**Carnegie Mellon**

73-390 Behavioral Economics
Spring 2006 Syllabus

Room: Schaife 219 — Lecture: T/Th 3:00 – 4:20
Professor P.J. Healy — Email: phealy@cmu.edu
Office: GSIA 331 — Office Hrs.: Mon 3–4 PM

# 1 About This Course

## 1.1 Books

There are *no requied textbooks* for this course. All required readings will be provided.

The following are good books for reference and understanding. Some of these should be on reserve at the library.

- *Advances in Behavioral Economics* by Camerer, Loewenstein & Rabin
- *The Handbook of Experimental Economics* by Kagel & Roth
- *Experimental Economics* by Davis & Holt
- *Games and Decision Making* by Aliprantis & Chakrabarti
- *Thinking Strategically* by Dixit & Nalebuff
- *Introduction to Game Theory* by Morris

## 1.2 Blackboard

This course will use the online Blackboard system to track grades, post announcements, and distribute required readings. If you cannot regularly access the Blackboard, please alert the professor.

## 1.3 Modules

The course will operate in a series of 13 'modules', one per week. In each module we will study a particular topic or phenomenon. The module will start on Thursday with an introduction to and demonstration of the topic and conclude the following Tuesday with presentations and discussions of related research and models. Specifically, the planned structure is:

**Thursday (45 min)** Participate in in-class experiments to demonstrate a phenomenon.

**Thursday (45 min)** Lecture on the predictions of the 'standard' theory. Some discussion.

**Weekend Homework** One or two readings are assigned. One student will prepare a presentation of the readings and an outline for discussion. All others must prepare a one-page summary of the readings, *including their own comments & opinions.*

**Tuesday (30 min)** Student presentation of the readings.

**Tuesday (30 min)** Open discussion on the readings, related ideas.

**Tuesday (20 min)** Lecture on additional related work, summarizes.

## 1.4 Assignments & Grade Weights

**40%** Each student must present and lead discussion at least once. Presentations will be graded.

**10%** Students should participate in others' discussions.

**20%** Non-presenting students must turn in a weekly readings summary with opinions.

**20%** All students must write a research proposal.

**10%** All students will give a short presentation of their proposal.

### 1.4.1 Presentation & Discussion

The second class of each module will open with a student presentation. This presentation and discussion outline must be prepared in advance.

1. The presentation should:
   (a) be well-prepared and organized,
   (b) summarize any important experimental or empirical fact (ie, data) of the readings,
   (c) summarize any theories and models provided by the paper,
   (d) discuss if/how the theory or data might apply to other situations,
   (e) and include some insights from related work (and perhaps knowledge from other courses).
2. The prepared discussion should:
   (a) include an outline or list of questions to be distributed to the class,
   (b) encourage the class to think about and discuss the 'sticky' issues related to the readings,
   (c) and help lead the class into unstructured (but constructive) dialogue.

Also note that audience members are graded based on their participation in the discussion.

### 1.4.2 Weekly Summaries

A healthy discussion *cannot* occur unless everyone does the required reading. The weekly summaries ensure that everyone has done their work and can contribute to the discussion. The summary should, *at least*:

- be at least a full page long,
- summarize the major points of the readings,
- and provide some of your own opinions about the research methodology, the general agenda, or the results. (Be careful to distinguish opinion from fact.)

Each weekly summary will receive a grade from zero to five.

### 1.4.3 Research Proposal & Presentation

Pretend this course runs for two semesters instead of one. In the second semester, we would work in two groups on new research papers that build on what we learned in the first semester. These research papers may build on one topic or bridge across topics. Your job is to 'sell' your idea to the professor and your classmates via a research proposal and presentation to be given the last week of the first semester. After the presentations, the class will vote on the best proposal and the top two proposals will be declared the winners.

A good proposal:

- is between 5 and 10 pages,
- is broken into 3 to 6 coherent sections, such as (for example) 'Introduction', 'Related Literature', 'Proposed Experiment', 'Hypotheses', and 'Conclusion',
- has a title page with the author's name, title, and an abstract of <100 words that summarizes the proposal,
- properly surveys related literature to show that (a) your idea hasn't been done, and (b) your idea fills an interesting hole in the literature,
- provides *in very clear detail* what exactly you propose to work on,
- has a list of hypotheses about what you expect to find,
- and provides *in very clear detail* what resources you'll need (money, subjects, computer lab, etc.).

Written proposals will be due 24 hours before the proposal presentations begin.

Each student must give a 15 minute proposal presentation during the last week of class. The proposal presentation must summarize your written proposal to the class.

