Brabazon, Philip G. and Woodcock, Andrew and MacCarthy, Bart L. Intelligent pipeline control - a simulation study in the automotive sector. In: Fifteenth International Working Seminar on Production Economics, March 3-7, 2008, Congress Innsbruck, Innsbruck, Austria. (Unpublished)

Intelligent pipeline control - a simulation study in the automotive sector

Philip G. Brabazon, Andrew Woodcock, Bart L. MacCarthy

Mass Customization Research Centre, Nottingham University Business School,
Jubilee Campus, Nottingham, NG8 1BB, UK
philip.brabazon@nottingham.ac.uk; andrew.woodcock@nottingham.ac.uk; bart.maccarthy@nottingham.ac.uk

Abstract

Automotive producers are aiming to make their order fulfilment processes more flexible. Opening the pipeline of planned products for dynamic allocation to dealers/ customers is a significant step to be more flexible but the behaviour of such Virtual-Build-To-Order systems are complex to predict and their performance varies significantly as product variety levels change. This study investigates the potential for intelligent control of the pipeline feed, taking into account the current status of inventory (level and mix) and of the volume and mix of unsold products in the planning pipeline, as well as the demand profile. Five 'intelligent' methods for selecting the next product to be planned into the production pipeline are analysed using a discrete event simulation model and compared to the unintelligent random feed. The methods are tested under two conditions, firstly when customers must be fulfilled with the exact product they request, and secondly when customers trade-off a shorter waiting time for compromise in specification. The two forms of customer behaviour have a substantial impact on the performance of the methods and there are also significant differences between the methods themselves. When the producer has an accurate model of customer demand, methods that attempt to harmonise the mix in the system to the demand distribution are superior.

Keywords: order fulfilment, automotive.

1. Introduction

The level of product variety on offer from the large automotive producers, particularly on passenger vehicle models, can be very considerable. In coping with a wide product range premium producers are moving to fulfil the majority of their customers by building to order (BTO) [1] but most of the mainstream large producers use several fulfilment mechanisms. Retail customers are served by dealers and in European markets it has become common practice for the dealer to be able to search the stocks of other dealers as well as their own, and search the vehicles scheduled for production [2]. If no vehicle is found they have the option to request a BTO vehicle. A schematic of this multi mode open pipeline fulfilment system is in Figure 1.
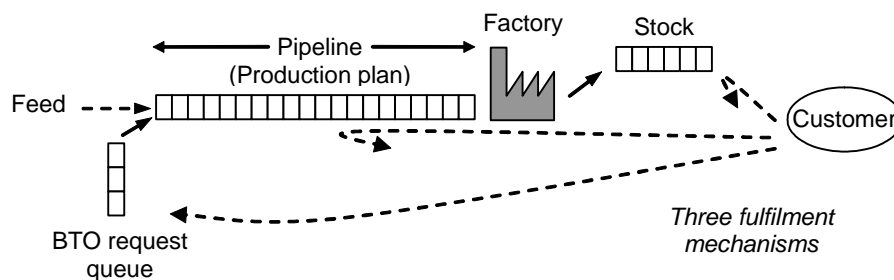

Figure 1. Schematic of the order fulfilment model with three fulfilment mechanisms

From an operations management perspective the multi-mode fulfilment system is interesting and potentially attractive to stakeholders in the system including the producer, dealers and customers. The system has a stock of unsold vehicles which is replenished from the factory, the production plan for which is typically mapped out for several weeks into the future. As can be expected, the producer is concerned with the volume and composition of stock, wanting these finished vehicles to be of an appropriate mix to satisfy as high a proportion of customers as possible. If the pipeline is closed from disturbance the mix in stock could be predicted using standard inventory analysis, assuming the customer demand for each product variant is known accurately. However, in multi-mode fulfilment the pipeline is open and hence a fraction of vehicles in the plan will be sold before they reach the

P G Brabazon, A Woodcock and B L MacCarthy

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008

factory and so do not replenish stock. Previous research has shown the volume and mix of stock in an open pipeline system is different from the stock in a conventional system with a closed pipeline [3]. That research uncovered inherent and fundamental behaviour of the fulfilment system but did not look at how the system could be controlled. This is the focus of the current research.

A producer may wish for customers to find the exact product variant they are seeking without waiting, i.e. the majority are fulfilled from stock. However, as variety increases, a point is reached where the volume of vehicles in stock and pipeline are fewer than the number of variants on offer. Whatever process or rule the producer uses to feed the pipeline, in this circumstance either some proportion of customers must wait for a BTO product, or they must be willing to compromise on vehicle specification. In this study these two behaviours are modelled explicitly to assess their impact on fulfilment performance across a wide range of variety levels.

The objective of this study is to test methods for selecting the product variants to feed into pipeline. To do so a discrete event simulation has been created which models the pipeline as a sequence of p products. At each time step of the simulation the products increment one position along the pipeline with one being fed into the upstream entrance of the pipeline and one leaving the downstream end. The exiting product goes into stock unless it has been sold already in which case it is removed from the system. Customer arrivals are synchronised with the incrementing pipeline, with one customer served in each time period. Every customer is allocated a product, either from stock or the pipeline or by requesting a built-to-order product. As customer arrivals are synchronised with production the number of <u>available</u> products in the system remains constant, implying that the producer's forecast is accurate in terms of volume. Therefore, if the system is primed with a pipeline full of p products and none in stock, the first customer will take one and reduce the count to $p-1$ (conditional on there being a match), but it will return to p when the next product enters the pipeline.

