Reflections on a Career as a Marketing Engineer

John R. Hauser, MIT Sloan School of Management, Cambridge, MA 02139

I was once an Editor-in-Chief for *Marketing Science* and founded (with Bob Klein) a company known as Applied Marketing Science Inc., but I am at my core a marketing engineer. I am driven to explore new ways to solve important and relevant problems. Along the way, I've written a few theoretical papers, but they were always driven by a desire to understand how things worked so managers could better serve customers (and make money doing it).

Throughout fifty years of solving problems and writing papers, the methods that we've used have evolved. We can do so much more with today's computers and today's mathematics than what we could do when we did Bayesian calculations by hand or estimated logit models on slow computers with almost no working memory. But better tools do not mitigate the need for human insight. Success, at least for the foreseeable future, will still require creative students and researchers to identify challenging problems, solve those problems with the best-available, and find ways to apply and diffuse the solutions and ideas.

I prefer to look forward, but we can look to history to gain insight on how ideas for research originate. I wish there were an easy answer, but there is not. My best (and worst) ideas have come from a variety of sources including talking to managers, attending seminars and conferences, teaching new courses, talking to colleagues, working with collaborators, and even reading the literature. It's hard work. We are bombarded with good problems, but we have to recognize them. Listening is important.

I get bored easily. I seek problems that excite me and others. When I hear about an interesting problem, I talk to other people. Do they find the problem exciting or boring? Do they have ideas to solve the problem—usually in the form of "did you read X?" Often a collaborator (and this includes students) bring a problem to me and it sounds interesting. We try it on for size, try to determine if it has the potential to excite, and research it enough to become convinced it can be solved. Sometimes those resources paper and pencil are sufficient to solve the problem; sometimes new methods and faster computers are necessary; sometimes data; and sometimes we need a partner firm to implement (or experiment with) proposed methods.

I have a few criteria. The problem should be new and not previously solved. Fortunately, there are many such problems. Methods must exist, or can reasonably be developed, to solve the problem. (MIT culture instills a certain arrogance about the ability to develop or learn new methods.) I seek talented and motivated collaborators. Glen Urban once said that the best part of MIT and the marketing profession is the ability to learn from students and other researchers. Most exciting problems are just too big to solve alone.

I stay away from popular topics. I learned this from John Little. It's better to work at the beginning of the S-curve where you can make substantial progress. It's less rewarding to work on the

upper, flat portion of the S-curve where you need to do substantial work to make incremental contributions. Besides, popular problems are burdened with dogma—you have to speak the language, use accepted methods, model the "right" phenomena, care about the "right" issues, and avoid challenging core beliefs. Occasionally, you can challenge the dogma and restart the S-curve, but the fight is not easy. I have many scars.

Enough with philosophy. Let's examine a few examples.[1] You can take from these examples any interpretations that help you.

Website Morphing. Glen Urban is one of the most creative people I know. If you get a chance, take a look at his book on his sculptures or his reflections on sailing. (Yes, he almost drowned me off New Harbor ME, but I am a better sailor because of the experience.) Glen had been talking to the BT Group about their website and had the vision that he could morph the website based on consumers' cognitive styles. His vision was a continuous morph—perhaps you've seen one face morph into another. Morphing excited us (and our collaborator Gui Liberali). Besides, BT Group promised to fund the research and run experiments.

But changing a website continuously proved a difficult nut to crack. Besides, it was likely that continual changes would give consumers whiplash. One day, after talking to Glen and Gui, I took a walk across campus just to think. I stopped at the Electrical Engineering and Computer Science (EECS) Department which is located in a Frank-Gehry-designed building known as the Stata Center. Warren Seering told me that MIT students call the Stata Center earthquake-robust. By this they mean that If there were an earthquake in Cambridge and the Stata Center toppled, you could not distinguish the before photo from the after photos. I'm no Frank Gehry, but perhaps the solution to website morphing was to redefine the problem. Perhaps we could solve a problem with a relatively few morphs and a relatively few cognitive styles. I had been reading Gittins' book on multi-arm bandits. He had an elegant solution that was relevant to a small number of actions. Many challenges remained, including Bayesian updating with uncertain cognitive-style assignment, but together the three of us developed the theory, implemented the theory, and published the paper. But the seed of the solution came from an earthquake-robust building.

