Inertia in the supply chain.

Synopsys

The notion that supply chains have 'inertia' is perhaps half understood. That is to say, it's well accepted that when sales are high it takes time for the supply chain to catch up. However, there's another side to this coin, that when sales slow down the supply chain 'keeps coming'. In the short term this overfills store shelves, robbing the store of space it could perhaps use to sell something else. In the medium term it leads to warehouses with unwanted product yet out-of-stock of crucial (other) products. In the long term, say a season, it leads to stock markdowns and write-offs. Using some remarkable tools to illustrate the case, consultant and researcher Bill Brockbank examines and quantifies the impacts before drawing some general conclusions.

For this paper we have used a retail supply chain, although the lessons apply to all ex-stock chains. Retail, characterised by low rates of sale (per SKU, per store) and the universal need to 'push' some stock in anticipation of sales is the clearest illustration of inertia. It's also the area with the biggest profit potential, much of it unrecognised.

To distinguish between 'speed up' and 'slow down' inertia, we'll call the latter **momentum**.

We'll examine the micro chain, between the store and its local (DC) distribution centre. The lessons are the same for the supplier >> DC echelon, for the manufacturer >> importer, for the manufacturer's supplier to the factory and so on. In one sense the momentum at these echelons is greater because the lead time is longer. On the other hand, the impact of momentum is less clear-cut. That's because the sole purpose of a retail chain is to have something in each shop. Everything else is, broadly, a step along the way. We shouldn't care about out-of-stock in the DC *providing we lose no shop sales.* Yet too many firms measure line fill (the converse of out-of-stock) in the DC because they can; for all sorts of reasons they don't measure out-of-stock in the stores, or they report something easy to measure which purports to represent OOS[1].

Our retail store also has 'a place for everything and everything in its place'. Corn Flakes can't overflow into the Rice Krispies shelf space or vice versa. Typically, mid priced goods with fixed shelf allocations work like this, while hanging goods do not.

In essence the micro chain contains:-

A number of retail outlets each with a target stock for each ranged SKU. We'll examine only one shop[2], from which the lessons are scaleable.

A DC (or factory, or agent, or importer or wholesaler) who supply, on demand, SKU's to replenish the shop.

A control mechanism with 'triggers' and 'rules' to make this happen.

We illustrate the simplest case: The store is replenished with singles through a BTL (Build To Level, i.e. the on shelf + in transit quantity is 'topped up' by the DC to the BTL. Under steady state this settles down to SOGO, 'sell one, get one') which stays constant throughout. The DC always has enough stock. The store sells the same total amount week after week, and the lead time of resupply (DC >> store) is constant. In the real world, everything is worse than this. Almost 60% of all sales take place on Fri/Sat/Sun; BTL's change like the weather; cheap items come in

---

[1] A leading high street chain thought its store out-of-stock was 11% when the customers saw 28%. The differences? Inaccurate store stock records; and counting backroom and in transit stock as 'on display'.

[2] It's a pretty busy shop! We ran over 2,000 combinations of BTL, rate of sale and lead time, each 10,000 times. Our "one store" not only worked extremely hard, it experienced every possible combination of circumstances.

Tel: 01279-726806, FAX 01279-726603; Mobile 07802-768078

box quantities, the warehouse shuts at the weekend so lead times go out, and so on. No matter, the purpose of this paper is to illustrate some fundamental truths about supply chain momentum. However alarming those results might be, in the real world things are worse. However large the opportunities under steady state, in the real world they are larger.

Configuring and running the tools

The tools are taken from the General Retail Model Suite, with capabilities including the trade off between service level and lead time, between lead time and inventory, between rate of sale and inventory or lead time or service level, between box quantity and service level, 7 day or 'skip day(s)' deliveries, shelf space and sales and so on.

Most of these are step relationships – at a fixed rate of sale and lead time the service level at a BTL of one might be 92% whereas at a BTL of 2 it might be 99%. There's no such thing as a BTL of 1.2 items! [3]

In the past that meant it was difficult to compare like with like … there was no obvious way to hold the service level constant and look at the inventory vs. service level trade off. GRMS now contains proprietary methods to bridge the gap, although these runs are on the earlier (integer) model. For that reason the illustrated rate of sale is higher than usual … at higher rates of sale the step increase in service level (between anyBTL and anyBTL+1) is smaller than at low rates of sale.

