Modifying the Package Size Decisions Paper

Daniel Guetta

(Dated: August 25, 2009)

We suggest a modification to the model of Koenigsberg, Kohli and Montoya, examining how a firm might choose the package size and price for a product that deteriorates over time. The current model assumes that the firm has nothing to lose by reducing the size of its packages, reasoning that the high-consumption consumers can simply buy a larger number of packages. We modify this model to take into account the fact that there is an inconvenience associated with buying a large number of packages. We also model the fact that consumers with high demand are usually willing to pay less per item obtained (because they expect a bulk discount, for example). With the help of approximate numerical analyses, we find that the first modification drives the optimal package size *up* whereas the second drives it *down*.

I. INTRODUCTION, PLAN & PREDICTIONS

A. Introduction

This proposal builds on a model examining how a firm might choose the package size and price for a product that deteriorates over time, proposed in a paper by Koenigsberg, Kohli and Montoya.

The said model makes a number of simplifications. Together, these imply that the firm can only gain by reducing package size − by making package size smaller, the company gains smaller consumers (who would not have bought larger packages) and does not lose larger ones (who can buy a large number of smaller packages). The paper then comes to an intuitively understandable conclusion; if the company has nothing to lose by reducing package size, the smallest package size is obviously optimal. (See section II B for a more formal discussion).

Intuitively, this result is not true − ad absurdum, it would imply firms should produce minuscule packages. The aim of this proposal is to extend the current model to make it somewhat more realistic in that respect.

We do this by revising two key assumptions made in the original model

1. We adapt the model to reflect the fact that consumers experience additional dissutility from having to buy multiple packages. This could be due to transportation, storage or handling costs (eg: a caterer does not want to buy 5 000 small packets of ketchup and have to open them all one by one)

   We do this by assuming that if the price of packages is p, a consumer buying i packages actually pays an 'effective price'

   $$\boxed{P' = ip + ai^k} \tag{1}$$

   Where a and k characterise the dissutility of buying a large number of packages.

   Note that the existing model does go some way to modeling the dissutility of producing many packages on the *firm* side, by investigating concave cost functions. Our aim, however, is to model the *consumer* side of this issue.

2. We adapt the model to reflect the fact that willingness of pay is generally inversely correlated with demand; a consumer intending to buy 10 000 packets of ketchup will, usually, expect to pay less per packet than a consumer wanting to buy 2 packets.

   We do this by assuming a joint distribution of demand and willingness to pay amongst consumers, which takes into account this inverse correlation. We introduce a parameter $\beta \in [0, 1]$ to parametrise the strength of this inverse correlation. (See the next section for mathematical details).

B. Plan

In **Section 2**, I will give a concise summary of the current model. The construction of the demand function is particularly tricky and involves a number of nontrivial constraint, which I will examine in detail.

In **Section 3**, I will rigorously construct the new mathematical model by introducing the new parameters outlined in the previous section. A key feature of the new model is that it reduces to the current model for appropriate values of the parameters.

In **Section 4**, I will present the results of the numerical analyses I have run on the new model, and compare them to my predictions.

In **Section 5**, I will outline some ideas to take this modification forward, as well as ideas for other modifications of the existing model.

C. Predictions

At its most fundamental level, the main effect of the first modification (using the 'effective price' in equation 1) is to add an additional disincentive for the company to reduce package size. Indeed, by reducing package size, the company now forces high-consumption consumers to buy lots of small packages. This might drive the 'effective price' up so much that the consumer will no longer buy, thereby reducing the company's profit. I would therefore expect this modification to drive the package size *up*, as the company tries to reach an equilibrium between low- and high- consumption consumers.

I would, however, expect the second modification (inverse-correlation of willingness to pay and demand) to drive the package size *down*, because the firm now has a reason to favour the low-consumption consumers (they pay more) over the high-consumption ones.

II. CURRENT MODEL

A. Parameters

The parameters and variables used in the current model are as follows

- Each package has fixed useable life T, contains s units of the product, cost $c(s)$ to produce and has price p.
- The consumer values each unit of consumption at γ dollars. In other words, the user needs to get at least $R = 1/\gamma$ (the **reservation quantity**) units for each dollar spent. This is the 'willingness to pay'.
- A consumer has a usage rate θ, buys i packages and uses $i-1$ of them completely and a quantity f units of the product from the last one.