Although the count of available products is constant their location in the system depends on the level of variety on offer to the customer. When only a few variants are on offer, many customers are fulfilled from stock, but when variety is high only a few will find a suitable product in stock and a high proportion will need vehicles built-to-order. In the former situation the level of stock is low and most of the available products will be in the pipeline. In the latter situation the available products are mostly in stock, and the pipeline is conveying BTO products. These conditions are illustrated in Figure 2.
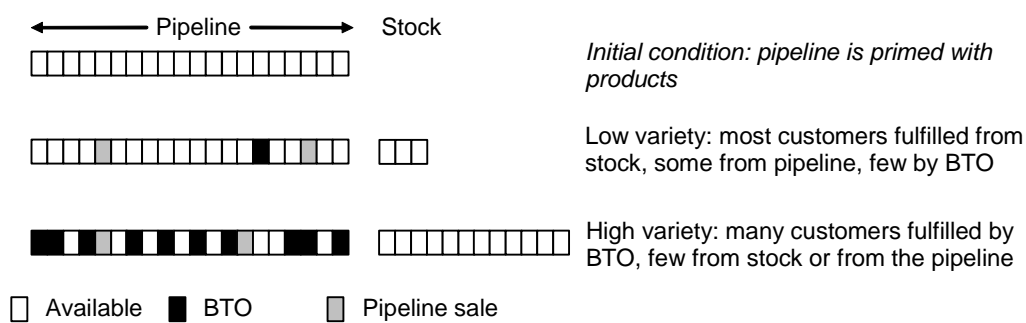

Figure 2. Indicative location of products in low and high variety conditions

Five Methods for selecting the next product to feed into the pipeline have been developed and are compared to a random feed. Four of the methods are based on comparing the mix of available products in the pipeline and in stock to a 'target' distribution. The fifth method is a simple but pragmatic rule, which is to feed in the variant the last customer wanted.

In reality the producer has the challenge of estimating the relative demand for each variant but in this study we make the producer's target distribution identical to the demand distribution. Each product variant has a unique number to represent its specification. The difference in specification between two variants is the difference in their numbers. To illustrate, the variant #47 is one step different from #48,

P G Brabazon, A Woodcock and B L MacCarthy

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008

and 51 steps different from #98. This property is used when customers are modelled as being willing to compromise.

The relative demand for each variant follows an 80/20 distribution, i.e. 20% of the variants account for 80% of demand as illustrated in the right plot in Figure 3. This is modelled in the simulation using a Beta distribution with the shape parameters set to 1 and 7.5.
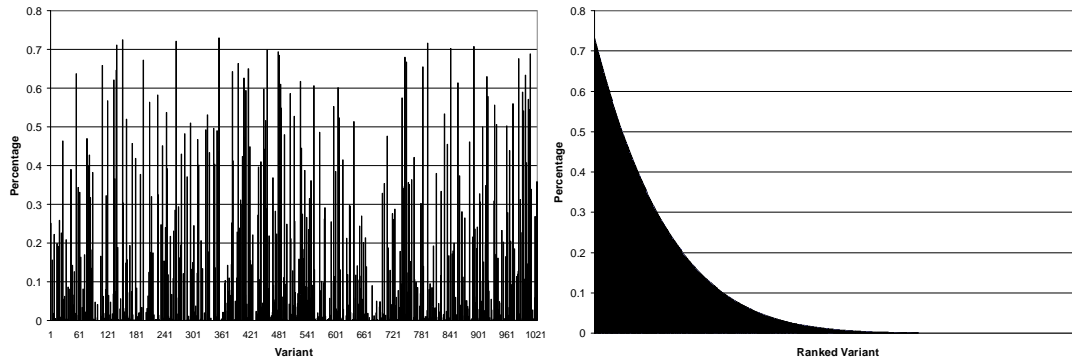

Figure 3. Demand for each variant in number order (left plot) and ranked by demand proportion to show the shape of the 80/20 distribution (product range 1024)

In the study a range of variety levels are simulated from 2 to 16,384 (i.e. from 2^1 to 2^{14}) and in all cases a skew equivalent to 80/20 is applied. Figure 3 shows the relative demand for variants when there are 1024 variants and it is important to note that demand per variant is randomised to avoid a correlation between variant specification and variant demand. This is to emulate the real world situation in which the most commonly requested variants from a product range differ greatly.

2. Description of the pipeline control methods

This section describes the control methods and how they are implemented. All the methods function in the same way, in that they select one product to feed into the pipeline. Common symbols are given in Table 1.

Symbol	Description	Symbol	Description
i	Variant identifier	l	Number of products held in the pipeline
s_i	Probability of stock out on variant i	p_i	Probability of a customer seeking a variant i
m	Number of variants	a_i	Volume of variant i in pipeline and stock
c	Number of customers	A	Volume of products in pipeline and stock, $A = \sum_{i=1}^{m} a_i$
		d_i	Volume of demand for variant i , $\overline{d}_i = p_i A$

Table 1. Symbols

2.1 Method 1: Random feed. The next variant to be fed into the pipeline is chosen at random from the target distribution, which is modelled as a Beta distribution.

2.2 Method 2: Reduce stockout probability. The next variant to be fed into the pipeline will be the one that has the highest probability of *stocking out*. Probability s_i of a variant stocking-out is calculated as the probability that demand for variant i will exceed its current availability a_i after c customers are processed (where c is equal to the sum of all products currently in stock and pipeline). The number of customers demanding variant i are estimated using a binomial approximation, i.e. $d_i \sim B(c, p_i)$, so the value of (1) is calculated using the cumulative density function of the binomial

3

P G Brabazon, A Woodcock and B L MacCarthy

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008

The University of Nottingham

distribution. The variant to be fed into the pipeline satisfies $\arg\max\{s_i\}$ i.e. the variant with the highest probability of stocking out.

$$s_i = 1 - P(d_i \leq a_i) \tag{1}$$

2.3 Method 3: Reduce weighted error from target distribution. This approach considers the error between the actual number of a particular variant in both stock and pipeline a_i and the expected demand for that variant $\bar{d}_i$. The error is weighted according to the demand for the variant in (2). The variant to be fed into the pipeline satisfies $\arg\min\{e_i\}$.

