CAMBRIDGE
Cambridge University Press & Assessment
978-1-009-42233-8 — Principles of Behavioral Economics
Sanjit Dhami
Excerpt
More Information

Introduction

Taught undergraduate and graduate programs in economics largely conform to the dominant paradigm, *neoclassical economics*, that models “real” humans through the lens of a “fictitious” being, *homo-economicus*. Throughout this book, we use the term “neoclassical economics” to refer only to its “actual practice” in mainstream economics rather than the original intent and motivation of its founding fathers.

Homo-economicus is an entirely *self-regarding*, *amoral*, *emotionless*, *utility maximizing* being that possesses *unbounded rationality*. Thus, homo-economicus will happily break a promise, lie, cheat, and violate social norms, if that leads to greater net expected monetary gains. Lacking emotions, this being also experiences no guilt or shame from such actions. Homo-economicus does not consider the morality of its own actions, or the actions of others, or the fairness of procedures, provided that its actions maximize material gains. In making decisions, homo-economicus believes that any other individual is also, well, another homo-economicus.

Unbounded rationality is slightly harder to define without invoking formal technical jargon to first define rationality, which we do later in this book. Unlike the usage of the term “rationality” in common discourse and in some of the other social and behavioral sciences, in economics rationality has a clear definition with falsifiable implications. Unbounded rationality requires *unlimited attention*; *unlimited and perfect memory*; *unlimited computational ability* that matches the fastest computers; *unlimited knowledge and use of existing mathematics and statistics*; and *unlimited ability to see through alternative methods of framing identical information*. But this does not exhaust the sense in which we may use the prefix “unlimited” to the mental gifts that homo-economicus possesses, as we shall see in various parts of the book.[1] These features of homo-economicus are typically considered relatively uncontroversial in courses in economic theory, and indeed in much research in economics.

All science proceeds on the basis of observations and intuition to formulate testable hypotheses. These hypotheses are then stringently tested with the aim of falsifying, not confirming, them. It was entirely reasonable for neoclassical economics to proceed in the manner laid out above, from the early part of the twentieth century onwards. Homo-economicus was a useful first hypothesis or model of actual humans. However, when tested stringently, many of its central assumptions and predictions failed the empirical test, often quite dramatically.

Neoclassical economics relies on *marginalist principles*. Roughly, the optimal action taken by homo-economicus in any given situation ensures that the marginal benefit of the action equals its

[1] Alfred Marshall’s *Principles of Economics*, published in 1890, and Paul Samuelson’s *Foundations of Economic Analysis*, published in 1947, are important early texts on neoclassical economics. These texts laid the foundations of what was to come later. While Marshall relied on a verbal/diagrammatic method, Samuelson used the language, methods, and techniques, from mathematics and statistics, that have since become characteristic of the subject.

marginal cost. This, of course, requires homo-economicus to be aware of the marginal benefits and marginal costs of all actions. The Hall and Hitch Committee Report was published in 1939 by a group of Oxford economists, and was based on interviews with UK businessmen. It showed that firms do not use marginalist principles (in this case, marginal revenue equals marginal cost) in deciding on their prices or quantities.

The evidence accumulated from the work of several Nobel Prize winners in the Economic Sciences has rejected the homo-economicus model and some of the central theories in neoclassical economics. We outline some of this work below and discuss it in the book in much detail, although a great deal of it finds, at best, an honorable mention in a footnote in many courses in economic theory.

Allais (1953) showed that under risk and uncertainty, humans do not follow *expected utility theory,* the main decision theory under risk in neoclassical economics. In his Nobel Prize acceptance speech in 1978, Simon (1978), who also invented Artificial Intelligence in his other avatar, pointed out that there is no evidence that firms produce at the point where marginal revenue equals marginal cost. This is almost as significant as a Nobel Prize winning physicist arguing in their acceptance speech that atoms are a figment of our imagination. Simon also made important contributions to the literature on bounded rationality; for a summary, see Simon (2000).

