UNLEARNING AND DISCOVERY

by Charles F. Manski*

I find introspection to be painful. Looking back, I recall doing little of it before age sixteen. Instead, I mainly did the things that I was instructed to do. I studied hard and performed well in courses and examinations, for no deeper reason than that this was what one was supposed to do. As an adult I have introspected when it has seemed constructive to do so, but I still try not to overdo it. I prefer to accomplish something tangible—to write a paper or wash dishes—than to contemplate.

Introspection is usually a private matter, but writing an essay such as this requires introspection for public consumption. Why agree to do it? Ego may play some role, but communication is the ostensible social objective. Humans learn from observing the experiences of others, and use the information to inform their own decision making. Thus, this essay provides data on my experiences that may perhaps prove useful to other economists.

There will be a theme to the essay. I have repeatedly found that I am able to make new discoveries only after I unlearn orthodoxies and go back to basics, with an open mind. I will relate a sequence of experiences of this sort. I begin at age sixteen, next describe my path through college and graduate school, and then discuss two intense mid-career episodes of unlearning and discovery. I conclude with a very recent episode whose implications are not yet clear to me.

1. Becoming Agnostic

If a person thinks at all, he or she must give some thought to the origin of the universe, the presence of some sort of god, and so on. My agnostic views on these matters have no significance per se. However, it may be informative to recount how I chose to be agnostic forty-five years ago. This experience set an important precedent for my research.

I grew up in Boston in a traditional but not devout Jewish family. I can't say why, but early on I took the religion and its rituals more seriously than my parents did. By age thirteen, I mastered the challenge of reading the weekly Torah passage from the ancient scrolls and I later led youth services each Saturday morning.

One such Saturday, in spring 1965, I found myself at the lectern chanting the Hebrew words of a prayer to the thirty or so other teenagers at the service. I write that I "found myself at the lectern" because my conscious mind had been elsewhere for several minutes, working through a math problem, while the words of the prayer were exiting my lips, on automatic pilot. When I left the math problem and rejoined mind and body, I experienced an epiphany. I recognized for the first time that the prayers meant nothing to me. I had memorized the words and could endlessly recite them by rote. However, I now was aware that I did not care what the words meant and that I had no god in mind as the recipient of my prayers.

This moment of revelation initiated a painful period of introspection. True believers in a religion or a scientific doctrine may shrug off episodes of doubt and hold on to their faith. But I was not a true believer. To be honest with myself, I decided that I would have to scrutinize the comfortable religious worldview of my youth. The outcome was a complete dismissal of the certitude of faith. It was not just that I rejected the tenets of religious Judaism. I learned to be skeptical of all dogma, of any doctrine proclaimed to be true without proof.

I also learned something about why dogmas can be tenacious and irreconciliable. Many doctrines pose nonrefutable hypotheses. That is, they make statements about the world that are impossible to disprove. For example, it is impossible to disprove the hypothesis that the god of the Torah created the universe in six days and then rested on the seventh day. It is similarly impossible to disprove the hypothesis that the universe was created by the Flying Spaghetti Monster. If an Orthodox Jew were to challenge me to disprove the literal word of the Torah, I could not do so. Neither could I disprove

* Department of Economics and Institute for Policy Research, Northwestern University February 2009

the tenets of the Church of the Flying Spaghetti Monster. If true believers in these religions were to debate each other for eons, neither would be able to disprove the beliefs of the other.

Finally, I learned to live with ambiguity. Recognizing that multiple doctrines pose competing nonrefutable hypotheses generates a choice problem. One might declare faith in a particular doctrine, rejecting all others. One might place a subjective probability distribution on the space of all doctrines and behave so as to maximize expected utility. Or one might reason only that the truth is unknowable, concluding that one must then make decisions under ambiguity.

