

Heuristics for Discovery

John Levi Martin

Published: July 26, 2018

John Levi Martin: University of Chicago (United States)

✉ jlmartin@uchicago.edu; <http://home.uchicago.edu/~jlmartin/>

John Levi Martin is the Florence Borchert Bartling Professor of Sociology at the University of Chicago. He is the author of *Social Structures*, *The Explanation of Social Action*, *Thinking Through Theory*, *Thinking Through Methods*, and the forthcoming *Thinking Through Statistics*, as well as articles on methodology, cognition, social networks, and theory. He is currently working on the history of the theory of social action.

I imagine that I am not alone in insisting that I do not choose my topics, but rather, the topics choose *me*. They do this in just the way that, say, a potato chip decides that you will eat it — by lying there, in plain view, looking all crispy and salty. But I still think I have some useful things to say, some prescriptive, some descriptive — some about good heuristics for sociology, and some about the heuristics I have used — which, unfortunately, were not always the same thing.

1 What is Discovery?

To do this, I will rely on an extended metaphor. I will try not to repeat arguments made in *Thinking Through Statistics*, *Thinking Through Theory* and especially chapter 2 of *Thinking Through Methods*. But first, I want to take seriously the charge to consider heuristics for *discovery*. There are, so far as I am aware, very few major discoveries in sociology — perhaps less than ten. A discovery isn't just an idea, or an argument. And simply being able to find data that don't contradict your claim doesn't mean you've discovered something. A discovery is something that involves some degree of surprise (even if the discoverer suspected it was there, other people didn't), and, after it happens, is the subject of relatively wide consensus. Of course, every now and then something that was "discovered" is later determined to be a mirage. But we call things discoveries when we think they aren't.

The social sciences as we know them really started with a discovery, one dating back to the political arithmeticians of the seventeenth century. This discovery was that certain aggregate ratios (such as the birth rate and the suicide rate) were extremely stable. This does not surprise now, but it did then, and has remained a very robust finding. Another example of such a discovery was Zipf's law of the distribution of city populations. For many years, this was scoffingly dismissed by sociologists as sheer coincidence, or, if robust, still irrelevant to theory. But it has remained an important fact and helped guide some of the most theoretically advanced work in social science, namely economic geographic/regional systems theory.

In sociology we have many claimed discoveries that unwind relatively quickly because they are cobbled together out of implausible theoretical assumptions. That is, it's only a "discovery" if you ignore all the unproven, supporting axioms. Of course, many scientific discoveries require a certain background acceptance of theory. For example, in geology, the pivotal discoveries that the magnetic orientation of the poles reverses over time, and that there is crust being produced at certain ridges under the sea, required people to accept some theories about the relation of current rock formations to historical processes. But those axioms had good justification, they've proven pretty robust, and still are widely accepted.

So first, let's hold ourselves to somewhat high standards when it comes to what constitutes discovery, and we won't count as discoveries claims like "social capital is important". Why not? Well, for one thing, "social capital" isn't a natural kind, but a hodge podge that no one's really thought through. It doesn't really hold together under investigation, such that we could make general statements about it that aren't either empty or wrong. But more important, other sciences don't want to give out prizes for things like "algae is important" (though it is). Second, if, we don't make discoveries, what *do* we do? I think most sociology is an attempt to shift the balance of the evidence one way among a set of contending explanations for a generally accepted phenomenon. That may not be discovery, but it still can be science.

Before discussing heuristics for sociology, then, we need to decide whether it is worth trying to maximize the chance of a real discovery, at the cost of being unable to do decent sociology. For the heuristics that are most likely to maximize the chance of making even minor discoveries may not be very good at anything else, and if the chance of making discoveries is very small, this might be a bad bet.

