of under-repair and over-replacement were not only replicated but amplified in this new task setting. Over the 300 trials respondents owned an average of 11.97 pets (compared to a normative mean for this task of 2.7), held them for a median lifespan of 25 periods, and with the majority (63%) of terminations occurring due to probabilistic failure of an ill pet. The number of owned pets ranged from as few as two to as many as 28, and life spans (again excluding immediate replacements) ranging from 2 to 198 periods. To illustrate this effect in its most dramatic form, in the simulation there were 265 cases where subjects were endowed with quality=100 pets—cases where the value of vigilant maintenance should have been most transparent to subject (recall that replacement was never mandatory). By age 30, however, less than half of these (128) were still actively owned, 22 having been voluntarily replaced, 95 suffering probabilistic failure, and 20 still living at the end of the simulation. Among these age-30 survivors—who were comparatively well-maintained--the average quality had deteriorated to a mean of 85.1, and had a median age of 29.

Alternative theoretical accounts

One limitation of the proposed algebraic theory is that it has a restricted contextual focus, and considers only a subset of the psychological factors that may influence maintenance decisions over time. A salient omission, for example, is learning. In the current formulation decision makers are assumed to hold beliefs about normal wear rates, normal maintenance requirements, and product –replacement intervals that are

invariant over a series of product-ownership cycles. In real settings, of course, these beliefs are likely to be ill-formed at the outset, evolve over time in light of decay rates that are actually observed, and thus be at least somewhat endogenous; maintenance policies are based on beliefs about product normal product wear rates, which are, in turn, formed by observing *actual* wear rates, which, in turn, are a function of past maintenance policies. Given these complexities the problem of how to best model these psychological dynamics is set aside for the moment, but its resolution is essential for any complete understanding of real-world product-care behavior.

To illustrate the challenges modeling learning potentially poses, in the current experiment we found evidence that respondent's attitudes toward the wisdom of investing in maintenance evolved over time, however it was not clearly in the direction of movement toward optimality in the task. Specifically, with increased experience of ownership subjects displayed an increased propensity to invest in the repair of the pets they owned (a move toward optimality on average), but they did not display evidence of learning about either the wisdom of retaining pets for longer periods of time before replacement or when it is optimal *not* to pay for repairs. Hence, it is possible that the increased tendency to invest in repairs with increased ownership experience did not accrue to learning but rather from other attitudinal dynamics, such as increased involvement with the simulation as it evolved (hence a desire to more fully utilize its available functions) .

Another caveat is that the theory provides only *an* explanation for the maintenance behavior observed in the experiment, and other processes may be at work. For example, one of the findings of the work that was supported for a prediction that

damage that occurs to goods that have survived in mint condition for long periods is more likely to be repaired than that whose arrival is more consistent with normal wear rates. The mechanism for this effect within the theory is that consumers have expectations for normal rates of spending for repairs, and funds that lie unspent for long periods eventually become reframed as if they were savings, reducing the psychic cost of repairs when they finally are needed (Axiom 2). One might argue, however, that the result may have accrued to—or at least been amplified by—a more mundane explanation. In this case may have taken the long delay in the need for repairs as evidence that the pet has a lower base accident rate than the norm, a belief that would rationally support more vigilant maintenance.

Conclusions

One often hears the accusation that as consumers we under-invest in maintenance, yielding a society where disposal and replacement have become the norm, repair and frugal care the exception. On the other hand, a case can be made that what might *seem* to under-maintenance is simply a rational response to the economic realities of modern markets; we replace rather than repair because new goods offer a higher benefit stream than repaired ones would, particularly given declining replacement costs.

Which one of these interpretations is right? The verdict based on the research reported here is seem mixed. On one hand, subjects in our simulation (as well as an extension) behaved in a way that would seem to support the popular accusation that consumers have a penchant for waste. Respondents paid for repairs less often than they should have, and replaced their computer pets too often. On the other hand, the real costs of these errors were comparatively small; while being demonstrably unaware of the

optimal policy that characterized the task, they made decisions that yielded an average performance score that was within 15% of the optimal expectation.

