inma
Sharing ideas. Inspiring change.

THE LATEST

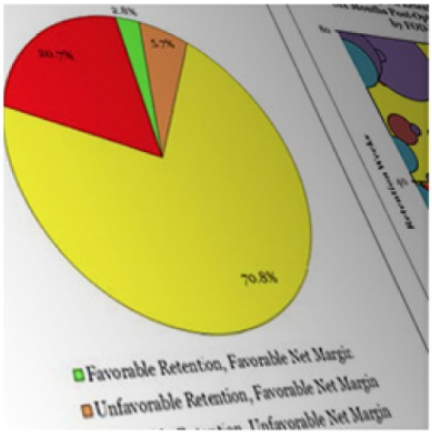

Keynote Blog | 13 May 2013

Why rifles are better than shotguns when it comes to subscriber acquisition

Each year at budget time, circulation executives ask how they should allocate their news media company's acquisition budgets among channels. Here are some solid theories explained, numbers crunched, and lessons learned regarding the value of micro-targeting subscriber acquisition. **By Jim Fleigner** ... More

---

Published by:

inma
Sharing ideas. Inspiring change.

Why rifles are better than shotguns when it comes to subscriber acquisition

13 May 2013 · By Jim Fleigner

Each year at budget time, circulation executives ask how they should allocate their news media company's acquisition budgets among channels. Here are some solid theories explained, numbers crunched, and lessons learned regarding the value of micro-targeting subscriber acquisition.

"How should we allocate our acquisition budget among our channels?"

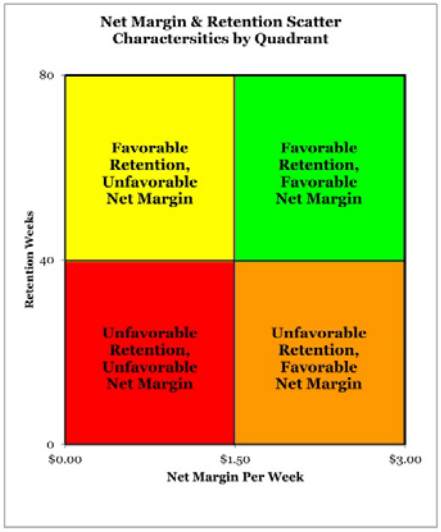

On the surface, it appears to be a reasonable question from circulation executives.

Many newspapers manage their acquisition budgets and strategies at the channel level because it is intuitive, simple, and transparent.

Perhaps as importantly, managing by channels is probably the way acquisition activity has always been managed, so the familiarity and comfort level in continuing to manage by channel is high. After all, this is why newspapers hire channel managers.

And yet, the question cannot be answered nearly as simply as it was asked, namely because acquisition dollars should not be allocated exclusively among channels.

Why? Because such an approach undercuts one of the key principles of strategy development and is inconsistent with true optimisation of subscriber acquisition performance.

The Matching Principle

A fundamental tenet of strategy development is the Matching Principle: Allocate investment resources in those segments that are most attractive in terms of their potential to generate a sufficient rate of return, and avoid segments that are least likely to generate a sufficient rate of return.

After all, most companies do a great job of employing the Matching Principle when compensating employees, as they pay the most compensation to their highest earners to prevent them from leaving for greener pastures.

Companies should employ the same principle when it comes to all investment decisions – invest the most where the prospective returns are highest.

Rate of return is driven by three factors:

1. Acquisition investment.
2. Net margin per week.
3. Weeks retained.

Thus, those segments that offer the best combination of lowest cost per start, highest weekly net margin, and most weeks retained are the segments with the highest rate of return. And these are the segments that warrant the greatest emphasis and investment per the Matching Principle.

One way to demonstrate this visually is through a Net Margin & Retention Scatter (top right of this article), where each quadrant represents a different degree of attractiveness.

The green quadrant is the most favourable (thus warranting the most investment); the red quadrant is the least favourable (thus warranting the least investment); while the yellow and orange quadrants possess a mixture of favourable and unfavourable characteristics, which make a determination of optimal investment more nuanced and conditional.

For decades, it was simple for a newspaper to apply the Matching Principle to its acquisition strategy. Because *all* starts were highly likely to generate a favourable net margin and retention (i.e., they were homogeneous in this regard), they would all appear in the green quadrant, regardless of how the starts were segmented.

This meant that the *optimal* acquisition strategy reduced to "sell as many starts as possible in as many channels as possible, because 100% of them are profitable regardless of their characteristics (e.g., delivery frequency, payment method, term, ZIP code, etc.)."

