BUTLER ANALYTICS

Prescriptive Analytics

A business guide

May 2014

Contents

The production of this document has been sponsored by FICO

The Business Value of Prescriptive Analytics

After fifty years of using information technology to increase the efficiency of business processes we are now firmly in the era where technology is also being used to provide us with information. Business intelligence allows us to establish what has happened and is happening in our business (often called descriptive analytics), and predictive analytics uncover patterns which can be useful in the prediction of future events. This doesn't complete the picture however. Descriptive and predictive analytics may tell us what has happened and what may happen, but they do not tell us the best way to deploy our resources to meet the demands of the future. An example will clarify. In a retail environment our descriptive analytics will tell us sales volumes, seasonal fluctuations and so on. Predictive analytics may give us insights into which products tend to be purchased together. Armed with this knowledge we then need to know how shelf space should best be allocated and more generally how resources should be utilised to maximise revenue and/or profitability. This is where prescriptive analytics fits in - think of it as a prescription for action.

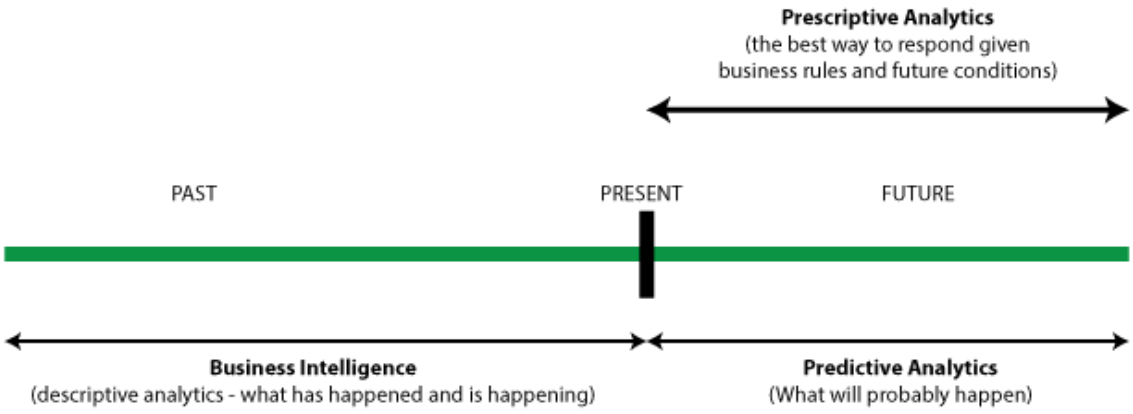

The major part of prescriptive analytics is concerned with resource optimisation given a set of business rules (constraints) and predictions relating to demand, customer behaviour, the success of marketing campaigns and so on. In real business problems, optimisation may involve thousands of variables and constraints, and finding the optimal use of resources, given an objective that is to be maximised or minimised, can only be achieved using powerful computerised optimisation software. Examples abound. Airlines use prescriptive analytics to determine the allocation of seats to each particular class. Vehicle rental businesses optimise the positioning of vehicles to maximise revenue and profitability. Energy companies increasingly use prescriptive analytics and especially with the unpredictable nature of renewable energy sources.

Of course this all assumes that business managers buy into the resource utilisation schedules created by prescriptive analytics techniques. As such the analytics initiative needs high level sponsorship and coordinated effort throughout the enterprise. Reporting mechanisms need to be put in place and procedures to deal with the inevitable changes of circumstances all businesses experience. To this end some businesses run some of their prescriptive analytics processes in near real-time to accommodate change, and such is the power of the optimisation algorithms and computer hardware that this has become possible for complex analytics tasks.

Prescriptive analytics is clearly not a trivial undertaking. It needs close liaison between analytics teams and business management, and an integrated analytics environment capable of integrating

business rules, predictive models and prescriptive analytics. The integration is important, and particularly in large complex businesses. Without such integration prescriptive analytics may be very difficult to achieve, if not impossible.

Expect to see prescriptive analytics technologies more widely used as the user interfaces become more user friendly, and business managers become empowered to address increasingly complex optimisation problems without recourse to teams of analysts. However for large, complex prescriptive analytics tasks the analytics teams are here to stay.

While most analytics technologies are concerned with **what** has happened or will happen, prescriptive analytics tells **how** to best deploy resources to optimise our operational activities - and the benefits are often substantial.

What is Prescriptive Analytics?

Optimisation sits at the heart of prescriptive analytics technologies, and specifically the computation of best resource usage given a set of constraints and objectives. Work planning problems represent a classic application, where work is allocated to limited human resources in a manner that meets constraints and optimises objectives.

