JOSEPH D. BLACKBURN

The value of time (and the value of waste): time-based supply chain design

The strategic concept of time-based competition was introduced by George Stalk (1988) in an article entitled 'Time—The Next Source of Competitive Advantage'. Stalk (1988) argued that an organisation could achieve a powerful competitive advantage through the speed by which it responded to customers, developed new products and services, and moved products through the supply chain. Books by Stalk and Hout (1990) and Blackburn (1991) expanded on this theme, by describing how firms could develop and implement time-based strategies.

Much of the early work on time-based competition focused on the benefits of speed and implied that faster is better (Schmenner 1988). Firms such as Dell Computer, Wal-Mart, and Zara were used as examples to show how a business model built upon speed could propel a firm to a pre-eminent position in its industry. Dell replaced the traditional distribution model in consumer electronics with a build-to-order direct-to-customer model which yielded dramatic increases in speed and reductions in cost. Wal-Mart employed a faster supply chain to obtain shorter replenishment cycles, gaining them the fastest product turnover and profitability among mass-market retailers. Zara used 'fast-fashion'—more frequent product introductions and shorter lifecycles—to become a dominant fashion retailer.

Is faster always better? In *Clockspeed: Winning Industry Control in the Age of Temporary Advantage*, Fine (1998) observed that the standards for speed vary across industries, and that competitive advantages based on speed are temporary. Blackburn (2012) found that there are limits to time-based competition and that these limits vary across industries. Just as there are limits on the speed of physical processes quantified by laws of physics, there are limits on the speed of business processes defined by economics principles. For business processes, the limits on speed are imposed by a metric called the marginal value of time (MVT) and defined as the value gained by an incremental increase in the speed of the process. For example, if a firm could increase future profitability by USD 10,000 by reducing the length of time to process orders by one day, then MVT for order processing is USD 10,000 per day. Generally, improvements in the speed of operations processes are desirable up to a point: the point at which MVT equals the marginal cost of additional speed. Thus, the firm in our example could spend up to USD 10,000 to improve order processing speed by one day. Such points are dynamic, not static, because they change over time with technology, the level of competition, and consumer preferences.

Service processes provide good examples of MVT and the existence of economic limits to speed. In the operation of a large telephone call centre, the time required to respond to customers by a service agent is a critical performance metric. Management could improve the speed of response by increasing the number of agents (but at a cost)—and, with enough agents on hand, it could reduce the response time to virtually zero! However, there are clear economic limits to how much a telephone call centre would be willing to spend to further reduce customer waiting time. The economic limit is set by MVT, the value to customers of an incremental decrease in average waiting time. Telephone call centre management would only want to add service capacity up to the point at which the marginal cost of additional service equals the MVT for the reduction in customer waiting costs.

This article summarises over two decades of research in the area by the author and his colleagues (Blackburn 1991; Blackburn et al. 2004; Blackburn and Scudder 2009; Blackburn 2012) and shows how MVT can be used as a tool to develop effective supply chain strategies. In a supply chain, the average lead time is the amount of time a product spends end-to-end. Since it measures the value gained by an incremental reduction in the average lead time, MVT is a useful design parameter that dictates when it is important to design a supply chain for speed—and when it is not. To illustrate this proposition, this article summarises three research cases—all drawn from the author's previous research—that show the role of MVT in optimal supply chain design. The first is a sourcing decision case, where a firm must decide how to configure the supply chain from suppliers to minimise total acquisition cost (Blackburn 2012). The second case concerns the design of a supply chain for perishable food products, where MVT changes along the supply chain (Blackburn and Scudder 2009). The third case concerns the design of a reverse supply chain for returned products (Blackburn et al. 2004).

The time-based approach to supply chain design suggested in this article is a modification of a supply chain taxonomy first proposed by Fisher (1997). He classified products into two categories, *functional* and *innovative*, and proposed a simple dichotomy for designing supply chains based on product characteristics. Functional products have predictable demand and relatively long lifecycles. Innovative products have highly variable demand and short lifecycles. Fisher argued that *responsive* (time-based) supply chains are appropriate for *innovative* products and that (cost) *efficient* supply chains are appropriate for *functional* products. All three research studies summarised in this article show that the value of time offers an alternative way to view Fisher's strategic construct. Innovative products are very time sensitive and tend to have a high MVT—functional products have a low MVT. For products with high MVT, the supply chain needs to be designed for speed (responsive)—for products with low MVT, cost efficiency is more important than speed. The following three cases illustrate the use of MVT in developing optimal supply chain strategies.

Designing supply chains for sourcing and offshoring component parts—the case of the automotive industry

The growth of global supply chains appears to violate the principles of time-based competition. Over the past several decades, supply chains for product sourcing have become longer, and slower, as US firms have moved production or sourcing of components to China and other offshore locations (Blackburn 2012). Time-based strategies would dictate shorter, faster supply chains to improve replenishment times and lower inventory costs.