In the current model, we assume that the consumption rate θ and reservation quantity R are uniformly distributed:

- $\theta \in [\theta_{\min}, \theta_{\max}],\ f_\theta(\theta) = 1/(\theta_{\max} - \theta_{\min})$
- $R \in [R_{\min}, R_{\max}],\ f_R(R) = 1/(R_{\max} - R_{\min})$

We also assume that each consumer can buy a maximum of n packages. This parameter is introduced to allow an elegant mathematical statement of the model, but can eventually be allowed to tend to infinity to accurately reflect the fact consumers are general able to buy an infinite number of packages.

B. Qualitative Description

Before we proceed, it will be useful to give a qualitative description of the trade-off the company faces when it chooses the package size.

In choosing a package size, the company attempts to serve as many consumers as possible. In doing this, it has the following trade-off

- Choosing a **small** package size loses some high-consumption consumers, who are limited to buying n packages.

  For example, if a consumer wants 200 units, each package contains 10 units and $n = 10$, that consumer will only be able to buy 100 of the total 200 units it wants. The firm loses 100 units of business.

- Choosing a **large** package size loses some low-consumption consumers, for whom it is not profitable to buy such a large number of units.

  For example, consider a low-consumption consumer that needs 10 units and is only willing to pay \$ 1 per unit. Consider further that packages are \$ 20 and contain 20 units. By buying this package, the consumer pays \$ 2 per unit for each of the 10 units he uses. It is therefore not economically viable for this consumer to buy the package, and the firm loses 10 units for business.

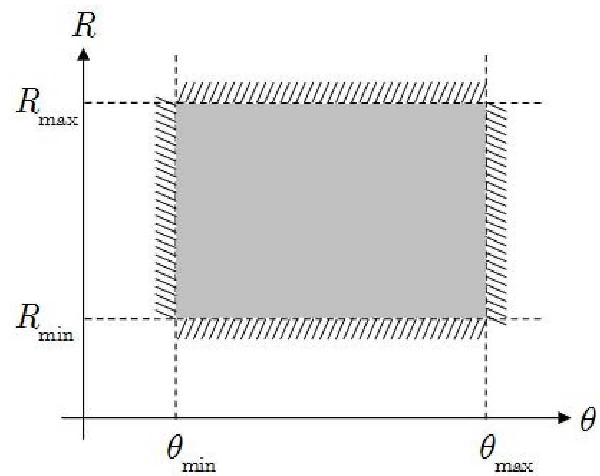

FIG. 1: The full set of consumers can be represented by the shaded rectangle. In this case, since the distributions of R and θ are uniform, area is proportional to number of consumers.

In reality, however, $n \to \infty$, and the first consideration disappears, because n can no longer limit the number of packages bought. This is what I was referring in the introduction when I said the company has no reason not to make the packages as small as possible. Thus, the package size tends to zero as n tends to infinity.

C. Demand Function

1. 'Consumer space'

Consumers are only heterogeneous in two parameters – $R \in (R_{\min}, R_{\max})$ and $\theta \in (\theta_{\min}, \theta_{\max})$.

We can therefore represent our customers on a two-dimensional 'consumer-space' diagram (Figure 1). Every consumer is included in the central rectangle.

In this case, since the distributions are uniform, area in 'consumer-space' is equal to number of consumers, but this need not be the case.

2. Subset of consumer space

As far as purchasing is concerned, the only heterogeneity in consumers is i – the number of packages each customer buys. Our aim in this section will be to find the subset of consumers (ie: the subset of the rectangle in Figure 1) that buy i products or more. This subset will be a function of *price* and *package size* – the two variables the firm can set.

Before we begin, consider that R is the number of units a consumer expects to obtain for each dollar he spends. Thus, if a consumer spends p, he expects to receive pR units.

We are now ready to find the said subset of consumer space by considering two obviously true facts about consumers

- **Consumers must obtain enough utility from the packages they use fully** – This means that the size of those packages, s, must be greater than the pR units the consumers expects. As such

$$\boxed{R \leq \frac{s}{p}} \tag{2}$$

- **Consumers must obtain enough utility from the last package, which they do not use fully** – This means that f, the amount used from the last package, must be greater than pR: $f \geq pR$.

  However, if the consumer buys i packages in total, and needs θT units, then $f = \theta T - (i-1)s$.

  This implies that

$$\boxed{\theta \geq \frac{pR + (i-1)s}{T}} \tag{3}$$

We now therefore have two constraints – one involving only R, and the other involving both R and θ.

How do these constraints partition the area in Figure 1? Let's consider each constraint

- **Equation 2** This constraint is simply an upper bound on R – a horizontal line in our diagram.

  It is clear that if our horizontal line is *above* the $R_{\max}$ line, we're not changing the demand or extending the area. Thus, without loss of generality, we can assume that

$$\boxed{\frac{s}{p} \leq R_{\max}} \tag{4}$$

- **Equation 3** This is a more complicated linear constraint. We only know that it has a positive gradient, and a positive intercept, but it could cross our consumer space in one of the four ways in Figure 2.

