



Institute of Actuaries of Australia

Determining an Optimal Distribution Strategy

Prepared by Mitchell Prevett and Greg Schneider

Presented to the Institute of Actuaries of Australia
5th Financial Services Forum
13 – 14 May 2010
Sydney

*This paper has been prepared for the Institute of Actuaries of Australia's (Institute) 5th Financial Services Forum.
The Institute Council wishes it to be understood that opinions put forward herein are not necessarily those of the Institute and the
Council is not responsible for those opinions.*

© The Quantum Group

The Institute will ensure that all reproductions of the paper acknowledge the
Author/s as the author/s, and include the above copyright statement:

The Institute of Actuaries of Australia
Level 7 Challis House 4 Martin Place
Sydney NSW Australia 2000
Telephone: +61 2 9233 3466 Facsimile: +61 2 9233 3446
Email: actuaries@actuaries.asn.au Website: www.actuaries.asn.au

Determining an Optimal Distribution Strategy

Table of Contents

1	Introduction	3
2	Key Drivers of Sales	3
3	Current Approaches in Determining Distribution Strategy.....	5
3.1	Rules of Thumb	5
3.2	Average Lifetime Value of a Salesperson	6
3.3	Issues with these Approaches	6
4	Conceptual Distribution Model	7
5	Sales Force Effectiveness.....	9
5.1	Measuring the True Impact of a Salesperson.....	10
5.2	Which Types of Sales People Perform Better?.....	12
5.3	Performance Benchmarking	14
6	Optimal Sales Force Strategy	14
6.1	Recruitment	15
6.2	Optimal Sales Force Allocation and Deployment.....	15
6.3	Optimal Channel Strategy	16
6.4	Practical Considerations	18
7	Conclusion	19
8	Bibliography.....	19

1 Introduction

Financial services companies have been applying a greater level of analysis to many aspects of their businesses. This now encompasses pricing, customer analytics, media and marketing decision making. For many organisations distribution is the primary driver of performance however very few are applying actuarial thinking and predictive models to solve pressing problems in this area. The purpose of this paper is to shed some light on recent advances in actuarial methodologies being applied to distribution strategy.

In section 2, we discuss the key drivers of sales and demonstrate how important it is to determine the optimal investment in each driver, namely Product, Price, Place (distribution) and Promotion. Following this, in section 3, we take a look at some of the current approaches used to determine distribution strategy and describe what we see as the major issues with these approaches. In section 4, we present a distribution model which will overcome these issues and appropriately allow for each of the main value drivers. In section 5, we describe in detail one of the key model components, the sales performance model, to demonstrate how this can be used to understand the individual drivers of performance. Finally, in section 6, we demonstrate how to use the model to derive an optimal distribution strategy.

2 Key Drivers of Sales

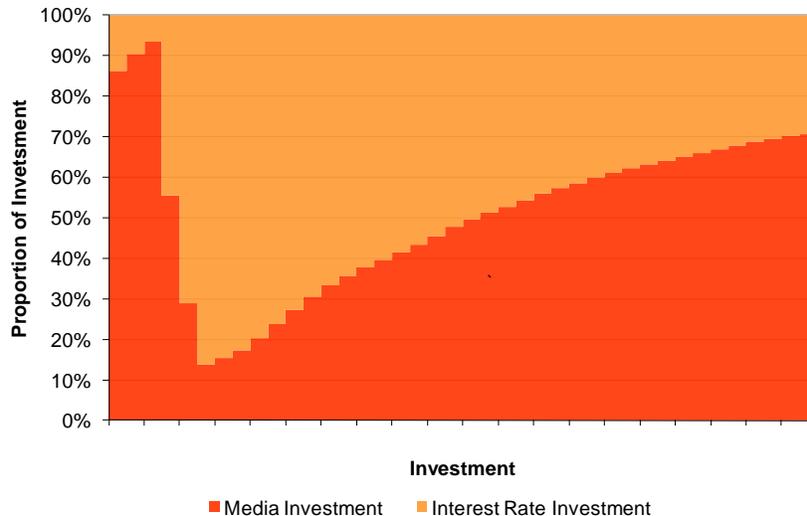
Product, Price, Place (distribution) and Promotion, the famous 4 P's of marketing, have long been regarded as the key drivers of sales. For each organisation the importance of these drivers will vary. The fundamental question any organisation needs to answer is how much it should invest in each of the P's based on the importance the factor has for the organisation. This obviously needs to be considered in the light of all the external factors that impact its industry – for example, the state of the economy, seasonality, competitor activity and so forth.

At the heart of answering this question is determining the pure effect of each of the drivers of sales, holding the other drivers constant. Given that these factors impact sales simultaneously and not in isolation, multivariate analysis is required to uncover these effects (discussed in section 4). Once these pure effects are understood a framework can be built that allows trade-offs between the 4 P's of marketing, leading to optimised investment decisions.

This is best illustrated with a case example from the retail banking industry where the key question for the bank is how much it should invest in Price (competitiveness) and how much in Promotion.

For this example the cost of investing in Price is the reduced Net Interest Margin (NIM – a measure of margin used for many retail banking product categories) – we will call this the interest rate investment. Figure 1 below highlights the results of the optimisation process. We have removed the size of the investment (Price and Promotion) in order to protect the company's confidential results.

Figure 1 – Optimal Mix of Media and Interest Rate Investments



At low levels of investment, the optimal outcome is found by having a higher weight in Promotional investment. This is due to the relatively high cost of improving Price at low levels of investment. Once past this peak, the best outcome switches around to Price related investment with a gradual increasing share in Promotional investment. At the highest levels, the majority is invested in media.

The conclusions illustrated are specific to the example, however it demonstrates that the level of investment in Price and Promotion varies considerably depending on the total investment budget. Determining this optimal investment mix requires robust and detailed analysis. A crude approach adopted by many organisations might be to split the investment based on some fixed proportion (say 50/50) which, at most levels of total investment, would be significantly sub-optimal.

These kinds of analyses are currently relatively rare although there are studies which explain the approach and benefits. “An Actuarial Approach to Optimising the Trade-Off Between Media and Price Promotions” prepared by Adam Driussi, Caroline Stevenson and Tony Davis (September 2007) will provide further details.