2 Backtracking

Let me give one example of an all-purpose heuristic that I have used upon occasion, and that does work for one kind of discovery. Sometimes we are sure that something *has* to exist, but we just don't know quite what or where it is. For example, think about setting out to discover the source of the Nile. This is very different from discovering America. The way you discover America is to be so bone-headed dumb that you literally run aground on something you didn't think was there. That's an unplanned discovery, and unplanned discoveries can have a lot of fun and drama to them. In contrast, if you want to discover the

source of the Nile, you can set out looking for it, because you're going to be pretty sure that it's *somewhere* — or at least, there are going to be *some* sources of the Nile. It's just ("just"!) a matter of tracing it out (and up — it's got to start somewhere higher than where you are now).

However, the reason that sociologists aren't thrilled about trying to use this heuristic ("just trace it backwards") is the same reason that Sherlock Holmes would rarely bother borrowing the bloodhound to trace back the steps of someone whose mucky shoe he found. And the reason was, you trace back that smell, no matter where you start, and you end up in the same huge pile of s—t. We can call it capitalism, modernity, hegemony, society, or whatever. Certainly, what an infinite number of mediocre papers on "reproduction" have shown us is that if you rush out randomly into the bushes and try to trace back, you end up with all the other lost students, out there on Hegemony Hill.

When we try to follow the heuristic in a more focused, and disciplined, fashion, and trace the *specific* source of a clearly delimited class of phenomena, we tend to be frustrated — we're trying to find the peak of a very craggy mountain, where we can't see far ahead, and constantly face forks where we have to go one way or the other. Not surprising that many methodologists increasingly tell us that the answer is, don't even *try* to go uphill — don't look for causes — *go downhill*, and look for *effects*.

But, if we are interested in discovering, we need to take the heuristic of back-tracing more seriously. In fact, killing off our enthusiasm for questions of the form "where did it come from" — a cauterization of our fundamental sociological curiosity — will devastate our field, since our strength must be, at least in part, in answering the sorts of questions that regular people have about the world, and those questions are rarely "how much will y change if we twist the x dial" and more like "where the heck did y come from, anyway?" Even a snoring drunk can roll down a hill — that shouldn't win any prizes.

But tracing back rarely involves just charging off into the first bushes you see with a machete. One has to start from the well-established parts of the Nile and work upwards, and then make choices when there are forks. There are three problems we run into. The first is when you come to an impassable zone — an unscalable height, an unfordable river, and so on. Most important here is when there are reproductive phenomena, or strong selectivity issues. The second is when you simply lose the trail. Sociologists can be better than they are at saying, "and here we must stop for the scent disappears." Finally, very often, if you do follow the trail with determination, you find that it leads you right out of the jurisdiction that you've been given permission to explore — that is, you need to start working with experts from other fields, such as psychologists, anthropologists, or (perhaps most importantly) historians.

The times when I have actually set out to trace back and discover something were ones in which I was brought up quickly to the limits of my training and ability. Some of those have been intellectual history: my first paper tried to figure out where the idea of a "sexual revolution" came from, and I probably got this wrong. At the time, I didn't know about the work of Otto Gross, and that almost certainly was pivotal (I dated it only to the end of the First World War). Much of the article is still right, but I should have turned to a German historian. My most solid claim to a discovery has to do with the relation between two different ways of interpreting Coombs's non-compensatory factorization. I had started to understand, with the help of James Montgomery, that there were two different versions of this, and figured that, there's got to be *some* place where these two things fork off each other, and if I'm patient, I can find it. In this case, I had to learn some math and demonstrate that it comes from the question of whether the lattice of all possible states is lower semi-modular (I published this last year and I have already forgotten what this all means). The problem is, this is probably something that mathematicians knew a long time ago.

Maybe that's a particular example, but I think it's not at all uncommon for serious backtracking to lead you into other fields. As far as I can see, Durkheim was just flat out wrong about this idea that a social fact always has to be traced back to another social fact. Good news for sociology if that were true, but there's no reason to think that the world works that way. So I don't think the "discover the source" heuristic is one to recommend (certainly, I've had no success with my most fundamental sociological impulse, the desire to "find the source of politics").