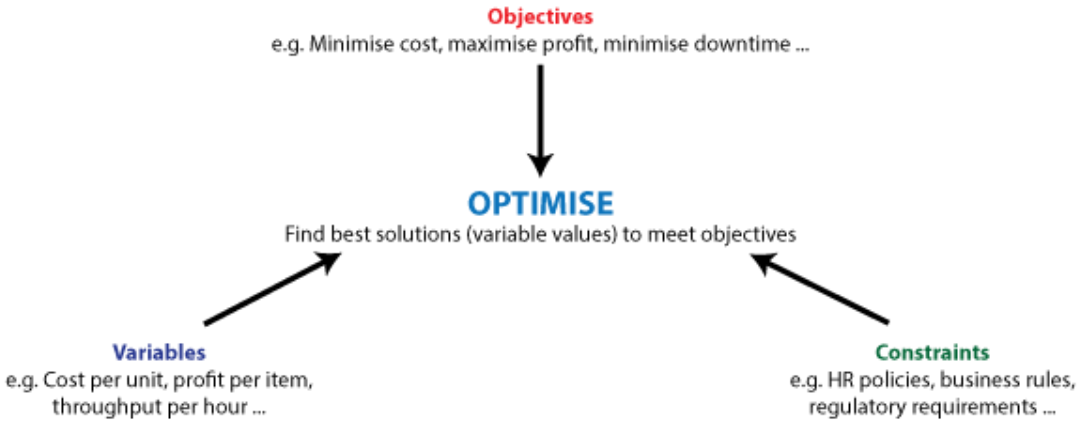

While optimisation has been used for decades in many large corporations, the compute intensive processing has traditionally been associated with very long compute times - typically days and weeks. This limited the application of the technology. However advances made in the mathematical algorithms and more powerful hardware mean that optimisation can be applied to a much broader range of problems, and in some instance execute on a near real-time basis.

The three essential components in an optimisation problem are variables, constraints and objectives. In a work planning problem the variables would typically represent the number of hours work allocated to various people from a given list of tasks. The constraints would limit the way the allocation of resources could take place - no more than 20% of the personnel from any department can be engaged on a project for example. Finally the objectives state what we are trying to achieve. Often this is simply to minimise costs, or maximise profits - or both. However in the work planning problem we might be most interested in minimising the time a project takes. Each optimisation problem has its own set of variables, constraints and objectives and much of the work goes into specifying what these are.

Prescriptive analytics can be divided into two primary activities. The first involves optimisation when the input variables are known (a stock count, or balances in accounts for example). The problem here is simply to establish the best outcome given these variables along with associated constraints and given objectives. A second set of optimisation problems comes under the heading of stochastic optimisation, a suitably off-putting name which simply indicates there is uncertainty in the input data - next month's sales for example. This more complex category of problems will attempt to find the best solution to a business optimisation problem for the most likely future situations. Obviously there is a strong link here with statistical modelling and other forms of predictive analytics, where probabilities are assigned to variables.

It is increasingly the case that prescriptive analytics is integrated with other systems. Optimisation has traditionally been an isolated activity, but today it can take inputs from business rules and predictive analytics processing, and benefits hugely from them. The business rules act as constraints (do not mail someone with an offer of a 5% discount when they have already been mailed a 10% discount - for example), and predictive analytics can provide inputs which predict variable values (the number of prospects likely to respond to a marketing campaign for example).

Prescriptive analytics is still relatively new (the term was first introduced about a decade ago) and only a handful of suppliers provide the integrated environment necessary to take advantage of outputs from other processes. However prescriptive analytics does complete the analytics picture - descriptive analytics (business intelligence) and predictive analytics say **what** has happened or will happen, while prescriptive analytics say **how** things should happen.

Prescriptive Analytics Methods

Optimising complex business problems requires sophisticated technology. Recent years have witnessed major advances in the speed of optimisation algorithms and in the complexity of problem that can be addressed. The net result is the proliferating use of optimisation technologies to address everything from marketing campaign optimisation to how many business class seats should be allocated on individual flights.

