HOW DO SUCCESSFUL SCHOLARS GET THEIR BEST RESEARCH IDEAS?

AN EXPLORATION

Cathy Cao, Xinyu Cao, Matthew Cashman, Madhav Kumar, Artem Timoshenko, Jeremy Yang, Shuyi Yu, Jerry Zhang, Yuting Zhu, and Birger Wernerfelt.[1]

September 5, 2019

Abstract

We interview 24 marketing professors to ask how they got the ideas for 64 of their papers. More than three quarters of the papers were inspired by holes in the literature, by a "stylized fact" that the current literature cannot explain, or by an interaction with a manager. The rest fall into several smaller categories that to a large extent can be seen as special cases of the three big ones. We describe how papers from each of the three big categories help move the literature forward. We also illustrate the range of situations contained in each category by way of several examples. Among the authors we interview, most do not use a single source. As these authors become more senior, managerial contacts play an increasing role, while the balance between literature and stylized facts appears to be unchanged.

Key words: PhD students, Research ideas, Marketing academia

"I don't know what the PhD students are doing in their third year, but it takes at least a year."

(MIT professor)

I. INTRODUCTION

Many PhD students have trouble finding good dissertation topics and some end up taking the better part of a year to get going on their job market papers. Since faculty can offer few principles in the area, PhD students often start generating research ideas with little or no help from the experiences of others. The purpose of this paper is to uncover and collect the experiences of some successful authors in the hope that we, and other PhD students, may learn how they identify successful research topics. A key finding of our study is that successful authors, despite their apparent heterogeneity in work habits and styles, get the majority of their ideas from the same very small set of sources.

We interviewed 24 marketing faculty members, asking about the origins of two, three, or four of their most important papers. We ended up with information about 64 papers. While the research had a wide variety of origins, there were several common themes. In particular, we were able to classify the inspiration for 63 papers into one of seven categories, with most falling in the top three: 41% were inspired by holes in the literature, 24% were motivated by a "stylized fact" that the current literature cannot explain, and 13% originated from interactions with a manager. Literature-driven papers improve on our knowledge of topics that are partially – but imperfectly – understood at the time of writing; papers in the stylized-fact category bring the literature in closer alignment with known facts; and papers inspired by contacts with managers help us focus on important problems. Phrased differently, papers in these three categories help the academic marketing literature become more complete, correct, and relevant.

Surprisingly few articles have addressed the question of research ideas and those we found have typically focused on a single source or a set of very similar sources. However, all three of our main sources, or something very close to them, have previously been identified in individual articles. The sociologist Davis (2001) proposes reading papers outside of one's primary area, looking at working papers by the leaders of a narrow field, attending seminars, and actively communicating with colleagues. Following these recommendations can predictably result in literature-driven ideas. In economics, Pischke (2009) proposes reading newspapers and talking to non-economists as two ways to identify unexplained stylized facts. In marketing, the empirical generalization program (Bass and Wind, 1995) focuses exclusively on the literature (in fact it is a requirement that there are at least two prior papers on

a topic). Finally, Roberts, Kayande, and Stremersch (2014) study highly cited papers that have been used for decision support by marketing managers and, perhaps unsurprisingly, find that most of them were inspired by managerial input.

The present paper stands apart from previous work in three ways. First, instead of focusing on a single source of research ideas, which implicitly suggests that any other sources are unimportant, we identify and compare several very different ways to get ideas. Second, we offer a number of examples to illuminate what it actually means to take inspiration from these different sources. Third, we link our findings to theories of scientific progress. As far as we can tell, ours is the first paper to offer a systematic, though preliminary, look at an issue of major importance.

After a brief description of our data collection methodology, we define eight sources of inspiration and illustrate each with several examples. We then anchor the findings in the problem-solving-based conceptual framework of scientific progress. Finally, we look at a couple of regularities in the behavior of authors.