$$e_i = p_i |a_i - \bar{d}_i| \tag{2}$$

2.4 *Method 4: Reduce distance (to reduce compromise).* This method applies the concept of compromise distance to select a variant. Consider the situation in which a producer can stock only one variant *k*. All customers will receive variant *k* regardless of which variant they request. To minimise the compromise of the customer population, the producer selects the variant which minimises the average expected distance defined in (3).

$$\overline{dist} = \frac{1}{c} \sum_{i=1}^{m} cp_i |spec_i - spec_k| \tag{3}$$

The procedure for implementing the method is as follows:

- The current holding of each variant in stock and pipeline is a_i. Sum to find total, A
- Add 1 to the holding and estimate the expected number of customers per variant, $\bar{d}_i = p_i(A+1)$
- Select the first variant and add 1 to the volume of this variant, i.e. $a_1 = a_1 + 1$.
- Then inspect each variant, attempting to fulfil the expected customer demand $\bar{d}_i$, firstly from a_i. The fulfilled volume is $f_{i,i}$. If $f_{i,i}$ is less than $\bar{d}_i$, then try to fulfil the remainder from a_{i-1} which will be $f_{i,i-1}$, and if some remains still then fulfil from a_{i+1}, and so on, following the general sequence of filling from a_{i-j} then a_{i+j} until $\sum_{j=0}^{\max\{i,m-i\}} f_{i,i+|-j} = \bar{d}_i$
- The 'distance' calculation sums the product of volume fulfilled and difference between variants, i.e. $dist_i = \sum_{j=1}^{m} f_{i,j} |i - j|$ noting that when $j = i$ the distance is zero, hence if the volume of stock and pipe is distributed over the variants in an ideal way, the 'distance' calculation would return 0.
- Repeat for all variants and select the variant that gives $\arg\min\{dist_i\}$.

The number of calculations is proportional to m^2 so to obtain results at higher variety levels a stopping rule is implemented but even with this included results have not been obtained for the two highest variety levels of 8192 and 16384.

P G Brabazon, A Woodcock and B L MacCarthy

The University of Nottingham

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008

2.5 Method 5: Increase forward sales coverage. The expected forward sales cover fsc_i of a variant is calculated using the binomial approximation, i.e. $d_i \sim B(c, p_i)$ and the property that the expected successes for an outcome is the product of the number of trials and the probability of success per trial. To calculate the forward sales coverage for each variant, 1 is added to the number of that variant currently in stock and pipeline. The variant selected to wholesale is the one which gives $\arg\min\{fsc_i\}$ where fsc_i is calculated using (4).

$$fsc_i = \frac{(a_i + 1)}{p_i} \tag{4}$$

2.6 Method 6: Follow the previous customer's request. In this method the sequence of wholesaled products repeats the sequence of customer orders.

3. Priming of the pipeline

Four of the methods (2, 3, 4 & 5) can be compared by how they prime the pipeline. The priming process starts with an empty pipeline and the method selects the first product variant. This is fed into the pipeline and is taken into account when the method selects the next variant, and so on until the pipeline is full.

The plots in Figure 4 analyse the characteristics of the pipeline distribution once primed by each of the methods. The pipeline holds 1024 products and the target distribution is from Figure 2 above. The plots rank the variants by their demand fraction in descending order and measure the difference in frequency from the target distribution.

Three of the methods exhibit a similar form of saw tooth pattern. This is the consequence of the actual number of a variant in the pipeline being an integer, whereas the target frequencies are in fractions. In Methods 2 and 3 the teeth alternate between over- and under- representation of variants in the pipeline (with a value above zero indicating over representation). In Method 5 the teeth to the left are all for over-represented variants, but then begin to alternate. In all three methods the lower demanded variants to the right of the plots are under-represented.

The overall pattern of Method 4 is also a saw-tooth but the over- and under- represented variants are interleaved. A second notable difference from the other methods is that many of the lower demanded variants are over-represented.

Statistics about these differently primed pipelines are given in Table 2. Method 5 feeds fewer than 30% of variants (296) while Method 4 feeds in just under 50% (504). Furthermore, the variants fed in by Methods 2, 3 and 5 are the highest ranked variants whereas Method 4 spreads its selections from across the product range. In terms of evaluating the shape of the distributions against the target distribution, using the measure of mean squared difference the methods are ranked from best to worst as 2, 3, 4, and 5.

Method	Number of variants represented in the pipeline	Lowest ranked variant represented in the pipeline (ranked by demand)	Mean Square Difference from the target distribution
2	337	337	4.81E-04
3	304	304	7.21E-04
4	504	1022	4.70 E-03
5	296	296	1.18 E-03

