Transforming the Warranty Sector Through Analytics

Mastering Analytics to Build a Stronger, More Profitable Sector

David King, Hongjie Wang, Paul Swenson, Richard Vermillion

Abstract

We provide a comprehensive assessment of the warranty sector on business model and processes, marketing paradigms, focusing on risk management issues from a practitioner's stand point. We argue that although this sector has enjoyed steady growth and profit, in many aspects of the business process, the standards and best practices are still evolving with many opportunities for significant improvements. With the increasing complexity of this business process, instead of optimizing individual areas, such as marketing or risk management separately, it is critical to integrate these pieces together. We draw lessons and analogies from the credit card industry and argue that it is imperative to leverage data, information and advanced analytics for companies to continue their success. Finally, a case study using advanced analytics that facilitates both service contract marketing and risk management decisions is presented.

Introduction

Warranty management, both for manufacturer limited warranties and the growing sector of extended warranties, is still an evolving field. Indeed, as we argue in this paper, it might even be viewed as immature when compared to other business processes. Fortunately, we believe that the next generation of management tools is now available, tools that will enable this sector to grow and prosper.

Our main focus will be on how the sector needs to develop a more sophisticated and multi-faceted approach to using data and information. As a steady stream of papers and books have argued over the past several decades, mastery of information and the use of advanced analytical methods is a requirement for being competitive.[1] The authors believe that many analytical practices used in the field are outmoded and were designed in an era where warranty was considered mainly an engineering problem or as a small part of the overall product marketing strategy, rather than a profitable business process in its own right. Such limitations constrain warranty product innovation and profitability.

We will also focus our discussion more on extended service plans (ESPs), both for the sake of convenience – we have several interesting data sets that help illustrate our points – and because there are more variables that interact in ESPs than in limited warranty programs. For example, consumer behavior is a much more important factor in ESPs. Since it is discretionary, rather than mandated by government regulation and/or competitive pressures, ESP purchase is largely determined by consumers' perceptions of

[1] For example, the recent popular book on this subject is Competing on Analytics: The New Science of Winning, by Thomas Davenport and Jeanne Harris (Harvard Business Press, 2007).

risk and benefit. On the cost side, the moral hazard for filing a claim and the risk of adverse selection depend heavily on consumer characteristics. However, many elements of this paper apply to limited programs, as well.

While we discuss certain specific analytical techniques, our purpose in doing so is not to promote one narrow technique. Rather, we think this sector's maturation requires that one consider and use information in more sophisticated ways. Thus, the technical part of this paper is designed more to show how a new level of analytical thinking will help promote the growth of and profitability of this sector.

The authors also believe that the ESP risk structure and management paradigm has changed significantly since the early 1990's, and this sector has not received the type of support needed in these areas to evolve manufacturer and retail ESP programs to new levels of profitability and compliance. Here, too, we believe new management tools are needed to understand risk in different ways, which in turn will lead to better strategies and options around underwriting and administration.

At the heart of our thesis is a belief that analytics has the power to transform companies and even entire sectors. We will begin by providing an overview of how this happened in another sector. Next, we will offer an overview of today's warranty analytics and provide an example of a new type of analysis. Lastly, we will show how such analysis can be used by senior business managers to drive new strategies.

A Useful Example

The credit card industry represents a good example of a business sector that underwent a similar transformation in the 1980s. At that time, credit card adoption and use by consumers was relatively low compared to today. The industry had existed for over thirty years, but its analytical, marketing, and risk practices were fairly basic and quite conservative: most marketing was aimed at relatively affluent, low-risk consumers and minimizing losses was still the primary goal. As a result, most Americans still used cash and checks as their primary payment methods.

This changed when a new set of credit card players, such as Capital One, Advanta, and MBNA, took a different approach. Behind their marketing efforts lay a significantly more sophisticated, multi-level analytical orientation. For example, card issuers knew that as they extended credit to more consumers, losses from charge-offs due to bad debt would also rise. Yet new analytical capabilities provided the confidence to marketers to be able to manage this increased risk, both at an individual consumer level and at the portfolio level.

Analytics also led to the development of new credit card products. For example, balance transfer offers were an innovation designed to help both acquire new customers and, more importantly, to activate the line of credit immediately. Likewise, convenience checks helped consumers use their credit in new ways.

These new players were also better in integrating processes and information systems. They recognized that a business could not be successful solely by having efficient marketing or solid risk management. Rather, these processes needed to be well integrated, which in turn meant that information and analysis had to be integrated. For example, MBNA relied heavily on affinity marketing, and its management recognized that card activation, credit risk and customer retention must be considered at the time of acquisition, rather than be handled as segregated decision processes.

There is typically an inherent tension between risk managers and marketers. If the business errs on the conservative side by being too risk-averse, growth will be constrained. If financial offers are extended too readily without a long-term consideration of risk, then the business may be threatened, something that happened on occasion during the first big expansion in credit cards and more recently in the mortgage lending sector.

