

Flexible-Duration Extended Warranties with Dynamic Reliability Learning

Guillermo Gallego

Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027, USA
gmg2@columbia.edu

Ruxian Wang

The Johns Hopkins Carey Business School, Baltimore, Maryland 21202, USA, ruxian.wang@jhu.edu

Julie Ward

Hewlett-Packard Company, Palo Alto, California 94304, USA, jward@hp.com

Ming Hu

Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada, ming.hu@rotman.utoronto.ca

Jose Luis Beltran

Hewlett-Packard Company, Palo Alto, California 94304, USA, jose-luis.beltran@hp.com

Frequent technological innovations and price declines adversely affect sales of extended warranties (EWs) as product replacement upon failure becomes an increasingly attractive alternative. To increase sales and profitability, we propose offering flexible-duration EWs. These warranties can appeal to customers who are uncertain about how long they will keep the product as well as to customers who are uncertain about the product's reliability. Flexibility may be added to existing services in the form of monthly billing with month-by-month commitments or by making existing warranties easier to cancel with pro-rated refunds. This paper studies flexible warranties from perspectives of both customers and the provider under customers' reliability learning. We present a model of customers' optimal coverage decisions and show that customers' optimal coverage policy has a threshold structure under some mild conditions. We further show that flexible warranties can result in higher profits and higher attach rates in a homogeneous market as well as in a heterogeneous market with multiple segments differing in various dimensions.

Key words: post-sales service; flexible warranty; reliability learning; heterogeneous market; threshold policy

History: Received: April 2011; Accepted: August 2013 by Michael Pinedo, after 3 revisions.

1. Introduction and Related Literature

As manufacturers face decreasing profit margins on hardware products, post-sales services like extended warranties (EWs) are increasingly important to their profitability. In addition to providing higher margins than hardware, services help to extend the useful life of products, generate a profitable revenue stream of consumables and accessories over the product's lifetime, and provide an opportunity to improve customer loyalty and brand image.

Attach rate refers to how many complementary goods and services (e.g., EWs) are attached and sold together with each primary hardware product. For electronic products, the attach rate of EWs at the market level is typically a single-digit percentage. Gallego et al. (2009) describe several challenges that manufac-

urers face in improving the attach rate of EWs for their base products. Among these is a perception for some customers that EWs are not very attractive although retailers are pushing hard to sell them (also see a survey on ConsumerReports.org 2009). This perception may be partly due to the fact that most warranties are offered at a uniform price regardless of how products are used. Some customers, because of their usage rates, operating environments, or other factors, may be susceptible to more product failures and thus be more expensive to support than others.

Another challenge to sell EWs is the combination of falling hardware prices and rapid technology improvements experienced in some industries. A good example is the personal computer (PC) industry. As prices decline, customers may find that replacing a failed product compares favorably to buying an EW

or paying out of pocket for repairs. Typically, the decision tilts more toward replacement when the replacement product has new features and is offered at an attractive price.

To address these challenges, service providers may want to consider introducing flexible-duration EWs that can appeal to a broader range of customers. Flexibility may be added to existing services in several different ways. For example, manufacturers or third parties could offer an EW with coverage commitments and payments made on a periodic basis, e.g., monthly or quarterly. Unlike a traditional, fixed-duration EW, which requires customers to commit to and pay up front for one or more full years of coverage, a flexible-duration EW allows customers to choose the duration of coverage with finer granularity and pay on a periodic basis for the coverage. Without creating much confusion, we call the new warranty a flexible EW from now on.

Flexible EWs are available on the market today. For example, Sprint offers an Equipment Service & Repair Program (ESRP) for smart phones at \$7.00 per month. ESRP can be added to devices within 30 days of activation or upgrade and can be canceled anytime at no cost but without possibility of reactivation in the future. AT&T offers a similar program at \$6.99 per month. Monthly WarrantyTM, a third-party service provider, offers EWs with monthly payments and with the flexibility to control their coverage duration for TVs, appliances, electronics, cameras, and PCs.

A flexible EW may appeal to customers for a variety of reasons. Periodic billing can attract customers with cash-flow constraints or high time-discounting rates, who would prefer low monthly or quarterly payments to a large lump-sum payment. There is evidence from the monthly prices offered by providers that customers are even willing to pay a premium for the combination of flexibility and monthly billing. Moreover, customers who buy low-priced products may have severe budget constraints and may opt to allocate all or nearly all of their budget at the time of purchase to the product itself. A flexible EW with deferred monthly payments does not compete directly with the product budget and may be more palatable than a warranty with a large up-front payment. Beyond its cash-flow advantages, a flexible EW allows customers to hedge against several kinds of uncertainty. Those who are uncertain about their future financial states will value the flexibility to terminate coverage at any time. Customers may also be uncertain of how much they will like the product. And in industries with rapid technological innovations, such as consumer electronics, customers may not know how soon they will wish to upgrade to a newer product with more features. Customers might also be unsure of the product reliability and therefore

the necessity of coverage; a flexible EW allows the customer to learn about the product's reliability over time without a costly long-term commitment.

This study focuses on two of the factors that make flexible EWs attractive: the customers' uncertainty in the timing of replacement and their ability to learn about product reliability. We discuss flexibility in the form of a monthly warranty, with monthly payment and month-by-month coverage commitments, although the results apply to any choice of period length in commitment and payment schedule and indeed to cancelable warranties that are paid with an up-front lump-sum payment and provided with prorated refunds at the time of cancellation. We illustrate that flexible EWs can result in higher profits and higher attach rates in a homogeneous market as well as in a heterogeneous market with multiple segments differing in various dimensions. In particular, we find that flexible EWs can be more profitable than traditional EWs in markets where some customers have low initial estimates of the product failure probability but adjust their estimates over time through their experience with the product.

In the warranty literature, research has largely focused on base warranties, although interest in EWs is now burgeoning as post-sales services are being recognized as important to manufacturers' revenue, margin, and customer loyalty (see, e.g., Cohen et al. 2006). Among studies on EWs, some address how EWs provide a signaling function for product quality, whereas others illustrate how heterogeneity among customers can enable segmentation of the EW market. For example, Padmanabhan and Rao (1993) consider pricing strategies in the presence of heterogeneous risk preferences and customer moral hazard. Lutz and Padmanabhan (1994) consider income variation among customers, whereas Lutz and Padmanabhan (1998) examine how customers' differing utility of a functional product enables market segmentation. Other studies discuss customer usage heterogeneity in the context of warranty pricing, including Hollis (1999), Moskowitz and Chun (1994), and Padmanabhan (1995).

The EW literature contains very little prior research on EWs with flexible duration. Jack and Murthy (2007) study the pricing of flexible EWs in a Stackelberg game setting, with the manufacturer as the leader and a customer as the follower. The manufacturer sets the price of the EW per unit time and the customer chooses the start time for the EW and replacement time for the product. Wang et al. (2010) present a model of flexible EWs in which customers dynamically choose how to respond to a failure (i.e., with a repair or a replacement) to maximize the total discounted utility derived from the products. The fact that customers can choose between repair and

replacement upon failure is one of the factors that impede fixed-term EW sales for inexpensive products, particularly in markets where technology is improving rapidly, and is part of what makes flexibility in warranty duration more attractive. Hartman and Laksana (2009) consider the design and pricing of a number of renewable EW contracts including unrestricted and restricted warranties. In their study, a customer is assumed to have perfect information about the product reliability, and her optimal strategy can be found by solving dynamic programs. The optimal menus of contracts are chosen to maximize the warranty provider's profits for homogeneous and heterogeneous markets.

Our study explores advantages of flexible duration in support services. Unlike the previous work related to flexible EWs, it models customers' dynamic learning of the product reliability over time, based on failure observations. In this setting, a monthly warranty premium can be thought of as an option value embedded in the flexible EW that enables customers to learn the reliability of the product over time without committing to buy a costly long-term warranty. The second contribution beyond the prior art is the incorporation of uncertainty into the timing of customers' product replacement. Replacement timing may be exogenously determined when market developments such as technology breakthroughs or dramatic price drops primarily drive product replacement decisions. The assumption of uncertain replacement timing is consistent with the uncertain timing of these events. In addition, this study analytically demonstrates the profitability of the flexible EW compared to the traditional EW in homogeneous and heterogeneous markets.

2. The Customer's Problem

In this section, we study a flexible EW with month-by-month commitments and a monthly premium m . The warranty provider (he) imposes two restrictions on a flexible EW purchase: (i) if a customer (she) purchases coverage, it must be started before the product reaches a pre-specified age; and (ii) the flexible EW cannot be resumed once it is discontinued. These restrictions are consistent with the practice of most warranty providers, who require extended coverage to begin when the product is still under the base warranty and require renewals to occur while the previously purchased flexible EW is in effect. Such restrictions are appropriate for a provider who finds it technically infeasible or cost prohibitive to verify that a customer's product is functional before allowing extended coverage. To simplify the exposition of the study, we further assume that the flexible EW must begin at the beginning of the planning horizon.

This assumption is not restrictive, as one can simply assume that the customer's planning horizon in this model begins at the end of the base warranty.

