

MODEL DEPLOYMENT: **THE MOMENT OF TRUTH**

Analytic Model Deployment
Best Practices & Case Studies

by Robin Way



CORIOS REDPAPER — Papers that get read.™

Key Findings

1. **Many companies are not realizing the full economic potential of their analytic model assets due to lack of adoption (largely the result of failing to deploy them properly).**
2. **Businesses need a common language for understanding and implementing analytic models that they can share with constituencies in information technology and the customer-facing field.**

Companies can directly measure the financial benefits of proper execution and deployment of analytic models.

Building analytic models for the sake of identifying analytics-driven business insight is a worthwhile exercise. However, the moment of truth occurs when the resulting models are used to drive customers' behavior change. To maximize the value from analytics, organizations need robust deployment practices that get models rapidly into the hands of the staff who interact with customers during those moments of truth.

Organizations need better governance and transparency in business processes. This shift will allow information sharing about how a model works, who built it and approved it, which decisions it informs, and who in the field should use it. Firms must also manage their model portfolio as it ages so they are calibrated, enhanced, and replaced as their value decays.

In this research paper, we provide specific steps to help you increase the value of your analytics activities. Author Robin Way draws on work with his firm, Corios, to provide practical lessons from financial services companies who've successfully translated their analytic models into the field, supported their governance, and monitored their performance.

Background

Corios has conducted more than 50 analytic model development and deployment projects, largely within financial services and energy companies, giving us a practical perspective on best practices. All projects have a broad applicability across industries. Analytic model deployment blends the world of analytics and information technology. It requires a shared language and business processes between two traditionally distinct teams.

While some common processes (e.g., SEMMA, CRISP-DM) now exist in the practice of model *development*, few common processes exist in the world of model *deployment*. This absence leads to long project cycles, increased risk, reduced quality, and reduced adoption of model-based insights. The multiple constituencies in analytics, information technology (IT), information security, and the customer-facing field lack the understanding about what is being deployed. To succeed, they need consistent conventions and processes with buy-in from all parties.

The Challenges: Business Adoption of Analytic Models

In order to increase business adoption of analytic models, there is a great deal of work that must occur in addition to model development, and it extends well beyond the model development team.

- First, businesses need to establish connections between model scores and **business decisions**. This connection usually takes place outside the analytics team building the model.
- Second, the **data structures and systems** used by model developers for building models are often different than the one that will be used for implementing in production. Adaptation of the model asset into production should incorporate these differences.
- Third, businesses must be able to easily interpret, assess, and catalogue the model scores and the **changes in scores over time** on an ongoing basis.
- Fourth, to **deploy and execute** these models in a production information technology environment and in the field requires diligence, planning, design, execution, and quality assurance practices that are not commonly adopted by model developers.

The Five D's of Model Deployment

The purpose of this research paper is to provide a set of best practices for analytic model deployment, organized into five phases, which we've nicknamed the "Five D's." They are:

1. **DEVELOP:** Developing and packaging models
2. **DECISIONS:** Tying operational business decisions to model scores
3. **DATA:** Operationalizing analytic model deployment in a specific data architecture
4. **DELTA:** Monitoring the workflow and numeric performance of analytic models in the field
5. **DEPLOY:** Implementing analytic models via a software development life cycle


1. Develop: Model Development and Packaging

The process of analytic model development has been active in practitioner circles for more than ten years. As a result, some standard practices have emerged. In the financial services industry, for instance, it is commonplace to find analytics teams recognized within risk management, marketing, pricing, and compliance functional departments. Analytic segmentation and prediction models are now relatively easy to build given the abundance of analytics tools available on the market and via the open-source movement.

Businesses typically populate analytics data structures with several years or more of historical information, and model inventories can climb to more than 50 models. The SEMMA and CRISP-DM process standards are two relatively similar formalizations of practices that analytics practitioners follow in the construction of most analytic models.

At the time these process standards were developed, few analytics practitioners had much experience or perspective in the skills, processes, and techniques needed to operationalize these models. Most of the focus in the academic and professional literature tended to focus on analytic algorithms, workbenches, and data structures. Researchers emphasized statistical model performance over computational speed or the feasibility of implementing those models in a commercial-class information technology infrastructure.

The SEMMA Process of Model Development



The Five D's of Model Deployment



Figure 1: Comparing the conceptual model development and model deployment processes.

A few years ago, at an industry analytics conference sponsored by SAS® Institute, I asked for a show of hands to the question: “How many of your IT departments support the use of SAS® code in your production computing environments?” Fewer than five out of roughly 80 people in the room raised their hands. This is indicative of the lack of support and shared practices for analytics-focused tools in IT departments of large businesses at that time.

To remedy this, Corios recommends model developers allocate more time and effort to consistently package the model assets they build. This includes documentation of model structure and performance at the point of development, as well as a more formal set of package standards that IT can use to receive and implement these models in a production environment.

Lessons learned from field experience with model development:

- Analytic model developers must understand the capabilities and limitations of the IT production facilities, standards, and field business process before they begin serious model development efforts that may be deployed.
- Take time to educate IT in the tools that the analytics team uses, as well as your assumptions about the field process generating the data and the nature of the data structures you build around them.
- Determine model development tools that are flexible for deployment in a number of contexts and under changing data standards. The tools you use and the models you build should be robust.

