Evaluating Solutions for Making Decisions Under Uncertainty

Rebekah Fearnhead

Making decisions is an important part of many people's lives, however there is often knowledge that would be helpful in order to make a decision that is not available when the decision is being made. For example, if you are wanting to go on a walk, you would first want to know if it is going to rain as this may mean that you would want to take an umbrella. Unless it is currently raining the moment you begin your walk, no matter what any weather forecasts predict, you can never be certain it is not going to rain. Whilst in this example the consequences of making a wrong decision are not big - either it rains when you did not expect it to and you get wet, or it doesn't rain and you have therefore carried an umbrella with you for no reason - in industry, problems with uncertainty can lead to much bigger consequences if the correct decision is not made.

Problems like this can be found in many different fields, including transportation, supply chain management and financial asset management [5]. This means that knowing how best to make decisions with uncertainty is an important problem to solve. When modelling decision problems with uncertainty, we are not looking to find decisions that perform the best for each given scenario - when making the decision we don't know what the exact scenario is going to be. We are instead looking to find a way to model decision problems in order to obtain solutions that perform well under the uncertainty we have.

The study of how to incorporate uncertainty into these decision problems is often referred to as **stochastic programming**.

Deterministic Problem

Imagine a person owns a stall at which they sell two different items, A and B. These could be newspapers, for example. Before each day, they need to decide how many of each newspaper they want to order and this stock is the only items that can be sold on that day. For each of the newspapers, there is a known price for buying the stock and a known price a customer will pay to buy each item. For both items, any stock that is not sold during the day is wasted. This is often referred to as the **newsvendor problem**. Assume that they know that there are going to be 80 customers who want newspaper A, and 70 who want newspaper B. Table 1 shows the prices for stocking and selling one newspaper of each kind along with the known

	Newspaper A	Newspaper B
Buying Price	£1.50	£1.20
Selling Price	£3.00	£4.20
Known Demand	80	70

Table 1: Buying and selling prices per newspaper, and the known demands for each product.

	Newspaper A	Newspaper B
Buying Price	£1.50	£1.20
Selling Price	£3.00	£4.20
High Demand	120	105
Medium Demand	80	70
Low Demand	40	35

Table 2: Buying and selling prices per newspaper, and the demands under each of the three scenarios.

demand.

In this case, there is no uncertainty and this is a **deterministic problem**, meaning that the best solution can be found. As the profit wants to be maximised, it is clear that the owner of the stall wants to order the same amount of each newspaper as the demand - 80 of newspaper A, and 70 of newspaper B. As there is a £1.50 profit on newspaper A and a £3 profit on newspaper B, the maximum profit that the person can make is:

$$£1.50 \times 80 + £3.00 \times 70 = £330.$$

Adding Uncertainty

This however, is not a very realistic scenario. In real life, they will not know exactly how many customers they are going to get each day. Assume that instead they think only a third of the time they will get 80 customers for item A and 70 for item B. They also think a third of the time they will get very few customers - 40 for item A and 35 for item B - and the other third of the time their stall will be very popular and get 120 customers for item A and 105 for item B. Table 2 shows the same information as Table 1 along with the additional scenarios for the possible demands.

In real life, there will be more than three different scenarios for the number of customers, and the popularity of the two items does not have to be dependent on each other, however, only considering 3 scenarios will make the **stochastic problem** easier to solve, and the same methods can be used for larger problems as well.

Stochastic Problem

For these two scenarios, the solutions to the deterministic problems are again easy to find. For each of the cases, they will want to order the same number of items as customers they are going to get. This will give profits of £165 for the low demand, and £495 for the high demand.

If the stall owner knew in advance what the demand would be, then the problem would be easy to solve - buy the same number of each item as the demand - however this cannot be accurately predicted. As it is not possible to make a perfect decision that would be best in all three of the scenarios, a solution needs to be found that will perform on average the best over all possible scenarios.

Expected Value of Perfect Information

If we want to **maximise the long-run profit** (the total profit made over multiple days), this can be done by **maximising the expected profit**. This assumes that we are neutral about risk, so are only interested in the profit expected to be made each day, rather than the range of different profits/losses that could be made on any one given day. To do this, we can use the assumption that the probability of each of the 3 scenarios is a third, and then solve the linear programming problem taking into account the chance of each scenario, to maximise profit.