A natural question that arises in light of these findings is the degree to which they offer a model of the consumer behaviors we might expect to see in real-world settings. On the surface the answer would seem to be the affirmative. Real-world car ownership, for example, displays many of the same biases that arose in computerized pet ownership: car owners often fail to undertake small acts of maintenance that would reduce the long-term costs of ownership (such as routinely checking tire pressure), and those who can afford it probably replace their cars more often than they really have to (though not as often as the auto industry might like). On the other hand, whether the *efficiency* of real-world decisions resemble that which was observed here is less clear. The laboratory simulation offered an idealized environment where all the factors influencing product deterioration were known and there was frequent performance feedback—favorable factors that would not be present in real-world settings. On the other hand, subjects did not have access to the kind of expert advice on maintenance that is routinely available in the real-world (e.g., a pet owner's handbook)—a factor that may have limited performance relative to real-world contexts.

As a final comment, while the substantive focus of the current investigation on maintenance decisions is relatively new, the theoretical ideas it draws on should be seen as far more familiar. The idea that consumers may index their perceptions of the disutility of owning a deteriorated good by considering normal wear rates, for example, can be seen as simply extending the widely-documented tendency for assessments of utility to be reference dependent (e.g., Thaler 1980) to the dynamic case. Likewise, the global bias of

under-maintenance follows from some of the same processes that underlie a range of better-documented biases in inter-temporal choice, such as an aversion for temporally accelerating payments for durable goods (e.g., Prelec and Lowenstein, 1998). The work also illustrates, however, the theoretical and empirical challenges that arise when trying to extend these ideas to the study of decision making in complex multi-period settings. The task of deriving a theoretical framework that adequately integrates these ideas with notions of strategic planning is not an easy one, and the work reported here is advanced simply as a first step toward this goal.

References

- Arkes, Hal R. (1996), "The Psychology of Waste," *Journal of Behavioral Decision Making*, 9 (September), 213-224.
- Assaf, David, Shanthikumar, J. George (1987), "Optimal Group Maintenance Policies with Continuous Periodic Inspections", *Management Science*, 33 (November), 1440-1452.
- Bayus, Barry (1991), "The Consumer Durable Replacement Buyer", *Journal of Marketing*, 55(1) 42-52.
- Barlow, R.E., and L. Hunter (1960), "Optimal Preventive Maintenance Policies", *Operations Research*, 8 (January), 90-100.
- Chao-Ton Su (2000), "Multi-action maintenance subject to action-dependent risk and stochastic failure", *European Journal of Operational Research*, Vol. 125 (1), 133
- Cripps, John, and Meyer, Robert (1994), "Heuristics and biases in timing the replacement of durable products", *Journal of Consumer Research*
- Gilboa, Itzhak, and David Schmeidler(1995), "Case-Based Decision Theory", *Quarterly Journal of Economics*, November, 605-630.
- Gourville, John T and Dilip Soman (1998), "Payment Depreciation: The Behavioral Effects of Temporally Separating Payments from Consumption," *Journal of Consumer Research*, 25 (September), 160-174.
- Hayre, Lakhbir (1983), "A Note on Optimal Maintenance Policies for Deciding Whether to Repair or Replace", *European Journal of Operational Research* 12 (February), 171-176.
- Loewenstein, George, and Prelec, Drazen (1992), "Anomalies in Intertemporal Choice:

- Evidence and an Interpretation, *The Quarterly Journal of Economics*, 107 (2) 573-598
- Okada, Erica (2001), "Trade-ins, mental accounting, and product replacement decisions" *Journal of Consumer Research*, 27 (4), 433-458.
- Pickering, J.F. (1984), "Purchase Expectations and the Demand for Consumer Durables", *Journal of Economic Psychology*, Vol 5, 4, 341-353
- Pham, Hoang, and Wang, Hongzhou (1996), "Imperfect Maintenance", *European Journal of Operational Research*, 94 (November), 425-439
- Prelec, Drazen and George Loewenstein (1998), "The Red and the Black: Mental Accounting of Savings and Debt," *Marketing Science*, 17 (Winter), 4-28.
- Thaler, Richard (1980), "Toward a Positive Theory of Consumer Choice", *Journal of Economic Behavior and Organization*, 1, 39-60.
- Usher, John S., Kamal, Ahmed, and Wasim Hashmni Syed (1998), "Cost optimal Preventive Maintenance and Replacement Scheduling", *IIE Transactions*, 30 (December), 1121-1128.
- Winer, Russell (1985), "A Price Vector Model of Demand for Consumer Durables: Preliminary Developments", *Marketing Science*, 4(1) 74-91
- Zhang, Fan, and Jardine, Andrew (1998), "Optimal maintenance models with minimal repair, periodic overhaul and complete renewal", *IIE Transactions*; Dec 1998