The customer may be uncertain about how long she will own the base product. This uncertainty may come from several sources. She may be unsure how much she will like the product or how soon a better product will be available on the market that induces her to upgrade. The randomness in her possession before a replacement depends also on the product reliability, as customers may be more likely to replace a failed product than a functional one. We model uncertainty in the length of the customer's possession as follows. Let $q_i(t)$ be the month t termination probability, where $i \in \{0,1\}$ denotes the state with $i = 0$ representing a failure and $i = 1$ representing a functional product. More specifically, $q_0(t)$ (respectively, $q_1(t)$) is the probability that the customer will abandon the product at the end of month t , given that she has not replaced it prior to that month and that the product fails (respectively, does not fail) in that month. Notice that this formulation essentially allows a random horizon that is bounded above by $T + 1$. Without loss of generality, T can be assumed to exceed the time that the customer abandons the product with probability one. A special case is the fixed time horizon model as $q_i(t) = 0$ for $t = 1, 2, \dots, T$, $i = 0, 1$, and $q_i(T + 1) = 1$ for $i = 0, 1$.

We assume that at most one failure can occur in each month. Although the results extend to the case of multiple failures per month, we make the single failure assumption to simplify exposition. Let I_t be a failure indicator random variable, taking value 1 if there is a failure in month t and taking value 0 otherwise. For notational convenience, let $\mathcal{I}_t = \{I_1, I_2, \dots, I_t\}$ denote the failure history up to month t . Let p_1 be the customer's initial belief of the failure probability in month 1. In the remainder of this study, failure probability refers to the probability that a product fails in a month unless otherwise stated. In each month, the customer updates her estimate of the product failure probability based on her initial belief and the product's prior failure history, following an updating scheme. Denote p_t as the estimated failure probability in month t . Let $p_{t+1}(\mathcal{I}_t)$ represent the estimated failure probability in month $t + 1$ for failure history \mathcal{I}_t , where $p_{t+1}(\cdot)$ is an updating scheme in month $t + 1$. We make the following assumptions about the updating scheme $p_{t+1}(\mathcal{I}_t)$.

ASSUMPTION 1. *The failure probability updating process is Markovian.*

More precisely, the failure probability update depends only on the present state, i.e., $p_{t+1}(\mathcal{I}_t) = H_{t+1}(p_t, I_t)$, where $H_{t+1}(\cdot, \cdot)$ is an updating

function. We define $p_{t+1}^+ = H_{t+1}(p_t, I_t = 1)$ and $p_{t+1}^- = H_{t+1}(p_t, I_t = 0)$. The quantity p_{t+1}^+ (respectively, p_{t+1}^-) represents the failure probability estimate at the beginning of month $t + 1$ when the month t failure probability estimate was p_t and a failure occurred (respectively, did not occur) in month t .

ASSUMPTION 2. *The failure probability updating method satisfies*

- (a) $H_{t+1}(p_t, 1) \geq H_{t+1}(p_t, 0)$ for any t and $0 \leq p_t \leq 1$;
- (b) $H_{t+1}(p_t, I_t)$ is increasing in p_t for all t and $I_t \in \{0, 1\}$.

Throughout this study we use “increasing” to signify “non-decreasing,” and “decreasing” to signify “non-increasing.” Assumption 2(a) is equivalent to $p_{t+1}^+ \geq p_{t+1}^-$, and it says that the estimated failure probability is higher for the next month if a failure occurs in the current month. Assumption 2(b) is equivalent to the condition $p_{t+1}^+ \geq p_{t+1}'$ and $p_{t+1}^- \geq p_{t+1}'$ for any $p_t \geq p_t'$, meaning that the estimated failure probability in a given month is increasing in the previous month’s failure probability estimate. Note that the updating scheme $H_t(\cdot, \cdot)$ need not assume stationarity and $H_t(\cdot, \cdot)$ may be different from $H_s(\cdot, \cdot)$ for $t \neq s$.

Customer risk aversion is typically assumed in determining premium insured for large losses, where customers’ utility functions tend to manifest concave curvature. However, the magnitude of losses from failures of many consumer products is relatively small, not enough to render concave curvature in customers’ utility function. In addition, warranties are different from other insurance in that the provider does not need to charge a risk premium to make a profit. This is particularly true when the warranty provider has economies of scale in repairing products. For these reasons, we adopt an environment where the customer is risk neutral and focus on other reasons for buying and selling warranties, such as information asymmetry in product failure probabilities and scale economies in providing repair services.

We consider a customer who has just purchased a new product and would like to minimize her total expected cost of supporting the product over a random planning horizon bounded by $T + 1$. She must choose whether to begin a warranty coverage in the first month and whether to continue the coverage for her product in each month thereafter. When the coverage is discontinued or the product is abandoned, there are no further decisions to be made for the product in the future. Thus, the model focuses on the customer’s decisions during the lifetime of a single product. The customer decides whether to continue warranty coverage based on her estimate of the product failure probability. She will continue coverage if

the total expected cost is not greater than that of canceling the flexible EW, which is often referred to as the participation constraints in the literature of principal–agent problems. For mathematical convenience and analytical tractability, we assume that repairs occur immediately and the product is returned to the customer by the end of the same month. This assumption may be more consistent with the actual repair time than the typical instantaneous repair assumption made in continuous time models. The repair cost of an uncovered failure in month t is a random variable C_t to the customer. We assume that its mean $c = EC_t$ is independent of t and is known to the customer and the warranty provider. Under Assumption 1, the total expected cost from the current month to the end of the horizon depends on the current estimated failure probability and is independent of the failure process realization before this month. Let $R_t(p_t)$ be the customer’s minimum total expected cost over the remaining horizon $\{t, t + 1, \dots, T + 1\}$, given that her estimated failure probability at the beginning of month t is p_t and that she does not buy coverage in this month.

Then, the customer’s total expected cost of terminating coverage at state (t, p_t) for any $t = 1, 2, \dots, T$ is

$$R_t(p_t) = p_t(1 - q_0(t))(c + R_{t+1}(p_{t+1}^+)) + (1 - p_t)(1 - q_1(t))R_{t+1}(p_{t+1}^-). \quad (1)$$

We assume that customers do not need to repair failed products after the planning horizon, so the boundary conditions are $R_{T+1}(\cdot) = 0$. Under some mild conditions, $R_t(p_t)$ in Equation (1) can be simplified.

LEMMA 1. *If $q_0(t) = q_1(t)$, denoted by $q(t)$, for $t = 1, 2, \dots, T$, and the updating scheme is a martingale with respect to the failure process, i.e., $E[p_{t+1} | \mathcal{I}_t] = p_t$, then $R_t(p_t) = p_t c \sum_{i=t}^T \prod_{j=t}^i (1 - q(j))$.*

All the proofs in this study are relegated to the online Appendix. Common updating schemes such as the Beta updating scheme (see, e.g., Stigler 1986) and the exponential smoothing mechanism (see, e.g., Brown and Meyer 1961) result in martingales that satisfy Assumptions 1 and 2. In the special case where the updating scheme is a martingale and $q_0(t) = q_1(t) = 0$ for $t = 1, 2, \dots, T$, $R_t(p_t) = (T + 1 - t)p_t c$.

Let $B_t(p_t)$ be the customer’s minimum total expected cost over $\{t, t + 1, \dots, T + 1\}$, given that her estimated failure probability at the beginning of month t is p_t and that she buys coverage for month t . Let $W_t(p_t)$ be the customer’s minimum total expected cost over $\{t, t + 1, \dots, T + 1\}$, given that she had warranty coverage in month $t - 1$ and that her estimated failure probability at the beginning of month t is p_t .

The customer’s minimum total expected cost of continuing warranty coverage is given by

$$B_t(p_t) = m + p_t(1 - q_0(t))W_{t+1}(p_{t+1}^+) + (1 - p_t)(1 - q_1(t))W_{t+1}(p_{t+1}^-). \quad (2)$$

The optimal decision is the action with the smaller expected cost, i.e.,

$$W_t(p_t) = \min\{B_t(p_t), R_t(p_t)\}, \quad (3)$$

with boundary conditions $W_{T+1}(\cdot) = 0$. For $t > 1$, Equation (3) reflects the customer’s choice in each month between continuing warranty coverage and incurring total expected cost $B_t(p_t)$ and terminating coverage and incurring expected support costs $R_t(p_t)$ over the remainder of the planning horizon. For $t = 1$, Equation (3) corresponds to the customer’s choice of whether or not to start coverage.

Without loss of generality, we assume that $(1 - q_0(T))c \geq m$; otherwise, the customer never buys coverage in month T because the expected cost of the pay-as-you-go service is cheaper. The quantity $R_t(p_t) - B_t(p_t)$ is the total benefit of continuing the flexible EW in month t , given that the estimated failure probability is p_t .

