1 Old Habits Die Hard

If talking about build-to-order could make it happen, many companies would already be building truckloads of custom cars. Indeed, many manufacturers have affirmed strategies to reduce costly finished inventory by shortening order-to-delivery time and building vehicles to customer order.[1] Something has obviously happened to dampen their enthusiasm. The variety of attempts and the scattered strategies reflect a rather half-hearted attitude, and the results often are more accidental than planned. Even at those manufacturers that are the furthest along, build-to-order strategies are not consistent across brands and models, and even vary within model production cycles. Companies tend to reserve build-to-order for luxury models and to use forecast-based production for most other products. Furthermore, build-to-order percentages are much higher at the start of a model's cycle than at the end.

Why haven't companies made the transition? Part of the problem is that no one seems clearly motivated. In the larger scheme of sales sourcing—determining how the customer's order is actually fulfilled—build-to-order is only one possibility, and owing to current long waits it is not a popular one. As table 1.1 shows, true build-to-order requests—the ones the company receives *before* building the car (new orders from a customer or orders already in the system amended to customer specifications)—do not typically amount to a significant percentage of sales sources. The high percentages for Germany and for Toyota in Japan are outliers rather than the norm.

Both sales from central stock and sales from dealer stock are build-to-forecast sales—the customer selects a vehicle from existing stock. The vehicle can come either directly from the stock at the dealer the customer visits, from another dealer, or from stock held at distribution centers awaiting dealer orders. To identify a suitable vehicle from the stock at a distribution center or at another dealership, dealers access the vehicle

Table 1.1
Sales sourcing in major volume markets, 1999–2000.*

	Europe	UK	Germany	US	Japan (Toyota)
Cars built to customer order	48%	32%	62%	6%	60%[a]
Sales from central stock (distribution centers) or transfer between dealers	14%	51%	8%	5%	6%
Sales from dealer stock	38%	17%	30%	89%	34%

*Sources: G. Williams, Progress towards Customer Pull Distribution, research paper 4/2000, International Car Distribution Programme, Solihull; H. Shioji, "The order entry system in Japan," International Symposium on Logistics, Morioka, Japan, July 2000; Office for the Study of Automotive Transportation, University of Michigan Transportation Research Institute, Delphi Study X, 2001.
a. Recent interviews at Toyota suggest that the current build-to-order level for Toyota in Japan is more like 50% across all sales channels.

manufacturer's stock-locator system. Once a dealer finds a vehicle, the distribution center transfers the product to the dealership. If the vehicle is at another dealership, the transport incurs an extra fee, since it is usually a single car on the truck. Ultimately this inefficient delivery mode translates into higher logistics cost, which all customers absorb. In the United Kingdom, this transfer cost is an estimated $180 for each delivery. The transfer from distribution centers is generally cheaper because the center can schedule more cars per truck.

Table 1.1 shows how radically three high-volume regions—Europe, Japan (Toyota), and the United States—differed in 1999 and 2000. We isolated Germany and the United Kingdom to show how even within a region, markets can differ significantly. We used Toyota to represent Japan. Even though it is considered the benchmark in the auto sector, it achieves only 60 percent build-to-order. This relatively high ratio is specific to the Japanese domestic market: all Japanese firms, including Toyota, are largely building to forecast in Europe and the United States.[2]

Much as in Europe, there is great diversity in the progress and effort Japanese manufacturers are making towards build-to-order in Japan. On one end of the spectrum are Honda and Subaru, which historically have made virtually all their cars to forecast. Honda bases its production entirely on forecast. Plants in Japan handle high levels of variety, but production is batched. Under *kanban* (a pull system with suppliers), all material is delivered by suppliers in the forecast batches, and this makes creating the pull required in a build-to-order system extremely difficult to attain. Toyota and Nissan, in contrast, have openly declared their intention to reduce order-to-delivery lead times and implement build-to-order. Toyota, known for *kaizen* (continuous improvement), has applied this strategy to its build-to-order efforts. Toyota initiated "daily ordering" in 1999, the objective being to reach an order-to-delivery target of 14 days. Currently, Toyota achieves an average of 23 days across all sales channels.[3] Nissan's build-to-order efforts in Japan are already well advanced. In 1991 Nissan implemented a new order entry system, called ANSWER,[4] that allowed dealers to schedule daily orders into the production schedule 6 days before production. (This was called "D-6 scheduling.") The number of days was reduced to 4 in 2001 with the launch of ANSWER II. Nissan's underlying philosophy is *douki-seisan*, which can be translated as "'the origin of everything is the customer." After trials in Japan, *douki-seisan* was announced on a global basis in 1997. With a 30-day order-to-delivery lead time (the target is 15 days), Nissan still lags behind Toyota. However, it has made tremendous progress in using build-to-order to cut inventory, and it holds only 20 days of inventory in Japan—half Toyota's average stock.[5]

