Paper 347-2010

Detecting Claim Fraud and Improving Product Quality: Case Studies in Reducing Warranty Costs

David Froning, SAS Institute Inc., Cary, NC

ABSTRACT

Manufacturers are utilizing analytics to reduce service claim fraud and increase product quality, yielding warranty cost reductions over 20%. This presentation highlights case studies of manufacturers using SAS® Suspect Claim Detection and SAS® Warranty Analysis to stop claim fraud, identify emerging issues sooner, and get to root-cause faster. In this session, you'll gain a greater understanding of these solutions and take away warranty best practices that can be implemented immediately.

INTRODUCTION

Globally, manufacturers spend over $70B to cover warranty expenses each year (Warranty Week, 2009). That's more that 2% of revenue for most companies, but it's just the beginning. The impact of indirect costs, such as increased government scrutiny, tarnished brand image, reduced stock prices, and lost sales can quickly dwarf the direct costs. The billions lost by companies such as Toyota, Firestone, and Ford clearly show this impact.

In good economic times, these costs are significant, but in the current economic slowdown they can be devastating. With margins tightening and sales dropping, the impact of warranty is even greater. Warranties for products that were sold months or years ago are compared to the declining sales of today. Even if your failure rates stay consistent, your cost as a percent of sales, the only metric stockholders see, gets worse and worse. At the same time, service providers are desperately looking for new sources of revenue. Service providers that were always thought to be trustworthy are now trying to slip things through the system, while the shadier ones are getting more and more sophisticated.

In order to reduce claim costs and increase customer satisfaction, forward-thinking manufacturers are applying automated analytics across the warranty chain. Detecting and preventing fraud, finding emerging quality problems sooner, and accelerating the problem-solving process can reduce warranty costs by over 20% while keeping customers happier than ever. Applying these analytics doesn't take a Ph.D. The power of sophisticated analytics can be surfaced in an interface that is appropriate for your auditors, analysts, and engineers. By building warranty business knowledge into the system and automating much of the work, problem solvers can do what they're meant to do – solve problems.

THE WARRANTY TIMELINE

Figure 2 shows the warranty process as a timeline. Smaller recalls, increased customer satisfaction, and reduced warranty costs are just a few of the positive benefits of shortening the timeline. Every day shaved off of this timeline is a day's worth of production built without that potential defect.

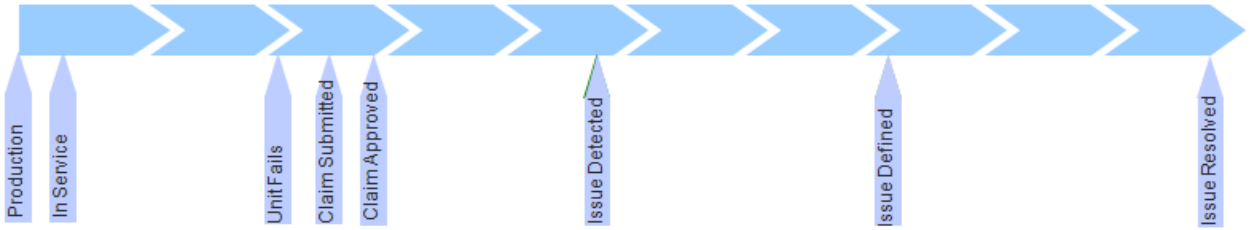

Figure 1. The Warranty Timeline

However, each step in this timeline can introduce lags and inefficiencies that cost time and money. For example, once units are produced, they might sit for weeks or months before going into service. Understanding when a unit

has gone into service is critical for understanding the number of units that are at risk. Unfortunately, many manufacturers don't have complete sales data. This sales lag can be known only months after the fact when a warranty claim occurs. In the interim, they are forced to ignore the lag or use broad estimates for how long units are in this transition. Using analytically derived sales lag distributions, sales lag can be modeled much more precisely. This enables more accurate understanding of failure rates and costs per unit.