II. METHODOLOGY

Our classification system and examples are based on interviews with 24 successful marketing academics. The whole process took a semester. We started by jointly conducting two-hour interviews with four subjects (two in the quantitative area, one behavioral, and one managerial). These were used to develop a tentative set of categories and an interview guide. The categories that ended up being important were not hard to identify; they were explicitly described – and even named – by several subjects. (One more category – papers aiming to understand a new phenomenon – was identified as we interviewed more authors, but we tried to balance the desire to retain information with the need to get reasonable cell sizes.)

Once we had an agreement on the initial set of categories and the interview guide, each student proposed two or three authors and two or three papers by each of those. The group discussed each proposal. In general, the selection of authors was based on overlapping research interests: We chose authors who had written the kinds of papers that we would like to write, typically coinciding with some of each author's most cited works. The idea was to look at papers that excited the youngest generation of scholars and thus were interesting to current PhD students and more likely to be interesting to future students. Given the strong vintage effect in research interests and methods, simply using citations or representative

topics would be backward-looking and give us less useful information.

A drawback of this way of selecting the sample is that it reflects the interests of nine MIT students rather than the "average current PhD student," meaning that some subfields are underrepresented. So while we find pervasive patterns that make sense, our analyses are best seen as exploratory and we have to be very careful about what conclusions we draw.

Respondents were contacted by email and interviewed by the student who selected them. Primarily, the interviews took place by phone, but a couple of respondents who were too busy to schedule phone interviews were instead sent a questionnaire over email. We then followed up by asking clarifying questions in later emails. One respondent was interviewed in person. All interviews were completed about six weeks after the initial contact. The interviews were focused on two or three specific papers and the students asked respondents to talk about how each of those came to be. The respondents were very generous with their time (the interviews took an average of about 45 minutes) and many found the question interesting.[2] We met as a group after each interview, and the interviewers summarized what was said about each paper. The papers were then classified by consensus.

III.   **EIGHT SOURCES OF INSPIRATION**

(1) Papers inspired by holes in the literature.

Papers in this category are motivated by gaps in the marketing literature. We have 26 examples of these and only in one case does the idea come from a single article. In every other case, the author got the idea after reading a lot of literature on a topic. Interestingly, not a single author mentioned finding inspiration in the "future research" section of a published article, perhaps casting doubt on the conventional inclusion of such sections.

For theoretical papers, this category of research typically introduces an assumption that is compelling, yet new to the literature. For example, Iyer was inspired to write "Coordinating Channels under Price and Non-Price Competition" after reading lots of papers on coordination in distribution channels and noticing that they all assumed that the same contract is offered to every retailer. Chen, Narasimhan, and

---

[2] The interviewed authors are listed in the Appendix and the interview guide is reproduced in the web appendix.

Zhang wrote "Individual Marketing with Imperfect Targetability" after noticing that all existing literature assumed perfect targetability.

Empirical papers in this category may improve on the statistical properties of earlier papers, check a widely held assumption, or test a theory. For example, Golder and Tellis wrote "Pioneer Advantage: Marketing Logic or Marketing Legend?" after realizing that the data used in earlier literature suffered from survival bias. Chevalier and Mayzlin took a fresh look at "The Effect of Word of Mouth on Sales: Online Book Reviews" because the literature at the time had focused on correlation rather than causation, and Montgomery, Li, Srinivasan, and Liechty got the idea for "Modelling Online Browsing and Path Analysis Using Clickstream Data" after reading the literature on clickstreams and realizing that it was descriptive rather than predictive. Finally, Godes and Mayzlin wrote "Firm-Created Word-of-Mouth Communication: Evidence from a Field Test," to test predictions developed in an earlier, theoretical paper, also written by them.

In two cases, one author became aware of an opportunity during his PhD program but only pursued it several years later with a coauthor. This process led to "Endogeneity in Brand Choice Models" by Villas-Boas and Winer and "Organizing for Radical Product Innovation: The Overlooked Role of Willingness to Cannibalize" by Tellis and his PhD student Chandy.