2. Decisions: Tying Operational Business Decisions to Model Scores

The typical output of a model, usually called a “score,” is essentially a ranking or estimation of the most likely business outcomes, such as the most likely customer behavior in response to a stimulus (or the absence of a stimulus).

A score is not a decision. A decision is the proactive response of the business to the prospective customer behavior, usually involving the expenditure of resources and the monitoring of the performance of this decision. In a financial services setting, examples of these decisions include marketing, sales and service contacts with customers, debt recovery activities, and financial crime prevention decisions. Outside of the financial services industry, comparable examples of decisions are markdown and assortment in the retailing market, pricing and contract design in the telecommunications industry, preventive care program development, and outreach in the health care provider industry.

Construct a virtuous, evidence-based cycle to make the best decisions. This is characterized by defining connections between decisions, treatments, contact strategies, test-and-learn controls, and execution mechanisms.

Test-and-learn-enhanced campaign optimization for a credit card issuer.

THE TOP LINE NUMERIC RESULTS

Generated 500% return on investment for an optimized marketing campaign.

THE CHALLENGE



The client is a leading credit card issuer. They developed a marketing campaign to shift consumer spending on recurring bills (i.e. cable television, mobile phones, garbage collection, utilities) to the credit card. This strategy builds balances and grows revolving finance charges for the issuer. The campaign offered targeted incentives (that could vary on a customer-by-customer level) to promote more sign-ups. The issuing company wanted to know the right incentives to allocate to the right customers so they could maximize the campaign's profitability.

THE SOLUTION



Corios developed a series of predictive response scorecards and a mathematical optimization routine to maximize the dollar-weighted incremental response rate for the marketing offer. Next, we built test-and-learn controls into the optimization scenario to measure the impact of all champion and challenger offers. This allowed the card issuer to adapt the campaign in the future since it was a fairly new offer.

THE RESULTS



On an eligible universe of 8.6 million accounts, we targeted slightly more than 700,000 offers. This generated over \$700,000 in recurring monthly incremental revenue on a one-time marketing investment of \$695,000 (representing contact costs and targeted incentives). Since customers stayed with the program for at least six months (a relatively conservative expectation), this resulted in over a **500% ROI** for the issuer.

Optimized next best offer for a retail bank.

THE TOP LINE NUMERIC RESULTS

Increased each sale per contacted customer by \$22. This resulted in a \$3.5 million increase in gross sales per month.

THE CHALLENGE



The client is the retail marketing group for one of the top 50 banks in the world, as measured by 2013 assets. The bank issues hundreds of offers every month via 30+ campaigns through its network of branches, call centers, teller machines, web, email, and direct mail channels. The bank's segmentation and analytic sophistication used to rank the top offers for each customer are well developed. However, marketing leadership knew that ranks and scores weren't enough to make the final decision about which offer to assign to each customer. They also had to balance product sales goals, cross-product halo effects, over-time contact strategies, and offer-cost and margin contribution alongside customer likelihood to respond. In short, which offer should they give to each customer in order to grow profitably?

THE SOLUTION



Corios implemented a mathematical campaign optimization routine to allocate the margin-maximizing offer per client for each monthly campaign wave. After only four months of design and implementation, the bank released their first optimized campaigns.

THE RESULTS



Comparing year-over-year campaign results, the optimized campaigns executed over a two-month period produced a **net increase of \$22** in gross sales per customer. In the aggregate, profitable sales growth averaged **\$3.5 million** in monthly incremental campaign-driven income. The bank attributed sales growth to an increase in both customer response rate and the increased profitability of accepted offers. The bank's marketers are now self-sufficient in executing optimized campaigns on roughly 80% of their campaign lead volume, several times every month.

Decisions, continued...

Since decisions based on models consume resources (i.e. budget, channel capacity, customer attention span, people in the field, etc.), firms should ensure these investments are prudent and the best decision has been made given the information available at the time. This prudence gives rise to A/B testing. However, most businesses in a practical setting have more than just two decision alternatives to make. Test-and-learn practices are the more mature option for businesses that are willing to make the tradeoffs between short term stimulus-response measurements and longer-term testing strategies that maximize the value of information.

Many organizations have developed rules-based approaches for capturing the decisions to be made in the field. Unfortunately, most rules-based approaches fall short on measuring decision trade-offs, optimizing for the ideal profit-maximizing outcome, and evaluating short-term versus long-term benefits of each decision. This recognition within leading analytics organizations has led to the adoption of mathematically-optimizing rule engines. These optimization engines place all possible decisions on a level playing field, and allow the business to observe the tradeoffs of competing decision strategies directly and simply as measured by a condensed, internally consistent set of objectives.

At design time, business analysts should specify the goal-seeking objective along with a set of rules regarding resource cost and availability, product goals, and customer contact strategies. From there, they can save and execute the collection of these rules as a scenario. Once analysts have designed and run several such scenarios, they can determine which scenario is the champion and use this scenario to allocate offers and treatments to all eligible customers. Alternative designs can allocate this treatment-allocation strategy as an on-demand scoring approach (read on for more coverage of on-demand scoring).

