

Optimising Business Rules in the Services Sector

Alan Dormer

Abstract—Business rules are widely used within the services sector. They provide consistency and allow relatively unskilled staff to process complex transactions correctly. But there are many examples where the rules themselves have an impact on the costs and profits of an organisation. Financial services, transport and human services are areas where the rules themselves can impact the bottom line in a predictable way. If this is the case, how can we find that set of rules that maximise profit, performance or customer service, or any other key performance indicators? The manufacturing, energy and process industries have embraced mathematical optimisation techniques to improve efficiency, increase production and so on. This paper explores several real world (but simplified) problems in the services sector and shows how business rules can be optimised. It also examines the similarities and differences between the service and other sectors, and how optimisation techniques could be used to deliver similar benefits.

Keywords—Business rules, services, optimisation.

I. INTRODUCTION

BUSINESS rules in the services sector are about guiding decisions. As information is processed, decisions are made and further information is generated. For example, a customer enters a bank and asks for a loan. Information is gathered about the customer and what they want, and the result of analyzing that information other information is generated, e.g., the customer is accepted for a loan of a certain value with certain repayment profile, etc. The correct application of rules enables consistency and management by exception. For example, a skilled and experienced person is not required to be present at every consultation, he/she only needs to engage where their judgment and/or approval is essential.

Business rules are essential to the services sector for consistency: if two customers present as the same it is not acceptable (or indeed sensible) to treat them differently. And in the last 20 years or so there has been an increase in computerisation that enables not only consistency but monitoring, enforcement and flexibility [1]. These software systems, called business rules engines, enable organisations to build and maintain sophisticated sets of rules that can control and monitor many thousands of staff and millions of transactions, in real-time. They also enable rules to be changed to reflect changes in business circumstances. But while business rules deliver consistency, they do not automatically deliver efficiency or maximise customer service or revenue. For further information see [2].

Business processes are also important to both the services and manufacturing sectors. This is particularly important in manufacturing as many important parts of the business are about processing physical items rather than information and rules themselves are not sufficient [3]. The human resources and physical resources which are used, both cost money. This has naturally led to the concept of business process optimisation where the sequencing of tasks, allocation of tasks to machines, etc, are planned to minimise costs and/or maximise revenue [4]. Indeed manufacturing often goes a lot further and uses optimisation and forecasting techniques to maximise profit based on variable demand and anticipated demand in quite sophisticated ways [5]. The key difference between manufacturing and services is that, to be efficient, manufacturing strives for repetition (i.e., we make large quantities of the same thing). In services the aim is more for consistency, that is, the customers are often different with different situations and needs but if two customers are the same, then we treat them the same.

Optimisation in the services sector is not so well advanced. Business rules are fundamentally designed for correctness and consistency. Two notable exceptions are staff rostering [6] and supply chain optimisation [7]. Rostering is widely used for the purposes of having the right number of people (and no more than is necessary) with the right skills, in the right place at the right time. This is motivated by cost reduction, because people costs are very significant in the services sector. Supply chain optimisation is about tasks, such as supply, storage and distribution, the allocation of resources to tasks, and the order that tasks are performed.

So, broadly speaking, an organisation may have right number of people, with the right skills, all doing the same thing (in the same situation) by using the same business rules. But could it do better, and serve its customers better by optimising business rules?

II. OPTIMISATION IN THE CONTEXT OF BUSINESS RULES

In the following section we look at three simplified examples where business rules have an impact of revenue, profitability and customer service. In each case a problem is presented and a solution is developed. The problems are real but have been deliberately simplified for the purposes of this exercise.

A. Loan Application Example

In many situations in the service sector actions are based on information which is either available or can be elicited from a customer as part of the interaction. A simple example is requesting a loan.

Alan Dormer is with CSIRO Mathematics, Informatics and Statistics, South Clayton 3169, Australia (phone: +613 9545 8082; e-mail: alan.dormer@csiro.au).