Yet, on the whole, the credit card sector provides a solid example of how these two potentially contradictory disciplines can be harnessed together to drive reliable growth. Indeed, as we will argue, if the ESP sector is to expand, these disciplines will both need to evolve and work together. A common understanding of consumers and the analytical tools to manage them effectively have allowed the "two sides of the house" to build highly effective businesses.

Why Analytics Matter

Analytics are often viewed by managers as something that is primarily of interest to technicians, whose job it should be to handle the technical details and summarize useful information for management. The complexity of today's analytical tools and – on occasion – the inability of analysts to convey business insights in a relevant way are partially to blame for this perception.

However, as we have shown, analytics can transform an entire business sector. And it is not just the credit card sector that has embraced this philosophy: major corporations, as diverse as Proctor and Gamble, Wal-Mart, and General Electric have embraced analytics as a strategic tool.

Within ESP marketing, there are a number of strategic opportunities that require new ways of thinking – informed by better business intelligence. For example:

- Product Development: Today's aftermarket ESP product is typically the same across all consumer products: repair or replacement for failures after the limited warranty expires. By having such a uniform offering the sector is probably appealing to a fairly uniform segment of consumers. In order to grow, we must have products that meet the needs of a wider set of consumers.
- Pricing: As we discuss in greater detail, pricing decisions often are made at a highly aggregated level and based solely at product-level analysis and generally arise out of the risk analysis. Particularly for ESPs, consumer choice is a big factor; incorporating a good behavioral understanding of consumers and the interplay with pricing will allow the development of products and pricing methods that generate higher perceived value and higher profits.
- Underwriting and insurance: With a few exceptions, the underwriting of risk is managed in one of two ways: companies self-insure and carry the ESP risk on their own balance sheet; or they work with one of a handful of third-party underwriters. There is a much wider range of insurance options available, but most companies do not possess sufficiently credible tools or knowledge to seek out and acquire alternative coverage. As a consequence, many companies may be paying too much for covering their risk – and they are not receiving detailed loss

cost information that is sufficient to help drive significant increased profitability and decision support.
- Self-regulation: While ESPs are lightly regulated today, the general trend in U.S. business is that growing, profitable sectors that sell complex financial products to consumers are facing increasing oversight. In the absence of solid, standardized business methods supported by validated analytical tools, regulators tend to devise their own rules. When a sector can bring its own tools, it is much more likely to be able to guide the regulation process or even be granted the power to regulate itself.

Toward the end of this paper, we provide a number of more specific ideas, presented in relation to our case study.

Today's Warranty Analytics

Companies marketing ESPs certainly all have some analytical capabilities. A brief examination of the most common practices will help illustrate where there are opportunities for improvement.

Risk and Underwriting

Risk management is a key function within ESP management and the industry would not be viable, if it were not for competent analysis of claims and losses.

Most ESP risk analysis is still performed by building loss triangles (see Figure 1), using an analytical method called the chain-ladder method to establish a plausible pattern from historical claims, and using this pattern for future projections. In the triangle, data are aggregated into rows called cohorts (for example, all the contracts sold in a single historical month) and normalized into development periods (the columns), which represent the terms of the contracts (for example, column 1 will represent the first month of coverage).

Actual Historical Claims

Cohorts	Total Contracts Sold	Development Periods → Month 1	Month 2	Month 2	Month 4	....
Jan-07	112	0	2	0	3	....
Feb-07	120	0	1	1	0	....
Mar-07	190	1	1	2	0	....
Apr-07	201	4	0	0	?	....
May-07	123	0	2	?	?	....
Jun-07	135	0	?	?	?	....

Future Claims We Need to Predict

Figure 1 - Overview of Loss-Triangle Format

Most often, a number of triangles will be built for a portfolio, usually around major product categories and contract terms (e.g. one-year and three-year contracts)[2].

It is important to note that when using chain-ladder, this aggregation and normalization of data is not just a way of summarizing data, it is required for chain-ladder to be used effectively. What has been forgotten today is that this collapsing of data and the relatively simple chain-ladder method were developed several decades ago, at a time when computing power was expensive and scarce, and individual contract level analysis was impractical

Given today's cheap and ever-faster computational power, collapsing data in this manner is unnecessary. From a statistical perspective, such aggregation destroys a significant amount of information.

Finally, we feel that the approach has a significant adverse effect on the sector's mindset. It promotes a narrow view solely focused on product and policy, with no regard to customer dynamics or marketing practices. Moreover, if we want to make changes to our products, we are constrained to evaluating potential outcomes against a narrow set of metrics. For example, pricing changes are usually made at a fairly well aggregated level. With a more granular understanding of claims, one could develop pricing schemes that

[2] For a more detailed discussion of chain ladder, refer to the case study later in this paper.

are more closely tailored to the expected loss patterns, with resulting increases in program profitability.

Marketing

More sophisticated marketers of ESP products also have used analytics in their marketing programs. Probably the most widespread analytical tool is the response model, which rank-orders customers on their likelihood to respond to an offer, such as an aftermarket or renewal mailing.

The use of such models generally has a positive effect on response rates and program revenue. And this approach certainly uses the differences between customers and products being covered, thus promoting a more customer-focused orientation.

