GIVING customers the car they want

Whenever vehicle manufacturers (VMs) ask their marketplace about the purchasing experience, invariably they find that customers would prefer to receive their vehicles in shorter time than the VM is able to deliver and that they would prefer not to compromise on the vehicle's specification. Consequently, VMs are constantly looking for ways to shorten the cycle from order taking to delivery and of removing compromise. The arrival of the Internet and integrated information systems has given rise to a form of order fulfilment termed virtual-build-to-order, the proponents of which see as offering a means of improving order fulfilment performance on both fronts — closer to the desired specification and delivered within an acceptable timeframe.

From the VM's perspective the problems are the level of variety and customisation potential to offer, and how to develop sufficiently flexible order fulfilment processes to maximise customer satisfaction while containing the costs.

THE VARIETY-CUSTOMISATION PROBLEM

The competitive automotive marketplace demands that all major vehicle manufacturers offer large variety envelopes across their major platforms. Although some manufacturers fare better than others in managing variety, both at the design stage and operationally, it is generally accepted that extensive variety results in additional operational costs.

When a customisation service is offered, the customer can select, combine and configure attributes of a vehicle rather than accept a selection from a pre-configured set of variants. However, allowing significant levels of customisation increases complexity further. For instance, forecasting and supply chain sourcing are more difficult. Methods have to be put in place to ensure that the customer specification is recorded accurately and is a valid configuration. On the positive side, the more demanding customer may generate higher margins and thus is particularly attractive in today's low-margin passenger car markets. So how much customisation should be offered? A conceptual model we advocate is illustrated in Fig 1.

PHILIP BRABAZON AND **BART MACCARTHY** DESCRIBE AN EMERGING AUTOMOTIVE ORDER-FULFILMENT STRATEGY

The figure illustrates that, as the number of customisable attributes increases the producer's costs tend to increase, slowly at first but then at an increasing rate. When viewed from the perspective of the utility value to the customer then certain customisation dimensions will be highly valued by a wide range of customers whilst others will be viewed as less valuable. The more value enhancing customisation dimension we term "key value attributes". Fig 1 is indicating that the point that maximises the differences between customer-perceived values and costs to the producer should be considered in judging how much customisation choice to offer.

For many reasons major automotive OEMs operate with extensive product pipelines and are forced to produce essentially to forecast across a theoretically large envelope of variety. Can this be exploited to satisfy a customer?

Given there can be very many possible vehicle specifications, even if most of these variants are seldom or never chosen by customers, the chance of a dealer having on its forecourt just the vehicle wanted by a customer is slim. Why not improve the chances by checking with nearby dealers? Better still, why not check with all the dealers around the country? Barring the challenge of getting competing dealers to collaborate, this method is catching on and is termed locate-to-order. Unfortunately, although a much bigger pool of vehicles is available to the customer, the amount of variety in the pool does not increase by a similar amount for the simple reason that most dealers are ordering a similar mix of vehicles.

THE VBTO CONCEPT

Before it became realistic to integrate business processes, production management tended to be segregated from the sales and order processing systems. The improvements in information technology and the internet are allowing them to be blended together so the ordering system can at least view, and perhaps even modify, the production schedule. This is a radical change and it has not taken long before a new name has been coined – virtual-build-to-order (VBTO). In the VBTO system a customer can be linked to a product that is not yet manufactured. This is nearly (i.e. virtually) the same as having a new product built to order, and the fact that the vehicle doesn't exist taps into the ideas of virtual reality. The key idea in VBTO is to exploit the variety that exists in automotive pipelines, both physical and planned, for individual customers, dealers, distributors and the manufacturer.

Introducing VBTO doesn't get rid of stock, nor does it eliminate fully fledged build-to-order (BTO). Stock is needed as some customers want to buy 'off the lot' as they say in the US. For these, speed of purchase is important and some compromise in the specification is tolerable. Other customers will have unusual tastes and the chances of finding a matching vehicle anywhere in the pipeline are low, hence BTO must remain an option. However, by opening up the pipeline VM are not only easing the way for customers to buy from future stock rather than from finished stock (which brings economic gains to the VM), they are also improving the chances of customers receiving the exact product they seek and, on average, sooner than is typical in current order fulfilment systems.

A VM's production pipeline can be long, of the order of three or more months and account for upwards of 100,000 vehicles. Nevertheless, a simple calculation shows that the chance of finding a match is only around 10% if the product range has in the order of one million variants – a figure that is realistic given the possible number of body shapes, powertrain types, wheel styles, gadget options and trim choices and so on, let alone paint colours. This probability is true only if the entire pipeline is as yet unallocated to other customers. If half of the products are sold already the chance of a match also halves. This simple calculation assumes all variants are in equal demand, but when a more typical Pareto type of distribution is used and the system is simulated, the likelihood of a match for a customer doesn't improve greatly. Furthermore, when you simulate this system, even when there is a matching vehicle available it is probably in the upstream portion of the pipeline since products that are reaching the end of the pipeline are likely to have been allocated to earlier customers.

