

Teaching a Behavioral Economics Elective: Highlighting the Science of Economics[†]

By TED O'DONOGHUE*

In this short paper, I provide an overview of how one might teach an advanced undergraduate elective on Behavioral Economics. I have taught such an elective at Cornell University since Spring 2005, and this paper will focus on the structure and themes from my own course. That said, I will attempt to highlight ways in which instructors might choose alternative structures for their courses.

I. Highlighting the Science of Economics

When I teach Behavioral Economics, the main goal that I have in mind throughout the course is that I want students to gain an appreciation for the science of Economics, and that I am really using Behavioral Economics only as a vehicle to accomplish this goal.

The typical textbook definition of Economics is often rather dry, something along the lines of “Economics studies the allocation of scarce resources.” But Economics is really a methodological approach for addressing real-world problems—an approach for how to make sense of the outcomes that we see in the world, and an approach to investigate what might be done to improve the outcomes that we see in the world.

This approach—which we often label the science of Economics—involves a strong interplay between theory and data. We develop theories for what we think might be driving economic outcomes, and we empirically test those theories on data. These empirical tests reveal the ways in which theories succeed, but also reveal the ways in which theories fail, which in turn leads us to

develop modified (or completely new) theories. And then the process continues. In this process, no theory can ever be proven “correct”—indeed, no theory is even intended to be a complete and perfectly accurate representation of the world. The best we can do is to assess the usefulness of a particular theory, or to compare the usefulness of two theories, in helping us to understand economic outcomes.

On the very first day of class, I discuss the science of Economics, and I highlight to students that, while Behavioral Economics contains very interesting material, what I really want them to leave the course with is an increased appreciation and understanding of Economics. As the semester proceeds, I repeatedly remind students of this broad goal.

II. General Course Structure

Even when I first designed my course for Spring 2005, there was easily enough material to fill two semester courses. Today, there is even more material available. Hence, anyone teaching a course in Behavioral Economics must make some decisions about what to cover.

For Spring 2005, my approach was to focus on the four main domains in which Behavioral Economics had made significant progress by that time:

- Choice under uncertainty (risk preferences).
- Choice over time (time preferences).
- Choice over interpersonal allocations (social preferences).
- Judgment under uncertainty (probability assessments).

As I reviewed the literature and assessed what I might cover, it occurred to me that if I divided the course into four equal parts, I would need to provide a very shallow treatment of each domain. In order to better accomplish my goal of

* Department of Economics, Cornell University, 482 Uris Hall, Ithaca, NY 14853 (e-mail: edo1@cornell.edu). For helpful comments, I thank Alex Rees-Jones, Brian Dillon, Saurabh Bhargava, and participants in our AEA Session.

[†] Go to <http://dx.doi.org/10.1257/aer.p20151048> to visit the article page for additional materials and author disclosure statement.

giving students an appreciation of the science of Economics, I decided that I should instead focus on two topics, and provide a detailed treatment of each. I also decided that I wanted to cover models of decision making that could easily be applied to a range of economic contexts, and to cover applied research that could be nicely tied to those models. Based on these criteria, I chose to cover risk preferences and time preferences. I have continued with this approach ever since, and it really permits me to demonstrate in detail the evolution of thought in each domain.

Within each of the two main topics, I follow the same general structure:

Step 1: Review/learn the standard model (e.g., expected utility or exponential discounting).—In order to appreciate the evolution of Behavioral Economics research, students must start with a strong understanding of the standard model, and the ways in which the standard model has proven very successful. In the process, I emphasize an important theme: behavioral economists do not think the standard model is bad; rather, we think it is great, but it could be made even better.

Step 2: Discuss evidence that contradicts the standard model (typically experiments).—I force students to be extremely precise about whether evidence does (or does not) contradict the standard model. Moreover, I highlight some evidence that does not contradict the standard model, but rather is useful in informing how one might develop an alternative model.

Step 3: Develop an alternative model motivated by that evidence (e.g., prospect theory or present bias).—As we develop an alternative model, I force students to assess carefully the extent to which the model does and does not capture well the evidence that we discussed. I also work through some general (and more abstract) predictions of the alternative model so that students gain an appreciation of the details—e.g., I have them work through some simple abstract examples of choosing between lotteries or when to complete a task.

Step 4: Apply that alternative model to economic contexts (experiments and especially field data).—For each topic, I cover a series of

examples where the literature has attempted to apply the alternative model to some economic context. For each, I encourage the students to think critically about how well the alternative model does and does not work in that context.

Step 5: Consider yet new models, if they exist (e.g., Köszegi-Rabin loss aversion).

In the current incarnation of my course, a detailed coverage of both risk and time preferences covers about 75–80 percent of the semester. At the end of the semester, I provide a brief overview of social preferences and probability assessments—just to give a flavor for those domains—and I also spend a little time discussing public policy implications (framed in the context of risk and time preferences).