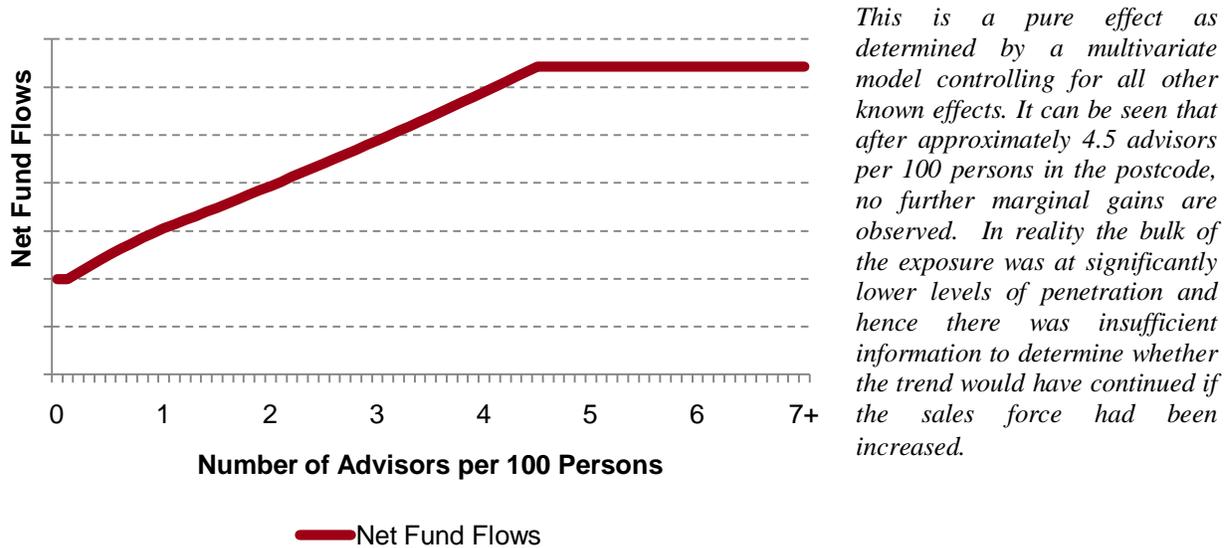
Another tension experienced by most businesses, is how to determine the investment in marketing versus distribution. Many companies often have very detailed strategies on marketing and distribution but the initial question of how important each is to the organisation is considered very infrequently if at all.

Consider another case example this time from the funds management industry. The organisation historically had a significant marketing budget to support its retail net fund flows (sales performance) from direct and advisor sourced business. It wanted to know how effective its media investments had been in generating net flows taking into account, among other things:

- The state of the economy and the overall funds management sector;
- The performance of its funds;
- Seasonality and socio-demographic factors ; and
- Advisor penetration.

One of the surprising findings of the econometric study for the organisation was the significant impact advisor penetration had on net flows, even though this was not the original objective of the study. Figure 2 below demonstrates the pure effect that the number of advisers had on fund flows and shows that as the number of advisors in each region increased, so too did the relative net flows.

Figure 2 – Fund Flows vs. Advisor Penetration



Although this finding is inherently obvious, it highlighted to the organisation just how clear and strong the relationship was to sales performance (net fund flows). The consequence of the analysis was that the organisation focused significantly more on distribution and less on marketing for an ultimately higher overall return. In reality, the importance of distribution and marketing varies considerably by business and industry and this result is by no means consistent across businesses, let alone across industries.

3 Current Approaches in Determining Distribution Strategy

Given that an organisation has actively or passively decided to focus its resources on distribution, how have they historically determined the required strategy in terms of the size and location of sales resources (sales force, branches or stores)? Generally, organisations treat these two questions independently, although they are intrinsically linked.

3.1 Rules of Thumb

It is common for businesses to determine the number of additional sales resources (sales staff, branches or stores) required based on analysis of which areas have the highest perceived need. The approach may be either top down, in which case the total number of additional resources is fixed and are allocated to the most needy areas, or bottom up, in which case the most needy areas are allocated resources to determined

the total requirement. Typically, the most needy areas are determined based on rules of thumb that are assumed to influence overall performance of the distribution network. These may include:

- Customer density in the region or foot traffic in the branch/store;
- Value and size of the customer base in the region or branch/store;
- The socio-demographic profile of the region;
- The number of customers per salesperson already servicing that area; and
- The number of branches per salesperson in that area.

These rules of thumb are usually not suitably tested against past performance to determine the weight each should have in the allocation decision. In the absence of analytically generated answers, “common sense models” or inherited weights are typically used.

3.2 Average Lifetime Value of a Salesperson

The average lifetime value or NPV for a sales representative is usually determined based on the average individual. That person’s productivity is found by an appropriately weighted analysis of the existing sales force across dimensions like tenure in the position, level of experience, location and so forth. Likewise, true marginal costs are averaged across the same group of sales staff.

To determine the total number of sales people required, the general approach would be to start with the growth and profit targets and calculate how many sales people are needed to achieve these targets. The approach also suffers from the fact that there is no way to determine where to allocate these additional sales people.

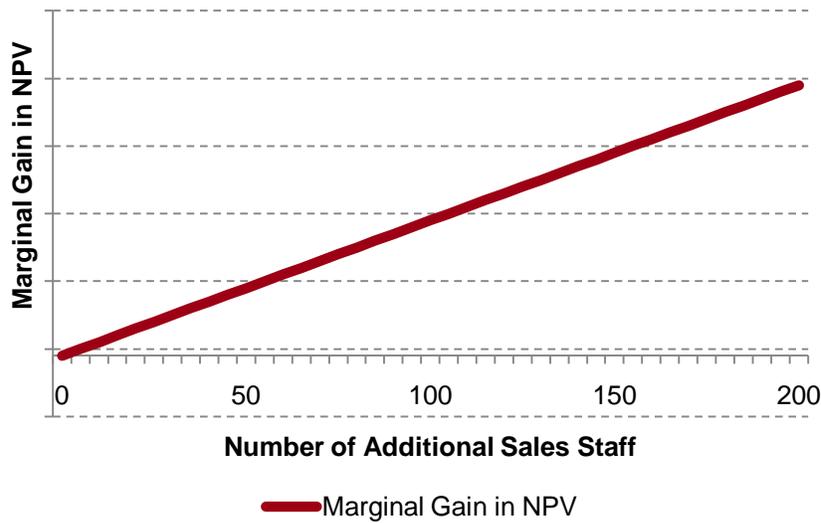
3.3 Issues with these Approaches

These approaches may vary from relatively simplistic to in some cases quite detailed. Regardless of the level of detail applied, there are a number of issues with using these types of analyses. The main issue is that there is no allowance for diminishing marginal returns from additional sales resources. This is due to the fact that there are no factors that take into account diminishing marginal productivity, for example:

- Productivity by number of regions, branches or stores serviced; or
- Productivity by number of customers allocated to each sales staff.