As a result, the need for newspapers to track performance simply boiled down to this: "How many starts did we sell?"

It was understood that more starts equaled more incremental net margin and higher rates of return. No prioritisations needed to be made. It was the "shotgun" approach — aim at the broadest target and blast away with the same strategy.

And it worked when all starts were effectively homogeneous.

Today, worldwide, the market could not be more different.

First, performance is no longer similarly favourable across all acquisition channels. It is typical to find that newspapers invest in a half dozen channels, yet only half or fewer of these channels generate starts with an average rate of return that is above sufficiency.

Furthermore, within each individual channel, not all starts are equally attractive. In fact, the average rate of return often masks a substantial dispersion of performance for individual starts within a channel.

As seen to the right, Channel #4 shows an average rate of return that was modestly below 0% (in the upper chart), yet an unbundling of these starts (in the lower chart) shows many highly successful starts with a rate of return in excess of 500% and highly ineffective starts with a rate of return of -100% (i.e.,

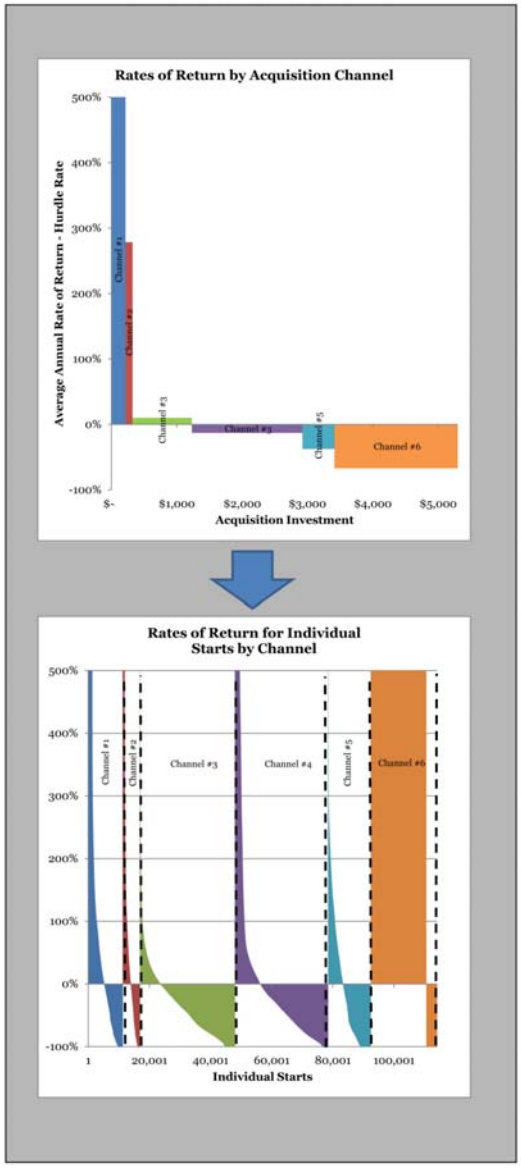

a complete loss on the upfront investment to acquire that subscriber).

This observation is important. For the Matching Principle to work properly, the ideal segmentation strategy is one where all starts within a segment are homogeneous, i.e., they demonstrate similar performance characteristics.

If this condition is present, then a strategy can be implemented for all starts within that segment that will be equally effective across those starts. Yet as we just observed, segmenting starts by acquisition channel in recent years has yielded a group of starts within each channel segment that is far from homogeneous.

As a result of these industry changes, the "shotgun approach" no longer works.

It is no longer adequate to develop strategy at the channel level, for it has become a blunt tool that pummels away at the target ("grab as many starts as possible"). It also hits not only the bullseye (i.e., sufficient starts) but also much of the area surrounding the bullseye (i.e., insufficient starts) that should be avoided.

That is an expensive form of "collateral damage" for any newspaper.

Thus, to optimise the acquisition of new starts, another segmentation scheme must be identified where the starts in each segment are sufficiently homogeneous, at least to the point where a strategy will work equally well for all (or at least most) starts within that segment.

Why shotgun segmentation schemes do not work

The underlying reason a shotgun segmentation scheme will not work is that the source of the segment performance varies. This prevents the collection of homogeneous starts in every segment, which in turn prevents the creation of effective acquisition strategy.