## 1.5 Tentative List of Topics

1. Decisions: Endowment Effect
2. Decisions: Expected Utility, Heuristics & Biases
3. Decisions: Overconfidence & Underconfidence
4. Decisions in Time: Myopia and undersaving
5. Regard for Others: Dictator, Ultimatum, Trust & Public Goods Games
6. Mixed Strategy Equilibrium: Tennis & Soccer
7. Repeated Games: Learning & Reputation
8. Labor Markets: Fairness & Reciprocity
9. Labor Markets: Incentive Schemes
10. Asset Markets: The Bubbles Mystery

11. Auctions: Overbidding & Revenue Equivalence
12. Auctions: The Winner's Curse
13. Optimal Contract Design: NASA

## 2 Defining Behavioral Economics: History & A Parable

> "All economics rests on some sort of implicit psychology. The only question is whether the implicit psychology is good or bad. We think it is simply unwise, and inefficient, to do economics without paying some attention to good psychology"
> Colin Camerer and George Loewenstein [2002, Advances in Behavioral Economics]

Behavioral economics is best defined not explicitly, but rather by how it deviates from 'standard' economic theory. To understand this, we should look at the historical development of economic theory.

Prior to the mid-1900s (and definitely before the late 1800s), economics was essentially a branch of philosophy. Writers who attempted to describe economic behavior and prescribe economic policy understood that behavior is linked to both extrinsic incentives (such as profit) and psychology. Even Adam Smith, who is generally credited for realizing that extrinsic incentives often lead to socially desirable outcomes, understood that humans are subject to emotions, biases, and irrationality. Although much progress was made in economic thought, theories were based on intuition and observation and not subject to rigorous proof or falsification.

Over the past century (and particularly over the past fifty years,) the powerful tools of mathematics have been brought to bear on economic problems. A key difficulty in this agenda is the development of a simple, parsimonious theory of human behavior without sacrificing external validity. Clearly, human behavior is sufficiently complex that no workable mathematical model can fully capture its many nuances. Instead, the authors proposed that economic agents are endowed with some intrinsic objective function (for example, utility or profit) that they seek to maximize. This approach is appealing for many reasons. First, it reduces behavior to an optimization problem, which mathematics is well equipped to handle. Second, we do not need to explicitly state what the objective function is or from where it comes – we only need to assume its existence, perhaps make some simplifying assumptions about its structure, and proceed with the math. Third, if the objective function is a reasonable approximation of the person's preferences, happiness, or goals, then any agent who doesn't aim to maximize this objective function would do better by following what the mathematical theory prescribes. Finally, the maximization approach is surprisingly accurate in certain settings and seems to be a model that isn't too far 'off the mark' in terms of describing behavior.

The development of the maximization paradigm received its biggest boosts in the development of game theory (starting in the 1940s), statistical decision theory, and mathematical economics (in the 1950s.) As these tools were initially developed, authors were very careful to note that they were making strong and clearly imperfect assumptions about human behavior. As the fields developed, however, the usual caveats about the inaccuracy of assumptions were dropped from the written literature and the maximization approach became a paradigm. Although most good theorists through the 1970s and 1980s understood that the huge body of mathemetical economics, decision theory, and game theory was developed on this imperfect model, few were successful in altering its core assumptions. The approach became standard and the assumptions unquestioned.

It is also critical to point out that the maximization paradigm, though simple in structure, is widely applicable because it is purposefully vague. Models often assume that agents have *some* utility function over possible outcomes and that they make choices to maximize the value of this function. To make the maximization problem tractable and to match some of the most basic observations from reality, the utility function will likely be assumed continuous, increasing, and concave. Beyond this, the theory is silent about the actual functional of the utility function. As a consequence, this approach is often open to the critique that virtually *any* observed behavior could be generated by *some* profile of utility functions that are continuous, increasing, and concave. What is phenomenal about the theory is that it makes a huge number of predictions despite the vagueness. The other benefit, of course, is that if one happens to know more structure about the utility function, then one can quickly generate unambiguous predictions.

While this theoretical paradigm was emerging and strengthening, a handful of economists and, separately, a handful of psychologists began to understand the value of using controlled laboratory experiments to test the assumptions and results of the newly developed theory. A key contribution is the development of the *induced value* approach, whereby human subjects make economic decisions and are paid for the simulated outcomes proportionally to the (hypothesized) value of the actual outcomes that the experiment simulates. For example, if we want to study bidding behavior in auctions, we can run a simulated auction in the laboratory where subjects are paid more money if they win the item than if they do not. Although the scaling of the incentives may be much smaller, paying subjects for performance allows the researcher to understand behavior in the face of incentives.

