Sawtooth Software

RESEARCH PAPER SERIES

Uncommon Choices: Novel Applications of Conjoint Analysis in Practice

Chris Chapman,
Google NBU (Next Billion Users)

3210 N. Canyon Rd., Provo, Utah
+1 801 477 4700
www.sawtoothsoftware.com

UNCOMMON CHOICES: NOVEL APPLICATIONS OF CONJOINT ANALYSIS IN PRACTICE

CHRIS CHAPMAN
GOOGLE NBU (NEXT BILLION USERS)

ABSTRACT

This paper is a managerial review of four innovations in conjoint analysis practice. Although each case stands independently, the combination of methods demonstrates how experienced conjoint practitioners may extend both their breadth and depth to tackle larger and more strategic projects. The cases include improved market prediction, anticipation of competitive response, and the application of choice-based conjoint analysis for psychographic segmentation. The descriptions here are high level, while technical details and code are provided in associated whitepapers from previous Sawtooth Software Conferences.

INTRODUCTION

In this paper, I review and link a series of prior work to demonstrate how relatively incremental innovations in conjoint analysis contribute to a vast expansion of applications for choice modeling practitioners. These innovations were presented at previous years of the Sawtooth Software Conference, among other venues. I recap the methods here, discuss how they fit together, and add new market observations and practical reflections for practitioners.

There are four core methods discussed here:

1. *Adaptive choice-based conjoint*, and how it may obtain high quality estimates
2. Using *game theory* with conjoint analysis to predict competitive response
3. Finding optimal product portfolios in the presence of competition, with *genetic algorithms*
4. Using choice-based conjoint to find *psychographic consumer profiles* (segments)

My goal here is to give an approachable introduction to the methods, explaining why each may be of interest to practitioners. Technical details, including links to R code for some methods, are available separately in the *Proceedings* of prior Sawtooth Software Conferences (see references in each section). For a more general introduction to the typical applications and problems for conjoint analysis with technology firms and products, refer to Love & Chapman (2007) and Chapman, Love, & Alford (2008).

For purposes here, I assume general awareness of the core concepts of conjoint analysis (Orme, 2014), types of conjoint analysis surveys (CBC and ACBC), and the foundations of market simulations. My aim is not to explain every application in depth but to discuss how and why each is useful, and how the methods and cases build on one another.

FOUNDATION: GETTING THE BEST INDIVIDUAL ESTIMATES WITH ACBC (CASE 1)

The first premise in this paper is simple: if we obtain the best data at the individual level, we will have the most and best options for later analyses. Aggregate, group-level analyses are generally OK with relatively imprecise individual-level estimates. However, when we have better estimates, we can make better predictions and improve other approaches such as segmentation and—as I argue in the following two sections—competitive and portfolio modeling.

In many projects, I have found that Adaptive Choice-Based Conjoint (ACBC) performs well at finding individual-level estimates that are effective for market share prediction (Chapman et al., 2009; Johnson & Orme, 2007).

Case 1: Background. We had developed and planned to launch a consumer electronics product "P1," when a competitor announced a product "C1" that would compete very closely. Arguably, C1 offered a strictly superior set of features at the same price point as P1. Our business leaders wondered whether P1 had any realistic chance of success vs. C1 in the market.

Case 1: Operational Research Question. The executive team decided that if our product achieved <25% estimated market share in a strict two-way head-to-head competition with competitor C1, then we would *cancel* the plan to launch P1. In my experience, it is rare for an executive team to give such a well-defined decision criterion. This reflects both the executive team's mature understanding of research, as well as previous experience and trust with this author's group and our research methods.

Case 1: Method and Results. We did not wish to make such a decision on the basis of a single method or estimate, and instead used 4 methods to estimate preference share for P1: (1) traditional choice-based conjoint (CBC); (2) a holdout task asked in CBC format, for the exact comparison of P1 and C1; (3) adaptive CBC (ACBC); and (4) a head-to-head offer for the respondent to choose one, either P1 or C1, presented in a richer, more descriptive style similar to retail marketing.

As shown in Figure 1, all 4 methods estimated P1 would attain preference share of at least 25%. We felt confident reporting that P1 would exceed 25% share, and thus there was no reason to cancel its launch. The firm subsequently released P1 as planned.

As detailed in the complete white paper (Chapman et al., 2009) we believed that ACBC provided the most credible estimate with a point estimate of 33% preference for P1 vs. C1. This was due to the greater amount of information collected in ACBC for each respondent, along with indicators of better internal consistency, such as the lack of attributes that showed level reversals. We reported the estimate of 33% to the executive team as our best prediction of consumer preference.

Figure 1: Estimated preference share for P1 vs. C1, using 4 methods.

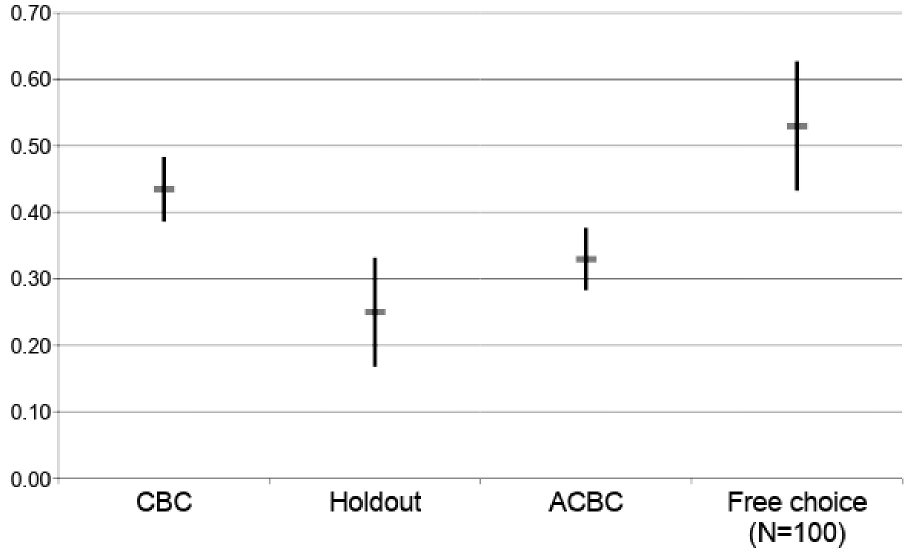

How did it perform? Several months after the release of P1 and C1, we observed an actual market share of 34.6% for P1 vs. C1, compared to our estimate of 33%—a difference of less than 2% absolute, which was well within the confidence interval in ACBC market simulation.