I chose the last option in religion, and have tended towards it in economics as well. My research on partial identification, to be discussed later in this essay, has emphasized that available data and credible assumptions often imply only weak conclusions. Observing that competing hypotheses often are nonrefutable, I have noted that debates between researchers who maintain different nonrefutable hypotheses can persist forever without resolution. And I have argued against resolution of ambiguity by declaring faith in one hypothesis, dismissing other plausible hypotheses that are consistent with the available evidence. Manski (1995, Chapter 1) and Manski (2007, Chapter 1) state these broad themes. The subsequent chapters of these books, as well as many journal articles, provide technical analysis that fleshes out the implications for empirical research and decision making.

2. Becoming an Economist

Most of my early research and some of my more recent work has concerned empirical analysis of choice behavior. I have sought to contribute both by developing new econometric methods and by studying particular choice settings. Among other things, I have stressed the importance of data that characterize the choice sets decision makers face and how they perceive their opportunities. I have devoted particular attention to analysis of schooling and career decisions.

With this background, I find it natural to want to provide the reader of this essay with an informative account of the sequential choice process through which I became an economist and developed my research focus. I do so here, covering the seven-year period from college entry in 1966 through completion of my Ph.D. in 1973.

I entered MIT as a freshman in fall 1966 with a vague intent to do some sort of work in the physical sciences. At the time, hardly anyone enrolled as an undergraduate at MIT with the social sciences in mind. I knew nothing about economics and was not aware that the Institute even had an economics department, much less one of the best in the world.

A central positive feature of the American collegiate system is that it encourages, indeed requires, students to explore multiple fields of study before choosing a concentration. In the spring of my freshman year, I enrolled in an elementary economics course to fulfill a social science/humanities distribution requirement. I found that I took well to the basic themes of resource allocation, analysis of choice behavior, and policy formation.

Not yet knowing what career direction to take, I declared a joint physics and economics major in my sophomore year, and simultaneously enrolled in the required courses in both disciplines. The summer following my sophomore year I realized that, although I had performed well in both sets of courses, economics seemed natural while physics did not. I was particularly bothered by the fact that I had received A grades in courses in relativity and quantum mechanics by successfully doing the math, while having no intuition for the subject matter. I then decided to continue in economics alone.[1]

As an undergraduate, I focused mainly on public economics. My interests were both theoretical and applied. I learned about normative welfare economics from Jerome Rothenberg, my undergraduate thesis advisor. I learned about pragmatic policy analysis from a summer job with a state higher education agency.

By the spring of my junior year, I had decided to pursue a Ph.D. in economics. However, for this to be feasible, the non-academic real world would have to cooperate. This was the spring of 1969. Much was happening in the United States at this time, but the aspect of the late ‘60s most on the mind of males of my generation was the Vietnam draft.

As an undergraduate, I opposed American policy in Vietnam but supported the broad need for national defense. I did not want to be drafted and sent to Vietnam, but I thought it honorable to serve in the military in a useful capacity. To avoid the former and accomplish the latter, I enrolled in Air Force ROTC as a freshman.[2]

I was for some time content with the idea of serving the required four years as an Air Force officer following college graduation, after which I could begin the rest of my life. However, as I began to develop plans for a Ph.D., the opportunity cost of military service became clearer. Yet dropping out of ROTC would mean subjecting myself to the draft after graduation, and I had a very low draft lottery number. I did not know what to do.

Then I experienced what an economist would term an "unanticipated positive shock," or an ordinary person would call "pure luck." I was due for a periodic appointment with a dermatologist to treat the chronic but minor ailment of psoriasis. While being examined, I jokingly asked the physician if my condition would render me medically unsuitable to be drafted. I was astonished when he replied in the affirmative. After I confirmed that he was correct, I formally dropped out of ROTC. Instead of seeking a temporary student deferment, I allowed myself to be reclassified as subject to the draft and waited to be called for a pre-draft physical at the Boston Army Base. Armed with a letter from the dermatologist and the empirical evidence that I did, indeed, suffer from psoriasis, I failed the physical and was free.[3]

It is rare for someone to do their Ph.D. at their undergraduate institution, but I remained at MIT. This was not according to plan. The faculty who advised me as an undergraduate thought it would be best if I were to move up the street to Harvard for the Ph.D. My academic record was such that they were pretty certain I would be admitted. However, Harvard rejected my application and I chose to stay on at MIT, after considering Chicago and Berkeley as alternatives. E. Cary Brown, the longtime chair of the MIT economics department, told me that Harvard had rejected my application, as well as those of all other applicants from MIT, as part of an ongoing feud between the two departments. I was thus introduced to faculty politics.