Method

We ran a GRMS tool at a range of rates of sale, lead time and inventory. The tool mirrors a real supply chain, so sales fluctuate about an average. Because sales of each SKU vary through pure chance and the BTL (on-shelf plus on-way) remains a constant the resultant on-shelf stock varies. With a one week lead time the on-shelf stock today is a consequence of yesterday's sales (luck, good or bad), the day before's sales (more luck) and so on. At this point we should talk about luck. The more we run demand data (real or simulated) through GRMS tools the more we are struck by how much a part luck plays in real world supply chain results. A large order (demand) can occur just before or just after a delivery; just before or after a sales lull or spike, just before or after we redid the forecast and adjusted the BTL; just before or after a previous spike washed out of the forecast history, just before or after the 'open to buy' switched in or out.

Since we now enjoy infinite processing power one way to insulate against luck is to run the data long enough for good, bad and average luck to stabilise[4].

We found that 1,000 & 2,000 runs per rate of sale + lead time + inventory combination gave 'lucky/unlucky' results. When we ran 10,000 of each combination the results were relatively logical and consistent.

For each run we measured the actual service level and the stock on shelf at a fixed point in the day.

Since we wanted to understand the relationship between shelf fixturing and the other factors, we measured just after the day's delivery has been stacked on the shelf. This is the daily stock peak. When there's too much stock for the shelf fixture, this is the 'why have they sent this/where the hell do I put it?' issue. (Conversely, when we look at lost sales we look at stock just before the new delivery, when the shelf is at its emptiest)

Backroom stock was not considered. Firstly, backrooms are expensive, are generally being phased out and the space given over to retail. Secondly, back rooms require the goods to be stored & picked again. Surely the DC already

[3] In passing, the conversion of decimal sales forecasts to integer BTL's is an area fraught with misunderstandings.

[4] Another way is to run the data both forwards and backwards.

did that (and likely did it a lot more consistently than hundreds of retail stores.) Thirdly, the greatest prize in retail is the extra sale made from existing space, staff and infrastructure. Conversely the greatest loss is the sale we didn't make because we paid all the costs but were out-of-stock – even if the stock was actually in the backroom[5]. Anything which infringes our display space (as backrooms often do) has to go. Lastly, stock in the backroom is too often 'pulled' onto display, not pushed. In other words, it's too vulnerable to human and control failings. There's a retail adage "If it's not on the shelf, it's not facing the customer[6]" – backroom stock rather fails on all counts.

There are many ways of looking at the results. We chose the one closest to the merchandiser and store manager's hearts. "For a fixed shelf frontage, how big must the shelf be; how often will it overflow?" One such graph is in Figure 1, which shows what happened with next day replenishment.