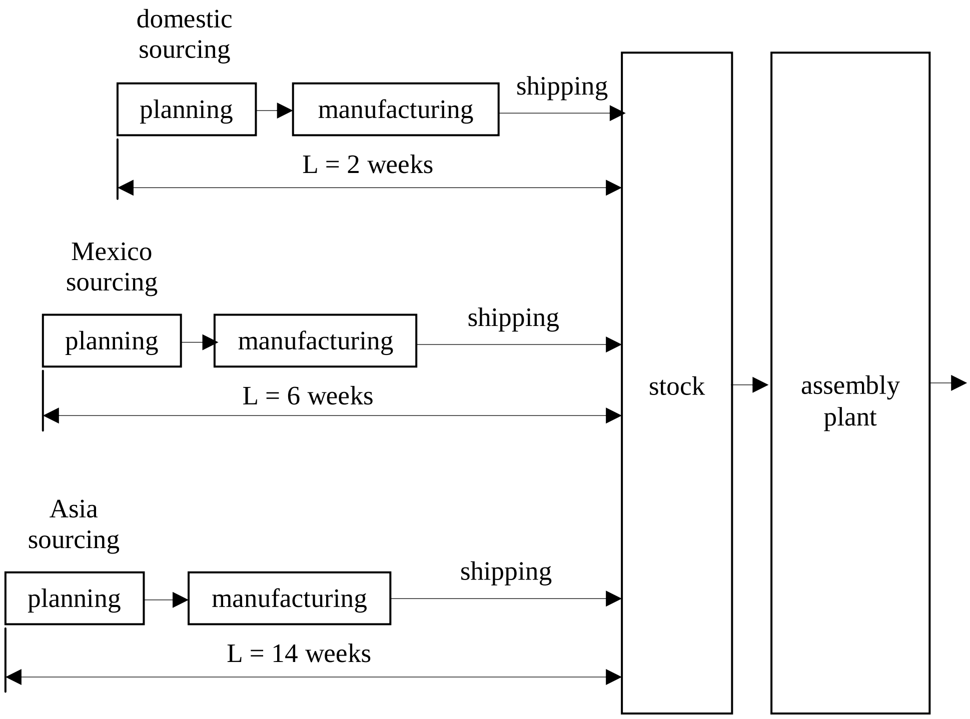

Figure 1: Component sourcing at Volunteer Manufacturing

If time is so valuable, then why are supply chains so long and time consuming? According to Blackburn (2012), the answer lies with the MVT for supply chains—

which is surprisingly low for *functional* products. Consequently, only a small reduction in manufacturing cost offshore is required to offset the increased inventory cost of a longer supply chain. The following example of a sourcing decision for automotive parts illustrates why we have observed such an explosive growth in global supply chains caused by firms moving production offshore in search of lower costs.

Volunteer Manufacturing sources components globally for assembly into automotive products in the US. In its efforts to reduce the total cost of sourcing components, the firm periodically evaluates alternative sourcing strategies. Figure 1 (p. 51) illustrates the available options for a typical component. Domestic sourcing has the shortest replenishment lead time (two weeks), but the highest component manufacturing cost. Mexico, with a slightly longer supply chain (six-week lead time), offers the advantage of lower manufacturing cost, but with higher supply chain inventory cost. Sourcing from Southeast Asia lengthens the supply chain (to about fourteen weeks), slows the response and further increases supply chain inventory cost, but has the lowest acquisition cost of components. For Volunteer Manufacturing, the sourcing decision involves a classic trade-off between manufacturing costs and supply chain costs, and the resolution of the trade-off depends on the MVT for the supply chain.

Automotive assembly at Volunteer Manufacturing is based on a make-to-stock supply chain: a stock of components is held in inventory and components are pulled as needed for assembly. To calculate the total cost of sourcing product from different locations, an analytical model is needed to calculate inventory cost as a function of the length of the supply chain lead time (L). Once the inventory cost has been determined, a total inventory cost model is needed to capture how cost changes with L, and to obtain the MVT for the supply chain.

The details of developing an analytical model to calculate the total inventory cost in a make-to-stock supply chain, such as the one at Volunteer Manufacturing, are omitted here, but can be examined in Blackburn (2012). The model is based on the following assumptions: the upstream producer of a product (or component) ships to a downstream inventory stocking point. The products are *functional* (in Fisher's terminology)—that is, the products have predictable demand and relatively long lifecycles. Demand at the stocking point is variable—and normally distributed. Inventory is managed by a standard order-up-to model with a replenishment lead time (the total time for upstream production and delivery of the product to the stocking point) of L weeks—the lead time may be fixed or variable. The total inventory cost is the sum of ordering costs, carrying costs of inventory (pipeline inventory, cycle stock, safety stock at the stocking location, and the average amount backordered), and penalty costs for shortages or backorders. Most inventories in supply chains are managed suboptimally, in the author's experience. Therefore, for greater generality, this analytical model specifically does *not* assume that the inventory ordering policy is optimal, only that the firm maintains a consistent inventory policy as lead time changes.

Although the analytical formulas for calculating MVT as a function of L are complex (Blackburn 2012), the basic, two-stage evaluative procedure is simple. First, the total annual inventory cost is converted to inventory cost per unit of product. Second, MVT is calculated analytically as the percentage change in unit product inventory cost per unit change in L. For example, if MVT were 0.5 per cent per week, then, implicitly, a one week decrease in lead time reduces the unit inventory cost of the product by 0.5 per cent. The conversion to inventory cost per unit of product is fundamental to this model, because, otherwise, considering MVT's sensitivity to the cost of the product would have to preclude making general statements about the MVT of the supply chain—including, specifically, making MVT comparisons across product categories.

