Embedded Entrepreneurship

Olav Sorenson

In this chapter, I review and reflect on my most influential stream of research, on how social relationships influence the location choices of entrepreneurs and the economic geography of industries. I also offer some thoughts about why some scholars end up being more influential than others.

Published as: Sorenson, Olav Sorenson (2021). "Embedded entrepreneurship" Pp. 106-123 in X.-P. Chen, and H.K. Steensma (eds.), *A Journey toward Influential Scholarship*, Oxford University Press: Oxford

Caveat emptor

In the 1980s and 1990s, Jim March taught a legendary course on leadership at Stanford.[1] Two elements of the course made it remarkable. First, he taught the course entirely from great pieces of literature. In the year that I took it, these included *Don Quixote*, *Othello*, and *War and Peace*. But legend has it that in many years, he simply asked students on the first day of class what they wanted to read and then taught the course from those literary works.

Second, he brought his own unique lens to the lessons of leadership. The lesson that comes to mind here came from George Bernard Shaw's *Saint Joan*. Jim would use the story of Joan of Arc to explore the distinction between geniuses and heretics. Both innovate. Both challenge the status quo. But while history smiles on some, those proven right, celebrating them as geniuses, it frowns on others, those who turn out to have been wrong, denigrating them as heretics.

But are geniuses and heretics really any different? In betting on a horse race, the stock market, or an idea, someone will always win. Did they win because they had some special insight – seeing a truth that no one else saw – or did they simply get lucky?

The pages of this volume have been filled with the words and wisdom of geniuses. But any attempt to learn only from success stories deserves a word of warning. These success stories have been chosen in hindsight. In methodological terms, listening only to these accounts amounts to selection on the dependent variable. Many others may have followed similar paths only to find themselves at dead ends—and perhaps defamed as heretics.

The reader should also exercise caution in reading these retrospections for another reason. Despite their best intentions not to revise history, people have a tendency to recall their actions as being the product of calculation and intention. In retrospection, accidents become experiments. Chance events become fated. Meanderings become plans and strategies.

The editors have done their best to guard against this revisionism. They have encouraged all of us to tell the "true" story behind our research paths, to discuss the missteps that have been made along the way. In my own attempts to recall the random and the regrettable, I have revisited many of the early drafts of these papers and a plethora of emails and notes related to them. Despite these attempts to avoid retrospection, however, I would hardly consider myself immune from either hubris or vanity.

Readers should therefore digest this story with a grain of salt. The history described below almost certainly has more coherence than one would find in a contemporaneous account of the evolution of these ideas. The lessons learned have often been inferred without a counter-factual. They detail patterns of behavior but these patterns may as easily reflect superstition as wisdom.

[1] Although no record can fully capture the experience of this class, those interested in learning more about the lessons that Jim taught in this course can read his book, *On Leadership* (March and Weil 2005). Some aspects of his lecturing style, moreover, have been captured in the film *Passion and Discipline: Don Quixote's Lessons for Leadership*.

Embedded entrepreneurs

Although I have written on a wide range of topics, from corporate strategy to the sociology of science, the research for which I have received the most attention has been on the importance of social relationships to entrepreneurs (Rickne et al. 2018). Many others have pointed to the same phenomenon (e.g., Aldrich and Zimmer 1986; Larson 1992). What has been distinctive of my approach to this topic has been to connect it to space and place.

Most social relationships connect people to those who live and work in the same places that they do. Overlapping in physical space provides opportunities for people to meet—whether at a school, a church, a club, or a workplace (Feld 1981). Being close also reduces the cost of developing and maintaining the relationship (Stouffer 1940). People from the same place have more common interests and acquaintances. Meeting for coffee or dinner requires less travel time with someone nearby. Meanwhile, the odds of unplanned meetings – happening to see a person at a store or on the street – decline exponentially with distance. These planned and unplanned interactions both strenghthen the relationship and provide opportunities for exchanging ideas and information. People therefore end up socially embedded in particular industries (social spaces) and places.

Social relationships influence entrepreneurial outcomes in at least three ways. The first, the demonstration effect, helps to determine who becomes an entrepreneur. The second, the mobilization effect, contributes to the ability of entrepreneurs to build organizations—to raise capital, to hire employees, to secure suppliers, and to attract customers. The third, the anchor effect, roots them to the places in which they have become embedded.