Case 1: Notes for success. At the conference, we discussed whether this is an expected level of accuracy for ACBC. That is impossible to answer, but I would make a few points. First, this is not a unique or exceptional case in my experience. I have had the opportunity on only a few occasions to compare conjoint estimates to actual market data in any clear way. On each occasion, conjoint has performed well (see Chapman, Alford, & Love, 2009). However, I believe it is not the expected performance for the *first* occasion to conduct a conjoint analysis study for a product line. Such studies require iteration to learn about the attributes that matter, how to ask them, and how to tune aspects of the model such as brand effects. Notably, although we had conducted CBC on many occasions in this product space, this was our first application of ACBC.

Overall, my conclusion is that ACBC is capable of delivering the best individual-level estimates, and these may lead to the best-performing preference estimates. This comes at the cost of greater survey complexity and respondent time. Table 1 summarizes some of the considerations between CBC and ACBC.

Table 1: Brief comparison of considerations for CBC vs. ACBC.

	CBC	ACBC
Survey length	**Shorter**, 4–10 minutes	~2x as long, 8–15 minutes
Sample	Larger standard deviations ~2x sample needed for same precision	Better with **smaller samples**
Individual-level precision	Moderate, depending on number of tasks	High. Especially good when **segmentation** is desired
Experimental design	**Full control**; easy to control for information density, nested attributes, prohibitions, etc.	Less control; depends on the ACBC process
Non-compensatory features	Limited assessment; relies on design matrix	**Moderate to high** ability to assess (higher=longer)

In cases where precision is crucial—as in the present business case—I highly recommend considering ACBC, and to make this possible by shortening other aspects of a survey. In the next section, we will see an analysis that benefits from the highest-quality individual estimates.

OPTIMIZE: RESPOND TO COMPETITION USING GAME THEORY (CASE 2)

A common question for any product manager considering an action is, “How will competition respond?” In this section, I consider a case where there was an opportunity to improve the feature set of a product, but at a higher cost of goods. Would it be worth it?

Case 2: Background. A consumer electronics product line “PL2” was presented with the opportunity to improve the nominal performance of a feature “F1” that was known to be highly salient for consumers. However, we knew that F1 would increase a product’s cost of goods and complexity of engineering, and believed that F1 might not yield any real improvement in the consumer experience. If our key competitor did not offer F1, and consumers ended up believing that F1 didn’t matter, then we would show lower profits for nothing. The executive presumption was that the competitor “Z” would *not* offer F1, and also that we should not; it was expected that our firm and Z would both benefit if F1 was in neither product line. This is structurally the same expectation as the so-called “prisoner’s dilemma,” where the overall best outcome is for both players to make the same choice of not “defecting”—in this case, adding F1—but only if both players make the same, independent decision (Myerson, 1991; Chapman & Love, 2012).

Case 2: Operational Research Question. Based on the success of prior analyses (such as Case 1 above), we expected that conjoint analysis was likely to yield a good answer as to the consumer value of F1 in the context of a competitive landscape. Unlike Case 1 above, in this case we were concerned with the effect on an entire product line, because F1 might be offered in multiple products and would be expected to shift preference within a product line.

We conducted a conjoint analysis survey and then estimated the net preference for preference of our line PL2 vs. the key competitor Z's product line "CL2" in four scenarios: that we offer F1 and they don't; that they offer F1 but we don't; that we both offer F1; and that neither of us offers F1 (for more on game theory and conjoint analysis, see Choi and Desarbo, 1993). In each case, we estimated the preference share for our line PL2 vs. their line CL2, in the presence of a "none" option estimate (see Karty, 2012, for discussion of the importance and difficulty of "none").

Case 2: Results. Figure 2 shows our model in the "extensive" format for a game theoretic analysis. First we consider what happens if we do *not* offer F1 in our line PL2, as shown in the two market simulation results on the left-hand side of Figure 2. The answer is that the competitor Z would see a large increase in preference share by offering F1—an increase from 44% preference to 72% preference vs. our line and the "none" option. Conclusion: if we don't offer F1, competitor Z will offer it.

Figure 2: Extensive form of the PL2 vs. CL2 competitive game for feature F1, with competition from brand Z.