The US market is a clear outlier, with most cars sold from dealer stock. In the United States, instant gratification rules. Customers have been trained to go to a dealer, find a car that roughly matches their needs and wishes, haggle for a good price, and drive away with a car the same day. Car buying is akin to grocery shopping: go in, take what is on the shelves even if it's not exactly what you want (after all, something is better than nothing), and drive home. Indeed, the average supermarket shopper hardly ever finds all the goods that are on his shopping list and commonly buys several items that aren't on it.[6] Apparently, car shopping is conducted in a similar spirit. At the other end of the gratification spectrum is Germany. The German market follows a tradition of long waits for custom cars,[7] so it is no surprise that it has the highest build-to-order content in Europe.

Toyota's build-to-order percentage is high in part because all its suppliers, retailers, and customers are conveniently co-located. That, coupled with overseas exports to provide a stable production base, makes it easier to implement build-to-order. This geographic concentration of value-chain stakeholders is common in Japan. The experience of other Japanese manufacturers shifting to build-to-order is similar. Nissan, for example, has been able to attain a 50 percent build-to-order rate on fairly short notice in the Japanese market.

In the United Kingdom, 19 percent of customers did not receive the exact specification they asked for, which relates to the high percentage of vehicles bought from existing stock. In Europe, on average, 20 percent of customers compromised on their vehicle's specifications, according to dealer perceptions. Despite Germany's high use of build-to-order, 24 percent of German customers felt the same way. These numbers strongly imply that no single market leads in proactively altering its processes to increase building to order. The high percentage in Germany is more of a cultural anomaly, as is the low percentage in the United States. Toyota's high percentage is due to its geographic concentration of supporting elements and exports.

A closer look at the built-to-forecast percentages also shows how aspects of the forecast distribution model affect the implementation of build-to-order. The trend in the United Kingdom is particularly revealing. When the technology for stock-locator systems first became available in the late 1980s, Ford and some other companies pioneered them in the United Kingdom, initially with great success. As figure 1.1 shows, in 1992, 45 percent of new vehicles were sales from dealer transfers. Soon, however, manufacturers realized that each sale from a stock-locator system came with a $180 transportation-to-dealer cost. As logistics costs skyrocketed and eroded profit margins, vehicle manufacturers backed away from stock locators and turned to distribution centers. By 1999, the 45 percent had dwindled to 15 percent.

But although distribution-center sales increased, build-to-order failed to take off. Part of the reason was that reliance on distribution centers introduced an additional objective: remove stock from the dealer and fulfill orders with less overall stock. Consequently, by the end of the 1990s sales from dealer stock had also decreased dramatically. Despite shifts in where product originates, the fundamental problem remains: all but one-third of the product is built to forecast demand rather than actual demand, and is shuttled around post factory to ultimately match up with a customer.

Figure 1.2 shows how the forecast-driven strategies set up two vicious cycles. Stock levels in the marketplace disconnect the entire value chain from the customer, which in turn frees factories to address a stable, long-term forecast without worrying about individual customer requirements.

No one will debate that the push strategy is tempting, but it is also a trap. An industry that supplies customized high-volume products ends up pushing finished goods into the market to get the revenue to offset bulk production and design costs.

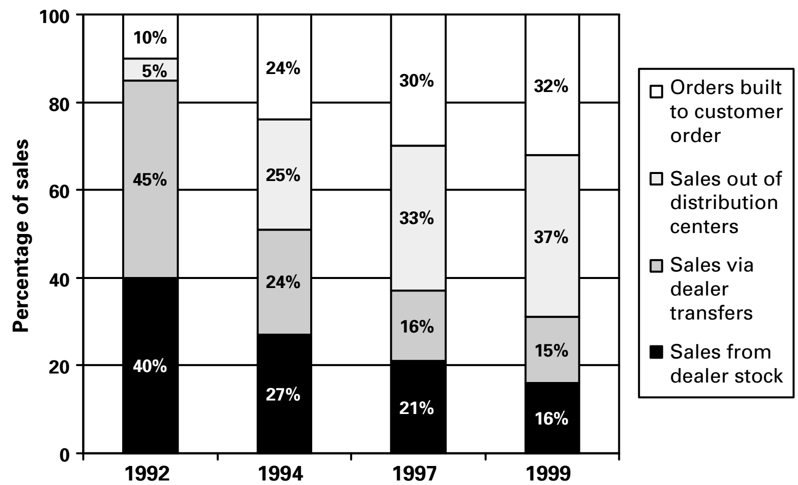

Figure 1.1
Sales sourcing in UK volume car market, 1992–1999. Sources: G. Williams, European New Car Supply and Stocking Systems, ICDP, 1998; Williams, Progress towards Customer Pull Distribution, ICDP, 2000.

