Classifying Merino Wool to Maximise Revenue

John F. Raffensperger, Ph.D.
Department of Management, Private Bag 4800
University of Canterbury
Christchurch, New Zealand
j.raffensperger@mang.canterbury.ac.nz

13 November 2000

Abstract

Wool prices depend on the wool's quality. The wool's quality depends on a variety of factors, a key factor being the diameter of the wool fibres. Thinner fibres feel more pleasant to the touch, and are more highly valued. Merino NZ's best and thinnest wool commands prices in the tens of thousands of dollars – *per kilogram*!

NZ PAC, a subsidiary of Merino NZ, classifies every fleece of Merino wool for each grower. Employees use sophisticated scanning devices to measure the wool's attributes. The most important attribute is the average fibre diameter. NZ PAC uses a mechanical sorting machine to sort each grower's wool into 16 bins. The problem is to choose attributes for each bin to maximise return for each grower. Usually, small quantities of wool with different diameters must be combined into a single lot to satisfy a minimum lot size constraint. The price of the lot depends on the average diameter of the wool in the lot. Fleeces can be combined into lots in different ways, resulting in different average diameters and quality, and therefore different market prices. This paper presents models to assign attributes to bins to maximise the revenue. The first model is a non-linear assignment problem, with blending constraints. The second is a large integer program with blending constraints. This work is being implemented by NZ PAC, and is expected to help increase revenue for growers.

1 Introduction

Merino New Zealand, Ltd., markets merino wool for New Zealand growers. Each grower owns and sells his own wool. One grower's production for a year is called a clip. The amount of money a grower receives for his clip depends on the wool's quality. Until recently, the growers themselves classified fleeces subjectively. In 1999, Merino NZ built a production facility called NZ Product Advancement Centre (NZ PAC), to classify wool with objective measurements. Each grower still owns and sells his own wool, but the wool can now be classified by NZ PAC.

NZ PAC uses sophisticated devices to classify wool based on each fleece's attributes. The most important attributes are the length and diameter of the wool fibres. The devices can accurately measure the length and diameter of each fleece. After measurement, the wool is sorted into lots and offered to the market for sale. Each lot must weigh at least 100 kilograms. The fleeces are then sorted mechanically. The sorting machine has only 16 bins, and a deeper sort would be difficult for production reasons. Also, since each grower owns and sells his own wool, each grower's wool must be kept separate.

When a lot is offered for sale, an independent laboratory tests the average diameter and length of the fibres. NZ PAC's measurements usually agree with the independent laboratory. The average diameter and average length of the fibres are the primary determinants of market price.

Currently, wool is classified by length into categories such as Long, Medium, and Short. Diameter is classified in 0.1 microns. Typical diameters are between 13 and 20 microns. The 13-micron wool can command prices in the tens of thousands of dollars per kilogram. At the other end, 20-micron wool may be sold for $10/kilogram or less.

Figure 1 shows a typical relationship between market price and diameter. It was necessary to produce a regression of market price to fibre diameter, since some market price data was missing, and because of differences in the timing of sales.