Lessons learned tying model results to decisions:

- The best rules systems implement business rules based on judgment and experience, as well as rules that prescriptively advise the SME on the trade-offs between alternatives in pure dollars and cents terms.
- The best performing offers and treatments are not the ones you will issue tomorrow. Instead, they will be the offers that your organization has refined over multiple waves of disciplined, rigorous trials paired with conscientious measurement.
- Business teams that commit resources to decisions will require the most convincing about the veracity of the analytic model results. Ensure your analytic story about the findings and recommendations of the model are well suited to these groups. Explain the nuances of customer behavior in terms that resonate with your audience. Find a willing listener from that constituency to help you develop your explanations in advance of the big presentation.

3. Data: Operationalizing Analytic Model Deployment in a Specific Data Architecture

When the field of analytic model development arose, the typical data architecture for model development emphasized the construction of offline, highly customized model development datasets, usually in analytics-centric platforms. This was driven by the preferences, skills, and experience of the model development team.

Since this analytics development-driven data structure didn't often coincide with the data architecture used by the rest of the business, model scoring activity was consequently owned by the model development team. Scores were delivered in batch cycles from the analytics team back to the IT team for loading to the warehouse and other production systems.

This arrangement worked when there were fewer than 10–20 models, the refresh cycles for new model scores fit into that batch window, and the data exchange mechanism for data file export-import was consistent with the use of these scores in the field. However, these conditions were not acceptable or sustainable for many leading organizations—they began to desire larger pools of models, more rapid refresh, and innovative methods of model scoring.

The retention risk of analytics professionals is a continuing challenge, whereas capital investments in data appliances and memory tend to stay level.

Four major data platform technology innovations have been introduced to meet these challenges: in-database analytics, in-memory analytics, over-time analytics, and on-demand scoring. These advances are compelling because the trend in cost for data storage, data processing, large-scale memory, and inter-application messaging have not increased at the same rate as the trend in cost for skilled, experienced analytics professionals. The latter resource continues to be in scarce supply compared to demand. The retention risk of analytics professionals is a continuing challenge, whereas capital investments in data appliances and memory tend to stay level.

In-database analytics are motivated by moving the analytics into the data warehouse. This is a contrast to the more conventional process of extracting data from the warehouse, moving it to the analytics platform, scoring records, and then moving the scores back to the data warehouse. Companies can focus the scope of in-database analytics on the construction of the model development and model scoring data tables, the construction of the model, and the model scoring routine. In-database analytics saves clock time, increasing customer responsiveness when rapidly-changing customer behavior (and consequently the data refresh cycle) demands moving from monthly to weekly to daily refreshes.




Figure 2: The evolution of data platform technology.

Massively parallel model development and deployment for an investments broker.

THE TOP LINE NUMERIC RESULTS

Reduced model build time to days instead of months; scored models in hours instead of days; increased model performance by 10–20% by incorporating transaction details.

THE CHALLENGE



The client is the data sciences team for one of the largest investment brokerage companies in the U.S. The brokerage has built more than 50 treatment response scoring models for their customers, which they use to advise their financial advisors in the field on the ideal recommendations for each client each month. Due to the massive array of information on customer positions, trades, and online transactions, it took the data sciences team months to acquire the data used to predict customer behavior.

Their platform could not physically accommodate all the desired information. When they had constructed a model to score in production, it required multiple days to execute all the model scoring routines for all their customers. The brokerage wanted to dramatically reduce the cycle time to develop new models and refresh their scores for customers. This would provide their financial advisors with daily updates on the best treatment for their customers, allowing them to grow their most profitable relationships.

THE SOLUTION



Corios designed and implemented a massively-parallel model training and scoring routine for three of the data science's team's models using the SAS® High Performance Analytics platform on a commercially-available database appliance. By dramatically modernizing the analytics development and deployment process, compute tasks that previously required hours to run now take a few minutes or less on billions of transactions.

THE RESULTS



The brokerage's data sciences team can now incorporate the full range of customer transaction behavior in their models. This resulted in a **10–20% increase** in model predictive performance. The brokerage can score their models within a nightly window, allowing financial advisors to get nightly updates instead of waiting for weeks or months for the next best customer treatment.

Data, continued...

The most compelling case for in-database analytics is when models directly take advantage of customer-transaction level details, capturing changes in customer behavior. When the volume of transaction details is in the tens of millions to billions of customer records, then extract-score-import cycles typically cannot scale computationally to the refresh cycle desired by the business.

In-memory analytics are a more recent technological innovation. This technology emphasizes loading data and the analysis function into memory in order to dramatically cut processing time. While some open source analytics platforms have always leveraged in-memory analytics for model development, they tend to be weak on the model deployment front. Hence, they aren't the focus of this discussion.

In-memory analytics are indeed blazingly fast, with processing times in the milliseconds. Nonetheless, analytic processing has always been relatively fast when compared with the gray-matter processing required by the analyst to specify hypotheses, render judgment, and balance science with art. The most compelling case for in-memory analytics appears to us to be when super-fast processing is embedded into a business analytics process that properly leverages it. One context is real-time numerical visualization. Another is real-time distributed business decision support.

