Introduction: A View on Experimentalism in the Social Sciences

By Alex Rees-Jones

Experiments are central to science. They are a crucial step of the scientific method on which all science, by definition, relies. And in that function, they have generated countless advances in knowledge across remarkably disparate fields of study.

Given the prominence of experiments in science, it should be no surprise that experiments play an important role in *social* science. Though some branches of social science were comparatively late adopters of experimentalism, the use of experiments is now common in the study of social behavior.

Despite this prevalence, academics in the social sciences often impose a peculiar social order surrounding experiments. Though there is near-universal appreciation that randomization is useful for determining causal effects, members of some fields are not near-universally considered "pro-experimental." Instead, an experimental subfield develops. My own subfield, experimental economics, serves as a clear example. To a certain degree, the members of these experimental subfields stand apart from some of their colleagues. At times, these members are even at odds with their colleagues.

Why do social scientists treat experimentalism in this way? Is it because substantial fractions of academics don't understand the scientific method? Many experimentalists would take comfort in thinking that our detractors were so obviously misguided. Unfortunately, I don't think this is the explanation. Instead, I think this opposition is a reaction to a simple but often misunderstood truth: in many social science applications, the experiments run by "experimentalists" often require the data analyst to adopt a number of views that that are more contentious than mere belief in the value of randomization. The label "experimentalist" conveys much more than the mere use of experiments. Understanding that is essential to becoming a good experimentalist, to understanding our role in the world, and to communicating with our critics.

Let me present a simplified example that helps illustrate my point. For brevity, I'll present non-social-science experiments as if they are more straightforward than they truly are. But hopefully it will be apparent that if each sentence of simplification were replaced with 10 pages of more realistic detail, the overarching point would still stand.

Imagine that we are tasked with safeguarding public health. In the course of our duties, we learn that some community of people has contracted some severe new illness. A recent nearby spill of Compound X—a previously unanalyzed chemical—leads us to suspect that it might be the cause. Being scientists, we wish to confirm our belief with the scientific method.

In a situation like this, we often rely on one of two conceptually distinct approach. The first is what I will call a *direct experiment*. The logic of a direct experiment is straightforward. We have some target: people. We have some outcome: the new illness. We have some candidate cause: Compound X. The most direct manner of experimentation is to randomly assign our candidate cause (Compound X) to our targets (people) and observe if this leads to a different rate of our outcome of interest (illness). Of course, in this case doing so requires the experimenters to themselves cause some people to be exposed to a hypothesized risk of harm through their treatment assignment. The most direct version of this experiment would have major ethical

concerns and in many cases would be prohibited. These challenges are not unique to this example: direct experiments are often made infeasible for one reason or another. However, when the results of a well-run direct experiment are available, we can normally count on them to be informative.

The second approach involves what I will call a *Petri Dish Experiment*. When pursuing a Petri Dish Experiment, our goals remain the same: we wish to understand the consequence of some treatment (Compound X) and see if it causes our target outcome (illness) in our target population (people). However, unlike in a direct experiment, the treatment will not be deployed to our target population, and our analysis will not focus on the true outcome of concern. Instead, we focus on some discrete stage of the process that terminates in our outcome of concern. When doing this, we create a microcosm that isolates this stage of the process in a manner that facilitates study. Say, for example, that we think this illness operates by interfering with a particular chemical or biological process in the body. Let's further imagine that this is a process that is well understood in current medical science. We know that if some measurable aspect of it falls outside certain parameters that people get sick. We also know that we can reliably recreate the process in a petri dish. Let's call this process the *causal chain* linking the contents of the petri dish to the illness. To test for a negative effect of Compound X, we may create petri dishes that contain the relevant stage of the process that we believe is disrupted, randomly assign Compound X, and then assess whether that stage is indeed disrupted more often in treated petri dishes as compared to untreated petri dishes. If we see such an effect, our causal chain lets us infer that this will influence the illness that we have been tasked to study. Beyond revealing a connection between Compound X and illness, this whole procedure additionally helps reveal *why* Compound X causes illness, and it does so despite being substantially easier and less contentious to execute than a direct experiment.

