White paper

Collections optimization

Table of contents

Introduction to optimization in collections

Debt collection has become a frontline profit driver for financial organizations, with a huge impact on the bottom line. As a lender, you need to maximize collections — which means you need to know your customers better, segment and target your approach more effectively, and automate as much as possible.

Optimization is a common term in everyday English, typically used to imply that something is being carried out in the best way possible. Inherent in the definition is the idea that there are multiple ways of doing things, and the most effective has been identified.

With regard to optimization in collections, this paper focuses on the actions that can be carried out on individual customers, and how these actions can be chosen to optimize overall performance in the collections environment. Back-office or process management required in the day-to-day running of a collections operation aren't discussed.

The following examples are collections situations in which customer-level decisions can be optimized:

- A lender can choose between sending a letter or making a phone call to a customer.
- There are many customers.
- There is a restriction on the number of calls that can be made.
- The objective is to maximize debt collected.

Optimization achieves an uplift in performance against traditional techniques by:

- Considering every combination of actions.
- Understanding the tradeoffs between the different actions (which are forced by constraints).
- Choosing the best set of actions to fit within the constraints.

When applying optimization techniques to the collections process, Experian® has observed Key Performance Indicator (KPI) improvements between 5 percent and 30 percent.

A simple example helps demonstrate this. The table below shows the expected debt collected and cost for three different customers and three different actions:

	Action	No action	Letter	Phone
Customer 1	Debt collected	$80	$100	$120
	Cost	$0	$2	$5
Customer 2	Debt collected	$30	$90	$60
	Cost	$0	$2	$5
Customer 3	Debt collected	$20	$250	$300
	Cost	$0	$2	$5

Choosing the best action for each customer results in $510 of debt collected, at a $12 cost.

If, however, the cost is constrained to maximum of $10, the best combination isn't obvious until all combinations have been calculated (it's $490 at a $9 cost). Similarly, if one can choose only one of each action, maximizing performance is driven more by the differential (or tradeoff) between actions, rather than the individual customer response to the actions themselves. The traditional hierarchical approach would be:

1. Choose the best return from the best customer first.
2. Move to the next best action for the next best customer.
3. Continue until a constraint is met.

This would result in $430 collected at a cost of $7 ($300 is picked first, then $100, then $30), missing the optimized maximum of $470 with a cost constraint of $7.

In a mathematical context, optimization is the term used to describe the task of maximizing a function, subject to constraints. This could be maximizing debt collected, subject to operational constraints such as call capacity and operational cost. These are often described as nonlinear integer programming problems, and there are a number of tools that help solve these problems.

This paper describes:

- The collections process.
- Core elements of optimization.
- The opportunities to optimize decisions made at various points within this process.
- Deployment of solutions.

One key question that hasn't been answered explicitly: Why do I need to optimize instead of just using better models? The key reasons for this are:

- The combinations of different options.
- The complexity of competing KPIs (e.g., bad debt, resource, cost, roll rates).
- Management of multiple constraints.
- Management of multiple customer-level dimensions (e.g., propensity to self-cure, propensity to right-party contact [RPC], propensity to pay, arrears balance, exposure).

This is why a structured approach with the backing of a suitable software tool is essential.

Optimization solutions from Experian

Experian offers optimization solutions that can be applied across the Customer Life Cycle, including various stages of the debt management process. The proven Marketswitch Optimization® solution uses patented technology that lets clients design, analyze and execute mathematically optimal strategies accounting for competing goals and operational constraints, as well as individual customer needs and preferences. Experian's decision optimization platform includes two software solutions: Marketswitch Optimization and Strategy Tree Optimization. Marketswitch Optimization develops optimal, data-driven customer strategies at the individual customer level. Strategy Tree Optimization relies on the same technology to help clients quickly design optimized decision strategy trees to deploy into an existing business rules engine.

Marketswitch Optimization® Strategy Tree Optimization

Credit risk	Marketing
Originations	Prospecting
Customer management	Up-sell Cross-sell
Fraud	Customer management
Debt Management	

Using Marketswitch Optimization and Strategy Tree Optimization, Experian has successfully completed projects in the various stages throughout the collections process, including those described in this paper, achieving significant benefits for clients. The resulting collections strategies have been implemented as individual optimization strategies as well as strategy trees.

Core elements of optimization

Optimization can improve performance for each of the collection processes, using all the tools presented in the Appendix. The challenge is to deliver improved performance in a simple, rapid fashion. Experian has completed a number of projects with clients across the process and identified demonstrable benefits.