PROPOSITION 1. *Under Assumptions 1 and 2, if $q_0(t)$ is increasing in t , then $R_t(p_t) - B_t(p_t)$ is increasing in p_t for $t = 1, 2, \dots, T$. Moreover, $R_t(p_t) - B_t(p_t) \leq ((1 - q_0(t))c - m)/q_0(t)$ if $q_0(t) > 0$.*

Proposition 1 says that the total benefit of continuing the flexible EW is increasing in the estimated failure probability p_t , and, moreover, it is bounded above. With these properties, we can now show that the optimal policy has a threshold structure.

THEOREM 1 (THRESHOLD POLICY).

- (a) *If the customer was covered in the previous month, it is optimal to continue coverage in month t if and only if $R_t(p_t) - B_t(p_t) \geq 0$. (She is indifferent between continuing coverage and terminating coverage when the inequality becomes an equality.)*
- (b) *Under Assumptions 1 and 2, if $q_0(t)$ is increasing in t , then there exists a sequence of failure probability estimate thresholds $\{p_t^* : t = 1, \dots, T\}$ such that it is optimal to continue coverage in month t if and only if the estimated failure probability $p_t \geq p_t^*$. Moreover, $p_t^* \leq m/((1 - q_0(t))c)$.*

The optimal thresholds are independent of failure realization, so they can be pre-computed. Proposition 2 shows that the threshold policy also holds under conditions beyond those specified in Theorem 1(b).

PROPOSITION 2 (THRESHOLD POLICY). *Under Assumptions 1 and 2, if $q_0(t) \leq q_1(t)$ for $t = 1, 2, \dots, T-1$, then $R_t(p_t) - B_t(p_t)$ is increasing in p_t and the customer’s optimal coverage policy also has a threshold structure: There exists a threshold p_t^* in month t such that it is optimal to continue coverage in month t if and only if $p_t \geq p_t^*$.*

In particular, if there is no difference between the termination probabilities for functional and failed products, i.e., $q_0(t) = q_1(t)$ for $t = 1, 2, \dots, T-1$, the optimal policy always has a threshold structure under Assumptions 1 and 2. However, the condition $q_0(t) \leq q_1(t)$ for $t = 1, 2, \dots, T-1$ may be less likely to hold in practice as customers are generally more likely to replace a failed product than a functional one, i.e., $q_0(t) > q_1(t)$. Moreover, the optimal policy does not have a threshold structure for general termination probabilities. The following is such an example, although somewhat contrived, in which the optimal policy does not have a simple threshold structure.

EXAMPLE 1. Consider a three-month problem ($T = 2$). Suppose the updating scheme is $p_{t+1}(\mathcal{I}_t) = p_t$ and thus the estimated probability in the next month is the same as in the previous month no matter what happens. Apparently, the updating scheme satisfies Assumptions 1 and 2. Assume $q_0(1) = 100\%$ and $q_1(1) = q_0(2) = q_1(2) = 0\%$, meaning that the customer replaces the product if it fails in month 1; otherwise the customer keeps it until the end of the horizon. Note that the termination probabilities violate the conditions $q_0(1) \leq q_0(2)$ and $q_0(t) \leq q_1(t)$ for $t = 1$. Let $c = \$100$ and $m = \$5$. Then, $R_2(p) = 100p$ and $W_2(p) = \min(5, 100p)$, so it is optimal to buy coverage in month 2 if and only if $p \geq 5\%$. According to Equations (1) and (3), the expected cost of not buying coverage in month 1 is $R_1(p) = (1 - p)R_2(p) = 100p(1 - p)$ and the expected cost of buying coverage in month 1 is $5 + (1 - p)W_2(p)$. It is optimal to buy coverage in month 1 if and only if $5 + (1 - p)W_2(p) \leq 100p(1 - p)$. Consequently, it is optimal to forgo coverage in month 1 if $p < 5\%$. For $p \geq 5\%$, it is optimal to buy coverage in month 1 if and only if $5 + 5(1 - p) \leq 100p(1 - p)$, which is equivalent to $10.6\% \leq p \leq 94.4\%$. So, the optimal policy in month 1 is to buy coverage if and only if $10.6\% \leq p \leq 94.4\%$. This condition does not have a simple threshold structure. The intuition is that for very low failure probability estimates, the customer will not buy coverage because she is unlikely to experience a failure, whereas for very high failure probability estimates, she will not buy coverage either because she is very likely to abandon the product in period 1 due to a

failure, so her monthly premium would be wasted. For more moderate failure probability estimates she will buy coverage.

2.1. Beta Updating Scheme

As a specific example of the failure probability updating scheme, consider the special case where the customer has a prior $\beta(a, b)$ and she updates after each month as follows. Let $a_1 = a$ and $b_1 = b$. Given a_t and b_t at the beginning of month t , then $a_{t+1} = a_t + I_t$ and $b_{t+1} = b_t + (1 - I_t)$. As a result, at the beginning of month $t + 1$ the expected failure probability is $p_{t+1} = a_{t+1}/(a_{t+1} + b_{t+1})$. The expected failure probability in month $t + 1$ is $p_{t+1}^+ = (a_t + 1)/(a_t + b_t + 1)$ if a failure occurs in month t ; $p_{t+1}^- = a_t/(a_t + b_t + 1)$ if no failure occurs in month t . This is the so-called Beta updating scheme (see, e.g., Stigler 1986). Note that the Beta updating scheme is non-stationary. The estimate $p_t(\cdot)$ depends on t , i.e., $p_t(\cdot)$ is different from $p_{t+1}(\cdot)$. It is easy to verify the following.

LEMMA 2.

- (a) The Beta updating scheme $p_{t+1} = a_{t+1}/(a_{t+1} + b_{t+1})$ satisfies Assumptions 1 and 2.
- (b) The Beta updating scheme is a martingale.

Let N_t be the total number of failures prior to month t . The following theorem describes the threshold structure of the optimal policy in terms of N_t under the Beta updating scheme.

THEOREM 2 (OPTIMAL POLICY). *If the optimal policy has a threshold structure with thresholds p_t^* , then there exist a sequence of thresholds $\{x_t^* : t = 1, 2, \dots, T\}$ for the Beta updating scheme such that it is optimal for the customer to continue coverage in month t if and only if $N_t \geq x_t^*$.*

2.2. Stationary Updating Schemes

In this subsection, we consider the case that the failure updating scheme is stationary, i.e., time independent, and derive further structural properties of the customer's optimal coverage policy under this assumption.

ASSUMPTION 3. *The failure probability updating scheme is stationary.*

Under Assumptions 1 and 3, $p_{t+1}(\mathcal{I}_t) = H(p_t, I_t)$, where $H(\cdot, \cdot)$ is a stationary function. Before considering the monotonicity of the thresholds in the optimal coverage policy, we present some monotone properties of the cost functions $W_t(p)$ and $R_t(p)$ with respect to t .

PROPOSITION 3. *Under Assumptions 1–3, if both $q_0(t)$ and $q_1(t)$ are increasing in t , then*

- (a) $R_t(p)$, $B_t(p)$, and $W_t(p)$ are decreasing in t for any $0 \leq p \leq 1$.
- (b) $R_t(p) - B_t(p)$ is decreasing in t for any $0 \leq p \leq 1$.

Proposition 3 says that the benefit of continuing coverage, which is captured by the term $R_t(p) - B_t(p)$, is decreasing when it approaches the end of the planning horizon. Theorem 3 below shows that the thresholds in a customer's optimal policy are increasing in time t .

THEOREM 3 (MONOTONICITY OF THRESHOLDS). *Under Assumptions 1–3, if both $q_0(t)$ and $q_1(t)$ are increasing in t , then the thresholds in the customer's optimal coverage policy are increasing in t , i.e., $p_1^* \leq p_2^* \leq \dots \leq p_T^* = m/((1 - q_0(T))c)$.*

In other words, the customer's likelihood of continuing the flexible EW coverage decreases as it gets closer to the end of the horizon. Although the thresholds are increasing, each month's threshold is bounded above by the estimated failure probability based on the threshold in the previous month. Let p_{t+1}^{*+} be the updated failure probability in month $t + 1$ if a failure occurs in month t given that the estimated failure probability is p_t^* at the beginning of month t , i.e., $p_{t+1}^{*+} = H(p_t^*, 1)$.

PROPOSITION 4. *Under Assumptions 1–3, if $q_0(t)$ is constant, $q_1(t)$ is increasing in t and $p_{t+1}^+ \geq p_t$ for all $0 \leq p_t \leq 1$, then $p_{t+1}^* \leq p_t^*$.*

Proposition 4 tells us that if it is optimal for a customer to buy coverage in month t and, furthermore, if a failure occurs in that month, then it will be optimal for her to continue coverage in the next month as well because her updated failure probability estimate will exceed the threshold p_{t+1}^* in month $t + 1$. This suggests that the customer would keep the flexible EW after a covered failure unless the failure occurs during the last month T of the planning horizon.