(2) Papers motivated by a "stylized fact" that cannot be explained by the current literature.

We define a *stylized fact* as phenomenon that is observed with some regularity. Authors are inspired to write this category of papers when they observe, or otherwise learn about, a stylized fact that current literature cannot explain. Papers in this category often propose one or more possible explanations and sometimes support them with data. In a sense they represent a scientific reaction to an apparent falsification of current theory.

In many cases, the basis for the paper comes from personal experience. For example, Dukes, Geylani, and Srinivasan were inspired to write "Strategic Assortment Reduction by a Dominant Retailer" when they felt that some large retail stores seemed to be reducing their assortment, and Norton, Mochon, and Ariely wrote "The IKEA Effect: When Labor Leads to Love" after noticing that even relatively unskilled origami makers were very protective of their creations. "Prominent Attributes under Limited

Attention" by Dukes and Zhu was based on personal shopping experiences, while Dunn, Aknin, and Norton got the idea for "Spending Money on Others Promotes Happiness" after noticing that many people talk about feeling good about helping.

Articles in the popular or scientific press are another source of inspiration. For example, Godes and Mayzlin wrote "Using Online Conversations to Study Word-of-Mouth Communication" after seeing one article about why TV shows were cancelled and another on how some conversations about cars suddenly attracted a lot more participants (became "buzz").

The inspiration can also come from a seemingly anomalous result in previous research. For example, Anderson and Simester got the idea for "Reviews without a Purchase: Low Ratings, Loyal Customers, and Deception" when they worked on a different project about consumer reviews and found that many reviews could not be matched with purchases, and Frederick wrote "Overestimating Others' Willingness to Pay" after finding the result in an in-class experiment.

It is not uncommon for the triggering observation to be brought out in conversations with peers. In fact, the authors of several papers in this category mentioned that they had conceived the idea when discussing the apparent mystery with one or more colleagues. For example, Bronnenberg, Dhar, and Dube wrote "Consumer packaged Goods in the United States: National Brands, Local Branding" after making an observation about an unrelated problem and then discussing.

(3) Papers prompted by contacts with managers.
This category consists of papers that were written after a question from one of the author's managerial contacts. The managers are looking for ways to improve their businesses and the papers follow the engineering tradition by proposing new and better ways to do or measure something.

Some papers in this category answer quite specific questions. For example, Rutz and Bucklin came up with the idea for "From Generic to Branded: A Model of Spillover in Paid Search Advertising" after a manager wanted advice on the appropriate expenditures on generic versus branded keywords. Ghose and Yang were inspired to write "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets" after a question about valuation, and Anderson, Lin, Simester, and Tucker wrote "Harbingers of Failure" after a manager asked about identifying customers who are predictive of product failures.

Another set of papers with managerial roots is concerned with whether a specific effect is something that could be bad or perhaps exploited in the future. For example, Goldfarb and Tucker wrote "Privacy Regulation and Online Advertising" after a business executive asked one of them about it, and "Exercise Contagion in a Global Social Network" by Aral and Nicolaides is a response to a question a manager asked about one of their other papers on peer effects.

(4) Papers motivated by teaching.

Papers in this category come in two subcategories. Some are initiated by a student question and others come about because the author, as an instructor, realizes that they cannot explain a particular point in a fully satisfactory way. As an example of the second subcategory, Chen, Koenigsberg, and Zhang came up with the idea for "Pay-as-You-Wish Pricing" when one of them, while teaching a case about it, felt the need for a theory describing the conditions under which pay-as-you-wish pricing is profitable.

(5) Papers resulting from a desire to apply a new research technology.
These "tool-driven" papers are results of an author looking for marketing applications of a research technology that has not seen previous use in marketing. Most of these involve statistical techniques, optimization techniques, or experimental procedures. For example, Horton, Rand, and Zeckhauser wrote "The Online Laboratory: Conducting Experiments in a Real Labor Market" to introduce and validate online labor markets as a tool for research.