Estimating sales lag is just one of many warranty-specific issues that must be understood in order to accurately analyze warranty data. Claim submission lag, usage profiles, seasonality, product maturity, and other warranty-specific issues must be incorporated to truly understand what your field data are trying to tell you. Figure 2 highlights three areas of the warranty timeline where analytics can have the most significant impact: Fraud detection, issue detection, and problem definition. The remainder of this paper will focus on these three key areas.

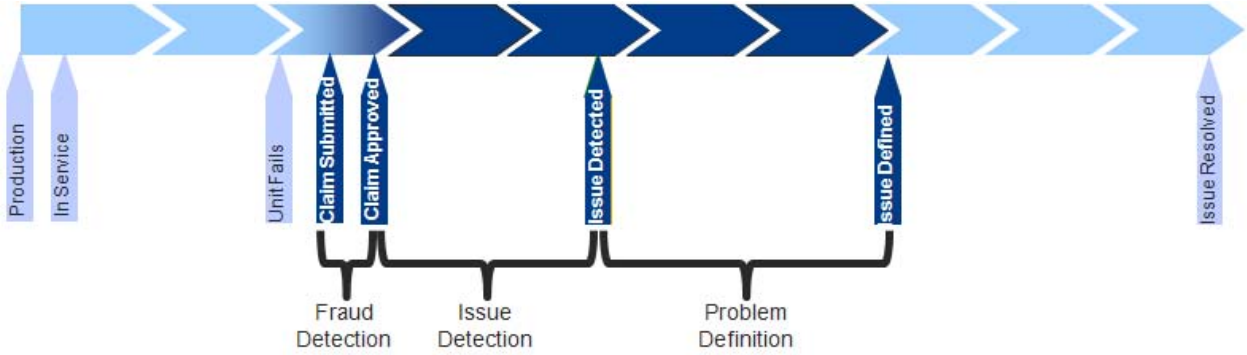

Figure 2. Using Analytics Across the Warranty Timeline

FRAUD DETECTION

According to Eric Arnum, Editor of Warranty Week, 10% to 15% of all warranty claims are fraudulent (Arnum , 2009). That represents over $7B in bad claims that are paid each year, but the costs don't stop there. Replaced parts being sold on the black market, insurance costs, and additional security needs make the problem even more expensive. Manufacturers have put many systems and processes in place to detect and reduce fraud, but current practices are not enough.

True fraud is most likely to occur when a third-party service organization is utilized. However, even in a captive service organization, detecting questionable claims can significantly impact claim costs and customer satisfaction. Identifying areas for training and policy enforcement can reduce the number of parts used for individual repairs and increase the percentage of successful service calls, regardless of the type of field organization.

An additional benefit of fraud detection is that once those claims are removed from the data, issue detection and definition become much more accurate. By eliminating this noise in the data, early warning systems can more precisely detect significant changes in claims rates. Problem definition is more effective because the claims being analyzed are actual customer issues, allowing product engineers to better understand failure modes and the factors that drive them.

Over the years, manufacturers have devised many approaches to combat claim fraud. The following list represents the most commonly used:

- Manual adjudication: Manually reviewing each individual claim is thorough, but extremely resource intensive. In addition, variation is high and patterns are often missed because different people are looking at different batches of claims.
- Automated (rules-based) adjudication: Rules-based systems are now common-place. Almost all claims transaction engines include this type of adjudication. This approach is effective for finding types of fraud that you're already aware of. However, these systems do not learn as new claims are processed. Through trial, error, and telephone calls to auditors, fraudsters quickly learn the rules and work around them.
- Return validation: Requiring servicers to return failed parts does reduce fraudulent behavior (at least for claims with parts), however the logistics and resources necessary to do this on a large scale is cost-prohibitive.

- Field audits: Auditing individual claims and service providers is a very useful way to find patches of fraud in the service network. However, it is often difficult for auditors to know where to focus. Should they look at the highest volume service providers? The claims with the highest cost? Claims with no parts? Focusing these key resources is critical in order to get the most value out of their work.
- Analytics: By scoring claims and service providers based on the likelihood of fraud, analytics can be used to focus field auditors where the most value exists. Analytics can learn from service provider and auditor behavior over time to find increasingly more fraud. Analytics can also identify new business rules that can be applied through the rules engine.