The customer can decide on a number of options (How much? How long? Secured or unsecured? Include repayment protection insurance? and so on). On initial application, some decisions can be automatically made either by an unskilled worker or the back end of a website using business rules. The possible outcomes are essentially decline, refer (by a more skilled and/or experience staff member) or approve. The business rules will use information (such as the customers income and/or credit history), and/or combinations of information (such as the ratio of loan to the customers income) to decide which category the customer should be placed in. See Figure 1 for a simple diagram which considers two factors: loan amount and customer income. The rules have an impact on the business in a number of ways. Consider the decline/refer interface. If the rule is too conservative then more customers are turned down leading to a loss of business and customer unhappiness; if the rules are more aggressive then more customers are referred (which costs money) and may ultimately be declined. Similar considerations apply to the refer/accept interface. So even in this oversimplified case, there are trade-offs and an optimisation problem to be solved.

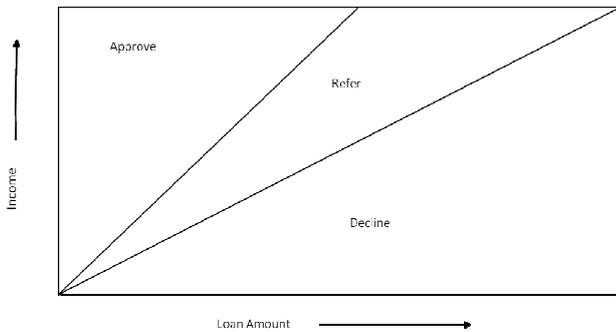


Fig. 1 Business Rule for Screening of Loan Applications

The business rule illustrated in Figure 1 may be stated as follows:

IF (loan/income) < accept_limit THEN approve = true
 ELSE IF (loan/income) < refer_limit THEN refer = true
 ELSE decline = true

In order to treat the loan acceptance problem as an optimisation problem, we need some further information such as:

- the probability of default as a function of loan value/income (x)
- the losses associated with default
- the additional cost referring an application
- the value of a customer
- the distribution of income/loan value across all customers

Let us assume that:

- $P(\text{default}) = x^2 / 100$ where x is the loan/income ratio
- Cost of default per customer = D
- Cost of referral per customer = R

- Value of customer = V
- x is distributed uniformly over $[0, 10]$
- Referral reduces the rate of default by 75% and rejects 50% of applications.
- The objective is to maximise revenue.

We have to choose parameters, x_1 and x_2 , for the interfaces between rejection and referral. The cost function is given by:

Value of customers automatically accepted

$$v_A = V \cdot x_1 / 10$$

Value of customers accepted after referral

$$v_B = (V - R) \cdot 1/2 \cdot (x_2 - x_1) / 10$$

Expected cost of defaults

$$c = D / 1000 \left\{ \int_0^{x_1} x^2 \cdot dx + \int_{x_1}^{x_2} x^2 / 8 \cdot dx \right\}$$

Total value minus costs

$$v_A + v_B - c$$

Differentiating, the value of a stationary point is given by the equations:

$$x_1^2 = 400(V + R) / 7D$$

$$x_2^2 = 400(V - R) / D$$

Intuitively these make sense. If the value of D increases, both parameter values go down; if V increases, both parameter values go up, and if R increases the gap between them decreases to the point where there is no gap when $R = 6V/8$. With $V = 15$, $R = 5$ and $D = 100$, the maximum value is achieved at (3.38, 6.32).

This simple example illustrates the general approach of finding the measure of system performance as a function of some parameters, then finding the best values for these parameters.

B. Debt Waiver Example

Another example that creates an optimisation problem is one of debt waiver. Data is normally available on the recovery rates of debts of various sizes and the cost of their recovery (administration, legal fees, etc), at the intersection of this line and the recovery cost, we can read off the value of the minimum debt that is economical to enforce.

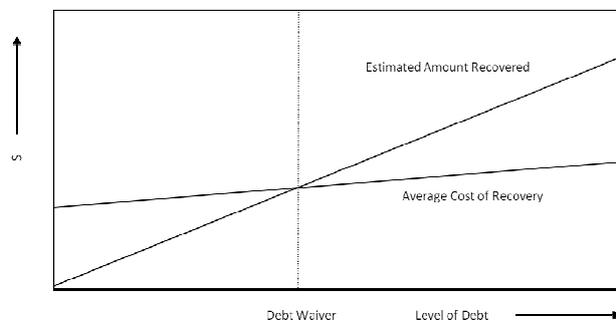


Fig. 2 Debt Waiver

C. Transport Rules

It is quite common in transportation systems to have rules governing disruptions, e.g., a train is late or disabled, or conflict, e.g., two trains requiring the same track at the same time [8]. A typical problem is what to do when train *A* is late and there is a connection to another train, train *B*. Typically *B* waits for up to a maximum cut-off (*c*) for *A* where $0 < c < f$, where *f* is the time until the departure of the next train.