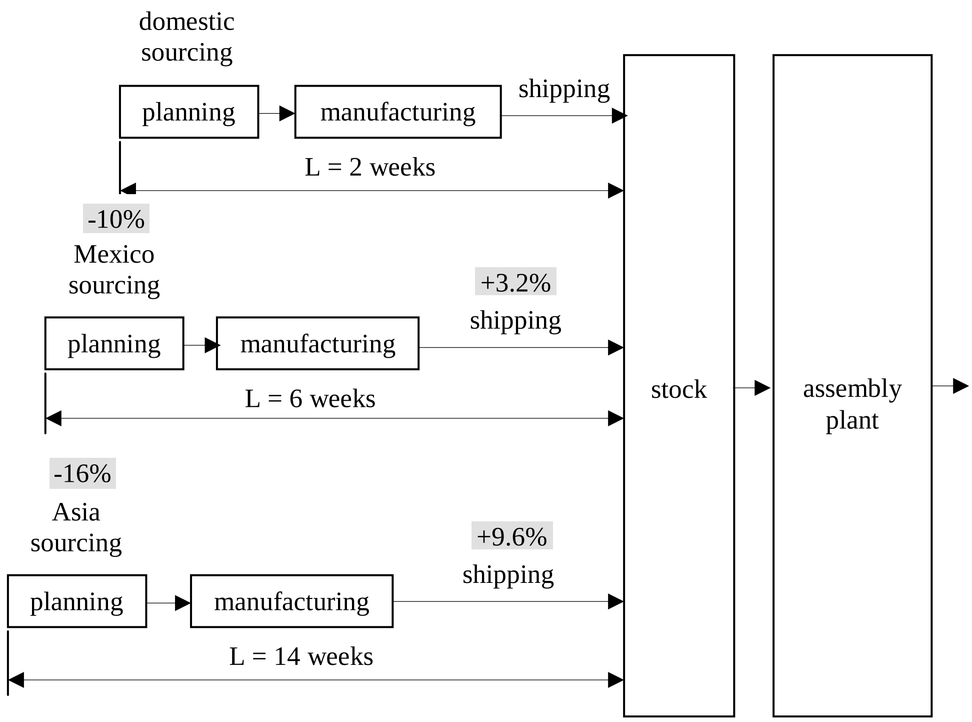

Figure 2: Comparative component sourcing at Volunteer Manufacturing

Once MVT is expressed as a function of L, it becomes possible to quantify how inventory costs change with L—and comparing sourcing alternatives becomes straightforward. Total sourcing cost per unit of product is the sum total of unit manufacturing cost, shipping cost, and inventory cost. To compare domestic sourcing with offshore sourcing, MVT and the difference in lead time (ΔL) are used to compute the increase in inventory costs effected by offshore sourcing (MVTxΔL)—offshore sourcing is only cost effective if the reduction in the sum total of unit manufacturing and shipping cost exceeds MVTxΔL.

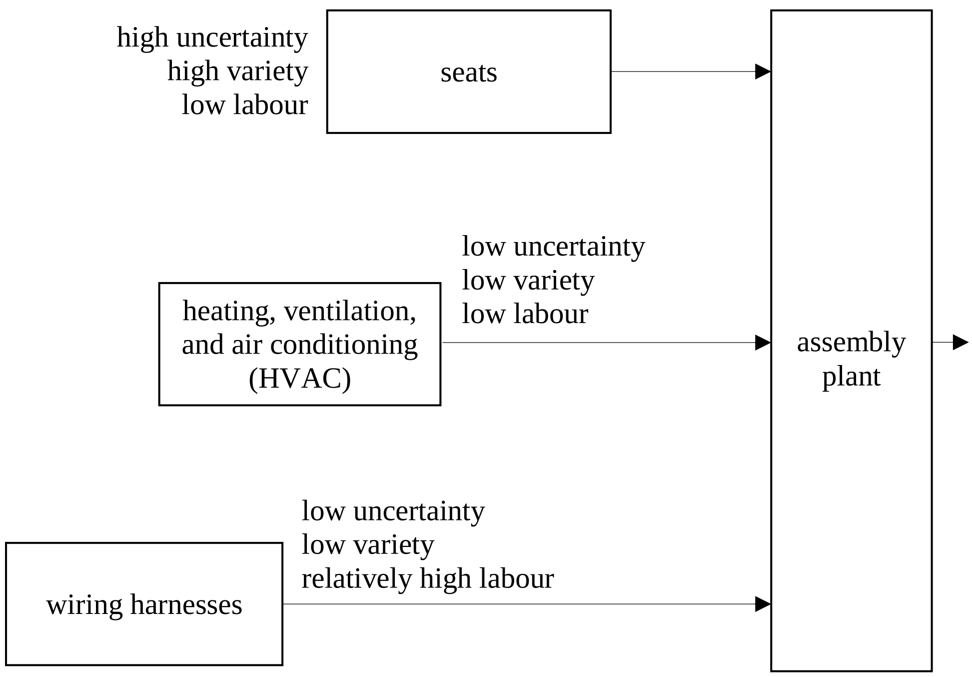

Figure 3: Response time strategies in supply chains in the automotive industry

For Volunteer Manufacturing, MVT is 0.8 per cent per week—that is, a one-week increase in supply chain lead time increases product cost by 0.8 per cent. Therefore, a 5 per cent decrease in sourcing cost (the sum total of manufacturing and shipping cost) would be roughly equivalent to a six-week increase in the length of the supply chain. Figure 2 (p. 53) summarises the results for the alternative sourcing locations—for comparison purposes, all cost differences are relative to domestic sourcing. Due to lower labour and related manufacturing costs, sourcing from Mexico reduces the sourcing cost by 10 per cent per unit. With a ΔL of approximately 4 weeks and an MVT of approximately 0.8 per cent, this longer supply chain increases inventory cost by approximately 3.2 per cent of unit cost. However, the net effect of sourcing from Mexico is a reduction in unit cost of 6.8

per cent. Similarly, sourcing from Asia reduces the sourcing cost by 16 per cent, but the increase in lead time increases the unit cost by approximately 9.6 per cent. All in all, however, the net effect of sourcing from Asia is a reduction in unit cost of 6.4 per cent. Based on the total cost of sourcing, Mexican sourcing is slightly more cost efficient than Asian sourcing—both alternatives offer a considerable advantage over domestic sourcing.