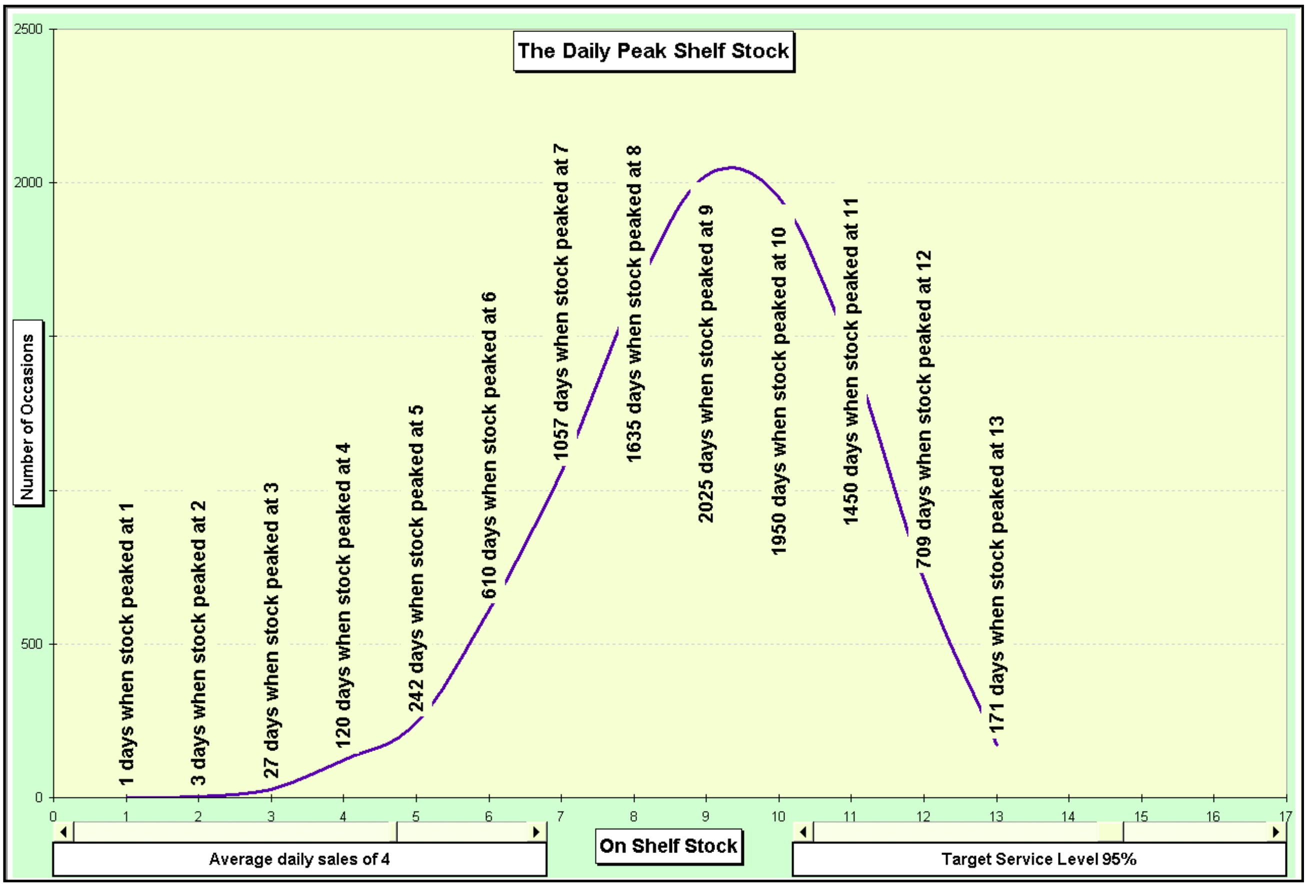

[5] No weighting is given to switching factors, e.g. the customer buys something else in our store, comes back another day, or shops elsewhere.

[6] Or, as the late Bill Bailey used to say about pushing all the stock into his shops "If you're gonna shit like an elephant, you gotta eat like an elephant"

Tel: 01279-726806, FAX 01279-726603; Mobile 07802-768078

The daily peak stocks vary across a wide range. Colloquially, "some days the shelf overflowed, other days it was almost empty".

The daily trough stocks (not shown) clearly go down to zero. There were many days when we ran out of stock, which is why the service level achieved was only 96.4% (This work was done on an old version of GRMS, which did not interpolate results between a BTL of 12 and 13.

At high rates of sale[7] the difference between integer an interpolated BTL's is small and accounts for some unevenness in the graphs. It's not substantive; the conclusions are still pretty brutal)