3 *Principal Cities*

I'd really like to stick with this metaphor — not necessarily that we are trying to discover the source of the Nile, but that we are trying to learn about a country. And by this, we mean to learn about its general

geomorphology — what is the land like? How is it structured and arranged? The metaphor can help to clarify the implications of a number of heuristics.¹

If you wish to “see France,” does that mean to go to Paris? In some ways yes, and in others no, in that while Paris is the apex of France, most of France is not Paris — and Paris is distinctive in a very particular way. It’s special, and it’s cool, I suppose, but it is, above all else, a concentration of *people*. You don’t go there to see the landscape as much as to see what people have done with it. An example of the Paris-equivalent in the realm of subfields of sociological endeavor is the problem of intergenerational mobility. It is a busy city of scholars in communication, and, for that reason, recognizably bad work is usually weeded out rather quickly. If you want to do good work, do you move to the city? There are advantages in starting where people have done a lot. The stumps are out, the swamps are drained. But, like the country kid moving to the city, it can be a lonely place for you to get started — the people already there seem to know what is going on, and you don’t. You sit in your little apartment alone at nights, watching Mr. Ed, while they are (or so you believe) at champagne bashes. Perhaps they aren’t, but it is true that those who come from clusters that have been working on these issues *are* ahead of you.

But there’s a deeper problem, or so I now believe, with the “principle cities” approach. While the work done is usually better than the work done elsewhere, it tends, like a city, to be highly artificial. Remember, in this metaphor, the thing we are trying to discover is less what the citizens do, than physical geology of the terrain. But it’s those sorts of people-doings that make cities what they are. Cities are certainly somewhere, but they tend to be in unusual places (usually by at least one river, sometimes a confluence of different waters). And sometimes they fundamentally change the terrain, or make it impossible to really see. So they’re relatively hard places to make geological discoveries.

I’m not saying there aren’t great archeological digs going on in cities. But just as in real life, when you do that sort of excavation, you’re still usually finding *earlier* cities, on top of which yours was built. Further, cities are relatively hard places to make discoveries, since there are so many watchful eyes about. When there are discoveries in cities, it’s often the result of being able to get new data.

To return to the case of stratification, a great deal of time was spent trying to estimate parameters from layered 7×7 tables. That was progressive in a way, but it was “indoor” research. It was arguing over relatively small issues of a model constructed for purposes of convenience, and one that didn’t always hold up to the “look out the windows” critiques. But it gave everyone something to do, and it was beautiful in an aesthetic way, so we stuck with it.

Cumulative social science is an important goal. But so is the reality check. I consider one of the few discoveries of sociology to be David Grusky’s work on micro classes. He had a talk in which he simply showed something like a 300×300 occupational inheritance matrix done as a 3-D frequency plot. It fundamentally changed my understanding of the American class structure in a way that none of those arguments about whether ϕ could or could not be compared across samples did. That’s what I mean by fundamental geomorphology.

So in cities, I think there are only two types of discoveries. One is the discovery of better data that allows you to avoid assumptions that everyone has had to make before. The other is the realization that you’ve been digging in the wrong place entirely, and, sadly, that does happen — you can be the killjoy who discovers that the city is built on a fault line or in the shadow of Vesuvius and has to be abandoned. That sort of negative discovery is a real contribution, but it’s not as satisfying as the positive discovery. And maybe it’s not even as satisfying as good old fashioned, Lewis and Clark, exploration and mapping. That’s often the first step to a serious geomorphology, and I don’t think we’ve done enough of it.

One key to a great Lewis and Clark heuristic is that you try to get somewhere that might be rather arbitrary (for example, it might come from political or identity concerns), but is still far away. To get there, you need to traverse lots of different types of difficult territory. You can avoid some of it, but not all. You learn about the nature of the Rocky Mountains when your goal is on the other side of them. Weber, perhaps, did some seriously good comparative sociology, because he was trying to walk by himself from coast to coast. He didn’t get there, but by gum, by the time he died, he had mapped a fair amount of what could be mapped at the time.

1. If you are impatient with metaphor, my dear little scientists, bear in mind that sociology itself has been, to a first approximation, a centuries-long attempt to think through the implications of the metaphor *society - organism*, and a great deal of insight has come via a critical examination of the strengths and limitation of that metaphor!