Volunteer Manufacturing is a typical example for supply chains for functional products. Regardless of unit value, MVT is surprisingly low (lower than 1 per cent, in most cases)—a very robust result that also applies to supply chains with variable lead times.

A low MVT implies that extending the supply chain to locations with lower manufacturing costs is not very costly. It is economical for a firm to extend their supply chain by up to five weeks to obtain sourcing cost savings of 5 per cent. A low MVT imposes strict limits on time-based competition, but provides a compelling explanation for the surge in outsourcing and longer supply chains. If time and inventory costs are ineffective barriers for safeguarding domestic manufacturing, the trade-off is easily tipped in favour of offshore manufacturing. It may be economical to go great distances to acquire products at a lower cost.

However, the limits on time-based competition do not exist for innovative, time-sensitive products such as fashion apparel, consumer electronics, and fruits and vegetables—the analytical model summarised here does not consider the cost of product obsolescence, supply chain disruption, and supply chain coordination. When these effects are equated for, MVT becomes significantly higher and favours shorter supply chains and domestic sourcing.

The effect of MVT on supply chains in the automotive industry results in different distances between suppliers and assembler. Figure 3 (p. 54) shows that different response time strategies in supply chains in the automotive industry lead to a mixture of sourcing strategies. For example, seat manufacturers are always located close to the assembly they feed, usually within a half-hour's drive time. Seat production has high variety and high demand variability because the seats are produced in sequence for assembly, and they vary both in colour and style within a given automobile line. Therefore, seats are a time-sensitive rather than functional product, and the supply chain for seats has high MVT. With low labour content and high MVT, the eventual benefits from extending the supply chain would not offset the accompanying increase in inventory cost. On the other hand, wire harnesses require a very high labour content and tend to be functional products. Because their MVT is low, manufacturers can afford to extend the supply chain to take advantage of low labour costs. For these reasons, wire harness manufacturing facilities for domestic US assembly plants have typically been located in Mexico. Manufacturers of heating, ventilation, and air conditioning (HVAC) products face lower levels of variety and demand uncertainty than seat manufacturers—each car model has a very limited set of HVAC options. The MVT for HVAC units lies between that of seats and wire harnesses. With low labour content there is little incentive to move production offshore, but distance to the assembly plant is less critical than with seats. HVAC manufacturers typically locate manufacturing facilities at sites that can conveniently serve more than one auto assembly plant, with supply chains that are longer than for seats. In most supply networks, and largely due to MVT, distance between supplier and assembler increases with increasing labour content and decreases with increasing variety.

For functional products, MVT tends to be constant along the chain. Changes in MVT along the chain require a more flexible design—a situation examined in the next section.

Designing supply chains for perishable products—the case of the fresh fruit and vegetable sector

In most supply chains, MVT is essentially constant along the chain because the value of the product remains stable throughout the supply process. This is not the case for the supply chains for fresh fruits and vegetables and other perishable products. The value of perishable products changes significantly over time, at rates that are highly dependent on temperature and humidity. This means that the MVT for perishable products changes along the supply chain, rendering conventional supply chain design strategies inappropriate.

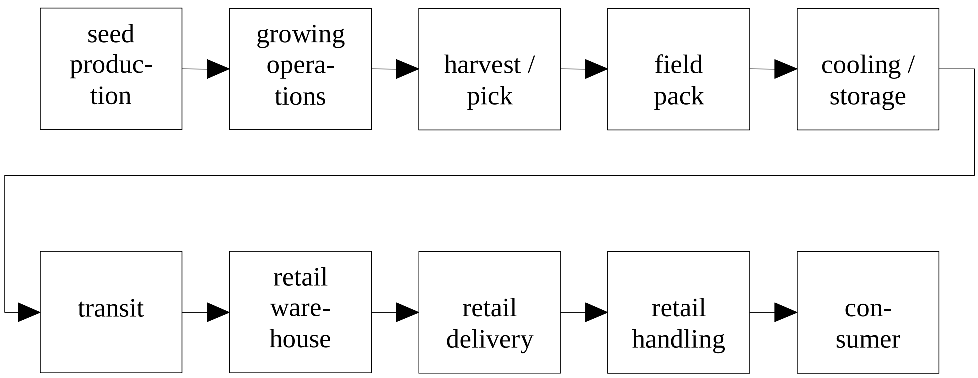

Figure 4: Stages in the melon supply chain

While Blackburn and Scudder (2009) studied perishable product supply chains in detail, this section focuses on the supply chain for melons, which is similar to that of other highly perishable fruits and vegetables. Figure 4 shows the sequence of activities in the supply chain for melons. For perishable products such as this, quality begins to deteriorate immediately upon harvest, and the problem is to choose a supply chain design that limits the loss in product quality in the stages between harvest in the field and customer.

Unlike the MVT for functional products discussed in the previous section, the MVT for perishable products is not constant along the supply chain, and the state of the product does not remain stable. A melon is at its peak value—in terms of freshness, taste, and texture—at the instant it is picked. As an organic product, it begins to degrade in value at an exponential rate highly temperature dependent, after picking, due to the process of respiration—respiration rates increase rapidly with temperature. Contrary to conventional perishable inventory models, the units of product lose value at different rates, depending upon the time and temperature since picking. Moving the harvested product to a nearby cooling facility, where the melon is cooled to a temperature just slightly above freezing, abruptly halts the loss in value due to product deterioration.