However, in numerous cases, we have seen the use of these models decrease program profitability even as they bring in higher revenue. For example, such models are usually developed with scant regard to claims risk (something exacerbated by the problems with today's reliance on loss triangles). As a consequence, the models generally recommend a high number of mailings, since each successive mailing may still produce incremental revenue. Yet, such a strategy increases acquisition costs and can also lead to significant adverse selection, in that late responders (those that required a higher number of mailings) may generate a disproportionate number of claims.

As a result, we have seen cases in which claims risks, or loss costs, are ignored completely when making these marketing decisions, with considerable adverse outcomes. Furthermore, even when loss costs are considered, they are almost always contemplated at an aggregate level, causing significant under-performance in program profit optimization.

Here is a case where risk managers and marketers often have worked at cross-purposes. As with the credit card sector before, building a common analytical approach between these groups will develop more growth-oriented and profitable operations.

The Analytical Future

It is fair to say that today's warranty and ESP sectors have an acceptable baseline of analytical tools. However, as we have pointed out, the current orientation leaves considerable room for developing new approaches. We believe that in order to represent a significant step forward and to provide the most information to decision makers, warranty and ESP analytics must have the following features:

- Data must be preserved. Unlike many business processes, warranty and ESP programs do not generate large and rich data streams. As a result it is even more important that we use all of the data that we have available. This will require the use of statistical techniques that can handle more dimensions and hierarchy to the data than current heuristic methods.
- In the absence of a good historical ESP and claim data capture strategy; the ability to aggregate and effectively build a database from several existing disparate sources is critical to immediate success. The ability to understand and generate go-forward data capture strategies will then serve the futures business.

- Methods must accommodate a dynamic environment. One of the great paradoxes of warranty and ESP analysis, at least the way it has traditionally been performed, is that we want both a stable model, but one that would allow us to observe evolving claims patterns quickly. Traditional methods emphasize the stability factor and will mask rapid changes. Even longer-term cyclical changes can be difficult to discern. Fortunately, there are good methods for monitoring both long- and near-term trends.
- Risk management and marketing must have a common understanding of risks and opportunities. We have already discussed the potential conflicts in this area, but companies need to pay attention to building capabilities that unify these disciplines. This requires centralized data storage, processes and analytical approaches that provide metrics that accommodate both needs.
- Analytics must be able to answer a broad range of questions, ranging from strategic to tactical. Too often, analytics are used to solve narrow, technical issues. Managers need to know the answers to big questions and there are often millions of dollars at stake. A mature analytical capability must be as good at answering these strategic questions as at resolving day-to-day details.

In the next section, we provide a detailed example of what is possible by using more sophisticated methods. As we have stated, our intention is not to promote one particular methodology, but provide a benchmark for the types of tools that should be available to all warranty and ESP program managers. Nor would we want to suggest that a single model could hope to solve every business need: analytics is a process, not an artifact.

Analytical Case Study

In this section, we present a detailed analysis of one analytical technique in the ESP sector. For illustrative purposes, we have focused on claim frequency in our case study. However, our approach is equally applicable and is used for severity estimations. Our purpose in presenting this material is two-fold:

- In order to transform an enterprise, and perhaps an entire sector, the analytical discipline must be real. Too much business writing today discusses the importance of analysis, but pretends – through the conspicuous absence of any sophisticated examples – that analytics consists largely of database reporting and spreadsheet analysis. While databases and spreadsheets are part of the continuum of analysis, there is also a need for statistically based models. Managers need not be able to build such models – analysts and vendors can be used for model construction – but they must be able to incorporate the resulting conclusions into tactics and strategies.
- Specifically, we want to provide an example of how we can use all of the information we have, notably ESP contract-level data, to develop better loss cost analysis. Since part of the challenge is moving beyond highly aggregated data analysis, such as loss triangles, the case study we provide shows one approach to achieving improved results. Again, our experience has shown that it is necessary to develop a variety of approaches for specific businesses in order to provide accurate loss cost analysis and forecasting.

Overview of Current Methods

As we have already discussed, most loss cost analysis is being performed at a relatively aggregated level, such as the entire portfolio or major product categories. The most common methodology is to employ chain-ladder algorithms at a general level. Here is a brief overview of the chain ladder method[3].

If there are $\boldsymbol{C}$ cohorts of ESP contracts of term $\boldsymbol{T}$, then $\boldsymbol{N_c}$ is the number of contracts written in each cohort, $\boldsymbol{c}$, and $\boldsymbol{x_{c,t}}$ is the number of claims paid in each lag month, $\boldsymbol{t}$.

We then define the **cumulative, normalized frequency**:

$$y_{c,t} = \sum_{i=1}^{t} \frac{x_{c,i}}{N_c}$$

And then define the vector $Y_t \in \Re^c$ to be a vector of cumulative normalized frequencies for each cohort. In other words, given a matrix of contracts organized into rows representing cohorts (groups of contracts sold during the same period) and columns representing lag months (the month from the effective date of the contract), we can calculate the cumulative frequency of claims over time (refer to Figure 1 to see how data are commonly formatted.).