Table 2. Analysis of how methods prime the pipeline, for variety 512

P G Brabazon, A Woodcock and B L MacCarthy

The University of
Nottingham

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck,
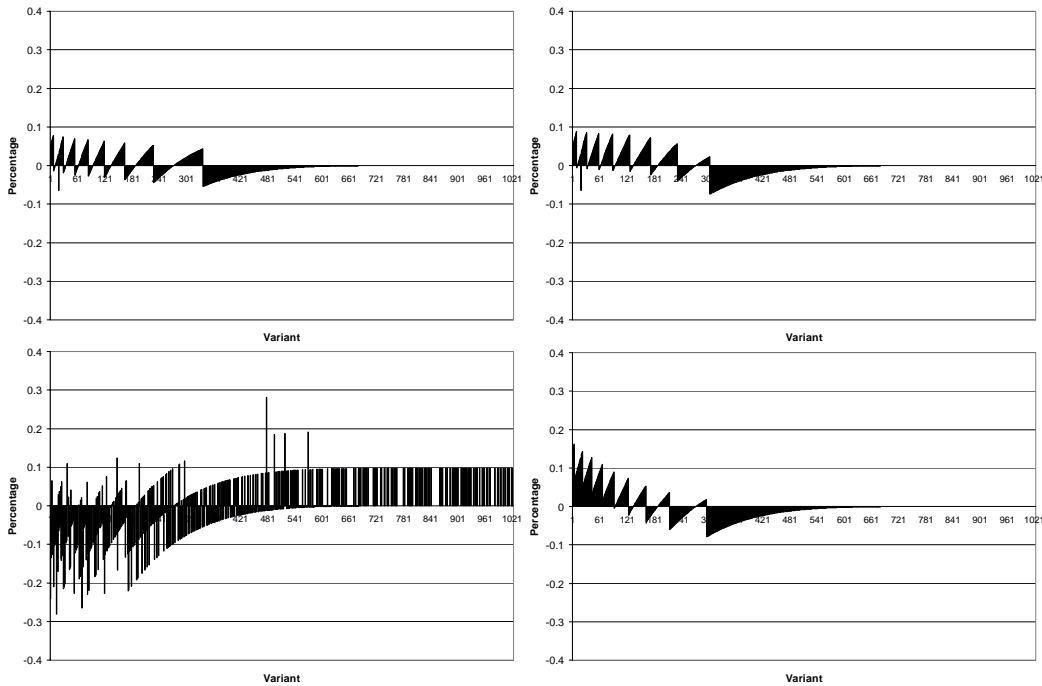*Innsbruck, Austria, March 3-7, 2008*

Figure 4. Differences in proportions of variants in a primed pipeline from the target percentage.
Clockwise from upper left: Method 2, Method 3, Method 5, Method 4

4. Analysis approach

The methods are tested under two customer behaviours. In one the customer must receive the variant requested, referred to as the *Exact match* search (and the oldest matching vehicle fulfils the customer). In the second the customer will trade-off waiting time and specification difference. This is denoted as the *Compromise search* and a customer is fulfilled by the oldest product giving the minimum value for (5).

$$score_i = \left| spec_i - spec_{requested} \right| + leadtime_i \qquad (5)$$

Results are collected using the batch means method with the initial transient discarded and data from 20 batches used to calculate metrics and confidence intervals (which are shown on many plots) as explained in [4]. Many of the results are plotted against the ratio of variety to pipeline length, denoted as the *v/p* ratio, which has been observed to allow fulfilment systems of different magnitudes to be compared [3].

5. Results

5.1 Zero initial stock. In this condition the pipeline is primed with products but there is no stock. The pattern of fulfilment when customers are fulfilled with the exact product they are requesting is given in Figures 5 to 8. Figures 5, 6 and 7 show the proportion of customers fulfilled from stock, pipeline and BTO, and Figure 8 plots their average waiting time. Evident from these four plots are:

- Methods 2, 3 and 5, which all try to match the target distribution, have similar performance and they are superior to the other methods. In the *v/p* range from 0.01 to 1 the contribution of each fulfilment mechanism is near constant. From ratios above 1, BTO fulfilment rises and other mechanisms reduce but are still significant.
- Each of Methods 1, 4 and 6 has a distinct pattern. Methods 4 and 6 have similar performance to 2, 3, and 5 at the lowest *v/p* ratios, but then diverge as *v/p* rises with lower stock fulfilment, higher BTO and longer waiting time. Method 6 diverges less compared to the substantial difference of Method 4.

6

P G Brabazon, A Woodcock and B L MacCarthy

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008

- The random feed (Method 1) has the poorest performance in terms of stock fulfilment which is less than half of the superior methods 2, 3 and 5. Average customer waiting time is an order of magnitude longer at low *v/p* ratios.
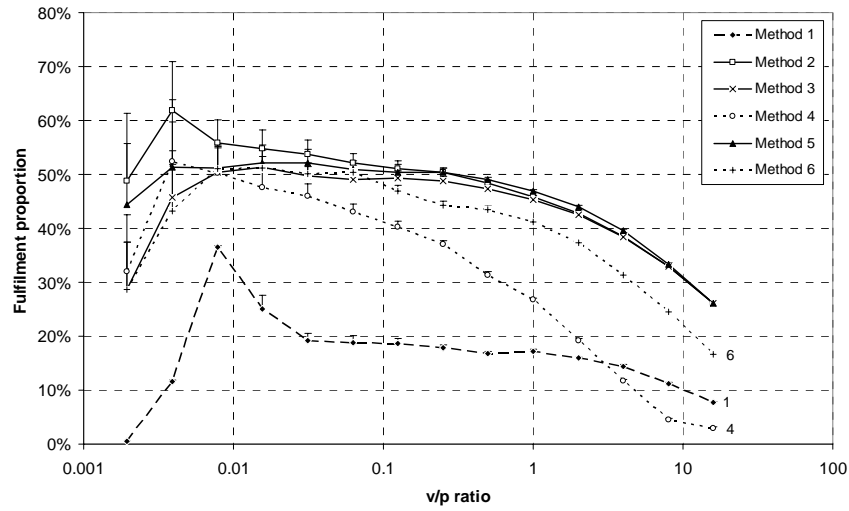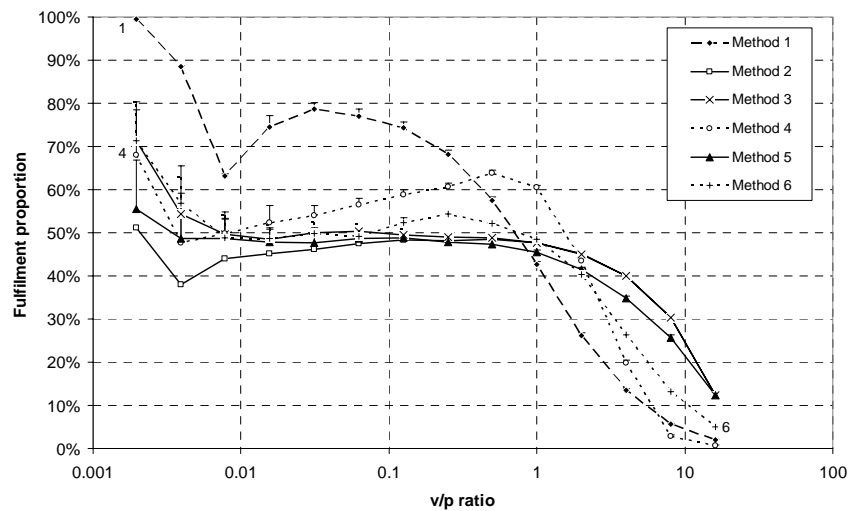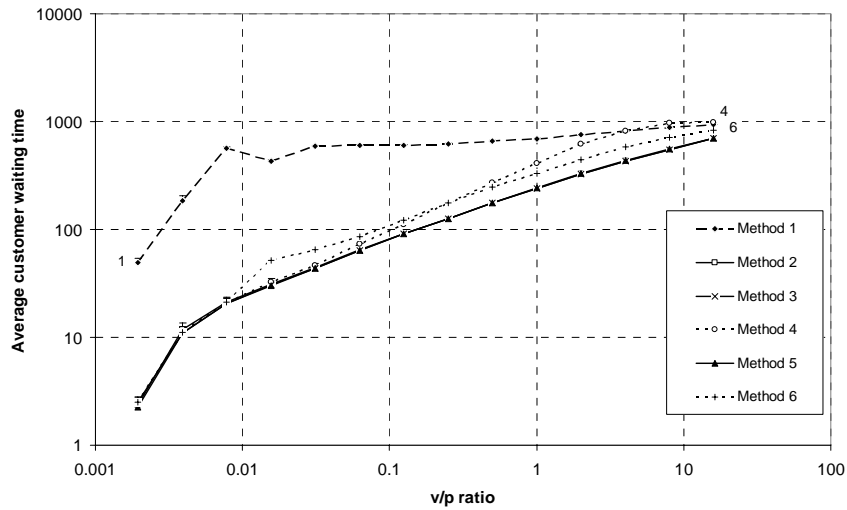