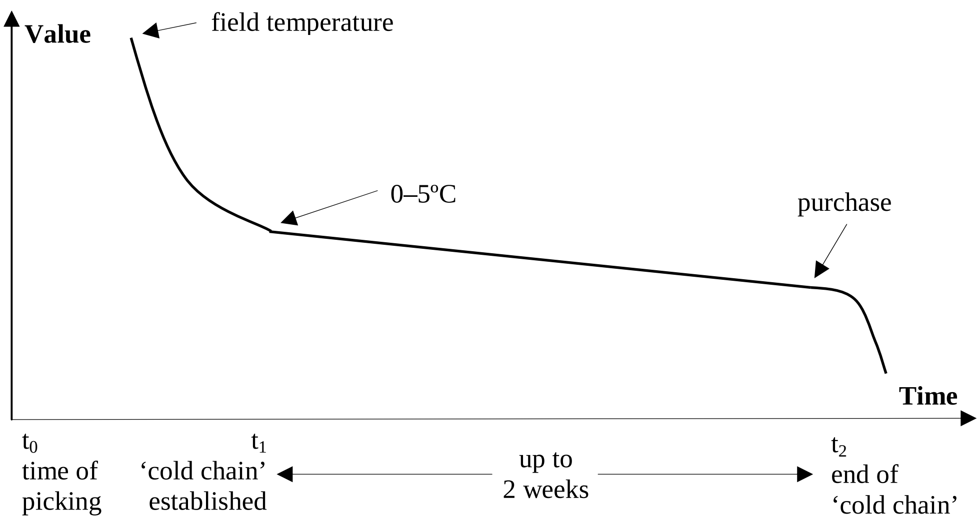

Figure 5: Decline of melon value over time

Figure 5 shows the MVT profile for the fresh melon supply chain. At time of picking (t_0), the product is at field temperature and at its maximum MVT—that is, it is losing quality and value at maximum rate. Value decreases exponentially, as the product waits for transfer to the cooling facility. Once there and cooled to a temperature a few degrees above freezing (time t_1), product deterioration is reduced to a much lower rate. If a 'cold chain' is maintained, the product value can be maintained for several weeks as the melon moves from the field to the market. Maintaining a cold chain is necessary to stabilise product value and quality.

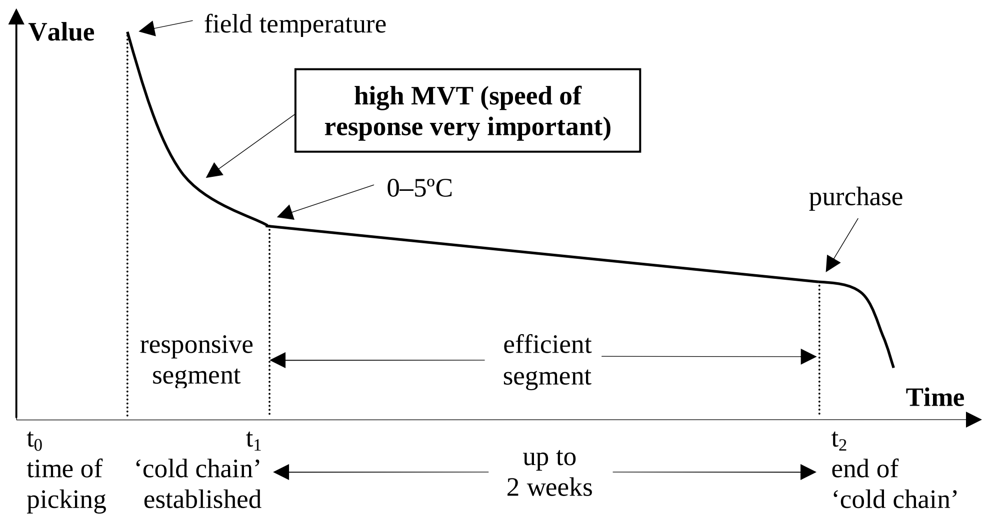

Figure 6: Hybrid supply chain: responsiveness and efficiency

The difference in MVT along the supply chain for melons dictates a hybrid design strategy (see Figure 6). Between t_0 and t_1, the supply chain must be designed for *speed*, because the product is losing value at a rapid rate and MVT is very high. From t_1 to t_2, the supply chain can be designed for cost efficiency—by maintaining the product at a cold temperature, MVT is kept low, so cost efficiency becomes more important than speed. In this part of the supply chain, costs can be kept low by choosing a minimum cost transportation strategy, while insuring that the cold chain is maintained throughout the transportation process.

Management of the first stage of the supply chain (between t_0 and t_1) is most critical because value lost in the product due to delays at this stage can never be recovered. In the field, product is typically transferred in batches to the cooling facility, and the time from harvest to cooling depends on the batch size. If the batch size is too large, the product sits too long at high temperatures, and significant value is lost. On the other hand, very small batch sizes incur higher transportation costs than large batch sizes. Blackburn and Scudder (2009) showed that the optimal batch size can be calculated using a model that is analogous to the economic order quantity (EOQ) model in inventory theory, altered to reflect an exponential loss in value of the product.

This example illustrates that there are significant differences between the supply chains for perishable products and conventional supply chains. Changes in MVT along the supply chain dictate a hybrid mix of responsiveness, in one segment, and cost efficiency, in the other. Fortunately, both segments can be managed separately, because little coordination is required between the two.

Designing reverse supply chains for product recovery—the case of time-sensitive technology products

The flow of product returns from customers has become a major issue for retailers and manufacturers. In the US alone, the total value of products returned in this *reverse supply chain* exceeds USD 100 billion annually. This value is growing rapidly with increasing on-line sales, which tend to have higher return rates than in-store sales. For products returned within 90 days of sale, the burden falls on manufacturers, who must credit the retailer for product value and dispose of the product through reuse, recycle, refurbish, or salvage.