Footnotes

¹Source: U.S. Census Bureau, *1997 Economic Census, NAICS 811 (Service Sector/Other Services/Repair and Maintenance)*

²As evidence of the pervasiveness of such ideas, a 2001 issue of *Consumer Reports* (October) published a large number set of age-dependent guidelines about when one should replace or repair a wide variety of durable goods. The magazine recommended, for example, that damaged computers should be replaced if they are over two years old, but damaged notebooks of the same age should be repaired.

³This symmetry would hold, of course, if respondents took the large previous repair expenditure as evidence of an enhanced attachment to the old good.

⁴This assumption is introduced to preserve linearity over the utility space, a feature that eases computation of the normative policy.

⁵In their October 2001 issue on repair advice *Consumer Reports* suggested a general “50% rule” for deciding whether to repair or replace an appliance: if the cost of repair is more than 50% the cost of replacement, replace the good. The origin and normative status—if there is any—of this heuristic, however, is unknown.

⁶Formally, $u'_o, u''_o > 0, u'_r > 0; u'_r(-x) > u'_r(x)$.

⁷To see this, note that as $A \rightarrow \infty$ the loss of a good that has been maintained in mint condition (i.e., $s_t = s^0$) converges to $-k_2 u_r(s^0)$. This value, in turn, would be similar to that suffered if the good was new ($A=0$), $-k_1 u_o(s^0)$. If $k_2 > k_1$, the loss of the older good would be perceived as greater.

⁸One might recognize this effect as akin to a dieter who feels justified in rewarding himself with a second helping of a favorite food if he had been “extra good” on earlier days.

⁹The task was inspired by the popular Japanese electronic-pets that first appeared in the late 1990's, such as Giga Pets and Tamagotchis. In these games players try to keep key-chain sized computerized pets alive for as long as possible through a regimen of care and feeding.

¹⁰ The actual generating distribution was a hybrid uniform that made pets of the two extreme qualities (10 and 100) slightly more likely than those of intermediate quality, with respective likelihoods of .165 (simple uniform would be .12).

¹¹This number reflects the four pets that would be owned for a full 75 periods plus the number of immediate replacements expected under the optimal replacement rule.

¹²To verify that the normative predictions held for shorter anticipated life-spans, we repeated the numerical simulations reported earlier for cases of life-spans of 20, 30, and 40 periods. As before, the normative policy for repairing increases in accident rates was that was optimal to do so until the 5th trial from the end, and optimal to pay for cleaning for up to the first half of trials.

¹³ To illustrate the problem using raw pet number as a measure of experience, for subjects who replace pets frequently the 4th pet owned may have arisen quite early in his or her experience hence be reflective of limited experience, while for those who replaced rarely the 4th may have come late in the simulation after considerable experience. In addition, we use pet ownership rather than simple time as a measure of experience to reflect the fact that within the course of ownership of a pet maintenance decisions tend to become routinized, and offer limited opportunities for learning.

¹⁴A more complete summary of this experiment and the data are available from the author on request.