Whereas my experience as an undergraduate in economics at MIT had been highly stimulating, I found the first semester of the Ph.D. program to be stifling. With the exception of a wonderful half-semester taught by Bob Hall, the core courses in micro and macro theory gave a perspective on economics that was decades out of date. Moreover, the material was taught as received wisdom set in stone rather than as the state of the art in a living science.[4]

The spring semester was a far better experience. Indeed, it introduced me to the two large ideas that have been most central to my research, identification and discrete choice analysis.

I was introduced to identification by Franklin Fisher, who taught the core course in econometrics. Fisher, who later chaired my dissertation committee, focused on the *simultaneity problem,* the classical econometric problem of inference on supply and demand functions from observations of market transactions. I was not fond of the dry tools of linear algebra used by Fisher and other econometricians to study identification of simultaneous linear systems. However, I was fascinated by the general problem of inference on quantities of economic interest by combining available data with suitable assumptions.

I was introduced to discrete choice analysis by Daniel McFadden, whose presence at MIT that spring was the second large unanticipated positive shock to my professional career, comparable in importance to release from the threat of the draft. Following my undergraduate summer job with a higher education agency, I had become interested in a specific applied problem of choice analysis: forecasting how tuition and financial aid affect college enrollment. The available econometric tools for analysis of individual choice behavior were linear demand models developed to study the neoclassical setting in which a consumer allocates a budget among various commodities, choosing how many units to purchase of each commodity. Such models had previously been used to analyze the "demand for higher education," but they seemed inappropriate to me. As I viewed the matter, students were not choosing how many units of a homogeneous good called college to purchase. They were making a discrete choice among various colleges or work options. However, the econometrics literature of the time provided no means to coherently analyze such choice problems.

It was serendipity when I saw a notice that Dan McFadden, then a young Berkeley faculty member visiting MIT, would give a short course on econometric analysis of discrete choice, with focus on maximum likelihood estimation of the conditional logit model. I sat in on the course and immediately saw that this was an appropriate econometric method for analysis of college enrollment decisions. Working with two co-authors, Meir Kohn and David Mundel, I subsequently performed one of the two earliest discrete choice analyses of college enrollment, the other being by Roy Radner and Leonard Miller. This work became the "job-market" chapter of my dissertation and a later version was published as Kohn, Manski, and Mundel (1976).

I learned much more from McFadden than a technical tool for discrete choice analysis. I learned that to make scientific progress, it is sometimes necessary to unlearn traditional research practices and go back to basics. Earlier efforts to study discrete choice had attempted to build on the existing econometric analysis of linear demand systems, which had been a centerpiece of econometrics for forty years. McFadden recognized that this would not work. Instead, he went back to the elementary principle of revealed preference analysis, which had been introduced decades earlier into economic theory by Paul Samuelson but which had, for whatever reason, not been directly exploited in econometric analysis of choice behavior. He drew from and re-interpreted past psychological research on random utility models, an idea which had remained largely outside the worldview of economics. And he characterized choice alternatives as attribute bundles, a simple idea that nevertheless was rarely used in economic practice, the main exception being in the construction of hedonic price indices. Thus, McFadden succeeded by shedding the current econometric approaches of his day and starting afresh, building on basic ideas in economics and psychology to develop a suitable new foundation for discrete choice analysis.

I went on to write my dissertation on discrete choice analysis. One chapter was the analysis of college enrollment mentioned earlier. A second was an investigation into the underpinnings of random utility models, later published as Manski (1977). The third was the earliest version of my work on maximum score estimation, which developed into my first published article (Manski, 1975). In retrospect, I think it fair to view it as a successful dissertation, with the maximum-score idea eventually considered to be a significant contribution to econometrics. However, my experience on the job market was anything but easy. I will explain, as the experience may be instructive.