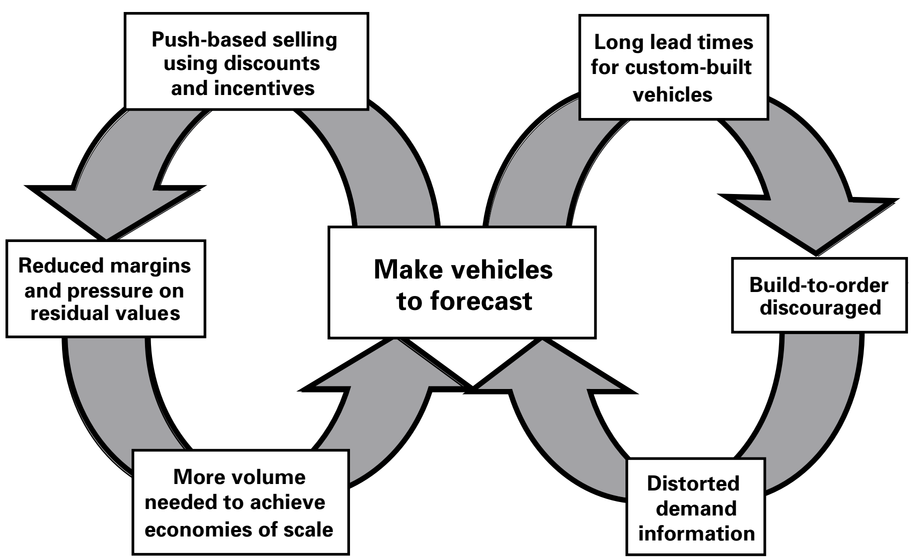

Figure 1.2
The vicious cycles of make-to-forecast. Source: M. Holweg and F. Pil, "Successful build-to-order strategies start with the customer," *Sloan Management Review* 43 (2001), no. 1: 74–83.

The focus turns to efficiency on the shop floor to get the unit cost down. The motivation to optimize the whole system from the customer to suppliers is lost in a frenzy of build, sell, and build more. Vehicles that must be built to a specific order are so unusual that the order-fulfillment lead times increase. The longer the wait, the less willing customers become to go through the gauntlet of getting exactly what they want.

In a push-driven world, customers have three basic choices: wait for a custom-built vehicle, settle for a vehicle in stock that isn't exactly what they wanted, or go to another brand with a similar configuration and better availability. The go-to-another-brand option translates to a lost sales ratio—total customers divided by the percentage of customers who bought elsewhere because they didn't get what they wanted in time. In 1999, this ratio for the United States was an estimated 36 percent.[8] In the United Kingdom, lost sales dropped from 10 percent in 1994 to 6 percent in 1999, and in the same period Germany's lost sales went from 11 percent to 3 percent. In France, the ratio stayed constant at 5 percent.[9] These numbers clearly imply that companies are altering their distribution systems to retain the customer.

Any lost sale represents lost revenue. Yet there is a larger issue. The longer it takes to provide customized vehicles, the more customers are actually discouraged to request a car built to their specifications. The manufacturer gets farther and farther from the customers, relying on sales forecasts that are based on artificial measures. For example, suppose a forecast contains X number of red cars. Now customers actually like blue cars, but dealers induce them to buy red cars with 0 percent financing, better trade-in offers, free upgrades, discounts, and other sales tactics. Eventually customers reluctantly give in to the pressure. Those evaluating the sales and generating the manufacturer's forecast do not capture customers' original desires; they see only that the red cars sold, so they again put red cars, perhaps even more this time, into the production program. Since the production program is generally based on historic sales information, the cycle is self-fulfilling. It breaks only when no amount of incentives will persuade a customer to buy a neon-pink car.

Incentives also end up increasing the cost of the sale because the manufacturer is essentially paying a customer to reduce its inventory. The more discounts are needed, the lower the per-vehicle profit becomes. The lower the profit, the less cost the manufacturer can recover on each sale, so it increases volume to recover cost through economies of scale. Higher volume then forces the manufacturer to push even more vehicles into the market, using still more incentives, and the vicious cycle gets more vicious.