Figure 5: Fulfilment from Stock (Exact Match search)

Figure 6: Fulfilment from Pipeline (Exact Match search)
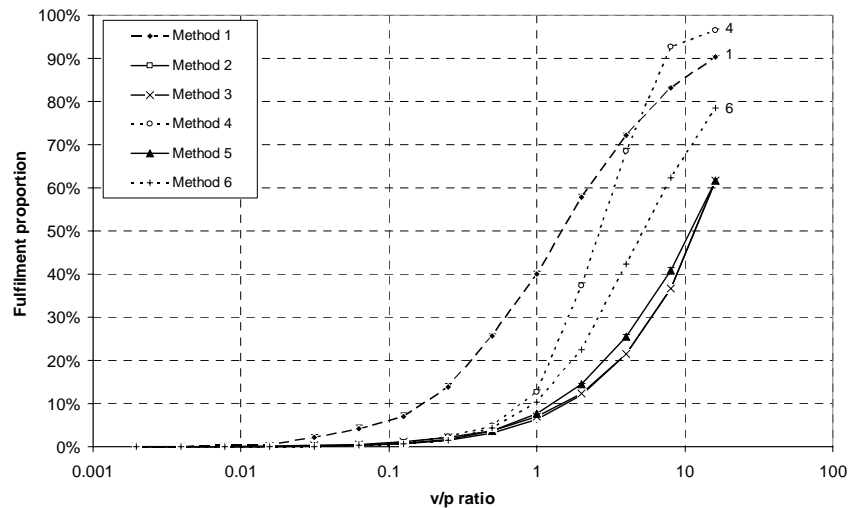
Figure 7: Fulfilment from Build-to-Order (Exact Match search)

7

Figure 8: Customer waiting time (Exact Match search)

Performance of the methods is greatly altered when customers are willing to trade-off specification and waiting time, as shown in Figures 9 to 11. In the variety range analysed there is no BTO fulfilment except for Method 1 at the highest variety studied (16,384 variants). In respect of the fulfilment mechanisms and waiting time all methods have similar performance, with Method 1 being the only one to trend away at higher *v/p* ratios. At the lowest *v/p* ratios only a small fraction are fulfilled from stock (Figure 9), but this is a little misleading since the majority of customers are being fulfilled from products just about to leave the pipeline which is evident from the plot of waiting time (Figure 10) which shows the average to be close to 1 time period at the lowest ratio (which is the waiting time for a stock vehicle).

The pattern of specification compromise depends on how it is measured. In terms of the difference in variant number, the amount of compromise rises as the *v/p* ratio rises (Figure 11, left plot) but when measured as a percentage of the product range, the compromise is greatest at a low *v/p* ratio and drops toward zero (Figure 11, right plot).
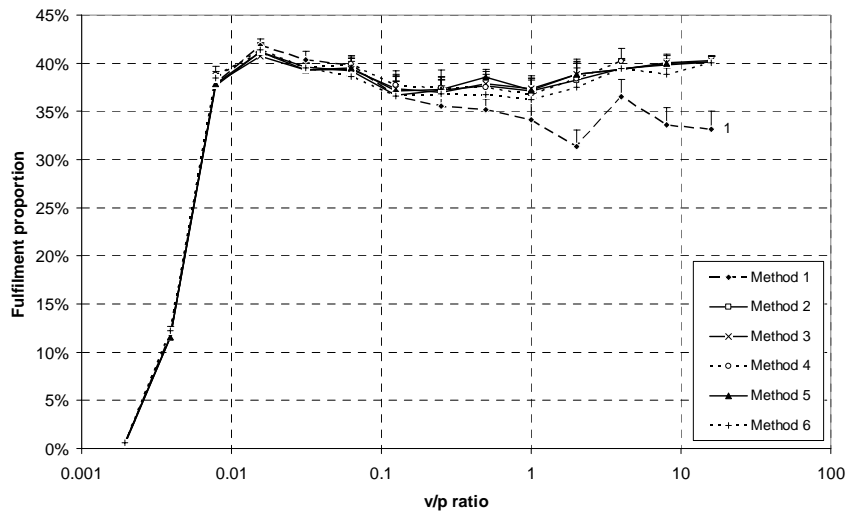