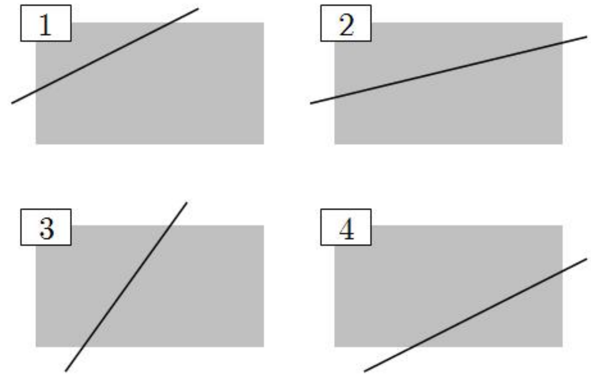

FIG. 2: The second constraint must have a positive gradient, and can therefore only cross our consumer space in the four ways above

We narrow these four choices down to one by making two crucial assumptions

$$\boxed{\min(R) = R_{\min} \geq \frac{\theta_{\min} T - s(i-1)}{p}} \tag{5}$$

$$\left(\boxed{\max(R) = \frac{s}{p} \leq \frac{\theta_{\max} T - s(i-1)}{p}} \right) \tag{6}$$

The right-hand-side of these two constraints are the least and most units-per-dollar that can be obtained from the last package. The two constraints therefore say the following:

- **Constraint 5** – there will be a subset of consumers whose consumption is so low that they will not buy i packages, regardless how low their reservation quantity is.
- **Constraint 6** – there will *not* be a subset of consumers whose reservation quantity is so high that they will not buy i packages, regardless how high their consumption rate is.

  This constraint is included in parentheses because it is somewhat redundant. High-consumption consumers are extremely valuable to the company since they buy many packages. The optimisation step, therefore, will always lead to a pricing structure that takes advantage of *all* its high consumption consumers. (Many thanks to OK, RK and RM for bringing this fact to my attention).

Together, these two constraints imply that *option 3* is the correct choice in Figure 2.

In summary, therefore, we have found that the region of consumer space occupied by consumers who are able to buy at least i packages is that in Figure 3.

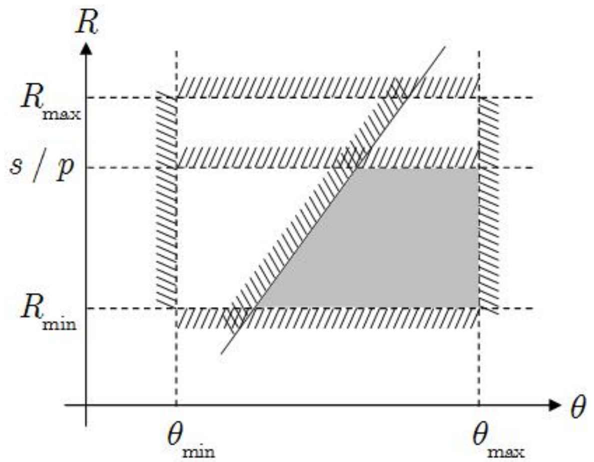

FIG. 3: The region of consumer space occupied by consumers buying i packages or more.

The area of this region of consumer space, which we denote D_i is given by

$$D_i = \int_{R_{\min}}^{s/p} \int_{\frac{pR+s(i-1)}{T}}^{\theta_{\max}} f_\theta(\theta) f_R(R) \mathrm{d}\theta \mathrm{d}R \quad (7)$$

Since the distributions are uniform, we can evaluate this integral by simple geometry

$$D_i = \frac{1}{(\theta_{\max}-\theta_{\min})(R_{\max}-R_{\min})}\left(\frac{s}{p}-R_{\min}\right) \times \left[\left(\theta_{\max}-\frac{is}{T}\right)+\frac{1}{2}\left(\frac{s-pR_{\min}}{T}\right)\right] \quad (8)$$

3. Finding the demand

We are, finally, in a position to find the demand, D.

So far, we know the number of consumers who buy at least i units. Now, imagine n is the maximum number of units each consumer can purchase, and consider the sum

$$D = \sum_{i=1}^{n} D_i$$

Note that

- Consumers that are willing to buy only one item are counted *once* in this sum.
- Consumers that are willing to buy two items are counted *twice* in this sum – once by D_1 (since they're willing to buy at least 1 item) and once by D_2 (since they're willing to buy at least 2 items)...