In the early 1970s, Tversky and Kahneman (1971, 1973, 1974), showed that many important predictions of the homo-economicus model were decisively rejected; for a lucid introduction, see Kahneman (2011). Kahneman and Tversky (1979) then showed, in the second most cited paper in economics, that expected utility theory is rejected, and proposed the best known alternative in behavioral economics under risk and uncertainty, *prospect theory*. Reinhard Selten, who made important contributions to classical game theory, also expressed an interest in experiments in the 1960s, and in modeling bounded rationality. He remained, throughout the rest of his life, a strong and vocal critic of the homo-economicus model (Selten, 1998, 2001).

In the 1980s and beyond, Richard Thaler showed that some of the key assumptions of the neoclassical model are violated. People do not calculate opportunity costs very well, and cannot write budget constraints in the usual way (total expenditure equals total income) due to *mental accounting*. For an engaging account of Thaler's diverse contributions to behavioral economics and behavioral finance, see Thaler (2015). In the early 1980s, Akerlof (1982) showed that fairness played a key role in the determination of market wages and tried to integrate theories from sociology into economics. Shiller (1981) showed that the efficient markets hypothesis, a cornerstone of finance, is rejected by the empirical evidence, so stock markets are not as efficient as they were thought to be.

It would be remiss not to mention the work of Smith (1962) who showed that double auction experiments lead to outcomes consistent with a competitive general equilibrium. However, Smith "assigned" values of objects to subjects; that is, an experimenter determines how much an object, say, a mug, is worth to a subject rather than the subjects determining the value of the object. By contrast, when subjects are free to assign values to objects, which appears to be the more natural method, systematic departures from the competitive equilibrium outcome are found, as in the classic *endowment effect* experiments (Kahneman et al., 1991; Thaler, 2015).

In addition, many influential pioneers within economics, several who are likely to join the elusive Nobel club, have conducted pathbreaking work in highlighting the serious shortcomings of the basic neoclassical framework. You will have ample opportunity to be introduced to their pathbreaking work over the course of this book.

In science, there is only one tenable response to the problem of lack of conformity of the data with the predictions of a theory, once all confounds have been accounted for. The theory must either be reformed, enriched, or altered, to account for the evidence. If it is non-salvageable, then it should be jettisoned. The empirical refutation of some of the key features of neoclassical economics set the stage for the development of *behavioral economics*, arguably the most significant development in economics, perhaps of all time.

The rest of this introduction attempts to put the reader in the right frame of mind to read this book. This does not necessarily take the form of explaining behavioral theories and the associated evidence; the book does enough of that. Rather, it takes the form of thinking through a few central questions. How should economists practice their craft? What is good science? What is the role of, and the relation between, theory and evidence? What are some of the potential misconceptions about behavioral economics?

In addition, we spend a bit of time outlining the philosophy behind the book and how one might design a good, well-rounded, course in behavioral economics. Lest some of the material in this chapter looks like sweeping assertions, note that references and details are kept to a minimum whenever the corresponding material is discussed and referenced later in the book.

0.1 The Broad Road Map of Behavioral Economics

In first year courses, economics is typically defined as *the science of allocation of scarce resources among competing uses*. I often cringe when I hear this – indeed I was taught the same definition in my first year in university. This is an unnecessarily narrow and mechanistic vision of economics; and a misunderstanding of the scope of the social and behavioral sciences, of which economics is an important part.

For the purposes of this book, the reader may use the following definition of behavioral economics.

Definition 0.1 *Behavioral economics is an interdisciplinary approach to the social and behavioral sciences that draws insights from economics, psychology, sociology, anthropology, evolutionary biology, sociobiology, and neuroscience. It aims to address the following two central questions.*

(1) How do humans exercise their judgment and make decisions?
(2) How and why do humans cooperate, compete, and coordinate, in small and large groups?