(6) Papers which come about when the authors get access to a new data set.

Papers in this category come to be when an author discovers, and gets access to, a new data set, which allows testing of a theory that previously could not be tested or only could be tested in a less appealing way. The new data often contain natural experiments and other features that allow identification. In many cases the discovery involves a significant element of coincidence or the kind of luck you have when you keep your eyes open. For example, Zhang wrote "The Sound of Silence: Observational Learning in the US Kidney Market" after an advisor made a suitable dataset available. Similarly, Van den Bulte was at a conference when he met a doctoral student who had access to a dataset on over

10,000 anonymized customers over three years from a leading German bank – a data set that Van den Bulte went on to use in "Referral Programs and Customer Value" (with Schmitt and Skiera).

(7) Papers aiming to understand a new phenomenon.

Today, these papers typically look at aspects of the digital economy. One might think of them as similar to those in the "stylized fact" category, but the difference is that these papers seek to understand the basics of how the new phenomenon works and what it can do for marketing. For example, "Click Fraud" by Wilbur and Zhu came about when the authors read about it in the news.

(8) Papers that introduce a new research technology.

There is only one paper in this category. It is different than the "applying new research technique" category, since in this case the author developed the technique. As far as we can tell, this is a very unusual situation and we will henceforth omit this paper from the discussion, leaving us with 63 papers in total.

IV.   MAIN FINDING AND POST HOC CONCEPTUAL FRAMEWORK

The relative frequency of the first seven categories is shown in Table 1 below.

Table 1

Frequency of Research by Source of Inspiration

	Literature	*Stylized Fact*	*Managerial Contacts*	*Teaching*	*New Technique*	*New Data Set*	*New Phenomenon*
# of Papers	26	15	8.5*	2	3	5.5*	3
% of Papers	41	24	13	3	5	9	5

*As mentioned earlier, a small number of papers were split between two categories.

The main message of the table is that *the first three categories, papers inspired by holes in the literature, papers motivated by a stylized fact that the current literature cannot explain,* and *papers conceived through contacts with managers* are the most common by far.

Relative to our field as a whole, behavioral research is underrepresented in the sample, accounting for just nine of the 64 papers. It is clear that our findings speak less to behavioral than to quantitative research and it is hard to judge how much less. A referee made the reasonable conjecture that behavioral research is less likely to be inspired by input from managers, implying that Table 1 overestimates the general importance of those. Consistent with this, none of the nine behavioral papers are classified in the "managerial contacts" category. Instead, four are literature-driven, four are inspired by stylized facts, and one springs from new technology. So while the numbers are small, there is no strong evidence that behavioral work is inspired by different sources than quantitative work (probably apart from the managerial category). We also doubt that the underrepresentation of behavioral research has caused us to miss any important sources of inspiration. However, it is likely that the relative importance of different sources would be different in a sample with more behavioral papers.

On the other hand, a similar study, conducted by another group of MIT PhD students in 2009, found very similar results in spite of the fact that about a third of the papers were behavioral. Specifically, the authors looked at 61 papers and found that literature-driven papers, stylized facts, and managerial contacts, in that order, were the three largest categories (Banker, Fong, Nguyen, Nistor, Selove, and Silinskaia, 2009).[3] The percentages in the 2009 (2018) study were 49 (41), 43 (24), and 13 (13). So two different sets of "raters" looking at two different samples came up with similar results.[4] This gives more confidence in both the absolute and relative importance of the three top categories.

The result nicely fits American philosopher Larry Laudan's definition of scientific progress.[5] Specifically, Laudan (1977) extended the Popper–Kuhn–Lakatos line by arguing that science progresses by solving what at the time are considered important "problems." These problems are of two kinds: Empirical problems are "anything about the natural world which strikes us as odd or otherwise in need of explanation." (p. 15), and conceptual problems are "higher order questions about the well-foundedness of the conceptual structures (e. g. theories) which have been devised to answer the first-order questions." (p. 48). Some problems are seen as more important by others, to a large extent as a function of extrascientific forces. Theories that solve more important problems are more likely to be preferred.