This continues up to n, where each consumer willing to buy n items is counted n times.

Thus, this sum provides a correct expression for the demand.

Carrying out the sum on our expression for D_i above (Equation 8), we obtain

$$D = \frac{1}{(\theta_{\max}-\theta_{\min})(R_{\max}-R_{\min})}\left(\frac{s}{p}-R_{\min}\right) \times \left[\theta_{\max}-\frac{pR_{\min}+sn}{2T}\right] n \quad (9)$$

This is the demand function.

D. Optimal decisions

The firm has two choices to make – package size, and price. We will consider a firm whose aim is

$$\max_{s,p} \pi$$

Where π is the profit, given by

$$\pi = [p - c(s)]\, D$$

We assume that the cost per package is a linear function of package size

$$c(s) = \alpha s$$

It turns out that

- The company should only produce if

$$\boxed{\alpha < \frac{1}{R_{\min}}} \quad (10)$$

  This is because $1/R_{\min}$ is the minimum any consumer is willing to pay for one extra unit of product. If it costs the firm more than that to add an extra unit to a package, then there's no point producing anything.

- The first order conditions $\partial\pi/\partial s = 0$ and $\partial\pi/\partial p = 0$ give

$$s^* = \frac{(1-k)\theta_{\max}T}{n} \quad (11)$$

$$p^* = \frac{k\theta_{\max}T}{R_{\min}} \quad (12)$$

  where

$$k = \left[\sqrt{(n+1)\left(\frac{n}{R_{\min}\alpha}+1\right)}\right]^{-1}$$

It is worth noting, at this point, that these results place some highly nontrivial constraints on n, because these solutions must satisfy the constraints in equations 5 and 4.

Feeding the optimal solutions in equations 11 and 12 into these constraints, we find that we require

$$n > \frac{R_{\min}\theta_{\max}}{\theta_{\min}R_{\max}} - \frac{R_{\min}}{R_{\max}} \quad (13)$$

and

$$n_{\min} < n < n_{\max} \quad (14)$$

where $n_{\max}$ and $n_{\min}$ are absolutely ghastly expressions, available from the author on request.

This seems to be in stark contradiction with the idea that we can let n tend to infinity to reflect the fact consumers can buy an infinite number of package. In practice, however, there is no contradiction, because once the $n \approx 30$ mark is passed, the model behaves, to all intents and purposes, as if $n \to \infty$.

These constraints, however, do make the numerical analysis of the model much more complicated, especially when the modifications in the next section are implemented.

E. Analysing the model

It seems clear, from equation 11, that as $n \to \infty$, the optimal package size tends to 0. This is in line with the predictions make in section II B, for the reasons described there. See the original paper for a more detailed analysis of the original model.

III. MODIFYING THE MODEL

A. Distribution of R and θ

The first modification involves the distribution of R (the willingness to pay) and θ (the consumer demand). We seek a distribution that reflects the fact these quantities are inversely correlated; in other words, those that buy more are willing to pay less for each item.

The joint distribution we use to achieve this is the following:

$$f(\theta, R) = N\left[1 + \beta\left\{\left(\frac{R - R_{\min}}{R_{\max} - R_{\min}} - \frac{\theta - \theta_{\min}}{\theta_{\max} - \theta_{\min}}\right)^2 - 1\right\}\right] \quad (15)$$

Where $\beta \in [0, 1]$ characterises the strength of the inverse correlation, and

$$N = \frac{6}{(6 - 5\beta)(\theta_{\min} - \theta_{\max})(R_{\min} - R_{\max})} \quad (16)$$

is a normalisation constant, to ensure that

$$\int_{\theta_{\min}}^{\theta_{\max}} \int_{R_{\min}}^{R_{\max}} f(R, \theta) \mathrm{d}R \mathrm{d}\theta = 1$$

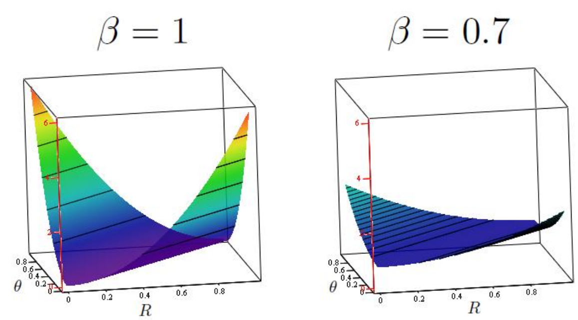

FIG. 4: The joint distribution, $f_\theta(\theta, R) f_R(\theta, R)$, for the cases $\beta = 1$ and $\beta = 0.7$, with $\theta, R \in (0, 1)$.

