PO Box 1127
Rotorua 3040
Ph: + 64 7 921 1883
Fax: + 64 7 921 1020
Email: info@ffr.co.nz
Web: www.ffr.co.nz

Theme: Radiata Management

Task No: 4.10
Milestone Number: 4.10.09

Report No. FFR- R042

Applications of System Dynamics in Forestry Supply Chains: A New Zealand Case Study

Author:
Thomas Adams

Research Provider:
Scion

Date: 2009

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Supply chains in forestry are beset by variation in both supply of raw materials and demand for processed wood products. Supply variation is due to discrepancies between expected timber yields from inventory sampling, whilst demand variation is due to varying consumer trends.

To ensure that every mill, plant and distributor has enough product to fill orders, they must hold a significant amount of stock which is often costly and can degrade if not used. Optimal stock levels are often found empirically through trial and error. If the business faces a change in production (for example through expansion) then new stock levels need to be re-estimated, which can be costly and wasteful until the optimum levels are reached.

Improving the efficiency of supply chains adds to the FFR initiative to improve this very costly part of the forestry value chain. Reducing supply and processing costs makes New Zealand timber internationally more competitive, which can lead to better prices for the tree grower and reduced costs across the supply chain.

This project uses a quantitative systems approach – known as systems dynamics – to model the flows of a single grade of timber at the Pan Pac timber mill in Napier. Using this model a formula was derived to estimate coverage (the number of days or weeks stock held based on expected production). This is based on the occurrence of a shortfall – when insufficient stock is held to immediately meet an order. The less stock is held, the more frequent these shortfalls occur.

The equation found is specific to this case study, as it was based on confidential production and delivery data. This project shows that the technique can be easily applied to a large range of industries in any supply chain, and that often improvements can be made on empirically estimated stock levels. This paper also shows that there is room for further work, for example considering multiple log grades or balancing the cost of missing orders against that of holding excessive stock. When viewed from a systems point of view, many aspects of supply chain management commonly taken for granted can be improved upon, and this project unveils an accessible methodology exactly suited to that task.

INTRODUCTION

System Dynamics in Supply Chains

Supply chains can be unrelenting places. The world's most successful businesses have a web of suppliers and distributors busy supplying products around the globe, and around the clock. Forestry is no different – for certain the world's forests are growing 24hours a day, but also the many subsequent steps along the supply chain, from raw logs to refined furnishings, are active at all times of day and night. With such constant pressure to supply, it is easy to become swallowed in the day-to-day tactical business of simply filling orders, making it hard to assess the enterprise from a more strategic viewpoint. When strategic changes are made, they are commonly made on a wholesale shift from one established system to another, in a somewhat 'blackbox' manner.

In this study we investigate a technique that permits a reduction of a supply chain into a representation of stocks and flows. The applicability of system dynamics compared to other modelling practices is explored in Adams, 2008. System dynamics has been widely used in forestry supply chains to assess the impacts of the bullwhip effect – a phenomenon where unexpected consumer demand variation leads to amplified demand oscillations for suppliers further up the chain (Fjeld & Haartveit 2002, Forrester 1958, Sterman 1984, Van Horne & Marier 2007). These studies explain an effect generic to all supply chains driven by variable demand. In Adams, 2008, the effect of variable supply (from the forest) was investigated, and found to be attenuated through the supply chain. Thus stocks for enterprises with uncertain supply (such as forestry) can calculate their minimum stocks based on the variance in their own supply, rather than having to additionally consider the tiers above them.

Beyond the bullwhip effect, there are few system dynamic studies of supply chains, and this literature review could not find any relating to forestry. This is because at a tactical level supply chains are uniquely different, and beyond indiscriminate effects such as the bullwhip, there is no such thing as a 'generic' model. At a tactical level supply chains must be considered on an individual basis. To that effect this study does not intend to demonstrate a universal model for supply chains, but instead will promote a technique that creates a common language and allows a rapid modelling of each individual case by utilising some common building blocks.

System Dynamics

System dynamics is a subset of the field of industrial dynamics pioneered by Forrester (Forrester 1958). System dynamics has been increasingly used for modelling business and environmental issues (Ford 1999, Maani and Cavana, 2000). A system is broken down into a set of stocks, flows and variables, all of which must be expressed numerically. The correlation may sometimes be tenuous, for example 'employee happiness' may be modelled with a crude numerical scale. Although the translation of a multidimensional measure to a single number is an extreme simplification, if the model is appropriately constructed around such inevitable simplifications insightful simulations may still be run (Cavana *et al*. 1999 and 2007).

Stocks and flows represent the accumulation and movement of measured quantities, whether they are physical entities, money, people, or less tangibles such as satisfaction or enjoyment. Variables are defined by equations to alter these flows.

METHODS

Case Study

To put this method to test, a case study was selected. Our case study is Pan Pac Forest Products Ltd's sawmill in Napier. Pan Pac is one of the most vertically integrated (encompassing most aspects of the supply chain) forest product companies in New Zealand, which tends to provide better data quality. The case study highlights a very common problem within supply chains – an uncertain supply and uncertain demand leads to a producer holding an empirically estimated amount of stock so that they can maintain production. This is summarised in figure 1 below.

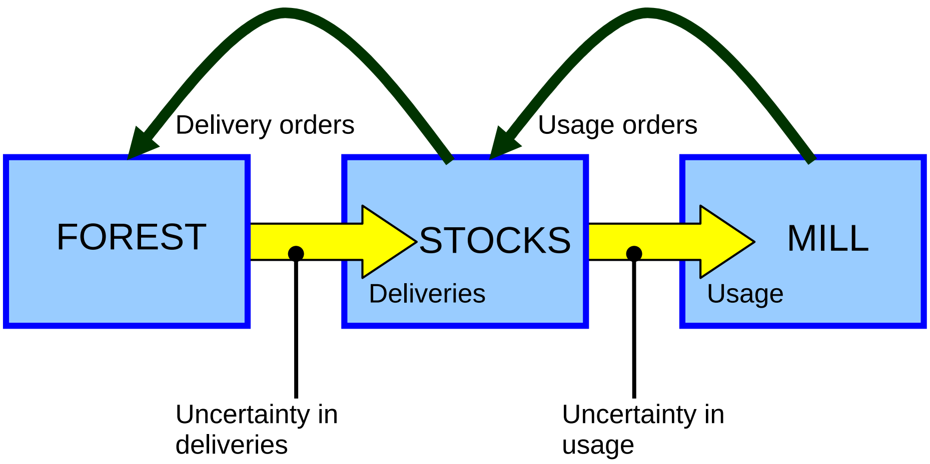

Figure 1. Flow of wood into and out of mill stocks.

Pan Pac have been extremely helpful by supplying us with data for their mill usage and deliveries, both planned and actual, for two years from 2007 to 2009. This is sensitive information, so in many of the following graphs scales will be removed or scaled to unity to protect the information.