Table One

Logistic Regression Analysis of Trial-By-Trial Repair Decisions

Parameter	Clean Decisions		Cure Decisions	
	Estimate	P< χ^2	Estimate	P< χ^2
Intercept	-4.897	<.0001	-5.087	<.0001
Starting Quality(Q)	.004	.1813	-.001	.8554
Sick (S)	2.118	<.0001	3.463	<.0001
Dirty (D)	3.273	<.0001	1.961	<.0001
Cleaning Cost (CC)	-.029	<.0001	-	-
Relative Age (RA)	.013	.1099	.007	.3048
Pet Experience (PE)	-.129	.5609	-.163	.4279
Q*D	.005	.1553	-	-
Q*S	-	-	.017	<.0001
S*D	2.264	<.0001	2.089	<.0001
D*CC	.006	.2900	-	-
D*RA	-.031	.0019	-	-
D*PE	.876	.0025	-	-
D*PE*CC	-.003	.7088	-	-
S*RA	-	-	-.057	<.0001
S*PE	-	-	.504	.0369
Model				
Likelihood Ratio	4293	<.0001	6394	<.0001
P ²	.340		.439	
N	3762		3762	

1c: Acquisition Quality Thresholds Given Optimal Maintenance

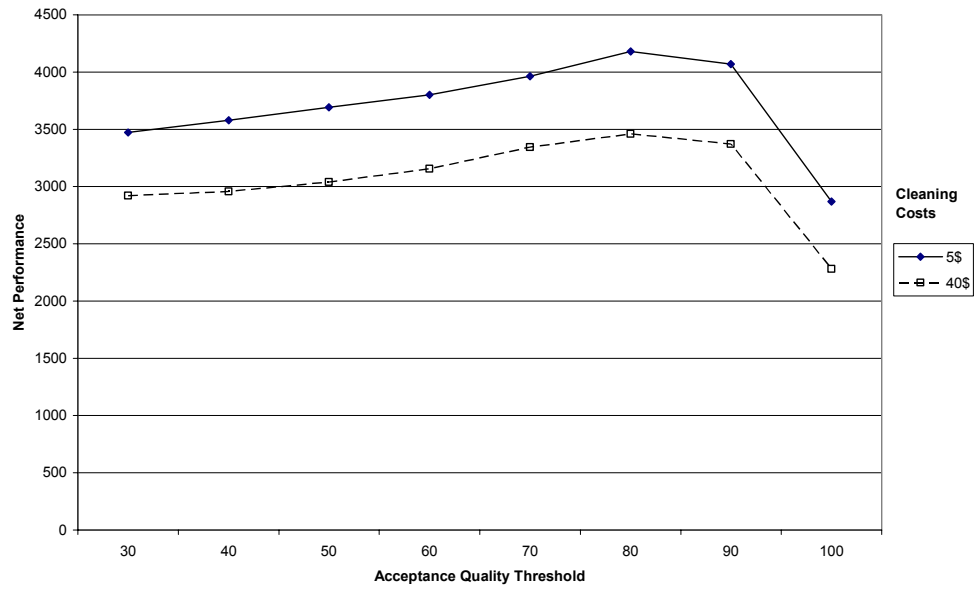



Figure 2: Empirical Repair Rates over Time

2a. Repair rates by Type and Relative Age



2c. Conditional Cleaning Probability by Relative Age and Marginal Cost





Figure 3 The Effect of Starting Quality on Repair Rates

3a. Cure Rate by Starting Quality



3b. Conditional Cleaning Probability by Cleaning Cost and Starting Quality




Figure 4


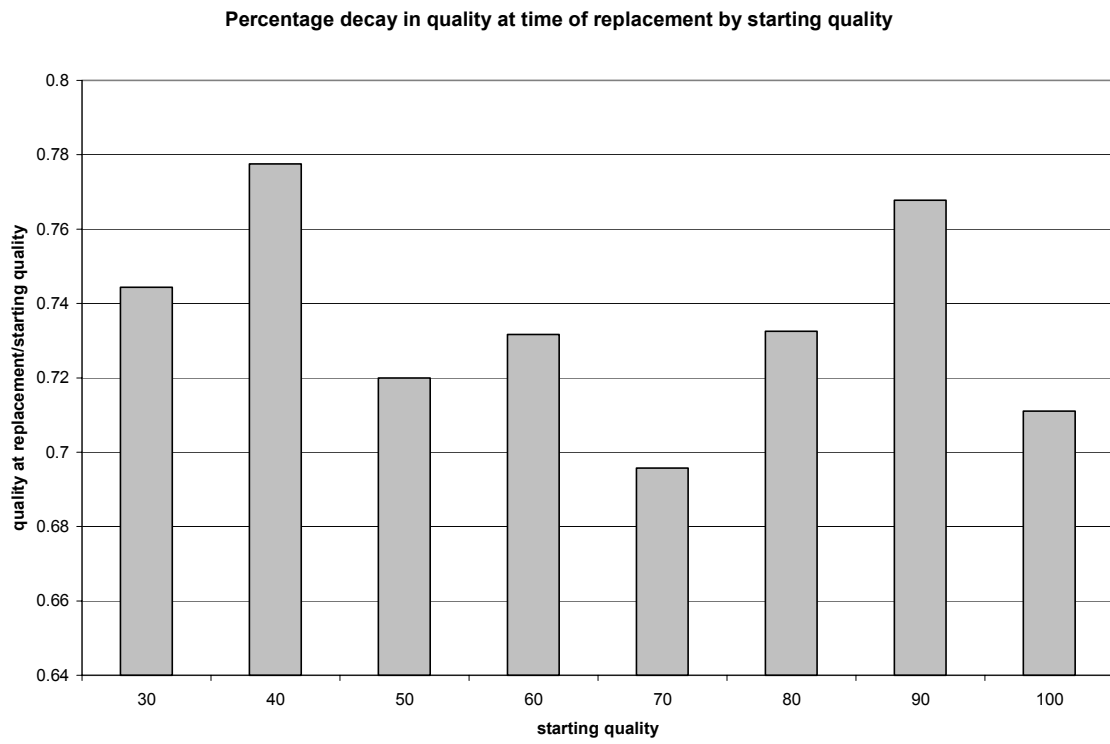


Figure 5



Appendix 1

The Optimal Maintenance Policy

We consider the general form of the optimal control policy for a finite-horizon optimal maintenance policy that has the following structure. At time $t=0$ a decision maker (*DM*) pays a purchase fee c^B to acquire a good whose quality s_0 is a random draw from a known distribution $G(S)$. On each subsequent occasion there is a constant known probability p that will cause the utility of the good to diminish by an amount δ that is a random draw from a distribution $F(\delta)$. The expected utility of the good thus evolves as $s_t = s_{t-1} - pE(\delta)$, where by assumption $\lim_{t \rightarrow E} (s_t) > 0$. On each occasion *DM* may pay a repair fee $c^R = r(s_0 - s_t)$ to restore the good to its initial quality level, purchase a new good for the price c^B , or consume the good in its current state. Each good has finite life expectancy of E periods, and *DM*'s goal is to maximize net utility over an undiscounted T -period time horizon, $T > E$.