First, I went on the market too early, in my third year. This would seem insane today, when very few Ph.D. students in economics finish in less than five years. However, the norm for completion of the degree was four years at the time, so the third year was only one year early. Nevertheless, I should have waited a year, given that the only work I had completed at the time of the job market was the paper on college enrollment. I wrote the second and third chapters after the job market was over.

Second, I suffered from the profession's inability to evaluate research that was so far out of the current mainstream. I was fully comfortable with McFadden's approach to discrete choice analysis and thought that it would be straightforward to exposit it to others. However, when I gave interviews and seminars, I found that most economists of the day could not think outside the traditional structure of linear demand systems and could not comprehend the idea of maximum likelihood estimation of a conditional logit model. McFadden's pathbreaking paper on conditional logit analysis had not yet been published, and it was as yet known to only a few economists. So I had to explain not only my own specific contribution but also the general framework within which I worked. I must not have succeeded, as I received only a moderate number of interviews at the AEA meetings, only a few flyouts, and no job offers at all.

One might draw the lesson that a student writing a dissertation should be careful to work within the current mainstream of economics. Indeed, I have heard many colleagues give this advice over the years. As I see it, this advice is professionally pragmatic and intellectually harmless if one thinks that the current mainstream is sensible and that one can contribute something original to it. However, it is professionally shortsighted and intellectually objectionable if, after careful scrutiny, one concludes that the mainstream is misguided.

In my case, I did not suffer too long for my heterodoxy. In April of my job-market year, after I had resigned myself to going on the market again the next year, two institutions that had not interviewed me at the AEA meetings contacted me. Both invited me for flyouts and made assistant professor offers. I accepted the offer from Carnegie Mellon University, made the transition from student to faculty status, and began my official academic career.

3. Studying Partial Identification

I will now jump forward to the late 1980s and early 1990s, when I wrote a set of papers that began two highly fruitful research programs. Of course much happened, professionally and personally, between completion of the Ph.D. in 1973 and the late '80s. However, it suffices for this essay to summarize the period by noting that I mainly performed empirical and methodological research on discrete choice analysis that developed naturally from my dissertation. I made a number of noteworthy contributions and, as econometrics broke out of its linear-model straightjacket, my work was increasingly viewed as within the mainstream. By the mid 1980s, I was a professionally successful econometrician with a comfortable position as professor at the University of Wisconsin-Madison.

At some point, I think in early 1987, I became concerned that my research was having diminishing returns. I had recently written a sequence of articles on semiparametric analysis of binary response, building on my early idea of maximum score estimation. However, I could see that this well was running dry. I could have continued for some time to write publishable articles on this and related topics, but I knew that they would make increasingly marginal contributions. This did not seem a good use of my time.

Instead, I decided to step back and explore some loose ideas that I had been carrying around for a while. Two such explorations initiated intense processes of unlearning and discovery. I discuss my work on partial identification here and on measuring expectations in the next section.

In the spring of 1987 Irving Piliavin, a social-work professor at Wisconsin with whom I was acquainted, asked if I had ideas on how to credibly handle a missing-data problem in his research on homelessness. He had surveyed a sample of homeless persons in Minneapolis. Attempting to re-interview them six months later, he had been unable to locate a sizeable fraction of the sample. Piliavin knew that researchers typically dealt with such attrition problems in one of two ways. Most would assume that attrition from the sample was random, while some would estimate a *selection model* that attempts to jointly explain respondent outcomes and attrition. Piliavin told me that he felt the assumptions underlying both approaches were implausible in his setting. He asked what analysis he might be able to perform with weaker assumptions.

This question intrigued me, so I agreed to give it some thought. I had always been troubled by the "assumptions of convenience" routinely maintained in empirical research. The main objective of my methodological work in econometrics had been to understand what inferences are possible under assumptions that are weaker but more credible. However, I had not previously worked on missing-data problems.