The model in Figure 1 is simple enough that it may be applied to numerous other instances in supply chains. To answer meaningful questions though, it must be tailored to a specific case through the equations that define:

- Characterisation of delivery and usage (trends, cycles, variance etc.)
- Definition of uncertainties
- Definition of orders (based on stock levels, previous orders, extrapolated order trends etc.)
- Timing of delays in delivery and orders.

In this case study we will initially only be considering one key log grade. Pan Pac group their log grades into key grades, and although logs may be downgraded within each key grade, there is no downgrading (or upgrading) from one key grade to another. As this model is a proof of concept, it could easily be expanded to include all log grades, but that is saved for a later work.

RESULTS

Characterisation of Data

We will characterise 7 main variables as follows in Table 1.

Table 1. System variables

Variable	Symbol	Definition	Explanation
Planned Usage	U	Exogenous	The planned amount of wood the mill will need
Actual Usage	A	A = f(U)	The actual amount of wood the mill uses
Planned Deliveries	D	D = f(U)	The amount of wood the mill orders in to keep up production
Actual Deliveries	S	S = f(D)	The actual amount of wood that gets delivered
Usage Difference	Δ	Δ = A - U	The difference between actual usage and planned usage
Delivery Difference	Γ	Γ = S - D	The difference between actual deliveries and planned deliveries
Stock	Q	$Q = \Sigma S - \Sigma A + c$	The net stock of wood left at the mill as yet unused by the mill.

Characterisation of Planned Usage (U)

Within this model, the planned usage of the sawmill is exogenous. We are not attempting to model the inner workings of the mill, merely to represent it as an exogenous demand. Figure 2 shows a plot of the mill's planned usage between 2007 and 2009 from the data supplied.

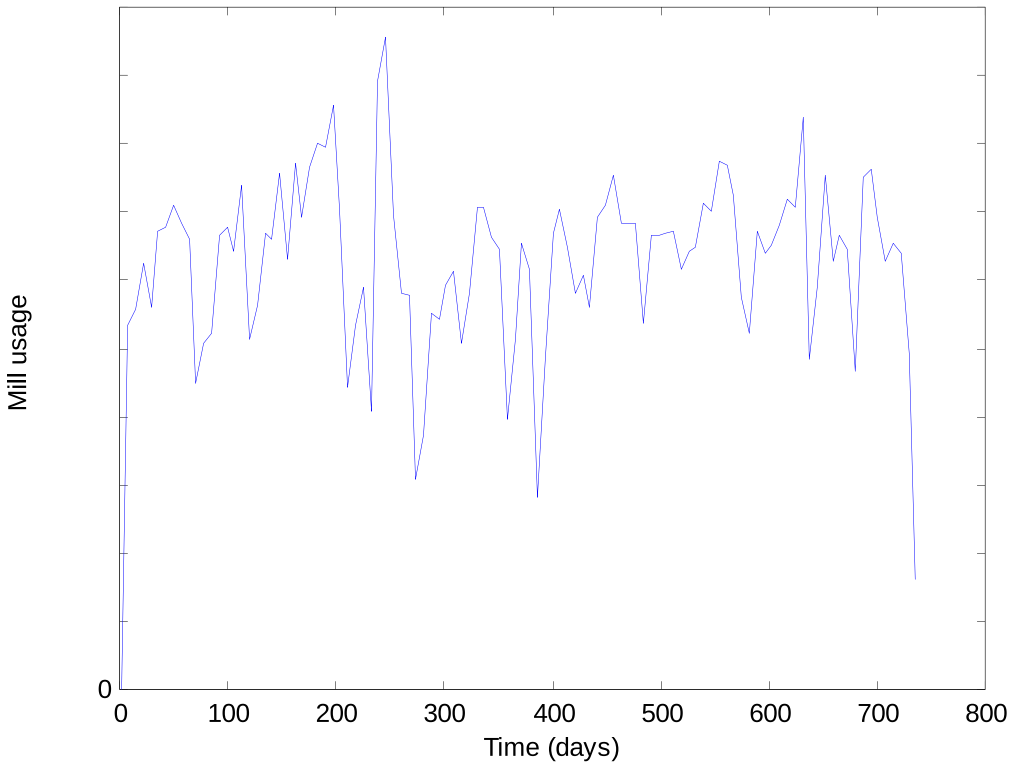

Figure 2. Planned Mill Usage of raw logs from 2007 to 2009.

From this graph it is reasonably apparent that there is no general upwards or downwards trend. However, it is possible that there is a kind of cyclic event creating periodic oscillations (such as the bullwhip effect). To assess this, the sequence was analysed using the Fourier transform, which is used to express a waveform in terms of its constituent frequencies. A graphic equaliser on a stereo is – crudely – performing this function by breaking down the sound into discrete frequency bands.

Expressed mathematically:

If X(n) is a waveform, sampled at N points in time t, with period dt, it can be expressed as a sum of sine and cosine components

$$X(n) = a_0 + \sum_{k=1}^{\frac{N}{2}} a_k \cos\left(\frac{2\pi k t_n}{Ndt}\right) b_k \sin\left(\frac{2\pi k t_n}{Ndt}\right)$$

For more on Fourier transforms the reader is directed to any mathematical textbook. We can take the amplitude of the frequency components F_k where

$$F_k = \sqrt{a_k^{\ 2} + b_k^{\ 2}}$$

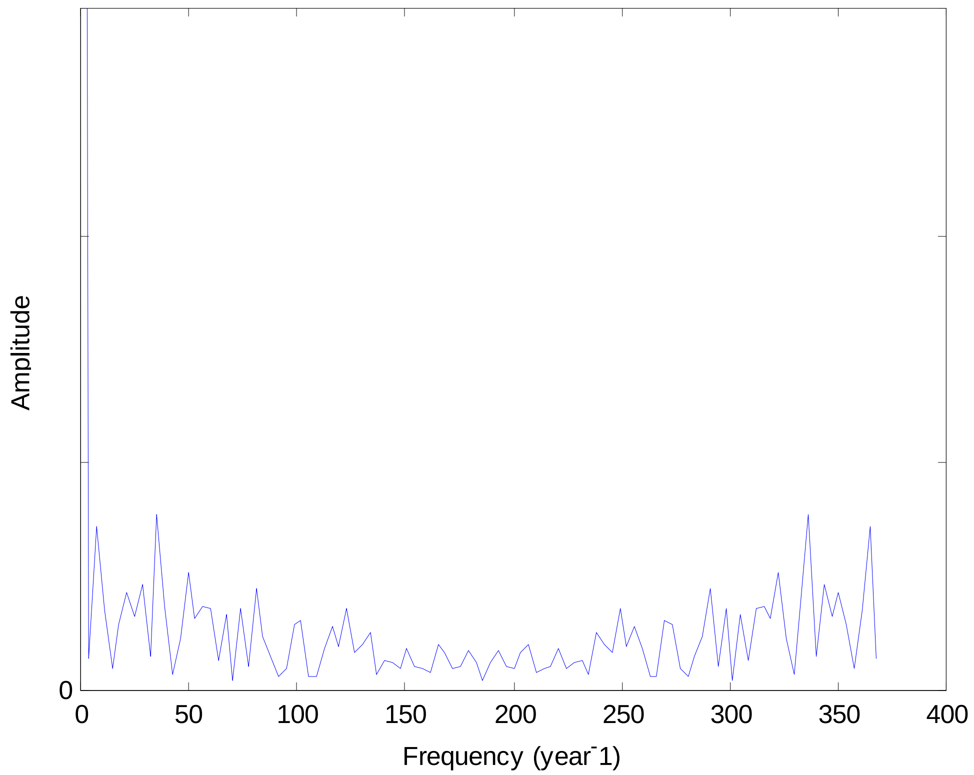

Figure 3. Fourier Transform of Mill Usage.