In each example, there are a number of core requirements to deliver an optimized solution. These are:

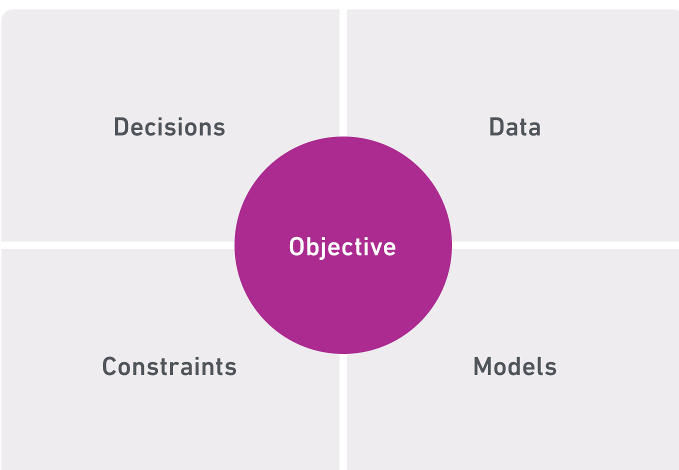

Objective (goal function)
What are we trying to optimize? Ultimately, for most projects this will be minimizing ultimate expected losses, net of cost (in other words, minimizing net impairment). But there may be subtly different themes — correlated, for example, to maximize cure rate or minimize roll rate.

Decisions
These are the actions that can be carried out on a customer. It could be a path, a time to call or a number of grace days. "Walk before you run" is always good advice, and it's no different with optimization. When considering decisions, it's best to start optimizing decisions that have occurred in the past rather than optimizing things that have never been tested before.

Data
Broadly, there are three types of data:

Operational data — Available data that can be used to make a decision (e.g., data that exists within the workflow management tool) and exists at a customer level.

Analytical data — Data that can influence a decision (e.g., outcome data that helps drive a model), but can't be used in an operational tool. This also exists at a customer level.

Portfolio or action-level data — Examples include the cost of a letter, average call duration and impairment ratios.

The operational and analytical data will include items such as:

- Demographic data.
- Behavioral data.
- Scores.
- Bureau data.
- Previous action data.

This paper isn't designed to assess the merits of these different pieces of data, but just to use them where they are predictive.

Models

Customer-level estimates for particular KPIs that influence the goal function will add a great deal of discrimination, and hence performance improvement, to an optimization problem. These customer-level estimates (e.g., propensity to self-cure) can be provided by:

- Existing scorecards that are predictive (e.g., behavior or collections score).
- Variables that are predictive or directly influence the KPI (e.g., balance).
- New scoring models (e.g., propensity to self-cure, given action A). Note that these new models are typically response-style models designed to highlight differences in outcomes based on different actions, rather than the traditionally labor-intensive credit risk or collections scorecards. As such, they could be uni-, bi- or multivariate models, and a payoff between time versus benefit often can be estimated.

Constraints

The constraints are the things that simply must or must not happen — for example, we must have a minimum of $x collected, or we must not have more than 20,000 outbound calls made this month. Constraints also can occur at lower levels of detail, such as the activity time period (e.g., 1,000 calls between 10 a.m. and 12 p.m.), and at a descriptive level (e.g., average debt per debt collection agency must be the same).

Opportunities to optimize within this process

It's attractive to consider optimizing across the entire end-to-end process for every single customer. This is an onerous task, however, both in terms of data requirements and analytical resources, and the robustness of such a solution can be questioned. Experian found that by taking discrete parts of the process and studying and optimizing them, significant benefits can be obtained at each stage — typically 10 percent to 20 percent improvement in target KPIs. The following examples describe some of these discrete points.

Precollections strategies

A recurring theme with all optimization opportunities is that no two problems are alike, because there are so many points at which variation can occur. This includes the obvious ones, such as the definition of a KPI-like impairment or the value of a call volume constraint, but also the "shape" of the problem itself, which is often determined by the actions.

Consider a call-barring decision in the telecommunications environment, for example. A decision can be made to fully/partially/specifically bar a customer. This can be made for fraud/risk reasons based on:

- A percentage/absolute amount.
- Over and above a set limit/previous usage.

And the customer can be called/SMS sent/directed to collections on inbound calls.

As such, flexibility and scalability of any tool used for optimization is a key requirement.

Experian developed a call-barring strategy for a telecommunications company. The key elements can be seen below.

The tradeoff between risk (propensity to enter collections, driven by behavior score) and balance (driven by spend) is the key to achieving the objectives and meeting constraints.

Customer spend behavior is particularly interesting in the telecommunications environment. Different customers present different behavior, but sometimes for the same causal factor. Consider a taxi driver taking a long weekend break versus a married couple with children taking a long weekend break … from the children. The former will probably have a drop in spend, the latter an increase.

Thus, when setting a limit on monthly spend it's more important to understand the individual customer's variability of spend, rather than the average or point estimate of spend. A limit can then be set to only be hit in x percent of cases.

This is analogous to the observation that the standard deviation for an individual customer is far greater than the standard deviation around the mean estimate of a customer. Models are built on thousands of customers, and as such they are predicting the average of those thousands of customers. The prediction for the group will be very accurate — but one shouldn't set a limit based on this accurate estimate of the mean, because the individuals within that group will vary, not just among the group, but also individually.