As one might expect, when $\beta = 0$, N and the distribution take on their standard, 'uniform distribution' form. The new model reduces to the old model for $\beta = 0$.

A three-dimensional plot of this joint distribution is provided in Figure 4.

B. Demand Function

In our new model, the consumer now pays an 'effective price' (equation 1) which includes a bulk utility, so we must review the subset of consumer space that will buy i products or more by re-formulating our 'two true facts':

- **Consumers must obtain enough utility from the packages they use fully** – Our consumer use $i-1$ packages fully, from which they obtain $(i-1)s$ units and pay $(i-1)p + a(i-1)^k$. They will only be willing to pay that amount if they get at least $\left[(i-1)p + a(i-1)^k\right] R$ items, and so

$$(i-1)s \geq \left[(i-1)p + a(i-1)^k\right] R$$

which implies that

$$\boxed{R \leq R_{H,i} = \frac{s}{p + a(i-1)^{k-1}}} \quad (17)$$

- **Consumers must obtain enough utility from the last package, which they do not use fully** – The price our consumers pay for their last package is $p + a\left(i^k - (i-1)^k\right)$. We must therefore have that $f \geq \left[p + a\left(i^k - (i-1)^k\right)\right] R$.

We have seen, however, that for a consumer who buys i packages and needs θT units, $f = \theta T - (i-1)s$. As such

$$\boxed{\theta \geq \theta_{L,i} = \frac{\left[p + a\left(i^k - (i-1)^k\right)\right] R + (i-1)s}{T}} \quad (18)$$

These two conditions are analogous to equations 2 and 3.

We must now derive constraints analogous to those in equations 4, 5, 6 and 10. Let's consider each in turn

- The equivalent form of equation 4 in this case is

$$\frac{s}{p+a\left(i-1\right)^{k-1}} \le R_{\max}$$

The largest possible value of the LHS occurs when $i = 1$, in which case

$$\boxed{s \le R_{\max} p} \tag{19}$$

This is our first constraint.

- The equivalent form of equation 5 in this case is

$$\min\left(R\right) = R_{\min} \ge \frac{\theta_{\min} T - \left(i-1\right) s}{p + a\left(i^k - \left(i-1\right)^k\right)}$$

The largest possible value of the LHS occurs when $i = 1$, in which case

$$p \ge \frac{\theta_{\min} T}{R_{\min}} - a \tag{20}$$

- The equivalent form of equation 6 in this case is

$$\max\left(R\right) = \frac{s}{p+a\left(i-1\right)^{k-1}} \le \frac{\theta_{\max} T - \left(i-1\right) s}{p + a\left(i^k - \left(i-1\right)^k\right)}$$

Unfortunately, I can't think of any efficient way of simplifying this constraint. A bit of re-arranging gives

$$\boxed{s\left\{ai^k + ip\right\} \le \theta_{\max} T\left\{p + a\left(i-1\right)^{k-1}\right\}} \tag{21}$$

This constraint is problematic, because it still involves i – it must therefore be satisfied for all values of i from 1 to n. We will have more to say about this in our conclusion.

- Finally, the constraint in equation 10 must apply in this case as well, for the same reasons.

$$\boxed{\alpha R_{\min} < 1} \tag{22}$$

The expression for the demand, D, in this modified model is then

$$D = \sum_{i=1}^{n} \int_{R_{\min}}^{R_{H,i}} \int_{\theta_{L,i}}^{\theta_{\max}} f(\theta, R) \mathrm{d}\theta \mathrm{d}R$$

Once again, this new model reduces to the old model when $a = 0$.

My preliminary investigations seemed to indicate that an expression for the demand could be obtained in closed form. The symbolic manipulator, however, indicated that the expression contained more terms than could be displayed. Given the complexity of the expression, I opted to begin with a numerical study of the system.

C. Qualitative Description

Referring back to section II B, in which we give a qualitative description of the current model, we find that the main effect of our modifications are the following

- The first modification favours the lower-consumption consumers, who are now willing to pay more. In other words, it adds an even greater disincentive for the firm to increase its package sizes.

- The second modification makes it very unlikely that high-consumption consumers will be willing to meet their demand by buying lots of small packages. This adds a disincentive for the firm to reduce package sizes.

  This point is crucial, because under the old model, once n tends to infinity, there is no longer any reason the company would no reduce package sizes for ever. This part of the new model ensures a disincentive remains for *all* values of n.

IV. NUMERICAL ANALYSIS

A. Notes

I begun the my numerical study naively – by simply feeding in the expressions for the demand and profit into Mathematica, and then using a simple multivariate optimisation function to find the combination of s and p yielding the highest profit for a given value of n.

