Vampire's Dilemma: How Hospitals Keep Blood in Stock

Fiona Wilson

In the UK approximately 5,000 units of blood are required every day across the NHS for lifesaving blood transfusions (NHS, 2024). This blood is used during surgeries, in emergencies and to treat long-standing health conditions such as cancer and blood disorders. Furthermore, blood is a perishable resource so it cannot be stockpiled indefinitely. The NHS guidance is that after donation, red blood cells have a shelf-life of just 35 days (NHS Blood and Transplant, 2024b).

All these factors put pressure on the hospitals to use blood supplies more efficiently and reduce wastage, while still having enough to save people's lives. There are additional challenges for hospitals when ordering blood supplies. Firstly, there are 8 main blood types and although generally someone requires a blood transfusion of the same type as themselves, there are some bloods types which can be substituted. For example, O negative is the universal blood type which most people can effectively receive. Secondly, blood supplies come from donations by the public. Hospitals cannot determine who will donate blood, so there is no guarantee they can order the supply they requested. Thirdly, demand fluctuates each day which means it is difficult to predict how much blood should be ordered. One method to improve the ordering process is to use an effective inventory management system.

What is Inventory Management and why is it important?

Inventory management is the system of tracking the amount of stock an organisation has and placing the required orders for more. Managing inventory effectively is important for all businesses with physical stock. An organisation wants to have enough products in stock to meet the customers' demand immediately, however there is usually a storage cost so they don't want to store too many products. This becomes even more important when stocking perishable products because if a product expires before you can use it, then you have wasted money and often there is a cost for disposal too. For hospitals it is very important that they do not run out of blood units because if there are insufficient supplies it isn't just lost business, but potentially lost lives.

Organisations use different models and simulations to predict what is the best number of products to order. They use a framework called a **Markov decision process** to model these situations. This is a type of model where decisions are made sequentially and there is some amount of uncertainty about these decisions. When hospitals order blood supplies, they don't just make one order and leave it there, instead they have to keep making orders sequentially as they need to restock. Further, hospitals don't know how many units of blood they will need on a particular day, so they have to use the past data and estimation in their models. This is the uncertainty in the model because the demand is **stochastic** (random). By running simulations before making a decision in real life, the business can test its plans without putting patients at risk or wasting money. The aim of these models and simulations is to produce an ordering strategy to help an organisation make better decisions. The strategy would tell the organisation how much they should order given the current stock level.

There has been a lot of research into inventory management systems and there are different models available. All models are a simplification of the real world, otherwise it would be too complicated to calculate. For example, when we later simulate the demand for blood to the hospitals, we assume that the amount of blood we order is actually available - this could be a reasonable simplification because blood supplies can be transferred between hospitals if needed. We also assume that hospitals use a 'first-in-first-out' policy for their blood supplies. This means that they use the oldest supplies first.

Our Hospital Scenario

We will simulate the ordering process at one hospital for one particular blood type (and we will ignore the possibility of substitutions). Using the NHS pricing we say that the price of one unit of blood is £180 (NHS Blood and Transplant, 2024a), the cost to dispose of one unit of expired blood at the end of each day is £0.30 (NHS Greater Glasgow and Clyde, 2024) and the storage cost for each day is £30 (National Clinical Guideline Centre (UK), 2015). The cost of not having enough blood could be a human life if someone dies which you cannot really put a price on, however we will use £1,000 as it is a high number and this forces the model to avoid running out of stock. We decide that the average demand for this blood type is 5 units per day using a similar study on Canadian hospital blood banks (Sarhangian et al., 2018).

For our simulations we assume that the decision making process is as follows. At the start of each day we observe the current stock level. We place an order and then proceed with the usual demands for the day. At the end of the day the ordered stock arrives and is added to our inventory ready to use the next day. This process is then repeated every day. We now look at some different models which simulate this hospital scenario.

The Newsvendor Model

The **newsvendor model** is believed to extend back to 1888 when economist Edgeworth used it to decide how big the cash reserves should be to satisfy depositors' demands (Edgeworth, 1888). However, it has since been studied in detail and applied to a range of situations. To best understand this model, imagine that you are a newspaper seller.