Behavioral economics draws on a rich body of interdisciplinary theory and evidence on human behavior. By contrast, neoclassical economics has been conspicuous in its neglect of the other social and behavioral sciences. Economists are sometimes surprised at how much work has been done in the other social and behavioral sciences on each of the central questions posed above. For instance, judgment and decision making, which is highlighted in the first question, is already an important and thriving part of psychology and management. The second central question, in particular, broadens the scope of behavioral economics beyond the allocation of scarce resources to make it consistent with the general scope of the social and behavioral sciences.

The broad approach, and methods, of behavioral economics may be described as follows.

Behavioral economic theories aim to derive clear, testable, falsifiable, predictions in a rigorous manner, employing as much mathematics and statistics as is required. Behavioral economics is committed to jettisoning theories that are rejected by the data, once all reasonable confounds have been accounted for. It uses data from the lab, field, surveys; observational data; and data using neuroscientific methods. It does not necessarily assign any hierarchical preference among the source of the data in forming inferences about human behavior, and in testing competing theories.

The first part of the approach outlined above shows that the methods of behavioral economics, that is, the mathematical and statistical models used to derive clear, testable, and falsifiable predictions, are rooted in neoclassical economics. However, the second part, a commitment to getting rid of rejected theories, is much stronger relative to the current practice within neoclassical economics (more on this below). The final part of the approach shows that behavioral economics relies on richer and novel sources of data relative to neoclassical economics, which typically uses only observational data, and sometimes views the other data sources with suspicion.

Behavioral economics is still relatively young. Some of the ideas can be traced back to a long lineage, extending back to Adam Smith's *The Theory of Moral Sentiments*, published in 1759, and developments in the related social and behavioral sciences; readers interested in a historical account can consult Camerer and Loewenstein (2004) and Ashraf et al. (2005). The work by Maurice Allais and Herbert Simon in the 1950s, followed by important work on behavioral theories of the firm in the 1960s by Cyert and March (1963), was largely ignored in mainstream economics, despite its importance and originality. However, the key developments that led to an explosion in the modern literature on behavioral economics occurred throughout the 1970s, in great part due to Daniel Kahneman and Amos Tversky, and gathered pace thereafter.

Behavioral economics does not always provide a fully satisfactory account of all the empirical evidence that we have at our disposal, nor has it been applied to every problem that neoclassical economics has been applied to. However, it has already been applied to a very wide range of problems, and when pitted against the neoclassical predictions, its record in explaining the data is significantly better. Its ability to explain the empirical evidence on human behavior is unprecedented in economics, and ignorance of its subject matter ought no longer be a viable option for any serious economist.

This book is an introductory account of the principles of behavioral economics, and a record of its predictions, when pitted against the data. It also demonstrates the fruitfulness of a constant feedback between the evidence and theory, which is one of the hallmarks of the natural sciences.

Sometimes a distinction is drawn between *behavioral economics* and *experimental economics*. Given the close feedback between theory and evidence in behavioral economics, such distinctions are artificial for most practical purposes. Experiments, unless they are underpinned by the relevant theory, may lack discipline and direction, although it is also true that in some cases there is no established theory yet to account for important experimental evidence. Conversely, behavioral theory, unless it is supported by the data, is not tenable.

0.2 What Constitutes Good Science?

The key to doing good economics, whether in research or teaching, relies on understanding what constitutes good science. I would encourage students and instructors to spend some time thinking about these issues that lie at the heart of their subject.

0.2.1 The Scientific Method

The scientific method is essentially a series of prescriptions in order to do good science. A researcher typically begins with some observations about the real world, or conjectures about a phenomenon. The researcher then generates predictions based on known theory, or new theory that the researcher may formulate. In order to test the predictions, the researcher then forms clear hypotheses to either explain the observations, or verify the conjectures. The hypotheses, based on the relevant theory, specify the circumstances under which the predictions of the theory will fail or succeed. This ensures that the relevant theory is falsifiable, a minimum requirement for any theory. Finally, the researcher will design experiments to collect data, or use existing data, to test the theory using appropriate statistical methods.