In Figure 3 we plot the magnitude of F_k against frequency. We see a large spike at k=1 (i.e., a frequency of once a year). This corresponds to the annual shut down over Christmas, shown in Figure 2 by a large drop in usage at t=0, 365 and t=730 days. Beyond this however, little appears in the frequency spectrum other than noise. From this we can assume that there are no cyclic components to the usage, which implies that no bullwhip effect is occurring. For comparison, figure 4 shows the simulated mill usage from the model developed in Adams, 2008, to show bullwhip effect. Even though this bullwhip effect is caused by a normally distributed random variation on the demand, it is apparent that there is a cyclic component in the mill's usage. Figure 5 shows the respective Fourier transform, which clearly shows a set of harmonics at around 4, 8, 12, 16 and 20 cycles per year.

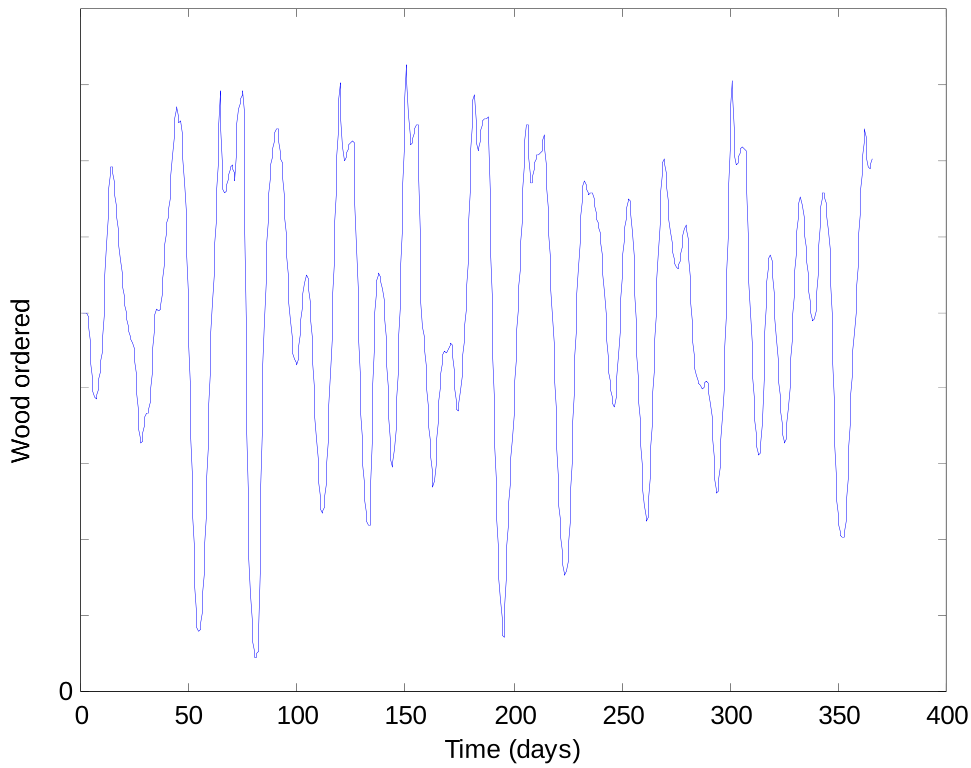

Figure 4. Simulated mill usage containing the bullwhip effect.

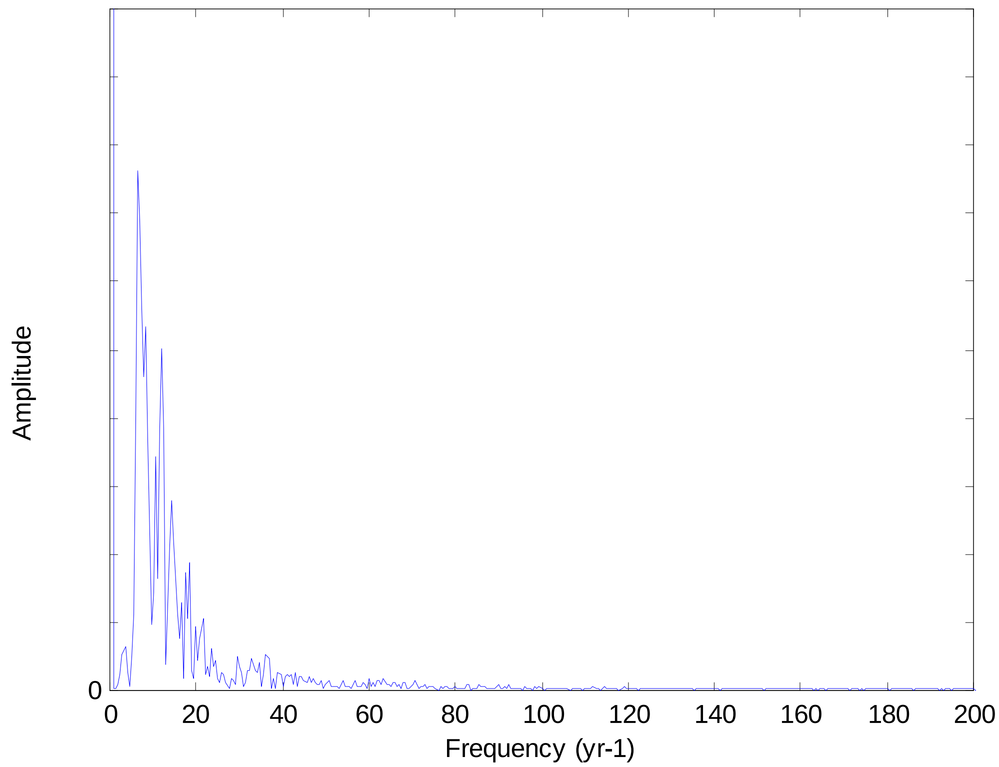

Figure 5. Fourier transform of simulated mill usage containing the bullwhip effect.