For example, consider a customer with a mean spend of $100 and a standard deviation of $20. The probability that their mean will be greater than $120 next year is 5,000-to-1. The temptation is to set a limit of $120 and expect only one in 5,000 customers to be barred. But on average a customer will exceed this twice a year — resulting in almost 1,000 bars!

Performance improvement

There are a number of swap sets with an optimized approach versus a traditional champion strategy that sets limits by risk. A few examples help illustrate where benefit can be achieved:

- Low risk, low spend, low variability = low limit (versus previous high limit).
- Low risk, low spend, high variability = high limit.
- High risk, high spend, low variability = medium/high limit (versus previous low limit).

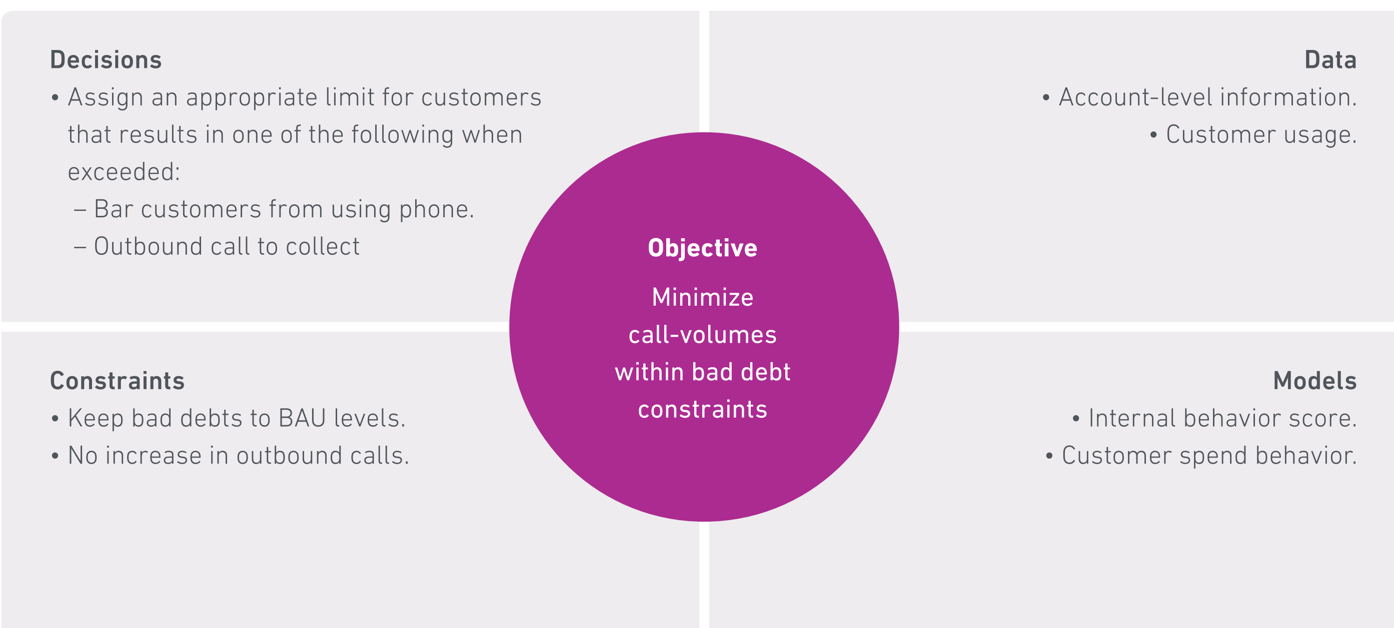

A reduction of more than 20 percent in the volume of bars can be achieved using optimization, with no impact on volume or value of debt entering collections.

Self-cure strategies

The key goal in self-cure strategies is to avoid spending resources on customers who will pay their arrears even if no action is carried out on them. As always, nothing is certain, so there is a payoff between the propensity of a customer to self-cure versus the cost of contacting them. The template structure of this decision can be seen below.

Modeling again forms a key part of the decisioning process, and consideration should be given to developing separate models for different days past due (DPD) or groups of days past due — e.g., propensity to cure, given no action, given x DPD. This obviously has to be traded off against the benefit that will be achieved on customers who wouldn't have self-cured but will make some payment if contacted.

It's also crucial to ensure that the objective is correctly aligned. "Maximizing debt collected" will simply result in customers likely to self-cure being contacted. At the very minimum, an "incremental debt collected" (versus no action being carried out at all) is required.

Interestingly, in this situation it may be worth having a minimum number of outbound calls carried out for the first challenger strategy. This ensures that the optimized solution doesn't leave call center staff unused. It also acts as a safety check against trusting the models implicitly. Once a new strategy has been demonstrated to improve on the champion, consideration can be given to reducing staff or employing them more effectively elsewhere.

Path optimization (early/mid/late)

Path optimization is the holy grail of collections optimization. Identifying the right communication for each customer and executing it at the right time, via the right channel, with the right terms is close to this holy grail … as long as this is then followed up with the correct set of communications, all working seamlessly together.