It is therefore not possible to determine the optimal number of resources. Figure 3 shows the implied results from analysis we have observed.

Figure 3 – Marginal Gain in Net Present Value (NPV) by Number of Additional Sales Staff



As the number of additional sales staff increased, the gain in NPV continued to rise in a constant fashion. The rational economic conclusion from this type of analysis would be to increase the sales force infinitely to secure an infinite increase in sales.

Another major issue is that the approaches assume that each salesperson, branch or store is average and is placed in a location that will be average. In this case, there will be an over allocation of resources to areas which have low potential. For example, an area may be dominated by a key competitor and the evidence suggests that there is low potential growth opportunity from additional resources. The somewhat naive approaches suggested above may continue to allocate resources to these areas to the detriment of other areas with high potential.

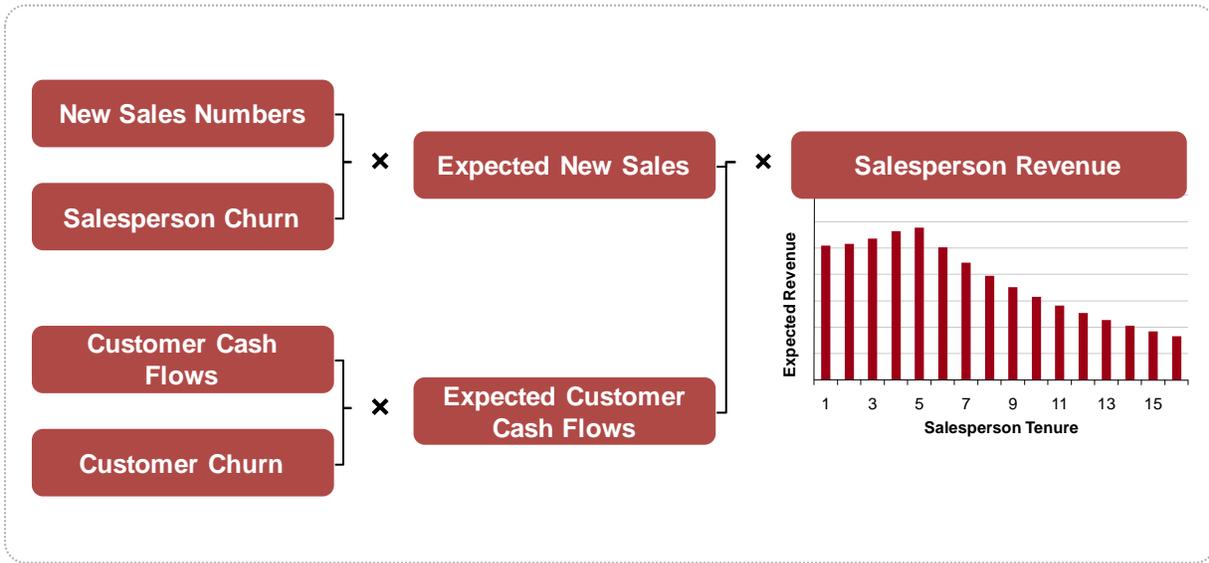
Further to this, the question of how many additional sales resources are required and where to locate them cannot be independently answered by these approaches.

Finally, there is no allowance for the changes in the sales force or distribution network for example, sales staff churn.

4 Conceptual Distribution Model

In the previous sections we discussed the ways that many organisations are assessing key distribution strategy decisions. These approaches are inadequate in most circumstances. So why do organisations continue to use them? We believe this is due to a lack of understanding of the underlying distribution model and the key components that are required. This section sets out an example distribution model for a financial services organisation however the model framework is applicable for any business.

Figure 4 – Conceptual Distribution Model

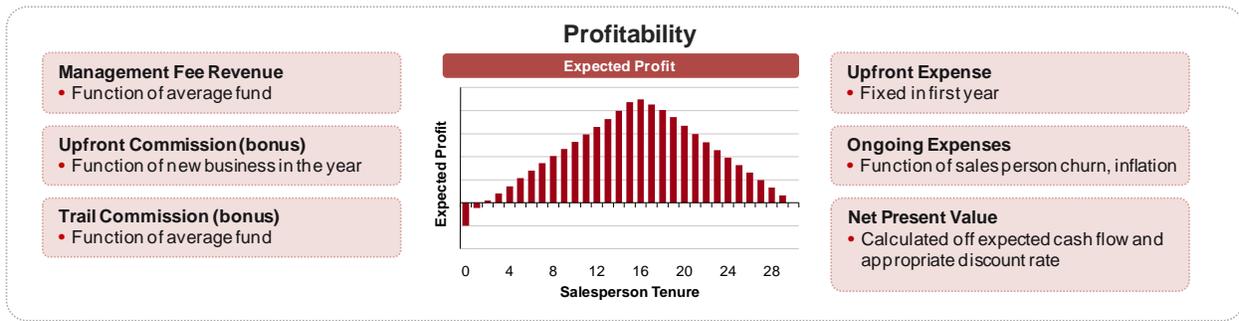


Although relatively simple, the model demonstrates the key value drivers for this example financial services product. The salesperson or advisor is crucial to the model and we begin at the top left with the estimation of salesperson effectiveness. This is the estimated number of sales generated by each salesperson and is typically strongly related to salesperson tenure. We also need to allow for the natural churn of sales people over time and combining this with the sales effectiveness we estimate the combined expected new sales.