With no trends, or frequency components, we can simply model the planned usage statistically. Figure 6 shows a histogram of the distribution over the two-year period. Note that actual volumes have been removed to conceal confidential information.

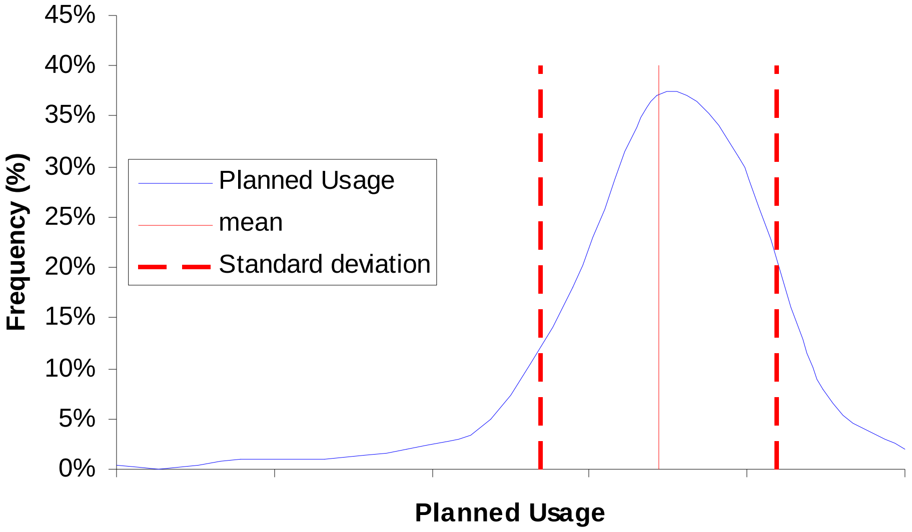

Figure 6. Distribution of planned mill log usage over a two-year period.

Given that this histogram is constructed on only 106 data points, it is not expected to be a perfect normal distribution. From the 106 data points, in a perfect normal distribution we would expect 68% of these to fall within $\pm$ 1 standard deviation. Here 79% do, suggesting that a normal approximation should only slightly overestimate the spread of the data. In conclusion, we will approximate the planned sawmill usage with a normal distribution based on a mean $\bar{U}$ and variance σ_U^2 found from the original data.

$U \sim N(\bar{U}, \sigma_U^2)$

Characterisation of Actual Usage (A)

It is reasonable to assume that actual usage is a function of planned usage

$A = f(U)$

Figure 7 plots the two against each other and confirms that $A = U + \Delta$ where Δ is the usage difference (see table 1).

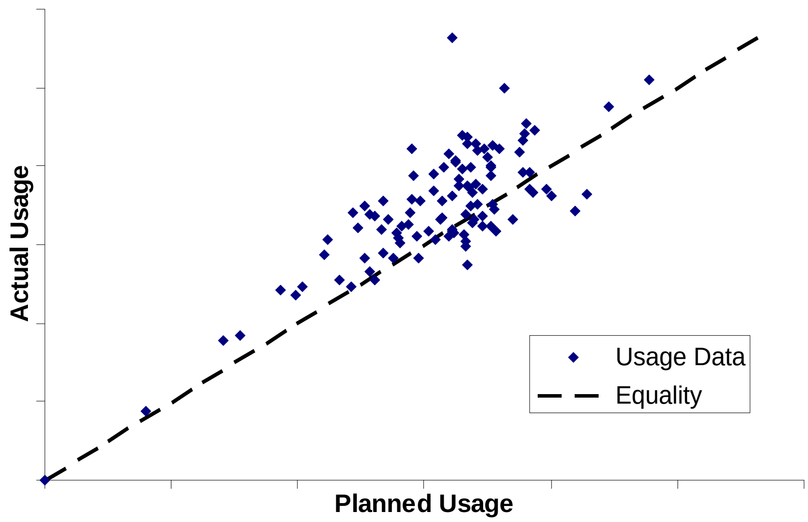

Figure 7. Actual Usage vs. Planned Usage.

Thus we simply need to define Δ, the usage difference.

Characterisation of Usage Difference (Δ)

There are two ways we can define Δ, either

$\Delta = A - U$ (1) Where Δ is the absolute difference

or

$\Delta = \frac{(A - U)}{U}$ (2) Where Δ is a relative difference

Equation 2 has the benefit that weeks of particularly high production will have no greater bearing on the overall distribution than the quiet weeks. However, the distribution is likely to be affected when S >> A or S << A, more so than equation 1. The following analysis was performed for both methods, and it was found that equation 1 gave a better fit between simulated and empirical data.

We can try and gauge the distribution in the usage difference in much the same way as we characterised the planned usage, with a histogram and (if applicable) a normal distribution. Figure 8 shows the usage difference as defined by equation 1.

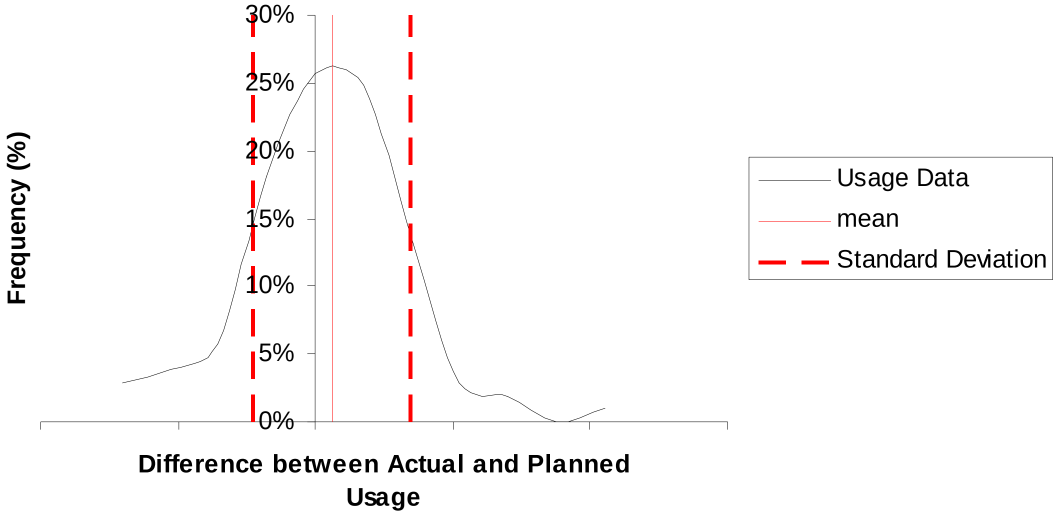

Figure 8. Histogram of Usage Difference.