Here's the picture at 2 day replenishment:-

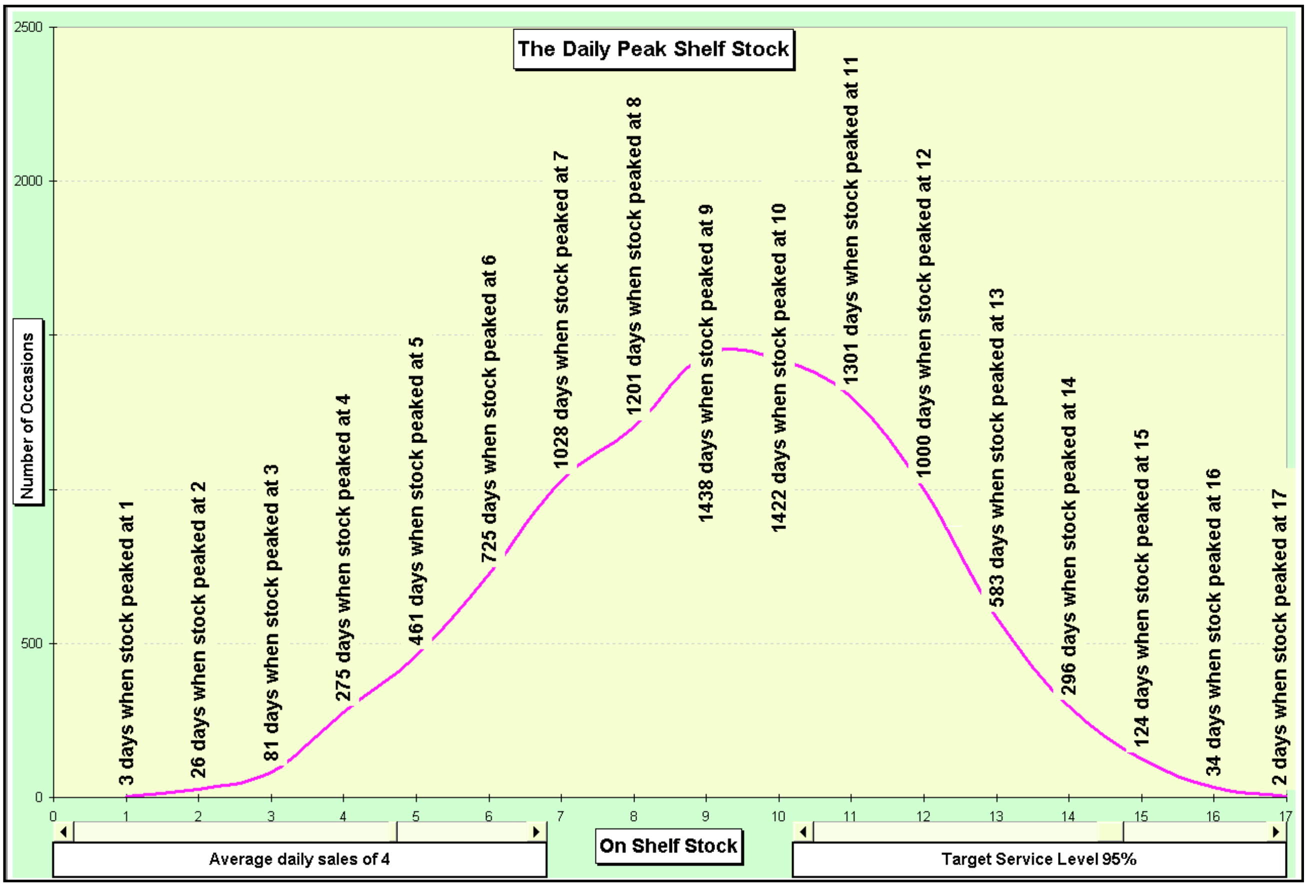

[7] 4 a day is a high rate of sale for retail. The average medium ticket (£25 to £150) item probably only sells once every 3 weeks. In hypermarkets, 90% of goods sell less than 10 a day.

Service level with a BTL of 17 was 95.2%

The average stock after each morning's delivery was 9.25, which compares with 9.03 at one day replenishment. Compared to the 1 day curve the 2 day curve has 'pancaked'. Since the left hand side of all these curves is pretty much anchored (we can't have a stock less than zero) the 2 day curve must 'spread right'. All of the averages (modal, mean and median) move right. More importantly, since shelf fixturing must store some if not all of the peak, the right extreme also spreads right *faster than the average*.

[Figure 4 will show all replenishment times from 1 through 10 days, at the same rate of sale and approximately equal service level.]