The chain ladder methodology assumes that the **shape** of the claim emergence pattern is constant. This is equivalent to estimating using the following formula:

$$\hat{Y}_t = \alpha_t Y_{t-1}$$

where a_t is the **development factor** for lag month **t** and is a scalar (and therefore the same for all cohorts). This development factor can then be applied to any cell or group of cells in the matrix to estimate the number of losses for future periods. For example, if we wanted to estimate the number of remaining claims we can expect for our current portfolio, we apply the development factor on all claims we have thus far and we obtain a projection of future claims.

Development factors can be calculated in a number of ways. For example, we can use a weighted average technique:

$$\alpha_t = \frac{\sum_{i=1}^{m} \left(r^i N_i \frac{Y_{C-i,t}}{Y_{C-i,t-1}} \right)}{\sum_{j=1}^{m} r^j N_j}$$

where $\boldsymbol{m}$ is the number of cohorts used in the average (e.g. 35 months) and $r \in (0..1]$ is the decay rate determining how much more recent cohorts count over old ones. If we

[3] As with any methodology, there are a number of variations of chain ladder. We present a fairly standard and basic implementation.

want to treat all historical cohorts as having equal weight then we can set ***r*** to the same value; otherwise we can decide to adjust the relative weighting up or down.

From a technical perspective, the methodology suffers from two main problems:

1. **Bootstrap Problem –** When new cohorts begin, often times there can be no (or very few claims) in the first month or two of the term, leading to a predicted cumulative claim count of zero for the entire term (since the formula is multiplicatively recursive).
2. **Explosion Problem –** Small variations in the initial lag months can lead to large differences in the performance of the cohort over the entire term. This is related to the first problem, but generally occurs the other way around with exaggerated or highly overestimated claim rates. This problem is exacerbated if the claim emergence pattern is back-loaded; that is claims tend to show up later during the contract term.

These problems can be reduced by using a direct cumulative frequency method that does not use the initial behavior of the cohort. Another more complex solution would be to develop a nonlinear function that models the claims emergence pattern more accurately.

However, such corrections, as well as the introduction of the weighting scheme for cohorts (***r***), rely on a considerable amount of judgment by the analyst or actuary. In addition, if claims emergence patterns change significantly in the future, then these handcrafted corrections may no longer be accurate.

Lastly, the authors have spent a considerable amount of time researching statistical methods that could be employed at the aggregated, loss-triangle level. Methods employed have included such advanced approaches as General Additive Models (GAM) and Bayesian models. Such techniques nearly always produce better estimates than chain ladder, but are still sensitive to the sparse data often inherent in the loss triangle – especially at sub-category levels, which essentially reduces our information to the number of claims for predetermined periods.

An Alternative Approach

One clear implication of the limitations of triangle-based analysis, which occur even when powerful analytical techniques are applied, is that one should avoid their extreme data reduction and use all the information we have about contracts and customers.

We first present a random-effect model formulated for individual contract data. This model assumes that for any contract the number of claims is a randomized variable. This is in contrast to traditional methods that assume claims rates to be relatively uniform across any unit of analysis (individual contracts are, of course, excluded from traditional claims analysis for this very reason).

Let $c_{i,j}$ denote the number of claims contract ***i*** filed in development period ***j***:

$$\log(c_{ij}) = \log(l_{ij}) + \alpha + \delta_i + \beta_i * j + \gamma_i * g(j)$$

Here, α is the population level claim rate baseline, δ_i (delta) is the deviation of contract ***i*** from this baseline, and β_i (beta) and γ_i (gamma) are parameters specific to contract ***i***, describing contract ***i***'s claim growth or emerging pattern over the development periods.

There is an important distinction to be considered concerning the inclusion of offset terms. When modeling at the triangle, or even the cohort-specific, level, one can use the number of contracts in cohort ***i*** as an offset. There may be significant variation in the number of contracts written for any given cohort, which is information that may be material to our prediction. For example, let us say we have two cohorts, one with a total of 100 contracts and another with 1,000. If they have claims in the first six months of 10 and 100 respectively, they are equivalent when we include the number of contracts; if we ignored this information, the second cohort would appear to have a claims rate of ten times the first.

Using the number of contracts does not make sense, if we use individual contract data directly. An additional complication is that when using contract-level data, we also do not have an automatic normalization of time periods: for any random contract, we will have a variable number of observational periods; and the data for many of our contracts will be right-censored.

Therefore, we use $l_{i,j}$ to denote the time contract ***i*** is subject to observation in development period ***j***. If we use the same time units for all development periods, then, this variable will be different only for the last development period where the contract is partially observed.

This random-effect model is theoretically appealing in that we can draw population-level, as well as contract-level inferences and forecasts. It also reflects the hierarchical structure inherit in our data. We have two levels of data:

- The lower level data is the time series data for each contract, corresponding to the growth path of a contract over time.
- The higher level of the data is at the contract level, describing the characteristics of the contract, as well as, potentially the contract holders.