It is interesting to note that, on average, the difference is positive – i.e., that the usage is slightly more than expected. This means that the mill must continually over-order to maintain stock levels. This distribution is a good fit for a normal distribution, so we can define Δ as a normal distribution of mean δ and variance σ_Δ^2

$\Delta \sim N\ (\delta, \sigma_\Delta^2)$

Therefore, if the actual usage A is the planned usage U plus an uncertainty Δ

$A = U + \Delta$ $\Delta \sim N\ (\delta, \sigma_\Delta^2)$
$U \sim N\ (\bar{U}, \sigma_U^2)$

Then A is the sum of two distributions. It can be shown that the sum of two normal distributions is a single distribution, so

$$A = N\ (\delta, \sigma_\Delta^2) + N\ (\bar{U}, \sigma_U^2)$$

$$= N\ (\delta + \bar{U}, \sigma_\Delta^2 + \sigma_U^2)$$

To evaluate this representation of A we can estimate the distribution of A from the original data as

$A \sim N(\bar{A}, \sigma_A^2)$

Validating this we find that

$\bar{A} = \delta + \bar{U}$ To within 99.9% accuracy
$\sigma_A^2 = \sigma_\Delta^2 + \sigma_U^2$ 90% accuracy

Notably we get the mean of A to be a very close fit to actual values, but tend to overestimate the uncertainty. This is to be expected given that our data were based on a relatively small number of data points, which can cause to lead to larger estimated variances due to a few anomalous data points.

Equation 2 above suggests an alternative method for describing Δ. Under this definition

$$\begin{aligned} A &= U(1 + \Delta) \\ &= U + \Delta U \\ &= N(\delta, \sigma_\Delta^2) + N(\bar{U}, \sigma_U^2)\, N(\delta, \sigma_\Delta^2) \end{aligned}$$

Thus we have a product of two normal distributions as well as a sum. This is best solved numerically, by creating a large set of sample points from this definition, and finding the mean and standard deviation of these. In this case we get an accuracy of 93% for the mean and 85% for the standard deviation. For this reason we use the definition for Δ in equation 1, with values for δ and σ_Δ^2 found from the data.

Characterisation of Planned Deliveries (D)

The deliveries (orders to the mill) must be a function of the orders. As a first logical step, figure 9 plots planned deliveries against planned usage.

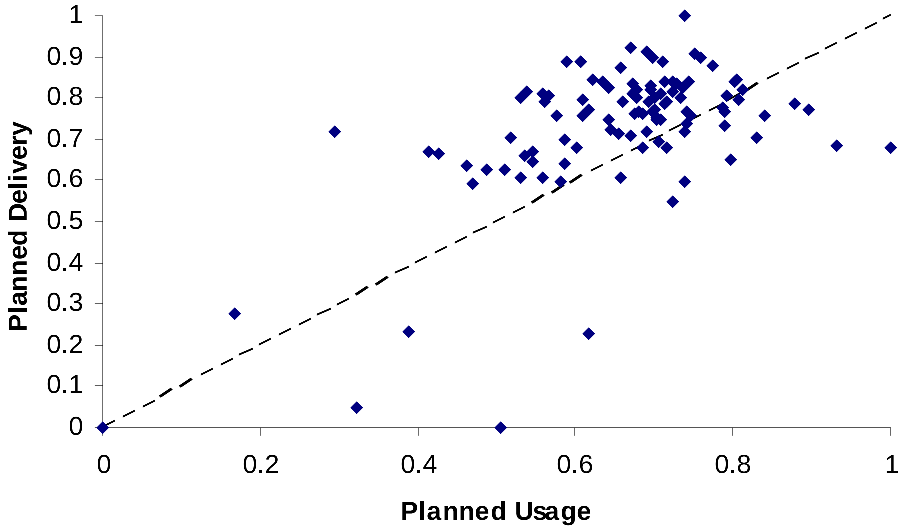

Figure 9. Planned deliveries against planned usage.

There appears to be a possible correlation in Figure 9. An alternative way to manage a company's orders is to order not on the immediate need, but instead order enough to keep a stock of wood at a desired level, to cover a given number of days' production should no deliveries arrive (Maani and Cavana, 2000). This requires us to characterise the stock.

Characterisation of Stock (Q)

We have already defined stock as

$Q = \Sigma S - \Sigma A + c$ Where c is a constant

c is the level of stock held at the start of the data. Every time the supply of wood to the mill outstrips demand (or usage) then this surplus wood will end up in the stock. Figure 10 shows how this stock varies over time, in terms of how many weeks' coverage it contains. We do not yet know c, so c is set to zero giving the stock relative to the start of the data. Absolute stock cannot go negative, but a negative relative stock simply means that stock levels are less than at our datum point, c. From this graph we can infer c is at least 2 weeks' worth of coverage. If we define coverage as the average expected usage ($\bar{U}$) then $c \geq 2\,\bar{U}$

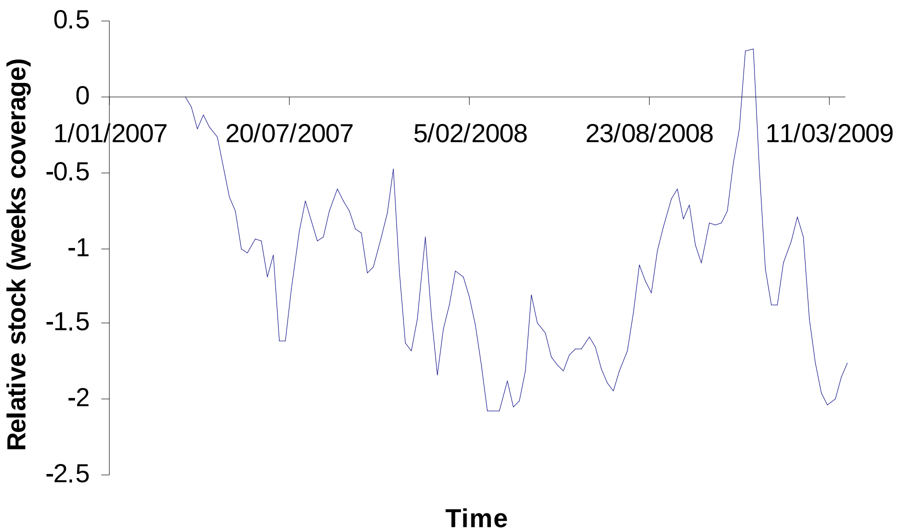

Figure 10. Relative stock (in weeks' coverage) between 2007 and 2009.

Plotting orders placed against the current relative stock will give a good indicator of whether orders are placed based on stock levels or immediate orders. Figure 11 shows that there is only a poor correlation between stock levels and orders placed, suggesting that orders are based on current usage alone.