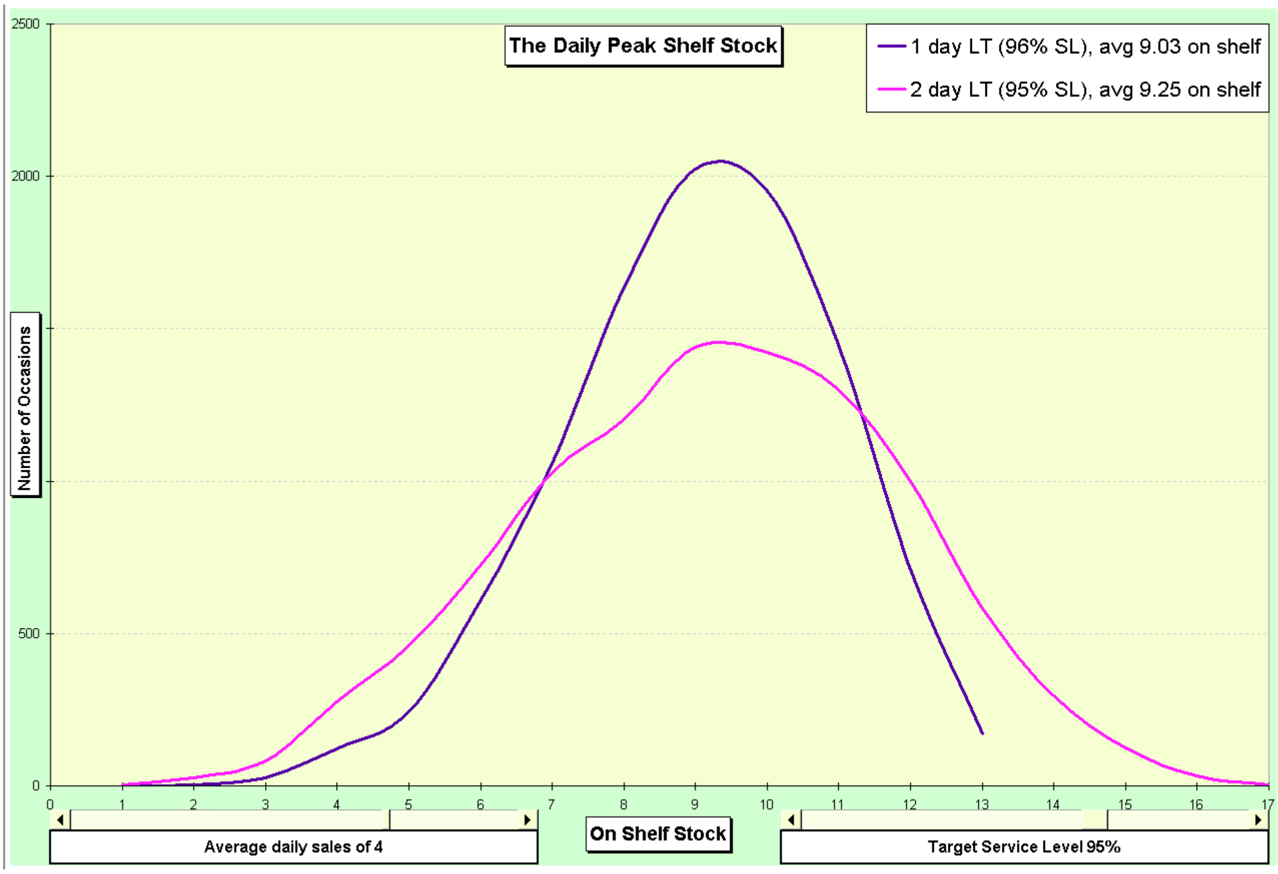

Figure #3 – Figures #1 and #2 compared, with the text labels removed.

Clearly there's an enormous difference between 1 and 2 day replenishment, the latter requiring more stock in store *(without any compensating increase in sales)* **and** overflowing the shelf more often.

Stepping right back from the graphs, let's try and put this into store manager's language.

[Store Manager] The difference between 2 day and 1 day replenishment (in this example) is that on average I need shelves to be 2.5% (9.25 plays 9.03) bigger?

[BB] Yes.

[Store Manager] Hold on, that's almost one third of a day's sales. Is that right?

[BB] No quite; the true picture is ***worse***. Across a very broad range of service levels and rates of sale the general rule of thumb is that each day's extra lead time adds 40% of a day's sales to the average on-shelf stock I've seen it as high as 80%, but 40% is a good average. Also, you sold 1.2% less from the larger shelves, which they were 30% bigger (holding 17 not 13) not 2.5%. On top of that it's likely that the business is holding a complete extra day's stock in the DC, hard allocated to your store. Alternatively the extra day's worth of stock might be in transit.

Either way the costs are real – increase lead time by a day and you'll have around 1.4 days extra stock in the pipeline.

[Store Manager] What about this overflow? How much bigger does the shelf have to be to always store the morning delivery?

[BB] At 2 day replen the shelf needs to hold 5 more items than one day replen.

[Store Manager] Wait a minute, that's a whole day's sales?

[BB] Yes – life's a bitch.

[Store Manager] What if I take, say 1 chance a fortnight that I'll overflow the shelf?

[BB] The (now smaller) 2 day shelf needs to be able to hold 1 more item (25% of a day's sales) more than the (also smaller) one day shelf.

[Store Manager] Is that a general rule of thumb?

[BB] No, you got lucky. The small 2 day shelf overflowed about 10% more often than the small one day shelf.