The lower half of the model demonstrates the typical customer modelling that would be used to estimate the net cash flows from each customer. In this case we have included customer churn models and cash flow projections. The combination of the expected new sales and customer cash flow results drives the total salesperson revenue.

Combining the estimated sales performance from the distribution model with the costs and other business cash flows, we can produce a estimate of the bottom line profit. Below is an example of the profit pattern (signature) for an individual salesperson in the financial planning industry. Each of the cash flow components presented below, is combined with the salesperson effectiveness and churn models to produce the individual’s expected profit by tenure.

Figure 5 – Example Profit Signature for Financial Planning Industry



5 Sales Force Effectiveness

Just as we build cash flow models for businesses, behavioural models for customers and media effectiveness models for marketing, the sales force can be modelled for distribution strategy. The sales force effectiveness is the expected new sales generated by the sales force and in the conceptual distribution model above is represented by the “New Sales Numbers” and the “Salesperson Churn” components.

The first step in the analysis is to determine which of the areas in the conceptual distribution model should be determined using a predictive model and for which areas assumptions will suffice. In general, the growth in sales (new sales numbers) is the key driver of the sales force effectiveness and this will be the focus of the following discussions.

Below is an outline of the factors that are often considered when modelling sales force effectiveness.

Figure 6 – Sales Effectiveness Modelling



In order to identify the cause and effect relationship between sales force strategy and sales performance we need to determine all the possible influences.

The measure of “success” needs to be defined and may take a number of different forms.

5.1 Measuring the True Impact of a Salesperson

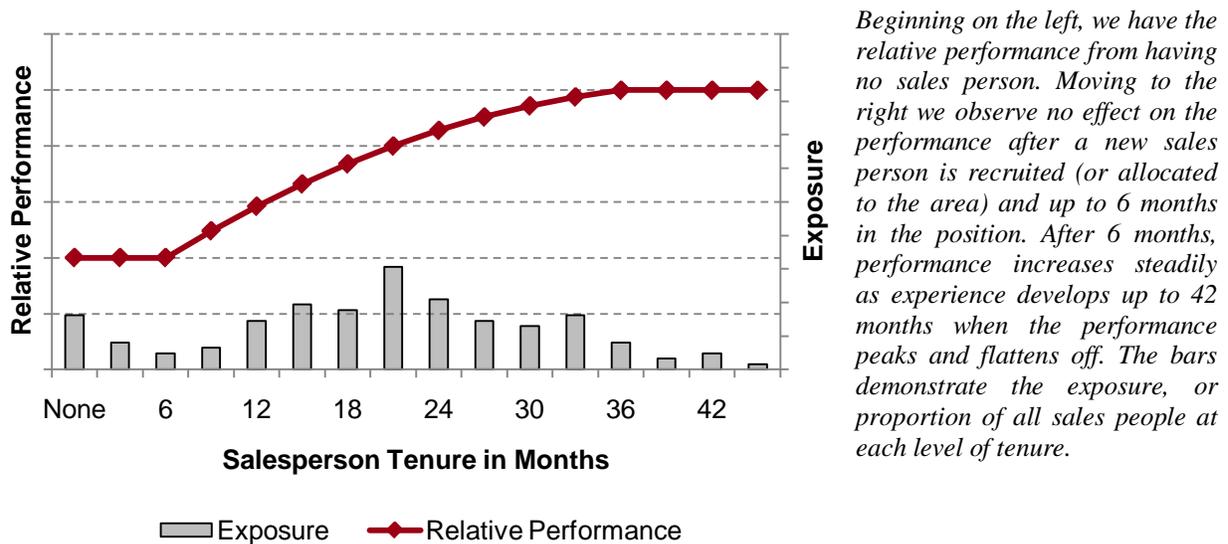
The theoretical value of a salesperson is determined by the acid test. What would the total sales be with and without each individual salesperson? In the case where we take the salesperson away, some of the resulting lost sales would be made up by other channels or sales people. In this way we are interested in the total net loss in sales from the reduction of each salesperson. Alternatively, when we add a salesperson to the team we need to measure the net increase in sales, allowing for some reduction in the sales of existing team members or channels.

As with any modelling or analytical task, we need to carefully structure the data and the modelling approach to answer the questions appropriately. Sales performance needs to be carefully defined and should be related closely to the financial performance of the business and, if possible, the remuneration structure of the sales force. We have used a number of measures including the number of new business sales, growth in FUM (funds under management), growth in customer base, etc. Typically the sales units or value generated will be the outcome of interest that should be modelled. It is also crucial to allow for past sales levels when predicting future sales and hence in many modelling exercises, growth in sales is the key outcome modelled. It is best to try a variety of measures and assess the suitability in each application.

Figure 7 demonstrates the pure effect of salesperson performance by tenure, as developed by a multivariate model. We typically use generalised linear models to develop the pure effects.

The pure effect results demonstrated in this section have been amended and de-identified to protect the organisation’s confidentiality.

Figure 7 - Pure Effect of Tenure on Individual Salesperson Performance



Beginning on the left, we have the relative performance from having no sales person. Moving to the right we observe no effect on the performance after a new sales person is recruited (or allocated to the area) and up to 6 months in the position. After 6 months, performance increases steadily as experience develops up to 42 months when the performance peaks and flattens off. The bars demonstrate the exposure, or proportion of all sales people at each level of tenure.