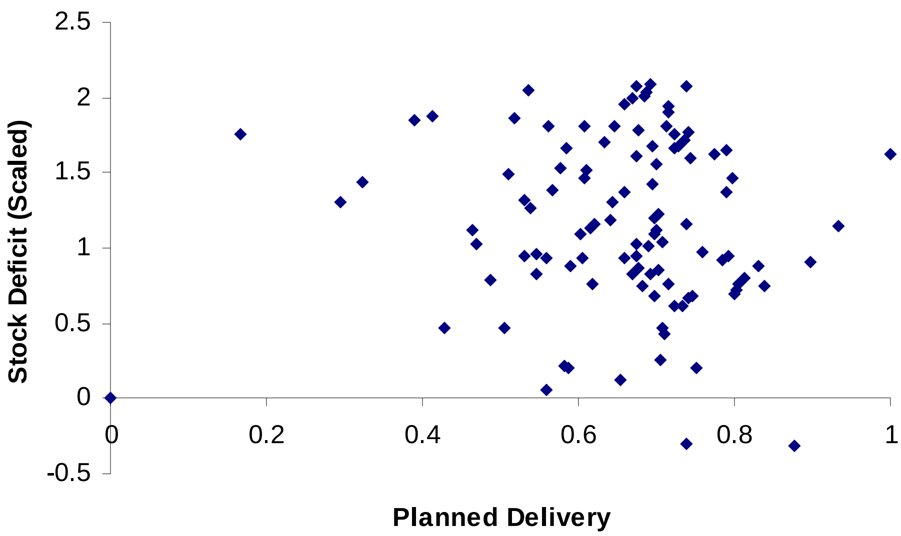

Figure 11. Orders placed vs. relative stock.

This means we can experiment with ordering policies based on the mill's current needs, but also build scenarios based around maintaining a given coverage. It is the major purpose of this paper to explore different formulations of S – which is referred to in system dynamics as an order policy – so the complete function D = f(U,Q) does not need to be fully defined yet.

Characterisation of Actual Delivery (S) and Delivery Difference (Γ)

Figure 12 shows that there is a strong correlation between planned deliveries and actual ones.

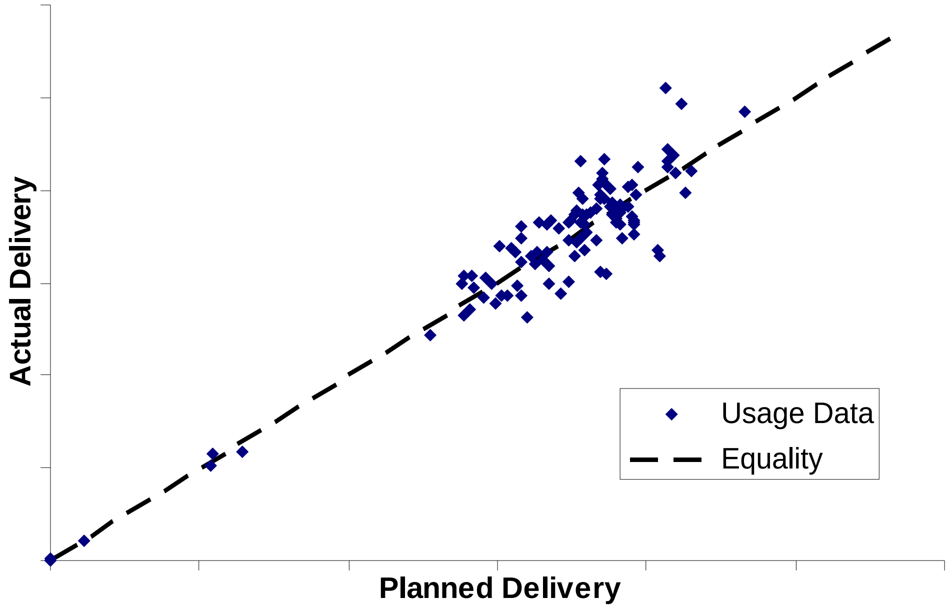

Figure 12. Actual delivery vs. Planned delivery.

Just as we then defined the actual usage in terms of the planned delivery and a difference term, we can do the same for the deliveries. Thus

S = D + Γ

We can also define Γ as previously, with a histogram and normal distribution. Figure 13 shows the distribution for Γ.

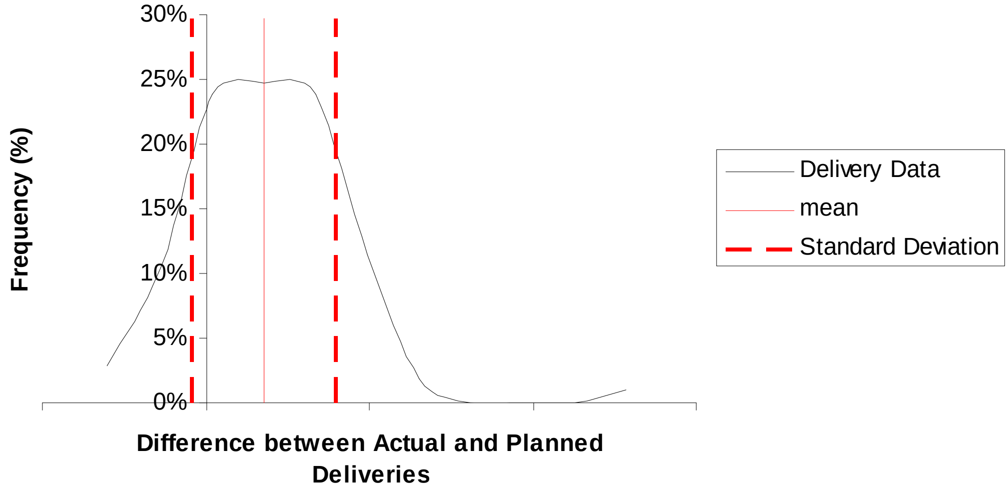

Figure 13. Histogram of delivery difference.

It is then a reasonable assumption to define Γ as

$\Gamma \sim N(\gamma, \sigma_\gamma^2)$

As when we characterised the uncertainty in usage, it is interesting to note that γ is positive, just as δ was. This means that, on average, deliveries are greater than planned.

In conclusion, we have

$U \sim N(\bar{U}, \sigma_U^2)$	Exogenous
$A = U + \Delta$	$\Delta \sim N(\delta, \sigma_\Delta^2)$
$D = f(U,Q)$	
$Q = \Sigma S - \Sigma A + c$	$c \geq 2\bar{U}$
$S = D + \Gamma$	$\Gamma \sim N(\gamma, \sigma_\gamma^2)$

MODELLING

In the previous section, we characterised our variables for the model. This is extremely necessary as the model should represent the variables as faithfully as possible.

The next step is to create a system dynamics model. In this study we used the software Powersim. Figure 14 shows the simple model.