		Driver of Rate of Return		
		Acquisition Expense	**Net Margin Per Week**	**Weeks Retained**
Segmentation Scheme	**Channel**	**Homogeneous Starts**	Based on mix of FODs	Based on mix of terms
	FOD	Based on mix of channels	**Homogeneous Starts**	Based on mix of terms
	Term	Based on mix of channels	Based on mix of FODs	**Homogeneous Starts**
	Payment Method	All three drivers vary based on mix of channels, FODs and terms.		
	Geography	All three drivers vary based on mix of channels, FODs and terms.		

As seen in the table above, three of the most common segmentation schemes used to manage acquisition activities are channel, delivery frequency, and subscription term. Two others are payment method and geography.

The most direct source of the performance of starts is usually within only one of the three drivers of rate of return (e.g., acquisition expense for starts from one channel), which allows for homogeneous starts (seen in the grey boxes) for that one driver.

But the most direct source will not be present for the other two drivers in the same segmentation scheme (e.g., net margin per week and retention). This is true regardless of which segmentation scheme we select.

The scheme may result in homogeneous starts for one driver, but the other two drivers will always be heterogeneous because their ideal segmentation scheme is not the same.

For example:

- Differences in *acquisition expense per start* are often greatest and most direct among starts when they are segmented by *acquisition channel*, since most channels have unique commissions for their vendors.

- Differences in *weeks retained* are often greatest and most direct among starts when they are segmented by the *original subscription term*, since the term sets a floor for eventual lifetime retention.

- Differences in *net margin per week* are often greatest among starts when they are segmented by *delivery frequency*, since subscriptions that have the greatest focus on pre-print revenue (i.e., Thursdays and Sundays) are almost always the most profitable delivery frequencies.

In sum, no single shotgun segmentation scheme can effectively uncover and isolate differences in performance across all three drivers of rate of return. Thus, **more than one segmentation scheme must be used to find homogeneous starts.**

This brings us to the concept of micro-targeting.

Micro-targeting

Micro-targeting stems from the recognition that managing at the channel level – or any single shotgun level – is inherently and highly imprecise.

Instead, strategy and investment must be explicitly targeted and managed at a much lower level, i.e., by multiple dimensions that drill down within a channel to a point where most starts within a segment are homogeneous, and then to target only the most sufficient segments.

Let's analyse a top-50 U.S. newspaper to see how a single shotgun segmentation scheme has the potential to generate a strategy that is far from optimised, and then iterate increasingly closer to the "ideal" segmentation scheme that allows for effective micro-targeting.

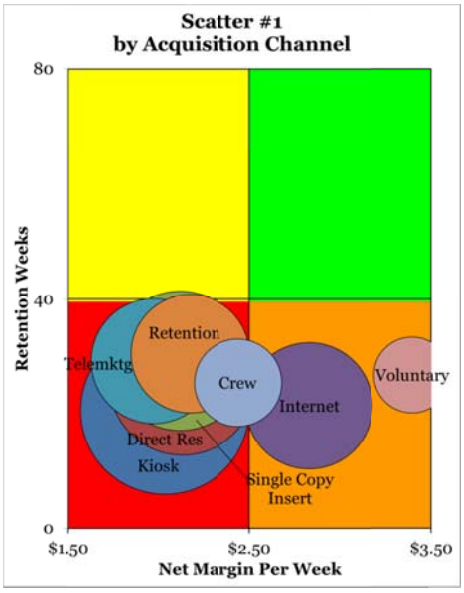

As discussed in my introduction, the most obvious segmentation historically used by newspapers is by acquisition channel. Scatter #1 shows the performance of one year's worth of starts for a major newspaper. The size of the bubble represents the number of starts from that channel.

Based on the results of this acquisition channel segmentation scheme, a channel manager might conclude that:

- All starts are expected to retain somewhere between 20 and 40 weeks.

- With the exception of the Internet channel, most paid starts generate less than US$2.50 per week in net margin.

- Not surprisingly, voluntary starts generate the highest net margin, although without greater retention.

In sum, Scatter #1 appears to support the contention that all starts are largely homogeneous for this newspaper. Because all channels (except for Internet) reside within the same (red) quadrant, it stands to reason that channel managers should not be doing anything differently and that all channels should receive the same level of investment to acquire these starts.

But this conclusion would be completely incorrect.

To demonstrate this concept, let's "drill down" within the crew channel.