BEYOND BUSINESS RULES

Although business rules and field audits are in place at most manufacturers, they are not enough to stop the majority of claim fraud. One of the most common calls in the claim processing department is a service provider asking why a claim wasn't paid. Although this is a legitimate question, many service providers are using this information to get a better understanding of business rules so that they can work around them next time. In order to effectively attack claim fraud, a broad, analytic-based approach is needed. Figure 3 shows examples of four analytic approaches that can enhance fraud detection.

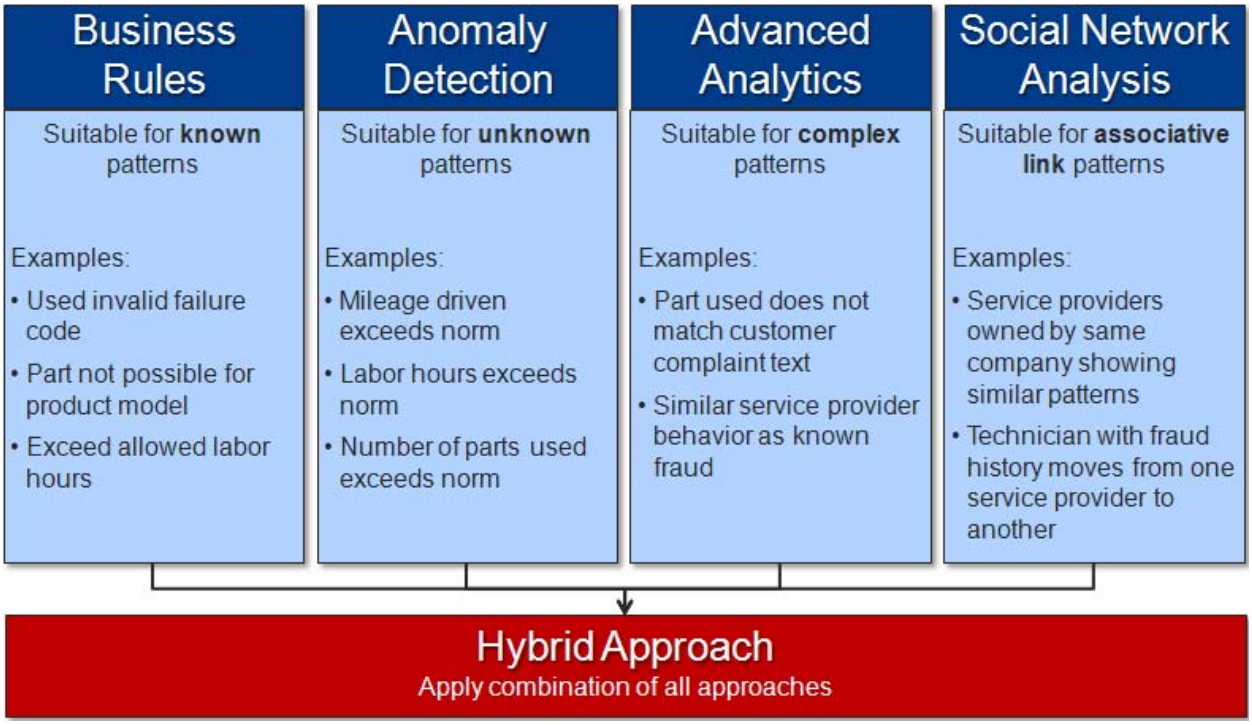

Figure 3. Recommended Approach to Fraud Detection

Business rules are essential for finding the low-hanging fruit. If you know the patterns that you're looking for (for example, invalid failure codes), then business rules are an effective way of weeding them out. When the patterns are unknown, anomaly detection can evaluate behavior compared to peers or across time. By detecting these outliers, anomaly detection can draw attention to questionable behavior. Advanced analytics and predictive models can be used to identify more complex patterns that might involve textual comments, look at multiple attributes simultaneously, or even learn from past decisions made by auditors. Finally, social network analysis can go beyond individual service providers to look for patterns across networks.