Of course, over the past decade, the Behavioral Economics literature has expanded even further, and there are new domains that could be covered in an undergraduate course. Given more time, I might also cover salience and attention, or identity, or research using happiness data. One might also choose to frame a course in terms of applications, perhaps focusing on retirement savings behavior, or behavioral public policy. In principle, in terms of the material covered, there are now many ways one could structure a Behavioral Economics elective.

III. Why is Behavioral Economics a Particularly Good Course in which to Illustrate the Science of Economics?

There are a number of reasons why a course on Behavioral Economics is a great place to showcase to undergraduates the science of Economics.

First, Behavioral Economics exposes students to a wide range of economics methodologies. They get to work through precise models of decision making, and to think very carefully about which types of assumptions drive which types of conclusions. They get exposed to experimental evidence, which then permits discussion of experimental methodologies and the advantages and disadvantages of experimental data. They get to develop applied theoretical analyses of economic applications—to see how one can take a model of decision making and embed it within an economic application and derive implications. They get to study field evidence,

which permits discussion of empirical methodologies and the advantages and disadvantages of working with field data. Finally, they get to investigate policy implications, and the connections between models of decision making and different policy prescriptions.

Second, and closely related, because Behavioral Economics often focuses on models of individual choice, students can be exposed to the various economic methodologies at a level that they can handle. In particular, they can be exposed to precise and rigorous methodologies that do not require graduate work.

Third, Behavioral Economics is a great domain in which to demonstrate the difficulties of doing applied theory, and the split personality that is required. Applied theory requires a first stage in which one reduces a complex field context into an economic model. This stage is often very loose and imprecise, and requires a lot of judgment and intuition (and sometimes “hand waving”). Once this first stage is complete, however, our other personality takes over, and economists are super precise in how we analyze the model that we have developed. Especially in Step 4 from Section II, students get to observe and discuss this difficult process over and over again.

Fourth, Behavioral Economics can help students appreciate that Economics is not rigid and narrow. Economics is constantly incorporating new ideas in a natural and scientific way, and can even be open to ideas from other disciplines. Moreover, Economics can be applied quite broadly, and not merely to the types of questions that have typically been considered within the domain of Economics.

Finally, because Behavioral Economics has blossomed so recently, students can observe the science of Economics at one of its current frontiers. Indeed, my students really enjoy the fact that they are reading very recent journal articles. Moreover, because the ideas are so new and often not fully fleshed out, students have the opportunity to think critically about the material—to not accept models and analyses at face value, and to throw out their own ideas for how to address open questions or improve economic analyses.

IV. Other Pedagogical Decisions

For those looking for guidance in designing their own courses, there are several other dimensions on which one must decide how to

shape a course. This section describes a number of such dimensions.

A. *Level of Methodological Rigor*

The course that I teach requires intermediate microeconomics with calculus, and thus students can be pushed quite a bit in the level of methodological rigor. However, I am confident that the general structure and themes that I describe above could easily be followed in a course that is somewhat less methodologically rigorous.

At the same time, my course does not require econometrics, and thus I am, unfortunately, forced to be somewhat less methodologically rigorous when covering empirical results. As more and more empirical analyses of field data appear in Behavioral Economics, it would be more and more useful to require econometrics as a prerequisite.

B. *Economics versus Psychology*

The course I teach is very much an Economics course, and it is very much not a Psychology course taught by an economist. While I do discuss evidence from Psychology at various points in the course, that discussion is really a rather small part of the course (maybe 10 percent). The emphasis of the course is using ideas from Psychology to do better Economics—indeed, related to my main goal of giving students an appreciation of the science of Economics, I repeat again and again that our end goal is “to do better Economics.”

In recent years, I have taken to explicitly emphasizing this point on the first day of class. During the early years in which I taught the course, there was a small group who expressed disappointment with the course. One of the two main sources of disappointment was from students who were expecting a course that would primarily cover Psychology research related to Economics. Such students complained that they did not expect to do so much Economics.

That said, each instructor of a Behavioral Economics elective must decide how much time to allocate to Economics versus Psychology research. Indeed, if I had the time—perhaps in some future course I teach, perhaps co-taught with a psychologist—it would be a valuable endeavor to get economics undergraduates

to critically read and assess research from Psychology.

C. Behavioral Finance

Closely related to the field of Behavioral Economics is the field of Behavioral Finance. Much as it is hard to draw a line between Economics and Finance, it is also hard to draw a line between Behavioral Economics and Behavioral Finance. Each instructor of a Behavioral Economics elective must decide how much time to allocate to Behavioral Finance.

In my course, I allocate relatively little time. As discussed in Section II, I wanted my course to cover applied research that could be nicely tied to underlying models of decision making. I cover a few Behavioral Finance papers that satisfy this criterion. However, I do not cover research from Behavioral Finance that is more focused on observed anomalies without a serious assessment of underlying decision making that might drive those anomalies.

I emphasize the decision of how much Behavioral Finance to cover because it is related to the second main source of student disappointment during the early years of my course. This came from students who were expecting a course in Behavioral Finance. Hence, much as for my comment about Psychology research, I also explicitly emphasize on the first day of class that the course will cover relatively little Behavioral Finance.

