Journal of Agricultural and Applied Economics, 35,2(August 2003):361–370

"It's for Teaching, not Believing": Comments on Teaching, Learning, and Problem Solving Through Economic Experiments

Robert G. Nelson and Norbert L.W. Wilson

This series of papers is an excellent opportunity to reacquaint agricultural economists in the Southern region with the exciting field of experimental economics and is indeed opportune in light of the recent awarding of the Nobel Prize in Economics to Vernon Smith, considered by many to be the father of experimental economics. Rather than try to share the limelight with the authors on their far-reaching and comprehensive topics, we plan to take this opportunity to share some of our views of the role of experimental economics in the research laboratory and classroom.

Problem Solving and Hypothesis Testing

In this session, Hudson argues that experimental economics offers a means to counter criticisms of economics as a science by addressing both the rigor of empirical tests and the quality and timeliness of data needed for policy prescriptions. He describes what experimental economics is and why we do experiments and suggests that experiments can be classified dichotomously according to whether subjects are paid or not. This, we believe, is not a particularly helpful distinction, because payment, and motivation in general, are better integrated under Smith's (1982b) criteria for a "valid, controlled microeconomic experiment," which we describe in a subsequent section. As an alternative, we suggest Davis and Holt's framework for classifying experiments in the dimensions of institutional and environmental complexity.

Robert G. Nelson and Norbert L. W. Wilson are associate professor and assistant professor, respectively, in the Department of Agricultural Economics and Rural Sociology, Auburn University, Auburn, AL.

Institutional Complexity

Confusion often arises with the term *institution,* because of its association with the field of institutional economics, which is concerned with the ways that institutions evolve in response to individual incentives, strategies, and choices and how institutions affect the performance of political and economic systems. For experimental economists, the term institution simply refers to the rules governing economic interactions in an experiment. In a recent interview (Lynch and Gillespie), Vernon Smith emphasized: "The thing that's not very explicit in much of economics is what the rules of trading are and how they affect outcomes. Experimental economics asks how the performance of a market is influenced by its rules."

The rules that characterize institutions include:

(1) The nature and timing of messages allowed between agents (e.g., Is face-to-face communication allowed? Who can make offers? Are offers sequential or simultaneous?);

(2) the kinds of decisions or actions that are observable (e.g., bids, offers, contracts, forecasts, draws from an urn, side payments);
(3) the mapping of these decisions into the payoff or incentive structure (i.e., the reward a subject gets for making a certain decision); and
(4) stopping rules for ending the session or trading period (e.g., random, fixed, known, unknown).

Game theorists also have an interest in precisely specifying institutional complexity, but the domain available for exploration by experimentalists is much more extensive and often more expedient.

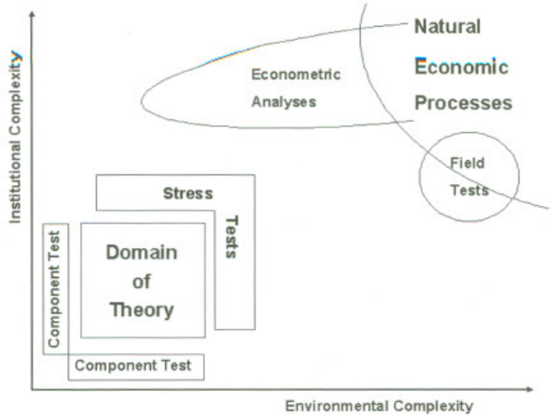

Figure 1. Economic Research Methods Positioned in the Dimensions of Institutional and Environmental Complexity (adapted from Davis and Holt)

Environmental Complexity

The term *environment* may also cause some confusion, because of its association with natural resources, ecology, and environmental concerns. For experimental economists, the concept is again far more simple and refers to the structural characteristics of the economic setting, including:

(1) the number of agents (e.g., monopoly, duopoly, oligopoly, "many" buyers and sellers);
(2) their initial endowments (e.g., money, information, experience, market power, property rights);
(3) production technologies (e.g., supply schedules derived from production costs; single vs. multiple units; constant vs. varying costs; storable vs. nonstorable units);
(4) demand conditions and structures (e.g., demand schedules derived from redemption values vs. home grown preferences; stationary demand vs. cyclical or aperiodic shifts); and
(5) the number of periods for continued interaction (e.g., "one-shot" games vs. repeated contact; 10 vs. 200 rounds).

An example of an institution with a high degree of complexity is a voluntary assessment mechanism for funding a public good like a church. An institution of modest complexity is the Dutch auction, in which the price for an item falls sequentially until it is sold to the first buyer who makes an offer. An institution of very low complexity is the mandatory assessment mechanism of a commodity check-off program, which has the authority of the police powers of the state and is functionally equivalent to a simple excise tax. Recently, an increasing number of commodity boards have had their check-off rescinded and are showing a renewed interest in and appreciation for the complexities of voluntary assessment mechanisms.

Market environments of graduated complexity could range from a simple monopoly setting with a single seller facing a computer-simulated buyer population with a fixed demand curve (Nelson and Beil 1994), to an oligopoly setting with two to four sellers (Nelson and Beil 1995), to a market with four buyers and four sellers, each with multiple units, trading in either forward or spot markets with random supply and demand shocks (Menkhaus et al.).