The tests are designed to *falsify rather than confirm* the theory; this is the sense in which they are "stringent tests of the theory." Any particular dataset may confirm the theory for now, but the theory might well be rejected by future data. Scientific progress is achieved as newer, and better, theories are formulated that explain everything that the older theories could explain, but also explain some new phenomena that the old theories could not explain. We may never discover the objectively true underlying theory, but so long as we successfully construct a succession of theories that do better at explaining the data, we make progress. Relative to the "true" underlying theory, the succession of intermediate theories may well be "wrong" but this is not a valid argument to avoid continually formulating better, or less wrong, theories. We come back to this point below because it has been misused in economics to protect existing theories from testing and refutation on the grounds that "all theories are wrong anyways."

The description of the scientific method above is roughly consistent with the justly celebrated work of Popper (1934, 1963). Subsequent work by Kuhn (1962) argued that scientific progress arises by periodic paradigm shifts (abrupt changes in theories) and science does not accumulate in a linear fashion. Kuhn describes a nonlinear process of scientific progress in which "paradigm changes" eventually occur as anomalies begin to accumulate against existing theories, leading to a tipping point where the existing theories become untenable.

Meanwhile several important developments took place in the Popperian position, which make it, in the view of this book, the only tenable framework to follow as good practice in economics. The Duhem–Quine thesis points out that a test of any theory is a joint test of the theory and the auxiliary assumptions. For instance, the data might reject a theory, not because it is incorrect, but because the experimental instructions were flawed, or the subject pool was inappropriate, or proper incentives were not offered. Hence, a single test of a theory may not reject it. One must, therefore, control for all reasonable auxiliary and confounding factors before one concludes that a theory is rejected.

In an important development in the Popperian tradition, Lakatos (1970) proposed *the methodology of scientific research programs*. He distinguished between the non-expendable hard core of a research program and an expendable auxiliary part. As anomalies to the research program begin to surface, the hard core is protected while the auxiliary assumptions are modified to address the anomalies. If such a process leads also to the generation of novel predictions, then it is referred to as *theoretically progressive*. If the novel predictions are not refuted by new data, then the research program is termed as *empirically progressive*. If, however, anomalies to the research program continue to accumulate and the program loses its theoretical and empirical progressivity, then the research program is doomed and likely to be replaced by a new research program.

0.2.2 The Scientific Method in the Natural Sciences

The scientific method, described above, is the critical blueprint on which the natural sciences are based and this should, in principle, apply equally well to economics. However, there is an unfortunate, and uncalled for, divergence of views between the natural sciences and economics on the purpose of the relevant theory. In the natural sciences, the purpose of the relevant theory is to explain the natural world, and researchers try hard to falsify existing theory by designing stringent tests to fail the theory. For, if a theory were to fail, then it provides scientists with an exciting opportunity to understand the real world better and drive further progress in their field. This takes the form of a relatively quick process of accepting refutations of the existing theory and then a stiff competition among competing research groups throughout the world to formulate new theories. For instance, there was a great deal of excitement in the scientific community when the large Hadron Collider was to begin running experiments, in the hope that the shortcoming in existing theories would be revealed, leading to important new advances.

In particular, there is no religious adherence to existing theories in the natural sciences. Let us see what this means. In one of the most celebrated results in physics, Einstein formulated the theory of relativity in 1915. He proposed three tests of the theory. The first, related to observations of the orbit of the planet Mercury, which Newtonian physics failed to explain and relativity could explain, was a known problem. The second, bending of light as it passes through the gravitational field of a stellar body, was a new prediction that needed to be tested, and it was clear how to test it. The third, gravitational redshift (electromagnetic waves traveling out of a gravitational field lose energy leading to an increase in their wavelength) was also a new prediction, but it was completely unclear how to test it.

