Building GIGO-Free Response Models

2002 PMSA Conference, New Orleans, April 28 – May 1, 2002.

Abstract

1. Overview

Establishing how your drugs respond to promotion is not necessarily confined to the realm of guesswork. While Delphi sessions can help guide the construction of response curves, they are not the only option. This article describes how to build response curves by aggregating responses of individual physicians. The ingredients to build the promotional response of individual physicians are readily available: IMS Xponent or NDC Source Prescriber on the one hand and the company call file on the other. The build-by-aggregation approach described here has three key advantages. First, the resulting promotion response is far more accurate than the best educated guess concocted by a group of experts. For one thing, it is anchored in reality. Second, response curves can easily be derived for specific territories, districts, areas, managed care organizations, or other physician segments. This is next to impossible if you only have a nationwide response curve to start with. Third, adaptive targeting is a reality. This means you can adjust in real time the number of details that go towards each physician individually, based on the promotion response feedback you get from the physicians.

Why is it then that build-by-aggregation is not the industry's gold standard? It will be. Up to recently, two factors have been holding things back. First, insufficient computing power to handle 800K+ physicians at one go. This is no longer true thanks to increased desktop computing power (Moore's law) and the ongoing price war in the PC industry. Second, technical difficulties associated with build-by-aggregation. For starters, how do we add two response curves of different lengths, say one that spans two details and another that spans twelve? The obvious $[f+g](x) = f(x) + g(x)$ does not reflect reality since it amounts to assuming the physician that receives the smaller number of details does not respond past a few details. Here is the larger issue: In what order should we add the individual response curves? If the order were not important, that would mean which physician gets detailed first is immaterial and that would fly in the face of targeting. The obvious schemes for adding curves, as we will see, do not work. Lastly, it is imperative that we be able to answer questions such as: What would have happened had we delivered an additional 250K details last year? Somehow, the response curve should overshoot the actual number of details delivered. This article addresses these technical issues head-on.

2. Response Curves

Promotion response curves describe how sales increase with promotional effort. Promotion is usually taken to mean detailing, but can also refer to sampling, meetings and events, journal ads, DTC, etc. Promotion response models play a pivotal role in sales and marketing since they allow us to establish the ideal promotional effort to maximize sales, long-term gross profit, or more importantly, any combination of the two. Promotional response models help us play out implications of what-if scenarios such as increases in standard costs, reductions in promotional budget, changes in promotional budget mix,

etc. Some people even call promotion response the holy grail of the pharmaceutical industry!

How do we ensure the promotion response is an accurate depiction of reality? Making the right decision based on the wrong model may cost tens of million of dollars. Indeed, GIGO (Garbage In – Garbage Out) may befall the promotion response model. Delphi sessions are a great forecasting technique (invented by the Rand Corporation during WWII to harness the wisdom of international experts) but are oftentimes misused and abused, probably for want of better alternatives. To that end, we developed a novel response building technique that takes the guesswork away and builds the aggregate response by combining promotion responses of actual physicians.

3. Build-By-Aggregation Approach: Issues and Solutions

Let's focus on the top two technical issues associated with building promotion responses by aggregation. First, the promotion response of some physicians (actually a lot) may be too short, i.e., the physician received only one or two details. If we simply add a short promotion response to a long one, we would be making the tacit assumption that the short response physician no longer responds past a few details, which is clearly incorrect. Second, in what order should we add the promotion responses of the individual physicians when aggregating the responses? The catch here is different ordering schemes lead to different aggregate promotion responses.

How do we address insufficient response information associated with under-detailed physicians? By leveraging the fact that we have detailed more profusely other physicians with similar characteristics. The key here is to assume the response of the under-detailed physician is a blending of the observed responses of comparable physicians: other physicians with similar specialties, managed care plan affiliations, or geographic locations. Indeed, we combine observed responses that we "graft" to the short response. That graft is represented not as a single segment of a response curve but as a cone, to capture uncertainty. The lower portion of the cone represents the low response scenario, the higher portion the high response scenario, and the distance between the two portions the magnitude of the uncertainty. This surgical grafting actually does two things. For one, it allows us to be more assertive regarding the response of individual physicians we barely detailed. For another, it allows us to construct an aggregate response that goes beyond the number of details actually delivered. This allows us to address questions such as "What would sales have been last year had we delivered an additional 250K details?".

The second issue is how to combine the individual responses to generate the aggregate response. Consider the most-responsive-first approach. This scheme sorts the physicians by decreasing responsiveness and connects their response curves starting with the most responsive one first. Because the physicians are sorted, the first half of the aggregate response will contain only half of the physicians detailed. This means reps will be granted the luxury of not detailing less responsive physicians even when they have free time on their hands. Rep compensation being sunk cost, this scheme is clearly unrealistic. Consider the round-robin approach instead. Under this scheme, it is only when all of the physicians have received their first detail that a second detail may be delivered, and so on. The first half of the aggregate response will contain all the physicians. Unlike the most-responsive-first approach, reach is preserved. Frequency, however, is butchered. The problem with the round-robin scheme is it treats every physician equally. It does not discriminate between very responsive and less responsive physicians, and discrimination

is precisely what targeting is about. Those two schemes just made two things very clear. First, the ordering scheme we are after must preserve both reach and frequency. Second, each response curve corresponds to an implicit targeting strategy. The ordering scheme we recommend is random summation. Like the round robin, random summation adds up segments of each individual response curve. Unlike the round robin, it picks physicians in a random manner. As a result, Dr. John Smith may get his third detail before Dr. Jane Doe gets her first, although large differences in details are most unlikely. The probability with which a physician is chosen is not inversely proportional to the number of physicians, but to the number of details yet to be delivered. This increases the odds of detailing a heavily detailed physician early on in the targeting process. As a result, both reach and frequency are preserved. In a variation of the random summation, effective details are distinguished from ineffective details (those delivered past saturation) and the probability is based on effective details only. Ineffective details are thereby deferred to the very end of the targeting process.

4. Implementation

To implement the build-by-aggregation approach above, we developed a physician promotion response database in Access and promotion response matching algorithms in VB. These tools allow us to retrieve all physicians whose promotion response starts with, ends with or contains a specific user-defined promotion response pattern. Those promotion responses are then combined and grafted to promotion responses of physicians that have not been adequately detailed.