The requirements of this perfect path (per customer) are onerous, as it isn't enough to understand the outcome of each combination. The interaction of the actions also must be understood, which is hard to assess, especially since:

- Relationships between actions can be very strong or very weak.
- Actions are often close together, which makes cause and effect hard to ascertain.

Decisions

- Grace days.

Data

- Behavioral data.
- Demographic data.
- Bureau data.
- Collections and behavior scores.

Objective

Minimize net impairment

Constraints

- No increase in bad debt.
- No increase in operational cost.
- Minimum number of outbound calls.

Models

- Propensity to cure given no action given DPD.
- Propensity to collect given action occurs at x DPD.

- There is typically a lack of testing (high-risk customers rarely receive no action).
- There are many possible outcomes, which aren't all clearly good or bad.
- Sequences of actions can be long — more than 90 days.

A more manageable challenge to optimizing paths is to simplify the problem to the point that it can be optimized. There are a few ways of doing this, including:

- Identifying the best action for a customer and moving that action to the front of the path.
- Optimizing smaller discrete periods of activity, such as week by week.
- Treating the path as a single action.

Experian carried out the last of these methods with a UK retail bank, and the methodology is described below.

In this example, there were three potential paths that could be assigned to a customer. It's easy to think about them as high-/medium-/low-risk paths. The champion approach assigned the high-risk path to high-risk accounts, and so on. There was also a challenger strategy, which tested some accounts of one risk through a different path.

This champion/challenger test was useful in that it allowed customer-level models to be designed for each path. Thus, one could estimate the expected impairment for each customer, for each potential path that could be assigned to them. It's equally possible to develop customer-level models to estimate the number of actions of each type that will occur for a customer, if they pass down a particular path.

So rather than being weighed down in the complexity of the movement of a customer through this path, evaluating which action is causing which reaction, etc., this methodology simplifies the problem to its core components:

- Which path should be applied?
- What is the expected impact on impairment?
- How much resource is required?

The improvements (of up to 5 percent without creating new paths) generated by optimizing in this way are coming from two key sources:

- The payoff between risk and balance:
  - High risk and low balance = low-risk strategy, hence less cost.

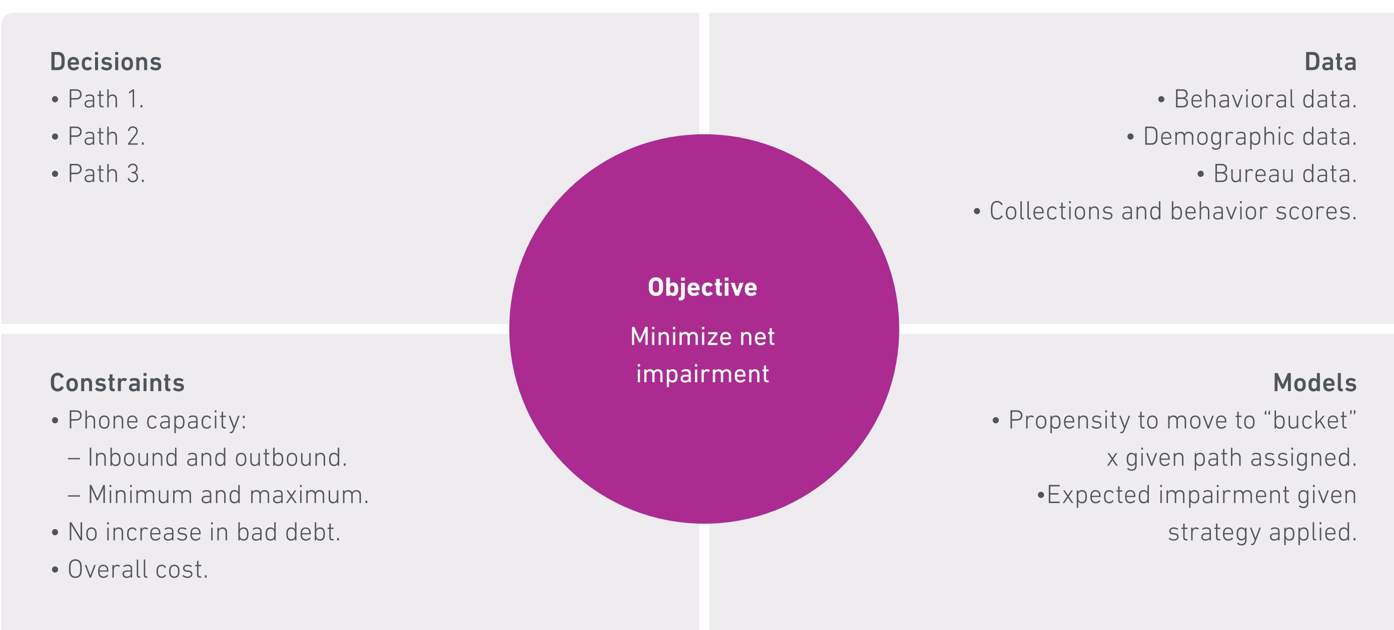

- The size of difference in outcome for individual customers for different paths:
  - Small difference between high- and low-risk strategy = apply low-risk strategy, hence low cost.
  - Large difference between high- and low-risk strategy = apply high-risk strategy to minimize impairment.