Doing this gives the optimal strategy of buying 80 of newspaper A and 105 of newspaper B. As this would give a profit of £435 with a high demand, £288 with a medium demand, and £21 with a low demand, the overall expected profit is:

$$\frac{1}{3} \times £435 + \frac{1}{3} \times £288 + \frac{1}{3} \times £21 = £248.$$

As expected, this is a smaller average than if the demand was known. If **perfect information** was available before making a decision, the expected profit would be:

$$\frac{1}{3} \times £495 + \frac{1}{3} \times £330 + \frac{1}{3} \times £165 = £330.$$

We can use these values to calculate a metric called the **Expected Value of Perfect Information** ($EVPI$)[2]. For a maximisation problem, this is calculated by:

$$EVPI = WS - RP,$$

where:

- WS is known as the **Wait-and-See Solution** and represents the expected average profit of the optimal solution, where we somehow know the scenario we are going to encounter before we have to make the decision [6].

- RP is known as the **Here-and-Now Solution** and represents the solution achieved corresponding to the recourse problem [4]. This is the linear programming problem where we take into account the probabilities of each of the different scenarios.

This value will always be greater than or equal to 0. If the $EVPI$ is negative this would mean that the value of the RP is bigger than the WS, however, as any RP solution can also be the WS solution, (but not the other way round), then the lowest value of the $EVPI$ is 0, when $WS = RP$.

In the newsvendor problem described above, the $EVPI$ is:

$$£330 - £248 = £82.$$

This value represents the maximum amount a decision maker would be happy to pay in order to receive accurate, and complete, information about the uncertainties in the future - for example the exact number of people wanting to buy each product. It can be seen that for any stochastic programming problem, this would be valuable information to have. The lower the value of the $EVPI$, the better the RP solution, which solves the problem whilst taking into account the different probabilities of scenarios, performs. A low $EVPI$ means that having full information compared to having uncertainty about the future, is not very valuable and not worth having unless it can be obtained very cheaply. On the other hand, if the $EVPI$ is high, then being able to obtain perfect information is valuable and will help to improve the average performance of the decision made, even if the information comes at a high price.

Value of Stochastic Solution

Even though, in theory, having access to perfect information would be very beneficial when making decisions with uncertainty, in practice, this usually isn't available. Finding the Wait-and-See Solution is also often a lot of work, or impossible to do if perfect information is not available for any price, especially as solving this problem produces a set of solutions, one for each scenario, instead of one solution that can be implemented for any possible scenario.

One popular solution is to instead solve a problem which is a lot simpler. This problem is obtained by replacing all the random variables (values with uncertainty) by their mean values. This is referred to as the **Expected Value Problem**, or mean value problem, and the value of the solution is represented by EV.

In the newsvendor problem with only 3 different scenarios, the mean values for the demand for each of the items is 80 for newspaper A, and 70 for newspaper B. This means that the expected value problem is the same as the deterministic problem that we solved at the beginning of the example, and therefore $EV = £330$

We can then define the **Expected Result of Using a EV Solution** (EEV) to be the value that measures how the solution for the mean values of the variables performs over the whole set of possible scenarios. In the newsvendor problem, the profits for each of the scenarios is £330 for both the high and medium demands, and £63 for low demand. This gives a value of:

$$\frac{1}{3} \times £330 + \frac{1}{3} \times £330 + \frac{1}{3} \times £63 = £241.$$

The **Value of the Stochastic Solution** (VSS) [2] can then be calculated for a maximisation problem as:

$$VSS = RP - EEV,$$

and this measures how good, or bad, using the EV solution performs, compared to the RP solution and can be thought of as the cost of ignoring the uncertainty when making a decision. For the newsvendor problem, this VSS is:

$$£248 - £241 = £7.$$

Like $EVPI$, the VSS must always be greater than or equal to zero, as the RP solution can always be the same as the EEV solution if this is the best. A low VSS is also preferred as this means that using the solution to the expected value problem for the full problem, is close to, or even just as good as the solution found by solving the RP. This is good as finding the EVV is a lot more computationally efficient compared to the RP solution, especially if there are a lot of scenarios that need to be considered.