In a letter to the *London Times* on 28 November 1919, Einstein wrote: "The chief attraction of the theory lies in its logical completeness. If a single one of the conclusions drawn from it proves wrong, it must be given up." In the event, the second prediction was confirmed by new experiments in 1919, but "proper" confirmation of the third prediction took many more years with significant breakthroughs in 1954 and continuing on to recent years. In every single test, not a single deviation from the predictions of the theory of relativity was found.

In economic theory, there is no comparable example to Einstein's clear road map to reject his own theory. Yet Einstein's approach is not uncommon in the natural sciences, and it is a major reason for the success of the natural sciences. In particular, natural sciences have inbuilt mechanisms that do not allow theories to be shielded from refutation, or to hide behind defensive positions.

It can take many years, decades, and sometimes centuries, to develop fully satisfactory theories in the face of constant empirical challenges. For instance, the germ theory of disease was initially proposed in 1549, and then through a series of discoveries, placed on a much firmer footing around the 1890s, particularly following the seminal work of Louis Pasteur and Robert Koch. The molecular basis of life, DNA, was first discovered by a Swiss chemist, Johann Friedrich Miescher, in the 1860s. Yet, he resisted in publishing his findings until 1874 as it was generally believed that the nucleus was too small to contain the building blocks of life. Finally, the structure of DNA was discovered about 80 years later in the 1950s by Francis Crick and James Watson, although several other teams, particularly the one led by Rosalind Franklin, appeared to be arriving at the same result. In other words, scientific discoveries are the results of long periods of persistence and grind.

One last example. The Higgs Boson particle was discovered in 2012, and hailed as the "God Particle" by the popular media. However, the theory behind it was originally published in 1964 by Peter Higgs, among others. However, it took decades to build the necessary equipment, in this case the large Hadron Collider, in order to test it. Peter Higgs received the Nobel Prize in physics in 2013, but he did not receive the prize until his theory had been stringently tested and confirmed. This is reflective of the standards for receiving a Nobel Prize in the natural sciences.

It would, however, appear from the track record of past winners of the Nobel Prize in the Economic Sciences that somewhat different standards of conformity of theory with the empirical evidence apply. A retrospective analysis of past winners in economics to examine what fraction of the winners have not had their theories rejected by the empirical evidence might not make for exhilarating reading. Curiously, in 2013, the Nobel Prize in the Economic Sciences was awarded jointly to Eugene Fama and Robert Shiller. Shiller received it for showing that financial markets are inefficient and Fama for showing that financial markets are not efficient; although both also made other notable contributions that were likely taken into account for the prize.[2] A situation such as this is unlikely to arise for the Nobel Prize in the natural sciences.

0.2.3 The Scientific Method in Economics

Many of you are likely to be pursuing an economics degree. Have you ever wondered why your typical microeconomics courses, or your game theory courses, hardly contain a shred of empirical evidence that offers "stringent" tests of the theory?

At best, unless you were taught by an exceptionally enlightened instructor, it might have been mentioned that people act "as if" they behave in conformity with the taught models. The "as if" condition does not require the factual accuracy of all the assumptions made in the model; such as checking if the relevant bordered Hessian in your optimization problem when you buy your next meal at the university cafeteria is negative semi-definite. The "as if" condition holds if, despite making assumptions that might not be literally true, the predictions of these models are indistinguishable from actually observed behavior. For instance, the moon, sun, and earth are

[2] To be sure, both winners also made sterling contributions in other areas too, deserving of the prize.

satisfactorily modeled for many purposes as just three point masses (the so called "three body problem"), even if this is not literally true.

The "as if" assumption is central to neoclassical economics, and it was forcefully advocated by Friedman (1953) whose views on methodology are widely subscribed to in economics. Friedman (1953, p. 14) argued:

> Truly important and significant hypotheses will be found to have "assumptions that are wildly inaccurate descriptive representations of reality, and, in general, the more significant the theory, the more unrealistic the assumptions .... To be important, therefore, a hypothesis must be descriptively false in its assumptions."