Consider first the demonstration effect. Deciding to found a firm often depends on seeing someone else, a peer, doing it. Although this effect could stem in part from being able to call on that peer, perhaps a friend or a colleague, for insight into the pros and cons of entrepreneurship, much of this effect probably occurs even without interaction. Seeing similar others becoming entrepreneurs, particularly successful entrepreneurs, gives people confidence that they too could do it (Sorenson and Audia 2000; Bosma et al. 2012). It also sends signals of the economic attractiveness of founding a firm and of the legitimacy of entrepreneurship as a career path (Sørensen and Sorenson 2003; Stuart and Ding 2006)

When individuals decide to become entrepreneurs, social relationships have a second, mobilization effect. Entrepreneurs must raise capital, recruit employees, and secure suppliers. Social relationships facilitate this resource mobilization phase. Most early financial capital, for example, comes either from friends and family (Ruef 2010), or from professional investors with whom the entrepreneur has a prior relationship or at least where the entrepreneur has someone who can introduce them to the investors (Sorenson and Stuart 2001). Early employees also frequently have prior relationships to the entrepreneur (Ruef 2010).

Relationships help to facilitate the raising of capital and the recruiting of employees, buyers, and suppliers in three ways. Most startups fail. Prospective investors, employees, and partners therefore rightfully approach these commitments with caution. When they have connections to the entrepreneurs, they may have better information about the startup and better insight into its odds of success. Or, they may simply overestimate the entrepreneur, believing the person more able and the ideas more promising than would an objective outsider (Sorenson and Waguespack 2006). Family and friends might even fund or collaborate with an entrepreneur simply for the satisfaction of supporting the person.

Finally, social relationships create an anchoring effect. Part of this anchoring comes from the resource mobilization effect. Entrepreneurs have the highest odds of success when they start their businesses in places where they have deep roots (Dahl and Sorenson 2012). But the anchoring effect goes beyond this simple profit calculation. People also value spending time with friends and family. Entrepreneurs therefore prefer to locate their firms in places where they have loved ones (Figueiredo et al. 2002; Dahl and Sorenson

2009). In fact, people sometimes even become entrepreneurs because they want to live in a particular place but cannot find a job there.

This perspective and these ideas have informed a stream of papers and projects. I will not discuss every one in detail but below I try to provide the "back story" on a few of them.

First, a failure

In 1995, when Netscape went public and the entire world became enamored with the Internet and its commercial potential, I had been studying at Stanford. Silicon Valley, much of it located on land adjacent to Stanford and owned by the University, sat at the epicenter of that boom (or bubble). It seemed as though everyone wanted to become an entrepreneur or to join a startup. Magazines, meanwhile, celebrated Silicon Valley as the future of the economy. Being in the midst of the phenomenon undoubtedly spurred my interest in both clusters and entrepreneurship.

Soong Moon Kang, a fellow graduate student (in the Engineering Economic Systems program), shared these interests. We spent hours discussing the ideas. Much of the classic literature on economic geography had argued that location choice stemmed from the need to access physical inputs, such as coal or iron ore (Weber 1928). But as production moved from being physical to digital, we felt that location choice should shift to depending on proximity to social inputs, such as the availability of specialized suppliers.

We embarked on a research project to explore this idea, focusing on "multimedia" manufacturers. In 1995, the machines used to press CD-ROMs cost tens of thousands of dollars. Hundreds of firms began to enter the business of "publishing" CD-ROM titles—copying, selling, and distributing discs. Many of the early titles essentially substituted for physical books. You could, for example, purchase a dictionary, an encyclopedia, or an industry directory on one of these CD-ROMs.

Multimedia production never became an industry. Technological advances soon allowed anyone with a personal computer to produce their own CD-ROMs. But at the time, the population seemed perfect to us. It had emerged recently, meaning that we could observe the entry locations of even the very first entrants. It also had a small enough number of firms that it seemed feasible to find all of the founders.

We developed a survey to find out what places founders had considered as potential locations for their firms. We conjectured that entrepreneurs would focus on places in which they had lived and places with large concentrations of multimedia manufacturers. We further speculated that those entrepreneurs who chose the places where they had the deepest connections would do best. With financial support from the Alfred P. Sloan Foundation (through a grant to Bill Miller), we mailed our survey to more than 500 entrepreneurs.

Despite receiving responses from more than 100 of them, we ran into an unexpected problem: We had been planning to use whether or not the entrepreneur moved to found their firm as our independent variable of interest but we had no variation on that variable. Almost no one had moved; 94% founded their firms in the places where they had already been living.