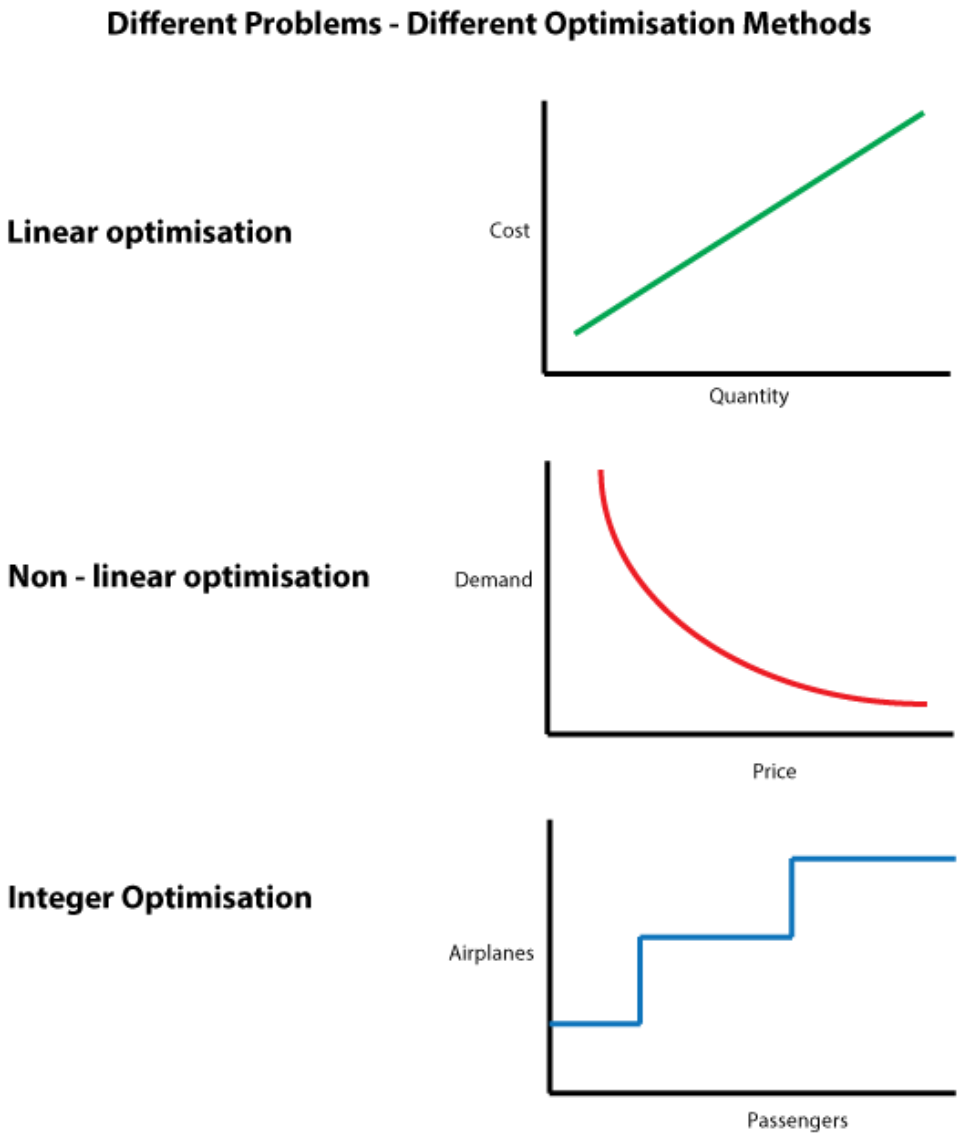

There are several well defined types of problem that optimisation techniques can address - and some they can't. The earliest and often easiest form of optimisation assumed that variables and objectives were related to each other in a linear manner. If a resource usage is doubled, so is its cost. While there are some problems that are well served by this model (the use of material components in a mix for example), many are not. To cater for more complex optimisation problems, non-linear relationships have been accommodated. A good example here is a price/demand curve where demand drops off rapidly as price exceeds a certain threshold, and increases exponentially as price drops below a critical level. The solution of non-linear optimisation problems is much more complex

than linear problems, but contemporary tools with good user interfaces help keep such problems manageable. Other problems require that variables can only take on integer values (we can't have 2.5 airplanes for example). Another class of problem makes use of network programming, where the aim is to minimise some function of the network. A good example here is minimising the cost of transport as a given number of trucks ship goods to a network of stores.

Other techniques are also finding their way into prescriptive analytics, in addition to the optimisation techniques mentioned above. Queuing problems are common in business and optimisation techniques are used to address problems from traffic flow through to minimising check-out queues in stores. Simulation is also used to model the performance of business systems and is a large domain in its own right. It is very often the case that the 'best' solution to various business problems simply cannot be found, and so looking for a good solution becomes necessary, and this is where both analyst and business managers need to really understand the problem they are attempting to solve.