2.2.1. Exponential Smoothing Mechanism. An example of a stationary failure probability updating scheme is the exponential smoothing mechanism (see, e.g., Brown and Meyer 1961). Assume that the prior failure probability estimate is p_1 and the customer updates it as follows:

$$p_{t+1} = p_t + \alpha(I_t - p_t) \text{ for } t = 1, 2, \dots, T - 1, \quad (4)$$

where α is constant and $0 \leq \alpha \leq 1$. The parameter α reflects the degree of inertia in the failure probability estimate; lower values of α correspond to more weight on the previous failure probability estimates. The probability estimate in the next month only

depends on the current probability estimate and the failure observation in the current month. The customer compares the failure observation with the estimated failure probability and makes a linear adjustment, so the exponential smoothing mechanism can also be called a linear adjustment scheme. Then, $p_{t+1}^+ = (1 - \alpha)p_t + \alpha$ if a failure occurs in month t ; $p_{t+1}^- = (1 - \alpha)p_t$ otherwise. Expanding the recursion (4) yields,

$$p_{t+1} = (1 - \alpha)^t p_1 + \sum_{i=0}^{t-1} \alpha(1 - \alpha)^i I_{t-i}. \quad (5)$$

Recall that in the Beta updating scheme, the estimate of the failure probability in the next month depends on the number of observed failures up to the current month, and all failures have the same weight regardless of recency. However, in the exponential smoothing mechanism, the more recent failures weigh more.

LEMMA 3.

- (a) The exponential smoothing mechanism satisfies Assumptions 1–3.
- (b) The exponential smoothing mechanism is a martingale.

Thus, all the results for general and stationary updating schemes hold for the exponential smoothing mechanism. As discussed before, if $q_0(t) = q_1(t) = 0$ for $t = 1, 2, \dots, T$, then $R_t(p_t) = (T + 1 - t)p_t c$ by Lemma 1 and it is independent of the adjustment factor α .

The following proposition demonstrates additional structural results for the customer’s coverage decision process under the exponential smoothing mechanism.

PROPOSITION 5. If $q_0(t) = q_1(t)$ for $t = 1, 2, \dots, T$, then the following results hold under the exponential smoothing mechanism.

- (a) Both $B_t(p_t)$ and $W_t(p_t)$ are decreasing concave in α for $0 \leq p_t \leq 1$ and $t = 1, 2, \dots, T$.
- (b) The thresholds p_t^* are decreasing in α for $t = 1, 2, \dots, T$.

Proposition 5 says that the higher the adjustment factor, the more likely the customer is to buy the flexible EW in the first month for any given prior probability estimate.

3. The Provider’s Problem

In this section we will study the provider’s problem assuming that customers use the optimal coverage strategy. We show how to compute the provider’s expected profit for any given flexible EW and then investigate its profitability against the traditional EW.

We assume that the provider has scale economies in repairs. Let βC be the repair cost to the provider for a failure with cost C to the customer, where β is constant and $0 \leq \beta \leq 1$. Denote termination probabilities from month t to the end of the horizon by $\mathcal{Q}_t := \{(q_0(i), q_1(i)) : i = t, t + 1, \dots, T\}$. Let $\pi_t(p_t, m, \mathcal{Q}_t)$ be the provider’s total expected profit per customer from month t to the end of the horizon, given that the customer’s estimated monthly failure probability at the beginning of month t is p_t , the monthly premium is m , and the termination probabilities are \mathcal{Q}_t . Note that $\pi_t(p_t, m, \mathcal{Q}_t)$ also depends on the true failure probability λ . The recursive equations for $\pi_t(p_t, m, \mathcal{Q}_t)$ can be formulated as follows.

If $R_t(p_t) - B_t(p_t) \geq 0$, which can be implied by $p_t \geq p_t^*$ from Theorem 1, then the customer continues buying the flexible EW in month t :

$$\begin{aligned} \pi_t(p_t, m, \mathcal{Q}_t) = & m + \lambda(1 - q_0(t))(\pi_{t+1}(p_{t+1}^+, m, \mathcal{Q}_{t+1}) - \beta c) \\ & + (1 - \lambda)(1 - q_1(t))\pi_{t+1}(p_{t+1}^-, m, \mathcal{Q}_{t+1}); \end{aligned}$$

if $R_t(p_t) - B_t(p_t) < 0$, which can be implied by $p_t < p_t^*$, then the customer stops buying the flexible EW in month t :

$$\pi_t(p_t, m, \mathcal{Q}_t) = 0,$$

with boundary condition $\pi_{T+1}(\cdot, \cdot, \cdot) = 0$. Then, the total expected profit from a customer is $\pi_1(p_1, m, \mathcal{Q}_1)$, where p_1 is the customer’s prior estimate of the monthly failure probability and \mathcal{Q}_1 represents the termination probabilities from month 1 to the end of the horizon.

The dynamic equations for the expected profit of the flexible EW are complicated in general, but essential analytical results can be obtained under certain mild conditions.

ASSUMPTION 4. The termination probabilities are independent of the product failure status, i.e., $q_0(t) = q_1(t)$ for $t = 1, 2, \dots, T$, denoted by $q(t)$.

We expect the termination probabilities to be lower if the product is working, so we can think of this assumption as providing both lower and upper bounds on expected profits. We obtain an upper bound on expected profits if the termination probability for a working product is inflated to be equal to that for a failed one. A lower bound on profits can be obtained when we lower the termination probability of a failed product to be equal to that of a working one.

3.1. Homogeneous Market

We first consider a homogeneous market, where all customers have the same characteristics in terms of prior estimate, learning scheme, true failure probability,

termination probabilities, etc. The problem, faced by the flexible EW provider, is to determine the monthly premium to maximize his total expected profit:

$$\max_{m \geq 0} \pi_1(p_1, m, Q_1). \tag{6}$$

Denote the optimal monthly premium by $m^*(p_1) = \arg \max_{m \geq 0} \pi_1(p_1, m, Q_1)$; then $\pi_1(p_1, m^*(p_1), Q_1)$ is the optimal expected profit. The profit function $\pi_1(p_1, m, Q_1)$ may not be unimodal, but the optimal monthly premium $m^*(p_1)$ is bounded so it can be efficiently found by some non-linear optimization algorithms. Moreover, it is straightforward to show that $\lambda \beta c \min_t (1 - q_0(t)) \leq m^*(p_1) \leq R_1(Q_1)$. This is because the customer will not purchase the flexible EW if the monthly premium is higher than the cost of not buying the warranty, and the flexible EW is not profitable if the premium is lower than the minimum support cost for each period.

Consider a traditional EW with duration T months in the same homogeneous market and let $S_t(Q_t)$ be the total expected support cost of the traditional EW to the provider for a customer from month t to the end of the horizon, given that the customer’s termination probabilities are Q_t . The recursive equations for $S_t(Q_t)$ are as follows:

$$S_t(Q_t) = \lambda(1 - q_0(t))(\beta c + S_{t+1}(Q_{t+1})) + (1 - \lambda)(1 - q_1(t))S_{t+1}(Q_{t+1}), \tag{7}$$

with boundary conditions $S_{T+1}(\cdot) = 0$.

Under some mild conditions, the support cost $S_t(Q_t)$ has a simpler expression, which significantly simplifies the profit comparison between the flexible and the traditional EWs.

LEMMA 4. Under Assumption 4,

$$S_t(Q_t) = \lambda \beta c \sum_{i=t}^T \prod_{j=t}^i (1 - q(j)).$$

In a homogeneous market where only a traditional EW or a pay-as-you-go repair service is available, a customer buys the traditional EW if and only if its price is not greater than the expected cost of not purchasing the flexible EW, which is $R_1(p_1)$ and can be found by solving the recursive equation (1) with termination probabilities Q_1 . The problem, faced by the traditional EW provider in a homogeneous market, is to choose the warranty price r to maximize his total expected profit:

$$\max_{r \geq 0} \pi^t(p_1, r, Q_1), \tag{8}$$

and the indicator function $\mathbf{1}(z) = 1$ if z is true; $\mathbf{1}(z) = 0$ otherwise. Apparently, the optimal price is equal to the customer’s willingness to pay $R_1(p_1)$

if the traditional EW is a profitable business in this homogeneous market, and the optimal profit can be re-expressed by $\max_{r \geq 0} \pi^t(p_1, r, Q_1) = \max\{R_1(p_1) - S_1(Q_1), 0\}$.

PROPOSITION 6. Under Assumptions 1 and 2, the optimal profit of the traditional EW is increasing in the prior estimate of the monthly failure probability p_1 . Moreover, if $q_0(t)$ is increasing in t or $q_0(t) \leq q_1(t)$, the optimal profit of the flexible EW is increasing in p_1 as well.

We are now ready to compare the optimal profits of the flexible EW and the traditional EW in a homogeneous market, where the customer updates the failure probabilities following certain updating schemes.

THEOREM 4. Under Assumptions 1, 2, and 4, in a homogeneous market where customers follow martingale updating schemes and $p_{t+1}^+ - p_{t+1}^-$ is independent of p_t , there exists a threshold p_1^0 such that the flexible EW is strictly more profitable to the provider than the traditional EW if and only if the prior estimate of the monthly failure probability p_1 is less than p_1^0 .