Figure 9: Fulfilment from Stock (Compromise search)

P G Brabazon, A Woodcock and B L MacCarthy

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008
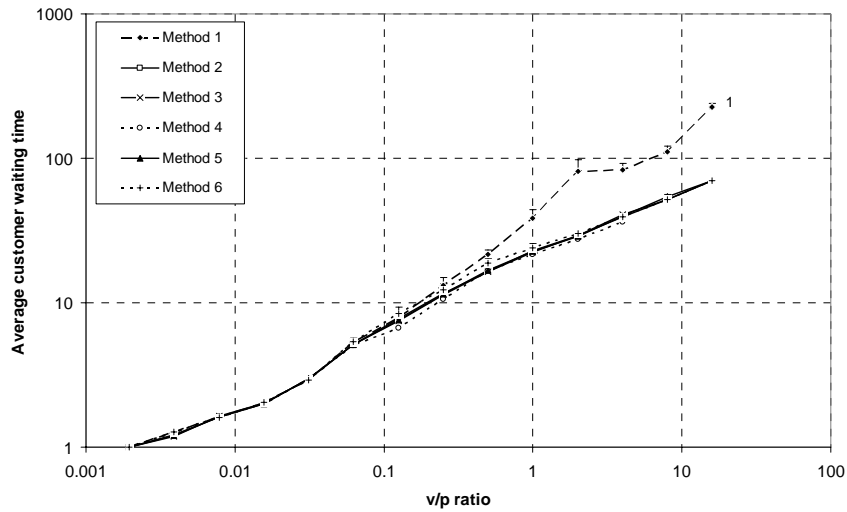
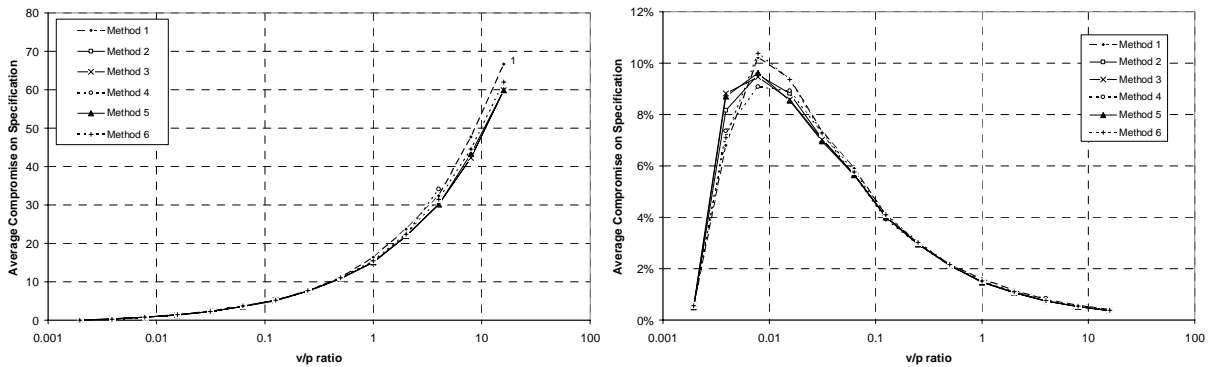*Figure 10: Customer waiting time (Compromise search)*

Figure 11: Customer compromise in specification, left: variant steps, right: percentage of product range

5.2   Increasing stock levels.   In this section the system is investigated further at the *v/p* ratio of 1 which corresponds to variety of 1024.   The issue considered is how the methods control stock mix, and this is investigated by priming the system with greater volumes of stock.   Once steady-state conditions are reached the methods are compared in respect of fulfilment from stock and customer waiting time.

Figures 12 and 13 plot the stock fulfilment proportions, the former for the exact match search and the latter for the compromise search. Figure 13 confirms the pattern observed above in that Methods 2, 3 and 5 are similar and they are superior to the other methods.   Method 6 is fairly close to them, differing by ~10 percentage points and converges to them at stock levels above 2000.   Method 4 also converges at this stock volume.   The random feed of Method 1 stands out as being a poor approach.   It does not achieve 90% stock fulfilment in the conditions analysed, but from extrapolation this method will require an order of magnitude more stock than Methods 2, 3 or 5.   Figure 13 also confirms the pattern observed earlier, with all but Method 1 having similar performance, though Method 6 has slightly less stock fulfilment when stock is below ~300.   In the exact match condition, the random method is poor, but it converges to the other methods at stock levels above 1000.   The data on customer waiting times in Figure 14 show similar differences between methods. Comparisons of the fulfilment curves in Figures 12 and 13 and the waiting time plots in Figure 14 highlight how substantial the impact on performance is of customer behaviour.   At the *v/p* ratio of 1, over 1000 products are required in stock for 90% of customers to be fulfilled from stock in the exact match search, but when customers compromise the same proportion is achieved with a stock of less than 200. In this condition the average compromise in specification is small at ~0.5% (Figure 15, left plot) and the maximum that any customer compromises is in the region of 5% to 6% (Figure 15, right plot).
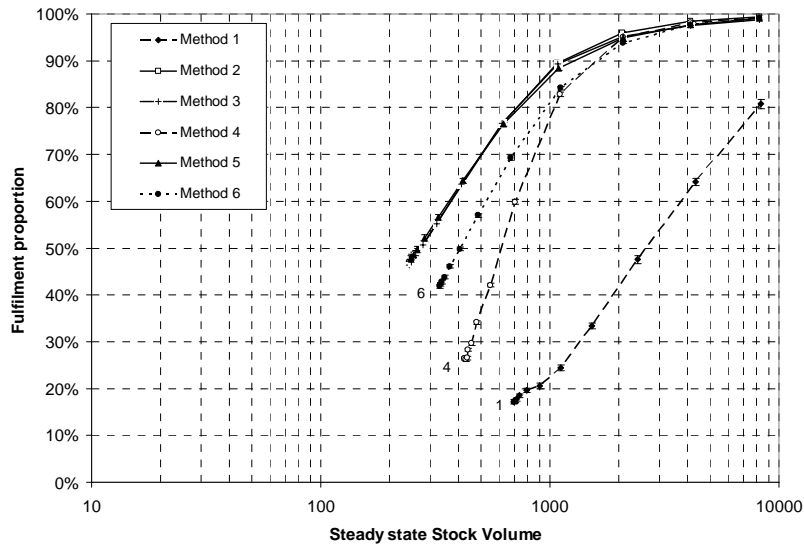
9