While this illustration is cartoonishly oversimplified, I hope you see the point: within the natural sciences, there are plenty of cases where we see some grand phenomenon or "big picture" like this illness, but rather than experimenting on it directly we instead experiment on some small, controlled piece of its casual chain. In that sense, the petri dish itself is not essential to being a "Petri Dish Experiment." The essential thing is that the experiment involves studying a microcosm of the world that isolates a piece of the causal chain. That isolation can happen in a petri dish, or a lab rat, or controlled bits of DNA, or a vacuum chamber, or a particle accelerator, or many other things. For our purposes here, when I write about a "petri dish," this is merely shorthand for such a microcosm.

Petri dish experiments can be extremely valuable. If one is trying to make a complete theory of a phenomenon—a complete causal chain in my language above—Petri Dish Experiments are how one would establish each link of the chain. Additionally, in cases where direct experiments are not possible—say, due to ethical concerns, or when the necessary scale is too large—the right sequences of Petri Dish Experiments can at times allow for experimental progress that could not occur otherwise.

While Petri Dish Experiments can be extremely valuable, these experiments have a special risk of not being valuable at all. This is because the practical value of these experiments is entirely contingent of their connection to a causal chain. To the extent that we have correctly identified a valid causal chain connecting our petri dish to the "big picture," we are in good position to learn from our results. If we don't have such a causal chain to rely on, however, it is much harder to discern the point the exercise. In the context of our example, if we see an effect of Compound X

in a petri dish but we do not know a path to relating the activity in that petri dish to human health, then it is entirely unclear how this finding helps us in our task to address this illness. Perhaps the findings could become useful with future advances in knowledge, but for now they are quite worthless for solving the problem at hand.

With this example laid out, we may now return to the motivating discussion. Why do the social sciences so often separate off an experimental subfield? If we all are scientists, why isn't everyone considered pro-experimental? And what does membership in an experimental subfield really mean?

Does being an experimentalist in the social sciences merely mean that you like or conduct experiments? While this might seem like the natural way to define the term, in my experience this way of defining the label would not successfully separate the experimentalists from the rest of the field. In the areas of social science I know well, there is widespread interest in a good direct experiment. Appreciating or conducting direct experiments does not appear to be considered a sign of being an experimentalist. This is easy to rationalize. The primary results of a good direct experiment are relatively immediately connected to the motivating topic of interest. Upon being convinced that the experiment was run appropriately, a reasonable person should then be willing to acknowledge the revealed relationship between the randomly assigned variable and the outcome of interest. Despite occasional evidence to the contrary, most of our colleagues are reasonable. And thus, for them, learning things from direct experiments is not an issue of tribalism—am I experimental or not?—but just good sense.

In my observation, the subgroups of social scientists branded as experimentalists are those who appreciate Petri Dish Experiments. They typically run comparatively small experiments that are optimized for measuring some important piece of a human or social process, with the understanding that the purpose of this experiment is to inform some phenomenon that is much bigger or broader than what is actually observed in their setting. To give just a few examples, the environment that the experimenter has constructed might involve a much smaller-scale market interaction than occurs in the full economy, or a stripped down and unusual decision compared to what we typically face, or a much smaller-scale and decontextualized social interaction than we see in daily life. Such experiments can be well designed to isolate a very specific piece of an overall puzzle. But of course, the results of these experiments alone are *not* sufficient to make inference about the "big picture" questions that motivate us. These results allow us to infer that an effect occurs in these petri dishes, but one still needs to see the causal chain that connects the findings to our true concern.

This brings us to the crux of the difference between the social science and the natural science applications that often come to mind when thinking about scientific experimentation.

In many areas of natural science, the field is sufficiently well developed than many important casual chains are already in place. We are able to make justified and informed forecasts of how the outcomes in a petri dish translate into impacts on a human, or how some physical interaction in an idealized frictionless environment translates into the predicted motion of an object in atmosphere, and so on, and so on. Even in cases where we do *not* yet have such causal chains fully in place, there is often a widespread confidence and optimism that *we will figure this out*. We have sufficient confidence in the existence of order in the world, and in mankind's ability to learn it, to think that the day will soon come when this question is within our grasp. This attitude has motivated a tremendous amount of immediately important Petri Dish

Experiments. It has also motivated a tremendous amount of “basic science” research that might not have been obviously useful at the moment it was completed, but which became extraordinarily useful later when new and often unanticipated causal chains were worked out. It is hard to look at the most focal areas of the natural sciences and not see evidence of enormous returns for Petri Dish Experiments.

