Opturion

Home Delivery Optimisation

Oct 2015

Alan Dormer

Background

Home delivery of food and groceries has increased rapidly over the last 5 years. Initially seen as a small add-on, its growing significance is prompting initiatives to streamline operations and increase competitiveness. These include enhancing operational effectiveness, reducing costs, and offering customers more choice and convenience, at a competitive price.

How Optimisation Can Assist

Optimisation is all about choice. We have decisions to make, each of which has consequences for the on-line retailer and the customer. These choices involve trade-offs, a large delivery fleet is expensive, but can cope with higher demand; customers like short lead times and small delivery windows, but these are more costly to serve.

Optimisation allows us to model these consequences and trade-offs and find the choices that are best for the retailer and the consumer. The choices can be strategic (size of fleet, regions supplied, price/time window services, etc), tactical (how many drivers do we need next month) to operational (which truck serves which customers on the day, management of premium, urgent deliveries, etc.).

Key Elements of any Solution

Efficient Delivery

The cornerstone of any delivery operation is efficiency and customer service. If we have a catchment area from a store (or dark warehouse) we will have a list of deliveries. This will include:

- Address
- Goods
- Delivery window

With this list in mind, and a fleet of delivery vehicles, we need to create a number of delivery rounds that satisfy all of our customers, at minimum cost (labour, fuel, vehicle depreciation, etc) taking into account such factors as:

- Drop off time
- Traffic
- Resupply
- Vehicle capacity
- Driver breaks

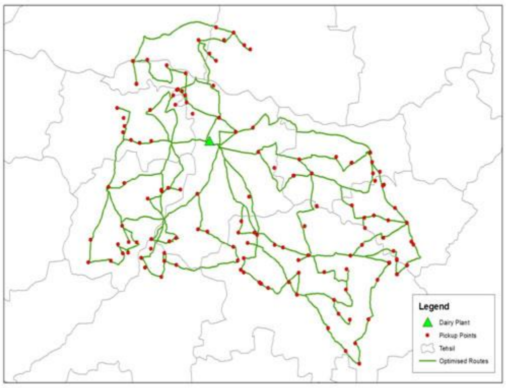

Ideally, we can calculate these rounds beforehand, and on demand, as things change. This ability to recalculate, quickly, is important later on.

Delivery Windows and Dynamic Pricing

Customers tend to have similar delivery preferences ; working people want deliveries shortly before and soon after work, for example. They also like short windows, ideally to nominate an exact time. Left unchecked, this would require a high concentration of resources over short periods, and a lot of downtime. This costs the retailer dearly. We need a way to encourage customers to elect for delivery slots that smooths demand and provide some room to manoeuvre for the retailer.

One way to do this is provide a restrictive number of delivery slots at any time during the day, and have variable pricing based on predicted demand and the length of the window. For example (given the assumptions above):

- A 1 hour slot between 7 am and 9 am or 5 pm and 6 pm might be the most expensive
- A 6 hour slow between 10 am and 4 pm, might be the cheapest

This is similar to the revenue management problem for airlines and hotels. Identical seats are priced differentially based on expected demand and perceived value for the customer. In the case of airlines, perceived value is mainly flexibility (book late, cancel or change without penalty). In our case, convenience is a combination of short lead time (last minute booking) and the size of the delivery window.

There are a couple of complications, however:

- Revenue management for airlines is about getting the highest price for the seats available. With home delivery, we want to offer prices that encourage customers to act in a way that enables the retailer to be efficient
- Home delivery customers are dispersed, so cost to serve is not the same for all of them

Cost to Serve

The cost to serve is how much each delivery costs the retailer. Clearly, we could divide the total cost of a day, week or month by our costs (labour, fuel, maintenance and depreciation) but that only gives an average, and makes no distinction between the location of the customer relative to the store (or DC) that serviced them. In addition, as identified above, there are synergy savings. For example, several customers in the same street or apartment block effectively subsidise each other. This is the same for public transport; population density reduces the cost for everyone.

In this case we are looking for a marginal or additive cost; we have the people and vehicles already. We are just looking for extra costs in terms of distance driven and fuel used. There is another consideration worthy of note: making a certain delivery may stop us making another later on. However, if we have our pricing model worked out sensibly, we can safely ignore this on the basis that deliveries should be largely interchangeable from a revenue perspective.