Reflections. We had to abandon the project. But we felt it important to understand better why it failed. Did we write our questions poorly? Did respondents misunderstand them? Did we ask the wrong questions?

We decided to debrief a few of the respondents. I vividly recall one conversation in Sunnyvale. We arrived and found the entrepreneur seated in a stand-alone garage that had been converted into an office. He sat at a desk piled high with colorful 3 1/2" floppy discs. Behind him, sat stacks of hardware in rack-mounted columns.

A few minutes into the interview, we asked him why he had started his firm where he did? He responded: "My wife did not want the mess in the kitchen." In other words, he thought our question had been about why he had his office in the garage instead of in the house rather than why he chose Sunnyvale instead of some other city.

In fact, as the conversation continued, it became clear that he had never even considered another location. He just founded his firm where he had been living. We heard similar stories from our other interviewees.

Those interviews changed my thinking. Although I had already expected social relationships to have a (resource) mobilization effect, I had not been thinking that they would have an anchoring effect as well. It also highlighted a problem with both the classic and contemporary literatures on economic geography. Both considered the question of location choice from the perspective of an omniscient and unconstrained social planner. But entrepreneurs face constraints, have interests beyond their businesses, and often do not even engage in active choice.

Sometimes the value of a project comes more from how it influences one's own thinking than in where it gets publishes or in how it influences others.

Persistence of clusters

The immobility of entrepreneurs has an interesting implication: If founders begin businesses in the places they live and in the industries in which they have worked, or in closely related ones, then industry clusters could persist for long periods of time even if they do not represent the most efficient geographic distribution for an industry. In clustered industries, the opportunities to gain industry experience occur primarily in places with existing concentrations of firms. So, the places with these clusters have the largest pools of would-be entrepreneurs in the industry.

Clusters themselves, moreover, might emerge from simple self-reinforcing stochastic arrival processes. By chance, one location could have a few more early entrants. That location then has a larger pool of able entrepreneurs in the industry, meaning that more future entry occurs in the same location. More entry again leads to more potential entrepreneurs, creating a loop of positive feedback.

That idea represented a major departure from the literature. Both the classic and the recent literatures on economic geography have argued that the geographic distribution of firms must stem from economic efficiency. In other words, industry clusters must emerge because firms benefit from being located near to large numbers of rivals (agglomeration externalities), perhaps because they share information, specialized suppliers, or an unusually high-quality pool of employees.

But the idea that industry clusters must reflect economic efficiency depends on two implicit assumptions: (1) Entrepreneurs choose the places that maximize their expected profits; and, (2) the places that maximize these profits do not vary from one entrepreneur to the next.

Although I had the intuition that these assumptions seemed wrong, I did not have the data to test my ideas. Soon, however, I had the good fortune to meet someone who did. In August of 1995, I decided at the last minute to drive up from Stanford to attend my first Academy of Management meeting in Vancouver. There I met Pino Audia, when I visited his poster session. We stayed in touch after the conference and I discovered that he had assembled an amazing data set, one with information on every plant in the shoe industry in the United States, from 1940-1989, including where each plant had been located. We agreed to collaborate.

The analyses took a long time. To test the effects of concentration, we introduced a novel type of measure, a distance-weighted density measure (i.e. a distance-weighted count of other plants in the industry). Because this measure required pair-wise calculations for each plant in the indusry, for each year, it took my 1997 vintage computer nearly two days to compute each version of the variable.

Pino and I found that footwear plants located in dense concentrations of footwear manufacturers failed at higher rates than those in more remote regions. But, because entrepreneurs in the industry entered these crowded places at even higher relative rates, the industry remained geographically concentrated (Sorenson and Audia 2000). Clusters persisted in the shoe industry because of where entrepreneurs entered rather than because clustering provided any performance benefits for firms.

In the course of presenting the paper, seminar participants at Berkeley, London Business School, UCLA, and the University of Chicago offered an abundance of alternative interpretations of the analyses, ones that could preserve the idea that industry clusters stem from positive agglomeration externalities. As we revised

the paper again and again, we tried to address each possible explanation—bringing additional analyses to the data or providing more details about the context that would allow us to dismiss them. I view this extensive consideration of alternative explanations as one of the strengths of the paper and a probable factor in its success.