In contrast to the most common set of cross-sectional classification and estimation models used by financial services companies, **over-time analytics** focuses on changes in patterns over time. Techniques that leverage this analytic approach include time trending, transaction sequences, and time series forecasts. Technologically, over-time analysis is distinct from in-database analysis because over-time requires physical sorting of the records in the analysis data file (a capability that is not possible with distributed databases). Over-time analysis is also distinctive because the dynamics of over-time model development and scoring are often more sensitive to small changes in the model specification than cross-sectional models. In addition, they require more intensive validation by analytics developers.

We've found that everyone's definition of real time varies as widely as from "20 milliseconds" to "nightly" with the greatest concentration of expectations in the sub-second to sub-five-minutes mark.

On-demand scoring is the final innovation worthy of note. While some business teams like to refer to this activity as "real time scoring," we have found that everyone's definition of real time varies as widely as from "20 milliseconds" to "nightly" with the greatest concentration of expectations in the sub-second to sub-five-minutes mark. On-demand probably deserves to refer to the customer's expectation given the context, with online securities trading authorization as an example of the fastest expectations, and in-branch offer customization via recommendation engine on the slow end.

There are several reference architectures and plenty of highly-customized approaches to delivering on-demand scoring, but the most popular and most flexible option is the web services approach. Web services are used for inter-application communication on a broad scale, and the skill set for designing interfaces that use them is relatively plentiful.

Optimized customer acquisition profitability forecasting for a credit issuer.

THE TOP LINE NUMERIC RESULTS

Reduced cash flow forecast time from weeks to minutes. The ideal product design enables growth of \$10–\$50 in net income per acquired customer.

THE CHALLENGE



The client is the acquisitions profitability team for one of the top credit card issuers in the U.S. (as measured by cards in force). The credit issuer launches 50+ new credit offers each month and they seek profitable growth by identifying the ideal card product for each segment of their target market. Their process of cash flow forecasting for each pool of cardholders required days of painstaking data gathering and spreadsheet model fact-checking to ensure quality projections. If an executive asked for a new spin on the cash flow projection, it often took weeks to re-work their spreadsheet models. They wanted to spend less time managing inputs and data and more time on defining the best card offer to build for their markets in order to actively grow profits.

THE SOLUTION



We implemented the Card Acquisitions Forecasting Engine (CAFE) process (developed by Corios) for this card issuer. CAFE is a process for predicting profits and is based on matching current performance for a pool of accounts with projected expectations for financial metrics that drive cash flow and risk-weighted profitability. The credit issuer's analysts can now design and run each cash flow forecast in just a few minutes - for both the consensus forecast and for economic stress test scenarios. This allows their executives to make product design decisions that pave the way for profitable growth.

THE RESULTS



This card issuer is now able to blend the benefits of running cash flow forecasts in a few minutes per scenario, evaluate the sensitivity of profit to economic stressors, and build the optimal offer design to maximize risk-weighted margin growth. This allows their company to confidently identify the top 10 card offers for each pool of acquisition prospects. For selected prospect pools, incremental net income estimates range from **\$10 to \$50 growth** per cardholder through the first 2 years of the customer lifecycle.

Real time price simulation and optimization for an auto lending portfolio.

THE TOP LINE NUMERIC RESULTS

Reduced P&L pricing model run time from hours to seconds, while simultaneously expanding portfolio segmentation and return on assets (ROA) by two orders of magnitude.

THE CHALLENGE



The client manages an automobile lending portfolio for non-captive dealers around the country. Their business-as-usual pricing model used only nine risk tiers and calculating the net present value margin contribution and return on assets for a single set of assumptions took roughly 15 minutes for each scenario. This kept the pricing team from developing more sophisticated pricing strategies and optimizing for the ideal price structure on a dramatically more segmented portfolio of dealers and loan prospects.

THE SOLUTION



Corios developed an automated approach for pricing profit & loss (P&L) modeling. This tied into the client's expanded hierarchical segmentation strategy and underlying data warehouse records. It allowed their dealers to create hundreds of segment-specific P&Ls in a few minutes or less. For each P&L, we developed an optimization approach to identify the ROA-maximizing discount rate on each segment of loans, effectively running hundreds of scenarios on each segment. A simple user interface was implemented to give the non-technical financial analysts in the team the ability to run their own assumptions and scenarios without technical assistance.

THE RESULTS



The portfolio financial analysts are able to **optimize pricing terms** at a more precise, highly segmented point of entry, in a **dramatically reduced time frame**. A conservative estimate of financial impact for this enhancement is to grow ROA by several percentage points at the portfolio level.

Data, continued...

Think of a web service interaction as a call and response. A point of sale system—such as a call center automation system—will send a message (the web service request) to an analytic server. It provides information such as the client ID, the type of model to score, and the attributes about the customer needed to calculate the score that aren't available elsewhere (and preferably, all the information about the customer that is needed to calculate the score). Once the analytic engine receives the request, it calculates the score, generates additional information such as treatment or offer configuration details, and sends this along other relevant information back to the call center system (this is the web service response). When properly designed, the call-and-response system can occur in a sub-second interval, depending on how much work the analytic server needs to do to calculate the requested information.