A Practical Approach

Delivery Windows and Dynamic Pricing

Let is consider the case where there are no deliveries booked. At this point, all we can offer is:

- Time of Day dependent (TOD) pricing
- Delivery Window Length Dependent (DLW) pricing

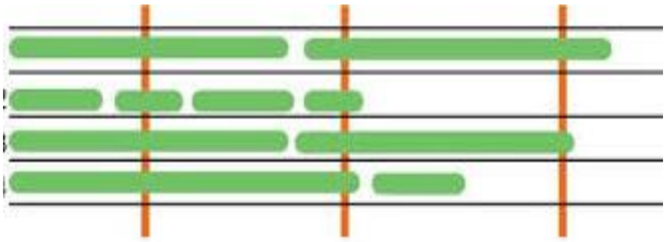

A simple approach to TOD is to assign a cost for every 1 hour segment, say low medium and high, based on experience and the average cost to serve. High might be 150% of the average and low might 50%. These numbers would be tuned based on customer behaviour to get a constant workload throughout the day. DLW pricing would be calculated as the minimum TOD price in that window, with perhaps a further discount to reflect the potential flexibility.

A classic revenue management approach would be to vary the prices of all deliveries to fill up all the slots at the desired rate, usually based on previous experience, whilst keeping back a proportion of medium and high priced slots. This prevents selling out high and medium value slots too early and losing revenue from people who would pay more for them later on. Whilst this makes sense for airlines and hotels, that aim to maximise revenue, the objective of home delivery is to encourage consumers to opt for slots that maximise overall delivery efficiency.

However this misses out an important factor. As orders come in, we can start to build routes based on those orders. At some stage, say when we have 80% of our orders in, these rounds start to take shape and we can start to factor in the marginal cost of additional orders, and be more analytic. That is, if I place an order and within my time window a delivery vehicle is *currently planned* to be passing close by, my delivery charge might be lower than it might have been using the formal above.

In fact, we can do better than this by adding the new order into the problem, and re-optimising. If existing the new and existing deliveries can all be made on time, the marginal cost of the new delivery is the difference.

Cost to Serve

Optimisation gives us an excellent approach to the question of cost to serve. When we have the data on a complete day of deliveries we have all the data we need. We have the ability to re-optimise; that enables us to systematically a delivery and re-optimise. We then put that delivery back and move on to the next one. The difference in cost is the true additional, or marginal, cost of that single delivery. Anything less is a potentially inaccurate, misleading and dangerous compromise. For example, it is possible to have a number of deliveries where their presence and absence impact the route of more than one vehicle. This is particularly true of areas that are geographically distributed, deliveries on the edges of catchment areas or where there are many one-way streets.

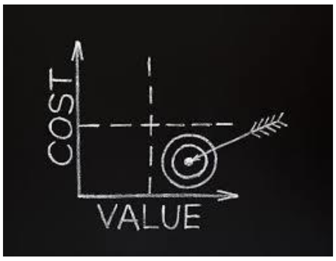

Strategic Issues

Optimisation also enables the retailer to explore various options of a more strategic nature. For example:

- If we are picking from stores (dark or otherwise), what are the optimal catchment areas?
- What is the optimal fleet size? Vehicle size?
- What delivery windows should we offer? Late night, early morning, weekends?
- Should we use contractors for deliveries? When?

The model that we build for operational or pricing purposes can also be used, with the right optimiser, to investigate these very important and cost sensitive issues.

Potential Benefits

Home delivery is poised to enter a new era. Volumes are increasing and supply chains are beginning to change:

- From 100% pick from store, retailers are moving to dark stores
- From home delivery only to workplace, lockers and other locations
- Order lead times and preferred delivery windows are shrinking
- Customers are less tolerant of non-delivery, for whatever reason

It is conceivable that, in the future, home deliveries could be handled by distribution centres.

All of this points to a world where customers have to be served, faster, more precisely and retailer has more options to consider.

What is Required

Home delivery is a challenging optimisation problem. It requires:

- The fundamental capability to allocate and schedule deliveries, and create routes simultaneously
- To do so in a reliable, scalable manner
- To be able to re-optimise, quickly, as things change (new deliveries, delays, traffic, breakdowns, etc)
- To deal with complex constraints: delivery windows, resupply, fatigue, enterprise bargaining rules, company policies
- The speed and reliability to back-end the customer ordering website to offer delivery windows and calculate prices
- The ability to re-optimise to calculate true cost to serve

The Opturion Platform has been proven in this harsh environment.

Getting Started

Home delivery is moving too fast for the traditional software development cycle of specify/tender/select/build. Two and three year build times are unacceptable (and quite ridiculous) is a business that is completely different to the one a few years ago, and in a few years it will be different again. By the time the system is working, it is either solving the wrong problem or will have gone through numerous costly changes.

We offer a different approach. Given the data, we can demonstrate all the features above for a catchment area in a matter of weeks, and get the system working, using real time data, in less than 3 months, with all customer feedback incorporated.

Once a single catchment is working, this can be quickly replicated across the enterprise. From zero to fully operational in a matter of months! Business evolution and change can be similarly handled; changes are rapidly prototyped and proven in one catchment and then rolled-out.