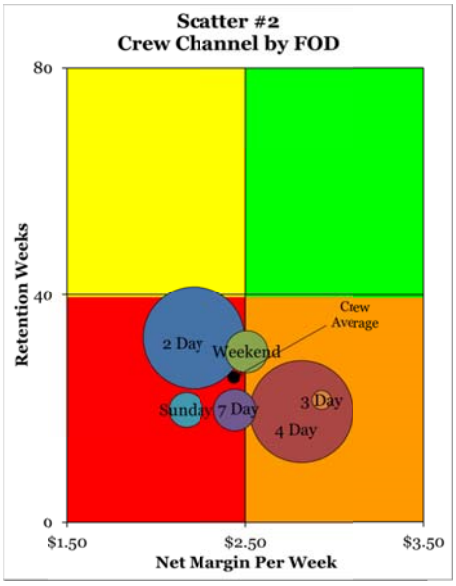

As seen in Scatter #1, its retention looks ordinary and its net margin per week is slightly above average.

But when we isolate and unbundle the crew starts and regroup them by delivery frequency in Scatter #2, the ordinary looking crew "average" (now shown as a black dot) is masking substantial differences in net margin per week, especially when one considers that three segments — four-day, three-day, and weekend – sit in the orange (and not red) quadrant.

Further, the two largest delivery frequencies – two-day and four-day – sit at the outer bounds of crew starts. The net margin for four-day is at US$2.81 per week, which is 27% higher than the US$2.21 per week that the two-day starts generate. However, the retention of the two-day starts is 32.4 weeks, which is 40% higher than the 19.4 weeks for the four-day starts.

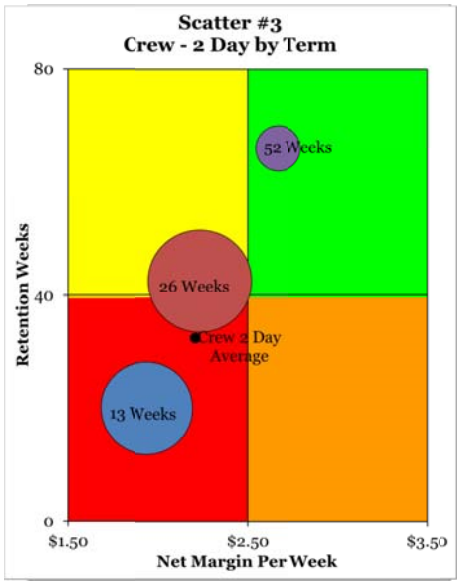

Now, let's go one level further and select one delivery frequency within crew starts: two-day (the bubbles have been resized).

As seen in Scatter #3, separating starts by subscription term shows greater differences in performance. The average retention for starts with a 52-week term outstrips the 26-week and 13-week starts that were bundled together in the average for all two-day crew starts.

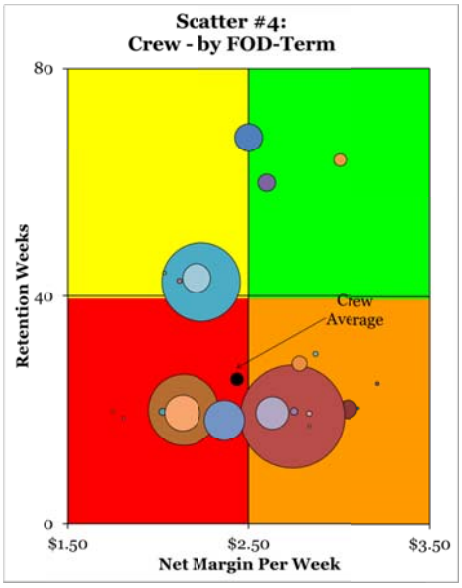

This has a substantial impact on the relative attractiveness of the two-day starts, since the 52-week starts have higher retention and net margin than the 13-week starts. Yet no one would know by examining starts performance at the channel level.

This can now be expanded to show a broader range of micro-targets. As seen in Scatter #4, drilling down by subscription term across each of the crew-channel delivery frequencies yields a distribution that looks very different than the single average crew bubble in Scatter #1.

Of course, the same analysis can be drilled down even further, whether by ZIP code, delivery route, payment method, etc. Usually, drilling down will

uncover ever greater dispersions in performance.