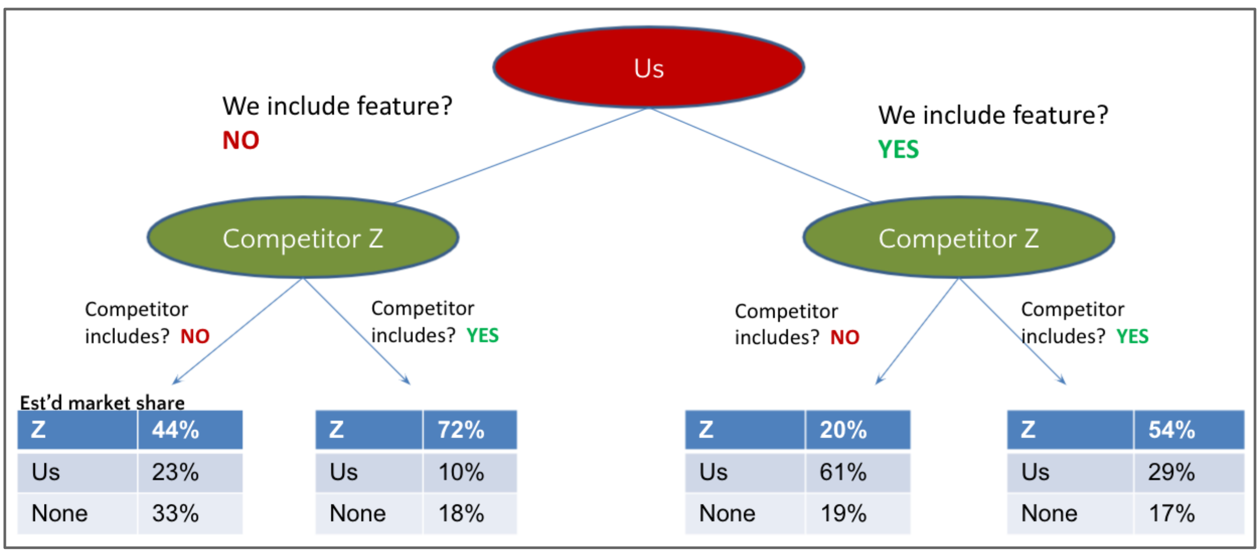

Next, we consider what happens if we *do* offer F1, as shown on the right-hand side of Figure 2. We see that if we offer F1, then competitor Z should also offer F1—otherwise they obtain a share of 20% vs. a possible 54%. Conclusion: if we offer F1, then Z also will offer it.

Now, given that Z should offer F1, regardless of what we do, which choice would be better for us? Comparing the two "YES" paths for the competitor, we see that if they offer F1, then we should also offer it—that increases our preference share from 10%

without F1 to 29% with F1. Conclusion: we should offer F1 if they do (and also if they do not).

Finally, what about the prisoner's dilemma? Comparing the left-most estimate in Figure 2 to the right-most estimate, we see this: if we both offer F1, we *both* expect to see market share increase. That is because the inclusion of F1 may pull in new purchasers vs. the "none" option.

The executive team was convinced by the data and analysis, and included F1 in our product line. *Were we right?* The product line share is difficult to estimate because the modeled portfolios for both us and the competitor rolled out over time. However, I can note the following: competitor Z initially did not offer F1 (in other words, the executive expectation about their likely action was correct), but our flagship product with F1 was highly successful, winning awards and achieving high sales. Z followed late and added F1 to their line many months later. In other words, the outcome closely matched the expectation from our game theoretic analysis, and our product line improved its competitive positioning. Finally, it turned out that feature F1 did in fact improve the user experience, which had been difficult to ascertain in advance.

Case 2: Keys for Success. As in Case 1, I cannot claim that such a successful outcome should be expected routinely, yet I also do not have examples of failure. Rather, my belief is that such success is attainable only if one has spent time to develop experience in a product area across multiple rounds of careful research and analysis. An important point in this case is that the "none" option turns out to be crucial: it is the part that breaks the prisoner's dilemma. That is a consideration that will need careful attention for such competitive modeling of a product line. Additionally, our experience in the space allowed us to add brand effects that produced better market share estimates, calibrated across multiple studies.

My recommendation from this case is simple: even basic game theory models may yield strong insight into likely competitive responses. In the complete paper, we consider a more complex model involving potential branding efforts (Chapman & Love, 2012). A corollary recommendation is this: don't bet against what customers are telling you. If customers want F1, then it's a good bet to give them F1, rather than attempting to outguess them.

This case considered a hand-crafted product portfolio vs. a competitor's expected portfolio, with regards to changing a single key feature. But what if we don't know what portfolio we should make, in light of many potential features? Can we get insight into an entire portfolio? The next section discusses how to optimize a portfolio using "genetic" search algorithms.

EXTEND: BUILD A PORTFOLIO USING GENETIC ALGORITHMS (CASE 3)

The previous sections discuss how to obtain improved estimates of product preference, and how to combine game theory with conjoint analysis to improve the definition and positioning for a product with regards to competition. In this section, I discuss how to generalize that analysis to an entire product line.

This extension answers three crucial questions. First is an obvious question: what is our optimal product line with respect to demand? Second, and closely related: how many products should we make? Finally, the third question may be more interesting: are there features that should be added to the product line?

Case 3: Background. For a consumer electronics product line (differing from Cases 1 and 2), the firm was making more than 20 SKUs, so many that engineering, marketing, and sales channels were excessively complex. The executive team wanted to know whether, and by how much, to reduce the product line size. One choice, of course, would be to cut the worst-performing products. However, consumer preference would shift around as the portfolio changes, and the popularity of products might not be related to their margin. For example, suppose a popular product has a lower margin. If we cut that product, would we lose sales to competition, or might we recapture them elsewhere in our portfolio, perhaps with currently less popular products that would have higher margin?

Case 3: Operational Research Questions. For purposes here, I will focus on two questions (see the complete writeup for others; Chapman & Alford, 2010). First, how many products should we make? We know that more products will always result in some additional share, but this should level off at some point where each additional product leads to a small incremental increase in share. In particular, if we find an optimal product portfolio for us (with the option to include a product portfolio for competitor Z), in the presence of the "none" option, how many products are needed to reach the point of strongly diminishing returns?

Second, is there a feature not currently offered in the product lines that often appears in optimal portfolios, and should be considered for addition to one or more products?

Case 3: Method. Our method to answer this question included four key elements. First, we wanted to obtain the highest-quality *data* (see Case 1 above). Not trusting a single method or survey, we opted to perform the analyses using two data sets from consumers that tested identical sets of attributes and levels; one using CBC and one using ACBC.

Second, we had to define an *outcome metric* to optimize. In this case, we optimized for the total preference *share* for any of our products—simply summed together, although one might instead optimize for revenue or profit, or a combination of metrics (Ferguson & Foster, 2013)—compared to the "none" option using randomized first choice market simulation. (Note that it is also easy to include competitors' products in the market simulation set, which then optimizes for share vs. competition, as in Case 2 above. In subsequent projects with this GA method, we have often done that.)