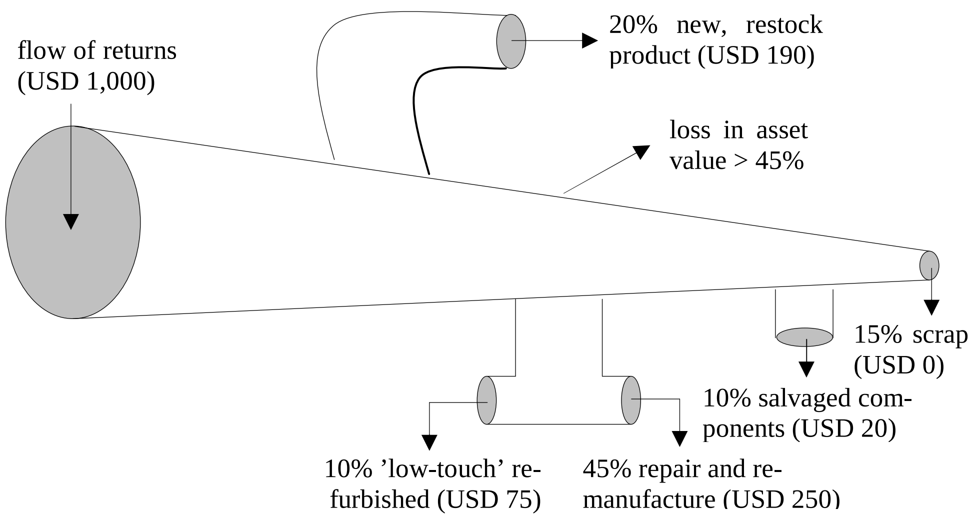

Figure 7: Shrinking pipeline for product returns

While focused on forward supply chains, many firms ignore their reverse supply chains and seek to handle returns at minimum cost. Cost-efficient supply chains tend to be slower than responsive ones, and delays in dealing with returned products can lead to increased losses in asset value and a diminished set of reuse options. This approach may reduce handling costs, but may result in sending a product to landfill, losses for the firm, and long-term damage to the environment.

Blackburn et al. (2004) illustrated how assets can be lost through mismanagement or neglect of the return stream for products whose value diminishes rapidly over time. Figure 7 (p. 59) represents the returns process as a shrinking, leaking pipeline. The percentage losses shown are averages representative of firms with time-sensitive technology products in the authors' own research database. For every USD 1,000 of product returns, approximately half the asset value (more than 45 per cent) is lost in the return stream. The explanation is twofold. First, the value of the product decreases with time (at rates as high as 1–2 per cent per week), as the product moves through the pipeline to its ultimate disposition. Second, the returned product, once new, must be downgraded to a lower-valued product—remanufactured, salvaged for parts, or simply scrapped as not reparable or obsolete. Much of the second type of loss is unavoidable, because only a fraction of returns can be restocked as new items (20 per cent, in the case of time-sensitive technology products). However, losses due to time delays represent a significant opportunity for asset recovery.

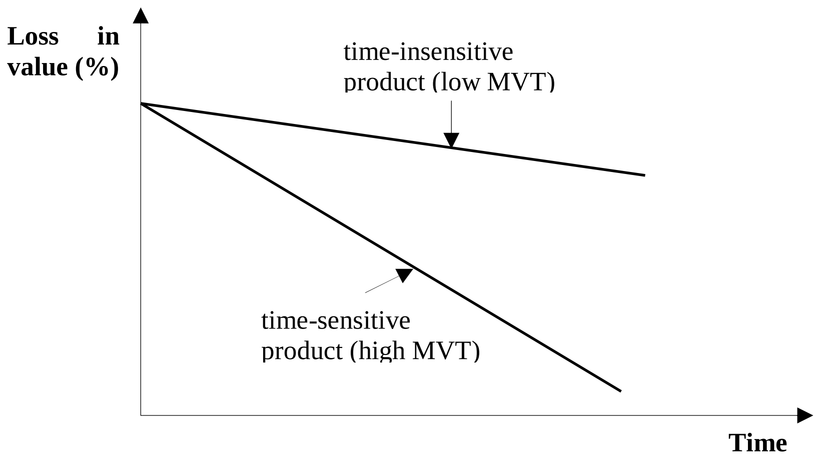

Figure 8: Differences in MVT for product returns

To control asset losses due to time delays in the reverse supply chain, managers must be sensitive to the value of time for product returns and use it to (re)design the reverse supply chain for asset recovery. The product's MVT is a simple, effective metric for measuring the cost of delay—for returned products, MVT varies widely across industries and product categories. Figure 8 shows MVT in percentage terms, to facilitate comparisons across product categories. Time-sensitive, high MVT products such as PCs and other consumer electronics may lose value at rates in excess of 1 per cent per week—at such rates, returned products may lose up to 10–20 per cent of their value simply due to time delays in the evaluation and disposition process. On the other hand, for low MVT products such as disposable cameras and power tools, the cost of delay is usually closer to 1 per cent per month.

Design choices

Because asset recovery depends strongly on reducing time delays, MVT is a convenient parameter to drive design decisions for the reverse supply chain. As in previous sections, it is useful to recast Fisher's (1997) taxonomy of strategic design choices for supply chains in time-based terms. Using Fisher's product classification, *innovative*, short lifecycle products (such as laptop computers) have high MVT, whereas products such as power tools and disposable cameras are more *functional*, less time-sensitive, and have low MVT. The resulting design dichotomy is expressed as follows: products with high MVT require time-based reverse supply chains designed for *responsiveness* and products with low MVT require reverse supply chains designed for *cost efficiency*.