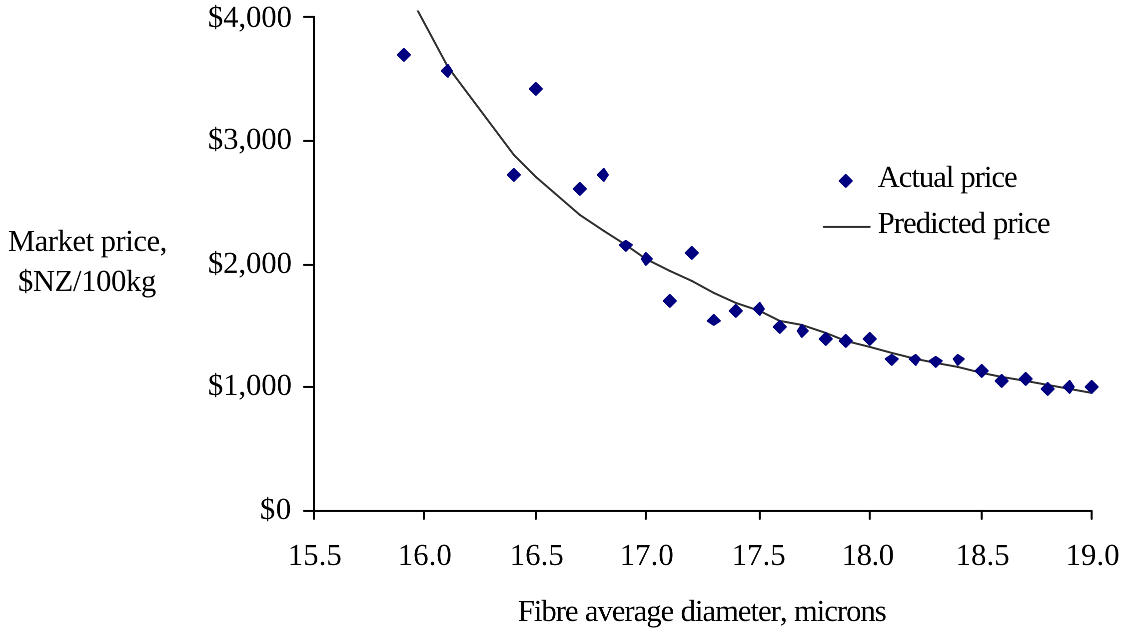

Figure 1. The relationship of market price to fleece fibre diameter. The predicted price function is price = $100{,}000/(21.91 - 21.60\, d + 1.36472\, d^2)$.

Fleeces of different diameters can be mixed to produce a lot of a given average diameter. A reasonable approach to maximising revenue is to sell every fleece at the highest price for its attributes. Since thinner fibres are more valuable, the grower can earn more money if the thin fibres sold in a lot with a smaller diameter. The ideal, then, would be to sell every fleece according to its true average diameter. However, buyers want each lot to weigh at least 100 kilograms. Furthermore, the NZ PAC sorting machine has only 16 bins. Also, each grower's wool must be kept separate. For these reasons, a lot is necessarily a blend of fleeces with different diameters.

The problem, then, is to assign fleeces to bins of the sorting machine, and to select attributes for bins, to maximise revenue for each grower. The weight by diameter and class for each grower's clip are forecasted based on the weight by diameter and class for the grower's clip from the previous year.

The literature has just one paper that described the use of optimisation to blend wool. This paper was in the 1986 Proceedings of NZ Operations Research Society, "Computer Blending of Wool," by Denis Maddever, Garth Carnaby, and Alec Ford [1]. They developed models to purchase wool at auction, and blend the wool for input to a carpet factory. They included constraints on attributes including colour, length, bulk, diameter, and vegetable matter. They used a dynamic model, with subscripts for month, to help minimise the holding cost of the wool, which suggests they must have used inventory variables, though this is not explicit. Under "Future Extensions," they reported work on non-linear functions related to colour and diameter. They used MINOS, a commercial non-linear math solver, and reported solver times of several hours. It appears their model would not fit NZ PAC's problem, because their model chooses wool for a single blend, while the NZ PAC problem is selecting wool for a sort into 16 bins. However, the work does seem to overlap. They also reported creating a database of market prices for input to their model.

This project has developed two models to assign attributes to bins on the PAC machine, to improve returns. The first model, "One-class bin solver," is a simplified non-linear model that assumes just one class of wool (such as Long, Medium, Short, or Colour). The second model is a mixed integer linear program that solves the problem for up to four classes.

The inputs for both models are market prices for each diameter and class (which can come from the regression model or actual market data), and the average and standard deviation of fleece diameter, and total weight, for each class (Long, Medium, Short, and Colour). A normal distribution was assumed, so quantities of wool could be assumed to be in categories, such as Long 15.0 to 15.1, Long 15.1 to 15.2, etc. Theoretically, negative or infinitely thick diameters could exist, but the formulas truncate these.