From here on, I'd like to make a case for heuristics of mapping. I'll talk about what has worked for me, what hasn't worked, and what I'm hoping to try soon.

4 Follow the Shiny

Of course, I must begin by admitting that if I do in fact decide what to do next, my consciousness has not been let in on the secret of that decision process. So rather than invent a plausible argument about how these decisions are made, I'll give a more accurate, if vague, report on the feeling I have, which is that rather than choosing my tasks, I merely accept them. That sort of relation to an external (possibly only virtually external) leading force — what Weber called an “object” — is easy to justify when (as in Weber's vision) it pulls you by the hair in a straight line. But it is no less a pull when it hops and skips all over the place. I trust that it knows what it is doing, because it hasn't let me down yet.

The way it appears to me is a bit like a halo of shininess that bobs beckoningly in certain areas, and I dutifully follow. While “mission from God” probably somewhat overstates the coherence, “Attention Deficit Disorder” probably understates it. Of course, most things I do branch off from previous projects. Finishing one isn't really finishing, and it points you in the direction to the next. But one branch may slowly ebb in force, and so something else starts to take its place.

At that point, I believe that I follow a simple three-part algorithm. First, look around for the shiniest thing you can see. Second, attempt to go towards it. Third, if it does not *obviously* increase in shininess, drop it and look around for something else that is shiny. This is not an altogether bad way of proceeding. While it does give disproportionate attention to things that are near, those are often ones we are in a better position to attack, as we have already some of the conceptual and methodological infrastructure in place. The emphasis on “obvious increase” in shininess as a key part of this heuristic means that one avoids chasing will-o-the-wisps — but it also means you aren't guaranteed to be moving across the terrain at any time. Still, to the extent that one spends a fair amount of time moving from one place to another in more or less the same area, one is presumably in a reasonably good place for research. On the other hand, the shiny does move, and leads to variation in the terrain one explores.

Following the shiny (FTS) then is a non-random if somewhat obscure way of traversing the countryside before us. And the random-walk aspects of the strategy can be improved upon if one has some general knowledge of the likely organization of the land based on past experiences. I'd like to contrast the pros and cons of this to other heuristics.

To some extent, FTS sits in the middle between two other heuristics. On one pole, we have the “hunker down” heuristic — wherever you are, that's where you stay. This is not at all a bad heuristic, especially if you happen to be in an area that has real scientific importance for a larger view of the whole. There is a reason why a geologist might go to the middle of the desert in Oman to do research. That's because it's one of the only places on the globe where ocean crust (which is qualitatively different from land crust) is above the water.² It's a strategic site for the study of the planet, and you could spend a career there.

When I was in college, my advisor always advised me that in every sociology paper, your second paragraph is about how your topic of study is a “strategic site” for examining the theoretical questions you are interested in. Because it wasn't true in my case — I was interested in the site for other reasons — I soured on the very notion, associating it (correctly, empirically) with dishonest salesmanship. But we shouldn't forget that it *is* great to choose a site because it has unique advantages in untangling a theoretical question. And for some projects, you really need to be there for a long time to make progress — if you want to study geology in the Antarctic, you can't just waltz in and waltz out. You need to go somewhere where others have already established a well-stocked base camp.

So I very much appreciate this heuristic, and I'm actually trying to move in that direction. I don't have as much to say about the opposite pole, which I'll call “airport science”. Here, like the traveling salesman who thinks he's been to all these different cities, but has only seen the airport and the Ramada Inn, we have people who think they've done many different projects because they've dealt with different *easy data applications*. I don't think we should be jealous of work that scholars from other disciplines do even when they poach on what was once our investigatory site, but in this case, I think there is a problem for sociology

2. I am grateful to Joe Martin, who did research in Oman, for telling me about this.

in the fact that we are increasingly expected to compete with pressures from what is euphemistically called “social science.” It’s great to have outsiders bring fresh methods and fresh eyes, but a lot of this is airport science, often done by modelers who just don’t know anything about what they are modeling. (Of course, some outsider modelers know a *lot* about the substance; I continually point to Dirk Helbing as an exemplar worthy of emulation.) But we should go to airports to study airport phenomena (like queuing) — not to do fundamental geomorphology. You study twitter data if you are interested in how people use narrow band broadcast communicative technologies — not to see where ideas come from.