A number of issues immediately arose, even when solving the original model:

1. The results were often completely nonsensical because they returned values inconsistent with the constraints in equations 20, 19 and 22. Some solutions even involved negative profits, package sizes and prices.

2. Mathematica was unable to find an appropriate starting point for its optimisation algorithm.

3. There are many parameters in the model that have to be chosen before the simulation is run ($R_{\min}$, $R_{\max}$, etc. . .). However, the large number of constraints in the model make it difficult to choose parameters that will eventually lead to a solution satisfying this constraints.

   For the original model, analytical solutions are available, and so we were able to derive analytic constraints for these parameters in equations 13 and 14. (And even with the help of these expressions, finding appropriate parameters is difficult).

   For the new model, not even those analytical expressions are available – this makes it impossible to know whether the values we choose for the original parameters will lead to a sensible solution.

I resolved these issues as follows

1. The first by carrying out a *constrained* optimisation using the said constraints.
2. The second by providing Mathematica with a starting point – namely, the *analytical* solutions to the original model, in Equations 12 and 11.
3. The third by trial and error.

These measures allowed me to numerically reproduce the analytical results obtained in equations 12 and 11. Unfortunately, there were some additional complications associated with the numerical analysis of the *new* model with non-trivial a, k and β. Not only did the algorithm take several hours to run, but it rarely converged to a sensible solution, even after 6000 iterations. The resulting graphs were rather ill-formed, and completely useless in terms of obtaining any useful information.

An alternative approach I experimented with was to produce three-dimensional plots of profit as a function of s and p, and use them to visually find the maximum profit. This approach proved to be very difficult; Mathematica took a very long time indeed to produce the plots, and it was somewhat difficult to take into accounts the constraints from point 1 above.

Given more time, I would have taken both these approaches further.

- I would have tried to remedy to the first by running the Mathematica workbook on a more powerful computer, and with more iterations.
- I would have tried to rememdy to the second by defining a three-dimensional function *including* the relevant constraints using Mathematica's 'if' function. Effectively, I would ask Mathematica to return 0 if (s, p) was outside the bounds defined by the constraints, and to return the value of the profit otherwise. This should have made the plot much easier to interpret.

As things were, and with time running out, I opted for an approximate solution. Instead of optimising the profit over *both* s and p, I only optimised over s, the package size. For p, I simply used the analytical optimal value from the *original* model in equation 12 (I used different values for each value of n – the optimal value in each case). This approach allowed me to gain some qualitative insight into how the model performs, but it significantly lacks rigour. Any further investigation of this model would require a more thorough investigation along the lines of the suggestions above.

B. Results

I ran my numerical simulations on a model with the following parameter values

- $R_{\min} = 100$ and $R_{\max} = 500$
- $\theta_{\min} = 1$ and $\theta_{\max} = 500$
- $\alpha = 0.01$
- $T = 1$

In each case, I assumed the price of each package was the optimal price given in equation 12, for the old model. I then used Mathematica's univariate optimisation function to find the optimal corresponding package size s. I ran my simulations for five values of n – 110, 120, 130, 140 and 150.

I ran three sets of simulations:

First set – $\beta = 0$, $a = 0.001$ and $k \in (0, 3.1]$.

As a reminder – k is the power to which the number of packages is raised in the extra dissutility incurred when buying many packages. So if $k = 2$ and i packages are bought, the dissutility increases as i^2.

It stands to reason that the larger the power, the steeper the dissutility's rise with number of packages, and this simulation investigates how this affects results.

Results in figure 5.

Second set – $a = 0$, $k = 1$ and $\beta \in [0, 1]$.

This set investigates how the inverse correlation of R and θ affects results.

Results in figure 6.

Third set – $\beta = 0$, $k = 2$ and $a \in [0.001, 0.1]$.

As a reminder – a is the weight of the additional dissutility incurred when buying many packages. So if $a = 2$ and i packages are bought, the dissutility increases as $2i^k$

This set investigates how the weight of the extra dissutility affects the results.

Results in figure 7.

In each case where $a \neq 0$, I chose $a = 0.001$, to make the bulk dissutility incurred significantly smaller than the price of the package, which was in the range $p \approx 0.01$.

In each case, I plotted optimal package size against n. As explained above, the package price used for these simulations was *not* necessarily the optimal price – I therefore did not produce any plots involving profits or price.

C. Conclusions

I now analyse each of the simulation sets, in light of my predictions.

The reader is direction to sections II B and III C, which qualitatively describe the old and new model. The results will be interpreted in light of these descriptions.