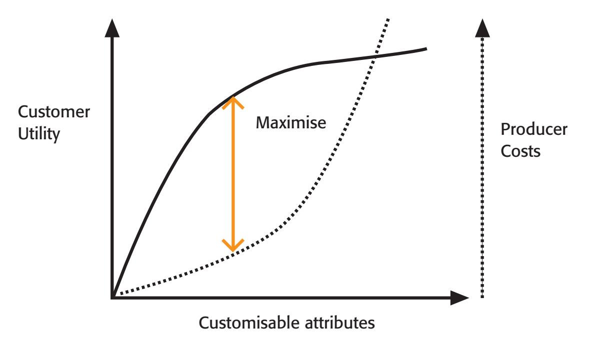

Fig 1: Comparison of fulfilment methods

FLEXIBLE PIPELINES

If there is some flexibility in the planning portion of the pipeline such that products can have their specifications reconfigured to some extent, the chances of finding a product for a customer greatly improves. It takes only a little flexibility for the matching probability to improve significantly. If the ability to reconfigure products can be sustained to a late point in the pipeline, say a day or a few days before vehicle assembly, customers can be matched to products nearing the end of the pipeline rather than to products at the start of the pipeline. →

Even if the pipeline is flexible, there will be times when a customer's request is so uncommon that it is unattractive to reconfigure any of the available products nearing production. In this situation one option is to request a BTO product, but there is a risk the customer won't be willing to wait. Other than persuading them to compromise on their specification, there may be two options open. One is to allocate them a product with the features they want but of a higher quality grade and the other is to allocate them a product that has the features they want but has also other features they have not asked for. The former is akin to a substitution strategy and the latter a redundancy strategy. If the customer cannot be persuaded to pay for the substituted or redundant features, the choice of how to fulfil this customer comes down to economics. Given the nature of today's automotive pipelines, the cost of reconfiguration is likely to increase as a product approaches production, whereas the cost of substitution and redundancy are constant. Hence, it is likely that as a product progresses along the pipeline there comes a point when it would be cheaper to leave a product unchanged than to incur the cost of reconfiguring it (see Fig 2).

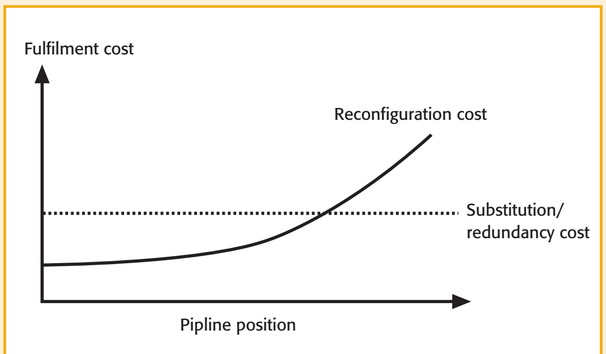

Fig 2: Comparison of fulfilment methods

From a cost perspective the early reconfiguration is attractive. But to the customer this means a longer wait. It would be ideal if the cost of reconfiguration remained low for as long as possible. This would be the case if the final identity of the product could be open until near the end the pipeline. If this were true, the product would have a 'postponed' reconfiguration cost curve rather than one of the decoupled, gradual or ingredient forms shown in Fig 3.

In a complex product like a passenger vehicle a reconfiguration cost curve can be plotted for each of the main assemblies or modules. The degree of independence between these curves will depend on the extent to which they share supply resources.

Although most customers may not be prepared to wait for a product to be processed through the entire length of the pipeline (i.e. a BTO product), the least cost strategy for the VM is to convince them to wait as long as possible so as to keep the cost of reconfigurations low. Of course competition scuppers this strategy and VM, in partnership with dealers, either have to sweeten the deal or lose the customer, unless they accept the reconfiguration hit or cleverly engineer the vehicle and the inbound supply resources to permit costless reconfiguration.