In the end, however, this paper also benefited from encountering a strong editor. Even in the third round of review at the *American Journal of Sociology*, one of the reviewers continued to complain about not believing our interpretation of the results—agglomerations, the reviewer argued, are efficient. The editor, Roger Gould, recognized that we could not do much more with the data and that the vociferousness with which the reviewer had been arguing signaled that the paper might garner a great deal of interest. He therefore accepted the paper even though one reviewer remained unsatisfied.

This paper led quickly to another on the same subject. When I presented the paper with Pino at the University of Chicago, Toby Stuart suggested that we could investigate the same processes using his data on the biotech industry. Although that paper began as a replication, Toby and I ended up building out the theory around the mobilization effect and testing it more directly by creating measures for each region of the distance-weighted availability of human, financial, and intellectual resources (Stuart and Sorenson 2003). That paper also struggled a little in the review process before finding a home at *Research Policy*.

Reflections. Despite academia's professed interest in novel ideas, ideas that go against the dominant beliefs are hard to publish. Papers that provide empirical evidence for something that most people already believe or that offer minor modifications of mainstream ideas seem to get the most attention.

Even though these papers have been highly cited, those papers that cite them often extract only part of the message. They rarely engage the idea that firms fail to benefit from locating in clusters. Instead these papers usually get cited for (i) the idea that spin-off processes play an important role in the creation of clusters, (ii) the idea that entrepreneurs rely on their social relationships to mobilize resources to build their firms, and (iii) the idea that regions rich in financial and human capital spawn more startups.

Relationships and resources

My most cited paper began as a footnote.[2] In the process of writing the paper on the geography of the biotech industry (Stuart and Sorenson 2003), Toby felt that we could not simply show that biotech firms entered in the places that had the largest concentrations of resources. We needed some direct evidence that resource mobilization depended on distance.

Through a grant from the Ewing Marion Kauffman Foundation (through Steve Kaplan's Center for Entrepreneurial Leadership at the University of Chicago), we purchased data on venture capital investments. Those data revealed that venture capitalists did indeed invest disproportionately in startups located near their offices. But through the course of analyzing the data, we decided that the idea could support a paper in its own right.

In the paper, we argued that venture capitalists would invest in companies located near them for two reasons (Sorenson and Stuart 2001). First, they rely on their social relationships to identify and evaluate promising startups. Second, post-investment, proximity facilitates the process of advising and monitoring their portfolio companies. Although a number of factors could explain the "home bias" in investing, our social relationship-based account had one further implication: Venture capitalists with more expansive social networks could invest at greater distances.

The paper therefore treats the fact that venture capitalists invest locally as a baseline. The analyses focus on demonstrating that social networks, not size or experience, explain which venture capital firms invest in more distant startups. In essence, the paper focuses on in-

[2]However, the footnote which began it all never made it into the final paper. Because the offshoot paper got published before the original from which it emerged, we simply cited our published paper instead of describing the analysis in a footnote.

teraction effects rather than on the "main" effects of distance.

One of the methodological innovations of this paper involved using a case-control design. The dominant approach to analyzing relationship formation up to that point had been to create a list of all the possible dyads that could exist and to treat each as an observation. In a population with N actors, that would mean $N^2/(N-1)$ observations. That approach had been feasible in the small, cross-sectional samples that had typically been studied. But with more than $6,000$ unique investors across many years, for us, it would have meant creating a data set with more than 6 million rows. Variable construction alone would have required weeks with the computing power available to us at the time.

We therefore created our sample by choosing all of the realized relationships – co-investment relationships that had actually formed – and matching them one-to-one with relationships that had not occurred. That approach meant that we would have 80 thousand rows of data instead of 6 million, saving an enormous amount of time without a large loss of information.[3]

The case-control design also has a second benefit. It limits the degree of spatial autocorrelation in the data. Most relationship formation studies end up with hundreds of repeated observations of each actor and analyzing many dyads that undoubtedly have correlated errors due to their similarity or even due to contagion through third-party connections. By sampling the matrix sparsely, the degree of autocorrelation between any two cases included in the estimation should remain small.

Reflections. All of the claims in this paper revolve around the interactions between distance and other variables. That approach to theory construction and testing has some advantages. Most notably, people find it far more difficult to come up with alternative explanations for interaction effects. Most endogeneity stories explain a single effect, selection into a treatment. They have a harder time accounting for heterogeneity in the treatment effects.

Papers can provide contributions in many ways: They may offer a truly novel theory, giving readers a new perspective on the world. They may present and analyze data that have unusual features, that even at a descriptive level have value. Or, they may offer a novel approach to analyzing data, whether in terms of measurement or estimation. Great papers seem to offer innovations on at least two or perhaps all of these dimensions.