Third, there must be a *method to search* the product space. A genetic algorithm (GA) is an optimization method inspired by an analogy to evolutionary biology, in which best solutions are found by recombining parts ("genes") from prior, less optimal solutions. Belloni et al. (2008) demonstrate that GAs may achieve near optimal search of complex product spaces. Goldberg (1989) is an excellent technical guide to GAs. For conjoint analysis data, the genes represent product attributes and levels (i.e., product features, brands, and prices). Figure 3 presents a schematic representation of a GA approach to finding an optimal portfolio with conjoint analysis data. We implemented this in R using a standard GA library (Mebane & Sekhon, 2011; also R Core Team, 2021).

Figure 3: Outline of a genetic algorithm approach to find an optimal portfolio.

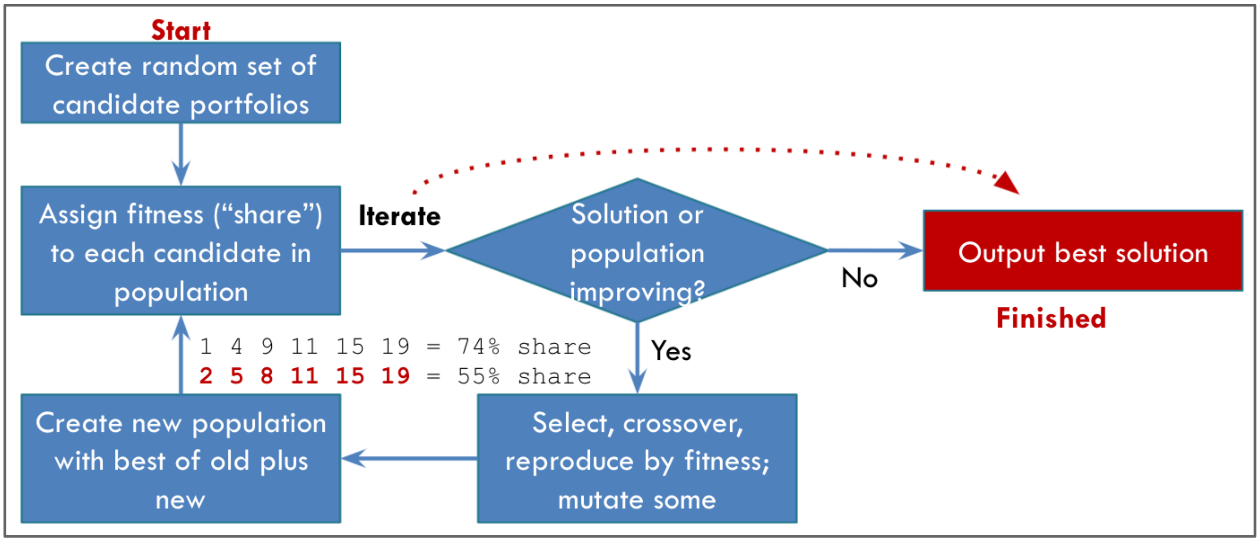

The fourth and final part of the method is *iteration* over the space of potential portfolio sizes. The GA method is stochastic; any single "best" solution is one view, but we need more comprehensive insight into the *distribution* of possible solutions. We accomplished this by repeatedly finding GA solutions for each portfolio size. Specifically, we examined portfolio sizes from 1 to 20 total products in the product line; and ran 50 iterations of the model to find 50 unique near-optimal portfolio solutions for each size of product line. This was repeated for the two sets of data, from CBC and ACBC surveys.

Case 3: Results. Our first question involved optimal portfolio (product line) size. Figure 4 shows the incremental change in total portfolio preference share (sum of all products vs. "none") as the portfolio size increases. In results derived from both the CBC and the ACBC data, there was very little expected incremental gain in total preference when lines exceeded 8 products. Each additional product above 8 gains share primarily at the expense of other products within the portfolio, and only about 1% additional preference vs. "none." This suggested that we were making far too many products.

Figure 4: Change in total share of preference for the portfolio, for each incremental product, as the portfolio grows from 2 to 16 total products.

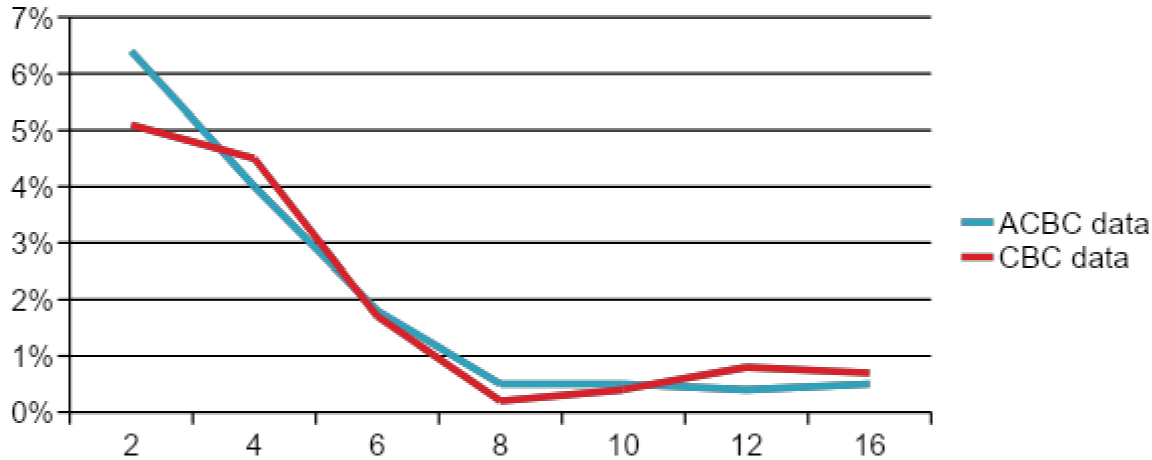

The second question was whether there was a key feature that should be in a product line but was not. Figure 5 shows the total preference share of the products in an optimal portfolio that include each CBC/ACBC attribute. For example, suppose there is an optimal line with 8 products, and 2 of those products include Attribute 2, Level (feature) 2, with shares of 10% and 5% each. We would say Attribute 2-2 thus has 15% total share.