A less well known failing of the push strategy is that it depresses the market for used cars. As manufacturers boost vehicle sales through fantastic discounts and incentives, the prices for used and nearly new cars plummet. This downward spiral not only lowers the residual values for specific vehicle lines; it also erodes the entire value perception of a brand. Such loss in perceived value greatly affects the purchase decision. Volkswagen, Mercedes, and Porsche capitalize on this relationship by touting the competition's loss of brand value in their marketing campaigns.

Plummeting residual values also have serious ramifications for financing companies. Leasing firms, for example, generally predetermine the price at which the vehicle will be bought back years later. In 2000, when market demand was less than expected in both the United Kingdom and the United States, leasing companies suddenly found drastic discrepancies between their book value and the actual prices they could achieve for their formerly leased vehicles. For example, a loaded 1997 Ford Expedition XLT with a sticker price of $36,580 was expected to be worth $25,606 after a 3-year lease. Instead, it fetched only $16,500 at auctions in 2000, and the bank or auto maker had to cover the $9,106 difference.[10] Two years later the situation was no different. The average auction price of an off-lease Taurus, which in October 2000 was $10,750, fell to $8,650 in March 2002. Other models fared no better—during the same period, the average auction price of a used Toyota Camry went from $11,475 to $10,250.[11]

In short, the forecast-driven model has dramatic costs, some of which are visible at the point of sale and some of which do not wreak their damage until many years later. The vicious cycle of build-to-forecast is hard to break out of. We will examine the order-to-delivery model associated with the vicious cycle of build-to-forecast, and from there will turn to the changes that are needed to introduce a build-to-order system in which cars are built as if customers mattered.

The patterns outlined thus far show that a transition to build-to-order is far more challenging than many companies had envisioned, perhaps in part because they lacked a deep understanding of how each stage of the order-to-delivery process affects build-to-order. In fact, it is worth pausing a moment here to clarify exactly what we mean by "order-to-delivery." Despite what some people think, order-to-delivery is *not* build-to-order. "Order-to-delivery" has two distinct meanings: (1) the process from the time the customer or dealer places the order until the vehicle is delivered to the customer or dealer and (2) the time this process takes. Thus, order-to-delivery and build-to-order stand in simple relation to one another: a manufacturer's order-to-delivery (process) time must be short enough to inspire customers to wait for their vehicles to be built to order. If the potential order-to-delivery time falls outside the customer's waiting tolerance, the customer can move to another manufacturer with better availability or shorter order-to-delivery time. When manufacturers talk about "reducing order-to-delivery," they generally mean shortening the process time so that they can build more vehicles to order and not lose customers to other brands with faster availability.

Another misunderstood term is "value chain." In fact, from this point on, we will begin using a term that is more representative of where value originates: "value grid." "Chain" implies that relationships and units within organizations contribute to profitability in a linear fashion. Nothing could be further from reality—despite what finance and strategy folks would have us believe. Relationships have complex interdependencies and overlaps that affect profits in far-from-straightforward ways. Equipment suppliers, for example, may have customers other than auto manufacturers. Even within a manufacturer, some activities (e.g., the development of fuel cells) can occur in parallel with other value chains. This is true of other industries too. In the wireless telecommunications industry, for example, Nokia is collaborating with other handset producers to maintain a sufficiently large user base of its software. That way, Nokia hopes to fend off Microsoft, which threatens to capture the value in the handset business much as it has done in the personal computer business. In contrast, Microsoft has penetrated the handset industry laterally, and views the handset operating software as a way of enhancing and reinforcing its core business in operating software for computers.

A prerequisite to understanding the order-to-delivery process and build-to-order capability is to have some measure of what is realistic. To satisfy this requirement, we evaluated the generic systems and benchmarked the order-to-delivery processes for six vehicle manufacturers—two European-owned and two US-owned vehicle manufacturers operating in Europe and two European divisions of Japanese car producers. Most of the research done in the past 20 years has focused on assembly operations, on the integration of component suppliers, and on distribution logistics. Relatively little is known on how the entire order-to-delivery process actually works, yet without viewing the entire process no company can identify and remove the obstacles to a build-to-order transition.