This paper did not necessarily offer a new theory but it did revive a literature that had been dormant for decades. It also used fairly novel data. Although papers on venture capital now number in the hundreds, at the time our paper came out, few had analyzed venture capital data and no one had connected it to geography. The case-control approach to estimating relationship formation, moreover, had been a methodological innovation.

Home sweet home

The final pair of papers represent a return to my first failed attempt at research in this area. Recall that Soong and I had expected entrepreneurs to found firms in places they had lived. We also anticipated that entrepreneurs who opened firms in places where they had deep connections would enjoy more success.

Although I did not have the data to investigate these ideas in 1995, a decade later I had found it. Jesper Sørensen had introduced me to the Danish employer-employee data (see Sørensen and Sorenson 2007). It had the information one would need. But Jesper did not seem interested in the project and I did not have direct access to the data. Fate again intervened. When discussing some other research ideas with Michael Dahl, I discovered that he too had access. Thus another productive collaboration began.

We originally wrote one paper. That paper first demonstrated that entrepreneurs tended to found their firms in the places where they lived (and that if they did move, they did not move far). It then showed that entrepreneurs who stayed in the regions in which they had

[3]With rare events, most of the information for estimation comes from the few cases with events rather than from the many without them (King and Zeng 2001).

been living performed better than those who moved.

We submitted that paper to the *Journal of Economic Geography* in October of 2007. Although the journal offered us the opportunity to revise and resubmit our paper, the revision appeared difficult and we never resubmitted it. During the time that our paper had been under review, Michelacci and Silva (2007) came out in the *Review of Economics and Statistics*. That paper demonstrated the same fact as the first half of our paper—that entrepreneurs move less than employees. The reviewers therefore felt that our paper merely replicated their results.

We decided to split the paper into two. The first half, on which we had been scooped, became Dahl and Sorenson (2009). Splitting this half off gave us space to do something more than just showing geographic inertia. We estimated entrepreneurs' revealed preferences for locations, characterizing places not just in terms of where they lived but also in terms of industry conditions and proximity to friends and family.

We found some interesting patterns: First, entrepreneurs appeared far, far more sensitive in choosing locations to where their friends and family lived than they did to industry conditions. Second, to the extent that they did weigh the local industry characteristics in their calculations, they seemed systematically attracted to places with more intense competition (see also, Sørensen and Sorenson 2003).

Splitting off the second half, the more novel part of the paper, also allowed us to extend that analysis. We demonstrated that the beneficial effects of having deep roots in a region held not only for firm failure rates but also for firm profits and for the total amount that the entrepreneur earned. These results, moreover, held even when using an instrumental variable to account for potential endogeneity in the choice of locations. That paper landed in *Management Science* (Dahl and Sorenson 2012).

Although both halves ended up becoming papers that could stand on their own, the splitting of the original paper had a casualty. The original paper had demonstrated the negative effects of being in regions with large numbers of competing firms—extending the results from the footwear and biotech papers to the entire range of industries in the Danish economy. But those results ended up being absorbed into industry-region fixed effects in the *Management Science* paper in the interest of focusing on the home court advantage for entrepreneurs.

Reflections. Nearly all of my papers have been written with coauthors. Those partners have played different roles. Sometimes they have had the data; sometimes they have done the analysis; sometimes they have had the idea. Regardless of the role played, my collaborations have almost always produced better papers than I could have written on my own.

Over the years, I have adopted a couple of rules for choosing collaborators. First, I need to enjoy interacting with the person. Coauthorships involve a lot of interaction. If you do not get along, it will become draining.

Second, I will only commit to doing one paper with a person until we have completed at least one project together. I have been fortunate. Although my collaborations have been more and less productive, I have not had any truly bad coauthors. Bad coauthors can cause misery in many ways: They might disappear for long periods of time. They could do sloppy work on their sections of the project. Or they might insist on first authorship despite having done little of the work. If you have a bad coauthor, you want to stop the bleeding as quickly as possible.

I also have one rule for writing with coauthors: If someone edits a sentence or a passage, another coauthor cannot just revert to the original text. The person who changed the text did it for a reason. Perhaps they found it unclear or misleading. It may remain unclear or misleading. Anyone should feel free to rewrite it again. But simply reverting to the original text signals either a disrespect for the person who edited it in the first place or that the original author has become too attached to a phrase or sentence.