The desired output is a list of bin splits. A split is a range of fleece diameters for one class that is assigned to one bin. For example, one split might be Long 15.0 to 15.7, which means that all fleeces with those attributes would go into one bin.

2 Non-linear model formulation: One-class bin solver

"One-class bin solver" is a simplified bin assignment model to find bin splits for one class. The model uses the market price as calculated by the regression model, though in principle it could use any market price function (depending on the solver).

Indices

b = bin 1,…, 16.

i = fleece diameter category, usually 1, …, 30.

Inputs

Q_i = quantity (kilograms) of wool in diameter category i.

$p(d)$ = a price formula that depends on diameter d, found with regression.

L = minimum lot size, usually 100 kg.

d_i = average diameter of diameter category i.

Decision variables

d_b = average diameter in bin b.

y_{ib} = kilograms of diameter category i to put in bin b.

Model

$$\text{Maximise } \sum_i \sum_{b=1}^{16} y_{ib}\, p(\mathrm{d}_b)$$

subject to

$$\mathrm{d}_b = \sum_i (d_i y_{ib}) / \sum_i y_{ib} \quad \text{for each bin } b,$$

$$\sum_b y_{ib} = Q_i \text{ for each diameter category } i,$$

$$\sum_i y_{ib} \geq L \text{ for each bin } b.$$

y_{ib} = 0, for all i, b.

This model was solved with What'sBest, a commercial solver. This model did not perform well. It was easy to find solutions by hand that were better than solutions found the non-linear solver.

3 Mixed integer linear programming formulation: BinOptimiser

BinOptimiser is an optimisation model to select a class for each of the 16 bins on the PAC machine, to maximise expected return. Unlike the model "One-class bin solver," BinOptimiser solves the problem for four different classes.

The solved model gives assignments of fleeces (by category) to bins, and assignments of market price to bins. If the model's recommendations are followed, the actual average fibre diameter in the bin should match the model.

Indices

b = bin 1,…, 16.

i, j = fleece diameter category.

s, t = class, such as Long, Medium, Short, Colour, ordered 1,…5, from highest to lowest quality.

Inputs

d_{si} = the diameter for class s and diameter category i.

L_{si} = minimum lot size (usually 100 kg) for class s and diameter category i.

p_{si} = market price, \$/100kg, for class s and diameter category i.

Q_{si} = quantity (kilograms) of wool of class s and diameter category i.

Decision variables

w_{si} = 1 if a bin is assigned class s and diameter category i, else 0.

x_{sitj} = the per cent of wool with class s and diameter i that is assigned to a bin with class t and diameter j.

Model

1. Maximise $\sum_s \sum_i \sum_t \sum_j p_{tj} x_{sitj} Q_{si}$

subject to

2. $\sum_s \sum_i w_{si} \leq 16$, which limits the number of bins to 16.

3. $\sum_s \sum_i Q_{si} x_{sitj} \leq w_{tj} \sum_s \sum_i Q_{si}$, for each class t and diameter category j.

4. $\sum_s \sum_i D_{si} Q_{si} x_{sitj} \leq d_{tj} \sum_s \sum_i x_{sitj} Q_{si}$ for each class t and diameter category j.

5. $\sum_s \sum_i Q_{si} x_{sitj} \geq L_{tj} w_{tj}$ for each market class t and diameter category j.

6. $\sum_t \sum_j x_{sitj} \leq 1$, for each class s and diameter category i.

7. $x_{sitj} \geq 0$ for all s, i, t, j.

8. w_{tj} = 1 or 0 for all t, j.

Constraint 2 will be satisfied as an equality in an optimal solution, so it may be written as an equality without loss. Also, one class (called Tender) may be given its own bin, so the right-hand side of 16 could be 15.

Constraint set 3 requires that a fleece is assigned to a market category only if a bin has been opened for it.

Constraint set 4 requires that average diameter in the bin must be less than or equal to the assigned diameter for the bin. If price decreases with increasing diameter, this constraint should be satisfied as an equality in an optimal solution.

Constraint set 5 is the minimum lot size constraint.