However, it is probable that not all customers need or want their vehicle as soon as possible. Some may plan ahead, some may not be in a hurry. If the VM could establish each customer's delivery expectations they could alter how they allocate product from the pipeline accordingly. Even though very few customers may be prepared to wait for a vehicle from the very start of the pipeline (a BTO product), a significant proportion may be willing to wait for product from the first third of the pipeline, and another sizeable proportion might wait for product from the middle section of the pipeline. If the VM

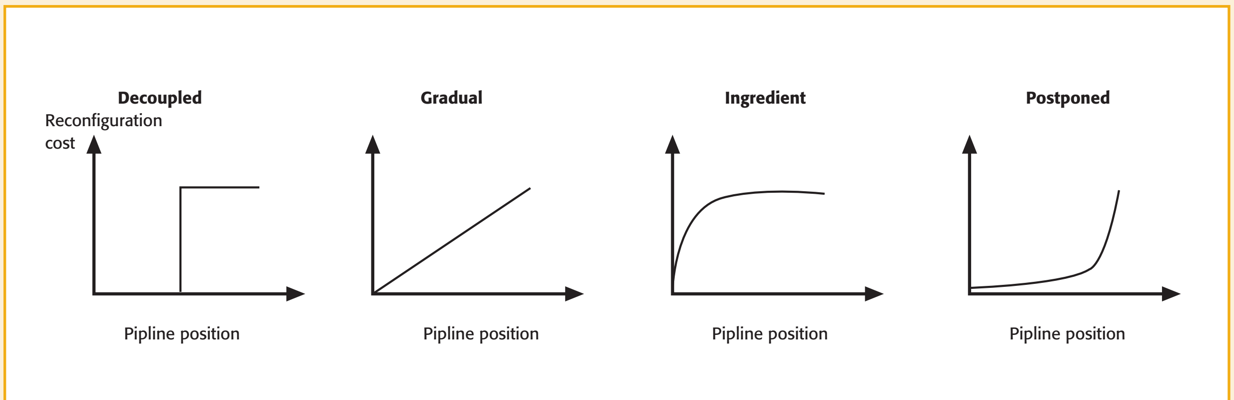

Fig 3: Signature reconfiguration cost curves

- **Decoupled:** a feature or product starts as generic and then at a point along the process becomes a specific variant, after which the cost of changing the specification is high;
- **Gradual:** as a feature or product progresses along the cost of changing the specification rises steadily;
- **Ingredient:** from an early point along the pipeline the cost of changing the specification is high;
- **Postponed:** not until late in the pipeline does the cost of specification change become significant.