Benchmarking order-to-delivery capability is tricky. We mapped the full circle of information flow, from order entry at the dealers through national sales organizations, production planning and scheduling functions at the vehicle manufacturer, and the material or physical flow from the suppliers through the assembly plants and back through the distribution channels to the vehicle's delivery to the dealer.[12] Our focus was on system capability[13]—the minimal system-related throughput time for a custom-built vehicle—rather than on average lead times for specific orders.[14] In using system capability, we assume that each subsystem provides the fastest currently feasible throughput, with no rework or other delays. Averages, on the other hand, typically include delays of various kinds—e.g., vehicles held in queues if the supply of an option many customers are requesting is limited, or if quality problems occur in a single order. By focusing on the critical path (the best possible time), we are accurately defining the best response of the current systems to a vehicle order. Most actual orders will take longer than this theoretically achievable time, but system capability portrays the manufacturer's and the supplier's basic ability to support build-to-order. Approaches based on actual average performance cannot provide this, since they would be distorted severely by the demand-supply scenario for a particular model.

Finally, in our benchmarking, we defined "delivery" as delivery to the dealer, excluding dealer preparation. Our research shows that customers rarely go immediately to the dealership to pick up their vehicles, but prefer to wait until the weekend, so it would be pointless to add the time the dealer takes to hand over the vehicle to the customer.

Figure 1.3 is a greatly simplified diagram of a typical order-to-delivery process. Of course the order-to-delivery process is more complex than any one diagram can convey, comprising dealers, national sales companies, head offices, assembly plants, tiers of component suppliers, and logistics companies in an arrangement that differs across manufacturers. Our goal in drawing the process map was to capture the underlying logic across order-to-delivery processes, which is surprisingly similar across manufacturers. From this map, we identified four basic stages in order-to-delivery, which we use as a basis for deriving order-to-delivery process times. These steps are order entry, production scheduling and sequencing, vehicle production, and vehicle distribution. We also

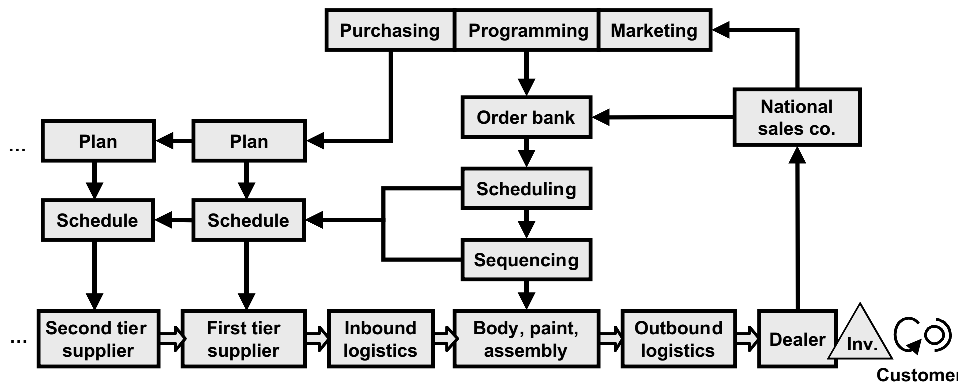

Figure 1.3
Map of typical order-to-delivery process.

identified four linked processes that are critical to understanding order-to-delivery: sales forecasting, production programming, supplier scheduling, and inbound logistics.[15] Thus, the order-to-delivery process includes the following stages and processes:

• *Sales forecasting* aggregates all dealers', national sales companies', and importers' forecasts and uses them as a basis for regional sales forecasts, which in turn become input for production programming.

• *Production programming* maps a consolidation of forecast market demand onto available production capacity to yield a framework that defines how many vehicles will be built in each factory, the vehicles' specifications, and the markets to which they will be delivered.

• *Order entry* is the stage in which orders are checked and entered into an order bank to await production scheduling.

• *Production scheduling and sequencing* fits orders from the order banks into production schedules, according to the production program. Once scheduled for a particular plant, these orders are used to develop the sequence of cars to be built on the scheduled date.

• *Supplier scheduling* is the process by which suppliers receive forecasts at various times, actual schedules, and daily call-offs—all of which can vary dramatically.

• *Inbound logistics* is the process by which logistics service providers (often third parties) collect parts from suppliers, consolidate them, and deliver them to the assembly plants at the required time.

- *Vehicle production* is the process of physically welding and painting the body and assembling the vehicle, including all testing and rework.
- *Vehicle distribution* is the stage at which the plant ships the finished vehicle to the dealer (or in rare cases, directly to the customer).

These eight processes and stages represent the building blocks most manufacturers rely on to manage the order-to-delivery process. The activities within each stage provide useful insights into the benefits and limitations of the current information and material flows in the supply chain.