Constraint set 6 requires that no more than 100% of the wool can be put into a bin. If prices are positive, this constraint will be satisfied as equality in an optimal solution.

Care is taken to define as few of the x_{sitj} variables as needed. For example, Short fleeces are never put in a bin of Long, so those variables are assumed to be zero, and the variables are not defined.

The typical number of variables is about 6,500, with 500 constraints, a small model. However, it is surprisingly hard to solve. Solution times are variable, but generally in tens of hours on a 277MHz Pentium II, though it finds a good solution in just a few minutes. Another solver such as CPlex may do better than What'sBest.

Additional constraints can be added to tighten the model:

9. $x_{sitj} \le w_{tj}$, for each class s, t, and diameter categories i, j. No wool goes in a bin if the bin is not open.

Adding these constraints will not hurt the true integer optimal solution, but it will reduce the (infeasible) upper bound from the continuous model. This reduces the number of branches that the solver has to search. Note that constraint set 9 dominates constraint set 3. While it is true that constraint set 9 tightens the model, each linear program in the branch and bound takes longer to solve. Constraint set 9 increases the number of constraints from 500 to about 8,000. Even though the search region is smaller, each search takes longer, so the effect is to make the whole search take longer.

It is easy to use BinOptimiser to find good feasible solutions quickly, since this is only an ordinary continuous linear program, with the w_{tj} variables fixed. By assigning attributes to the bins by hand, the model will find the splits almost instantly. For example, Bin 1 could be assigned to Long 17.4, Bin 2 to Long 18.8, and Bin 3 to Medium 17.4, etc. The model will quickly find which splits will best satisfy those assignments. This solves in just a few seconds, and gives a valid solution quickly.

4 The solution

It was expected that an optimal solution would have an uninterrupted range of splits in a single bin. For example, Long wool of diameter 17.0 to 17.5 microns would go in bin 1, Long wool of diameter 17.6 to 18.0 microns in bin 2, etc. This was not the case. For example, for fleeces that are Long 17.0-17.5, half may go in a Long 17.4 bin, and half in a Medium 17.2 bin. Another example: Long 18.0 fleeces and Long 18.5 fleeces may be assigned to a Long 18.3 bin; but fleeces Long 18.1 to Long 18.4 are assigned to a

Medium 18.2 bin. In effect, the splits were split! This may make the solution hard to implement.

The model sometimes put valuable fleeces in lower value bins. This is partly due to the 100-kg lot size constraint, but not completely. The small quantities do not justify a bin. The high-value fleece's small diameter slightly lowers the average diameter for the lower quality wool. So diameter and volume are both important. More bins should be used for classes with larger quantities, such as Medium. More revenue may obtained with Long in 2 bins and Medium in 6 bins, than the other way round, because there is more Medium wool.

Sensitivity analysis suggests that the 100-kg lot size constraint lowers the objective value only a little, less than 1% for the sample data. The 100-kg constraint does not hurt the model much, because the sample clip had a large total weight, so the 100-kg constraints are easy to satisfy. Even without the constraint, many bins have more than 100 kg.

It was expected that NZ PAC should be able to improve revenues somewhat by combining growers' clips to avoid the 100-kg constraint. However, combining wool from different growers may not always improve revenue. Sorting two clips separately is like having 32 bins for one combined clip, which results in a finer sort.

The value of combining clips depends on each clip's average and standard deviation of diameter. If two clips are large and have significantly different average diameters, they should be sorted separately. A small clip should probably be combined with another clip that has approximately the same average diameter. The 100-kg constraint will lower the value for smaller clips by a larger percentage than for large clips.

5 Conclusion

This project developed two models to help assign fleeces to bins. The most important model is BinOptimiser. This model uses What'sBest, a commercial solver, to find good assignments of wool to bins.

Acknowledgements

Thanks to Andrew Harris and Tony Hewitt at Merino NZ for their help and support in developing this model.

References

[1] Denis Maddever, Garth Carnaby, and Alec Ford, "Computer Blending of Wool," *1986 Proceedings of NZ Operations Research Society*, vol. 14., no. 2, pp. 168-171.