Additional thoughts

The reflection sections above include many thoughts on the research process. But they have not necessarily engaged the central question of this volume: What leads to influential scholarship?

The notes below reflect my thoughts on this question. I have not done systematic research on any of them. But I feel that they all relate to the production of influential scholarship.

Choosing projects

Papers that address important topics with novel theory and/or high-quality data get the most attention. They influence both other researchers and often inform the pubic policy debate. But what counts as important? What qualifies as high-quality data?

Pick important problems. Theories and theoretical perspectives come and go. Although scholars continue to build on the central insights of these theories, history forgets most of the nuances of these theories and the internal debates within these perspectives. Most papers motivated by the need to elaborate a particular theory therefore end up limited and short-lived in their influence.

But important problems persist. Although importance may seem subjective, some problems still enjoy greater consensus in being perceived as important. Global warming. World peace. That's aiming a little too high for most of us. But important problems usually affect a lot of people. They often stem from fundamental aspects of human behavior.

In choosing projects, one should therefore ask: How many people does this problem affect? How large of an effect does it have on them? Does it shape their health, their wealth, or their well-being?

Important problems often attract a lot of attention but influential scholarship usually brings a different perspective to understanding these problems. That's the theoretical lens.

The most satisfying research on important problems, moreover, has implications for policy. By implications for policy, I do not mean policy evaluation—research oriented toward assessing the efficacy of a particular intervention. Rather, I have in mind research where the theory and results could inform a wide range of potential policies or interventions.

In choosing projects, one might therefore ask: Who, if anyone, would change their behavior if they had awareness of the research? Would bureaucrats and legislatures pursue different policies? Would executives or entrepreneurs choose different strategies? Would employees manage their careers differently?

Choose quality over quantity. Most research in management involves the analysis of data. Of late, people have become enamored with big data. What "big" means varies. Datasets with tens of thousands of cases used to seem large. Many large datasets today include information on millions of something—people, patents, or transactions. As sensor data become more widely available, millions will give way to billions.

But the value of data stems more from their quality than from their quantity. High quality data have high fidelity, the information corresponds to reality. They offer valid measures, connecting closely to theoretical constructs rather than relying on distant proxies. They accurately represent some population, allowing generalization from the sample to the population.

Digital technologies and platforms have dramatically reduced the cost of acquiring data. But the ease of collecting data can become a trap. One could, for example, conceivably email a survey to millions of people, meaning that even a 1% response rate would yield tens of thousands of cases. With a 1% response rate, however, the survey has limited value since one has no idea how well it might represent the sampling frame, the population.[4]

Mechanical Turk and other platforms allow for cheap experiments. Usefully, these platforms often allow for including large numbers of participants in each treatment, hundreds instead of dozens. But, again, one has limited

[4]With low response rates, selection effects become such a large concern that one should probably just treat a survey as a "convenience" sample even if the original sampling frame might have been representative of a population.

information on who participates in the experiments. In some cases, experiments on these platforms may simply assess the behavior of bots.

Screen scraping has similarly enabled the construction of large datasets at low cost. Often, these data allow researchers to study novel phenomena. But in too many cases, researchers scrape the data and then search for a problem that it could address. As a result, the connections between the variables and the theoretical constructs end up being distant at best.

People also pay too little attention to the provenance of scraped data. Who collected the information originally? What effort has been put into verifying the accuracy of the data?

Do not fear the identification police. "Identification" concerns whether variance driving the estimates might allow for a causal interpretation of the coefficients. Strategies that support a causal interpretation under certain assumptions would include estimation with instrumental variables and differences-in-differences designs.[5] Although management came to the "identification revolution" later than economics, these concerns have now become central to editors, reviewers, and readers.

But the concerns regarding identification often go too far. All else equal, I too would prefer a study that allows for causal inference. But all else is rarely equal. In the search for causal inference, researchers often sacrifice the picking of important problems. Perhaps the worst form of this malady involves choosing problems on the basis of finding a plausible source of exogenous variation. Many studies that rely on "natural" experiments are brilliant. But others feel like a solution in search of a problem.

People also forget that providing evidence in favor of a theory, even in support of a causal effect, does not necessarily require a source of exogenous variation. Another effective and compelling strategy involves triangulation: Being explicit about what alternative interpretations might threaten the researcher's preferred account for the results, and then using those alternative theories to develop critical tests that could distinguish between the competing accounts.