Constraints also can be applied by the different channels, as there are estimates of volume of customers receiving each communication based on the strategy applied. This can allow for cost and resources to be managed effectively.

It should be noted that deployment of optimized strategies should be no different from normal challenger strategies — implementing on a small test group is recommended to allow for data gathering, clear test results and risk mitigation.

Debt collection agency (DCA) optimization

Once the decision has been made to use an external agency to collect the debt, assuming there is more than one to choose from, optimizing between agencies becomes a very attractive, quick-win option. This is because different DCAs perform differently on different types of customers. This can be driven by:

- Geography.
- Customer type.
- Specialism in unable to contact/trace, etc.
- Actions taken (phone versus door knocking).

The other advantage of this decision point is that collections departments will often have a wealth of data available, as DCAs will have been tested against one another, with a random set of accounts being passed to each. This is perfect for developing propensity-to-collect models that are specific to each DCA.

Again, the core elements required to develop the optimized strategy are outlined below.

Clearly an internal agency can be included in the optimization problem to minimize cost, as can second placements of debt and even debt sale.

The results of such a problem can be very rewarding, with improvements of up to 10 percent observed. Equally satisfying is that optimization delivers a win-win:

- Each DCA gets debt that is more aligned with their collections ability, so they collect more and earn more commission.

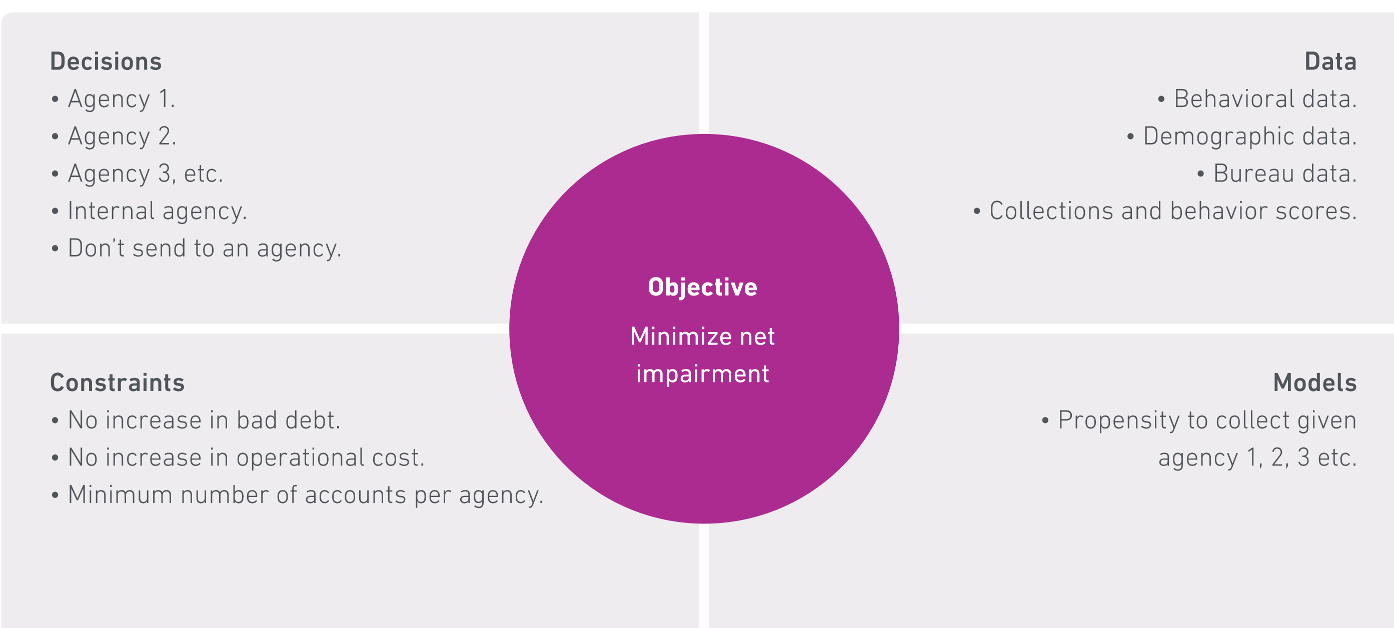

- The collections department sees an increase in total debt collected, improving the core KPI.

In this example, consideration also can be given to batch processing DCA files, as the visibility that a strategy tree delivers isn't necessarily needed by either the collections team or the DCAs. As such, a further improvement in performance can be gained by using individual (customer-level) optimization tools.