The expression of the threshold p_1^0 can be found in the proof of Theorem 4 in the online Appendix. Both the Beta updating scheme and the exponential smoothing mechanism are martingales, and $p_{t+1}^+ - p_{t+1}^-$ is independent of p_t , so in a homogeneous market where the customer has a low prior failure probability estimate, the customer’s willingness to pay for the traditional EW is low as well and the traditional EW may result in little or no profit, so the flexible EW can be more profitable. The following example illustrates the comparison between the flexible EW and the traditional EW.

EXAMPLE 2. We consider the warranty provider’s problem with a time horizon $T = 12$. The customer updates the estimate of the failure probability according to the exponential smoothing mechanism with different adjustment factors $\alpha = \{0.3, 0.5, 0.7\}$. The expected repair cost is $c = \$100$ to the customer and is $\beta c = 0.6 \cdot \$100 = \60 to the provider. The termination probabilities are $q_0(t) = q_1(t) = 0$ for $t = 1, 2, \dots, 6$ and $q_0(t) = q_1(t) = 5\%$ for $t = 7, 8, \dots, 12$.

For the monthly failure probability $x\%$, we define its equivalent nominal annual failure probability as $12 \cdot x\%$. Assume that the true monthly failure probability is $\lambda = 2\%$ or the equivalent nominal annual failure probability is 24% (more precisely, the equivalent actual annual failure probability is equal to $1 - (1 - 2\%)^{12} = 21.5\%$). Solving optimization problems (6) and (8) gives the optimal profits for the warranty provider from the flexible and traditional EWs, respectively.

Figure 1 illustrates the profit comparison for the traditional EW and the flexible EWs with different exponential smoothing adjustment factors $\alpha \in \{0.3, 0.5, 0.7\}$. If the prior estimate of the monthly failure probability p_1 is lower than 1.25%, the traditional EW earns no profit. In contrast, the flexible EW always results in profit from each customer, and the profit is non-trivial even for a very low prior failure probability. The profit of the traditional EW is strictly increasing in p_1 at a constant rate for $p_1 > 1.25\%$ because $\max_{r \geq 0} \pi_t(p_1, r, Q_1) = R_1(p_1) - S_1(Q_1)$, $S_1(Q_1)$ is independent in p_1 , and $R_1(p_1)$ is linearly increasing in p_1 by Lemma 1. The profits of the flexible EWs with different adjustment factors are all increasing in p_1 but at a lower rate than the traditional EW. The flexible EW is more profitable than the traditional one for sufficiently small prior failure probability. For example, if the prior monthly failure probability is less than 1.88% or the prior nominal annual failure probability is less than 22.56%, the flexible EW is strictly more profitable than the traditional one in a homogeneous market, where customers update the failure probability estimate following the exponential smoothing scheme with adjustment factor $\alpha = 0.3$. From Figure 1, we see that the profitability thresholds for the flexible EW against the traditional EW are $p_1^0 = 2.13\%$ and $p_1^0 = 2.10\%$ for adjustment factors $\alpha = 0.5$ and $\alpha = 0.7$, respectively.

A profit comparison with different adjustment factors shows that the profit is increasing in the adjustment factor if the prior failure probability is relatively low (lower than 1.75% in this example). The monotonicity of the profit is not clear if the prior failure probability is relatively high. Figure 1 demonstrates that the profit of the flexible EW for

adjustment factor $\alpha = 0.5$ is higher than the profits for $\alpha = 0.3$ and $\alpha = 0.7$ when p_1 is greater than 2%.

The profit comparison for different adjustment factors seems to contradict Proposition 5, which indicates that the customer’s purchase threshold is decreasing in the adjustment factor. However, Proposition 5 shows the static monotonicity of the customer’s optimal policy, while the provider’s profit not only depends on the optimal policy but also relates to how the failure probability may dynamically evolve. The evolution estimate of the failure probability may be decreasing or lower than the threshold according to the updating scheme (5) if the true failure probability is very low and the adjustment factor is relatively high, so customers may drop the flexible EW relatively early.

Unlike the prior estimate of the failure probability, increasing true failure probability adversely affects the provider’s profit. The following example illustrates the comparison between the flexible and traditional EWs with respect to the true failure probability in a homogeneous market.

EXAMPLE 3. We continue with Example 2 and compare the profits of the flexible and traditional EWs with respect to the true failure probability. Assume that the prior estimate of the monthly failure probability is 3%, i.e., $p_1 = 3\%$. Figure 2 illustrates the profit comparison.

The optimal profit of the traditional EW is $\max\{R_1(p_1) - S_1(Q_1), 0\} = \max\{(p_1 - \lambda\beta)c \sum_{i=t}^T \prod_{j=t}^i (1 - q(j)), 0\}$ by Lemmas 1 and 4, so it is linearly decreasing in λ for $\lambda \leq p_1/\beta$, i.e., $\lambda \leq 3\%/0.6 = 5\%$; it is equal to zero for any $\lambda > 5\%$.

Figure 1 Profit Comparison w.r.t. Prior Estimate in a Homogeneous Market

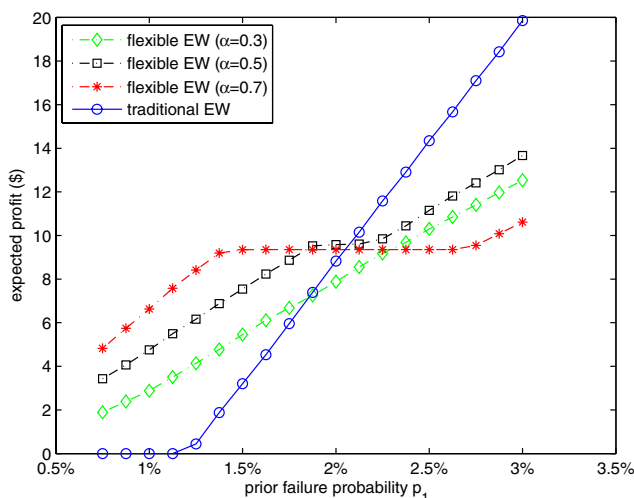
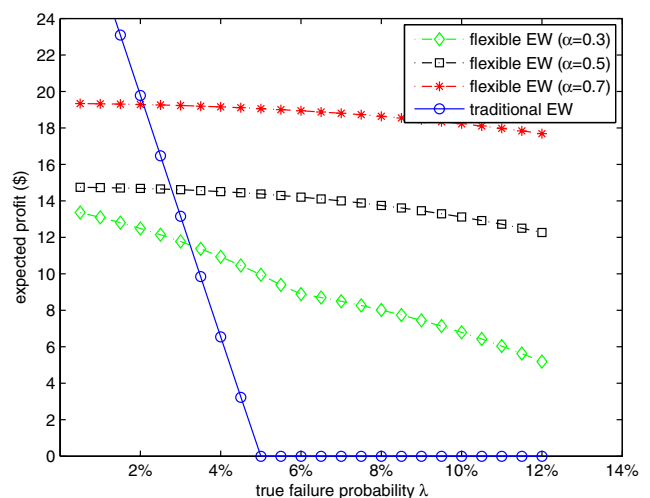


Figure 2 Profit Comparison w.r.t. True Failure Probability in a Homogeneous Market



The profits of the flexible EWs are all decreasing with respect to the true monthly failure probability but at a much lower rate than the traditional EW for all adjustment factors $\alpha \in \{0.3, 0.5, 0.7\}$ in the exponential smoothing mechanism. The flexible EW is strictly more profitable than the traditional EW for any $\lambda > 3.2\%$ with the adjustment factor $\alpha = 0.3$. The thresholds of the profitability for the flexible EW are 2.75% and 2.2% with adjustment factors $\alpha = 0.5$ and $\alpha = 0.7$, respectively. Figure 2 also demonstrates that in this example the profit of the flexible EW is increasing in the adjustment factor α for any $0 \leq \lambda \leq 12\%$.

If customers underestimate the true failure probability, their willingness to pay for warranties is low, so the traditional EW may earn little or no profit. However, the profit of the flexible EW does not decrease much because customers correct their estimates based on observed failures, and they may keep buying the flexible EW for longer period of time.

We also discuss the profitability of the flexible EW against a menu of multiple traditional EWs with different coverage durations. Our studies demonstrate that the single flexible EW may outperform multiple traditional EWs. We leave the analysis to the online Appendix.

3.2. Heterogeneous Market

Next, we consider a heterogeneous market with heterogeneity along various dimensions. We illustrate the profitability of flexible EWs through analytical results and numerical examples. For tractability, assume that there are two market segments: type- L and type- H customers. Market segment n , $n \in \{L, H\}$, has a prior failure probability estimate p_1^n , termination probabilities $Q_1^n := \{(q_0^n(i), q_1^n(i)) : i = 1, 2, \dots, T\}$, and proportion γ^n with $\gamma^L + \gamma^H = 1$. Without loss of generality, assume that type- H customers have higher prior failure probability, i.e., $p_1^H \geq p_1^L$. The true monthly failure probabilities are denoted by λ^L and λ^H for type- L and type- H customers, respectively.