The trade off in choosing *c* is the inconvenience to passengers in train *B*, for waiting, against the inconvenience of passengers in train *A* for having to wait for the next train.

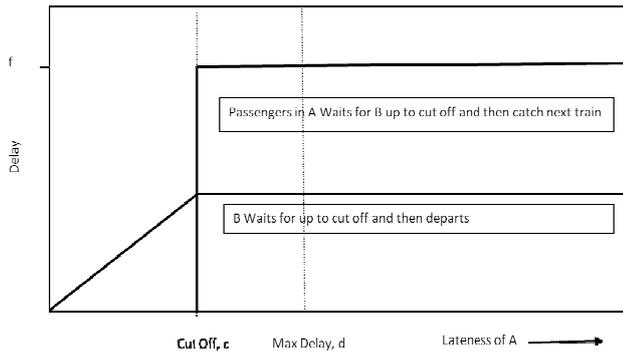


Fig. 3 Transport Rules

If we wish to be fair to passengers in *A* and *B* then the best value for *c* depends on the distribution of delays to *A* and the interval between trains, *f*. If we assume that the distribution of delays is uniform between $[0, d]$, then we have:

$$\text{Exp}(\text{additional delay } A) = (d - c) \cdot \left(\frac{f}{d}\right)$$

$$\text{Exp}(\text{additional delay } B) = \frac{1}{2} \cdot \frac{c^2}{d}$$

For equal inconvenience, we have:

$$(d - c) \cdot f = \frac{c^2}{2}$$

$$c^2 + 2cd - df = 0$$

This is a quadratic equation and can be solved for *c*. For example, when *d* = 10 mins and *f* = 20 mins, *c* = 8.33 mins, close to the maximum delay anticipated.

In addition, there may be benefits to considering the number of passengers in trains *A* and *B* to genuinely make the rule fair.

D. Other Examples

There are many other examples of rules that affect the performance of an organisation. For example in public services we have the concept of triage where cases come in and the problem is to determine the urgent cases. This is particularly important in health and social service, such as child protection. But a good definition of urgent requires a rule and parameters. It is no good classifying too many cases as urgent if that is out of step with the available resources; you simply create a queue later on.

There is also the problem of fluctuating demand. Rules must remain relevant even when demand is changing.

III. THE ADDITIONAL OPTIMISATION CHALLENGE

As mentioned above, work has been done in business process optimisation. That involves the structure and ordering of tasks in order to minimise time, minimise cost or maximise throughput. The key feature in the examples above is probability. In the first, it is the probability that a customer, if pursued, pays up. In the second it is the probability of default, and in the third it is the probability of delay. Where there is a range of customers, or requirements, or external factors that impact the business problem needs to be considered in the optimisation of rules.

A. Real Life Problems

Real life problems are more complicated. For example:

- In the loan example there will be more than one criterion. Credit history or value of security may also be taken into account [9]. The interest rate may vary with credit rating or size of loan, affecting affordability.
- With debt waiver there may be additional rules (also subject to optimisation) that govern recovery and affect the cost. For example it is not unusual for debts above a certain size to be subject to mandatory legal action
- In transport there are additional complications. People travel in both directions, there are knock-on effects of delays around the network, and the concept of fairness may be difficult to sell to customers [10].
- The order in which rules are applied is also important. For example, servicing customers costs money so finding out sooner rather than later that a customer does not qualify for a loan is an advantage
- The rules that need to be applied to minimise cost or maximise profit may not be evident or obvious. If we have not explored every possible rule, in every possible combination, then we cannot be sure that we have the best set.