The major structural difference between cost-efficient and responsive reverse supply chains lies in the supply chain positioning of returned product testing and evaluation, where product condition is to be determined. For low MVT products, cost efficiency is the objective, and the returns supply chain should be designed to *centralise* the evaluation activity. With high MVT products, responsiveness is the goal, and the evaluation activity should be *decentralised*, to minimise time delays in processing returns.

The centralised model for cost efficiency

Figure 9 (p. 62) shows a generic model of a reverse supply chain built around centralised testing and evaluation. To achieve cost efficiency, the returns supply chain is designed for economies of scale. Every returned product is sent to a central location for testing and evaluation, to determine condition and issue credit. Product returns are shipped in bulk, usually, to minimise shipping costs. Once evaluated for condition, the product is distributed to the appropriate facility for disposition: restocking, refurbishment or repair, parts salvaging, or scrap recycling. The centralised reverse supply chain is designed to minimise processing costs, often at the expense of long delays, so it should only be used for functional, low-MVT products.

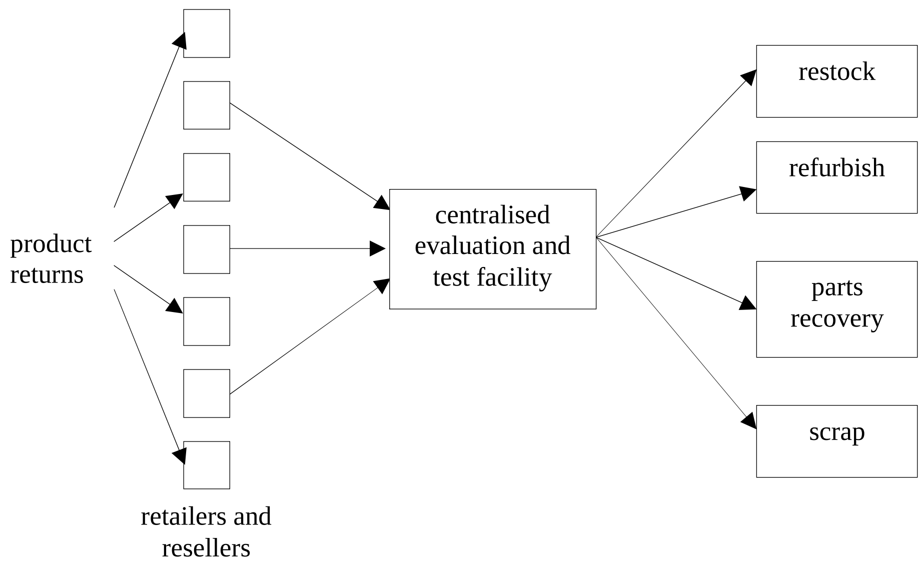

Figure 9: Centralised, cost-efficient reverse supply chain

The centralised model adopts a fundamental principle of forward supply chain design strategy, *postponement*. Postponement—or delayed product differentiation—is an effective strategy for dealing with the cost of variety, reducing inventory in forward supply chains (Feitzinger and Lee 1997). However, postponement is less effective as a strategy for reverse supply chains. The key return decision is based on an initial evaluation of product condition in order to make a disposition choice. With returns, little is gained from postponing product differentiation, because product variety and condition are predetermined at the time of receipt.

The decentralised model for responsiveness

In the reverse supply chain, there are significant time advantages to early, rather than late, product differentiation—a design principle called *preponement*. Early diagnosis of product condition maximises asset recovery by fast-tracking returns to their ultimate disposition and minimising delay cost. Upon return of consumer electronics products such as PCs, printers, and phones, a field test can be conducted to categorise products as new, refurbishable, salvageable for components, or scrap. Then, new, unused products can be restocked without any time delay (and asset value loss), scrapped products can be recycled, and the remaining products can be sent on for further evaluation. As Figure 10 illustrates, to achieve preponement and reverse supply chain responsiveness, product testing and evaluation must be *decentralised*.

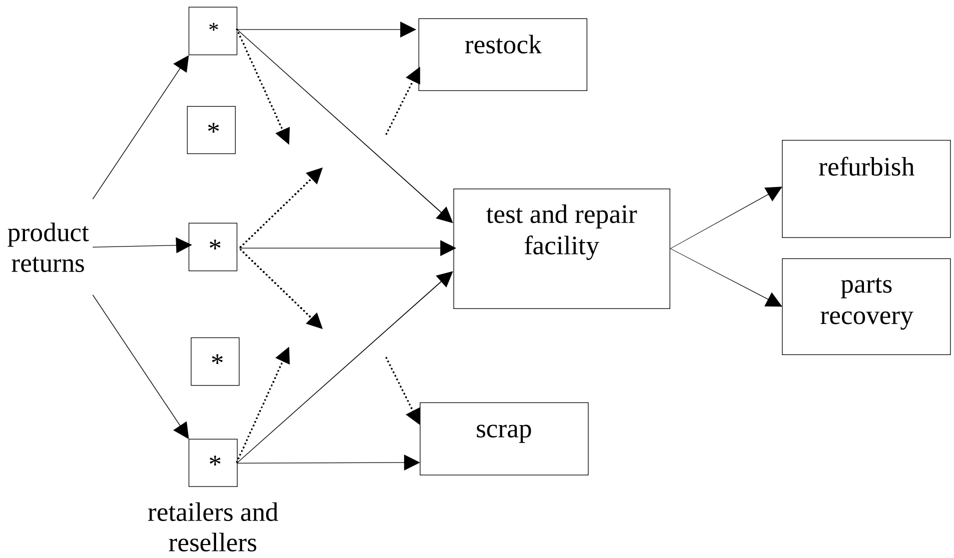

* = evaluation of product at retailer or reseller

Figure 10: Decentralised, responsive reverse supply chains

Decentralised preponement improves asset recovery by reducing time delays in two ways. First, it reduces the time delays for disposition of new and scrap products—new products tend to have the highest MVT and stand to lose the most from processing delays. Second, it speeds the processing of products that need further testing and repair. By diverting new and scrap products from the main returns flow, overall congestion is reduced and the remaining products flow faster, thereby reducing asset loss further. For products with high MVT, preponement may increase asset recovery speed dramatically.