Recoveries optimization

Recoveries optimization is similar to DCA optimization, in that there is typically just one action and a simple 0-to-1 outcome. It differs in that there is limited challenger testing between strategies; data is limited, and because of smaller volumes, complexity is lower and constraints fewer. This isn't to say that it can't be optimized, but most uplift in performance will be driven by taking the most profitable action, if there is one.

The elements of a recoveries optimization opportunity are outlined below.

Dialer optimization

The examples given above all lend themselves to being implemented in a decision agent or a workflow management tool that can score and segment customers.

However, optimization also can be applied to dialer activity. Again, this isn't optimization of the dialer process, but rather optimization of which customers get called and when.

On the next page is a schematic of the "dialer pyramid," outlining how the ultimate goal of obtaining payment from a customer is driven by a number of outcomes that need to occur.

Fortunately, there is typically a lot of data available to model each of these events, although not every individual outcome needs to be modeled. There will be correlation between outcomes, so one or two key drivers can be modeled, such as propensity to get RPC and propensity to get successful payment.

The approaches to using dialers vary significantly between collections departments, emphasizing once more the need for flexibility in a software tool. Outlined below are the key requirements for taking an optimization approach into the dialer environment.

Significant benefit has been observed in the optimization of dialer activity. The improvement comes from understanding the differential between different actions and selecting

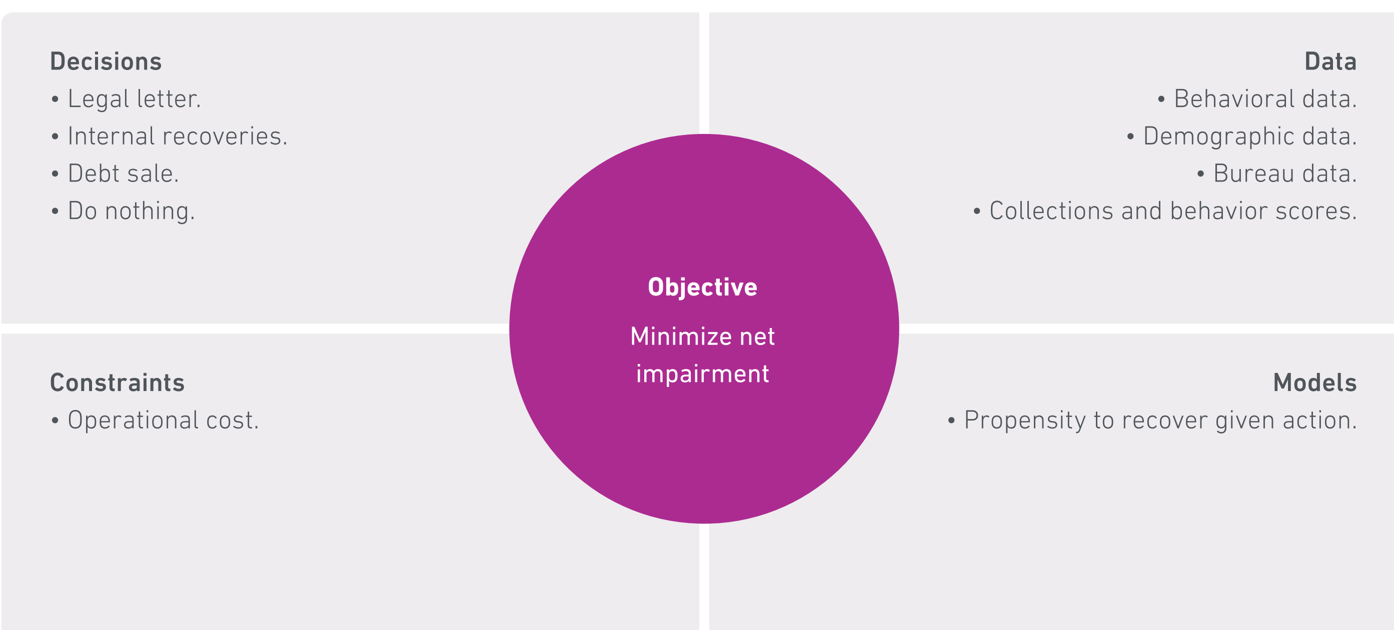

customers appropriately, rather than just choosing the best action per customer. For example, consider the following two customers:

- **Customer A:**
  - Propensity to collect (a.m.) = 70 percent.
  - Propensity to collect (p.m.) = 50 percent.
- **Customer B:**
  - Propensity to collect (a.m.) = 40 percent.
  - Propensity to collect (p.m.) = 10 percent.

A hierarchical approach (with constraints on call volumes) will pick a morning call for customer A and an afternoon call for customer B. But swapping these around results in a higher total response. When considering the impact of balance, RPC rates and so on, it's clear that significant uplift is achievable by measuring each accurately and optimizing them.

Similar to the DCA optimization, individual optimization can be more readily considered for this, and where modeled previously, uplifts in performance can be up to 35 percent.