One feature that was not currently present in any portfolio, for us or key competitors, was Attribute 2, Level 2. In Figure 5 we see that an optimal portfolio that includes Attribute 2-2 would have about 29% (range 14–43%) of preferred products with that feature. We concluded that this feature would be popular with consumers, as it should appear in roughly ⅓ of the products chosen in this product space, and we recommended consideration of including it in the product line.

Figure 5: Total preference share of products that include each feature (attribute level), in optimal portfolios.

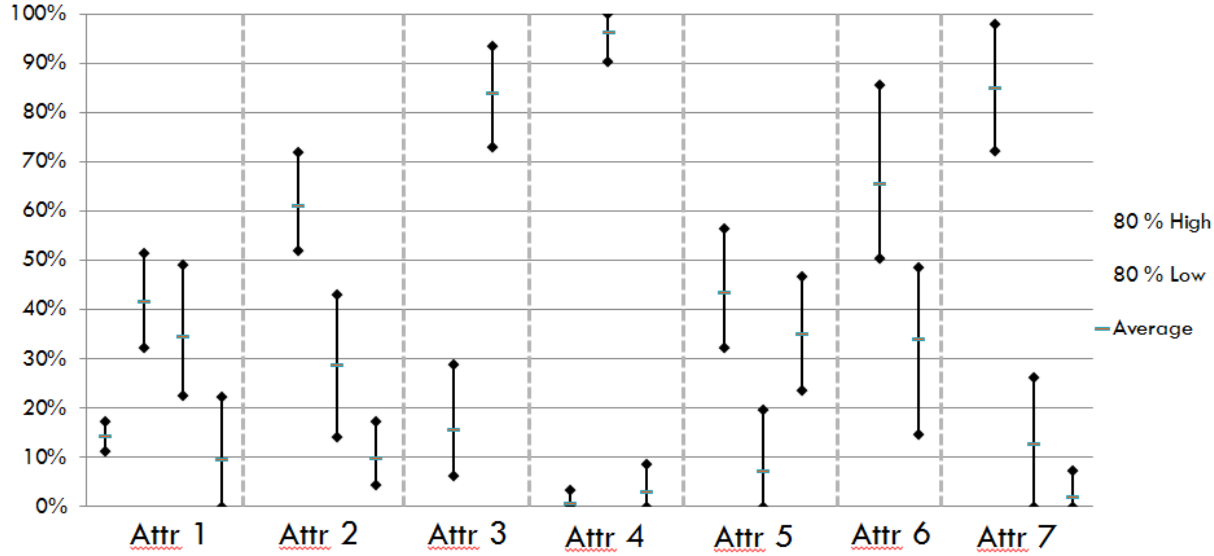

What actually happened? First, the product line was shrunk. Similar to Case 2 above, it is difficult to determine the exact extent to which market results corresponded to our analysis, because the portfolio changes took significant time to achieve in the market, while other market variables continued to change. However, subsequent years did not see the firm return to a larger, expanded product line, so we can at least note that it was a stable strategy.

Second, because of its complexity, the product team decided not to include Attribute 2-2 in the product line, despite the expected consumer popularity. We view that as an analytic success because it was strongly considered, even if ultimately out of scope. We also note that 1.5 years later competitor Z introduced Attribute 2-2 in its product line. Several years later, it remains a core part of the product line for Z as well as other brands in the market. I would consider this also to be a success for our analytics team because we forecasted the desirability of that feature far in advance of its introduction, and thus anticipated a likely change in direction for consumers.

Case 3: Keys for Success. Similar to the preceding cases, the most essential factor in success for this case was high quality consumer data built on experience in this product space. We also conducted the analyses with data gathered through two methods, CBC and ACBC, and the agreement in analysis increased our confidence in the results. Another important aspect, similar to Case 2, is careful attention to the choice task design and the nature of the "none" attribute (Karty, 2012; Dotson et al., 2012; Huber, 2012; Chapman, 2013). Finally, it required innovation in method and a substantial degree of customized code (see the original whitepaper for more discussion, including availability of the R code; Chapman & Alford, 2010).

As I have noted already, I would caution an analyst not to expect such success in general; yet I also have no particular reason to expect less success. The case reflects our direct experience without "file drawer" selection issues. At the same time, our depth of experience in the product line, and our closeness to the engineering and executive teams were probably unusual, compared to a typical analyst's or external supplier's position. Such experience, communication, and trust no doubt contribute to the success of these kinds of strategic projects.

UNDERSTAND: DEVELOP CONSUMER SEGMENTS WITH PROFILE CBC (CASE 4)

Cases 1, 2, and 3 demonstrated a logical progression in the depth of questions that one can answer with conjoint analysis data: good data leads to exceptional prediction of consumer preference; this may be used for competitive insight; and this can be broadened to answer highly strategic questions about an entire product line.

In the final Case 4, I turn to a different question and ask, if one becomes an expert at conjoint analysis (Orme & Chrzan, 2017), what other, highly novel applications might one tackle? In this case, we look at the problem of consumer segmentation, specifically the problem of building psychographic profiles that are sometimes known as "personas."

Case 4: Background. Marketers, designers, product managers, executives and other stakeholders often request descriptions of typical customers derived from segmentation or similar analyses. Often called personas or profiles, such descriptions are often produced by non-replicable, overfit, high-dimensional quantitative processes or qualitative methods (Chapman & Milham, 2006). Also, the descriptive dimensions are typically selected post hoc by the analyst and are not necessarily salient or particularly relevant to the user, and thus of little value for qualitative understanding.