To summarise, product returns represent a value stream, not just a waste stream, and the reverse supply chain should be managed to recapture that value. Like that of forward supply chains, the design of reverse supply chains involves a trade-off between efficiency and response. To improve asset recovery in reverse supply chains, firms must consider time (as measured by MVT) as a key performance and design metric. If the product has a high MVT, then the reverse supply chain should be decentralised, to achieve preponement and speed. If the product has a low

MVT, then the reverse supply chain should be centralised, to achieve cost efficiency.

Conclusion

This article has important implications for managers seeking the right supply chain strategy for their product. The three cases discussed here are typical of most supply chains, in that selection of the appropriate design involves a trade-off between speed and cost efficiency. MVT is an ideal tool for resolving this trade-off, because it succinctly captures the relationship between cost and time in a supply chain. In supply chains for functional products, with stable demand and relatively long lifecycles, the value of time is much lower than intuition would suggest, and firms can extend their supply chains long distance effectively, to obtain lower sourcing costs. When the value of time varies along the supply chain, as it does with perishable products, managers should consider flexible designs, to combine speed and cost efficiency. By ignoring their reverse supply chain, many managers miss an opportunity to capture asset value. As is the case with forward supply chains, the design of the reverse supply chain pivots on MVT. With time-sensitive products, early identification of product condition (preponement) is necessary to maximise asset recovery.

References

Blackburn, J. D. (ed.) (1991). *Time-Based Competition: The Next Battleground in American Manufacturing*. Homewood, IL: Business One Irwin.

Blackburn, J. D. (2012). 'Valuing Time in Supply Chains: Establishing Limits of Time-based Competition', in *Journal of Operations Management*, 30/5: 396–405.

Blackburn, J. D. and Scudder, G. (2009). 'Supply Chain Strategies for Perishable Products: The Case of Fresh Produce', in *Production and Operations Management*, 18/2: 129–37.

Blackburn, J. D., Guide, V. D., Souza, G. C., and Van Wassenhove, L. N. (2004). 'Reverse Supply Chains for Commercial Returns', in *California Management Review*, 46/2: 1–17.

Feitzinger, E. and Lee, H. (1997). 'Mass Customization at HP: The Power of Postponement', in *Harvard Business Review*, 75/1, 116–21.

Fine, C. H. (1998). *Clockspeed: Winning Industry Control in the Age of Temporary Advantage*. Reading, MA: Perseus Books.

Fisher, M. L. (1997). 'What is the Right Supply Chain for Your Product?', in *Harvard Business Review*, 75/2: 105–16.

Schmenner, R. W. (1988). 'The Merit of Making Things Fast', in *Sloan Management Review*, Fall 1998: 11–17.

Stalk, G., Jr. (1988). 'Time—The Next Source of Competitive Advantage', in *Harvard Business Review*, 66/4: 41–51.

Stalk, G., Jr. and Hout, T. (1990). *Competing Against Time*. New York: The Free Press.

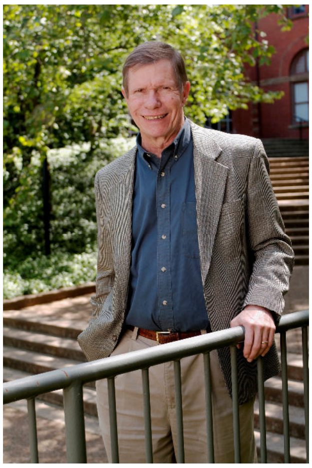

Joseph D. Blackburn is James A. Speyer Professor of Production Management Emeritus at the Owen Graduate School of Management, Vanderbilt University.

Joe received his PhD in operations research from Stanford in 1971. His research and teaching in operations management focuses on time-based competition—how organisations can use faster response to customers for competitive advantage. In 1995, Joe was elected Outstanding Professor in the Executive MBA Programme, where he continued to teach until recently. He is the editor of *Time-Based Competition: The Next Battleground in American Manufacturing*, and is an authority on accelerating new-product development, streamlining supply chains, and reverse supply chain strategy.

Joe served as Associate Dean (1988–1993), Acting Dean (1999–2000), and Senior Associate Dean (2005–2007) at the Owen Graduate School of Management, Vanderbilt University. Prior to coming to Owen, he was Associate Professor of Operations Management at Boston University's School of Management, Assistant Professor of Management Science at the University of Chicago's Graduate School of Business, and Visiting Assistant Professor of Decision Sciences at Stanford University's Graduate School of Business.

Before moving to academia, Joe was a chemical engineer for Eastman Kodak. He is a Fellow of the Production and Operations Management Society and serves as Senior Editor for *Sustainable Operations, Production and Operations Management Journal*.

Joe can be contacted at joe.blackburn@owen.vanderbilt.edu, at the Owen Graduate School of Management, Vanderbilt University.