The development of experimental economics was slow at first, but grew quickly through the 1980s and 1990s. Although a huge number of interesting results have been generated from these induced value laboratory tests, the entire body of literature can be roughly summed up in one conclusion: The maximization paradigm is a surprisingly good predictor in competitive market-like environments, but often becomes inaccurate when the environment allows for strategic gamesmanship, when one's decisions directly affect another's payoffs, when the decision task is complex and limitations on rationality make maximization difficult, or under any of a variety of other conditions that have been identified. In some cases, the most basic predictions of maximization were violated in the laboratory. On the whole, the maximization approach seemed to be a good (and in market-like situations, a very good) approximation of behavior, but it became obvious once again that the theory was predicated on inaccurate assumptions.

The body of experimental evidence created an urgency to re-examine the basic model. As a consequence, theorists and experimentalists began to develop new theories of behavior to explain the various phenomena observed in the laboratory. It is this research agenda that is now called *behavioral economics.*

The appeal of behavioral economics is that its models and theories are able to capture the phenomena that the maximization paradigm could not. Furthermore, a good behavioral economics model properly incorporates knowledge from psychology – something that the standard approach almost universally ignores. Unfortunately, behavioral economics is a terribly fractured field. A different model was developed to explain each phenomenon, but rarely was one model applicable to domains outside those for which it was developed. This difficulty usually comes from the model making strong assumptions that generate the desired prediction, and those strong assumptions make the model too inflexible to be ported. For these reasons, many economists resist the behavioral economics approach. While the maximization paradigm is clearly flawed, it has been very well developed and its domain of application is nearly universal.

In retrospect, the fractured approach of behavioral economics is perhaps not surprising. The psychology literature teaches us that modeling human behavior naturally leads to the development of context-dependent theories. Models that seem to capture behavior in one setting often have little to say about behavior when that setting is changed. Any theory that attempts to bridge various contexts is likely less accurate in each context than the individual theories that were developed for each.

Consider the following playful story that illustrates how research proceeds in the face of these difficulties. Imagine a set of data on a variable $x$ and a variable $y$. In reality, $x$ and $y$ are related by the formula $y = x^5$. Researchers observe various values of $x$ and $y$ and attempt to develop a simple model of their relationship. Suppose the 'standard' model specifies that $y = f(x)$, where $f$ is an arbitrary polynomial. Because the mathematical tools of the day are limited, our researchers cannot work with models that are concave for some values of $x$ and convex for others. So, they usually proceed by assuming $f'(x) \geq 0$ and $f''(x) \leq 0$ for evey possible $x$. Clearly, the standard model is incorrect for $x > 0$, where the actual data is convex.

Now suppose a pioneering researcher gathers values of $x$ and $y$ for values of $x$ between -10 and -20. She finds that the data are increasing and concave – perfectly consistent with the standard model. Years later, a second researcher gathers data for $x$ between 0 and 5. He finds the data to be convex, declares the standard model to be incorrect, and proposes that $y$ must equal $x^2$ for $x > 0$. This model provides a better fit for these new data and is easy to work with because its second derivative does not change signs, but it simply does not apply to $x < 0$. Thus, he is forced to study the $x^2$ model only for $x > 0$. Subsequent researcers study more extreme values of $x$ and find that the $y = x^2$ model predicts values of $y$ that are far too low when $x$ gets large. Furthermore, they find that the relationship between $x$ and $y$ is approximated well enough by a line with very steep slope. Alternative models are developed that predict $y = x^2$ for $x$ near 0, but then that $y$ is a steeply-sloped linear function of $x$ for extreme values of $x$. All of the researchers continue to assume that the standard model is accurate for $x < 0$. Research continues in this way, building new simple models for various ranges of $x$ and sticking them together to give an incomplete picture of the true relationship.

Are our hypothetical researchers flawed in their approach? With enough domain-specific models they will approximate the true relationship, although without the full model, it is not likely that they will completely understand the exact implications of the true relationship. On the domain of negative values of $x$, our researchers continue to apply the standard model despite the fact that it only works if we add the caveat of 'for $x \leq 0$'. Although this all seems off-base, our researchers have their hands tied. Until their modeling techniques enable them to deal with both concavity and convexity, they must be content with the current agenda. The real benefit of their explorations and model-building is that they have begun to identify exactly where their standard model falls short and in what directions it should be re-built. In the mean time, only more data gathering can help them approximate the true relationship.

This parable roughly describes the current state of economics and behavioral economics. Experimental methods identify cases where the standard model works and those where it does not. Where it doesn't, simple models are developed to help understand how behavior works on that limited domain. With enough studies and models, we begin to paint a rough picture of reality. Simultaneously, these explorations invite discussion (and argument) about the strengths and weaknesses of the standard model. While more and more data is collected and behavioral models constructed, a new push develops to broaden the standard model to help explain what has been observed. The latter process is slow, indeed, as it is not immediately clear (yet) what needs to be discarded, what needs a slight adjustment, and what works perfectly well. Ideally, this interplay between the two approaches will benefit the discipline by improving both our understanding and our models.