It is typically impossible to say whether a persona is accurate, replicable, or descriptive of many—or any—customers (Chapman et al., 2008). Why? Because such a profile typically includes many descriptors; yet, as more descriptors are added, the proportion of users or customers who match the combined description will drop. This is one version of the "curse of dimensionality." Figure 6 shows this effect in several real and simulated data sets. In consumer data sets, a segment profile with 7 or more attributes is likely to be an exact match to *almost no one*.

Figure 6: Proportion of respondents who match the combination of values in a profile, as the number of variables in the description increases from 2–11 (Chapman et al., 2008).

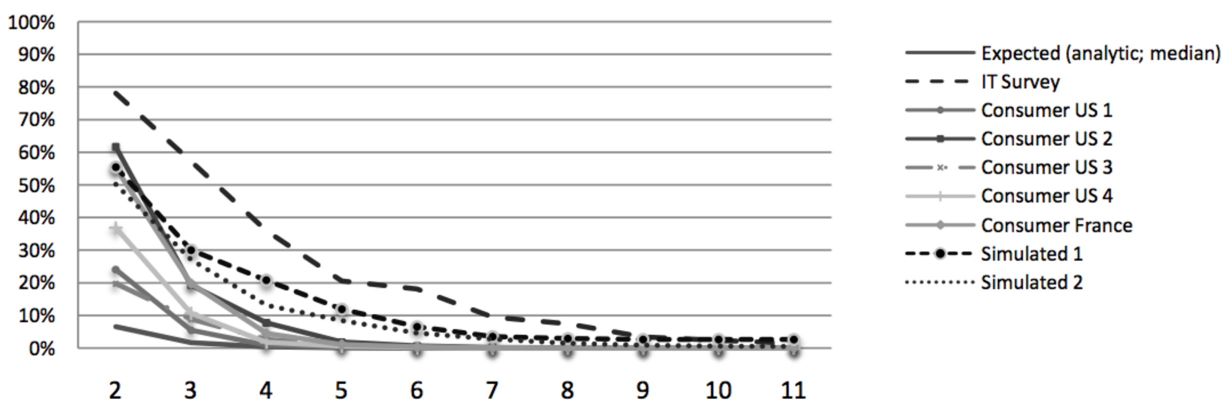

In Case 4, our research team was asked for user profiles, with regards to civic engagement, of adults in the US. We expected, on the basis of qualitative research, that there were many adults in the US who might be regarded as "interested bystanders," who are generally interested in civic information yet are not actively engaged in civic activities (Krontiris et al., 2015). The executive stakeholders wanted to understand these users better, and crucially, to know "How many interested bystanders are there?"

Case 4: Operational Research Question and Method. In light of the curse of dimensionality, our team sought a method that would do two things: (1) focus on identification from a *user's* point of view rather than our post hoc selection of variables, and (2) be less subject to the curse of dimensionality by assigning users probabilistically rather than categorically. We realized that choice-based conjoint analysis was exactly such a method (thanks to a suggestion from Greg Allenby, personal communication). We could find no prior example of psychographic segmentation using conjoint analysis, yet believed it was highly promising, and we differentiated it from general CBC by calling it "Profile CBC."

We selected attributes and levels to describe civic engagement, based on qualitative interviews conducted across the US. These were gathered into 8 attributes with typically 3 levels each, and used to form a CBC task asking users, "Which profile is more like you?" We found that the task was best constructed with 3 cards (3 profiles), without a "none" option, and with no more than 3 attributes shown in a partial profile format (Chrzan & Elrod, 1995; Patterson & Chrzan, 2003). Figure 7 shows an example task, as seen by a respondent (Chapman, Krontiris, & Webb, 2015).

Figure 7: Example task for Profile CBC, used for psychographic segmentation.

Thinking about civic or community engagement, which one of these PROFILES is more like YOU?

Choose *which profile is more like you* by clicking one of the buttons below:

	Profile 1	**Profile 2**	**Profile 3**
Career Involvement	I'm not working or in school right now.	My career or education is my main priority right now.	I balance my career or education with other obligations and pursuits.
Civic engagement (volunteering or community activity)	When I have free time, I spend it on civic or community activities.	I don't have time for civic or community activities.	I try to do as much civic engagement as I can, but I have other obligations.
Family involvement	I balance family time with career and social pursuits.	I don't spend very much time with my family.	I spend as much time with my family as I can.
	○	○	○

Note that these design options are *unlike* the typical recommendation for CBC tasks that involve products. The choices of partial profile and omission of the "none" options would be unusual and not recommended for most product choice tasks, yet we found them to be warranted for a psychographic profile task. Otherwise, the task was too complex for respondents.

After collecting the CBC data on respondents' identification with the randomized profiles, we used latent class analysis (LCA; Sawtooth Software, 2021) to identify and size the potential segments for civic engagement.

Case 4: Results. Figure 8 shows the result from LCA analysis, in which an optimal solution identified 6 segments, with estimated sizes that ranged from 11–23% of the sample in each segment. In answer to the research question, we identified 3 segments—"absentees," "issues-aware," and "vocal opinionators"—who jointly comprised 49% of the sample and matched the concept of interested bystander. This answered the core executive question as to how many potential users we would target with our engineering and design efforts.

Figure 8: The psychographic segments and their sizing, as found by Profile CBC.