Lessons learned from experience with innovative data architectures include:

- Each of these innovations brings a great deal of promise, assuming that the fit-to-purpose of the technology to the business process is a good one.
- In-database analytics is typically a good fit when the organization maintains very large data files (i.e., customer records in the tens of millions and/or transaction records in the hundreds of millions) and has already made a commitment and/or investment in commercial-class data appliances. The value of in-database analytics applies equally well to the functions of analytic database development, model development, model scoring, and model deployment.
- In-memory analytics is likely a good fit when the organization has many different choices to make across many scenarios, needs a low-latency connection to a large pool of historical data, prefers to work with the data visually, and needs to embed substantial analytic richness and power in the analysis process. Good examples include identifying the drivers of product profitability or conducting scenario-based uncertainty analysis.
- Unfortunately, over-time analytics seems to receive fairly limited visibility despite its usefulness. Time series analysis and forecasting have been available techniques for decades but in the financial services industry, at least, they appear to have been relegated to the economic forecasting department. More recently, delinquency and loss forecasting tied to stress testing have emerged driven in part by the Federal Reserve's mandate for capital liquidity analysis. This type of analysis is simply good business, and should have broader adoption by more companies. These techniques, for instance, support analyses like customer lifetime value (by way of survival analysis) as well as new product adoption (by way of Bass adoption curve analysis).

- On-demand scoring is the least well understood technique in the bunch, due in part to its technical complexity, but these barriers are disappearing fast. In addition to dealing with the technology issues, businesses also need to construct quality assurance programs for measuring the flexibility and robustness of real time scoring mechanisms to situations they didn't anticipate or capture in the historical data used to train those models. In order to increase adoption, analytics teams should invest in significant outreach to explain how their models work to the field teams that will use them.

4. Delta: Monitoring the Workflow and Numeric Performance of Analytic Models in the Field

Recently, a sharp uptick in the demand for analytic model inventory tools and practices has emerged in the financial services business. This has been driven in part by recent legislative and regulatory pressure. Many of these organizations had been following a relatively informal approach of tracking their model inventories with spreadsheets, but found that keeping up with periodic updates of their model assets, ongoing review and approval, and robust monitoring of performance became too much to track in such a fast and loose manner. Furthermore, regulators and internal reviewers expect to see evidence of well-built practices and processes for transparency and governance, with which informal tracking systems are incongruent.

Businesses must be able to document the scope, ownership, quality assurance, approval and sign-off, and lifecycle of a model asset. These lifecycles are often characterized by a standard workflow within an organization consisting of multiple steps of review at multiple levels in the organization. If a model fails a review, an iterative process of findings and limitations identification, resolution, documentation, re-review, and re-approval follows. Companies should track all of this activity so there is a good audit trail and reporting mechanism, tailored for the process. A typical workflow is illustrated in **Figure 3** below.

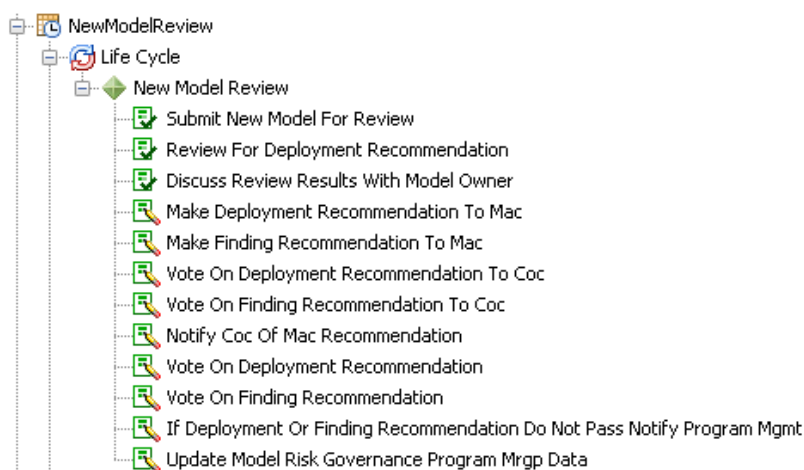


Figure 3: Model approval workflow from representative financial services company.

Delta, continued...

Businesses also need to understand the model's insights, including the interpretation, defensibility, and identification of the best business actions to take based on a model score. The Bank of International Settlements published guidelines in 2005 related to credit risk model validation and measurement. This provided detailed specifications of metrics and measures for ensuring sound internal rating systems for lending policy. These measures are now well understood, and these standards have broad applicability for any predictive modeling effort in the banking industry and elsewhere.

As an example, Credit Risk model guidelines require that model results should demonstrate sound measures of the stability of the population of customers for which loans are being rated, the performance of a model in terms of discriminating between good and bad historical outcomes, and calibration of the model to fit the historical event rates accurately.

The examples in **Figure 4** and **Figure 5** are excerpted from a 10-page standard model performance report that Corios developed for the banking industry. This particular report tracks a model's predicted scores, expected outcomes, and actual outcomes. This is an example of model calibration.