My first instinct, following a common practice among econometricians, was to start from the existing literature and explore how inference would change if one were to weaken conventional assumptions in particular ways. After some exploration, I found that I was not getting to the bottom of the missing-data problem with this investigative approach. I then decided to pose a basic question that appeared not to have been addressed previously: What inference would be possible if one were to make no assumption at all about the process generating missing data?

It turned out to be extraordinarily easy and revealing to answer this question, using no more than the elementary tool of the Law of Total Probability. The broad answer was that, although assumptions are necessary to fully reveal the probability distribution of outcomes, informative conclusions can be drawn in the absence of assumptions. That is, the data alone partially identify the outcome distribution. Going further, I found that the data alone yield simple bounds on key parameters such the mean or median outcome in the population, the width of the bound depending transparently on the prevalence of missing data.

My analysis made plain the extent to which conventional approaches to research with missing data rely on the assumptions imposed. A selection model or a missingness-at-random assumption picks out one specific outcome distribution from all those that are consistent with the data. Hence, one should never downplay such an assumption as "convenient" or an "approximation." The assumption is a critical determinant of the inferences drawn. Particularly pernicious is the fact that many commonly used assumptions are nonrefutable. For example, it is easy to show that missingness at random is a nonrefutable hypothesis. Thus, an assumption about missing data could be as ludicrous and as incontestable as the theological assumption of the Flying Spaghetti Monster.

Missing outcome data in survey research is a ubiquitous problem, but somewhat mundane. The scope of my new work expanded enormously when I soon recognized that my findings applied to the fundamental scientific problem of inference on treatment response from observation of the outcomes of realized treatments. Here the inferential problem is the unobervability of outcomes under treatments that a person did not receive. One may be able to observe the outcomes of realized treatments, but the outcomes of counterfactual treatments are logically unobservable.

I introduced my perspective on missing outcome data in Manski (1989), extended it to analysis of treatment response in Manski (1990), and presented many further findings in an invited presentation at the 1990 World Congress of the Econometric Society, subsequently published as Manski (1994). By the early 1990s, I was confident that I not only had achieved a set of important specific findings but, more importantly, that I had opened a broadly useful new approach to empirical inference in the social sciences. The basic idea was to proceed conservatively, first asking what can be learned from data given only knowledge of the sampling process, and then determining the identifying power of weak assumptions having high credibility. I developed this worldview for a broad audience in *Identification Problems in the Social Sciences* (Manski, 1995). When this short monograph was published, I felt that I finally had made a contribution of truly lasting value. I still consider it my finest work.

I will not take the space here to describe the considerable expansion of research on partial identification, performed by myself and many others, that has occurred since the early 1990s and that continues without letup today. Interested readers who are not familiar with the field may want to begin with *Identification for Prediction and Decision* (Manski, 2007), my recent textbook that updates and greatly expands the coverage of the 1995 monograph. What I do want to discuss is the extraordinary difficulty that I faced for many years in getting the profession to pay attention to my work and take it seriously.

The ideas introduced in my early papers on partial identification and in the 1995 book are simple and have wide applicability. The work appeared in accessible publications and I presented it at many seminars and conferences. Nevertheless, econometricians and empirical researchers long considered partial identification to be a curious "niche" topic distant from the mainstream. Some viewed the topic with considerable hostility. Why so?

One problem was that many researchers were (and continue to be) much more comfortable with building on existing literature than with unlearning and discovery. As a general matter, researchers often find it psychologically difficult or even impossible to question received wisdom and go back to basics. In the present case, the received wisdom was that identification is an all-or-nothing proposition, with available data and maintained assumptions either fully revealing a population parameter or revealing nothing at all. There is no logical reason to think this, and scattered examples of partial identification were shown as long ago as the 1930s. Nevertheless, econometricians and empirical researchers were long fixated on what we now call point identification.