Stochastic optimisation takes prescriptive analytics into a realm where many uncertainties in business can be accommodated. Employee attendance, future sales, the response to marketing campaigns, wastage and hundreds of other variables are inherently uncertain in nature. The variables can be treated as random in many ways, with limits on how much they can vary. The stochastic optimisation algorithms will find the best, or at least a good, solution for the most likely outcomes where uncertainty is present.

Such is the advanced nature of some prescriptive analytics tools and solutions that near real-time optimisation can occur to accommodate changing business conditions. For the very largest optimisation problems this still is not possible, but the frontier is being pushed forward rapidly and in volatile markets real-time optimisation can, and often does deliver significant benefits.

Prescriptive analytics technologies will advance rapidly over the next four to five years with new entrants and new capabilities. As always integration is largely the determiner of how successfully prescriptive analytics can be used in a live production environment. Building prescriptive models is one thing, using them in a production environment requires extensive integration capabilities and good management and control tools.

Integration

Prescriptive analytics, and specifically optimization, has traditionally been treated as a stand-alone domain. This has meant that the inputs for optimization have been manually created, and that the outputs have been produced in isolation from all other systems. In practical terms this equates to increased cost, delays, errors and a frustrating lack of flexibility. For these reasons it is important that the prescriptive analytics environment is integrated into the overall systems environment as fully as possible.

One of the most significant overheads associated with performing prescriptive analytics is the creation and maintenance of business rules, or in the terminology of optimization, the constraints. Even modest optimization projects might involve hundreds or thousands of rules, and recreating them for every optimization problem is a heavy overhead prone to errors. Ideally the prescriptive analytics tools should have access to a business rules data base and management system and be able to convert them into a format that makes sense. In fact for large optimization projects such a facility is not really an option.

Business rules typically express how resources can be combined. It obviously does not make sense to offer a customer a five per cent discount after offering ten per cent on the same product. Neither is it likely to be acceptable that the whole of a workforce is laid off periodically. These and tens of thousands of other constraints typically influence the way large businesses operate, and since prescriptive analytics makes extensive use of them there should be a high level of integration.

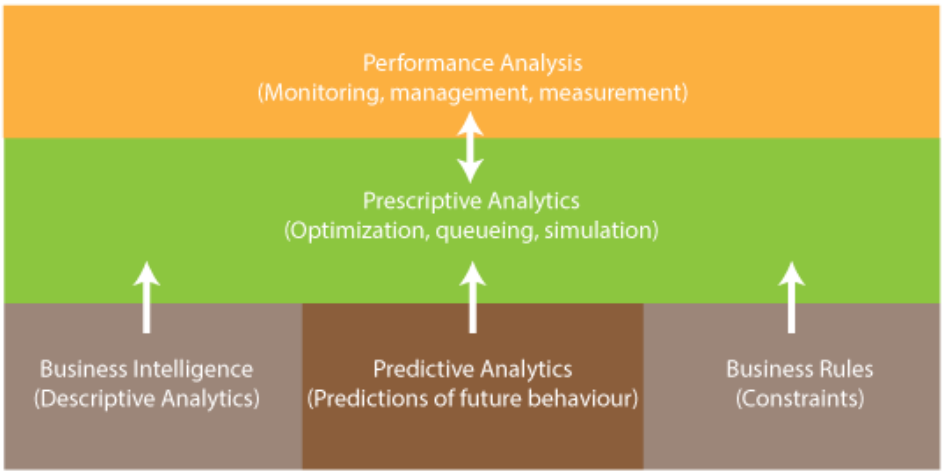

However it is not just business rules which should be integrated. Many of the inputs involve the specification of forecasts and other types of prediction. Predictive analytics will generate many models with predictive power - the propensity of customers to respond to an offer for example. These can be used as input to a optimization project, and once again it is extremely useful if the predictive and prescriptive analytics environments are integrated. Optimization caters for probabilistic inputs (a sales forecast for example) through stochastic optimization techniques, and since these facilitate much greater sophistication, the integration with probabilistic predictive models is very important.

Business intelligence, or descriptive analytics, can also provide key inputs for prescriptive analytics, providing information on what has happened and is happening. Information such as the sales of items by region and period might constitute key information in production planning optimization, along with many other metrics. Again it is important that output from reporting and analytics are available to optimization projects.

Finally the results of optimization need to be deployed, and done so in an environment where performance can be monitored, measured and managed. The results of optimization may have a very short life, and it is essential that changes in circumstances can be compared with assumptions made when an optimized model was built and new models built in a timely manner.