Sales effectiveness increases with the tenure of the salesperson. For some industries this period may be short (a few weeks or months) whilst in others it may be longer. Early periods are generally categorised by up-skilling and learning the ropes. The most rapid growth periods will occur in the middle when the

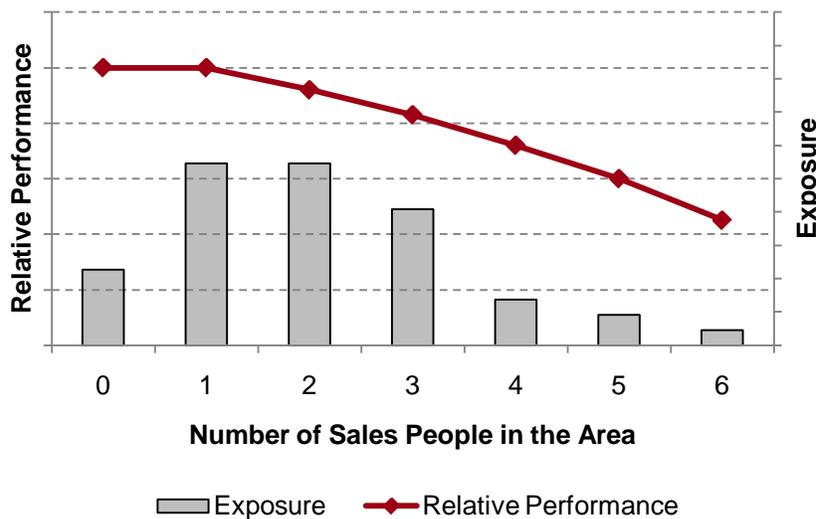
salesperson starts achieving targets. At some point the improvement plateaus as they reach upper limits of capacity or the limits of the region. Another reason for the increase by tenure is that poor performing candidates typically drop out and the remaining group have resulting higher average productivity.

This is a generalisation across all sales people - they are not all equal as we will discuss later. In each individual case the pattern of performance will be very different. Some sales people will quickly reach peak performance whilst others will never reach it.

It is essential that the modelling includes factors that control or constrain growth. Growth (above industry average) cannot continue unhindered in a finite market. This is the failing of many current approaches to distribution strategy assessment. There are many factors that will reduce the growth prospects including: the number of sales people in the area (internal competition), the market size, market share, and the external competition.

Figure 8 demonstrates the pure effect on performance from the number of sales people in each area. The total sales for the area will typically increase as each person is added but the productivity per person typically falls due to the internal competition effect.

Figure 8 – Pure Effect of Internal Competition

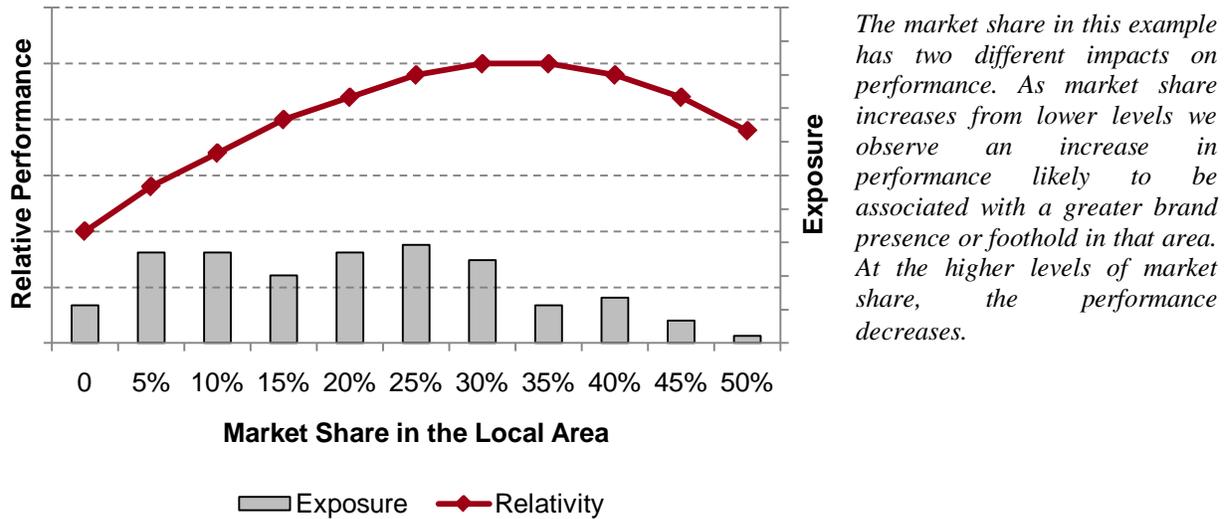


When the first person is allocated to an area, we will see an impact on sales based on the tenure effect discussed in Figure 7 but there will be no internal competition effect.

As each additional person is allocated to the same area, we will still observe an increase in performance over time due to the tenure effect but we will also observe a reduction in performance per sales person due to the internal competition effect.

Figure 9 demonstrates the pure effect on performance from the market share in the area.

Figure 9 - Sales Performance by Market Share



The market share level at which performance turns down will be different for each situation and is dependent on the strength of other competitors. In theory it may be possible to reach 100% market share however, there are typically other competitors with significant customer bases that make this unrealistic.

5.2 Which Types of Sales People Perform Better?

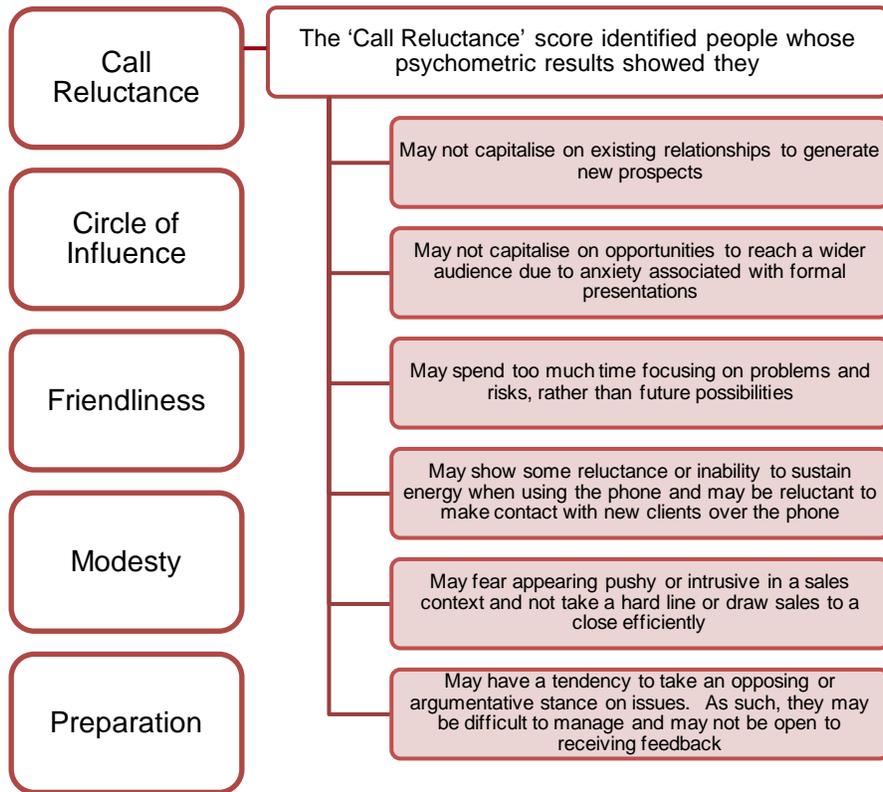
As part of this analysis we can include characteristics about the individual sales people to uncover which types of sales people perform better. This information is of primary use in recruitment decisions to select the right people up front, and in training and reviews to encourage sales people to develop more productive behaviours and approaches.