The University of Nottingham
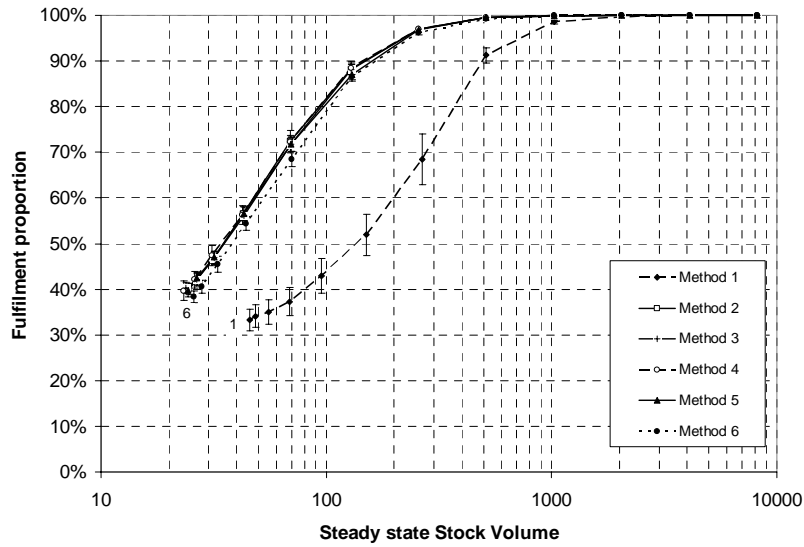
Figure 12: Fulfilment from Stock (Exact Match search)

Figure 13: Fulfilment from Stock (Compromise search)
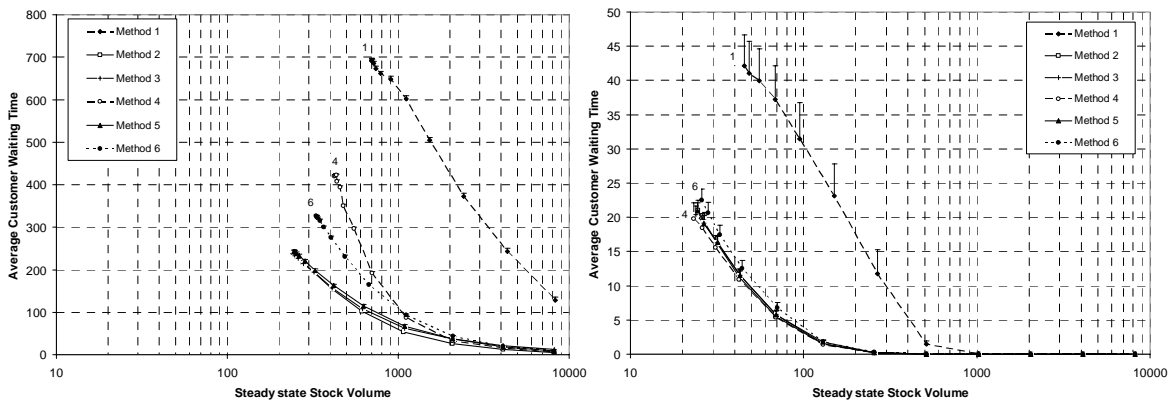
Figure 14: Customer waiting time, left: Exact Match search, right: Compromise search

P G Brabazon, A Woodcock and B L MacCarthy

*Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck,
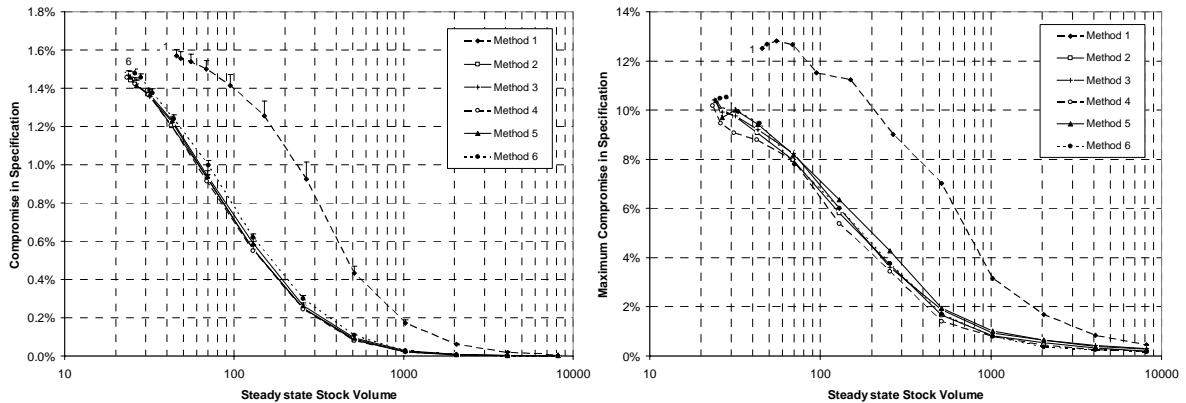Innsbruck, Austria, March 3-7, 2008*