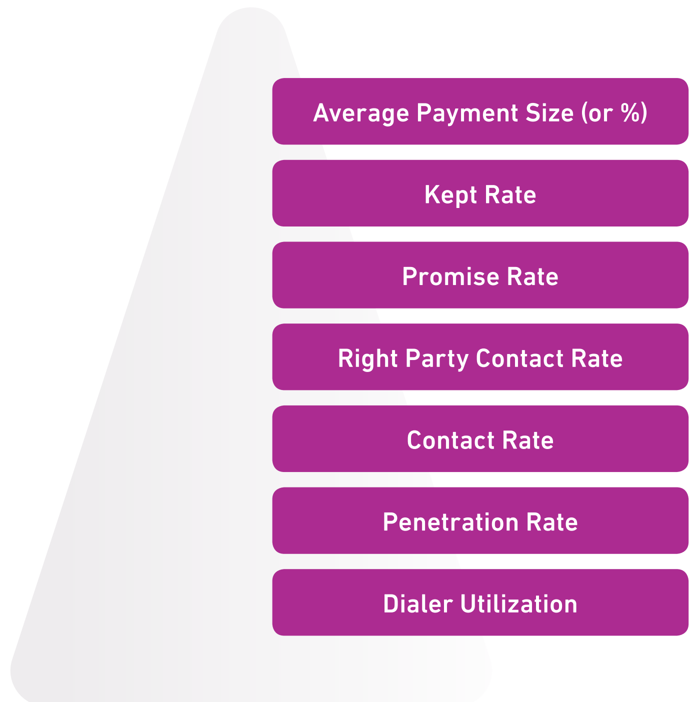

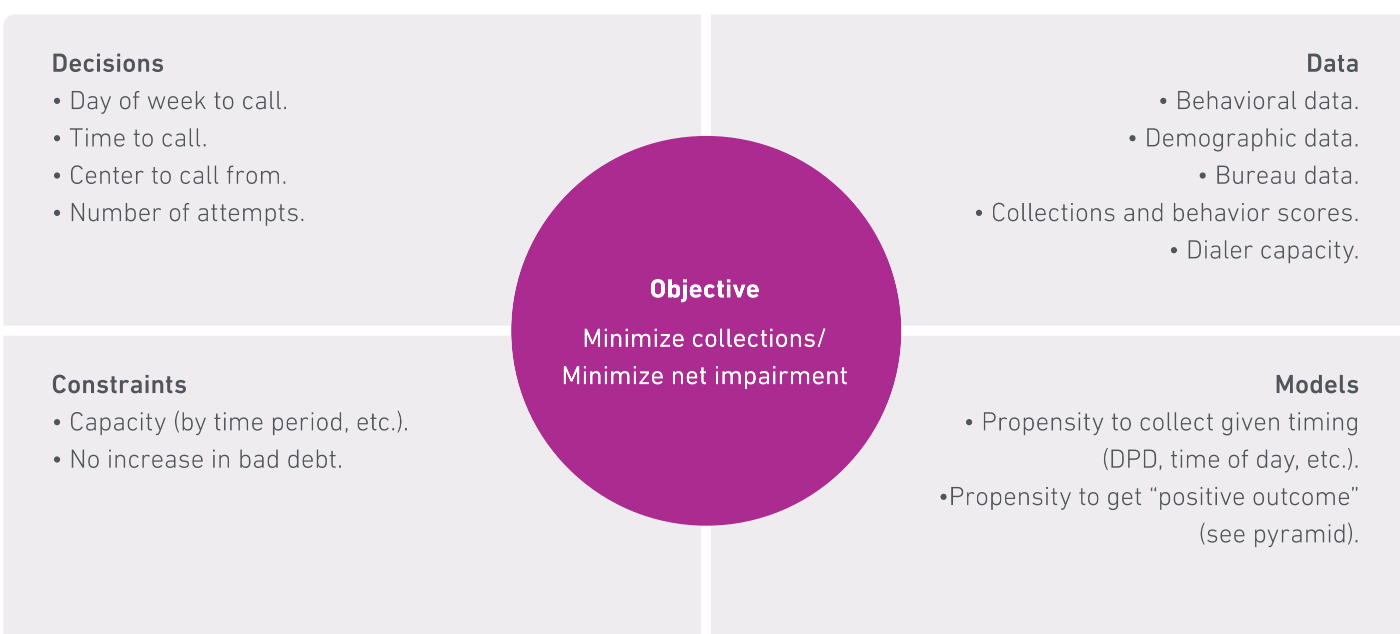

Summary

Optimization is a mathematical methodology that has been demonstrated to add value to many decisioning problems across many industries.

The collections process is no different. There are many different actions that can be taken on many different customers, and clearly there are opportunities to improve performance by taking an optimization approach.

The key challenges to achieving the benefit are:

- Distilling the complexity of the problems into manageable, discrete opportunities.
- Creating customer-level drivers that predict the truly desired outcome.
- Having a sufficiently flexible tool to manage the multiple drivers, KPIs and constraints.