But having said this, these simplified real world examples demonstrate that:

- There are optimisation problems even within construction of simple rules
- They can be posed and solved
- One or more probability distributions are often important

IV. PARALLELS WITH MANUFACTURING AND OTHER NON-SERVICE INDUSTRY

As previously stated, manufacturing has typically focussed in business processes rather than business rules per se. Take up of optimisation techniques is common with a number of key elements [11] as follows.

A. Long and short term demand forecasting

Demand forecasting is a vital element of any forward looking optimisation of decision support. Techniques have been developed to combine stochastic and deterministic variations in demand that could be applied to the services sector.

B. Strategic optimisation: process design and investment in resources

The key issue here is to match equipment capacity to the anticipated long term demand. This includes investment and capacity planning.

C. Tactical optimisation: production planning and resource allocation

At this stage the rates of production are decided and the production task is allocated to the most appropriate, efficient capital and human resources.

D. Operational optimisation: production scheduling

This deals with short-term variations in demand and schedules production to meet current and short term anticipated demand.

The provision of services consists of a series of related tasks that are carried out to satisfy the needs of a customer (or customers). We can draw some similarities between these issues and the services sector. There are even some deeper similarities. For example in manufacturing there is the so-called lot sizing problem which is to determine the optimal production run taking into account start-up/changeover cost and the cost of holding stock.

If lots are too small then start-up/changeover costs are too high; too large and more stock is produced, and the cost of holding it is too high.

There is a similar problem in the services sector where the same customer service staff can serve customers, act in the call centre and process paperwork in the back office. Demands for all three can vary and to be efficient staff allocation is based on demand or expected demand.

There are also changeover costs as staff shut down one activity and start another. This is similar to the lot sizing problem; shorter periods enable maximum flexibility and higher overall efficiency, but make them too short and changeover costs become unacceptably high.

The transport problem, which is case of disruption management, is typically solved by using optimisation techniques to determine a new schedule.

This is a complex and expensive process and, like any optimisation problem, heavily dependent on the availability of large amounts of accurate data.

Rule-based approaches have been used in timetable generation, particularly for conflict resolution, and these could be developed further to inform the rules that transport system controllers (implicitly) use to resolve disruptions to create greater consistency. The simple example could be used as a way to manage delays in the system.

V. HOW TO ADDRESS OPTIMISATION PROBLEMS

In order to optimise we need to identify the decision variables, x , create the cost and constraint functions, $f(x)$ and $c(x)$. The latter two are essentially models which may be based on a fundamental understanding of the process, empirical data or a combination of the two. In manufacturing models, they are readily available as the performance of machines and processes is generally understood. In the service sector this is more difficult due to the variability in staff performance and customer behaviour. For example, for the loan application example, there are several relationships that need to be determined:

- What is the relationship between loan/income ratio and the probability of default?
- How good is the referral process (e.g., proportion of false positives; acceptances that should have been rejections; and false negatives, rejections that should have been acceptances)?
- What is the average and variability of cost of the referral process?

Similarly in the debt waiver problem, there are:

- Costs to recover debts of different sizes
- The expected recovery rate for debts of different sizes

These sorts of relationships can be determined by analysing past data and fitting statistical models, but there are some complications such as the rules (and the consistency of their application) that were in use at the time the data was gathered. This may require some experimentation in addition to data analysis.

A. General Rule Parameterisation Problem

The general problem is, given a business process, a set of rules we can define an objective function $f(x)$ that will calculate costs for any given set of parameters that drive the business rules, x . The optimum value of x is that value that minimises $f(x)$ over the set of all expected states of the business, weighted by their frequency or probability.

The optimisation method to solve the parameterisation problem depends on the nature of the objective function and constraints. A linear problem may be solved by Linear Programming [12]. For other problems (i.e., non linear), there are a range of methods including direct search, gradient methods and methods that utilise information about second derivatives, such as Sequential Quadratic Programming (SQP) [13]. The choice between these methods depends very much on the nature of the objective function. Gradient methods are efficient but they require the function to be continuously differentiable as a minimum; direct search is less efficient but imposes fewer conditions on the objective [14]. Different methods may be appropriate depending on the rules and the underlying objective function relating to the business. For example, the cost function may be continuous, continuously differentiable or even linear. What happens when we impose a set of business rules with parameters and then optimise those parameters? Can we draw any conclusions about the nature of the new objective function?