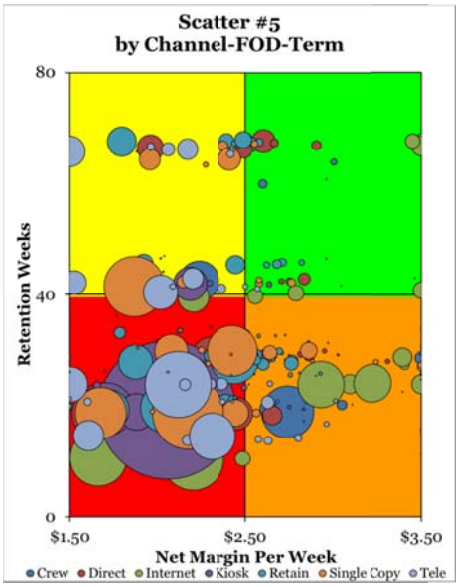

The goal is to drill down far enough such that each bubble consists of homogeneous starts.

Finally, applying this technique across all channels yields Scatter #5, which shows a very high dispersion in the relative attractiveness of each channel-delivery frequency-term sub-segment. But each bubble is much closer to homogeneity.

It would be impossible to deduce this from Scatter #1, yet this is what most newspapers implicitly do when they manage by channel.

In sum, starts segment management must be more granular to identify segments with homogeneous starts, which facilitates effective strategies consistent with optimisation.

Depending on size, scope, and complexity, newspapers can have dozens or even hundreds of segments to consider and manage.

Using micro-targeting to migrate starts

A key component of strategy development in subscriber acquisition is to find ways to migrate as many of the bubbles over time as far upward and to the right as possible into the green quadrant.

Let's examine another top-50 U.S. newspaper, both before and after a six-month acquisition optimisation effort, including micro-targeting, was undertaken:

Roughly 62% of this newspaper's pre-optimised starts came from the lower left hand (red) quadrant. With an average net margin per week of less than US$1.50 per week and an average retention of less than 40 weeks, this quadrant yields starts with the lowest average prospects for sufficiency.

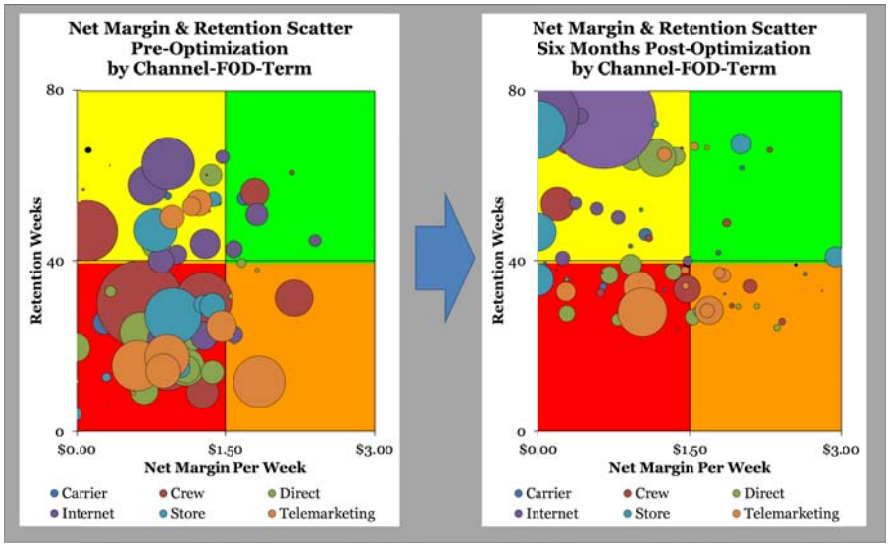

However, over the course of six months, this newspaper was successful in shifting a large portion of these insufficient starts to the upper left hand (yellow) quadrant, through a combination of new channel growth, changes to commission levels, vendor bonuses for hitting starts targets in preferred segments, and other tactical shifts.

As a result, more than 70% of the total starts came from this quadrant, with the number of starts in the worst quadrant scaled back substantially to less than 21%.

It is unlikely that such a substantial shift in overall performance would have been possible without the active use of micro-targeting.

With a conventional channel-based approach, any strategy changes at the channel level would have been highly ineffective as the newspaper would have been unaware that it had so many smaller pockets of opportunities and challenges. It would have not known where or how to apply the Matching Principle.

Summary

Newspapers are looking for every possible edge to wring maximum productivity from their acquisition investment. By embracing the concepts of micro-targeting, channel managers can use their rifles to pinpoint the greatest areas of opportunity within its sphere of influence and surgically strike.

And by avoiding as many insufficient starts as possible, the collateral damage that a shotgun approach inevitably brings can be minimised.