You might think I am simply saying, my wandering is okay, and others’, not. But I know that there is a difference between FTS and airport science, and it is objectively measured in the effort/publication ratio. You know you’re not going from airport to airport, but wandering through swamps and hedges, where you need to put in 1000 hours to get a third tier publication! The key to the airport strategy is that you write papers, but you don’t really solve problems. You can perhaps make a real contribution with this strategy when there is a data set for which slapping an existing method on it will yield a significant substantive finding that others have missed. But that’s pretty rare. Airports, like cities, are pretty crowded places. Even if someone did drop a twenty dollar bill on the floor, odds are you won’t find it.

5 Advantages and Disadvantages of FTS

The “follow the shiny” heuristic, however, has its own problems. Moving from one place to another meant that I spent a fair amount of time doing entry-level work in a host of different tangential fields. That has been immensely rewarding as a way of life, and I can’t say I regret that. And there’s an intellectual upside to this as well. The savage man, says Adam Smith, does many tasks in an uncertain environment, and accordingly is forced to exercise his ingenuity over and over again. Perhaps what he makes is not quite as well made as that composed in a society with a mature division of labor, but he has a great deal more fun.

I’m not claiming that FTS leads to a well-rounded personality (if there is such a thing, I certainly don’t have one); the heart of Smith’s point is a bit different. True, you aren’t completely one-sided, but what is essential about the average person in a non-industrial society, according to Smith, is that he or she is constantly being forced to respond to problems, as opposed to the factory worker incessantly repeating a simple action. I have always preferred to do a project that required a new solution, than one that involved application of well-known solutions — and I think this was an error I am trying to correct. But I want to take with me this key principle. If you choose your projects by your tools — which is an extremely sensible heuristic — you can get very far. But it’s a bit like having one sort of vehicle (such as a big motorboat), and being able to go anywhere that it is suitable for. You can definitely outrace the canoes where things are good for you. The problem is, if the river gets shallower, you get stuck, and if you have to actually port across land, you’re totally out of luck. You have nothing to do but retrace your steps and go home.

An amazing thing to learn is that you aren’t actually a slave to your own technology. I always tell the story of my friend Andy Helck, teaching me auto-mechanics as we worked together on his slant 6 Dodge engine. “We can’t get that out unless we have a tool that has a box wrench on one side, is 25” long, with a 30 degree bend.” Where can we get one? “We can’t.” I assumed that this meant we had to abandon the project. But he continued, “We’re going to have to make one.” Sometimes to solve a problem, we have to visualize a second problem.

Now I’m dubious about our capacity to come up with heuristics for discovery, or even for generating problems, by simply using our noggins. We can come up with quests, but impossible quests aren’t problems. Nor is it really a problem if everything around you is lousy. Problems are things with solutions. And I think that the way we come up with problems is not from deduction or reflexion, but from the *Gestalt* experience of “insight”. Here one sees all at once time a solution, often an indirect one, and therefore a problem appears as a true problem, a *solvable* problem. Part of what the Gestalt theorists identified as intelligence involves the capacity to substitute one problem for another. To return to the geographical metaphor, if I can’t get there from here, from where *could* I come closest? Can I get *there*?

Sometimes we can handle this in a straightforward fashion (thus Polya teaches us to solve proofs by attempting to find the problem that we *can* solve that is closest to this problem). But in other cases, getting to the place from which you can solve the problem involves taking a “detour” (*Umweg*) whereby we seem to go *away* from our goal to get closer. To do this, it appears that it requires that sort of flash of insight —

an appreciation for what the overall constellation of the field affords us — that is resistant to being turned into an algorithm (as C.S. Peirce emphasized).