Firstly, we discuss the opportunities for using this analysis in recruitment. In many companies psychometric testing is undertaken during the interviewing stage to profile the candidates. Some of the challenges in using this data in isolation include:

- The data is difficult to interpret alone as there are many distinct questions some of which are very similar leading to confounding results;
- It is difficult to determine which questions are the better differentiators; and
- We must make assumptions regarding which types of sales people will be better performers (i.e. we must assume which questions are the best indicators of future success, rather than knowing this based on actual evidence).

Significant value can be obtained by analysing the psychometric test results in combination with sales performance data (an example is presented later). To do so, we firstly decompose the test results using Principle Components Analysis, a dimension reduction technique. The aim is to separate the raw test results into key “themes” which are independent of each other and together retain the maximum information from the raw data. Figure 10 below which shows the results of such an analysis.

Figure 10 - Psychometric Testing Themes



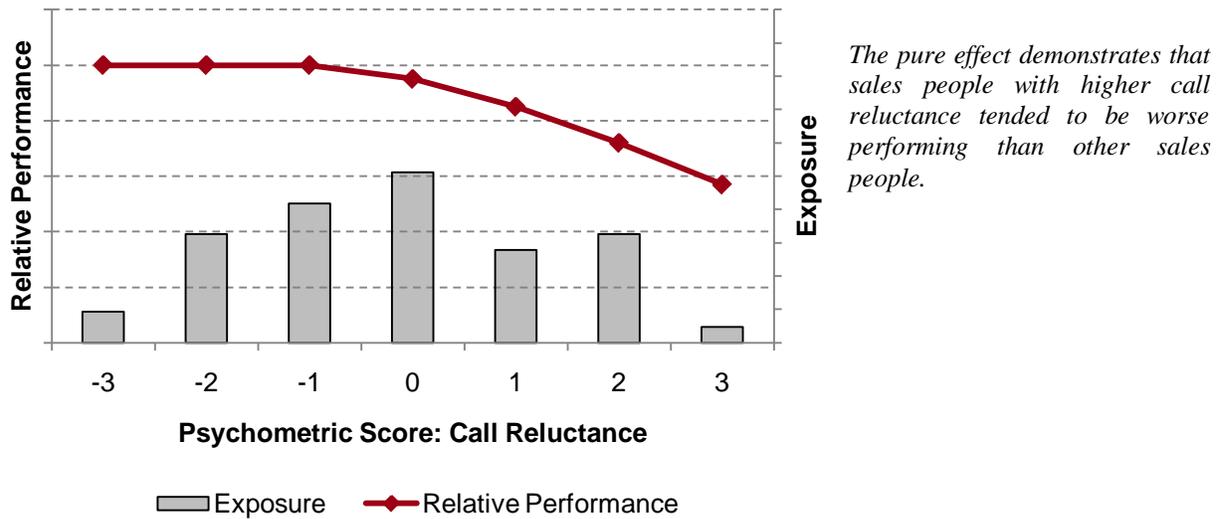
In this case, we summarised the internal psychometric testing that was conducted for new hires into five broad factors enabling us to draw out the key insights of the tests.

Using these Principle Components has many benefits from a modelling perspective:

- The components themselves are much clearer and stronger than each individual test result;
- They are independent which means that the factors will work together in the model and not impact on each other; and
- There are fewer modelling factors to search through and test.

Below is an example of a psychometric factor used in sales performance modelling.

Figure 11 - Impact of Call Reluctance on Sales Performance



5.3 Performance Benchmarking

The performance model provides an estimate of the sales results for each person, controlling for environmental and market impacts. This provides a fair and unbiased benchmark of the sales results for comparison to actual results. The performance of each salesperson can then be benchmarked against this model prediction to determine the individual performance rank. For example, if we are benchmarking two individuals in different locations the distribution model will appropriately allow for the impact of the area on sales performance. We also need to exclude any salesperson specific factors from the benchmark (model prediction), such as the psychometric scores, to give appropriate credit to the salesperson for having these attributes.

The difference between the actual sales performance and the benchmark provides us with a ranking from the best to the poorest. Using these rankings we can “drill-in” further to uncover why particular areas / sales people are either under or over performing. We can do this by investigating the combination of model factors driving the benchmark and overlay the qualitative information from the sales managers, about the areas or individuals.

6 Optimal Sales Force Strategy

Once the distribution model is well understood and parameterised, it can be used for assessing the optimal sales force strategy. Optimal strategies can either be determined by exhaustively testing all possible scenarios, by random simulation, or by selective scenario testing, to determine the most attractive strategy against a defined metric.

We note that it is crucial to appropriately allow for dynamic factors in the simulation to ensure it does not produce unrealistic scenarios. For example, when simulating additional sales people, we need to update all factors regarding salesperson saturation, the type of new sales people available for recruitment and the impact on market share.

6.1 Recruitment

The inclusion of psychometric factors in the sales performance model allows us to identify the best performing sales people which can be used to inform the recruitment process. Firstly, the Principle Components Analysis tells us which questions and themes are the strongest differentiators of candidates.