Writing papers

Given the way that most graduate programs train their students, the writing of papers appears almost an afterthought. Students take multiple courses on research methods, on econometrics, and on the relevant past research. But few, if any, Ph.D. programs include formal training in writing. Instead they rely on students learning on their own or from informal mentoring.

But what becomes influential and what does not may depend more on how the paper describing the research has been written than on anything else. Writing determines who will read the paper, whether they will read the whole thing or simply skim the introduction and the tables, and what they will take away from it.

Frame around the phenomenon. Framing helps to determine the audience for the paper. Who will find it interesting? It also tells readers why they should pay attention to the paper. How will it change their thinking?

Authors usually have choices about how to frame papers. I tend to frame my papers around the phenomenon, the empirical puzzle. My papers on the persistence of clusters, for example, begin by noting that many industries cluster but that the traditional explanations for clustering fail to fit many of the cases where we see it. The home sweet home paper opens by noting that entrepreneurs remain rooted in place, even when other regions appear to offer more favorable environments.

Framing around the phenomenon will usually attract a larger audience. Only those interested in a particular theoretical perspective will read on when a paper has been framed in terms of its theoretical contribution. But people from many different perspective often care about the same phenomena, particularly if an important problem has been chosen.

Mind the gap. Among the various ways that authors may open papers, I find the "gap" ap-

[5] For an excellent yet practical guide to these methodological issues, see Angrist and Pischke (2009)

proach to framing the least compelling. This framing usually comes in one of two forms. Either the author claims that theories x and y have been studied independently but never together. Or, the author argues that theory x has never been applied to the particular setting, and then proceeds to apply theory x to that setting.

This approach to framing has at least two problems. First, as noted above, this approach limits the audience to those who care about theory x. Second, it does not set up a puzzle, a question to be answered. Who cares? Or, perhaps more accurately, why should anyone care? Why would we not simply expect theory x to have the same effects in another setting, or why would theories x and y not simply have additive effects?

Write well. It often surprises me how poorly many academics write. But then again, despite the fact that we probably spend more time writing than on any other activity, most of us have never had any explicit training in it.

I, too, have had no formal training but I have tried to train myself. During my first four months as an assistant professor, I spent a large share of my time reading articles carefully, not for what they said but for how they said it. I chose about a dozen articles that I liked and proceeded to read them first at the level of the paragraph and then at the level of the sentence. What did the author hope to achieve with this paragraph or sentence? Why did it appear here instead of earlier or later in the paper? At a structural level, writing well means telling the story effectively.[6]

I have also read many a book on writing. One of my favorites is *On Writing Well*, by Zinsser. These guides include a great deal of advice on how to improve one's writing. They often provide exercises for practicing these principles. At the sentence level, writing well means writing sentences that readers can easily consume. The prose need not be stylish. Few can emulate a Jim March or a Steve Barley. But it should be clear. I will not elaborate further other than to say that I consider learning to write better to be time well spent.

Writing means revising. My papers often go through 20 or more drafts. Sometimes I will do an entire round of revision focusing simply on eliminating sentences in passive voice or unnecessary adverbs or on trying to ensure that I have variation in the lengths of my sentences. This process takes time. But, as with anything, you get better at it, the more that you do it.

Building an agenda

Influential scholars are never one-hit wonders. They rarely have portfolios of one-off papers on a wide variety of subjects. Instead, they write a sequence of papers that connect to one another. Those papers add up to more than the sum of their individual contributions.

In other words, influential scholars build research programs. That's not to say that they have a series of research papers and projects mapped out from the start. Influential scholars typically develop these programs in two main ways, dynamically responding to the attention and criticism that their papers receive.

Replicate. First, they replicate their successful papers. They do not replicate them in the literal sense of simply rerunning their code. Rather, they engage in conceptual replications, demonstrating that the results do not depend sensitively on a particular setting or a specific way of measuring one of the core constructs.

In studying the persistence of clusters, for example, after first examining the patterns in the footwear industry, I then studied them in the biotech industry. I also wrote a book chapter that extended my dissertation data to study the persistence of clusters in the computer workstation industry. Each of these "replications" also offered an opportunity for building the theory out in more detail. Toby and I, for example, fleshed out the resource mobilization effect when examining these processes in the biotech industry.