In the social sciences, we also have a significant history of successes of the scientific method. However, despite our successes, we still lag far behind the natural sciences. And of course, how could we not? Compared to the natural sciences, most social sciences are quite young. We have not had comparable time to develop and to learn from the mistakes of youth. And furthermore, many of the central phenomena that social scientists study have a fundamental complexity far beyond that seen in many natural-science causal chains. Our ideal causal chains would often require a complete theory of how a particular social interaction would scale to a large society in richly varying circumstances that go beyond what exists in our petri dish. Broadly speaking, our theory is not there yet, and it is not clear how soon it will be. Given these challenges, I take no shame in admitting that social scientists generally have fewer fully vetted causal chains to work with compared to their natural science colleagues, and in general those that we have are shakier and vulnerable to attack. Plenty of people—and indeed, plenty of our colleagues—take the pessimistic view that the state of our knowledge is not yet at a point for many applications of the petri dish approach to be sensible.

So what is essential to the identity of experimentalists in the social sciences? Perhaps most important is adopting the belief that, at least for some questions of social importance, there is indeed enough of a causal chain to latch onto. That it is indeed possible to set up your petri dish, your microcosm, in a way that can both yield valid results on that link of the chain, but also in a way that connects to a longer chain necessary to make your petri dish directly valuable. Adopting this view requires optimism, in that you must take a comparatively rosy view about the order inherent in human and social interactions and about our ability to discern that order. *Rationally* adopting this view requires a great deal of subject-matter expertise, developed for carefully chosen topics. It requires picking some phenomenon of interest and developing full command of the associated literature. This allows you to map out the causal chains that have been proven, or that seem on the verge of being proven, or that are often assumed but in reality are probably invalid, and so on. For some topic areas, you will find that the area is still insufficiently developed to provide the type of causal chains you need to make a career as an experimenter. For some other topic areas, I believe, you will find a rational foundation to proceed. And thus, through the careful and informed selection of an area of expertise, you can put yourself in the position to rationally adopt the optimism you need to pursue Petri Dish Experiments and execute them well.

In my view, this way of thinking about experimentalism is embodied beautifully in this Handbook. This work contains many chapters on valued topics, all of which I think are top quality. And each chapter communicates the key pieces of the authors’ experience designing successful experiments in that topic area. As you will notice, there are many commonalities in the attitudes and approaches of the different authors. But more interesting are the cases where authors disagree. One might initially look at this negatively: how can these people disagree on the basics? Is the field is not worked out? I look at it differently: people working on different topics, or working within different camps, are running experiments that connect to different causal chains for relevance. And the way you design an experiment must, of course, depend on that.

When preparing a petri dish, a growth medium that is necessary for one organism might be poison to another. The preparation of the petri dish is dependent on the question that is being studied and the causal chain that will be latched upon. So too for experiments in the social sciences.

The fact that the best advice about conducting experiments depends on your topic of study presents a major challenge to training new experimentalists. In practice, students are often trained by someone who learned how to design good experiments in their topic area, but who does not have complete knowledge of how to design good experiments in other settings. And furthermore, there are relatively few resources available to get such advice. In my view, this results in experimental practices passing down a little too much like folk history and too little like organized science. In some ways it is quite useful learn how to run experiments in this way, as it inherently results in one learning how to run experiments optimally in some setting. But a reasonable student should find ways to avoid being too dependent on the views of one person latched to a particular causal chain.

This handbook is meant to help fill this need. The overarching goal has been to collect articles about a variety of issues that are common when running experiments in the social sciences (with some significant tilt towards economics). In doing so, I've asked authors to write a bit less "academically" than usual. Rather, I have asked authors to frankly explain their thinking about experimental design in their area. What issues do you view as critical in a design, what issues less so? How do you approach the whole enterprise? All the chapters in this handbook do a great job of addressing these issues. If one's goal is to learn how to run an experiment in an area, this lets you see how successful experts in that area think about it, and what particular arrangements they view as necessary to set up a successful petri dish and to map their type of petri dishes to the big picture.

Whether you are at the beginning of a path to becoming an experimentalist or already well along it, I hope seeing these authors' approaches helps you along the way. I encourage you to proceed with the optimism that experimentalists need, but to make sure it derives from broad assessment of existing knowledge and not faith. The work the authors have shared here should help you along that path. They have certainly helped me. Good luck.