---

[3] None of the 19 authors were included in the present study.
[4] The 2009 paper did not consider categories (7) and (8).
[5] Interestingly, Laudan is also cited extensively in Anderson's seminal 1983 article on the scientific status of marketing.

Seen in this light, papers inspired by holes in the literature are solving conceptual problems, papers motivated by stylized facts are taking on empirical problems, and inputs from managers influence the relative importance of different problems. So we will argue that *our top three categories are exactly those one would expect to see.*

The above argument raises a question about why we have papers in the other four categories: teaching, new techniques, new data, and new phenomena. It is, however, possible to argue that these categories can be viewed as combination of special cases of the first three categories and are more or less related to problem-solving.

Teaching: The "student" who is asking questions could also be a manager or scholar seeking answers to either conceptual or empirical problems and faculty seeking to complete lectures; the questioner is clearly trying to solve a problem.

New techniques: These papers use work from another field to solve marketing problems.

New data: Access to new data is only useful if the scholar in question already has identified a problem that they would like to solve, presumably from one of the other sources.

New phenomena: Papers in this area can be seen as serving the same function as those motivated by stylized facts. They increase the scope of the literature by explaining odd observations of the "real world."

V.   COMPARING AUTHORS

Since our sample is anything but random, we cannot draw any statistical inferences about papers from the different categories. However, if we are willing to assume that our 24 authors form a representative sample of "successful authors," we can report a couple of tentative findings about them.

First, many of us thought that individual authors would have a consistent "style" in the sense of repeatedly using the same source of inspiration. To look at this, we cross-tabulated how each of our authors got the idea for two of the papers we interviewed them about. (For authors who talked about more than two papers, we arbitrarily chose to focus on the two oldest.)

Table 2

Source of Inspiration for Oldest and Second Oldest Paper*

Source of Inspiration	Literature-Driven	Stylized Fact	Managerial Contacts	Rest
Literature-Driven	6	2	0	3
Stylized Fact	3	2	1	1
Managerial Contacts	0	0	1	1
Rest	1	1	1	1

*Row refers to the oldest paper and column to the second oldest.

Only 10 of our 24 authors are on the diagonal, indicating that they used the same source for these two papers. Since random choice according to the distribution in Table 1 yields a mean of 7+, there is *no strong evidence to suggest that individual authors confine themselves to a single source of inspiration.*

A second appealing conjecture, proposed by an interviewee in the 2009 study, is that more senior authors rely less on the literature and are more inspired by stylized facts. We looked for this by measuring seniority by the number of years between an author's PhD graduation and the publication date of a paper. Table 3 gives the mean "author PhD age" of papers in each of the three big idea generation categories.

Table 3

Mean Author PhD Age by Source of Inspiration

Literature-Driven	Stylized Fact	Managerial Contacts	Rest
8.0	9.8	15.8	6.3

The variances are large, so any *evidence that authors rely less on the literature and more on stylized facts as they become more senior*, is at best extremely weak. On the other hand, there is, not surprisingly, a tendency toward *a larger role of managerial contacts among more senior scholars*.[6]

---

[6] Recall that the article by Roberts et al. (2014), which essentially looked at papers inspired by managers, found that their authors had a lot of seniority.

V. CONCLUSION

We wrote this paper to help aspiring researchers generate successful ideas. While our authors are quite different, most of them used one of very few ways to find inspiration: In particular, the evidence suggest that PhD students and other young scholars simultaneously should pursue ideas in two places. One, read the literature in your subfield (go to seminars/conferences) aiming to identify holes in the overall conceptual structure. Some holes cannot be filled at the moment, but those that can will automatically be of interest to the field. Second, scan the business press for events you find puzzling, claims about regularities in seeming conflict with theory, and unexplained successes and failures. Upon further scrutiny, you may find that some of these are covered in the literature, but those that are not will give you your research topic if you can identify them.