The two charts below (**Chart A** and **Chart B**) indicate how the original calibration of the model in time period 0 (indicated by the black indicator **Chart A** and the first time period on the x-axis of **Chart B**) has worsened over time, by shifting into the red zone in both charts, and the time period in which the greatest dilution occurred (i.e., time periods 2 through 4 in **Chart B**).

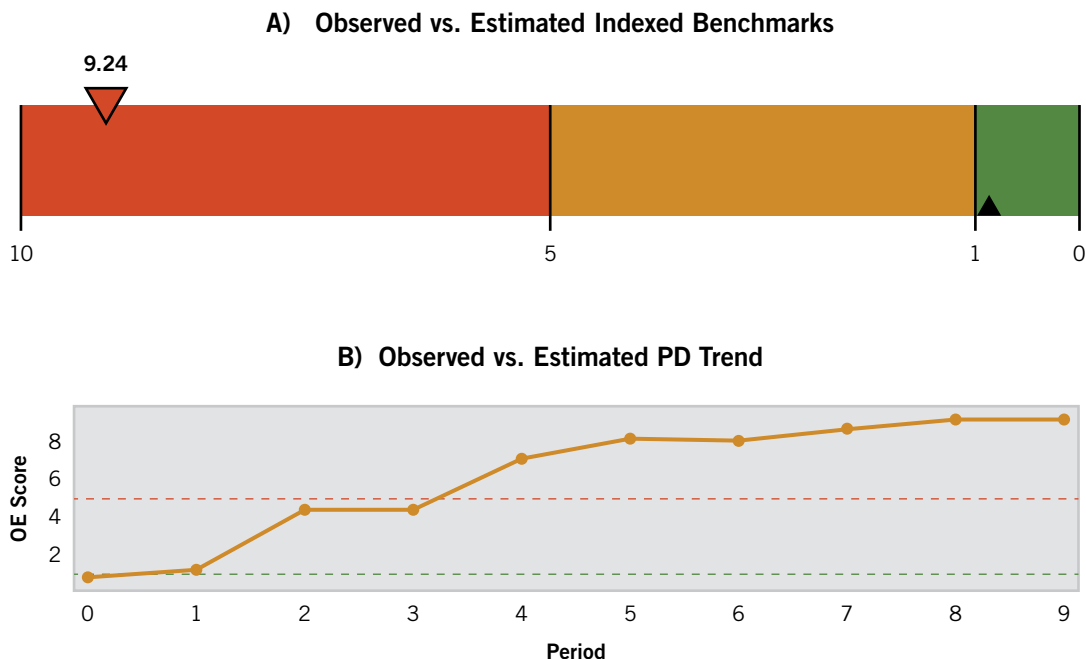


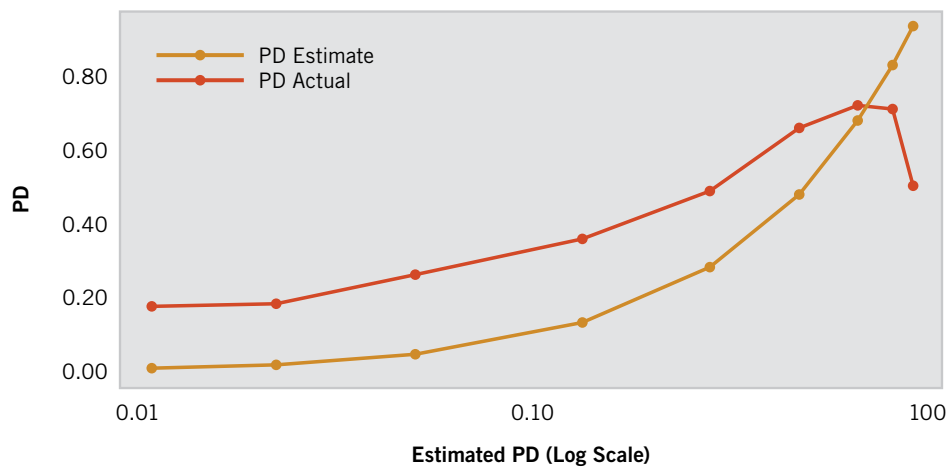
Figure 4: Sample model calibration report for a Probability of Default ("PD") model on a loan portfolio.

Delta, continued...

The pair of charts below in **Figure 5** indicate where in the risk distribution the greatest loss of calibration occurred. In this case, it occurs in the most risky part of the customer distribution. Specifically, on the right hand side of **Chart C**, where the expected and actual scores diverge in opposite directions, as well as where the Observed-Expected (OE) score is highest, as reflected by the green line in **Chart D**.

Combing the model workflow reporting with the model performance reporting, and active mechanisms for keeping these reports up to date, keeps all constituencies in the organization well apprised of the state of the businesses' model inventory.