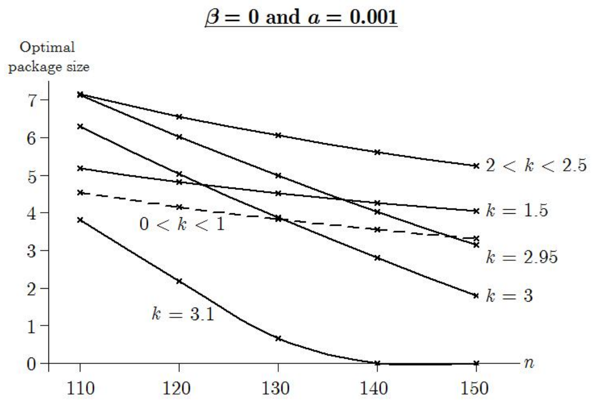

FIG. 5: Optimal package size for various values of n and k. These simulations were all carried out with $\beta = 0$ and $a = 0.001$. In some cases, several values of k produced very similar looking plots, and these are therefore labelled with a *range* of values. *The plot produced by the original model is indicated by a dotted line.*

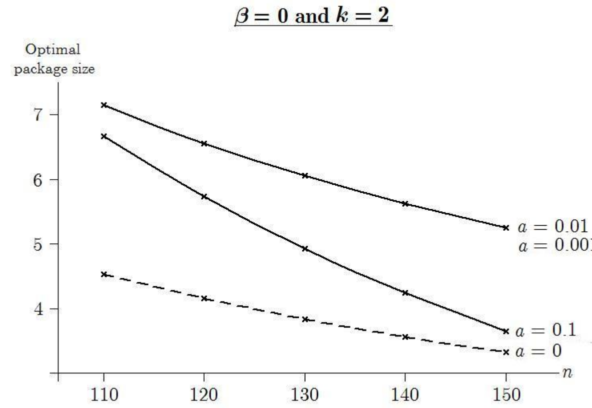

FIG. 7: Optimal package size for various values of n and a. These simulations were all carried out with $\beta = 0$ and $k = 2$. *The plot produced by the original model is indicated by a dotted line.*

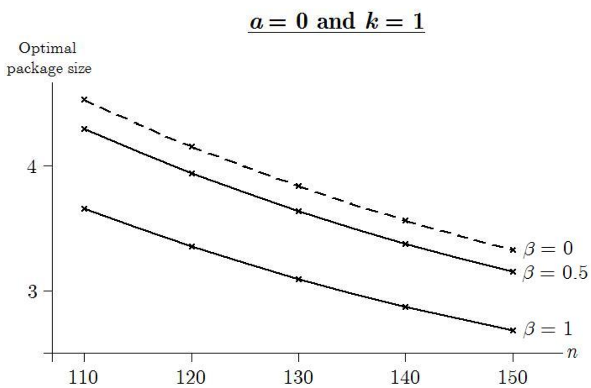

FIG. 6: Optimal package size for various values of n and β. These simulations were all carried out with $a = 0$ and $k = 1$. *The plot produced by the original model is indicated by a dotted line.*

1. *Set 1*

There are three clear region in this plot, in which changing k affects results very differently

- For $k \leq 1$, the model behaves very similarly to the old model – the additional dissutility does not affect the optimal package size in any way.

  This result is sensible. If $k \leq 1$, then the additional dissutility per packages *decreases* as more and more packages are bought. (For example, if $k = 0.5$ and $a = 1$, then the dissutility-per-package for 25 packages is $\left(25^{0}.5\right)/25 = 0.2$, whereas the dissutility-per-package for 100 packages is $\left(100^{0}.5\right)/100 = 0.1$). This simply means that by reducing package size, the company does *not* lose higher-consumption consumers (indeed, it gives then an additional advantage – they now need to buy more packages to fulfil their demand, and so pay a lower per-package dissutility).

  We are therefore back to the situation in the old model, where the only disincentive the company has to increase package size is the limit imposed by n. Thus, the results are very similar.

- For $k \in [1, 2.5]$, two interesting features emerge
  - The optimal package size increases with k.
  - As n increases, the optimal package size seems to fall faster than in the original model.

  The first result is easily understood. If $k \in [1, 2.5]$, the dissutility-per-package *increases* with the number of packages sold. This leads to larger package sizes, for reasons we have already discussed at length in our predictions.

  I cannot, however, think of any explanation for the second observation. It is perhaps due to the fact that the true optimal price was not used in this simulation.

- For $k > 2.5$, two interesting features emerge, once again
  - The optimal package size plummets rapidly
  - Initially, for k close to 2.5, this only happens for larger values of n. At larger values of k, this occurs for *all* values of n.