As mentioned throughout, the study is an exploratory cut at a very important topic. We are only looking at the best papers by the most successful authors, our sample is not representative, and we are not even defining "best" and "success." A more comprehensive study would correct these three limitations.

First, it would include the best papers of less successful authors and the less good papers by authors with varying degrees of success. Second, it would involve a much larger and more representative sample of authors and articles. Third, it would formally define formal measures of the qualities of papers and authors, presumably by age-adjusted citation counts.[7] The last two steps would be difficult. Since the worst papers may be unpublished and the least successful authors may have left the field, it is likely that a realistic study would involve "good, but not great" papers and authors with some, but not overwhelming success.[8] So it might be hard to get enough variance to run a regression analysis.

The construction of a representative sample will also be a challenge, as this will require identification of a typology for both authors (behavioral, economic, statistical, etc.) and articles (lab experiment, field experiment, causality structural econometrics, other causality identifying tests, etc.). On the other hand, once the categories have been identified, it may be possible to collect the data through a questionnaire, rather than the much more expensive interview method.

---

[7] Although we did not set out to follow these criteria, the outcome of our process looks almost as if we did.
[8] It may be also be hard to recruit authors if they are to be used as examples of "what not to do."

REFERENCES

Anderson, Eric, and Duncan Simester (2014), "Reviews without a Purchase: Low Ratings, Loyal Customers, and Deception", *Journal of Marketing Research* 51: 249-69.Anderson, Eric, Song Lin, Duncan Simester, and Catherine Tucker (2015), "Harbingers of Failure", *Journal of Marketing Research* 52: 580-92.

Anderson, Paul F. (1983), "Marketing, Scientific Progress, and Scientific Method", *Journal of Marketing* 47, No. 4, pp. 18-31.

Aral, Sinan, and Christos Nicolaides (2017), "Exercise Contagion in a Global Social Network", *Nature Communications* 8: 14753.

Banker, Sachin, Nathan Fong, Han Nguyen, Cristina Nistor, Matthew Selove, and Daria Silinskaia (2009), "Where Do Successful Research Ideas Come From? Young Scholars Interview Experienced Scholars", Working Paper, MIT.

Bass, Frank, and Jerry Wind (1995), "Introduction to the Special Issue: Empirical Generalizations in Marketing", *Marketing Science* 14, No. 3: G1-5.

Bronnenberg, Bart, Sanjay Dhar, and Jean-Pierre Dube (2007), "Consumer packaged Goods in the United States: National Brands, Local Branding", *Journal of Marketing Research* 44: 4-13.

Chandy, Rajesh, and Gerald Tellis (1998), "Organizing for Radical Product Innovation: The Overlooked Role of Willingness to Cannibalize", *Journal of Marketing Research* 35: 474 – 87.

Chen, Yuxin, Oded Koenigsberg and John Zhang (2017), "Pay-as-You-Wish Pricing", *Marketing Science* 36: 780-91.

Chen, Yuxin, Chakvarathi Narasimhan, and John Zhang (2001), "Individual Marketing with Imperfect Targetability", *Marketing Science* 20: 23-41.

Chevalier, Judith, and Dina Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews", *Journal of Marketing Research* 43: 345-54.

Davis, Donald (2001), "PhD Thesis research: Where do I start?" www.columbia.edu/~drd28/Thesis%20Research.pdf

Dukes, Anthony, Tansev Geylani, and Kannan Srinivasan (2009), "Strategic Assortment Reduction by a Dominant Retailer", *Marketing Science* 28: 309-19.

Dukes, Anthony, and Yi Zhu (2017), "Prominent Attributes Under Limited Attention", *Marketing Science* 36: 683-98.

Dunn, Elizabeth, Lara Aknin, and Michael Norton (2008), "Spending Money on Others Promotes Happiness", *Science* 319: 1687-88.