C) Estimated vs. Observed PD



D) Observed PD and OE Index by Score Points

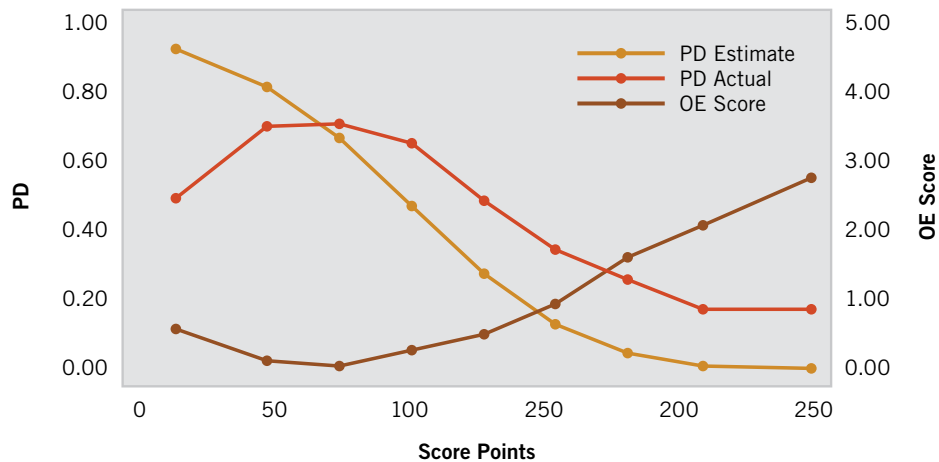


Figure 5: Drill down detail for a sample model calibration.

Model governance for a property and casualty insurer.

THE TOP LINE NUMERIC RESULTS

Migrated 15 underwriting models into automated model governance system in 2 days.

THE CHALLENGE



The client is a property and casualty insurer that struggled to manage the workload of model performance review for their underwriting and actuarial models. Like other financial services companies, they struggled with using spreadsheet based tools for model inventory management but had not developed or implemented a fit-for-purpose tool that would satisfy their internal and external review requirements.

THE SOLUTION



Corios implemented an automated system based on SAS® Model Manager at the firm. This was enhanced with custom-built reporting components for workflow and model inventory. The actuarial team was trained in a hands-on lab setting. By end of the two-day session, all the actuaries had implemented at least one model each.

THE RESULTS



In only two days, never having seen the new automated system before, each of 15 actuaries had **successfully implemented** automated reporting for at least one of their underwriting models.

The table on the next page (**Figure 6**) provides some directional guidance from within the financial services industry to characterize demand for model inventory management. These cases are based on Corios' direct experience in supporting these firms. Hence, this evidence is only directional and not representative of the industry as a whole. These companies maintain model inventories that are fairly substantial in terms of the number of model assets being managed, compared to similar inventories five years ago. Governance, regulatory compliance, and performance tracking factors are neck and neck as the most compelling drivers of adoption of these practices.

Client	Banking	Insurance	Brokerage/ Other	Number of Models	Model Deployment	Governance, Regulatory, & Compliance	Performance Tracking	On Demand Scoring
A	✓			70		✓		
B		✓		20	✓		✓	
C	✓			10		✓	✓	
D	✓			70	✓	✓	✓	
E			✓	5		✓	✓	✓
F	✓			300			✓	
G	✓			100	✓	✓	✓	
H			✓	50				
I			✓	10	✓			✓
J			✓	5		✓		
K		✓		15			✓	
L	✓		✓	200			✓	
M	✓			50		✓		
N			✓	300		✓	✓	✓
O	✓			50		✓	✓	
P	✓			200	✓	✓	✓	
Q	✓			300		✓	✓	

Figure 6: Model lifecycle management reference cases.

Lessons learned from experience with model monitoring and performance indicates:

- Senior leadership in financial services firms have identified the lack of visibility into model inventories as a source of regulatory pressure that needs to be met now, rather than “a few years from now.”
- Leading organizations have developed detailed processes and workflows for tracking the model asset through its stages of development, review, approval, and retirement. They have also staffed the organization properly in order to deliver on those processes.
- Regardless of the presence of regulatory pressure, leading organizations have realized that models are worth hundreds of times (and sometimes more) over their development cost. They should have the proper mechanisms in place to manage these assets, comparable to managing other assets such as human resources or physical plant.

5. Deployment: Implementing Analytic Models via a Software Development Life Cycle

Many businesses struggle with implementing analytic models in production systems. This is often because analytics and IT tend to have few shared conventions or practices, despite each functional unit having relatively technical backgrounds. The companies that tend to excel in developing shared conventions and practices have cross-trained key staff in both functions. To excel, analysts should develop and maintain essential IT capabilities such as relational and distributed data modeling, application development, unit and system testing, version management, authentication systems, and program management skills.

The process of migrating a predictive model into production systems should follow a well-known practice used by IT teams, namely the software development life cycle (SDLC). This process consists of specified phases, each consisting of tasks, with each task being iteratively developed, tested, documented, and hardened.

When applied to predictive modeling, we recommend adding an additional phase on the front end, called “Explore,” followed by the traditional Development, Test, and Production phases. The rationale for the additional work is that the analytics team should pursue the heavy emphasis on data exploration, hypothesis generation, and evaluation as normal. This work should not be overlooked for its value in the credibility of the resulting model.

At the outset of the Explore phase, there is a great deal of uncertainty in the design and composition of the eventual model. As the process moves through Explore to Development and then on to Test and Production, the final model asset to be implemented in production is much like the phrase attributed to Albert Einstein: “Everything should be kept as simple as possible, but no simpler.”