Each morning you have to decide how many newspapers to order for tomorrow. These newspapers will arrive in the evening and will be sold for just one day. If we buy more newspapers than there is demand, then the leftover newspapers have to be binned because no one wants to buy an old newspaper, so we have lost money on them. On the other hand, if there is more demand than we have newspapers, then we have missed out on potential profit. In a perfect world, we would order exactly the number of newspapers that are requested, but in real-life we cannot know in advance how many we will need.

We want to order enough newspapers to maximise the expected total profit we get each day which is calculated as follows:

$$\text{Total profit} = \text{number of papers sold} \times \text{price} - \text{number of papers bought} \times \text{cost}.$$

We maximise this by looking at past data to determine what the probability is that we sell a particular number of newspapers. For example, if over the last month the average number we sold each day was 5, then we want the probability that we sell 5 papers to be the highest. But it is also quite likely we sell 4 or 6 papers, so we want them to have a high probability too. There is still a small chance that we sell 0 or even 15 newspapers on a day, though these are very unlikely. We need to have a probability for each possible number of newspapers we could sell. The easiest way to do this is using a **probability distribution** which assigns a probability to each number. Figure 1 shows an example of a Poisson probability distribution we might use in this case. The highest point (the most likely demand) is 5, but the demand varies and there is a small chance it could even be 0 or 15.

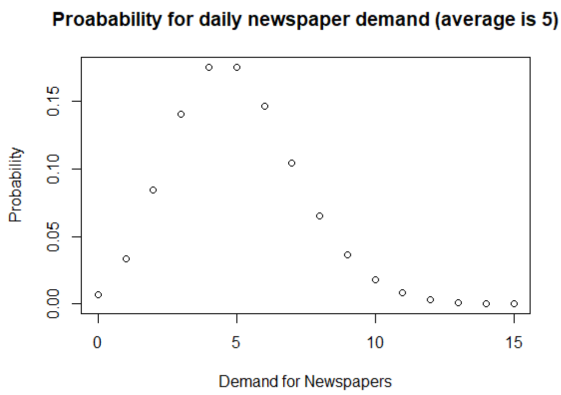

Figure 1: Poisson probability distribution with an average value of 5

With a suitable probability distribution we can find the number of newspapers to order which maximises our total profit. This model has an exact solution because it is quite simple. First we need to calculate the probability of meeting our demand:

$$\frac{\text{cost of understocking}}{\text{cost of understocking} + \text{cost of overstocking}}.$$

Then using a probability distribution like the one in Figure 1, we determine how many we should order each day (Rossi, 2021).

If we return to our hospital scenario we can calculate the optimal order quantity. Using the costs discussed previously, the cost of overstocking is the disposal cost (£0.30) plus the storage cost (£30). The cost of understocking is £1,000. Therefore the probability we want is

$$\frac{\text{cost of understocking}}{\text{cost of understocking} + \text{cost of overstocking}} = \frac{1000}{1000 + (30 + 0.3)} = 0.97.$$

Now we need a probability distribution to compare this to and we will use a Poisson distribution like the one discussed previously. Using our Poisson distribution with a mean of 5 since (our average daily demand), if we want a 97% chance of meeting our demand, then we need to order 10 units of blood each day.

We can run a simulation of this to see how often we satisfy the demands. If we simulate 100 days using the previously mentioned costs, then we know we order 10 units each day. In Figure 2, the left graph shows the first 30 days of this simulation. In these 30 days, only once is the demand greater than the stock, and over the whole 100 day simulation the demand is only greater than the stock 4 times. This means that overall we satisfy the demand on 96% of the days which is very close to our target of 97%. The difference is due to the randomness of the demands and because we only ran one simulation.