If this is achieved, optimization can add benefit across collections:

- Precollections.
- Self-cure strategy.
- Early/Mid/Late paths.
- DCA allocation.
- Recoveries.

By developing offline strategy trees that can be deployed in various decision agents and workflow management tools, the result is limited IT impact and no IT integration. A similar approach also can be used to optimize dialer activity.

The benefits are significant at each stage: Possible improvements of between 5 percent and 30 percent, correspondingly significant impacts on KPI performance and, ultimately, improved profit for an organization.

Appendix: Optimization across the collections process

The below diagram broadly describes the various stages in the collections process. These are points at which individual optimized decisions can be made, namely:

Precollections strategies

These are actions that occur prior to a customer becoming delinquent. Examples of precollections strategies that can be considered include:

- Limit management.
- Authorizations strategies.
- Preventative calls.
- Triage.
- Reaction to inbound calls.
- Prebill reminder messages.

Self-cure strategies

The idea of a self-cure strategy is to provide a grace period during which a customer will return to order, or cure, without any actions occurring. Often a grace period will be as long as a full month or cycle, though it also can include a number of days, during which no collections action will occur.

Path setting

"Collections path" describes a sequential group of actions that are carried out. Other terminology includes "collections module" and "contact schedule." Typically a customer will remain on the path for its duration (e.g., 30/60/90 days), unless they return to order or are assigned to another path or hold queue (e.g., payment plan).

Optimization across the collections process

Pre-delinquent period	Early arrears	Mid arrears	Late arrears	Recovery
• Credit risk assessment • Proactive monitoring • Early advice and control	• Leave self-cure • Act on repeat debtors • Early commitment • Customer retention	• Reduction in roll rates • Reduction in DSO • Limit collection costs • Rehabilitate	• Pursue full balance • Litigation • Internal collection agencies • External collections agencies	• Debt sales • Write off

Debt management timeline

So, for example, a collections path could be:

- SMS at five days past due (DPD).
- Letter at 15 DPD.
- Outbound call at 25 DPD.

Path setting is probably the most complex of points within the collections process to consider. It can be segmented into three stages:

- **First action optimization —** Choose the preferred (optimal) contact methodology for each customer.
- **Path segmentation —** When there are multiple paths to choose from, consider choosing the right path for each customer (and consider using only existing paths for greater simplicity).
- **Optimal path development —** This is the most complex, aiming to create a path that modifies both the actions and the timing of these actions, which requires significant data and modeling.

Debt collection agency (DCA) optimization

A debt collection agency is a specialist collections organization that is (usually) an external function from the organization the debt is owed to. Typically, a creditor will first decide that a customer needs to be passed to a DCA for resource issues or because the DCA will have greater chance of success of collection. A panel of four or five DCAs is typically used to ensure good performance through competitive tendering and contingency in terms of volume.

Recoveries

The final step in the collections process is typically the recoveries step. This usually involves more intensive or escalated activity, such as legal activity, going to court and ultimately debt sale.

Customer collections versus account-level collections

The term “customer,” which is used frequently in this paper, is a recurring theme of many debates in the collections industry. Here it’s used in its loosest sense — a person on which a debt is to be collected. More analytically, it means a “row” in a decisioning problem, so each problem described here can be applied equally at either the customer or account level.

Tools

There are a number of tools that can assist credit risk and collections departments in their ability to collect well. These include:

- **Strategy design and decision optimization tools —** Analytical tools that employ mathematical techniques to develop data-driven decisioning strategies.
- **Decision agents —** Typically a strategic tool used to score and segment customers and apply different treatments to them.
- **Workflow management tools —** These tools also can score and segment customers, but typically are designed to work in a real-time environment so they can update actions based on activity occurring on the account (e.g., ensuring that if a customer makes a payment in the morning, the planned dialer call isn’t made in the afternoon).
- **Predictive dialers —** These tools are used to make efficient use of operational calling resource, when a list of customers needs to be called.

About Experian's Global Consulting Practice

Experian's Global Consulting Practice is dedicated to creating measurable and sustainable value for organizations around the globe. We specialize in delivering deep analysis and recommendations to help drive profitable business decisions. Our consulting team is composed of former business leaders and consultancy advisors with years of operational and industry experience. Having overcome many of the same challenges you face today, we have real-world experience to help you overcome these challenges by combining our knowledge of global best practices and competencies.

We take pride in helping you not only meet expectations, but also exceed them.

- Average of 20 years' experience per consultant on the Global Consulting Team.
- 43 consultants based throughout Asia Pacific, Europe, Middle East, Africa, India, North America and the United Kingdom.
- More than 700 engagements across the globe.
- Teams with local knowledge of best practices.

Ready to make your debt collection a front-line profit driver for your organization? We can help.

Let's get started

Experian
475 Anton Blvd.
Costa Mesa, CA 92626
T: 1 888 414 1120
www.experian.com

11/17 • 2000/1301 • 1122-DA