An example of the SDLC applied to predictive model development and deployment appears in **Figure 7** below.

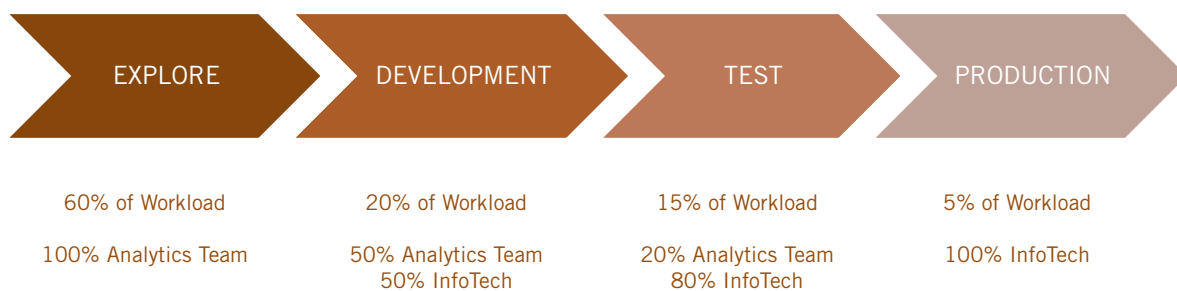


Figure 7: Software development life cycle applied to analytic model development and deployment.

Lessons learned from experience with model deployment provide best practices including:

- As much work as possible in model development and scoring table build should be moved into the database or data appliance. This will produce the best data development speed at production run time.
- Where possible, select a single platform for model development and deployment in order to reduce the quality assurance challenges that arise when translating from one platform's code base to another.
- A key transition occurs in the Development phase, from validating the model training process into the model scoring process. The chief purpose of Test and Production is in validating and hardening the scoring task itself.
- Track model performance for both model stability and whether accounts that receive on-demand scores are fundamentally different compared to the average or typical account. This should occur both on accounts that get scored on-demand during the day, as well as on all accounts on a periodic basis.
- Improve quality assurance by using peer review, version and change management practices, and structured expectations for model documentation.
- Improve model interpretation by adding explanatory factors to the calculated score. These factors should explain to a person in the field why a given score is high or low so they can use the information to render a more effective application of the score in the decision the field person makes regarding the account.

Conclusion

Properly deployed models create measurable and significant financial benefits for businesses of all sizes and across all industries. However, many companies are not realizing the full economic potential of their analytic model assets due to lack of adoption and failure to deploy them properly.

The fact is that your models should do more than improve understanding; they should drive business performance. The moment of truth occurs when analytic models are deployed and are used to drive customers' behavior change. In order to make the most of their model assets, businesses must develop the common processes for communicating and integrating model deployment practices across multiple constituencies in analytics, IT, information security, and the customer-facing field.

In this RedPaper, we're taking a first step towards crafting a common language for understanding and implementing analytic models. We've drawn insights from practical lessons, and identified specific steps that will help you improve model performance and efficiency, increase the value of your model assets, and improve your analytics ROI.

Interested in learning about how Corios can help you get more from your models? Give Robin a call at **503.295.1685** or email him at **deploy@coriosgroup.com**.

About the Author



Robin Way

Robin Way is the Lead Faculty Member for Banking at the International Institute of Analytics, and is the founder and CEO of the management analytics consultancy, Corios. He has over 25 years of experience in the design, development, execution, and improvement of applied analytics models for clients in the credit, payments, lending, brokerage, insurance and energy industries. Robin was previously employed with SAS® Institute's Financial Services Business Unit as a managing analytics consultant for 12 years, in addition to another 10+ years in analytic management roles for several client-side and consulting firms.

Robin's professional passion is devoted to democratizing and demystifying the science of applied analytics. His contributions to the field correspondingly emphasize statistical visualization, analytical data preparation, predictive modeling, time series forecasting, mathematical optimization applied to marketing, and risk management strategies. Robin's undergraduate degree from the University of California at Berkeley and his subsequent graduate-level coursework emphasized the analytical modeling of human and consumer behavior.

About Corios

Corios is a leader in the discipline of management analytics consulting focused on helping clients across industries to bridge the gap between their data and their business decisions. The company's custom offerings provide analytical solutions for risk management and compliance, marketing, pricing, and big data initiatives. Corios' solutions have identified business improvements worth hundreds of millions of dollars for corporations across North America.

Corios believes that data-driven techniques are the key to making profitable business decisions, and that data should be simple, approachable, and implementable for both analysts and decision makers. The Corios team is persistent in their efforts to promote transparency and common process across the analytical discipline.

To learn more about how Corios is bridging the gap between data and business decisions, please visit coriosgroup.com.

Model Deployment:
The Moment of Truth

Copyright © 2013

Published by Corios
133 SW Second Avenue
Suite 202
Portland, Oregon 97204



CORIOS

WE TELL THE **STORY**
NUMBERS CAN'T

All rights reserved. Except as permitted under U.S. Copyright Act of 1976, no part of this publication may be reproduced, distributed, or transmitted in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

Visit our website at: coriosgroup.com