Again, in traditional analysis, the lower-level data would have been aggregated into cohort-level data, while many characteristics in our higher level would not have been available at all. Indeed, even in a cohort-level model, we would have assumed that the characteristics of all contracts and consumers represented in a cohort are the same.

Computationally, such a random-effect model may impose a significant burden – even on powerful processors – when we apply it to a large number of contracts, rather than cohorts[4]. In light of this and in order to use as much information as possible, we propose

[4] For example, on a sample of twenty thousand contracts, we have run a Bayesian hierarchical model on a reasonably new personal computer with a dual-core processor and 2 GB of RAM, a configuration that could be replicated by any one. Using the methods described here, the model took about 28 hours to converge. Naturally, the use of more powerful multi-processor computers with higher RAM would reduce computational time.

a latent class model that is structurally analogous to a random-effect model, but is much more practical to estimate and deploy.

Similar to the random effect formulation, for each contract, we have:

- A distribution that models the number of claims in each period. In this paper, we have used a Poisson distribution, but others, such as Negative Binomial or Zero Inflated Poisson, could also be used. Within the same development period, we assume the Poisson rate is constant. Therefore, the addition of observation duration as an offset is justified under this assumption. However, we allow the Poisson rate to evolve over different development periods.
- For each contract, the Poisson process can be characterized by two parameters: the baseline claim rate, and the functional form of development periods, which is assumed to be nonlinear.

Unlike a random effect model, we do not assume the Poisson process is unique for each contract. Instead, we assume that there are latent segments of contracts and the contracts in the same latent segment will have the same Poisson process. Contract's latent segment membership is determined via a multinomial model using contract level information. In a sense, the latent class formulation can be considered as a discrete random effect model.

Instead of assuming contract-specific random effects that are drawn from multivariate continuous distributions, latent class specifies a set of discrete supports. The number and the proportions of such discrete supports are determined, along with the class specific Poisson process parameters, simultaneously within one integrated EM algorithm.

One of the appealing features of a random-effect model is contract-specific predictions. This feature is preserved in our proposed latent class formulation. Theoretically, all the contracts in the same latent segment will share the same Poisson process. However, latent class is a probabilistic model and therefore we will have probabilistic weights assigned to each contract for each of the latent classes. As such, no two contracts with different characteristics will follow the same Poisson process.

Analytical Data Set

For this case study, we have used a random sample of 30,000 five-year contracts from a manufacturer in the marine industry. Unlike the chain ladder method, there is no need to pre-determine any sub-grouping based on product types (for example, two-cycle engines versus four-cycle engines). The development period time unit in this example is defined as a year. We typically use month or quarter in applications. However, here, for illustrative purposes, we wanted to keep things relatively simple. Moreover, our exploratory analysis showed that for this dataset, the claim emergence pattern is relatively flat within the same year and majority of the variations over time are coming from the yearly cycles. In the case where smaller time units are used, instead of using indicators, we would use GAM type of smoothing to accurately capture and characterize the nonlinearity in the emerging pattern.

At the contract level, we have used the following variables:

- Engine types- we used three variables (indicators) to indicate if the engine covered is a 2-cycle, 4-cycle or inboard/outboard. Contrary to prior business hypotheses, we found no evidence that 2-cycle engines are more likely to suffer a failure.
- Contract types-
  - Dealer contract – we found dealer-generated contracts incur fewer, but more expensive claims. This suggests that there may be differences in the characteristics of consumers buying ESP policies or in the aftermarket service practices of dealers.
  - Pro vs. other types of contracts with different features, primarily on coverages.
- The number of days between the service contract starting date and the expiration of the original manufacturer warranty (covering engine) date. Our hypothesis was that customers with a smaller remaining coverage are more likely to use ESP. However, empirical study shows this variable is largely irrelevant.
- Contract starting year, using the hypothesis that the more recent the contract, the more reliable the products and therefore, the less intensified the corresponding claim Poisson process. This hypothesis is validated by our study.
- Engine horsepower- this is the most important variable in the model and the study shows the larger the horsepower, the higher the claim frequency. This result actually ran counter to our initial hypothesis, which was that lower horsepower engines would have higher failures due to technology differences.

We have some owner demographic data and their purchase data on other products, which we could easily incorporate into the model. But we mask this part of the model to protect client confidentiality.

Methodology

The chief methodology we employ in this example is called latent class analysis. This approach assumes that in our data we have what is called "unobserved heterogeneity." That is, there are differences in the wider population that cannot be determined directly from our data. For example, we observed a difference in the frequency in and size of claims between dealer-generated and other types of contracts. In a traditional approach, we might even choose to analyze these types separately. However, these differences may in fact be due to other factors that we can never observe directly (e.g. consumers who buy an ESP from the dealer have a different perception of ESPs than aftermarket buyers).

Another key distinction from traditional methods is that we also assume that we have varying amounts of information for different contracts. For example, we may know relatively little about new contracts. As time goes, the absence or presence of claims adds information. And for those contracts with claims, we have yet more information, including claim amounts, frequency, affected components, and the like.