These insights form the basis on which most neoclassical economic theory has been constructed and it is the basis on which many economic theorists practice their craft. In principle, the "as if" assumption is not problematic, provided, and this is the critical caveat, the predictions are stringently tested against the data.

One of the enduring contributions of Kahneman and Tversky in their work in the 1970s was to show that the predictions of neoclassical economic models do not hold even in an "as if" sense (Tversky and Kahneman, 1971, 1973, 1974; Kahneman and Tversky, 1979). A mountain of evidence has supported and added massively to the work of Kahneman and Tversky in subsequent years. This book gives an account of this evidence. However, this work is conspicuous by its absence in the typical course in economic theory.

Consider again the failure of most of your microeconomics and game theory courses to mention the relevant empirical evidence. How are corresponding "theory" courses in the natural sciences taught? Should you expect a course in, say, theoretical physics to stay relatively silent on the relevant empirical evidence? Here is a quote from one of the leading social and behavioral scientists of our times who is also well versed in physics. Gintis (2009, p. xvi) mentions:

> Economic theory has been particularly compromised by its neglect of the facts concerning human behavior … I happened to be reading a popular introductory graduate text on quantum mechanics, as well as a leading graduate text in microeconomics. The physics text began with the anomaly of blackbody radiation, …. The text continued, page after page, with new anomalies … and new, partially successful models explaining the anomalies. In about 1925, this culminated with Heisenberg's wave mechanics and Schrödinger's equation, which fully unified the field. By contrast, the microeconomics text, despite its beauty, did not contain a single fact in the whole thousand-page volume. Rather the authors built economic theory in axiomatic fashion, making assumptions on the basis of their intuitive plausibility, their incorporation of the "stylized facts" of everyday life, or their appeal to the principles of rational thought …. We will see that empirical evidence challenges some of the core assumptions in classical game theory and neoclassical economics.

What accounts for this unusual state of affairs? The chief contributory factor has been the inability of neoclassical economics to internalize and practice the scientific method. In its place, the economics profession has adopted a range of homespun methodological positions. In some cases, these homespun positions reflect a misunderstanding of the purpose of economic theory; in other cases they are overly pessimistic of the possibilities for economic models. In yet other cases, they appear to serve a vested purpose in protecting existing economic models from refutation. Whatever the underlying reason, this has in many cases been pernicious for progress in developing neoclassical economic theories that are in sound conformity with the evidence.

Here is a set of quotes, taken from published papers by some of the leaders in economic theory, that amply demonstrate the range of views, yet they share a common intellectual position that goes back to Friedman's "as if" position. It is worth stating these views in some detail because they are so widely subscribed to in economics, yet they have little basis in the scientific method.

> Hence the choice of a model will depend on the purpose for which the model is used, the modeler's intuition, and the modeler's subjective judgment of plausibility .... One economist may reject another's intuition, and, ultimately, the marketplace of ideas will make some judgments. Dekel and Lipman (2010, p. 264).

> In particular, we agree that: economic models are often viewed differently than models in the other sciences; economic theory seems to value generality and simplicity at the cost of accuracy; models are expected to convey a message much more than to describe a well-defined reality; these models are often akin to observations, or to *gedankenexperiments*; and the economic theorist is typically not required to clearly specify where his model might be applicable and how. Gilboa et al. (2014, p. F. 516).

> As in the case of fables, models in economic theory are derived from observations of the real world, but are not meant to be testable. As in the case of fables, models have limited scope. As in the case of a good fable, a good model can have an enormous influence on the real world, not by providing advice or by predicting the future, but rather by influencing culture. Yes, I do think we are simply the tellers of fables, but is that not wonderful? Rubinstein (2006, p. 882)

> Theories are never correct, and in the case of the social sciences they tend to be almost always wrong. The question, therefore, is not whether they are right or wrong, but whether they are wrong in a way that invalidates the conclusions drawn from them. Gilboa (2009, p. 104, in his Section 8.1 titled "Theories Are Always Wrong")

In these views, the arbiter between competing theoretical models is not the empirical evidence, but the "marketplace of ideas." Economic models are also justified as having a "special status" distinct from models in the natural sciences in several respects. They are expected to convey a message, not describe a well-defined reality, or even be testable; theorists are not even required to specify where their models might be applicable; and economic models are akin to fables that are designed to influence culture. I suspect that most neoclassical economists will agree with these views to varying degrees.