And so I think that the advantage of doing FTS for at least a while is that by forcing a spirit of invention, doing a little sociological “outward bound” (can you survive in the wild with nothing but an 80486-based computer and FORTRAN?), one sees a wider class of problems. I’m increasingly horrified that, perhaps motivated by the admirable concern for rigor, more and more sociologists seem to think that it is acceptable not only to choose their problems by what their tools are good for — that makes a lot of sense — but even to prefer the wrong tool if it is fancier or of greater precision. People will go to great trouble estimating what they swear up and down is the “right” model, even though it requires that they ask the “wrong” question — for example, they’ll blatantly ignore the most obvious confounders, or use a model that makes strong and false *substantive* assumptions, because this allows them to use a technique that has been declared “best” by someone who doesn’t care about getting the right answer to the question at hand. And often, other methods are going to do far better at approximating the right answer, even if they are cruder. I’m delighted for you if you have a *Lamborghini Centenario*, but if you need to get up a dirt road and it can’t make it, what the hell’s wrong with taking your dad’s old *F250*?³

So what makes something shiny to me isn’t that the substance is “cool” or that the methods are “state of the art” or that if you write about it, you can sell a lot of books. It’s seeing that there’s a way to solve a problem — and for this reason, very often, my ideas come from a particular data structure — in a few cases, actually stumbling across data in books (like a new piece forthcoming, or so I am told, in the *Journal of Social Structure*, analyzing the social networks produced by a sufferer of multiple personality disorder).

But that meant that I got better at knowing how to solve problems in general, than about any particular thing. I don’t feel that bad about this, because so many of the “things” that sociologists seem to know about aren’t really about the social world anyway, in the same way that, in our metaphor, the city isn’t really the same thing as the country — often we sociologists don’t know the world, we just know *knowledge* (at best — putative knowledge at worst). But still, there are some sociologists who really know a lot about something, the way Claude Fischer really knows about Americans’ social relations.

To return to the metaphor, the FTS heuristic can lead one to become a good woodsman — to have a general sense of how to get by outdoors without losing one’s orientation, to have a sense of the regularities in the organization of land and landscape. But one lacks the detailed knowledge of the particular soil conditions and flora that someone who grew up there and did subsistence agriculture would have. And, not having spent enough time in the cities to really have “the knowledge” (as cab drivers in London call a total memorization of all streets), one isn’t really likely to enjoy or contribute to the projects in the good cities. And for the bad ones — the Sodoms and Gomorrah of social investigation — well, they deserve what’s coming to them, and the sooner, the better, I say. So one can be a good teacher, with a stock of valuable lessons disproportionate to one’s years, without yet making any particular discovery.

6 Correcting the Heuristic

And that would be fine, except that there is a question, I suppose we could call it a theoretical question, that’s been nagging me since I started in on sociology in college. It’s a question in what was once understood as the sociology of knowledge but I think is better understood as political psychology — how people understand the social structure they are in, and develop insightful beliefs that allow them to navigate these structures in ways that are ecologically rational, in that they have a correlativity to that structure. (Of course, locally rational beliefs can be 100% wrong and in fact suicidally insane. They’re no less impressive for all that). To me, this is the great wonder of social life, and I don’t believe we really understand the first thing about it. And I don’t think we currently can. Our methods, and our theories, build in assumptions that prevent us from understanding the true meaning of political knowledge (and in this way all social knowledge is political knowledge).

3. As an aside, I am always struck by the way that I get accused of telling people they shouldn’t study, for example, influence, if I say that even with longitudinal data, identification assumptions are too implausible to be worth doing. They are shocked when I suggest that they might need to use an entirely different method, such as participant observation, if they really want to learn about influence. The truth is, probably they don’t. That’s like having to take that hot, dirty bus out of the air-conditioned airport.

So a lot of what I've spent my time on are attempts to figure out what are the theoretical and methodological blocks to us even formulating the right questions here. And while I am no closer to an answer, I do think that, in part because of the new attention to cognition in sociology, and the perestroika of political science, we are in a better position to collectively make substantive progress on what might well be a very good, if difficult, problem. So now it might be time to hunker down.