Each of these approaches has benefits and is capable of detecting certain types of fraud. By taking a hybrid approach that combines these methods, manufacturers can implement an effective fraud detection strategy. Auditors can be focused more effectively if they are able to use a hybrid approach that ranks claims and service providers based on the likelihood of fraud and calls out the specific reasons that they were flagged.

The following best practices are the basis for an effective fraud detection approach:

- Utilize automated, rules-based adjudication.
- Use multiple text and data-based analytic models.
- Look for anomalies across claims, service providers, and networks of service providers.
- Score claims before they are paid.
- Score claims in real time.
- Rank claims and service providers based on likelihood of fraud and indicate type of fraud suspected.
- Focus field auditors on the right claims and service providers.

BEST PRACTICES IN ACTION

General Electric has implemented a process and system based on many of these best practices (SAS Institute Inc., 2010). By utilizing a hybrid analytic approach, GE was able to catch over $5M in claim fraud each year. Those savings have continued year over year despite the reduction in audit staff. GE is able to do more with less.

GE's one million annual service claims are processed through a third-party transaction system that includes a business rule engine. The claims that pass through the business rules are submitted to SAS every fifteen minutes. These claims are run through a set of thirty-six hybrid fraud algorithms in order to look for suspicious claims and service providers. Claims are scored, flagged, and made available for pickup within a couple of minutes. Claims with low probability of fraud are set to be paid automatically, while the more questionable claims are audited.

By scoring and ranking the claims in near-real time, GE is able to make decisions about claims before they are paid. This allows them to adjust or deny payment before the money changes hands. This approach is much easier than trying to charge back a service provider after payment has already been made. For claims that show suspicious behavior, auditors are now able to understand the likelihood of fraud as well as the specific reason that individual claims were flagged. With this additional knowledge, the warranty fraud detected by auditors has gone up each year. At the same time GE has reduced the number of auditors. The dollars of fraud found per auditor has increased by 340% since the system was implemented.

A COMPLETE FRAUD DETECTION SYSTEM

Flagging and prioritizing suspicious behavior is the core of an effective fraud detection system. However, case management and dashboarding provide additional value that can enhance the effectiveness of the system across the organization. Figure 4 shows a diagram of a complete fraud detection system.