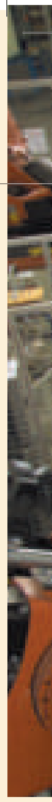

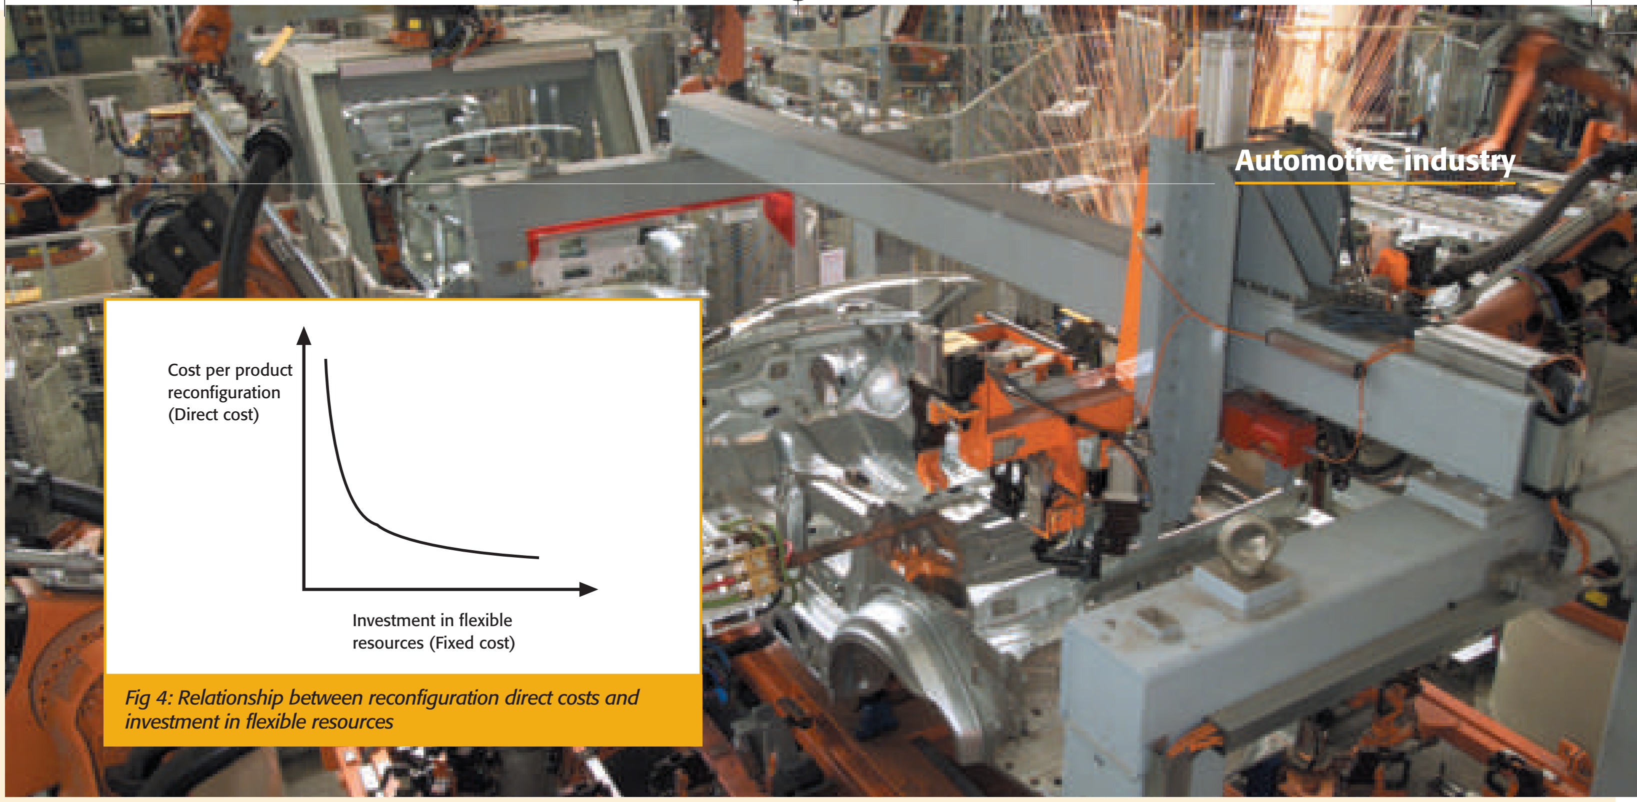

Fig 4: Relationship between reconfiguration direct costs and investment in flexible resources

could distinguish these customers, any unallocated product in the final section of the pipeline could be reserved for the group of customers who are most intolerant to waiting. Therefore the objective for managing the pipeline in a VBTO manner is to satisfy customers in terms of matching their specification and lead-time expectations while keeping on top of the costs of fulfilment. From an operations management perspective this is a cost minimisation problem with multiple constraints. The pipeline is dynamic with products moving along the pipeline with new products entering and others leaving. The constraints on the pipeline are complex. For example, there may be limits on the proportion of products in any given period that can have a certain feature due to supply capacity ceilings, or rules governing how products can be sequenced through assembly. Even more complexity can be added by some constraints changing over time. All of this means that finding the best vehicle for a customer is not straightforward, and requires complex optimisation logic and algorithms.

CUSTOMER SERVICE

One of the operational questions for the VM is whether they should handle each customer in turn or pool customers into a batch and rank them according to their fulfilment preferences before allocating them to products in the pipeline. From the point of view of customer service the first is attractive as in this scenario the customer, whether in the showroom or at their computer at home, is found a product as soon as they have entered their selection. This immediate response would not be given in the second scenario. Instead the customer is added into the pool and only when the pool is sorted and searches made would a firm link be formed between customer and vehicle. While less attractive, pooling customers may offer benefits to customers as well as the VM, particularly to the customers wanting a short delivery time.

The VBTO system is being researched at the University of Nottingham, with the aim being to study and quantify the operational benefits it brings, a significant part of which is involving the development of the simulation and analytical tools to analyse and optimise its operation. A practical problem that has to be solved is how to quantify the cost of reconfiguration?

Products and supply chains have not typically been viewed in this way. It is not a simple matter of collecting existing cost data. The first challenge is to ascertain how reconfiguration is being enabled. One or more methods may be being used e.g. having a standby stock of components ready for swapping, or overtime capacity, or process flexibility, or highly modular product architecture. What becomes apparent is that there can be two components to reconfiguration cost. One is the variable cost that is incurred when a reconfiguration is enacted, the other is the fixed cost of being ready to reconfigure. The balance between fixed and variable will differ for each method and there is likely to be a relation between the investment in resources to enable reconfiguration and the cost per reconfiguration (Fig 4). If this is the case there will be an optimum investment that balances the capital and operating expenditure and the objective for an enterprise is neither to over-invest nor under-invest in resources to enable reconfiguration. It would be undesirable if an enterprise reduced the cost of reconfiguration through investment, only for the sales function to persuade customers to compromise on their specification. In the same vein, if the cost of reconfiguration is high it would be inadvisable for a customisation service to be over-played to customers. One of the lines of research at Nottingham is to develop the protocol for profiling and quantifying the reconfiguration cost curve at the level of an assembly and hence open the way for investment analysis. ■

Philip Brabazon is Research Fellow in Mass Customisation and Bart MacCarthy is Professor of Operations Management at Nottingham University Business School. For further details contact Philip Brabazon or visit www.mcrcnottingham.org