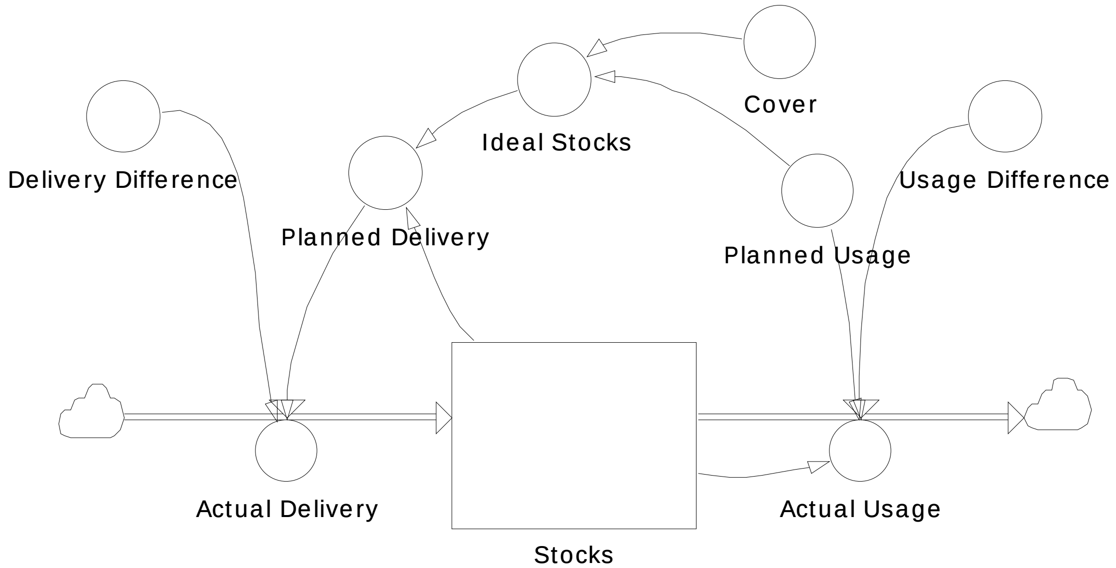

Figure 14. Basic Powersim model.

The model is set to work on a daily timestep, as in real life. By adjusting the equations controlling the variables, we can trial several different scenarios.

Scenarios

Table 2 shows the scenarios that were trialled using the model in Figure14.

Table 2. Scenarios used in Powersim model

Scenario number	Name	Description
1	Just in time	Deliveries are planned to exactly match that week's usage. Stock is not taken into consideration.
2	Never short	Order policy is set to keep the stocks at a level in which they never run out in 10 years.
3	1:10year	Stock levels are set so that in 10 years production is only halted due to lack of stock once.
4	1:365days	Stock levels are set to be too low only once every 1 years.
5	1:91days	As above, except the stocks fall short once a quarter.

Scenario 1
In this scenario stock is considered irrelevant, so is initially set as the average weekly actual usage ($\bar{U}$). The planned deliveries are set to be exactly equal to the current day's planned usage, i.e.,

$D = U$

We make the assumption that deliveries must arrive and be catalogued, therefore deliveries planned to arrive on day X do so, but only form part of the stock on day X+1.

In this scenario we ran 100 runs, and on average stock was insufficient to meet usage on 100 days in the year. Figure 15 shows the stock levels.

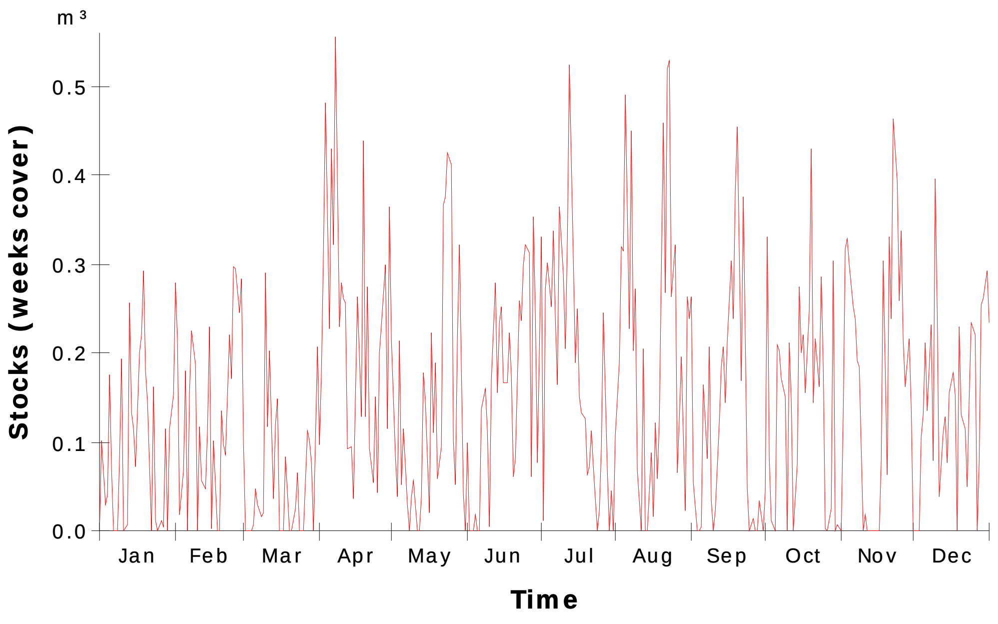

Figure 15. Stock level from scenario 1.

It is interesting to note that the stocks do not appear to trend up over this year. γ (the mean amount of excess delivery) is positive, so on average deliveries will be greater than expected. This will act to raise stocks. However δ is also positive (usage is on average more than planned), which acts to bring the net stock down, and $\delta \approx 6\gamma$. This means that usage dominates and the stocks should head towards zero. This keeps the stocks at their perpetually low levels. This scenario is not ideal, as on average production is frequently limited due to wood shortage. Some form of stock management can prevent this.

Scenario 2 – Never short
In this scenario we set the stock to be high enough that in ten years it never runs to zero. Through a sensitivity analysis, it was found that as long as the order policy acts to maintain a stock of around 9 weeks, the likelihood of ever running short is <0.01 over ten years.

This is an expensive way to run a business though, as the excessive stockpiles are costly in terms of space and capital, and there is a chance that some wood would be unused sufficiently long that it would be spoilt and downgraded to a lower grade. Incorporating a small amount of flexibility into the order policy so that stocks are run to zero on rare occasions, can have great savings in stock costs.

Scenario 3 – 1:10years
This is still a fairly extreme scenario, but using the same sensitivity analysis as above, it was found that a stock pile of 3.33 weeks worth (based on a sliding average of the last 5 weeks orders) is sufficient that stocks fall to zero only once every 10 years. Repeating this run 100 times, naturally on some runs the stocks fall to zero more than once, and on others not at all. However the standard deviation on *n*, the number of times stock falls short is relatively small at 2.01.

Scenario 4 – 1:365 days
Again using the same principle, with a stock set at 2.78 weeks' coverage, the stock fell to zero on average 10 times over a ten year run, with a standard deviation of 4.93.

Scenario 5 – 1:91 days
This is the most likely. Here stocks must be set at 2.64 weeks' coverage to achieve only 40 shortfalls in a year (standard deviation 7.08). It is worth noting that a relatively small reduction in cover (e.g., here between scenario 5 and 4) has a marked effect on the frequency of shortfalls.

This is plotted in figure 16.