A lack of integration is not so problematic for small, one-off projects, but it becomes a major headache as projects grow in size and frequency. Integration should be in the top three or four requirements when selecting a prescriptive analytics platform.

Business Application

Prescriptive analytics applications embrace most aspects of business operations. In fact whenever there is a resource allocation problem with constraining rules, many variables and well defined objectives, then it is likely that prescriptive analytics can be applied.

It's application in marketing is gaining momentum, particularly combined with predictive analytics where customer response to various initiatives can be predicted. Resource allocation can then be optimized based on these predictions and the relevant business rules which determine how customers are to be treated. Of course optimization has been applied for decades to workforce and other resource deployment problems, but it is only now with much faster execution that other problems are becoming feasible. Even sports bodies are using optimization to maximize television ratings while catering to the needs of players and fans.

Optimization is used widely in the airline industry. Since margins are so thin frequent optimization can help turn unprofitable scenarios into profitable ones. Bad weather will mean re-optimizing resource allocation, and as seats are sold on a flight it is desirable to re-optimize the allocation of various classes of seat. Unscheduled maintenance can also cause resources to be re-allocated and optimization is essential if inefficiencies are to be ironed out of operations. In this particular application it is very important that optimization can be performed in less than 24 hours, and often in much shorter times.

The energy industries have used optimization for decades, but increased optimization speed and capability have meant new applications. With new energy sources, and particularly renewable ones, it is desirable to combine prescriptive analytics with predictive analytics to create long term energy scenarios. The predictive analytics are used to create long term energy forecasts and optimization is then applied to explore optimization of energy production.

An application ideally suited to optimization is that presented by vehicle rental firms. With thousands of vehicles assigned to hundreds of locations there is an obvious need to make sure that each vehicle is located where it will generate most revenue.

The earliest application of optimization included problems such as 'least cost mix'. Here the aim is to produce a mix of component materials with a particular specification, given that each component has its own constituent properties. Animal food mix was the classic application where overall nutrient levels had to be met from mixing various ingredients, which each had their own nutrient levels. This and other similar problems are still solved using optimization technology, but look fairly primitive compared with the real-life problems that are being solved today.

Optimization technology will grow in use as it becomes more user friendly, executes with greater speed, and is more tightly integrated into the over systems environment. For many businesses optimization will provide an edge that cannot be achieved any other way.

Strategy

Realizing the considerable benefits that prescriptive analytics can bring to an organization requires that several issues are adequately addressed. These include technical, organizational and operational factors and include:

1. Large, complex optimization problems require a team of professionals trained in operations research or some related discipline. Some smaller problems can be addressed by non-specialists using technology such as Excel Solver, but most real-life problems will need an experienced team and significantly more sophisticated technology.
2. High level sponsorship is needed, since prescriptive analytics usually span functional silos. The optimal solution to enterprise problems may seem sub-optimal at the department level

for example, and so there will have to be mechanisms put in place to allow such issues to be resolved.

3. The technology platform must scale and offer high levels of performance. While initial projects may be comparatively modest the scale and scope will rapidly grow as benefits are realized. Performance and scaling bottlenecks will be experienced if the supporting technology is architecturally weak.
4. Integration with existing analytics tools and business applications means inefficiencies can be kept to a minimum and errors largely eradicated. Business intelligence, predictive analytics, rules based systems and some transactional applications will need to be integrated with the prescriptive analytics platform. Unless this can be achieved the speed and accuracy of optimization will be compromised.
5. Obviously there needs to be adequate monitoring and management of prescriptive analytics projects, with effective reporting mechanisms so that changes in the business environment can be responded to in a adequate manner, and changes in business strategy quickly implemented.

It should be clear that prescriptive analytics is deeply concerned with the operation efficiency of an organization and needs to be integrated into the information systems environment. Supporting information needs to be extracted from other systems and sent to operational systems to implement the resulting solutions. Parallel with this is the need for management and reporting structures so that associated issues can be resolved. Without this enterprise support the prescriptive analytics efforts tend to remain isolated and inefficient.

Finally it is necessary that the whole prescriptive analytics effort is business driven, with a good understanding of where the major payoffs are and how projects should be prioritized. For organizations inexperienced in the domain this may mean using external resources (such as consultants and experienced suppliers) to formulate a strategy. In some industries it may be possible to buy solutions to specific problems, and inevitably the options here will grow rapidly over coming years. However it really is very important that organizations do not end up with multiple point solutions, and worse still with solutions that will not scale. And so the issues listed above are just as applicable to solutions as they are to deploying a prescriptive analytics platform.