Given the penetration of business rule engines in the market, another challenge is to how to populate the software with new parameters, and then run a simulation to determine the impact on the objective function.

VI. FURTHER WORK

This work could be extended and generalised in a number of ways:

- Exploring the relationship between business rules and the business objective function. As discussed, how what is the relationship between business objectives, the rules and the objective function that drives the parameterisation problem?
- Can we go further than just calculate optimal parameters? What about the re-organisation of rules, or their automatic generation?

How do we integrate business rule optimisation with other types of optimisation that are well established in manufacturing such as strategic, tactical and operational?

VII. CONCLUSION

We have presented some simplified (but real world) examples where choice of business rules has a real impact on the financial and/or customer service performance of an organisation. In all cases there is an optimisation problem, which has been posed and solved. One important conclusion is that probability plays role in all of the examples. It is possible to optimise business processes without regard to probability since the inputs and outputs are can be considered deterministic. But with rules, the process itself is driven by external events and actions by customers, and information received from customers. The rules must work over a range of inputs and responses and when we optimise them it must be for those probability distributions.

We believe that this approach could be extended and generalised so that organisations could determine the optimal parameters in any set of business rules in order to satisfy business goals such as lower costs, greater revenue or better customer service.

Further work could be directed towards more general problems, understanding the nature of the parameterisation problem and ultimately generating optimal rules.

REFERENCES

- [1] Andreescu, A, Methodological approaches based on business rules, *Informatica Economica Journal*, Vol XII, Issue 3, pp 23-27, 2008
- [2] Taylor, J, *Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics*, IBM Press, September 30, 2011, ISBN 0-13-288438-0
- [3] Harmon, P, *Business Process Management Today and Tomorrow*, M. Dumas, M. Reichert, and M.-C. Shan (Eds.): BPM 2008, LNCS 5240, p. 1, 2008, Springer-Verlag
- [4] Vergidis, K, *Business Process Optimisation Using and Evolutionary Multi-Objective Framework*, PhD Thesis, 2008
- [5] Dormer, A, *Hybrid Optimisation System for Solving Planning and Scheduling Problems*, COR/INFORMS, Banff, 2004
- [6] Ernst, A, et al, *Staff scheduling and rostering: A review of applications, methods and models*, *European Journal of Operational Research* 153 (2004) 3–27, 2004
- [7] Zachary, H et al, *Supply-Chain Optimisation- players, tools and issues*, *OR Insight* (2001) 14, 20–30
- [8] D'Ariano, A., *Conflict Resolution and Train Speed Coordination for Solving Real-Time Timetable Perturbations*, in *IEEE Transactions on Intelligent Transportation Systems*, Vol 8. Issue 2, 208–222, 2007
- [9] Federal Deposit Insurance Corporation (FDIC), *Risk Management Examination Manual for Credit Card Activities*, Chapter VII, FDIC-Division of Supervision and Consumer Protection, March 2007
- [10] Dormer, A, *Innovations in Recovery Optimisation*, AGIFORS, Sao Paulo, 2005
- [11] Lawler, E, et al, *Sequencing and scheduling: Algorithms and complexity*, in *Handbooks in Operations Research and Management Science*, Volume 4: Logistics of Production and Inventory. Elsevier, 445–522, 1993
- [12] Guéret, C, Prins, C, and Sevaux, M, *Applications of optimization with Xpress-MP* Revised translation from the French language edition of: *Programmation linéaire* by, Editions Eyrolles, Paris, France. Translated and revised by Susanne Heipcke, Dash Optimization Ltd, 2000
- [13] Nocedal and Wright (2006). *Numerical Optimization*. Springer. ISBN 0-387-30303-0 *Journal of Global Optimization*
- [14] Ali, M, Törn, A and Viitanen, S, *A Numerical Comparison of Some Modified Controlled Random Search Algorithms*, *Journal of Global Optimization*, Volume 11, Number 4 (1997) S. P. Bingulac, "On the compatibility of adaptive controllers (Published Conference Proceedings style)," in *Proc. 4th Annu. Allerton Conf. Circuits and Systems Theory*, New York, 1994, pp. 8–16.