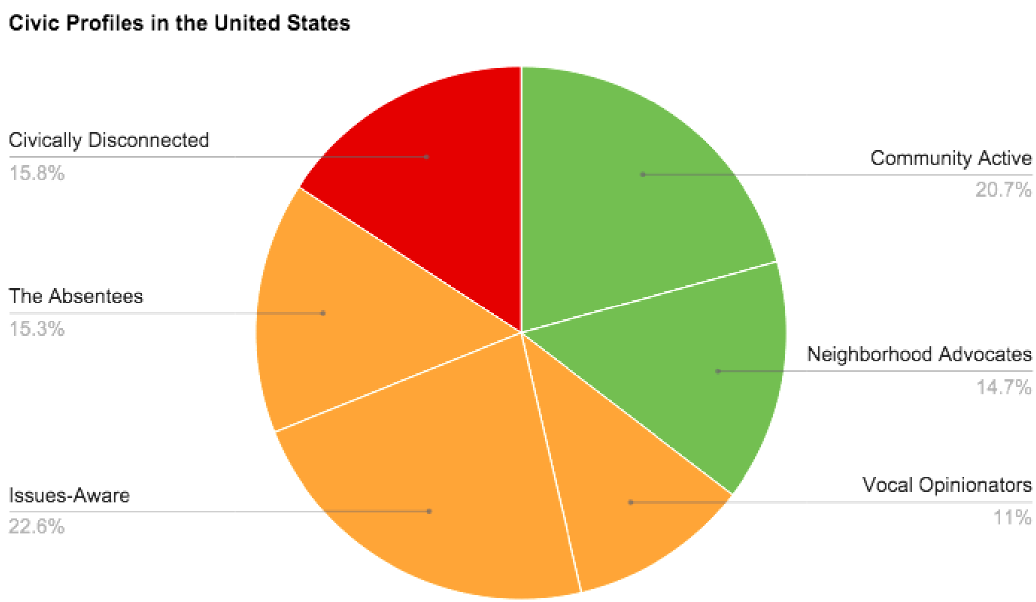

Additionally, we found that the segments were not only differentiated on the psychographic basis variables used for segmentation (i.e., the CBC attributes), but were also highly differentiated on other demographic and socioeconomic variables that were not used in the segmentation (for details, see Chapman, Krontiris, & Webb, 2015). In other words, the psychographic segmentation showed strong external validity with expected covariates and reported civic behaviors.

In short, the Profile CBC method achieved a useful result for psychographic segmentation. It answered the question of segment composition and size and did so on the basis of users' own reports about their identification rather than post hoc variable selection. Because the important identifiers were selected by respondents themselves, and applied through a probabilistic method (LCA, and conjoint analysis utilities in general), it was free from the most detrimental aspects of the curse of dimensionality. It also afforded the opportunity to explore options to recombine the segments on the basis of the underlying attitudes, i.e., to do "market simulation" in psychographic space.

Case 4: Keys for Success. There are two crucial aspects for Profile CBC: the attributes and levels must be appropriate, and respondents must be able to do the task. It is more difficult to specify appropriate attributes and levels than in a product-focused CBC, because the attributes are psychographic and attitudinal rather than a direct reflection of the product. The best approach would be a combination of qualitative and ethnographic research, plus any available baseline work on the prevalence and distribution of individual attitudes.

The other key concern is construction of the task such that it makes sense to respondents. In our trials, we found—again, *unlike* a product-focused CBC—that it was best to eliminate the "none" option and to force a choice among the cards; to limit it to no more than 3 cards at a time; and to use a partial profile approach with 3 attributes. Pre-testing with an in-person, think-aloud protocol is even more important than it is for product-focused CBC. For more design details and discussion of task format, see the full paper (Chapman, Krontiris, and Webb, 2015).

Conclusion

Taken together, the cases here demonstrate a "T-shaped" path for conjoint analysis practitioners, that extends both their breadth and depth of research offerings. With better respondent utilities (Case 1), consideration of competitive response (Case 2), and exploration beyond a product to an entire product line and brand (Case 3), skilled conjoint practitioners will be able to answer deeper, more important, and strategic questions accurately.

We have seen that conjoint analysis skills can also extend outside the product space, to conduct breadth research into consumers' attitudes, self-identification, and profiles (Case 4). I believe this "T" offers a much larger, more interesting, and more impactful space to inform decisions than simpler, single-point studies of product optimization (which remain highly important).

All four cases here demonstrated an arguably strong degree of external validity, ranging from successful market share prediction (Case 1), to anticipation of competitive responses (Cases 2 and 3), to alignment between psychographic and demographic variables (Case 4). I believe that such validation is possible only because of the attention to iteration within a product space, where the studies build on prior understanding of product features and attitudes. At the same time, the cases here—although novel—are not especially unique; I suspect that such success should be attainable in many product spaces, if one iterates and builds foundational knowledge.

In short, once one has developed expertise in product-focused CBC and the trust of executive sponsors, I highly encourage innovation! However, in such innovation, I strongly recommend building on existing best practices and methods (as with Case 1 and market simulation, or Case 2 and game theory) rather than creating novel statistical methods. In my experience, ad hoc tinkering with statistical methods is more likely to be a mistake than an advance.

If the reader is an R practitioner, I invite you to follow the development of our open source R package "choicetools" for conjoint analysis, MaxDiff, and related methods (Chapman & Bahna, 2019; Chapman, Alford, & Ellis, 2021). The package is in early development, and future releases will incorporate code for methods such as those in this paper (e.g., we will soon add code for the GA approach in Case 3, which is available separately today).

Finally, please consider sharing your cases, successes, and perhaps especially non-successes, as talks in future Sawtooth Software Conferences. None of the cases here would have been possible without the support, interchange, and inspiration my colleagues and I have received over the years from this community. It will make you a better practitioner!

Chris Chapman

References

Belloni, A, Freund, AR, Selove, M, & Simester, D. (2008), Optimizing product line designs: Efficient methods and comparisons. *Management Science* 54:9.

Chapman, C. (2013). 9 things clients get wrong about conjoint analysis. In B. Orme (ed), *Proc. 2013 Sawtooth Software Conference*. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/sawtooth-conference-2013.pdf

Chapman, C., & Alford, J. L. (2010). Product portfolio evaluation using choice modeling and genetic algorithms. In B. Orme (ed), *Proc. 2010 Sawtooth Software Conference*. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/dc4cc660-880a-4e98-9049-32a3beb935d8.pdf

Chapman, C., & Alford, J. L. (2010, December). A genetic algorithm system for product line exploration and optimization. In *Proc. 2010 Second World Congress on Nature and Biologically Inspired Computing* (NaBIC) (pp. 268–273).