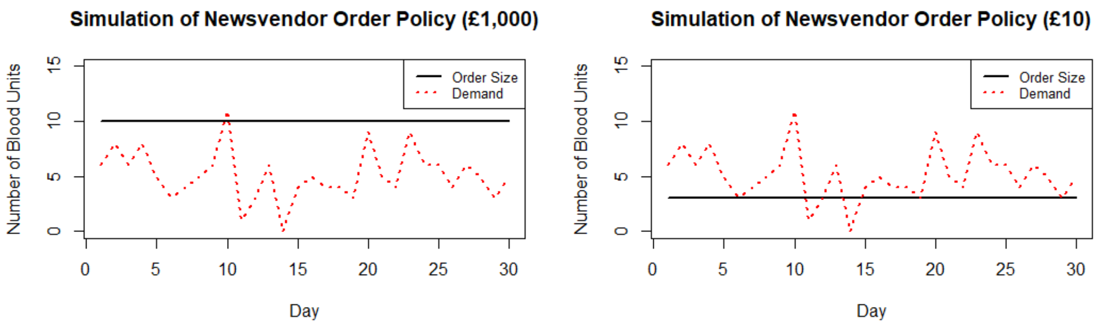

Figure 2: The order size and demand for the newsvendor model where the understocking cost is (a) £1,000 and (b) £10

We set the cost of understocking to be very high (£1,000) because human lives are important. However, if we no longer care as much about unsatisfied demand and set its cost to be just £10, then it all changes. The target probability of satisfying demands becomes just 25% and as the right graph in Figure 2 shows, the order quantity is now reduced to just 3 units, so we often don't have enough to satisfy the demand. Over the whole 100 days, we are missing 226 units of blood and can't satisfy the demand on 27% of days.

The Periodic Review Model

There is another standard model called the **periodic review model (PRV model)** which is almost the opposite of the newsvendor model. The newsvendor model assumes that all stock expires at the end of each day, whereas the PRV model assumes that stock never expires. In this situation it would no longer be optimal to order the same amount each day because if our demand is always less than that, then we would have to store more and more stock. Instead we need to change how much we order each day depending on how much we currently have in stock (Clarkson et al., 2023).

Recall that for the newsvendor scenario we previously calculated the optimal amount of blood units to have in stock at the start of each day was 10 units. In a perfect world, we want to order enough units so that we always start a day with 10 units. Let us consider an example. Suppose that on Monday morning we start with 6 units of blood. We have to place an order now which will arrive on Tuesday morning. If we knew that we would use 3 units during Monday, then on Monday morning we would order 7 units of blood. Although on Monday evening we will only have 3 units left in stock, once the delivery arrives, we will have 10 units ready for Tuesday morning. However, in reality we don't know exactly what the demand will be, so we have to estimate it choosing the value we most expect. We would have to run lots of simulations and eventually we could average them to give us an ordering strategy. This strategy would tell us how many units to order given the current stock level. This doesn't guarantee we will always start the day with at least 10 units of blood, but on average it is the best policy.

The Age-Dependent Periodic Review Model

Although models are never a perfect reflection of reality, neither the newsvendor model nor the PRV model are good enough for our hospital scenario. Blood doesn't just expire after 1 day, nor can it be stored forever, instead it has a maximum shelf-life of 35 days, though it could expire before this. Another model was developed called the **age-dependent periodic review model (ADPRV model)** to reflect this situation (Clarkson et al., 2023). It was originally developed to model perishable products ordered by retailers, however we apply it here to our hospital scenario. The ADPRV model assumes that every day, each individual unit of blood in stock has a certain probability of expiring. Studies have shown that older blood can be less effective at treating people than newer blood (Flegel et al., 2014). To mimic this decay as well as the random chance that blood could expire sooner, we assign a probability that the blood expires on any day within those 35 days. It would be extremely unlikely for blood to expire on day 1 so we set this to have probability 0, but it is certain that the blood expires on day 35 so we set this to have probability 1. As the blood gets older, it becomes more likely to expire, so we assign evenly-spaced increasing expiration probabilities for all the other days. We can use this information along with the previously mentioned costs to run another simulation of the stock levels in a particular hospital.