To investigate the profitability of the flexible EW, we evaluate the performance of the traditional EW in a heterogeneous market where only a traditional EW or pay-as-you-go repair service is available. A type- n customer buys the traditional EW if and only if its price r is not greater than the cost of not purchasing warranty $R_1^n(p_1^n)$, which can be found by solving recursive equation (1) with termination probabilities Q_1^n , $n \in \{L, H\}$. The support cost for the warranty provider is $S_1^n(Q_1^n)$, which can be found by solving (7) with termination probabilities Q_1^n and true failure probability λ^n , $n \in \{L, H\}$. The problem faced by the traditional EW provider is to choose the warranty price r to maximize his total expected profit over the two market segments:

$$\max_{r \geq 0} \{ \gamma^L (r - S_1^L(Q_1^L)) \cdot \mathbf{1}(r \leq R_1^L(p_1^L)) + \gamma^H (r - S_1^H(Q_1^H)) \cdot \mathbf{1}(r \leq R_1^H(p_1^H)) \}. \quad (9)$$

If the termination probabilities are the same for type- L and type- H customers, i.e., $Q_1^L = Q_1^H$, then the type- H customer's willingness to pay for the traditional EW is higher, i.e., $R_1^H(p_1^L) > R_1^L(p_1^L)$. The optimal price of the traditional EW should be the willingness to pay of type- L or type- H customers if offering the traditional EW is a profitable business, so problem (9) becomes

$$\max \{ R_1^L(p_1^L) - \gamma^L S_1^L(Q_1^L) - \gamma^H S_1^H(Q_1^H), \gamma^H (R_1^H(p_1^H) - S_1^H(Q_1^H)), 0 \}.$$

PROPOSITION 7. *Under Assumptions 1, 2, and 4, in a heterogeneous market of two market segments with martingale updating schemes and the same termination probabilities, the optimal traditional EW captures both segments L and H if and only if $p_1^H < (p_1^L - \gamma^L \lambda^L \beta) / \gamma^H$.*

Proposition 7 says that the optimal traditional EW captures both segments if and only if the segments are not very different in their prior probability. Note that Proposition 7 is independent of type- H customers' true failure probability, and segments L and H may follow different martingale updating schemes. If the provider offers a traditional EW with a uniform price to both market segments, the customers with low priors may not buy coverage. We remark that when customers are heterogeneous to a large degree (e.g., $p_1^H \geq (p_1^L - \gamma^L \lambda^L \beta) / \gamma^H$), it is optimal for the traditional EW only to capture market segment H . Moreover, the condition $p_1^H < (p_1^L - \gamma^L \lambda^L \beta) / \gamma^H$ can be rewritten as $p_1^H < (p_1^L - \lambda^L \beta) / \gamma^H + \lambda^L \beta$ by using $\gamma^L + \gamma^H = 1$. Then, we see that the higher the proportion of type- H customers, the more likely the optimal traditional EW only captures type- H customers; the lower λ^L or β , the more likely the optimal traditional EW captures both the market segments because the coverage cost is lower.

For the flexible EW in this heterogeneous market, the expected profit to the flexible EW provider with monthly premium m from market segment n is $\pi_1^n(p_1^n, m, Q_1^n)$, which can be found by solving the dynamic programs discussed before with λ^n, p_1^n, Q_1^n , $n \in \{L, H\}$. The provider's problem is to choose the monthly premium m to maximize his total expected profit over the two market segments:

$$\max_{m \geq 0} \{ \gamma^L \pi_1^L(p_1^L, m, Q_1^L) + \gamma^H \pi_1^H(p_1^H, m, Q_1^H) \}. \quad (10)$$

Denote the optimal monthly premium to problem (10) by $m^*(p_1^L, p_1^H)$, i.e.,

$$m^*(p_1^L, p_1^H) = \arg \max_{m \geq 0} \{ \gamma^L \pi_1^L(p_1^L, m, \mathcal{Q}_1^L) + \gamma^H \pi_1^H(p_1^H, m, \mathcal{Q}_1^H) \}.$$

A flexible EW offers an opportunity to all customers to obtain coverage while learning about the product reliability at a reasonable cost and thus can appeal to a broader market and earn more profit for the provider.

Customers can differ in multiple dimensions, including true failure probability, prior failure probability estimate, termination probabilities, and learning schemes. To isolate various effects, we consider the flexible EW's profitability in a heterogenous market with respect to one of the factors at each time.

3.2.1. Heterogeneity in Prior Estimate. We first investigate the heterogeneity in the prior estimate of the monthly failure probability, assuming that all other factors including the true failure probability are the same for both type- L and type- H customers. So we omit superscripts L or H in the true failure probabilities and termination probabilities. As customers buy EWs to cover the risk of future repair cost in the event of product failures, their prior estimates about the product reliability are among the most important factors influencing the warranty purchase decisions. According to the literature (see, e.g., Erdem et al. 2008 and Chen et al. 2009), customers' prior estimates about the product reliability depend on brand value, product price, customers' usages, habits, gender, income level, etc.

THEOREM 5. *Under Assumptions 1, 2, and 4, in a heterogeneous market of two segments L and H with martingale updating schemes and where $p_{t+1}^L - p_{t+1}^H$ is independent of p_t , there exist two points \hat{p}_1^H and \check{p}_1^H such that the flexible EW is strictly more profitable to the provider than the traditional EW if and only if type- H customers' prior failure probability estimate p_1^H falls into an intermediate range $\hat{p}_1^H < p_1^H < \check{p}_1^H$.*

Note that \hat{p}_1^H and \check{p}_1^H may depend on type- L customers' prior failure probability p_1^L , market segment proportion γ^L (or γ^H), and other model primitives; their expressions can be found in the proof of Theorem 5 in the online Appendix. From Proposition 7, the traditional EW captures only the type- H customers if and only if $p_1^H \geq (p_1^L - \gamma^L \lambda \beta) / \gamma^H$, so there are two scenarios for comparing the profitability of the flexible EW to the traditional EW in this heterogeneous market. If $p_1^H < (p_1^L - \gamma^L \lambda \beta) / \gamma^H$, the traditional EW captures both the market segments, and the optimal price is the will-

ingness to pay of type- L customers. The profit of the traditional EW is independent of the prior estimate of the failure probability p_1^H of type- H customers as long as $p_1^H < (p_1^L - \gamma^L \lambda \beta) / \gamma^H$. The profit of the flexible EW is increasing in p_1^H by Proposition 6, so there exists a threshold \hat{p}_1^H such that the flexible EW is strictly more profitable than the traditional EW for any $p_1^H > \hat{p}_1^H$. If $p_1^H \geq (p_1^L - \gamma^L \lambda \beta) / \gamma^H$, the traditional EW only serves type- H customers. The profits of the traditional EW and the flexible EW are both increasing in p_1^H , but the former is increasing faster, so there exists an upper bound \check{p}_1^H on p_1^H below which the flexible EW is strictly more profitable than the traditional one.

Theorem 4 indicates that in a homogeneous market with only type- H customers, there exists an upper bound p_1^o such that the flexible EW is strictly more profitable if and only if $p_1^H < p_1^o$. In a heterogeneous market, Theorem 5 states that there is an upper bound \check{p}_1^H on p_1^H for the profitability of the flexible EW over the traditional EW. Comparing the two upper bounds, it is straightforward to see that $\check{p}_1^H \geq p_1^o$, because type- L customers may contribute some profit to the flexible EW although they do not purchase the traditional one.

EXAMPLE 4. We extend the profitability analysis of the flexible EW to a heterogeneous market of two segments with proportions $\gamma^L = 25\%$ and $\gamma^H = 75\%$. Both type- L and type- H customers follow the exponential smoothing mechanism with adjustment factor $\alpha = 0.3$. Assume that the true monthly failure probability is $\lambda = 2\%$, the same as in Example 2. We will fix the prior monthly failure probability of type- L customers at $p_1^L = 2.5\%$ (its equivalent nominal annual failure probability is 30%) and study the profitability variation with respect to p_1^H , the prior monthly failure probability of type- H customers.

In this example, we allow p_1^H to vary in a large range; we do not restrict that $p_1^H > p_1^L$. From Proposition 7, the optimal traditional EW captures both segments if and only if $\gamma^L p_1^L + \gamma^H \lambda \beta < p_1^H \leq (p_1^L - \gamma^L \lambda \beta) / \gamma^H$, i.e., $1.53\% < p_1^H < 2.93\%$. In Figure 3, for any $p_1^H \leq 1.53\%$, the traditional EW only captures type- L customers, and the profit is constantly equal to \$3.5; for $1.53\% < p_1^H \leq 2.5\%$, the traditional EW captures the two segments and the price is equal to the willingness to pay of type- H customers, so the total profit is increasing in p_1^H ; for $2.5\% \leq p_1^H < 2.93\%$, the traditional EW also earns both segments but the total profit is constant because the price is constantly equal to the willingness to pay of type- L customers; for $p_1^H \geq 2.93\%$, the traditional EW only captures type- H customers and the profit is linearly increasing in p_1^H .