From this data, we developed a three-segment grouping of contracts using our methodology. The following table reports the segment-specific Poisson processes (*all numbers are masked to protect client confidentiality, but the patterns are preserved*):

	segment	A	B	C
	Size	71.17%	25.80%	2.40%
Poisson Process	intercept	-12.450	-13.456	-10.365
	time (year=1)	0.000	0.000	0.000
	time (year=2)	1.093	3.488	2.092
	time (year=3)	0.731	3.622	2.361
	time (year=4)	-0.097	3.401	2.514
	time (year=5)	-0.146	3.220	2.366

The Poisson process is characterized by the intercept (baseline claim rate) and the functional form of the development periods. For the latter, we have used five dummy variables for 5 contract years. We noticed that all the intercepts are negative. This implies the baseline claim rates are very low which is consistent with the claim frequencies we observe in the database.

Segment A is the largest segment (71.17%) and has a very low baseline claim rate (the second largest negative intercept). The claim rate goes up in years 2 and 3, but goes down in 4 and 5. The magnitudes of the time variable coefficients are small, which implies the claim rate increases relatively slowly over time.

Segment B is the second largest group, but with the lowest baseline claim rate. It has the largest coefficients for the time variables, which implies that claim rates increase quickly.

Segment C is the smallest segment, but with the highest baseline claim rate. The emerging pattern for this group is monotone going up and peaks at year 4 and levels off at year 5.

Notice, for the latent class regression, we only use the time-series portion of the data to construct the Poisson process characterizing contracts' baseline claim rates and the claim emergence pattern over time. The way we introduce contract characteristics as cross-sectional effects is to use these contract feature variables as covariates to predict the prior probabilities of the latent segments. As such, this model is a special case of a so-called concomitant latent class regression model.

The following table provides the multinomial model describing the latent segment priors as functions of the contract characteristics (Segment A serves as the reference category):

Segment	A	B	C
intercept	0.000	53.400	122.670
dealer	0.000	-0.628	-0.178
2-cycle engine	0.000	-0.796	-0.850
Outboard engine	0.000	-1.261	-2.880
Pro type contract	0.000	-0.653	-1.110
manu exp- contract start (days)	0.000	0.000	-0.001
contract years	0.000	-0.270	-0.615
horse power	0.000	0.008	0.029

Holding other factors equal, we can draw the following inferences:

- If a contract is a dealer contract, it is more likely to be in segment A and therefore, incur a modest number of claims.
- A two-cycle engine is more likely to be in segment A and less likely to be in segment C.
- An outboard engine is more likely to be in segment A, and experience fewer claims.
- The larger the contract year (more recent the contract), the more likely it is in segment A, and least likely to be in segment C. Therefore, more recent contracts have lower baseline claim rates.
- The larger the horsepower, the higher probability this contract is in segment C, with higher claim rates.

Notice by using a regression model, we are able to accommodate many contract characteristics variables simultaneously without artificially dividing the contracts into subgroups, which would lead to smaller sample sizes per subgroup, which in turn leads to less reliable conclusions. One could also build a simple Poisson regression model where we treat the contract variables (such as horsepower) and time-series variables (such as yearly indicators) in the same way[5]. One may argue that this would be a more direct use of contract variables, since the latent class model uses the contract variables indirectly: rather than using them to predict the claim rates. We have used them to predict the segments only, which indirectly are related to different claim rates.

We argue that such a hierarchy is justified, since we explicitly acknowledge the hierarchy in the data. From business stand point; the hierarchy in the data reflects the sequence of how the information and data becomes available. Acknowledging such a sequence allows the model to help decision makers at different stages of the process using the information available so far.

- <u>A new contract is sold</u>. At this point, we have no claim data whatsoever for this contract. Using our model, we can use the contract characteristics as covariates to compute the prior probabilities of this contract belonging to any of the latent classes (these latent classes in a sense are risk classes). For example, we can get Prior1, Prior2 and Prior3 by entering the covariate values of a contract into the multinomial model we have listed previously. Once we have these prior probabilities, we can perform two tasks:
    - We can use the prior probabilities to determine which latent class this contract is more likely to be in and what kind of risk class this contract is likely to be part of.
    - We can use the prior probabilities as weights to compute the expected number of claims for this contract. This is straightforward, since, for each segment, we have a regression equation for the Poisson process characterizing the claim rates. The expected claims will be the weighted average of these three rates.

[5] Indeed, the author has fitted a non-hierarchical Poisson model with comparable prediction performance. But latent-class analysis gives us more flexibility in the future, should the underlying processes or heterogeneity change.

We can perform such predictions for different development periods so that the results can be used to compute inflation adjusted total future run off.