Frederick, Shane (2012), "Overestimating Others' Willingness to Pay", *Journal of Consumer Research* 39: 1-21.

Ghose, Anindya, and Sha Yang (2009), "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets", *Management Science* 55: 1605-22

Godes, David and Dina Mayzlin (2009), "Firm-Created Word-of-Mouth Communication: Evidence from a Field Test", *Marketing Science,* 28: 721-39.

Godes, David, and Dina Mayzlin (2004), "Using Online Conversations to Study Word-of-Mouth Communication", *Marketing science,* 23: 545-60.

Golder, Peter, and Gerard Tellis (1993), "Pioneer Advantage: Marketing Logic or Marketing Legend?" *Journal of Marketing Research* 30: 158-70.

Goldfarb, Avi, and Catherine Tucker (2011), "Privacy Regulation and Online Advertising", *Management Science,* 57: 57-71.

Horton, John, David Rand, and Richard Zeckhauser (2011), "The Online Laboratory: Conducting Experiments in a Real Labor Market", *Experimental Economics* 14: 399-425.

Iyer, Ganesh (1998), "Coordinating Channels Under Price and Non-Price Competition", *Marketing Science* 17: 338-55.

Laudan, Larry (1977), *Progress and Its Problems,* Berkeley, CA: University of California Press.

Montgomery, Alan, Shibo Li, Kannan Srinivasan, and John Liechty (2004), "Modelling Online Browsing and Path Analysis Using Clickstream Data", *Marketing Science* 23: 579-95.

Norton, Michael, Daniel Mochon, and Dan Ariely (2012), "The IKEA Effect: When Labor Leads to Love", *Journal of Consumer Research* 22: 453-60.

Pischke, Steve (2009), "How to Get Started on Research in Economics?" www.econ.lse.ac.uk/staff/spischke/phds/get_started.pdf

Roberts, John, Ujwal Kayande, and Stefan Stremersch (2014), "From Academic Research to Marketing Practice: Exploring the Marketing Science Value Chain", International Journal of Research in Marketing 31: 127-40.

Rutz, Oliver, and Randolph Bucklin (2011), "From Generic to Branded: A Model of Spillover in Paid Search Advertising", *Journal of Marketing Research* 48: 87-102.

Schmitt, Philipp, Bernd Skiera and Christophe Van den Bulte (2011), "Referral Programs and Customer Value", *Journal of Marketing* 75: 46-59.

Villas-Boas, Miguel, and Russel Winer (1999), "Endogeneity in Brand Choice Models", *Management Science* 45: 1324-38.

Wilbur, Kenneth and Yi Zhu (2009), "Click Fraud", *Marketing Science* 28: 293-308.

Zhang, Juanjuan, (2009), "The Sound of Silence: Observational Learning in the U.S. Kidney Market**",** *Marketing Science* 29: 315-35.

APPENDIX

LIST OF INTERVIEWEES

We are indebted to the following scholars for generously donating their time to our project. They are, of course, in no way responsible for any errors and inaccuracies in this report.

Eric T. Anderson, Northwestern

Sinan Aral, MIT

Alessandro Bonatti, MIT

Bart J. Bronnenberg, Stanford

Randolph E. Bucklin, UCLA

Yuxin Chen, NYU Shanghai

Anthony Dukes, USC

Shane Frederick, Yale

Anindya Ghose, NYU

David Godes, University of Maryland

Peter N. Golder, Dartmouth

Avi Goldfarb, University of Toronto

Raghuram Iyengar, Wharton

Ganesh Iyer, UC Berkeley

Dina Mayzlin, USC

Alan L. Montgomery, CMU

Michael I. Norton, HBS

David G. Rand, MIT

Gerard J. Tellis, USC

Catherine E. Tucker, MIT

Christophe Van den Bulte, Wharton

Russell S. Winer, NYU

Juanjuan Zhang, MIT

Yi Zhu, University of Minnesota