Figure 3 Profit Comparison in a Heterogeneous Market Differing in Prior Estimate

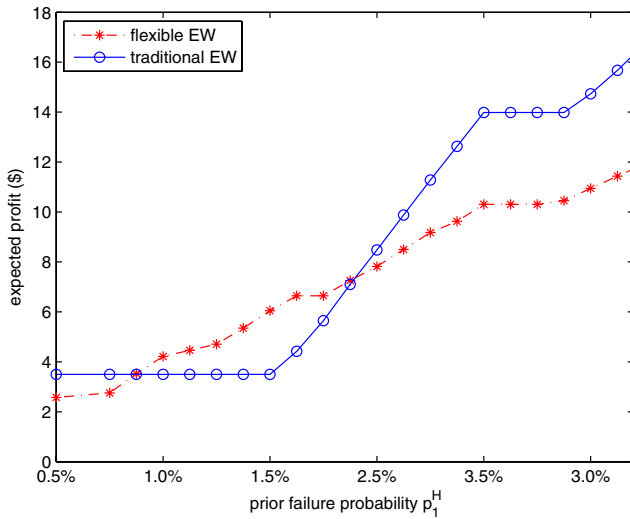


Figure 3 also illustrates that the profit of the flexible EW is increasing in the prior estimate of the monthly failure probability of type-H customers. The flexible EW is strictly more profitable to the provider than the traditional EW if and only if the prior failure probability p_1^H falls into an intermediate range $0.88\% < p_1^H < 2\%$. Recall in Example 2 that the threshold for profitability of the flexible EW is $p_1^H = 1.88\%$ in a homogeneous market, which is lower than the threshold 2% in this heterogeneous market. This is indeed consistent with the above argument.

3.2.2. Heterogeneity in True Failure Probability. We have studied a heterogeneous market with two segments differing in the prior estimates of their failure probabilities. In reality, customers may also be heterogeneous in other dimensions, like the true failure probability, learning schemes, and termination probabilities.

Heterogeneity in true failure probabilities may arise when customers vary in their usage intensity or operating environments. To analyze the profitability of the flexible EW in a heterogeneous market of two segments differing in the true failure probability, we assume that all other factors are the same. So we omit superscripts L or H in the prior estimates and termination probabilities. Note that while the expected costs are perceived the same to the two segments, their actual support costs are different to the warranty provider. Similar to Theorem 4, the following result holds for this heterogeneous market.

THEOREM 6. Under Assumptions 1, 2, and 4, in a heterogenous market with two segments differing only in the true failure probability, where customers follow martingale updating schemes and $p_{t+1}^+ - p_{t+1}^-$ is

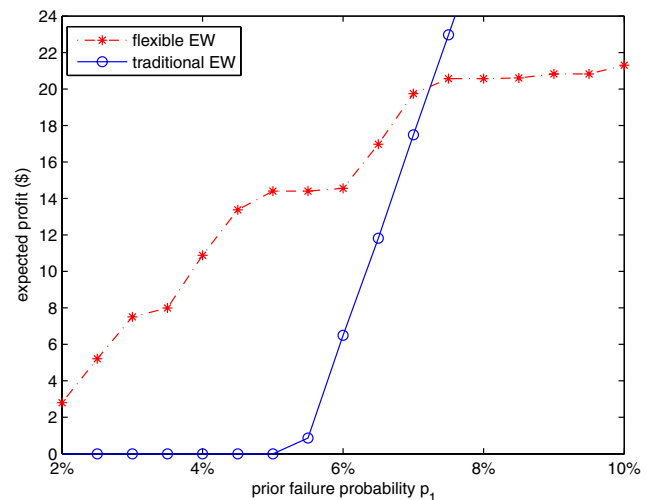
independent of p_t , there exists a threshold $p_1^\#$ such that the flexible EW is strictly more profitable than the traditional one if and only if their prior failure probability p_1 is less than $p_1^\#$.

The following example illustrates the profit comparison in such a heterogeneous market, where customers differ only in their true failure probabilities.

EXAMPLE 5. We modify Example 2 to account for a heterogeneous market of two segments with proportions $\gamma^L = 75\%$ and $\gamma^H = 25\%$. Type-L and type-H customers update their estimates following the exponential smoothing mechanism with the same adjustment factor $\alpha = 0.3$. The true monthly failure probabilities are $\lambda^L = 3\%$ and $\lambda^H = 10\%$ for type-L and type-H customers, respectively. (Equivalently, the annual expected numbers of failures are 0.36 and 1.20 for the two types of customers, respectively.) We assume that the prior estimates are the same for type-L and type-H customers, denoted by p_1 . In this example, we will vary the prior estimate p_1 to compare profits of traditional and flexible EWs in this heterogeneous market.

If the prior estimate of the failure probability is less than 5%, the traditional EW earns no profit because type-H customers underestimate the failure probability and their willingness to pay is not enough to cover their support cost. However, the flexible EW always receives non-trivial profit by offering reliability learning opportunities to customers, especially to the type-H customers who would realize that they were initially too optimistic and would continue coverage for a relatively long period of time. Figure 4 demonstrates that the flexible EW

Figure 4 Profit Comparison in a Heterogeneous Market Differing in True Failure Probability



is more profitable than the traditional one in a fairly large range of prior estimates in this heterogeneous market. There exists a threshold $p_1^\# = 7.25\%$ such that the flexible EW is more profitable than the traditional one for any prior estimate less than 7.25%.

In a heterogeneous market with segments differing in the true failure probabilities, their perceived support costs are the same, so they will behave the same if the manufacturer only offers the single traditional EW (or multiple traditional EWs; see the discussion in the online Appendix). In contrast, the flexible EW can achieve *market segmentation* through customers' dynamic reliability learning and coverage continuation decisions based on their own experience with the product.

3.2.3. Warranty Menu. Because a flexible EW offers many advantages, it may attract a broad range of customers. But for a warranty provider who also sells a traditional EW with one or more fixed durations of coverage, the introduction of a flexible EW will likely cannibalize some or all the demand of the traditional EW. Therefore, the flexible EW should be carefully designed and priced to avoid eroding profits of existing traditional EWs and indeed to improve profitability.

We continue our discussion of the heterogeneous market and consider an EW menu that includes a traditional EW and a flexible EW. Under the assumption of individual rationality, each customer will select the option with the lowest total expected support cost. Moreover, if the customer selects the flexible EW, she will make a coverage decision in each month strategically to minimize her expected cost as discussed before. Therefore, the problem for the warranty provider is to find the flexible EW premium m and the traditional EW price r to maximize the total expected profit:

$$\begin{aligned} \max_{m \geq 0, r \geq 0} \{ & \gamma^L ((r - S_1^L(Q_1^L)) \cdot \mathbf{1}(r \leq W_1^L(p_1^L))) \\ & + \pi_1^L(p_1^L, m, Q_1^L) \cdot \mathbf{1}(r > W_1^L(p_1^L))) \\ & + \gamma^H ((r - S_1^H(Q_1^H)) \cdot \mathbf{1}(r \leq W_1^H(p_1^H))) \\ & + \pi_1^H(p_1^H, m, Q_1^H) \cdot \mathbf{1}(r > W_1^H(p_1^H)) \}, \end{aligned} \quad (11)$$

where $W_1^n(p_1^n)$ is the minimum cost of the flexible EW to a customer in market segment n and can be found by solving recursive equations (1–3) with termination probabilities Q_1^n , $n \in \{L, H\}$. In problem (11), we break the tie in the way that a customer will buy the traditional EW if she is indifferent between the flexible and traditional EWs.

Note that the minimum cost under the flexible EW $W_1^n(p_1^n)$ has taken into account the non-purchase alternative, so problem (11) considers four possible cases:

Each type of customers may buy either the traditional EW or the flexible EW. For instance, if the manufacturer intends to sell the flexible EW to type- L customers and sell the traditional EW to type- H customers, the problem can be formulated as follows:

$$\begin{aligned} \max_{m \geq 0, r \geq 0} & \gamma^L \pi_1(p_1^L, m, Q_1^L) + \gamma^H (r - S_1^H(Q_1^H)) \\ \text{s.t., } & r > W_1^L(p_1^L) \text{ and } r \leq W_1^H(p_1^H), \end{aligned}$$

where $r > W_1^L(p_1^L)$ and $r \leq W_1^H(p_1^H)$ are the corresponding incentive compatibility constraints for this particular case, indicating that type- L customers will purchase the flexible EW and type- H customers will buy the traditional EW.

The optimization formulations for the other three cases are similar, and all the four cases are addressed by the indicator functions in problem (11). In general, it is hard to tell who will buy which EW in the optimal EW menu under any arbitrary market heterogeneity or any market segmentation because customers' self-selection and the calculation of the flexible EW profit under each EW menu involve solving several dynamic programs. The following example numerically compares the profits of the EW menu to the traditional or flexible EW alone.