Relationship Between $EVPI$ and VSS

As the newsvendor example has shown, the quantities $EVPI$ and VSS can be different, however they both show some similar behaviours. For both of these metrics, the values must be greater than or equal to zero, and we also want them to be as low as possible. If the $EVPI$ is low this means that we don't gain much performance if we are given perfect information, something that is rarely easily obtainable. If the VSS is low then it means that the EEV solution performs close to if not as well as the RP solution, and is preferable to use as it is a lot quicker and easier to compute, even for complicated problems.

Whilst these quantities can be different, there are some relationships that exist between the two measures of the effects of uncertainty. It can be shown that both of these values are bounded by $EEV - EV$ which is easily computable, and when $EV = EEV$, both the $EVPI$ and VSS are equal to 0. This can occur when the optimal solution is not dependent on the values of the random elements.

As well as it being possible for both of these quantities to be equal to zero, as they are not equivalent, it is also possible for each of them to be equal to zero, while the other is positive. This also shows that knowing the value of one does not help to inform anything about the value of the other.

$EVPI = 0$

The occurrence of $EVPI = 0$ while $VSS > 0$ is relatively common, and it is possible to predict its occurrence. This can happen when there are multiple solutions that have the same value as the optimal WS value which can lead to it being possible to find a RP solution such that $RP = WS$. If however, the chosen solution for the EV problem is different to this, this could lead to $EEV < RP$ in a maximisation problem. This would mean that $VSS > 0$.

$VSS = 0$

On the other hand, the occurrence of $VSS = 0$ while $EVPI > 0$ is not as easy to predict. If there is no single decision that is optimal across all the scenarios then the $EVPI$ is likely to be non-zero, however it is not as easy to define the conditions for which $VSS = 0$. It is thought however to be relatively frequent in practical problems [1].

For example, in the newsvendor example mentioned above, if the selling price for item B is decreased to £3.60, then the RP and EEV solutions both have the optimal decision being to buy 80 of item A and 70 of item B, meaning that $RP = EEV$ so $VSS = 0$. On the other hand, $EVPI$ is greater than zero as the profit is not as high as it would be if the correct number of items for the demand could be bought.

Predicting Values of $EVPI$ and VSS

Being able to estimate the values of $EVPI$ and VSS for a stochastic problem, without having to solve the problem its self would greatly simplify the process of stochastic programming. If we were able to know if these values for a problem were going to be high or low, we would only need to find the solution to the stochastic problem if one of these were large. If the values of the $EVPI$ and VSS are both small, then using the solution to the expected value problem would be favoured. The optimal value of this problem would not be much less than the solution to the stochastic program, and could be achieved a lot more easily.

Whilst in theory this would be useful, in practice it is difficult to do. It has been shown that there is **no general rule** to when high or low values occur, however there are some comparisons available for very specific problems [3].

It is usually felt that using stochastic programming is more valuable when there is **high amounts of randomness** in the problem. This can lead to the expectation that for a given problem, the values of $EVPI$ and VSS would increase when the variability of the uncertain variables increase, however it has been shown that this is not always the case.

Summary

We have seen that being able to predict the relative sizes of the values of $EVPI$ and VSS is useful when finding good decisions when making decisions under uncertainty. In theory, being able to calculate how useful additional information about the variables would be, or taking into account the best decisions for every possible scenario, would be beneficial to the problem solver. This would allow computational time to be saved when the benefit gained from the new information is small or complexity of the model is large. However, in practice, these values are not possible to estimate, and it can be difficult to predict patterns in the values, even if they have been calculated for similar problems.

References

[1] John Birge. The Value of the Stochastic Solution in Stochastic Linear Programs with Fixed Recourse. *Mathematical Programming*, 24, 12 1982.

[2] John R Birge and Francois Louveaux. *Introduction to Stochastic Programming*. Springer Science & Business Media, 2011.

[3] Hung-Po Chao. Exhaustible Resource Models: The Value of Information. *Operations Research*, 29(5):903–923, 1981.

[4] Julia L Higle. Stochastic Programming: Optimization When Uncertainty Matters. In *Emerging theory, methods, and applications*, pages 30–53. Informs, 2005.

[5] Alan J King and Stein W Wallace. *Modeling with Stochastic Programming*, volume 4. Springer, 2012.

[6] Alan S Manne. Waiting For the Breeder. *The Review of Economic Studies*, 41(5):47–65, 1974.