We show that the optimal maintenance policy is a control-limit rule of the form,

1. *On the initial trial, $t=0$, observe the quality of the initially-drawn good, s_0 . If s_0 is less than a threshold acquisition quality Q^* , immediately pay c^B for a replacement, and continue doing so until a good of quality Q^* or higher is obtained;*
2. *Once an acceptable good is acquired, proceed with its ownership for the duration of the life expectancy E , paying for the repair of all damages until a termination age of K^D and none thereafter.*

Repair decisions

Consider a *DM* who entertains two policies at time t for repairing an expected series of decreases in quality over time, $E(s_{t+N}) = s_0 - NpE(\delta)$: one that repairs each decrease δ as it arises (vigilant repair) versus one that repairs the good periodically, say just in period $t+N$. Because the expected cost of implementing both policies is identical (

$r(s_0 - E(s_{t+N}))$, it should be transparent that a policy of immediate will always dominates that for periodic or delayed repair by virtue of offering a superior benefit stream (Ns_0 versus $(Ns_0 - \sum_{i=1}^{N-1} ipE(\delta))$). The existence of a terminal care period follows immediately from this result. If a decision maker concludes that it is *not* optimal to pay c^R to restore a pet at t , it cannot be optimal to do so at any later period $t+i$. For a given marginal repair cost, accident rate, accident magnitude, and expected ownership duration, the optimal repair policy will thus be that of a finite duration rule: the optimal decision maker will undertake all repairs as needed until a threshold time K^D , and none thereafter.

