

Maxims for Experimenters

Martin Dufwenberg*

2/15/10

How do you produce a good experiment? How should you write up the results? The answers may depend on the goal of the exercise. While there may be more than one sensible formula, I have my clear favorite. I present it here, alongside four corollaries on what not to do if one follows the formula. I also make a related proposal for how to submit papers to journals, with the results in a sealed envelope. I suspect that my call for such submissions will not be heeded, but I have a back-up plan and therefore finally make a request to editors & referees who evaluate experimental work.

I. The formula

Economists are interested in how the world works. They wish to understand for example the causes & effects of unemployment, how the rules of the Nasdaq effects stock prices, and how vengeance may influence hold-up in business partnerships. Since economic outcomes are shaped by how humans reason as well as the principles by which groups of individuals reach outcomes in games, economists care about these matters too.

Theorists tell stories about how the world works. It is natural to wonder what stories are empirically relevant. In this connection experiments may be useful, as they can help provide answers. My advice to students who want to do such experimental work is this: Follow the

$$I \rightarrow D \rightarrow H \rightarrow R$$

formula! Here...

- *I* stands for *idea* – a statement about how the world works, often but not necessarily derived from theory.
- *D* stands for *design* – a set of experimental conditions and treatments meant to be optimal for shedding light on the empirical relevance of the idea.
- *H* stands for *hypothesis* – a statement formulated in terms of the data generated by the design, which would support the empirical relevance of the idea.

* Department of Economics & Economic Science Laboratory, University of Arizona; martind@eller.arizona.edu. I thank Gary Charness for helpful comments.

- Finally, *R* stands for *result* – the evaluation of whether or not the data generated by the design supports the hypothesis & the idea.

II. What not to do

The formula comes with four corollaries regarding what not to do:

- Don't get (too) curious about changes to the design!

The point here is not to confuse interest in the design with interest in the idea. Under the formula, the design has its *raison-d'etre* only as a tool for shedding light on the empirical relevance of the idea. Therefore, making a change to the original design to see what happens is not useful, unless one somehow can argue that this change would shed light on the idea.

- Don't think it is always desirable to 'explain' the data!

Under the formula, the goal of your endeavor is to test the idea not to understand the data. Suppose, for example, that the idea you test is a theory which abstracts away from some aspects (or 'frictions') known to be relevant to some extent. Suppose the data provides some but not perfect support for the theory. You may be done, if explaining the frictions were never part of your research goal. Or suppose the data does not support the theory. That may be an interesting finding, and again you may be done. I'm not saying that explaining the data is never useful, only that one shouldn't take for granted that it is.

- Don't decide on the idea after seeing the data!

Since the design is a response to the desire to test an idea, it is meaningless to run a design without an idea to test (if one follows the formula). And if ideas are specified *ex post* there are two problems: First, it is unlikely that the design will be optimal for testing the idea, since the design was not constructed with the idea in mind. Second, there is risk of outcome bias in the perception of causality.

- Don't worry about the data!

Running an experiment is similar to decision-making under uncertainty. One wants to make decisions that maximize expected utility, and in an uncertain world one can't rationally always hope to make the decisions that turns out to be best ex post. For example, drawing to an inside-straight at poker without proper odds is a sucker play that loses money in the long run, even if every now and then the straight is made. Similarly, experimenters should run the experiment they deem to have the greatest scientific merit viewed from an ex ante perspective. With respect to testing theory, the value attached to the research you are doing (in terms of publishability) shouldn't depend on the nature of the data.

This last corollary comes with a caveat. What if the editors and readership of our journals are more interested in certain kinds of results than others, say results that are deemed surprising? This may skew researchers' incentives, as maximizing the probability of a surprising result need not be the same thing as choosing the project with the highest expected scientific value from an ex ante point of view. And if non-surprising results do not make it into journals this may skew the picture of how the world works that emerges from published research. These problems lead me to a proposal:

III. The sealed-envelope-submission strategy

When you write up a research paper and submit it to a journal for publication...

- Don't mention results in the abstract!
- Don't mention results in the introduction!
- In fact, don't mention the results at all in the submitted paper. *Put the results in a sealed envelope!* Ask the editor and the referees to make their call before opening the envelope. Only once they have decided whether or not to publish the paper, they may open the envelope, study the data, and read your summary.

IV. A request to editors & referees

Will journals give serious consideration to submissions with the results in sealed envelopes? Will editors resist the temptation to sneak a peak? Might they even be convinced to insist that manuscripts be submitted this way? I believe the answer is: probably not.... If so, then my call to editors & referees is this: *Evaluate the research you consider as if you had to form an opinion before looking at the data!*