A second problem was that researchers often took my findings on the dependence of conclusions on assumptions as bad news that they would rather not hear. In Manski (1995, p. 3), I wrote:

> "Researchers sometimes do not recognize that the interpretation of data requires assumptions. Researchers sometimes understand the logic of scientific inference but ignore it when reporting their own work. The scientific community rewards those who produce strong novel findings. The public, impatient for solutions to its pressing concerns, rewards those who offer simple analyses leading to unequivocal policy recommendations. These incentives make it tempting for researchers to maintain assumptions far stronger than they can persuasively defend, in order to draw strong conclusions."

Considering what one can learn from data alone, or with only credible assumptions, makes plain that one should be wary of inferential approaches that use strong assumptions to draw strong conclusions. The methodologists who developed these approaches and the empirical researchers who applied them often were not pleased to hear this, and some reacted with hostility. I will always be grateful to my colleague Arthur Goldberger for his support and encouragement during a difficult period.

The above problems are slowly fading as a generation of younger researchers emerge who evaluate the science without preconceptions or vested interests. Much credit goes to a sequence of Ph.D. students whose dissertation research on partial identification I have been privileged to advise since the early 1990s, first at the University of Wisconsin-Madison and then at Northwestern. These persons have made their own important contributions to the field and have become valued colleagues and friends.

4. Measuring Expectations

I wrote above that I have always been troubled by assumptions of convenience in empirical research. When performing my empirical studies of choice behavior in the 1970s and 1980s, I was concerned with the strength of the assumptions I needed to maintain about how decision makers perceive their choice sets. I was similarly troubled by the assumptions on choice-set perceptions made by other researchers. We made strong assumptions to compensate for the weakness of available data.

Revealed preference analysis aims to infer decision rules from inequalities asserting that the action a person chooses is preferred to the other options in his choice set. To exploit these inequalities, the researcher must know the identity of the chosen action and the composition of the choice set. In discrete choice analysis, where alternatives are characterized as attribute bundles, the researcher also must know how decision makers perceive the attributes of their alternatives.

The problem has been that population surveys typically provide only a small part of the data needed to support revealed preference analysis. Surveys often ask respondents to identify their chosen actions, but they rarely ask them to state the composition of their choice sets. Surveys do not ask much about the attributes of alternatives, and they often ask nothing at all. Given this, researchers have long imputed choice sets and attributes. They then have assumed that their imputations match how decision makers actually perceive their choice sets.

Consider, for example, analysis of schooling decisions. Essentially all economic models suppose that a youth compares the outcomes that he expects would occur if he were to choose each of the available schooling, work, and other nonschooling options. Labor economists use the broad term *returns to schooling* to describe the expected outcomes from schooling relative to other activities. The surveys available to study schooling decisions have documented schooling choices but have not asked respondents how they perceive the returns to schooling. In the absence of data, researchers have made assumptions about the expectations that youth hold. It has been particularly common to pose so-called *rational expectations* models, in which the researcher develops a model of outcome determination and assumes that youth believe the same model.

Contemplating the state of practice in the late 1980s and early 1990s, I first explored how revealed preference analysis might proceed if one were to weaken conventional assumptions about choice-set perceptions in particular ways. I wrote a few methodological articles of this type, but I found that the assumptions required for informative inference were too strong to be credible. As in my dissertation research and in my study of inference with missing outcome data, I eventually decided that building on existing literature would not work. I would again need to go back to basics.

Considering the inferential problem anew, I concluded that making progress requires new data rather than new methodology. Instead of making assumptions about how decision makers perceive their choice sets, we should ask them. I thought it particularly important to obtain data on the expectations that decision makers hold for the outcomes of alternative actions. Considering research on schooling decisions, I proposed in Manski (1993) that the profession should develop new survey instruments asking respondents how they perceive the returns to schooling.

This idea seemed obvious enough, but it went against a deeply entrenched conventional wisdom among economists, one that I myself had previously accepted without question. Early in my career, I had been taught that a good economist believes only what people do, not what they say. I was indoctrinated to be especially skeptical of subjective statements. Such statements are incapable of external validation and, hence, constitute a prime example of what game theorists call "cheap talk." Accepting this perspective, I had viewed collection of expectations data as a useless activity.