Defensive Marketing Strategy. Laboratory test markets (LTMs) continue to be one of the great successes of marketing science. Few consumer product firms would launch a product without an LTM. LTMs are themselves a major industry and have been extended to durable goods and virtual reality. As junior faculty, Glen Urban and Gerry Katz gave Steve Shugan and me access to historic data on applications of the first LTM – the Assessor model developed by Glen Urban and Al Silk. We were impressed with a stylized fact in database. Most LTMs were funded not by the innovator, but by a

[1] These are but a few examples, Others, not mentioned in this essay, include work with Rene Befurt, Don Clausing, Min Ding, Songting Dong, Daria Dzyabura, Siham El Kihal, Rosanna Garcia, Marat Ibragimov, Frank Koppelman, Jim Orlin, John Roberts, Alice Tybout, Bruce Weinberg, Ken Wisniewski, Michael Yee, Florian Zettelmeyer, Juanjuan Zhang, and a variety of valued collaborators. For each there is a great fireside tale.

competitor who would be impacted by the innovator.

Back then, you had to present a paper in order to get funding to travel to a conference. The INFORMS (then ORSA/TIMS) conference was in Hawaii. Steve and I needed a topic. We proposed a decision support model to help competitors impacted by innovators. We called them defenders. Our first attempts were in the tradition of John Little's decision calculus, but Steve noticed that, independently of the parameters of the model, the recommendation was almost always the same. This was an "aha" moment for me. Not only was the topic new, but there were no papers that we knew in marketing that were based purely on theory. Applications and data were minimum requirements for publication.

We worked hard. The pricing theorem alone took many person-months of effort—we got to know arctangents like the back of our hands. We finally were able to prove all of the theorems and had a rudimentary equilibrium argument. We it off to *Marketing Science*. The editor was supportive, but we were challenged by reviewers on the lack of data. We went through multiple revisions while we convinced the editor and reviewers that theory papers could be published. It took a deep dive into the history of science. I had just gotten married when we began work on "Defender." Steve was married, but not a father. By the time the paper was published, Steve and Irene had three children; Marija and I had two. To be fair; the final paper was much superior to the first draft.[2]

Fast Polyhedral Adaptive Conjoint Analysis. Duncan Simester and I were working with a US Army Research Development and Engineering Center (RDEC)—the RDEC was responsible for nutrician and clothing. It invented meals-ready-to-eat for the US Army. It was clear that the scientists and engineers at the RDEC were dedicated professionals who believed in their mission. They were motivated more by the non-monetary aspects of compensation than the monetary aspects. The RDEC asked us to optimize those non-monetary aspects. Cool. Conjoint analysis applied in a new way.

Upon looking at the data, it was clear that respondents became bored with the conjoint analysis tasks. After a certain point, fit degraded and the estimates became less precise. Surely there must be a better way. We had to get the most information possible from each question asked.

Olivier Toubia was at the MIT Operations Research Center and was taking a course on math programming by Rob Freund. Duncan, Olivier, and I had a visual geometry model of how to ask questions efficiently, but it was Rob who suggested that we use polyhedral methods and interior-point solutions to solve the geometry. Meanwhile, Ely Dahan had joined MIT and was developing some of the first web-based, highly-visual, interactive questionnaires. He had contacts at Timbuk2, who wanted to help us apply the methods. (We also had funding from the National Science Foundation for

[2] Later, with Steve Gaskin, we applied Defender with real data. Managers did not understand probability density functions, so Steve G explained them with Smurfs and Legos. On one project, we held our ground when our model predicted a new product should not be launched. (Yes, coming full circle, Defender was now being used by an innovator.) We were fired, but ultimately vindicated. The client was not happy with its new consultant.

programming and a very ambitious incentive-aligned test of the methods.) We developed the methods, published the theory and application, extended and applied polyhedral methods to adaptive choice-based conjoint analysis, and then to probabilistic interpretations. We published two additional theoretical papers on experimental design and endogeneity.[3]

Identifying Customer Needs from User Generated Content. The voice of the customer (VOC) has a long history in marketing and in industry.[4] It has helped many firms design products and services that consumers actually want to buy. But VOC applications are labor-intensive. On the other hand, user-generated content (UGC) is ubiquitous and, mostly, free. Consumers evaluate products all of the time. A variety of firms contacted me asking if I knew some way to use UGC for VOC applications.

We started to explore methods to do so. Latent Dirichlet Allocation (LDA) and its related method LDA HMM (Hidden Markov Models) was popular in machine learning and marketing and seemed to be a great solution. LDA identifies bundles of words (called "bags of words") that are close together in sentences and, perhaps, represent the same construct. We tried LDA. It wasn't bad, but the managers to whom we showed the word bundles did not agree. Besides, word bundles do not provide the rich context of VOC customer needs. They don't solve the managerial problem.