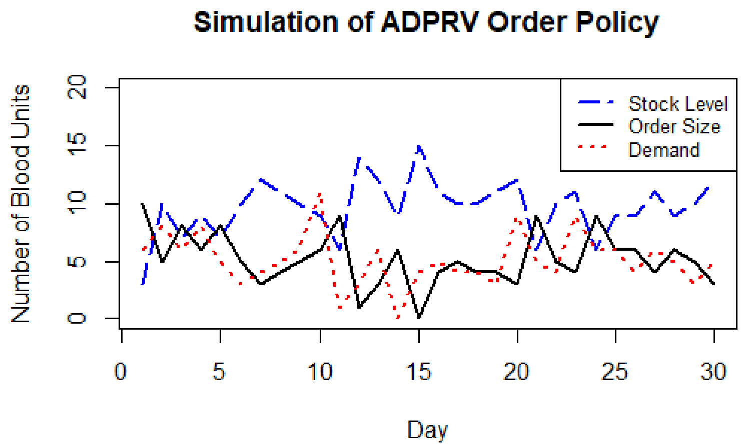

Figure 3: The stock level, order size and demand for the ADPRV model

The graph in Figure 3 shows the first 30 days of a 100 day simulation. We used the same demand in this simulation as we did for the newsvendor model, again shown in red. Now we also have the stock level in blue which changes over time. If the demand is above the stock level, then there will be unsatisfied patients. Over the full 100 day simulation, we were short of 9 units of blood which is slightly worse than for the newsvendor problem. If we wanted to further reduce the unsatisfied demand, we could increase the understocking cost to be more than £1,000 which would force us to order more. When the cost is increased to £100,000 the unsatisfied demand is just 3 units of blood. However, the biggest benefit of the ADPRV model over the newsvendor model (because it is more realistic) is that it reduces the spending costs as well as the wastage of blood units.

To show this improvement more clearly we repeated both the newsvendor and the ADPRV simulations 200 times so that each replication had different demands. The optimal ordering policy would be if the demand and the order quantities were equal, because then there would be no wastage at all and no extra costs. By calculating the cost of the optimal solution and the costs for each of these models, we can compare the efficiency of these models. On average over the 200 simulations, the newsvendor model's ordering policy was 97.5% more expensive than the optimal solution. However, the ADPRV model was only 14.4% more expensive! This is a large improvement and illustrates how effective this ADPRV model can be.

Summary

It is difficult for any organisation to decide how many perishable items they should order at a time. Different organisations have different priorities which also impact these decisions, for example supermarkets want to maximise their profits, while hospitals need to ensure they always have enough blood supplies for patients while minimising their wastage and costs. There are different models to help with this decision making, but we've seen that the ADPRV model works quite well for hospital blood supplies which gradually expire. To make this model more realistic, we also need to have models working together to allow substitutions between blood types. Furthermore, extending the model to allow for hospitals ordering ahead several days would provide them with extra supply security.

References

Clarkson, J., Voelkel, M. A., Sachs, A.-L., & Thonemann, U. W. (2023). The periodic review model with independent age-dependent lifetimes. *Production and Operations Management*, *32*(3), 813–828.

Edgeworth, F. Y. (1888). The mathematical theory of banking. *Journal of the Royal Statistical Society*, *51*(1), 113–127.

Flegel, W., Natanson, C., & Klein, H. (2014). Does prolonged storage of red blood cells cause harm? *British Journal of Haematology*, *165*(1), 3–16.

National Clinical Guideline Centre (UK). Blood Transfusion. 2015. November. Accessed: 26/02/2025. National Institute for Health and Care Excellence (NICE), NICE Guideline No. 24, Appendix N, Unit costs. https://www.ncbi.nlm.nih.gov/books/NBK338809/

NHS. Partners: Saving lives together. 2024. Accessed: 25/02/2025. https://www.nhsbt.nhs.uk/how-you-can-help/partners/

NHS Blood and Transplant. Price List 2024/25. 2024. Accessed: 26/02/2025. https://nhsbtdbe.blob.core.windows.net/umbraco-assets-corp/34945/price_list_bc_cost-per-item_2024-25.pdf

NHS Blood and Transplant. The journey of blood – after you donate. 2024. Accessed: 25/02/2025. https://www.blood.co.uk/news-and-campaigns/the-donor/the-journey-of-blood/after-you-donate/

NHS Greater Glasgow and Clyde. Watch Your Waste to clean up on costs. 2024. Accessed: 26/02/2025. https://www.nhsggc.scot/watch-your-waste-to-clean-up-on-costs/

Rossi, R. (2021). *Inventory Analytics* (1st) [Pages 115-149]. Open Book Publishers. https://doi.org/10.11647/OBP.0252

Sarhangian, V., Abouee-Mehrizi, H., Baron, O., & Berman, O. (2018). Threshold-Based Allocation Policies for Inventory Management of Red Blood Cells. *Manufacturing and Service Operations Management*, *20*(2), 347–362.