I came to rethink my preconceptions about measurement of expectations for two reasons. First, I became aware that there existed no persuasive body of empirical evidence showing that expectations data are uninformative. Second, I became increasingly concerned that, in the absence of data, economists were making unsubstantiated and often fanciful assumptions about expectations. I reasoned that collecting data from survey respondents should be better than creating assumptions from thin air.

I decided that the most effective way to make progress would be to initiate collection of expectations data in a suitable survey. The University of Wisconsin Survey Center operated an ongoing national telephone survey of Americans. As a faculty member, I was able to place questions on the survey at low cost. I worked with Jeff Dominitz, whose Ph.D. dissertation I was then advising, to develop a module of questions asking respondents to state the likelihood that they would experience various events in the year ahead. For example, we asked working persons "What do you think is the percent chance that you will lose your job during the next 12 months?"

As economists, Jeff and I thought it perfectly natural to ask "percent chance" questions of this type. After all, economic models of decision making under uncertainty routinely assumed that persons place subjective probability distributions on uncertain outcomes. However, many of the attitudinal researchers associated with the UW Survey Center were highly hostile to this question format, preferring instead to ask questions seeking verbal responses. For example, they would ask "How likely is it that you will lose your job during the next 12 months?" with the permitted responses being (very likely, fairly likely, not too likely, not at all likely).

These attitudinal researchers warned me that many respondents would not understand the concept of percent-chance, would not answer the questions and, even worse, might find the questions so disagreeable that they would hang up without completing the remainder of the questionnaire. I asked for references to empirical research demonstrating these difficulties and learned that no research had been undertaken. The conventional wisdom among attitudinal researchers apparently was so strong that they thought empirical investigation was unnecessary.

Undaunted, Jeff and I initiated our Survey of Economic Expectations in 1993. It became clear early on that the qualms of the attitudinal researchers were baseless, and we continued the survey through 2002. We wrote several articles analyzing the data, beginning with Dominitz and Manski (1997a, 1997b). We also performed a separate exploratory study asking students to state in percent-chance form their expectations for the returns to schooling (Dominitz and Manski, 1996).

Independently, Thomas Juster at the University of Michigan, who had proposed probabilistic measurement of expectations in surveys in a long overlooked article in 1966, began the highly important Health and Retirement Study (HRS) in 1992. The first administration of the HRS asked a broad set of expectations questions and later administrations have added more, making the HRS the pre-eminent source of data of this type. Data collection in other major surveys has followed.

A decade after the new literature had come to life in the early 1990s, enough progress had been made by a growing set of researchers that I was able to write a fifty-page review article describing the state of the art and calling attention to important open questions (Manski, 2004). Important further advances have been made in the past five years. Yet I still consider measurement of expectations in surveys to be a work in progress and an exciting area for future research. Most of the work to date has posed fairly simple unconditional expectations questions. The objective of using data on choice-specific expectations in revealed preference analysis has begun to be realized only in the past few years.

5. Questioning Consistency

To be clear, my research has not all been a matter of unlearning and discovery. Considering my time allocation over the past forty years, I can see that my hours worked have been devoted primarily to the ordinary-science activity of building on what I considered to be secure foundations for research. This is as it should be. As indispensable as the episodes of unlearning and discovery have been, I have known that they could not yield anything lasting unless I were to make the effort to work out their implications thoroughly. Thus, I worked for about fifteen years on discrete choice analysis in the early part of my career and continue today to make occasional fresh contributions. I have now worked for about twenty years on partial identification and measuring expectations, and as yet see no diminishing returns to research in these fields. Some of my most gratifying ordinary-science experiences have occurred when I have been able to perform fruitful crossover research, synthesizing ideas from my various research streams.