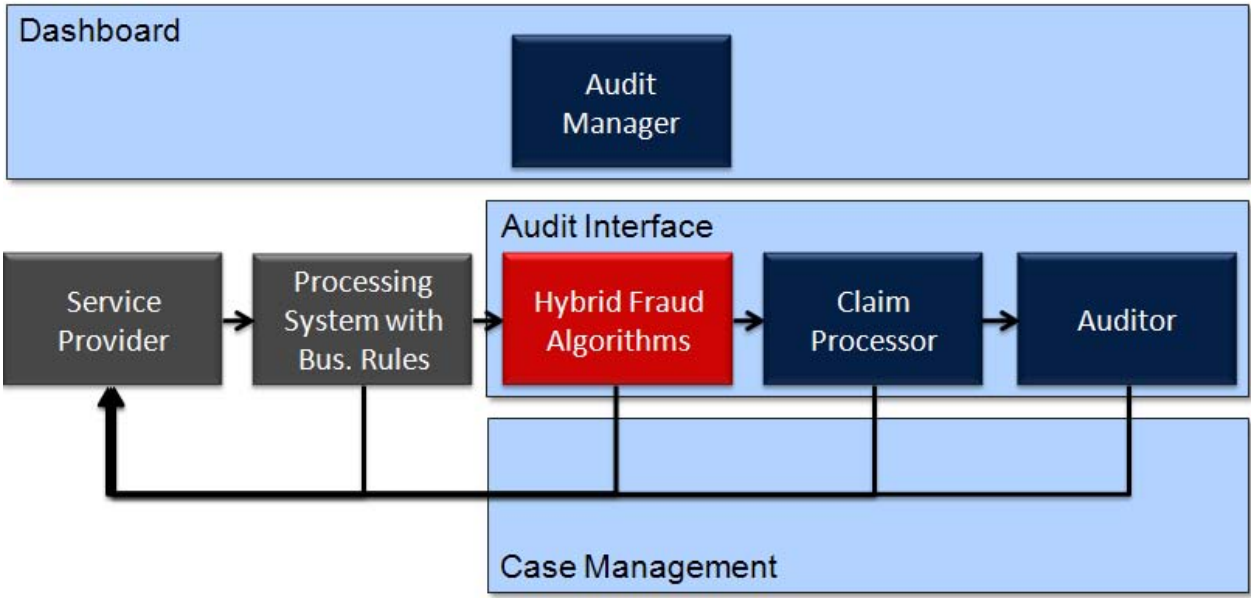

Figure 4. The Complete Solution

Integrating case management into the fraud detection process allows claims processors and auditors to provide more efficient feedback to service providers. It also gives them the ability to build more comprehensive cases before moving forward with a service provider audit. For example, an auditor might build a case for auditing a specific service provider based on their overall monthly performance. The auditor might then add multiple specific claims that will be investigated during the audit. This information is then available as a package when the auditor is ready to go on site to the service provider. The results of the audit can then be maintained in the case management system as a historical record. These results can also be used by the analytics to develop more advanced algorithms for scoring new claims.

Creating a dashboard to measure and monitor the entire process gives management the ability to better understand the process and make adjustments when necessary. The dashboard can be broken down into two basic areas: service processes and audit processes. Key performance indicators (KPIs) for service processes such as volume of claims submitted, percent of claims rejected, percent of claims adjusted, percent goodwill, percent no parts, and others can help you understand service provider behavior and identify service providers that might need special attention. It is also important to monitor the audit process itself. Monitoring KPIs such as number of claims audited, percent denied, amount recovered, and average time to close can inform decisions regarding workload distribution and training.

ISSUE DETECTION AND PROBLEM DEFINITION

Detecting an issue in the field is the first step toward fixing it. Although this statement seems obvious, most recent attention has been on reducing the detection-to-correction cycle. The reduction of this cycle is important, but early detection of emerging issues allows the entire process to begin sooner. Once the issue has been detected, it is critical to use all available data to understand which attributes are driving the issue. Is it a specific combination of supplier, assembly location, and usage? Does the problem only occur in certain geographic regions? Understanding what is driving the issue focuses engineers on which parts to tear down, which processes to examine, and whether a recall might be necessary. Each day saved out of the issue detection and problem definition process reduces the number of potentially defective products that are introduced into the field. With fewer defects in the field, warranty costs will be reduced, customer satisfaction will be preserved, and recalls, if necessary, will be smaller. By introducing effective systems to support these processes, manufacturers are shaving months off of their issue detection process.

An effective issue detection and problem definition system is based on the following three essential components:

- **Data integration:** Warranty and service contract claims must be integrated with sales data in order to understand failure rates. Product information such as build location, model, part suppliers, customer selected options, and important dates greatly enhances the precision of issue detection and problem definition. Additional data sources such as call center data, customer surveys, and field service data can

give a more complete picture of field performance, and in many cases are leading indicators of warranty issues. Bringing these data together and applying warranty-specific business rules (for example, claim lag distributions, usage profiles, and sales lag distributions) are the first steps to an effective identification and definition system.

- **Automated analytics:** Many companies try to identify new issues through brute force. Analysts combing through thousands of claim records, Paretos, and trend charts will eventually identify new issues. However, their valuable time is much better spent fixing problems instead of searching for them. Automated analytics can be put in place to detection significant changes in failure rates, costs, and so on. It is important to look for changes across multiple attributes. For example, production period, claim period, time in service, and usage each yield a different perspective on field failures. Once these issues are identified, analysts and engineers must be able to quickly isolate failure modes and understand what factors are driving the claims. By providing a standard set of analytics that identify patterns within text comments and quantitative fields, problem definition can be greatly accelerated.
- **Communication:** Issue identification is useful only if someone does something about it. For this reason, communication is essential. When the system detects a new issue, it should notify the appropriate engineer immediately. Engineers should be able to subscribe to specific sets of potential issues (for example, brake issues across all product lines and electrical issues on a specific model), allowing them to focus on specific alerts that are relevant to them. During problem definition, communication continues to be important. Updating issue status and sharing analyses among team members reduces wasted time and accelerates the problem solving process. Once the issue has been defined, effective communication allows all of the stakeholders to understand the current status and the expected outcome of the fix.