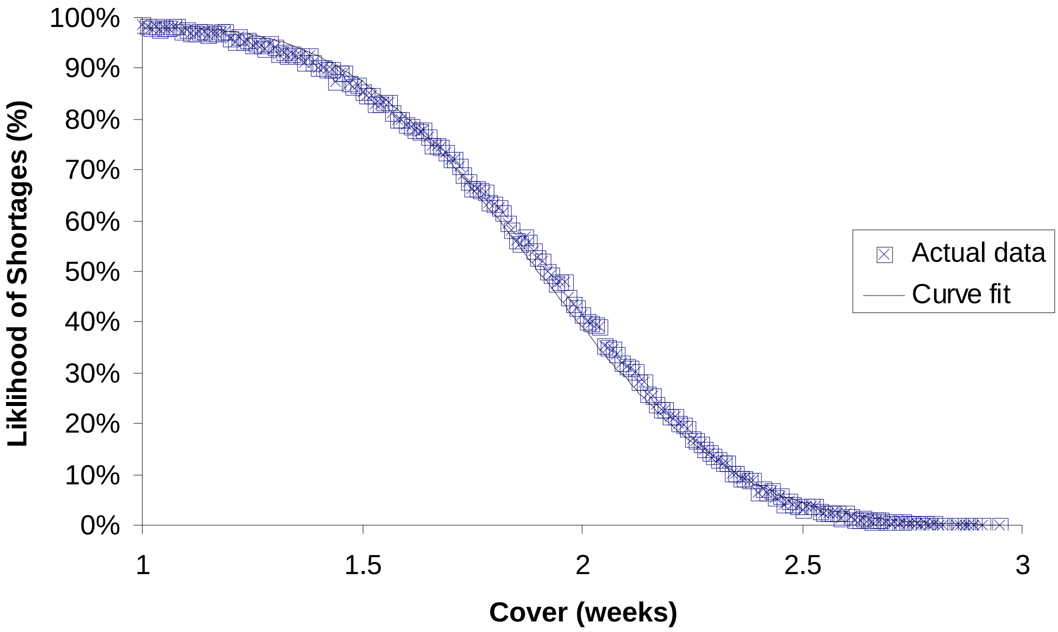

Figure 16. Plot of Likelihood of Shortages vs. Cover.

This sigmoidal curve shows that a coverage of 1 week is almost definitely going to lead to shortages. Under this definition of coverage, a coverage of 1 week equates to just-in-time, i.e., in one week we will use a given amount of wood, so we will order in exactly that amount over the week. Under this scenario we showed that the mean over-usage (δ) is greater than the mean over delivery (γ, in fact δ ≈ 6γ), which leads to a constant risk of shortage. If δ = γ, we would only expect the mill to be short on 50% of days, and if δ > γ this would be less.

As the amount of stock increases to more than just the immediate need, the liklihood of shortages decreases. A cover of three weeks (this week's orders plus two weeks' extra) almost guarantees that shortages won't occur. This curve can be approximated by the equation

$$p = \alpha e^{-\beta(c-1)^3}$$

Where α = 0.980199
β = 0.972558
p = proportion of days when stock was below planned usage (shortfall)
c = cover in weeks

FURTHER WORK

This model could be extended to include degradation of the wood when held in stock, which would act to increase the amount of stock needed. We could also factor in the cost of having stock against the cost of missing production. There would be a trade-off between the two and an economically optimum stock level could be found. Were we also supplied with data on other key log grades, it would be very possible to extend the model to including the other grades.

For Pan Pac, the model could be used to investigate scenarios where uncertainty in supply and demand are reduced. This would change variables α and β in the above equation, and reduce the necessary stocks. By including costs it would be possible to evaluate the savings that these improvements could make.

Applicability to Other Cases

The model in Figure 14 is sufficiently generic that it could be applied to other instances in a supply chain that suffer from an uncertain supply and demand. By varying the uncertainties and definitions given in Table 1, the model can be tailored to explore this problem for a wide range of businesses. Small differences, such as the polarity and dominance of the mean difference in actual and planned supply and demand, can have profound effects on the results (see scenario 5).

CONCLUSION

System dynamics has been explored as a tool to assist in forestry supply chain decisions. In previous reports the effects of the bullwhip effect have been investigated, but it was found that although the bullwhip effect can be driven by a variable demand, it is not caused by a variable supply, such as common in forestry.

To take this further, data from a specific New Zealand case study has been collected and characterised. No periodic oscillations were found, showing that the bullwhip effect was not occurring in this supply chain. The supply (usage) was found to not contain any trends, and could be approximated by a normal distribution. The demand (deliveries) were found to be loosely correlated to planned usage, but could also be related to stock levels. The discrepancies between planned supply and demand and the actual figures were also described statistically.

A basic model was then populated with the characterisations, and scenarios were run exploring different ordering strategies. It was found that a strategy that acted to maintain a given stock level worked best, and this ideally would contain the current week's worth of wood plus an additional two weeks' supply in case of reduced supply and increased demand. This level of stock almost entirely guarantees sufficient stock for the demand. The relationship between the amount of stock held by the company, and the occurrence of shortfalls (when demand is greater than stock) is expressed as a sigmoid curve.

REFERENCES

Adams, 2008. *Impact of Improved Inventory Precision on the Forestry Supply Chain: Literature Review* Contract report for Future Forests Research by Scion, Rotorua.

Adams, 2008. *Supply Chain Modelling – A Forestry Perspective.* Contract report for Future Forests Research by Scion, Rotorua.

Cavana, R.Y., Boyd, D. and Taylor, R. 2007. A systems thinking study of retention and recruitment issues for the New Zealand Army electronic technician trade group, *Systems Research & Behavioral Science*, 24(2): 201-216.

Cavana, R.Y., Davies, P.K., Robson, R.M. and Wilson, K.J. (1999). Drivers of quality in health services: different worldviews of clinicians and policy managers revealed. *System Dynamics Review*, 15(3): 331-340.

Fjeld, D. E., Haartveit, E. Y. 2002 Experimenting with industrial dynamics in the forest sector – A beer game application. *International Journal of Forest Engineering* 21-30.

Ford, A. 1999. *Modelling the Environment: An Introduction to System Dynamics Modelling of Environmental Systems*. Island Press, Washington D.C.

Forrester, J. W. 1958. Industrial dynamics: A major breakthrough for decision makers. *Harvard Business Review* 36: 37-66.

Maani, K.E. and Cavana, R.Y., 2000 *Systems Thinking and Modelling: Understanding Change and Complexity*, (Auckland, Pearson Education New Zealand Limited)

Sterman, J. D. 1984. *Instructions for running the Beer Distribution Game D-3679*, Systems Dynamics Group, MIT, E60-383, Cambridge, MA 02139

Van Horne, C., Marier, P., 2007 *The Quebec wood supply game: An innovative tool for knowledge management and transfer*. Forac Research Consortium.