We might add that because the decay parameters p and $F(\delta)$ are independent of the initial (or restored) quality of a good s_0 , the optimal control parameter K^D will also be independent of s_0 . By invoking the optimal replacement policy (below) it follows that once one concludes that a good is worth owning, one should adopt the same posture toward repairing damage regardless of its initial quality.

The optimal replacement policy

Whereas the repair policy holds for any undiscounted linear evolution of repair costs and benefits, the replacement policy is restricted to the case where p and $F(\delta)$ are such to insure that $s_t > 0$ for all $t < E$; i.e., the good never completely fails before it reaches the end of its expected life expectancy. With this in place the replacement policy follows using much the same line of argument above. A replacement offers only one advantage over a repaired incumbent: it provides the consumer with a chance to initiate a new stream of consumption utility with a good that has a higher base quality level than the current one, by paying for a new draw from the distribution of new-product qualities, $G(s)$.

Consider a consumer who owns a good of age $0 < t < E$ that has an initial (restored) quality s_0 , and who concludes that it is worthwhile to pay $c^B > c^R$ to replace a good at that time. By the assumed stationarity of $G(s)$ the single-period benefits of a replacement over a repair at t will be the same as they were at $t-1$. The optimal age of replacement will thus be $t=0$ by backward induction. By translation there will thus exist a critical control parameter Q^* that defines the quality at which a rational decision maker would be indifferent between accepting a good at the time of purchase and consuming it for E periods versus paying for another immediate replacement..

General Comments

It is important to stress that this simple maintenance policy accrues to a number of problem simplifications that would not arise in more general maintenance problems, such as constant linear costs. To illustrate, in the case where a good may deteriorate to an unusable condition before it reaches the maximum limit of its lifespan (E) and accident rates increase with chronological age, there will also exist an optimal replacement age less than E —a common feature of many optimal maintenance models.

Appendix 2

Relation of the Choice Axiom to the Optimal Maintenance Policy

The normative decision problem faced by the consumer is to find the set of control parameters Q^* and K^D in the optimal policy that, when applied to make a sequence of maintenance actions k over a time horizon, satisfy the recursive optimization criterion (Bellman's equation):

$$V_t^*(s_t^i) = \max_{k_t} \{u(s_t^i | k_t) - c_t^k + \beta \sum_p V_{t+1}^*(s_{t+1}^p) p(s_{t+1}^p | k_t)\} \quad (A1)$$

where s_t is the state or condition of the good at time t , c_t^k is the cost of maintenance action k at t , $p(s_{t+1}^p | k)$ is the probability of observing the good in state p at time $t+1$ given that action k is take at t , and β is a temporal discount rate ($\beta=1$ in the actual reward structure). The hypothesized choice axiom can be derived as the special case of (A1) that arises if consumers are assumed to make two simplifying assumptions about how the performance of the good evolves over future time periods:

- 1) Once a decision k is undertaken at time t no other action will be taken for M^k periods—i.e., the action has a finite expected duration; and
- 2) The future performance of a good undergoes the deterministic aging process $s_{t+1}^i = H(A)(s_t^i | k_t)$ as in expression (3).

The key implication of these assumptions is that they allow (A1) to be re-written in the form

$$V_t(s_t) = \max_{k_t} \{u(s_t | k_t) - c_t^k + \sum_{q=1}^{M^k} \beta q [(s_t | k_t) H(A)] + [\beta(M^k + 1)] V_{M^k+1}\} \quad (A2)$$

where the term $\sum_{q=1}^{M^k} \beta^q [(s_t | k_t) H(A)]$ is the discounted future benefit of pursuing action k

with expected duration M^k . If we further assume that consumers utilize a hyperbolic subjective discount function $f(\beta|t)$ that gives little consideration to future decisions beyond the duration of repair and maintenance actions (i.e., $f(\beta|t)=0$ for $t > M^R, M^B$),