As such the psychometric test can be focused to concentrate on these questions to reduce time and complexity. We can also combine the questions into a candidate ranking based on their likely sales performance impact. Overall, the psychometric testing is only one element of the selection process and it cannot replace other selection criteria however, armed with the model findings we are much more focused in interpreting the test results and can be sure to use them in the best way.

6.2 Optimal Sales Force Allocation and Deployment

Sales force allocation and deployment typically involves answering the following questions:

- How big should the sales force be in total?
- Where should they be located and how many in each area?
- What will be the impact on sales performance for each strategy?

As previously discussed, the current approaches to sales force allocation are generally basic and rely heavily on judgement. The distribution model provides us with a precise basis upon which we can quantify the value generated by each strategy. The first approach is to work with the business to consider alternative options for sales force allocations and quantify the lift in performance of each. Through iterative testing of alternative options an appropriate strategy can be derived to best meet or potentially exceed the targets for growth and profit and hence generate significant value for the business.

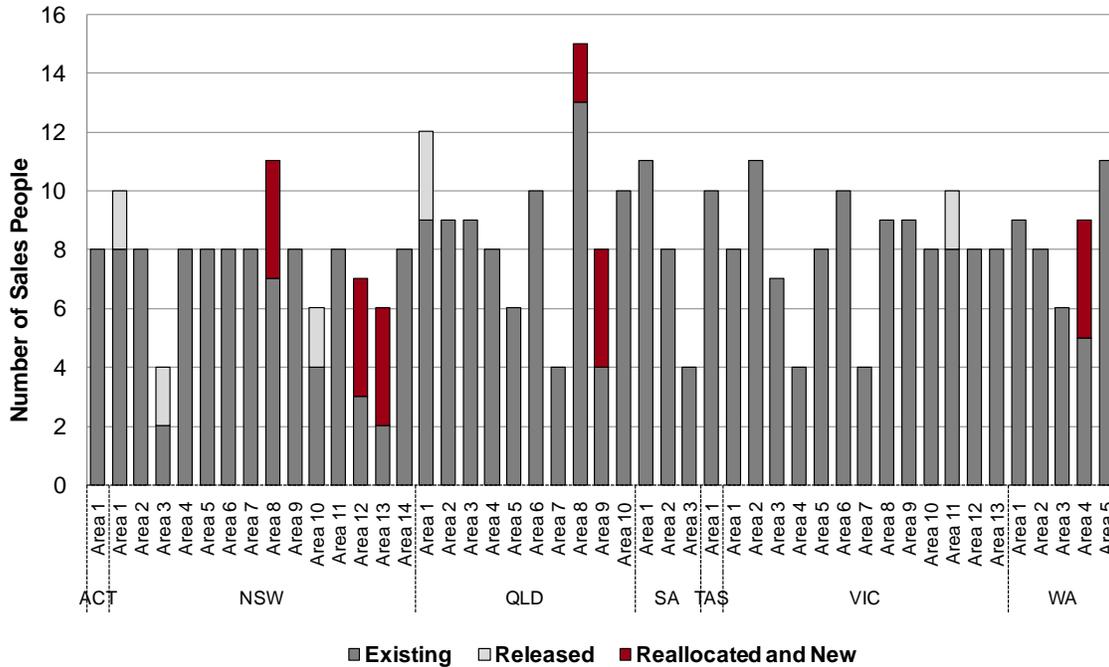
Alternatively, an analytical optimisation approach can be applied to search for the best sales force strategy. The approach involves the following steps:

- Defining the evaluation criteria (lift in sales performance);
- Incorporating constraints as required (for example the total budget for new recruits);
- Reallocation of existing resources;
 - Each salesperson is “released” from the area to determine the resulting drop in sales performance and identify areas which are over allocated. For example, internal competition may be high in some areas, so that releasing sales people will result in only marginal losses or even a lift in performance;
 - The released resources are then tested in each area to determine which area / salesperson combination will result in the greatest lift in performance. This is repeated for each remaining resource in a sequential way until they are all reallocated. Obviously practical constraints need to be allowed for such as the ability to move sales people across regions or states; and

- Testing new sales people in each area to determine the combination with the greatest lift in performance. This is repeated for all new sales people. We note that there needs to be realistic assumptions about the types of sales people who will be recruited in each area. Not all recruits can be of the highest calibre.

Figure 12 demonstrates where current resources could be released and additional sales people allocated, using an analytical optimisation approach.

Figure 12 – Optimal Reallocation of Current Resources and Allocation of 20 New Sales People

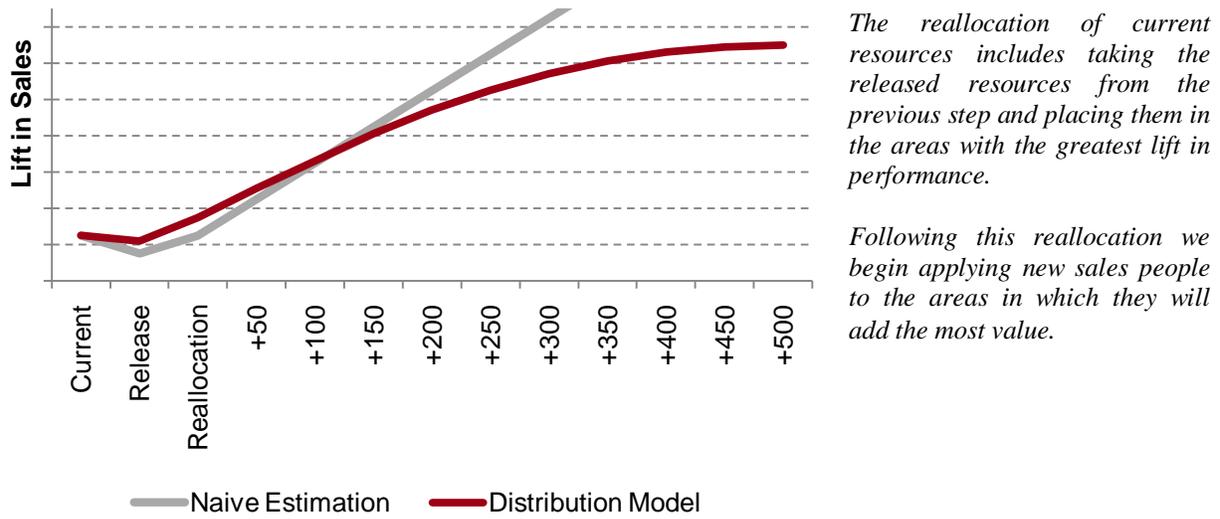


Using the approaches as described in section 3 will often result in allocation of new sales staff to areas that are relatively under resourced even though these are the areas with the lowest growth opportunity. We see that with the analytical optimisation there are actually resources released from the areas with the lowest numbers to begin with, recognising that these are worse areas for growth and hence investment should be limited. Likewise, staff are allocated to areas that are already relatively well resourced, recognising that these have the highest growth opportunity.