  In this case, the dissutility experienced by consumers having to buy a large number of packages increases dramatically. This helps us explain both these observations:
  - The dissutility experienced by high-consumption consumers is now so large

that the extent to which the firm would have to increase package size to serve these consumers adequately would mean the firm would have to lose too many of its low-consumption consumers. As a result, the firm ignores its high consumption consumers, and drops package size dramatically to favour the low-consumption consumers.

– Initially, this only occurs at high n, because at low n, the high-consumption consumers are already restricted in the number of packages they are able to buy by n's low value. For so few packages, this extra dissutility doesn't affect them much. For much higher n, however, the massive dissutility starts to kick in.

2. Set 2

The second simulation set investigated the effect of inversely-correlating R and θ on the optimal package size.

The results clearly indicate that the larger the inverse correlation (the larger β), the smaller the optimal package size.

This is, once again, sensible. Increasing β implies that higher-consumption consumers are less inclined to pay large amounts of money for each package. Thus, as β increases, the firm has more to gain by favouring low-consumption consumers, who pay more.

The firm therefore reduces package size, even though this means it will lose some of its higher-consumption consumers, due to the n-package limit.

3. Set 3

The third simulation set investigated the effect of changing the weight of the extra dissutility term (a).

The results were somewhat confusing. For very low a, the model behaves like the old model, as might be expected. As a is increased, the optimal package size increases dramatically. This can be explained using similar arguments to those made for the first simulation set.

The confusing behaviour occurs when a is increased even further, to $a \approx 0.1$. Under this regime, we find old-model-behaviour for large n (small optimal package sizes) but new-model-behaviour for small n (large optimal package sizes).

We attempt to explain these observations as follows:

- For low n, the very high consumption consumers are excluded in both the old and new model by n. The only consumers that remain are the reasonably-high-consumption consumers. Even for these, however, a is so high that the additional dissutility does have some effect, and it serves to push the optimal package size up – as in the new model.
- For high n, those consumers that make the difference are the very-high consumption consumers. However, in this case, a is so high that the additional dissutility to those consumers is so large than the firm simply ignores them – in the same way the company ignore high-utility consumers when $k > 2.5$ in the first simulation set. Thus, all that remains are the low-consumption consumers, and the optimal strategy is to choose small package sizes, as in the old model.

It is also possible, of course, that this anomaly arises because the price used in these simulations is not the optimal price for the new model, but an approximation to it.

V. SUGGESTIONS

In conclusion, the model behaved more or less as expected, and with some additional work, it promises to provide a valuable improvement to the existing model.

I would offer the following suggestions to anyone wishing to take this model further

- Implement one or more of the suggestions in section IV A, to carry out an accurate multivariate optimisation and find the optimal price *and* package size.
- Pay more attention to the constraint in equation 21. When analysing the old model, we mentioned that this constraint need not be taken into account because it would automatically be satisfied when the firm made its optimal decision.

  It is not clear that this is also true for the new model. In fact, preliminary investigation into this issue seems to indicate that it is *not* true for the new model. In an adequate treatment of the new model, this constraint will need to be considered explicitly.
- Extend the model to be valid even outside the bounds defined by the constraints in equations 5, 6 and 4. As we mentioned above, these constraints are very stringent indeed, and considerably restrict the range of applicability of the model. By changing the form of the integral in equation 7 and taking into account the other three possibilities in figure 2 for the topology of the consumers space, it should be possible to find an expression for the demand even outside the said bounds.
- An additional limitation of the current model, which we did not treat in this paper, is that the firm is only allowed to sell *one* kind of package with *one* size. This is clearly completely unrealistic – most firms offer a number of options, and are therefore able to cater to both high- and low- consumption consumers at the same time.

If this suggestion were to be implemented, the use of a nonzero β and of inversely correlated R and θ distributions would be particularly interesting, because a firm producing two different package sizes would be able to adapt the price of each package to the relevant consumer. Thus, one would expect the large package size to have a low price per unit, and the small package size to have a high price per unit.

Acknowledgements

This investigation was carried out in June, July and August 2009, during my time as a visiting scholar at the marketing department of the Columbia business school. I am most grateful to Prof Olivier Toubia for inviting me to the department and to everyone there for their warm welcome.

I would especially like to thank Professors Oded Koenigsberg, Jacob Goldenberg and Olivier Toubia for being so patient, and for always having seemingly limitless amounts of time to answer my questions, however trivial. Many thanks also to Professors Rajeev Kohli and Ricardo Montoya for their valuable advice in understanding the paper on which I based this study.