[Store Manager] OK, let's up the chance of an overflow to once a week. Does that make it better?

[BB] No, worse. You have more than twice as much restacking to do (the frequency of restacking doubled, but there are now more times when you restack 2, 3 or 4 than before). In addition, the 2 day shelf still needs to be able to hold 1 more item.

[Store Manager] Whoa, this is getting silly! Under either average or peak fixturing we're talking about allocating shelf space to products under the 2 day scenario that, with 1 day replenishment I could use to extend the store range?

[BB] In principle, yes. Using these example rates of sale and service level in a 'pigeonhole' store where every slot holds one item and can't be used for any other item, a quarter of the display space in your 2 day store would be wasted[8].

[Store Manager] Some of my space is hanging, does that make a difference?

[BB] You still need to hold more *average* stock under 1 vs 2 day replenishment; that penalty remains the same as for boxes. Where hanging lessens the penalty is when we look at the *peak*. So, taking our example, we still 'waste' a minimum of 2.5% of the retail space with 2 day replenishment vs one day, and sell less. Where hanging (and other 'generic' stock space) helps is that an overflow on (say) 36 waist/31 leg jeans can go into a hole left by unusually high sales of 38/29 size.

[Store Manager] So for fewer sales I always lose 2.5% space and sometimes lose 25% - what do you think the average is?

[BB] It's between the two, but all sorts of factors mean that it's closer to 25%. The only stores that seem immune are those who make a virtue of stock in depth – giving the customer a 'warm, cuddly feeling' that they have more stock than you could ever buy. Other than real fast moving grocery the only one I can think of is Bed, Bath and Beyond (the US chain)

[Store Manager] Why them?

[BB] I can only guess, but it would suit high margin cheap items in continuous (non-seasonal) supply. And cheap land and building costs. That's pretty much no-one, at least in Europe.

[Store Manager] Hey, I almost forgot. Our DC shuts down at the weekend. Saturday's sales only get replenished Tuesday running into Wednesday, with Sunday's sales replenished Wednesday and Thursday. Because weekend sales are so high there's not enough picking or packing capacity to replenish Monday's sales until Thursday. On Friday we get Tuesday's and Wednesday's

[8] Assuming no backroom or shelf reallocation.

Tel: 01279-726806, FAX 01279-726603; Mobile 07802-768078

replen. I've been thinking we had 2 day lead time, but that's only 1 day a week. All the rest of the time it's 3 to 4 days. You're going to tell me that's worse, right?

[BB] Yes. How depressed do you want to be? The calculation is a bit tricky because service level falls off a cliff as lead time extends, but if you take the rule of thumb 0.4 days average stock in store for each extra day's lead time, then you are erring on the safe side.

The calculation might look like this:-

Day of sale	Proportion of sales	Time to replenish	Extra days stock (over 1 day replen)	Weighted by sales that day
Monday	8%	3 days	0.8 days	0.06 days
Tuesday	10%	3 days	0.8 days	0.08 days
Wednesday	12%	2 days	0.4 days	0.05 days
Thursday	13%	4 days	1.2 days	0.16 days
Friday	16%	3 days	0.8 days	0.13 days
Saturday	23%	3.5 days	1 days	0.23 days
Sunday	18%	3.5 days	1 days	0.18 days

Average additional on-shelf stock	**0.89 days**

[Store Manager] Let me get that straight, the true benefit is greater than the rule of thumb?

[BB] Yes

[Store Manager] Does anything make this better?

[BB] Apart from shorter lead times, there's only one thing, and it's not really 'better', merely 'less bad'. That's where you already have free shelf space. So if a 'perfect' BTL only fills 2/3rds of the Snickers shelf there's no opportunity to stack (say) overflowing Mars Bars in front – you might as well load the shelf up with more Snickers than you really need. The same applies to slot hung blister packs; if all the hooks in a fixture are the same length and the items are cheap, then you can sometimes afford to 'load them up'. Mind you, you need to keep this 'something for nothing' argument in perspective. One of the big sheds has refixtured their stores to halve the length of the slot hooks and still make the store look full. There's fresh air behind the fixtures!

[Store Manager] OK, but my inventory goes up, right?

[BB] Right, but it may not be a lot. We're talking penny items here, and they may be a very small proportion of the stock.

[Store Manager] Yes, they're about a tenth of my range and 1% of my value. What about, say, car batteries?