(A1) further simplifies to

$$\text{MAX} \begin{cases} V_t^D = u(s_t) + f(\beta) \text{MAX}(V_{t+1}^D, V_{t+1}^R, V_{t+1}^B) \\ V_t^D = u(s_0) - c_t^R + \sum_{q=1}^{M^R} f(\beta|q) u(s_0) H(A) \\ V_t^D = u(s_0^*) - c_t^B + \sum_{q=0}^{M^B} f(\beta|q) u(s_0^*) H(A) \end{cases} \quad (\text{A3})$$

If we let $k_D = f(\beta) \text{MAX}(V_{t+1}^D, V_{t+1}^R, V_{t+1}^B)$, $c_t = \psi(c_t)$ be the subjective cost of action k as in equation (3), and $v_k(\cdot)$ be an empirical scaling function, then (A3) can be written, for a finite planning horizon T ,

$$\text{MAX} \begin{cases} V_t^D = k_D + k_{1D} u(s_t) \\ V_t^D = k_R + k_{1R} v_R \left(\sum_{q=0}^{\text{Min}(T-t, M^R)} f(\beta|q) u(s_0) H(A) - \phi(c_t^R) \right) \\ V_t^D = k_B + k_{1B} v_B \left(\sum_{q=0}^{\text{Min}(T-t, M^B)} f(\beta|q) u(s_0^*) H(A) - \phi(c_t^B) \right) \end{cases} \quad (\text{A4})$$

where (A4) is now the hypothesized choice axiom.

Note that (A4) implies that the scaling constant in assessments of decisions to defer maintenance k_D has a specific meaning in the context of a normative decision model: it is the expected value of facing a repair or replacement decision tomorrow rather than today; i.e., the psychological benefit of choice deferral.