- Established contracts. As the claim history of a contract becomes available, we can use the additional information to reassess the latent class membership of this contract. Such a step is called posterior probability based inference.
  - For example, we may have a contract whose prior probabilities are 0.5, 0.4, 0.1 for each of the three segments based on the covariates. Two months later, we have collected one claim for this contract already. We can obtain the posterior probabilities of this contract in any of the three latent classes, given the new information. We will not get into the mathematical details here, except to give some conceptual overview of how we would do this.
  - Based on covariate information, we initially believed this contract more likely to be in class A (0.5 probability). Class A has a modest claims rate and is a slowly emerging segment. However, given the new claims information, we noticed that the actual data from this contract deviates from the Poisson process for segment A. Using Bayes' formula, we obtained the new probabilities to be (0.4, 0.15, 0.45). These new probabilities allow us to obtain new projections for this contract. This process would be repeated for each new time period.
  - Such updating can be achieved by traditional statistical modeling, if time-varying variables are included. For example, if we use the prior claim data as predictors. However, such an approach amounts to recursive modeling, which may cause significant error in sparse data. It can also cause error propagation.
- The latent class also provides a very convenient and useful dashboard-like metric. For example, a snapshot of the latent class distributions provides a concise summary of the overall portfolio risk. Continuous monitoring of the distributions (including new contracts just sold recently) provides a useful mechanism to keep track of the portfolio risk evolution and migration.

Model Validation

As with any regression model, we can use the latent class model to make predictions. For example, using the prior probabilities, we can compute, at the contract level, the total number of claims we expect to see within the observational window (using the offset). We then compare the prediction with the actual observed claims. We noticed that the point estimates are reasonably accurate, which means we can use this model to obtain the portfolio level forecasting.

In addition, we notice that the model provides good lift, which means, if we decided to use this model in connection with an ESP response model, we stand a good chance to identify these few contracts that are responsible for a disproportional amount of claims. It is also a good indication that the sub-category derived loss costs will be an accurate component in the overall profit model that will drive marginal marketing decisions.

Finally, this model is dynamic. For example, the lift table below is based on prior probabilities. If we would use the posterior probability based forecasting, the lift for the top decile would be 690%. This implies that as we learn more about a contract, the model can dynamically take advantage of such information to shape future predictions. It is important to note that the absence of a claim also generates new information that can be used by the model.

contracts	predicted avg. claims per contract	actual avg. claims per contract	% of total contract	cum. % contracts	total claims	% of total claims	cum. % claims	lift
2844	0.5532	0.5545	9.9%	9.9%	1577.00	28.6%	28.6%	289.5%
2925	0.31466	0.31145	10.2%	20.0%	910.99	16.5%	45.1%	225.2%
2875	0.25138	0.25913	10.0%	30.0%	745.00	13.5%	58.6%	195.3%
2882	0.20722	0.22207	10.0%	40.0%	640.01	11.6%	70.2%	175.4%
2839	0.17435	0.17788	9.9%	49.9%	505.00	9.2%	79.3%	159.1%
2924	0.14611	0.15048	10.1%	60.0%	440.00	8.0%	87.3%	145.5%
2882	0.11482	0.10097	10.0%	70.0%	291.00	5.3%	92.6%	132.2%
2881	0.08655	0.08053	10.0%	80.0%	232.01	4.2%	96.8%	121.0%
2881	0.05475	0.05554	10.0%	90.0%	160.01	2.9%	99.7%	110.8%
2882	0.01022	0.00625	10.0%	100.0%	18.01	0.3%	100.0%	100.0%

Since the proposed model is strictly at the individual level and could incorporate not only product and contract characteristics, but also consumer level attributes, we can link this model directly with response models to enable total contract long-term value optimization, rather than simple response optimization.

Reshaping the Business

Our central tenet has been that improved analytics represents a vital competitive capability that can help the performance of individual companies and promote the growth of the sector. We also stated that such analytics must preserve data, accommodate program dynamics, permit risk and opportunity to be evaluated, and provide managers with answers to important questions.

Thus far, we have demonstrated the first two elements successfully. This approach uses much more data than traditional claim analysis; the resulting model in turn leads to better prediction. Even if we accomplished nothing else, we would now have a better method to calculate such metrics as future claims and the ultimate run-off of our in-force portfolio.

And we have also shown that the model itself is dynamic. While we are always constrained to using historical data to build such models, we now have a mechanism that should alert us more quickly, should behavioral patterns begin to change.

This leads us, finally, to a discussion about a few examples of how we can use this type of approach on the more strategic questions.

ESP Pricing

If we had sufficient price variation in the data, we could use that information in predicting the segments. Since the segments represent different risk classes, the model will inform us the degree of adverse selection due to price changes. For example, the model might inform us that customers who are less sensitive to price are more likely to be in a high claim rate segment.

Some further mathematical manipulation can numerically quantify the extent of such adverse selection bias in the portfolio. Similar analysis can be carried out to analyze the pricing impact on consumer response too. Finally, the latent class approach and Bayesian hierarchical approach can be applied in conjoint analysis data to design the best product feature and pricing combinations. The uniqueness of these approaches is that instead of finding one global optimal combination for all consumers, they are capable of finding alternative optimal products and prices for consumers of different characteristics and contracts covering different products.

Lastly, limited warranties, although free to the consumer, are part of consumers' overall purchasing decision and are part of the consumers' perception of price and value. While the broad trend in limited warranties has been to reduce coverage terms, some companies have used limited warranties as a competitive differentiator.