[BB] For medium movers the answer is always fast replenishment.

[Store Manager] Bikes?

[BB] For medium movers, fast replenishment. For very slow movers, where the BTL is always 1, you might not see any difference between 1 and 2 days.

[Store Manager] What about between 1 day and 7 days? Bikes come through a different channel.

[BB] If a bike only sells once a year then the service level at 1 day replen is so close to 100% it doesn't matter. The thing which might catch you out would be twins wanting a matching pair, and you're never going to stock 2 against that chance! If the replen time goes out to 7 days, service level falls by a couple of %.

[Store Manager] I'll cope with that! So what you are saying is that it's the 'soft middle' where all the money is made or lost?

[BB] Yes. For slow costly items the answer is always 1 (or 2 or 4 if it's tyres) and reasonably fast replen. For fast, cheap items the answer is 'enough'. It's the middle movers where fortunes are being lost through slow or erratic replenishment.

[Store Manager] What about the P&L?

[BB] As a generality, store costs are fixed. So any extra sales you make are at gross margin, while the sale you were already making covered all the overheads, so only made nett margin. In most retail outlets the biggest prize is to sell more with the same space and staff. If slow replenishment is robbing you of that opportunity this 'margin multiplier' effect can be enormous. I've seen 15:1, but 10:1 is certainly common.

On the transport side you have to look on a case by case basis, the same at the DC.

What we normally find is that the change in DC costs are trivial. Firstly, it costs the same to pick something today as it would to pick it tomorrow. Weekend working costs more, but on the other hand the asset has 40% more capacity (7 days vs 5) so lasts longer or is able to operate at a more even pace rather than a 'Monday peak' which only clears just in time for the next Monday!

On transport it's worth getting a quote. Every time we've looked the margin multiplier on the extra sales would more than pay for the whole logistics chain, let alone the *increase in the transport element alone* through delivering more often.

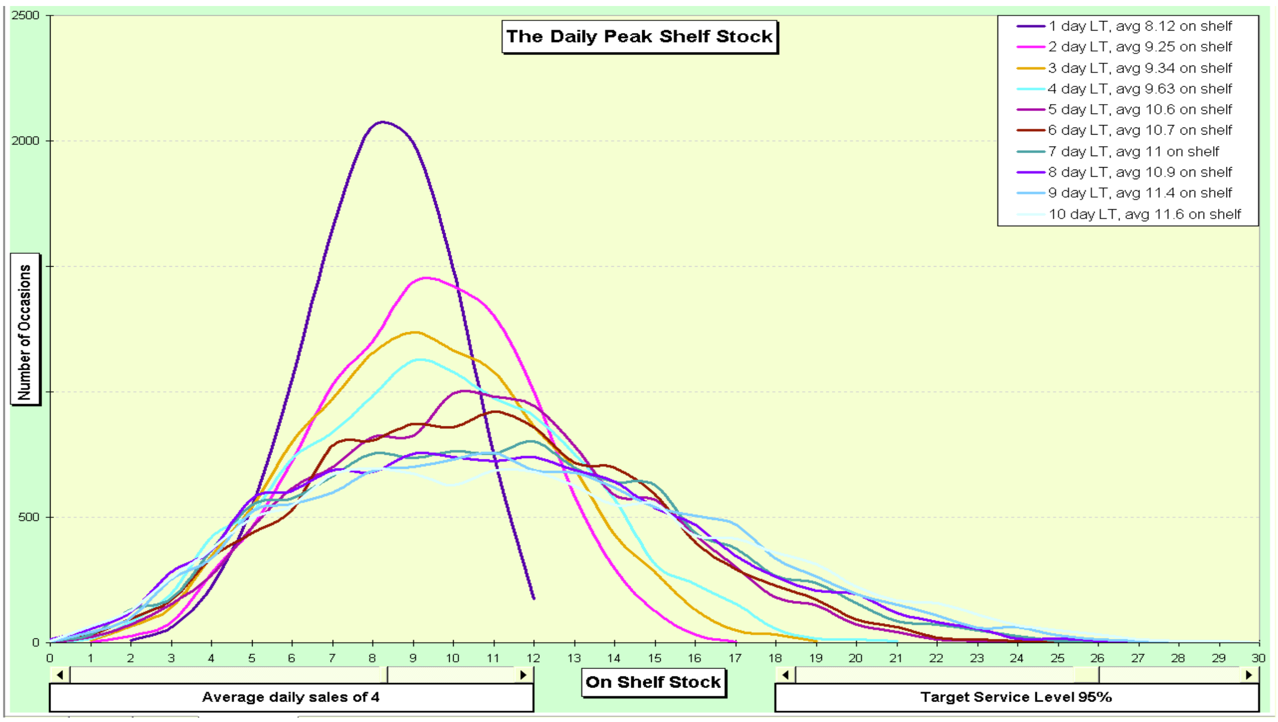

Figure #4 - Also shows the sliders used to browse different rate of sale and/or target service levels