Appendix 3 :Screen Shots of the Pet Ownership Simulation

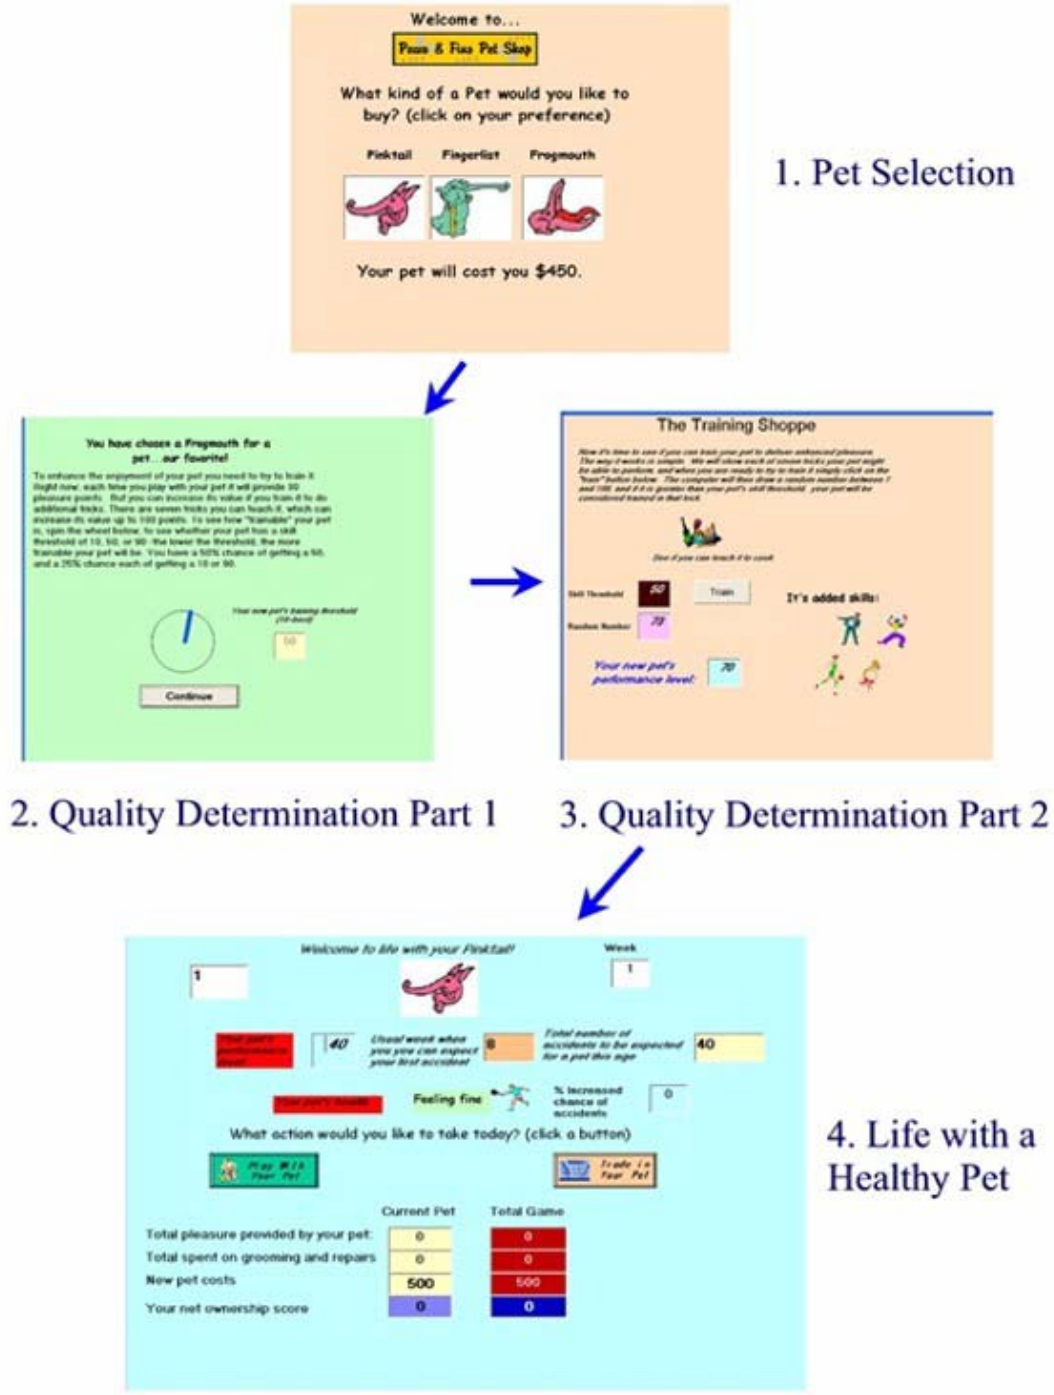


Figure 1: Screen Shots of Pet Acquisition and Play

Welcome to life with your Frogmouth

Week 11 23

79 Number of accidents you have had so far

2 Total number of accidents to be expected for a pet this age

1

26 % increased chance of accidents

Feeling Ill

What action would you like to take today? (click a button)

Play With Your Pet Restore its Performance Cure its Illness Trade in Your Pet

	Current Pet	Total Game
Total pleasure provided by your pet:	791	1261
Total spent on grooming and repairs	320	370
New pet costs	450	960
Your net ownership score	21	-69

An Accident that Decreases the Pet's Pleasure Level

A Dirt and Sick Pet

A Sickness that Increases the Probability of an Accident

Caps, your pet has had an accident... it's not life-threatening, but it has become a bit dirty!

Pleasure has decreased by **10** units

Continue

Your pet is feeling under the weather

The percent chance of it having an accident has increased by **26** %

Continue

Welcome to the Grooming Salon

Your pet's current pleasure level: **72**

The cost of restoring its pleasure to its original "like-new" level is **320** Dollars

(the cleaning fee is **5** per lost pleasure unit)

Do you want us to groom your pet?

Yes, Groom it! No, send me home

Repair: A trip to the grooming salon

Welcome to Countyland Hospital

Your pet has an illness that has increased the odd that it will have an accident in a given week by **26** Percent

The cost of restoring you pet to full health is **50** Dollars

Do you want us to restore your pet's health?

Yes No

Repair: A trip to the Hospital

The Process of Accidents and Repairs