D. Use of Experiments in Class

In principle, a Behavioral Economics course lends itself well to conducting experiments in class—especially for Step 2 in my general structure from Section II. While I used to run a simple time preference experiment in my class, in recent years I have chosen not to incorporate any experiments into my class. In part, the reason is that I am not an experimenter myself, and I always felt that my time preference experiment was run in a very sloppy way (of course, this might also have been a good lesson for the students). More importantly, because I emphasize so much the goal of doing better Economics, and because I view the main goal of Economics to be understanding field behavior, I wanted to tilt the class more in the direction of discussing field behavior (especially as more Behavioral

Economics research using field data has become available).

Each instructor of a Behavioral Economics elective must decide how much time to allocate to experiments. Especially if one wanted to give students a deep appreciation (and perhaps even some training) in experimental methodologies, one could devote a large chunk of time to experiments. But if one does choose to incorporate experiments, it is worth the investment to do them well.

E. Ancillary Lessons for Students

While I did not design the course with these lessons in mind, over the years student feedback has pointed out two important lessons that students have taken away from my course. First, many students felt that the course gave them a much better understanding of their own behavior, along with strategies to help them alter their behavior in ways that they viewed as desirable. This lesson is most obvious for the topic of intertemporal choice and procrastination, but it also seemed to show up in how people think about the risks in their lives. Second, many students felt the course gave them a better appreciation of others' behavior, which they found helpful in their interpersonal relationships.

In recent years, I have acknowledged these lessons—I briefly mention them on the first and last day of the semester. But I do not emphasize such lessons much as we move through specific topics. I could imagine instead teaching a Behavioral Economics elective that emphasizes and focuses more directly on these types of lessons.

F. Textbook

When I designed the course in Spring 2005, there did not exist a textbook that covered Behavioral Economics. Hence, I instead asked students to read primary journal articles. For some research, I was able to find journal articles written at a level that undergraduates can understand well. For other research, I told students to focus on a subset of the paper—often the introduction and initial setup. Most important, I told students that they would need to rely heavily on their class notes, and that attendance would thus be quite important. While I have received occasional complaints about the lack of a textbook, in general this approach seems to have worked well.

There now exist a few textbooks on Behavioral Economics (e.g., Angner 2012; Wilkinson and Klaes 2012; and Just 2013). While none of these textbooks match my vision for my course, they might prove useful for others in designing their courses.

G. Assignments

Students really seem to love Behavioral Economics, and, despite it being a lot of work, my course is typically quite large (100–150 students). Because my course is so large, I limit myself to assigning problem sets and giving exams. I am quite demanding on problem sets, where I ask students to work through some complex models. On exams, I include a combination of working through simple models (much simpler than on the problem sets), along with short questions that require verbal intuitive discussions. I emphasize to students that, much as Economics requires a combination of rigorous analysis and eloquent verbal arguments, I require both on my exams.

More generally, I wish I could assign more written work. Behavioral Economics lends itself well to having students provide a critical assessment of recent research, or having students discuss how some aspect of Behavioral Economics is relevant to some aspect of their own life. For those who can teach a smaller course, such assignments could be very productive.

H. A Graduate Course

Before I conclude, I briefly discuss how one might teach a graduate elective in Behavioral Economics—in part because doing so further illustrates the nature of my undergraduate course. Much as for undergraduate courses, there is easily enough material in Behavioral Economics to fill several semesters of graduate electives. The graduate course that I teach covers a similar set of topics as my undergraduate course—I start with risk preferences, then cover time preferences, and finally cover a few other topics. The main differences between my graduate and undergraduate courses are:

- I teach the graduate course at a higher level of methodological rigor.

- I teach the graduate course with even more focus on the historical evolution of the literature (since I feel it is important for graduate students to have some historical perspective).
- I teach the graduate course with even more focus on the flaws in existing research and open questions to be addressed.
- The assignments in my graduate course are more geared toward helping students make the transition from being a student to being a researcher (e.g., referee reports, research proposals).
- In my graduate course, I feel less need to convince students to be excited by Economics.

V. Final Thoughts

Behavioral Economics is an ever expanding field, and there are a multitude of ways in which one might take a chunk of the literature and turn it into a Behavioral Economics elective. In this short paper, I have described the way I organize my course and the major themes of my course. While this approach has worked well for me, others will need to make their own choices to create a course that fits their own styles. Toward that end, I have also discussed a number of dimensions on which one must make a pedagogical decision when designing a course.

My own course is constantly changing. I try to cover one or two new pieces of research every year, and every few years I am forced to squeeze out certain topics. In fact, I have hit the point where I hope in the very near future to develop a second Behavioral Economics elective with a somewhat different focus—exactly what, to be determined.

REFERENCES

- Angner, Erik.** 2012. *A Course in Behavioral Economics*. New York: Palgrave Macmillan.
- Just, David R.** 2013. *Introduction to Behavioral Economics*. Hoboken: Wiley.
- Wilkinson, Nick, and Matthias Klaes.** 2012. *An Introduction to Behavioral Economics*. New York: Palgrave Macmillan.