Tel: 01279-726806, FAX 01279-726603; Mobile 07802-768078

Figure 4 shows the 2 original lines and the additional lines for days 3 through 10. The broad picture is the same, that as lead time increases at constant service level the average stock on shelf goes up, and at any reasonable shelf size the number of occasions when it overflows also goes up. Here we see that the average stock on-shelf increased by 3.5 items as we moved from 1 to 10 days lead time. That's 0.4 additional *items* per additional day of lead time, or 0.1 days of stock per added day of lead time. Does this disprove the rule of thumb, which was 4 times higher?

No. At high rates of sale & low-ish service level (I'd want better than 95% service level on something selling that fast!), 0.4 additional stock per day of extra lead time is an overestimate. At higher service levels **or** lower rates of sale, the penalty is higher, often much higher.

The conclusions make eminent sense – at any (constant) service level and BTL combination, shorter lead times give lower average stock and fewer shelf overflows (peak stocks) for any fixed shelf size. The broad conclusions remain the same whatever the rate of sale or service level target. The specific conclusions are subject to luck, both good and bad. For every slow mover where an extra day's replenishment makes no noticeable difference there are typically 5 to 8 medium movers where the extra day may (or may not) tip the balance between BTL and BTL+1

What is remarkable is the extent of the overflow, especially at the 1 to 2 day jump. When selling 4 a day at 95% service level, the difference between 1 and 2 day replenishment is that the latter needs a shelf 4 or 5 items larger.

This is not an area of diminishing returns, quite the opposite. I think that's part of its fascination, and perhaps one of the reasons it's not widely understood. The biggest prize is 'squeezing the last hour' out of the replenishment cycle.

SUMMARY

Momentum in the supply chain is an important concept, particularly to retailers.

Extra shelf space required to store 'momentum' goods is greater than is generally realised. The extra space escalates sharply as the lead time increases. Larger shelves usually rob the store of the opportunity to sell more. Some other alternatives (overflowing shelves, backroom stock) are at best unattractive and at worst unworkable.

Long or uneven (e.g. over the weekend) lead times lead to valuable store space being wasted.

The extra potential sales come from reducing the out-of-stock or (more usually) by increasing the range. These extra sales are at full margin. Because so many distribution and store costs are fixed, the profit from regaining the wasted space (and thereby selling more) is much greater than the profit on goods already sold.

For mid-priced goods this multiplier effect makes ultra-short, reliable lead times one of the biggest prizes in retail.

For high margin medium movers the benefits seem to follow a law of increasing (not diminishing) returns, the shorter the lead time the greater the gain.

This can perhaps be illustrated with a bizarre notion, a negative lead time. If the lead time were (say) 'minus one minute', then the goods would arrive immediately after the customer wanted them.

The store would carry no stock, service level would be 100%. There would be no unsold or over-stock, no markdowns. The store would be tiny, likewise the space and fixturing overhead costs. The range could be infinite, and the profit huge.

Author notes

Bill Brockbank (B.Eng., FILT, FCIT, CMC) has been involved in logistics across the Northern hemisphere for 30 years. An engineer by accident and mathematician by instinct, he has developed and published new methods for warehouse location (Distribution magazine), forecasting slow movers (Institute of Business Forecasting conference, Chicago) and safety stock (Cranfield Supply Chain Knowledge Conference)

Upcoming articles include a re-examination of the Forrester effect; the (tiny) impact of DC out-of-stock on well tuned supply chains; and the damage done by many forecasting systems.

Concerned at the poor communication between logisticians and their peers and 'customers' he founded Supply Chain Tools Ltd to encapsulate and animate supply chain principles in engaging, interactive and graphic ways.

He can be contacted at 01279-726806; FAX 01279-726603; Mobile 07802-768078 or bill@supplychaintools.co.uk