Artem Timoshenko saw deeply into the problem and proposed using advanced machine learning methods to uncover rich and complex customer needs. Initially, we failed. Despite the power of deep learning and despite the almost mystical properties of word embeddings, the "machine" could not automate the entire process. Artem's true vision was to recognize that parts of the VOC process could be automated and parts needed human judgment. The resulting hybrid is substantially better than either a machine-alone or a human-alone system. We applied the hybrid methods using donated professionals from Applied Marketing Science, Inc. and compared our method to a human-only system. The hybrid system continues to be applied and continues to help firms develop better products.

I am gratified that papers are now appearing to automate more steps in the VOC process. Some work well and some remain works in progress, but the field advances. Ultimately, it will be almost automatic for any firm to query UGC as a virtual customer.[5]

---

[3] The editor of this volume asked me to articulate a disappointment. In the endogeneity paper, we knew the full set of adaptive conjoint-analysis questions was finite. Typical applications involved a very small number of questions, hence we did not rely on asymptotic methods. Later, we were vigorously criticized by a published paper that solved a related problem using asymptotic methods. The endogeneity issue remains in the originally-defined problem. I never got around to writing a reply. Similarly, I once used conjoint analysis to compute willingness to pay, not a market equilibrium, for a patented feature. Matt Selove's thesis had earlier demonstrated that a conjoint-analysis-based market-equilibrium calculation was neither practical nor stable for realistic problems. But at least Matt, Felix Eggers, and I were able to publish a later paper exploring the issue.

[4] I would be remiss without mentioning Abbie Griffin, a pioneer in the voice of the customer, or Bob Klein and Applied Marketing Science, Inc. who perfected the application of the voice of the customer.

[5] Ely Dahan, before his untimely death, had the vision of a virtual customer, that is, a database that could be queried to gauge reaction to potential new products. We are getting closer every day. Alex Burnap, in particular, along with Artem are getting extremely close on a number of dimensions.

How an Engineer Came to Marketing. As an undergraduate, I was a student in Electrical Engineering at MIT and spent two summers working at General Electric's space center in Valley Forge, PA. I liked EECS, but I did not like getting to the lab at 7 am and I did not like working without windows in a secure facility. When I returned to MIT, I because interested in public transportation (similar math, by the way).[6] My first ever paper, available on my website, is an algorithm to route dial-a-ride buses. We were working on a US Department of Transportation grant to implement dial-a-ride in a "demonstration project." The algorithm worked, but ridership was poor. As the most-junior author, I was tasked to run a survey. I sought advice from John Little, then head of the Operations Research Center. (I had moved from EECS, to transportation engineering, to OR—kind of a nomad before being a nomad was Academy Award worthy.) The survey made it obvious why consumers were not riding the system—the algorithm, which minimized average travel time, assigned "tours" by putting hard-to-serve riders at the "end of tour." Some riders, hoping for a short trip, sat on the bus seemingly forever. Nigel Wilson, one of my many mentors, recognized that the end-of-tour problem was due to the objective function in the algorithm. We changed the objective function to overweight long tours (minimize the square of travel time). Ridership improved—there was something to this marketing.

John Little encouraged me to speak with Glen about a project to design MIT's HMO using measures of consumers' perceptions and preferences. At the time, I was inspired by Frank Bass and Jerry Herniter in marketing, by Moshe Ben-Akiva in Civil Engineering, and by Dan McFadden in transportation economics. Demand modeling, as it is called in transportation planning, was interesting, fit my skills, and, at the time, was on the beginning of the S-curve. Working with Glen led to my ScD thesis supervised by John, Glen, and Moshe.[7] We ultimately expanded the paper to a text on the *Design and Marketing of New Products*.

That year I was late on the marketing market, seeking a job in spring with the rest of the operations research students. Glen's advisor, Phil Kotler, called Glen because Northwestern University had a slot that was joint with marketing and transportation. Glen told Phil I was the best student he had ever had, neglecting to mention I was the only student he had ever had. (Glen was creative as always.) I interviewed, Northwestern hired me, and I spent five wonderful and productive years at Northwestern with colleagues whom I respect and value.

[6] I was motivated, in part, by a desire to "give back" and design public transport that served consumers.
[7] My thesis used von Neumann-Morganstern (vNM) utility theory rather than conjoint analysis to measure preferences among product features. vMN theory is particularly good at modeling risk. Although I published a few papers on vNM theory with Glen and later with Josh Eliashberg, conjoint analysis proved the more popular method.