Yet there is no denying that the episodes of unlearning and discovery have been the exciting periods, during which nothing else in the world seems to matter. My most recent such episode began in summer 2008, as an outgrowth of my ongoing program of research on planning under ambiguity, which in turned developed from my work on partial identification. Seeking to characterize reasonable/rational behavior under ambiguity, I found myself deeply at odds with the long influential worldview of axiomatic decision theory, which maintains that a decision maker is rational if he would behave consistently in particular senses when facing various hypothetical choice sets. I will not take the space here to describe this ongoing process of unlearning and discovery. The interested reader should see Manski (2010), my first attempt to present my evolving perspective in writing.

Notes

1. Years later I mentioned this episode to a physicist. He replied that one is not supposed to have intuition for relativity and quantum mechanics. I have since wondered what career direction I would have taken had I attributed my lack of intuition to the nature of these subjects rather than to a personal failing.
2. The Air Force had particular appeal because I long was fascinated with flight. If my eyesight had been up to it, I would have contemplated a career as a pilot. Although I could not fly professionally, I did obtain a private pilot's license while in graduate school and later an instrument rating while an assistant professor.
3. This episode may appear cut-and-dry when summarized in brevity, but two aspects of it still weigh on me. First, when dropping out of ROTC, I deeply regretted having to break the nearly paternal tie that I felt to the MIT unit commander, Lt. Col. George P. Gamache, whom I considered a person of great integrity. I regretted this even more several years later, when I learned that Colonel Gamache, after completing his MIT assignment and returning to his normal duty as an aviator, had been killed in a B-52 crash in Florida.

   Second, I am well aware that others were not so easily able to escape the draft and get on with their lives. The morning of the draft physical, I waited for the Army bus with a group of ordinary Boston guys from the neighborhood where I grew up, including many whom I had not seen since elementary school. Most of them were not there with minor ailments expecting to avoid service. They were expecting to pass the physical and be drafted.
4. I was not the only student troubled by this. Our class formally protested to the department and collectively refused to take the semester final examinations in micro and macro. This action was not as foolhardy as it may appear. We were careful to read the departmental regulations, from which we learned that the course examinations were not mandatory.

References

Dominitz, J. and C. Manski (1996), "Eliciting Student Expectations of the Returns to Schooling," *Journal of Human Resources,* 31, 1–26.

Dominitz, J. and C. Manski (1997a), "Using Expectations Data to Study Subjective Income Expectations," *Journal of the American Statistical Association,* 92, 855–867.

Dominitz, J. and C. Manski (1997b), "Perceptions of Economic Insecurity: Evidence from the Survey of Economic Expectations," *Public Opinion Quarterly,* 61, 261–287.

Kohn, M., C. Manski, and D. Mundel (1976), "An Empirical Investigation of Factors Which Influence College Going Behavior," *Annals of Economic and Social Measurement,* 5, 391–419.

Manski, C. (1975), "Maximum Score Estimation of the Stochastic Utility Model of Choice," *Journal of Econometrics,* 3, 205–228.

Manski, C. (1977), "The Structure of Random Utility Models," *Theory and Decision,* 8, 229–254.

Manski, C. (1989), "Anatomy of the Selection Problem," *Journal of Human Resources,* 24, 343–360.

Manski, C. (1990), "Nonparametric Bounds on Treatment Effects," *American Economic Review Papers and Proceedings,* 80, 319–323.

Manski, C. (1993), "Adolescent Econometricians: How Do Youth Infer the Returns to Schooling?" in C. Clotfelter and M. Rothschild, eds., *Studies of Supply and Demand in Higher Education,* Chicago: University of Chicago Press.

Manski, C. (1994), "The Selection Problem," in C. Sims ed. *Advances in Econometrics, Sixth World Congress*. Cambridge: Cambridge University Press.

Manski, C. (1995), *Identification Problems in the Social Sciences,* Cambridge, Mass.: Harvard University Press.

Manski, C. (2004), "Measuring Expectations," *Econometrica,* 72, 1329–1376.

Manski, C. (2007), *Identification for Prediction and Decision,* Cambridge, Mass.: Harvard University Press.

Manski, C. (2010), "Actualist Rationality," Department of Economics, Theory and Decision, forthcoming.