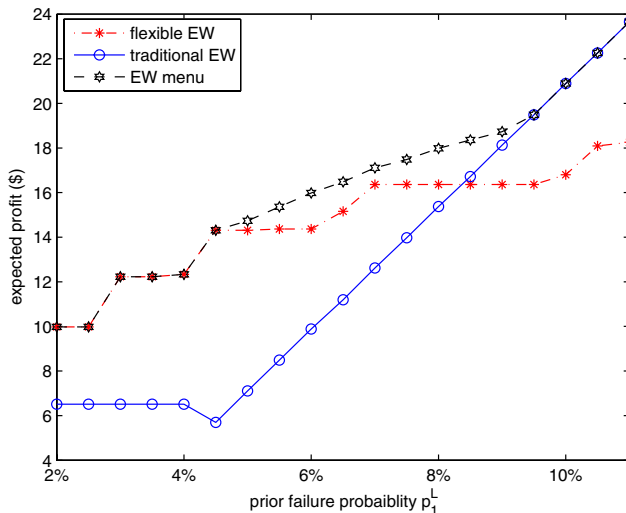
EXAMPLE 6. Again, we consider a heterogeneous market consisting of two market segments, type- L and type- H customers, with different true failure probabilities, prior failure probabilities, and termination probabilities. The time horizon is $T = 12$ and the expected repair cost is $c = \$100$ for customers and $\beta c = 0.6c = \$60$ for the provider. The termination probabilities for type- L and type- H customers are $q_0^L(t) = q_1^L(t) = q_0^H(t) = q_1^H(t) = 0$ for $t = 1, 2, \dots, 6$; $q_0^L(t) = q_1^L(t) = 5\%$, $q_0^H(t) = q_1^H(t) = 15\%$ for $t = 7, 8, \dots, 12$. The true monthly failure probabilities are $\lambda^L = 4\%$ and $\lambda^H = 8\%$ for type- L and type- H customers, respectively. (The equivalent nominal annual failure probabilities are 48% and 96% for the two types of customers.) The type- H customers' prior estimate of the monthly failure probability is $p_1^H = 5\%$ (the equivalent nominal annual failure probability is 60%). The type- L customers' prior estimate of the monthly failure probability varies between 3% and 10% (see Table 1). Type- L and type- H customers update their estimates using the exponential smoothing mechanism with the same adjustment factor $\alpha = 0.3$.

The provider offers an EW menu consisting of a flexible EW and a traditional EW; customers choose the option (including non-purchase) with the lowest total expected support cost. Table 1 shows the optimal EW menu and customers' choices with respect to different prior estimates p_1^L of type- L customers'

Table 1 Optimal EW Menu

p_1^L	3.0%	3.5%	4.0%	4.5%	5.0%	5.5%	6.0%	6.5%	7.0%	7.5%	8.0%	8.5%	9.0%	9.5%	10.0%
m^*	\$9.5	\$9.5	\$8.8	\$9.5	\$9.5	\$9.5	\$9.5	\$9.5	\$9.5	\$9.5	\$9.5	\$9.5	\$9.5	n.a.	n.a.
r^*	n.a.	n.a.	n.a.	n.a.	\$45.5	\$48.0	\$50.5	\$52.5	\$55.0	\$56.5	\$58.5	\$60.0	\$61.5	\$104.4	\$110.1
Type-L	F	F	F	F	T	T	T	T	T	T	T	T	T	T	T
Type-H	F	F	F	F	F	F	F	F	F	F	F	F	F	N.B.	N.B.

Note: m^* (respectively, r^*) is the optimal monthly premium (respectively, the optimal traditional EW price); n.a. represents that this type of EW is not profitable; F (respectively, T and N.B.) stands for that customers who buy the flexible EW (respectively, the traditional EW and neither of them).

Figure 5 Comparison among a Flexible EW, a Traditional EW, and an EW Menu

monthly failure probability, in this heterogeneous market of two segments differing in the true failure probability, prior estimate, and termination probabilities. For $p_1^L \leq 4.5\%$, the optimal EW menu only consists of the flexible EW and both segments buy it; for $5.0\% \leq p_1^L \leq 9.0\%$, the optimal warranty menu includes both the flexible and the traditional EWs, and the type-L customers buy the traditional EW and the type-H customers purchase the flexible one; for $p_1^L \geq 9.5\%$, the optimal EW menu only consists of the traditional EW and only the type-L customers buy it.

Figure 5 demonstrates the profit comparison among a flexible EW, a traditional EW, and an EW menu. Obviously, the EW menu performs at least as well as each EW alone. In particular, when the type-L customer's prior estimate is in an intermediate range, i.e., $5.0\% \leq p_1^L \leq 9.0\%$, it is optimal to offer a combination of a traditional EW and a flexible EW, targeting type-L and type-H customers, respectively. Moreover, the EW menu generates *strictly* more profit than a flexible EW or a traditional EW alone.

4. Conclusion

As product margins decline in increasingly competitive hardware markets, EWs with high margins

and high revenues are becoming critical to manufacturers' profitability. Beyond direct profits, post-sale services are a critical lever for influencing customer loyalty, and in commodity product businesses, service quality and variety are important competitive differentiators. Through the sale of innovative and attractive post-sales services, a manufacturer can enable its customers to reduce product support cost and can increase customers' loyalty. Attractive EW offerings can increase service attach rates and enhance profitability.

In this study, we have demonstrated that a flexible EW is attractive to a broader range of market segments and can significantly improve the provider's profit through market expansion to customers who upgrade more frequently or abandon the hardware earlier. Moreover, unlike the traditional EW that only captures customers whose prior failure probability is high, a flexible EW is also attractive to customers with low prior failure probability, as it offers an opportunity to learn the true failure probability and correct their estimation at a reasonable price while being covered under warranty.

Through analytical results and numerical examples, we show that if the provider only offers one type of warranty coverage—either the traditional EW or the flexible EW—the flexible EW can capture more customers and earn more profit than the traditional one. Furthermore, we also show that if the provider offers a menu of both traditional and flexible EWs, it is optimal to offer the traditional EW at a relatively higher price to customers who keep the hardware longer and offer the flexible EW at a premium price to customers who will upgrade the hardware sooner. The menu of traditional and flexible EWs can capture even more market segments and earn higher profit. The superiority of the EW menu over the individual offerings is robust to the market composition.

Acknowledgments

The authors thank Enis Kayis and Shailendra Jain for numerous discussions in Hewlett-Packard Labs. They are very grateful to a senior editor and two anonymous referees whose thoughtful comments and constructive suggestions have helped to improve the quality and exposition of this study. In addition, the fourth author is also grateful for

financial support by the 2008–2010 Innovation Research Program Award of Hewlett-Packard Labs.

References

- Brown, R. G., R. F. Meyer. 1961. The fundamental theorem of exponential smoothing. *Oper. Res.* **9**(5): 673–687.
- Chen, T., A. Kalra, B. Sun. 2009. Why do consumers buy extended service contracts? *J. Consum. Res.* **36**(4): 611–623.
- Cohen, M. A., N. Agrawal, V. Agrawal. 2006. Winning in the aftermarket. *Harv. Bus. Rev.* **84**(5): 129–138.
- ConsumerReports.org. 2009. Why you don't need an extended warranty? Available at <http://news.consumerreports.org/money/2009/11/why-you-dont-need-an-extended-warranty.html> (accessed date March 20, 2011).
- Erdem, T., M. P. Keane, B. Sun. 2008. A dynamic model of brand choice when price and advertising signal product quality. *Mark. Sci.* **27**(6): 1111–1125.
- Gallego, G., M. Hu, J. L. Beltran, J. Ward, S. Jain. 2009. Channel management for basic products and ancillary services. Working Paper, HP Labs, Palo Alto, CA.
- Hartman, J., K. Laksana. 2009. Designing and pricing menus of extended warranty contracts. *Naval Res. Log.* **56**(3): 199–214.
- Hollis, A. 1999. Extended warranties, adverse selection and aftermarkets. *J. Risk Insur.* **66**(3): 321–343.
- Jack, N., D. N. P. Murthy. 2007. A flexible extended warranty and related optimal strategy. *J. Oper. Res. Soc.* **58**(12): 1612–1620.
- Lutz, N., V. Padmanabhan. 1994. Income variation and warranty policy. Working Paper, Graduate School of Business, Stanford University, Stanford, CA.
- Lutz, N., V. Padmanabhan. 1998. Warranties, extended warranties, and product quality. *Int. J. Ind. Organ.* **16**(4): 463–493.
- Moskowitz, H., Y. H. Chun. 1994. A Poisson regression model for two-attribute warranty policies. *Naval Res. Log.* **41**(3): 355–376.
- Padmanabhan, V. 1995. Usage heterogeneity and extended warranties. *J. Econ. Manag. Strat.* **4**(1): 33–53.
- Padmanabhan, V., R. C. Rao. 1993. Warranty policy and extended service contracts: Theory and an application to automobiles. *Mark. Sci.* **12**(3): 230–247.
- Stigler, S. M. 1986. *The History of Statistics*. Harvard University Press, Boston, MA.
- Wang, R., J. Ward, G. Gallego, F. Balestrieri, M. Hu, E. Kayis, J. Beltran, S. Jain. 2010. Increase attach rates and profitability through flexible-duration extended warranties. Technical Report, HP Labs, Palo Alto, CA.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix A. Proofs.

Appendix B. Single Flexible EW vs. Multiple Traditional EWs.