THE VALUE OF SHORTENING THE WARRANTY TIMELINE

Most organizations believe that they have a handle on issue detection. When asked in surveys, manufacturers often report initial detection within a couple of weeks. After more closely reviewing their processes, the estimates often go up to several months. However, when actual data are analyzed and specific examples are averaged, we have seen that the actual time to detect issues is often six to twelve months. In order to quantify these observations, we began a study of twenty manufacturers from a range of industry segments.

The companies involved in the study participated in one to two day workshops where processes for issue detection and definition were documented and analyzed. These companies were using an assortment of stand-alone analytic and reporting tools (for example, desktop analytics, OLAP, spreadsheets, and mainframe reporting) to detect and define warranty issues. The twenty companies broke down as follows:

- 10 automotive/industrial manufacturers (warranty spend from $.8M to $3B)
- 7 high-tech manufacturers (warranty spend from $3.5M to $300M)
- 3 white goods manufacturers (warranty spend from $18M to $99M)

Once business processes and systems were mapped and analyzed, we worked with the manufacturers to understand the value that data integration, automated analytics, and communication would have at each step of the process. The results of these workshops are displayed in Figure 5.

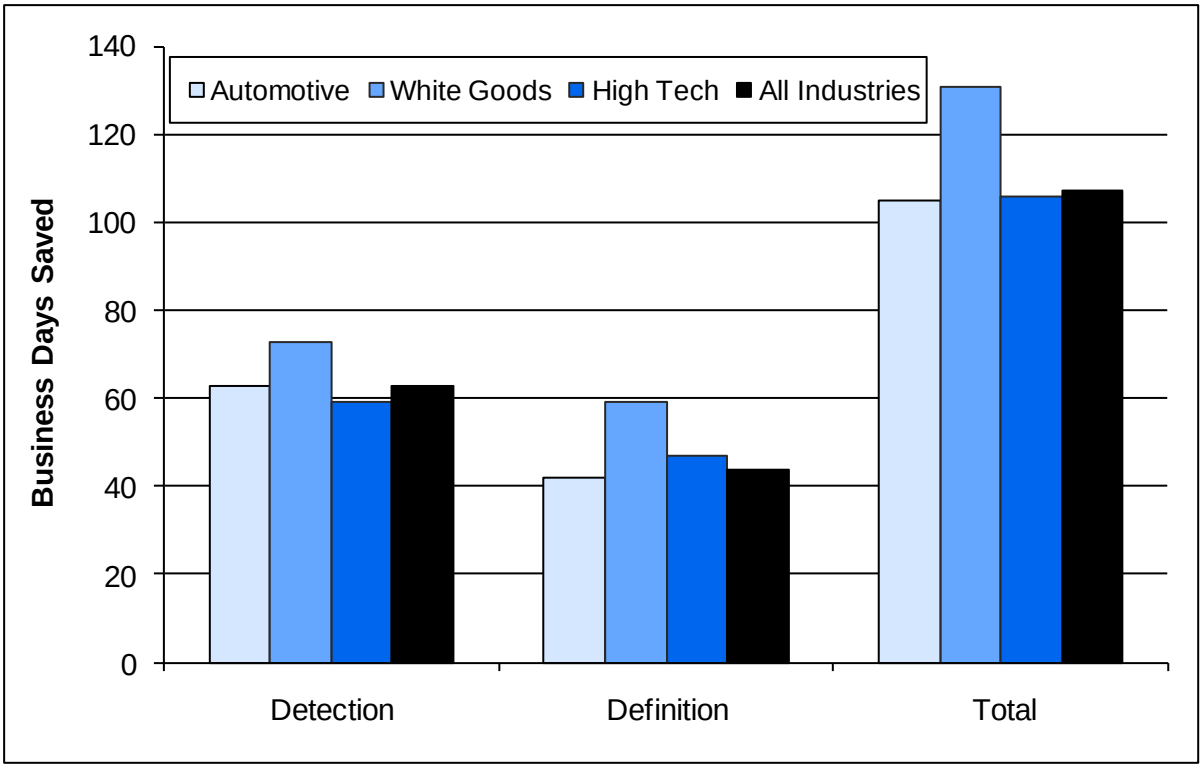

Figure 5. Warranty Benchmark Results

Overall, companies estimated that they would save 107 days off of the identification and definition process. That represented a 58% reduction in cycle time and a 15% reduction in warranty costs. As you can see, these benefits were significant regardless of industry segment. Although the estimated savings were impressive, they were still estimates. For the next phase of the study, we ran actual data through the analytics in order to understand the true benefit.