All this makes it very difficult, if not impossible, to distinguish between competing theories designed to explain the same reality. This is a recipe for the proliferation of "styles" and "fashions" in economic models; ad-hoc untestable explanations of supposed stylized facts; and loose hand-waving to cherry pick supposed stylized facts that are consistent with a given theory. None of these views wishes to expose economic theories to stringent tests designed to reject theories, as compared to tests that confirm them. Indeed, since all theories in the social sciences are presumed to be wrong anyways, and only required to be internally logically consistent, the claim is that there is no point at all in stringently testing them with the data. This book does not subscribe to any of these views.

One argument for demanding a special status for economic theories has been to plead that they are difficult, if not impossible, to test. Furthermore, it is argued, testing of theories is easier in the natural sciences where the subjects are atoms and natural phenomena, rather than idiosyncratic humans and human behavior. Nothing could be farther from the truth. These arguments reveal ignorance about empirical testing in the natural sciences and the development of empirical regularities in behavioral economics.

Testing can be incredibly difficult and challenging in the natural sciences, as evidenced by our brief discussion in Section 0.2.2. But scientists exploring, say, the chemical composition at the core of a distant star, or the origins of life, or the origin of the Universe, or what causes particular diseases, did not plead for a special status for their subjects. They got on with the very difficult and dogged job of finding the answers, sometimes taking centuries in the process to arrive at a satisfactory answer.

It is true that human behavior is very different from the behavior of atoms. It is, for instance, more heterogeneous. But that does not prevent us from constructing economic theories that allow for this heterogeneity. Physicists did not throw up their hands in despair when asked to model particles in Brownian motion. They got on with developing the appropriate theoretical tools that would allow them to predict the probability that a particle was at a particular location at a given time. This can be tested, and when tested, the theory was confirmed. Nothing prevents economists from engaging in such an exercise.

Economics can, and should, do better in two ways. First, by accepting that the purpose of economic models is to explain a well-defined reality. Second, by showing willingness to subject their theories to stringent empirical testing, with a variety of methods and sources of data, and by jettisoning empirically rejected theories. This allows one to make a choice between competing theories, and point to the way forward, based on the only tenable arbiter – the relevant empirical evidence.

But this is largely what the emerging field of behavioral economics is already engaged in, and this task is facilitated by the close feedback between theory and experiments. Students taking a *well-designed course in behavioral economics* (more on this below) are often palpably excited by the ability to understand human behavior in such a comprehensive manner.

0.3 Data and Methods in Behavioral Economics

The ability to perform controlled experiments is the hallmark of all good science. Experiments not only inform us which of the competing theories explains the data, they might themselves provide the basis for future theoretical developments. Economics has been a relative latecomer to controlled lab and field experiments. It has typically relied on using observational data to draw inferences about actual behavior, or to test theories. Such data is often collected for entirely different purposes (e.g., government statistics collected for annual reports), hence, it is a challenging task to control for all possible reasonable confounds and to ensure exogenous variation in just the right variables. Despite the use of sophisticated econometric techniques, in actual practice it is doubtful how well all possible reasonable confounds can be controlled for.

Data collected from lab and field experiments provides the gold standard in science that gives rise to true exogenous variation in one variable while keeping everything else fixed. *Randomized control trials* (RCTs) in medicine are an example of such field data. This process was played out in public during the recent developments of new vaccines for Coronavirus as clinical trials for several candidate vaccines began in 2020. A control group was given a placebo and a treatment