*Figure 15: Average (left) and maximum (right) customer compromise in percentage specification
(Compromise search)*

6. Discussion

The results show in most conditions studied that intelligent methods are superior to the baseline random feed. In the more demanding situation of customers requiring an exact match some of these methods far exceed the random feed in terms of fulfilling from stock and customer waiting time. As can be seen in Figures 5 to 8, the benefits from these methods is seen across the full *v/p* range studied. Although the experimentation here extended to a v/p ratio of 16, the plots suggest that all methods may converge at a *v/p* ratio of 100 or higher. In contrast, when customers compromise, all of the methods, including the random feed, have near identical performance at *v/p* ratios below 0.1 and the divergence above this ratio is small. The comparison of the two customer behaviours shows how important their decisions are to system performance. It is anticipated a real customer population will have a mix of customer types and the implications of the relative proportions is an issue for further study.

Three of the methods – 2, 3 & 5 – attempt to harmonise the mix with the target distribution and they have similar performance. The results show these methods improve the stock mix compared to the random method and the greater proportion fulfilled from stock shortens customer waiting times. A further contributor to the reduced waiting is an improved mix in the pipeline, evidenced by histograms of where along the pipeline products are allocated to customers (Figure 16). The left plot is with the random feed and the bias is toward the upstream start of the pipeline while the plot for Method 2 is a near mirror image, with the bulk of allocations in the downstream half of the pipeline.
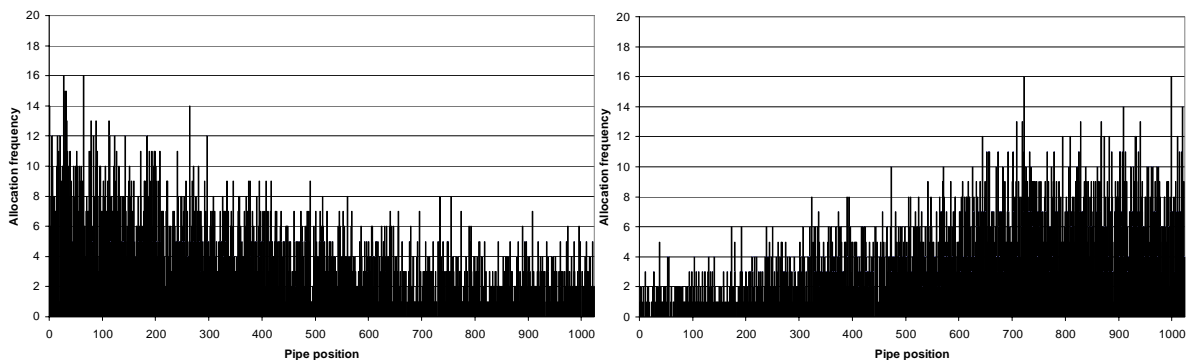

*Figure 16: Pipeline slot at which products are allocated to customers, left: Method 1, right: Method 2
(Exact Match search, variety 1024, 10,000 customers)*

Method 4, which looks to create a mix in the system to cover the full range of products, performs poorly at higher variety levels when customers must have the exact specification they are seeking. It is a method conceived for the compromise situation but it does not stand out as a superior method in those conditions.

11

P G Brabazon, A Woodcock and B L MacCarthy

Presented at: Fifteenth International Working Seminar on Production Economics, Congress Innsbruck, Innsbruck, Austria, March 3-7, 2008

Method 6 is notable in that although it underperforms the best methods when customers demand exact matches, its performance is considerably better than the random method and it would seem to be straightforward to implement. It is conceivable Methods 2, 3 and 5 will lose their superiority if the producer has an inaccurate forecast of the customer demand distribution. Method 6 is robust to this as it does not require a forecast. It will lag behind any change in customer tastes but this is a challenge to all forecasting techniques.

A question worth dwelling on is why Method 6 is superior to the random feed (Method 1). In Method 1 the sequence of variants requested by customers is random and the sequence of variants fed into the pipeline is also random. In Method 6 the pipeline feed follows the customer sequence, hence it can be considered to be random also. However, because the feed follows the customer sequence it avoids a phenomenon observed in earlier research on open pipeline systems [3] in which the mix in stock becomes unrepresentative of the mix fed in to the pipeline. Consider a situation in which two dice are thrown several time and one has a sequence of four or more '5s' while in the other's sequence there is no '5'. When the equivalent occurs in the open pipeline system, the result is that stock and the pipeline are stripped of '5s' and from then on any '5' entering the pipeline is likely to be sold to a customer before it can replenish stock. By following the customer Method 6 duplicates the run of '5' and prevents the stripping effect.

7. Conclusion

A simplified version of an open pipeline system used in the automotive sector has been studied using a discrete event simulation model. A set of methods for selecting products for manufacture have been developed and implemented. Using a number of performance metrics clear differences have been observed in the methods. When the producer has an accurate forecast of customer demand the performance attained by some methods is very much better than a random feed. The approach of producing the products requested by recent customers does not achieve the best results but may be a more robust method. Further research can study the implications of forecast error on the performance of the methods.

8. Acknowledgements

9. References

1. Meyr, H., 2004. Supply chain planning in the German automotive industry. *OR Spectrum*, 26 (4), 447-70.
2. Holweg, M., Pil, F.K., 2004. *The second century: reconnecting customer and value chain through Build-to-Order*. MIT Press, Cambridge, Mass.
3. Brabazon, P.G., MacCarthy, B.L., 2006. Fundamental behaviour of Virtual-Build-to-Order systems. *International Journal of Production Economics*, 104 (2), 514-24.
4. Law, A.M., Kelton, W.D., 2000. *Simulation modeling and analysis*. McGraw-Hill, Singapore.