Chapman, C.N., Alford, J.A., and Ellis, S. (2021). choicetools. R package for conjoint analysis, maxdiff, and composite perceptual maps. Ver 0.9076. [R code]. Available at https://github.com/cnchapman/choicetools.

Chapman, C., Alford, J. L., Johnson, C., Weidemann, R., & Lahav, M. (2009). CBC vs. ACBC: comparing results with real product selection. In B. Orme (ed), *Proc. 2009 Sawtooth Software Conference*.

Chapman, C., Alford, J.L,. and Love, E (2009). Exploring the Reliability and Validity of Conjoint Analysis Studies. Poster presented at the 2009 Advanced Research Techniques Forum (ART Forum).

Chapman, C., and Bahna, E. (2019). choicetools: a package for conjoint analysis and best-worst surveys. Presented at UseR! 2019, Toulouse, France. Presentation; GitHub package.

Chapman, C., Krontiris, K., & Webb, J. (2015). Profile CBC: Using Conjoint Analysis for Consumer Profiles. In B. Orme (ed), *Proc. 2015 Sawtooth Software Conference*. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/40d2acfc-5474-48a6-a81a-aafbf4ea2686.pdf

Chapman, C. & Love, E. (2012, September). Game theory and conjoint analysis: using choice data for strategic decisions. In B. Orme (ed), *Proc. 2012 Sawtooth Software Conference*. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/1f6470c2-ddbb-4783-88f5-8c9b9b4d9c8f.pdf

Chapman, C., Love, E., Milham, R.P., ElRif, P., & Alford, J.L. (2008). Quantitative evaluation of personas as information. *Proc. Human Factors and Ergonomics Society (HFES) 52nd Annual Conference*, New York, NY, September 2008. http://goo.gl/4rLYEO

Chapman, C., Love, E., & Alford, J. L. (2008, January). Quantitative early-phase user research methods: hard data for initial product design. In *Proc. 41st Annual Hawaii International Conference on System Sciences* (HICSS 2008) (pp. 37–37). IEEE.

Chapman, C., Milham, R.P. (2006) The personas' new clothes: methodological and practical arguments against a popular method. *Proc. Human Factors and Ergonomics Society Annual Meeting*, 50:5, 634–636. SAGE Publications. https://goo.gl/szQ54E

Choi, S.C. and W.S. Desarbo (1993). Game theoretic derivations of competitive strategies in conjoint analysis. *Marketing Letters*, 4 (4), 337–48.

Chrzan, K., & Elrod, T. (1995) "Partial Profile Choice Experiments: A Choice-Based Approach for Handling Large Numbers of Attributes." Presented at the 1995 Advanced Research Techniques Forum, Monterey, CA.

Dotson, J., Larson, J., Ratchford, M. (2012). Maximizing purchase conversion by minimizing choice deferral: examining the impact of choice set design on preference for the no-choice alternative. *Proc. 2012 Sawtooth Software Conference*, Orlando, FL. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/1f6470c2-ddbb-4783-88f5-8c9b9b4d9c8f.pdf

Ferguson, S., and Foster, G. (2013). Demonstrating the Need and Value for a Multi-Objective Product Search. *Proc. 2013 Sawtooth Software Conference*, pp. 275–304. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/sawtooth-conference-2013.pdf

Goldberg, D.E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley.

Huber, J. (2012) CBC Design for Practitioners: What Matters Most. *Proc. 16th Sawtooth Software Conference*, Orlando, FL, March 2012. http://goo.gl/ieqgMK.

Johnson, R.M., and Orme, B. K. (2007). A New Approach to Adaptive CBC (Whitepaper). Sawtooth Software, Provo, UT. https://sawtoothsoftware.com/resources/technical-papers/a-new-approach-to-adaptive-cbc

Karty, K. (2012). Taking nothing seriously: a review of approaches to modeling the "none" option. *Proc. 2012 Sawtooth Software Conference*, Orlando, FL.

https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/1f6470c2-ddbb-4783-88f5-8c9b9b4d9c8f.pdf

Krontiris, K., Webb, J., Krontiris, C., Chapman, C. (2015). Understanding America's "Interested Bystander:" A Complicated Relationship with Civic Duty. Technical report, Google Civics Research Workshop, New York, NY, January 2015. https://research.google/pubs/pub44180/

Love, E., and Chapman, C. (2007). Issues and Cases in User Research for Technology Firms. In B. Orme (ed), *Proc. 2007 Sawtooth Software Conference*. https://sawtoothsoftware.com/uploads/sawtoothsoftware/originals/dc4cc660-880a-4e98-9049-32a3beb935d8.pdf

Mebane, W.R., Jr., Sekhon, J.S. (2011). Genetic Optimization Using Derivatives: The rgenoud Package for R. *J. Statistical Software*, 42:11. http://sekhon. berkeley.edu/rgenoud.

Myerson, R. (1991), *Game Theory: Analysis of Conflict*. Harvard Univ. Press, Cambridge, MA.

Orme, B. (2014). *Getting Started with Conjoint Analysis: Strategies for Product Design and Pricing Research*, 3rd edition. Madison, WI: Research Publishers.

Orme, B, and Chrzan, K. (2017). *Becoming an Expert in Conjoint Analysis: Choice Modeling for Pros*. Sawtooth Software.

Patterson, M., & Chrzan, K. (2003). Partial Profile Discrete Choice: What's the Optimal Number of Attributes. *Proc. 2003 Sawtooth Software Conference*. https://goo.gl/uNqjTt.

R Core Team (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Sawtooth Software (2021). The Latent Class Technical Paper (version 4.8) (whitepaper). Sawtooth Software, Provo, UT. https://sawtoothsoftware.com/resources/technical-papers/latent-class-technical-paper