Five of these companies took an additional step and submitted data to be run through the automated analytic process. Each company provided a set of historical data with known issues and asked us to find the issues without disclosing any details. The historical data sets were run through the analytics iteratively in order to understand the date on which the analytics would have uncovered each issue. The analytic detection dates were then compared to the actual detection dates from the companies. This comparison allowed us to assess the true benefit of having an automated analytic detection system in place. The results from this phase are displayed in Figure 6.

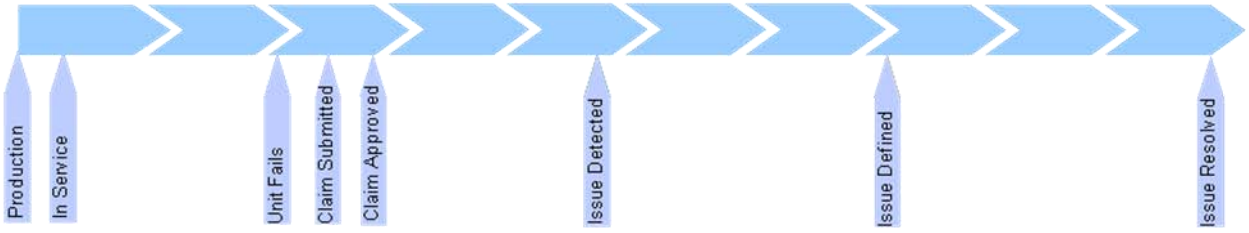

	Detect	Define	Total
Current Process (days)	90	95	185
With analytics	27	51	78
Estimated Savings	63	44	107
Actual improvement*	121		

Figure 6. Estimated Savings Versus Actual Savings

Based on the initial workshops, the estimated improvement in detection time was 63 days. Based on the analysis of real data, the actual improvement was 121 days. That 121-day improvement was even greater than the original estimate of the current detection process, indicating that the actual current process was even longer than the estimate. A 121 day improvement for these companies equates to an 18% reduction in warranty costs attributable to the detection phase alone.

CONCLUSION

In the current economy, keeping customers satisfied and reducing costs are top objectives. Warranty claims dissatisfy customers, and at over 2% of revenue, warranty is a significant area for potential cost reduction. Applying analytics across the warranty timeline can reduce warranty costs by over 20%. Today, manufacturers are applying analytic best practices to deny the 10% of those claims that are fraudulent and reduce costs associated with legitimate claims by 15% to 20%. These savings are realized by improving the tools that problem solvers have at their disposal and automating much of the work so that they can focus on solving problems instead of finding them.

By applying analytics at multiple points across the warranty timeline, companies are seeing that the results are more than just additive. Fraud detection can weed out bad claims that used to be noise in the data. Now issue detection can be more accurate and problem definition can be more precise.

REFERENCES

Arnum, Eric. 2009. Keynote Presentation at Warranty Chain Management.

SAS Institute Inc. 2010. “Putting a Stop to Service Fraud with SAS.” http://www.sas.com/success/GEconsumerindustrial.html.

Warranty Week. 2009. “Worldwide Auto Warranties, Part One.” http://www.warrantyweek.com/archive/ww20090702.html.