A famous example is Hyundai, which had grown rapidly because of its products' low prices, but was considered to be of relatively poor quality. A sea change in its business came through dramatically lengthening its limited warranty while improving build quality. More recently, Hyundai has redesigned and re-priced its product line to compete with category leaders, such as Toyota and Honda. However, it has retained its long limited warranty as a key element of its pricing strategy, believing that consumers will view their products as better values even if their sticker prices are now much closer to their competitors.

Channel

In our example, our analysis revealed some interesting differences between dealer- and aftermarket-generated ESP contracts. While the analysis we performed is not conclusive, it suggests that there may some opportunities to investigate additional opportunities within the dealer network.

The dealer and agent network is an important sales channel – sometimes the only channel – for many businesses. As a result, companies such as auto, truck, and other manufacturers spend a large amount of money of their dealer programs: sales incentives, drive-to-dealer marketing programs, and many others. ESP programs are often part of the overall mix of compensation available.

What our case study suggests, although even in this case one would need to perform additional analysis, is that increasing attachment rates through dealers might be quite beneficial to the program. If incremental dealer sales of ESPs will hold to the current pattern of lower claims rates (offset somewhat by a higher severity), the program might be able to grow profitably. In addition, the resulting value created in the dealer relationship – higher revenue per sale and more aftermarket service revenue from claims and resulting service work coming to dealer's service bays – may have a beneficial effect on dealer loyalty.

Risk Structure and Cost

The more accurate loss cost forecasting provided by improved analytical techniques such as the example discussed in this paper enables additional program optimization strategies and tactics, as well as more transparent risk structures with lower costs. The authors have found that underwriters and insurance companies often are willing to lower the fees associated with risk in return for the accuracy and piece of mind provided. This includes sharing all or a significant part of the underwriting profits required by any safety margin, as well as sharing some or all of the investment income on the unearned premium.

Since the cost of underwriting often consumes more than fifty percent of ESP retail, obtaining better terms for structure and cost of the risk can yield significantly higher profits. Even for companies that manage their own risk, understanding the patterns in claims permits a more fine-tuned approach to maintaining reserves, recognizing income, and making marketing decisions.

Program Compliance

The more advanced loss cost techniques discussed in this paper help promote SOX and GAAP compliance, since the reserve amortization schedules can be directly derived from the accurate claim emergence patterns. Risk exposure and income recognition strategies also benefit from more accurate loss cost analytical techniques. Since FASB 90 provides no specific, detailed direction on calculating potential exposure, auditors are increasingly demanding that companies be able to provide more sophisticated risk analysis for both limited and extended warranty programs.

In addition, the authors have seen cases in which more accurate ultimate run-off calculations result in additional available income recognition, as emerging trends in loss patterns are much more favorable than the historical average. Conversely, they have also observed cases in which better intelligence has permitted firms to address potential reserve shortfalls more nimbly, thus avoiding painful adjustments to programs that had been underfunded for much too long.

ESP Profit Optimization

Armed with the sub-category loss cost information and forward trended go-forward contract loss cost estimates, Fulcrum has found that significant profit gains can be obtained immediately, and that optimization can be further refined in the future – leading to substantial and sustained gains in profits.

Summary and Conclusions

In this paper, we have presented a new paradigm for the ESP industry. We argue that this sector can gainfully learn from other industries to leverage data, technology and advanced analytics more effectively to stimulate and sustain growth and profitability. We also point out the paramount importance of taking a holistic view of the business process by integrating otherwise discrete and segregated decision components such as risk management and marketing together. Our case study showed how more sophisticated analytics can facilitate such integration and generate substantial business values. We hope our paper will contribute to the on-going research and improvement of this sector and will stimulate more productive discussions in the near future.

Bibliography

Blischke, W. and D. Murthy, “Product Warranty Management”, *European Journal of Operational Research*, Vol. 62, 127-148, (1992)

Caudill, D, J. Littlefield and E. Garman, “A Comparison of Buyers and Nonbuyers of Extended Service Contracts: Ethical Implications for Service Providers”, *Service Marketing Quarterly*, Vol 22, Issue 4, (2001)

Hayne, M. R., “Extended Service Contracts, An Overview,” *Variance*, Vol.1, Issue 1, 18-39, (2007)

Karim M. and K. Suzuki, “Analysis of Warranty Claim Data: A Literature Review”, *International Journal of Quality & Reliability Management*, Vol. 22, No. 7, 667-686, (2006)

Maronick, T., ”Consumer Perceptions of Extended Warranties”, *Journal of Retailing and Consumer Services*, Vol. 14, 224-231, (2007)

Schilling, T. “The Challenge of Pricing Extended Warranties”, *Casualty Actuarial Society Discussion Paper* (1990)

Taylor,G., G. McGuire and A. Greenfield. “Loss Reserving: Past, Present and Future” (Research Paper Number 109, September 2003) Centre for Actuarial Studies, Department of Economics, The University of Melbourne, Victoria, Australia