Extend. Second, those building research programs look for the next logical question. Each set of papers described above – except of course for the first one – emerged to some extent from

[6]The Heath brothers have written an excellent book, *Made to Stick*, on the elements of effective story-telling (Heath and Heath 2007).

questions left open or that became salient in an earlier paper. The failed multimedia project led me to the idea of cluster persistence. Examining cluster persistence pushed me to examine the resource mobilization effect directly. The home sweet home papers examine the assumptions underlying the cluster persistence papers. Most research programs have this sense of moving in a series of logical steps.

Enjoy. To close the circle, let me end with another lesson from Jim March. Jim often discussed the importance of deriving pleasure from the process, of leadership, of research, of writing.

Attempting to measure success in terms of having influenced others seems a losing game. We have little control over the ends, whether we achieve fame and fortune. Research is consumption. Most of us became academics because we enjoy doing research. Above all then, it's important to study topics that we find interesting, where we win even if no one else pays attention.

References

Aldrich, Howard E. and Catherine Zimmer. 1986. "Entrepreneurship through social networks." *California Management Review* 33:3–23.

Angrist, Joshua D. and Jorn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton, NJ: Princeton University Press.

Bosma, Niels, Jolanda Hessels, Veronique Schutjens, Mirjam Van Praag, and Ingrid Verheul. 2012. "Entrepreneurship and role models." *Journal of Economic Psychology* 33:410–424.

Dahl, Michael S. and Olav Sorenson. 2009. "The embedded entrepreneur." *European Management Review* 6:172–181.

Dahl, Michael S. and Olav Sorenson. 2012. "Home sweet home? Entrepreneurs' location choices and the performance of their ventures." *Management Science* 58:1059–1071.

Feld, Scott L. 1981. "The focused organization of social ties." *American Journal of Sociology* 86:1015–1035.

Figueiredo, Octavio, Paolo Guimaraes, and Douglas Woodward. 2002. "Home-field advantage: Location decisions of Portuguese entrepreneurs." *Journal of Urban Economics* 52:341–361.

Heath, Chip and Dan Heath. 2007. *Made to Stick: Why Some Ideas Survive and Others Die*. New York: Random House.

King, Gary and Langche Zeng. 2001. "Logistic regression in rare events data." *Political Analysis* 9:137–163.

Larson, Andrea. 1992. "Network dyads in entrepreneurial settings: A study of the governance of exchange relationships." *Administrative Science Quarterly* 37:76–104.

March, James G. and Thierry Weil. 2005. *On Leadership*. Malden, MA: Blackwell Publishing.

Michelacci, Claudio and Olmo Silva. 2007. "Why so many local entrepreneurs?" *Review of Economics and Statistics* 89:615–633.

Rickne, Annika, Martin Ruef, and Karl Wennberg. 2018. "The socially and spatially bounded relationships of entrepreneurial activity: Olav Sorenson—recipient of the 2018 Global Award for Entrepreneurship Research." *Small Business Economics* 51:515–525.

Ruef, Martin. 2010. *The Entrepreneurial Group*. Princeton, NJ: Princeton University Press.

Sørensen, Jesper B. and Olav Sorenson. 2003. "From conception to birth: Opportunity perception and resource mobilization in entrepreneurship." *Advances in Strategic Management* 20:89–117.

Sørensen, Jesper B. and Olav Sorenson. 2007. "Corporate demography and income inequality." *American Sociological Review* .

Sorenson, Olav and Pino G. Audia. 2000. "The social structure of entrepreneurial activity: Geographic concentration of footwear production in the United States, 1940-1989." *American Journal of Sociology* 106:424–462.

Sorenson, Olav and Toby E. Stuart. 2001. "Syndication networks and the spatial distribution of venture capital investments." *American Journal of Sociology* 106:1546–1588.

Sorenson, Olav and David M. Waguespack. 2006. "Social structure and exchange: Self-confirming dynamics in Hollywood." *Administrative Science Quarterly* 51:560–589.

Stouffer, Samual A. 1940. "Intervening opportunities: a theory relating mobility and distance." *American Sociological Review* 5:845–867.

Stuart, Toby E. and Waverly W. Ding. 2006. "When do scientists become entrepreneurs?" *American Journal of Sociology* 112:97–144.

Stuart, Toby E. and Olav Sorenson. 2003. "The geography of opportunity: Spatial heterogeneity in founding rates and the performance of biotechnology firms." *Research Policy* 32:229–253.

Weber, Alfred. 1928. *Theory of the Location of Industries*. University of Chicago Press.

Zinsser, William. 2006. *On Writing Well*. New York: HarperCollins.