6.3 Optimal Channel Strategy

The optimisation will estimate the lift in sales performance for each step in the allocation process. This lift allows for all the impacts of releasing, reallocating and adding new sales people. It also appropriately allows for the contribution to new sales, cannibalisation of sales from the existing team and channels, market share impacts and ageing of the current sales force and each step in the process. Figure 13 demonstrates the relative lift in performance from the analytical optimisation process.

Figure 13 – Lift in Sales from Optimal Sales Force Allocation

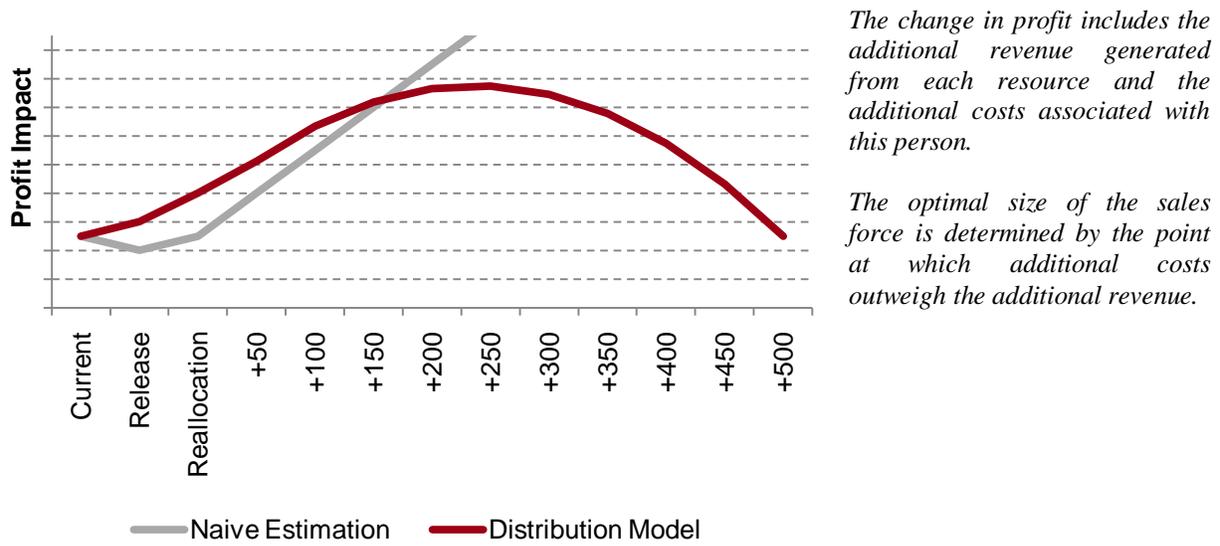


In this example, the distribution model demonstrates that the release of current resources would result in only a marginal decrease in sales whilst the reallocation of these resources would generate a significant increase. Following this, the relative sales performance would continue to increase with diminishing returns. This is due to the fact that as more sales people are added, they will be allocated to areas that are already highly saturated (and hence they cannibalise sales from other staff) or areas that have low potential opportunity for growth.

A naive estimation approach would most likely result in a decrease in sales from the release of existing resources and a corresponding increase for the reallocation with no net gain from these two steps. Sales would then increase in a linear fashion as new resources are allocated.

When the lift in sales is combined with the marginal costs we can estimate the profit impact from each step in the optimisation. In Figure 14 below, the optimal number is found by continuing to add sales staff while the marginal revenue generated by the additional sales staff outweighs the marginal cost.

Figure 14 – Impact on Profit from Optimal Sales Force Allocation



The profit impact of the distribution model demonstrates that the release of current resources would actually result in an increase in profit due to the cost savings from releasing these resources outweighing the marginal loss in sales. Following this, the reallocation step produces the greatest lift in profit and subsequent allocation of new sales people increases profit until the optimal size is reached.

As with the sales impact, a naive estimation approach would most likely result in a decrease in profit from the release of existing resources and a corresponding increase for the reallocation with no net gain from these two steps. As discussed in section 3.3, profit will increase in a linear fashion as new sales people are allocated resulting in no allowance for diminishing marginal returns and hence the optimal sales force size cannot be found.

6.4 Practical Considerations

As we mentioned previously, practical considerations need to be applied for the model to be implemented appropriately. Some of these include:

- The ability to geographically relocate resources;
- The size of the area each salesperson can service;
- The ability to recruit appropriate resources in each area; and
- Overlaying the soft information regarding the existing sales people.

We recommend implementing the optimisation process in a collaborative and user friendly way to allow these practical considerations to be appropriately incorporated.

7 Conclusion

The distribution strategy is paramount for many financial services companies. In comparison to marketing, there is considerably less analysis and rigor applied to making crucial decisions such as how much to invest and in what areas.

We have demonstrated some of the shortcomings of the current approaches applied to assessing these questions and presented a distribution model which will appropriately allow for each of the main value drivers. We have shown how the sales performance model can be parameterised on past data to understand the individual drivers of performance including which types of people perform better.

We have also discussed how one can use the distribution model to develop the optimal allocation of resources and decide what level of investment produces the greatest profit impact. Relying on actuarial methodologies, this analysis enables us to assess and determine the distribution strategy and add significant value in this field.

8 Bibliography

Adam Driussi, Caroline Stevenson and Tony Davis, 2007, *An Actuarial Approach to Optimising the Trade-Off Between Media and Price Promotions*, The Institute of Actuaries of Australia