Figure 1 is adapted from Davis and Holt to illustrate these relationships by positioning various research approaches in the two dimensions of institutional and environmental complexity. Natural, "real world," economic pro-

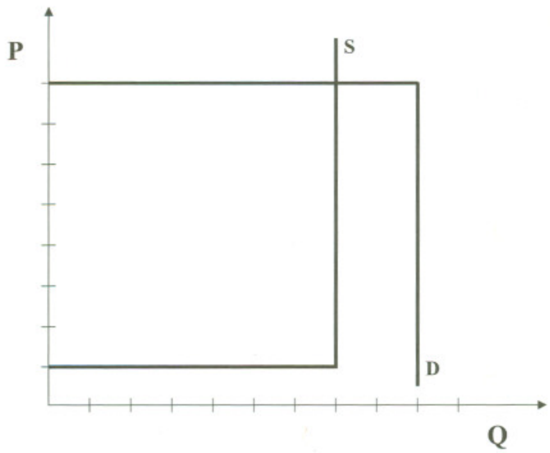

Figure 2. Example of a Rectangular Market

cesses appear in the upper right corner because of their extreme complexity in both dimensions. The flattened parabola encompasses econometric modeling of natural processes. These models are frequently subject to questions of whether the dimension of environmental complexity has been adequately specified or whether important relevant variables have been omitted.

In the lower left corner of the graph, the "domain of theory" does not extend all the way to the origin if we allow for "component tests" that reduce standard elements of theory to even simpler conditions. For example, standard market theory predicts market clearing at a single price in competitive markets, typically with upward-sloping supply curves and downward-sloping demand curves. A component test in the environmental dimension might use a "rectangular market" that consists of a horizontal supply schedule with seven units, each of which costs $1 to produce, and a horizontal demand schedule with nine units, each with a redemption value of $8 (see Figure 2). Standard theory would say that, in perfect competition, seven units would trade for $8 each, and all the surplus would go to the sellers: but who believes this? Component tests have revealed that the result depends critically on the institution used: the double-auction market strongly supports the theory (Holt, Langan, and Villamil), whereas the posted-offer market does not (Cason and Williams).

In contrast, "stress tests" explore the robustness of theory applied beyond the maximum limits specified as its domain. For example, Hayek (1945) speculated that perfect information is not necessary for competitive markets. In fact, strict *privacy* of production costs and redemption values might be necessary to decrease opportunities for collusion. Such an increase in the uncertainty of market structure constitutes a more environmentally complex condition. Smith (1982a) showed experimentally that Hayek was in fact correct: competitive equilibrium is regularly achieved in double-auction markets with as few as two buyers and two sellers having only private information. On this general finding, subsequently repeated in hundreds of experiments, Smith (1989, p. 101) wryly commented: "This result was not as expected. The conventional view among economists was that a competitive equilibrium was like a frictionless ideal state which could not be conceived as actually occurring, even approximately. It could be conceived of occurring only in the presence of an abstract 'institution' such as a Walrasian *tâtonnement* or an Edgeworth recontracting procedure. It was for teaching, not believing."

A comprehensive series of experiments would begin with a test of some theory within its accepted domain. If a number of tests consistently supported the predictions of the theory, then the researcher might try some stress tests to explore the limits of the theory's predictive capacity in either the institutional dimension or the environmental dimension, or both. Conversely, if the theory failed a crucial test, then a series of component tests might be conducted to determine what conditions the theory is most sensitive to or how much simpler or more specific the conditions have to be for the theory to predict well. To most researchers in the natural sciences, this sort of research agenda would be a legitimate, systematic, and familiar approach to the enterprise of knowledge acquisition. Unfortunately, this sentiment is not as common in the economics profession, wherein the experimental approach is often far outside of many researchers' comfort zone. More than a few experimental economists who have found evidence that appears to disconfirm a theory have received reviews with comments such as "the test failed to capture the complexity of the real situation," or "you can't test an economic theory in a laboratory." On the other hand, those who have experimentally corroborated a theory are subject to comments such as "this was a waste of time and money, because we already knew the theory was true."

Field tests can be characterized as an extreme form of experiment. Compared with laboratory tests, field tests are usually larger, more costly, and more difficult to control and analyze. With the exception of commercial, proprietary marketing research in full-scale test markets, economic field tests are not commonly attempted in the United States. However, as Lusk points out in this session, the agricultural economics profession appears to be the bearer of the torch in using such techniques as experimental auctions for marketing applications, which approach the scale of field tests.

Examples of early field studies include investigation of peak load pricing of electricity (Battalio et al., 1979); income maintenance and the negative income tax (Kershaw and Fair; Pechman and Timpane); token economies in prisons, psychiatric wards, aircraft carriers, and similarly isolated communities (Kagel); and allocation of housing statistics in Sweden (Bohm). Characteristics that distinguish field studies from laboratory studies are largely related to expense, logistics, and manpower requirements. Some field studies cost millions of dollars, last several years, and may employ large staffs of researchers, administrative personnel, and even medical support. In the natural progression of scientific inquiry, the most cost-effective use of field experimentation is in validation of results from laboratory studies, i.e., in examining the generalizability of bench-scale results to more complex economic environments. However, pressures to produce policy prescriptions, as well as concerns with sample representativeness, often promote field experiments ahead of lab experiments. In our experience, there appears to be a bias among agricultural economists that field studies are preferred over laboratory experiments. Unfortunately, the consequence of this bias is that fewer experiments of either kind are conducted, in the former case because of expense and in the latter because of doctrinaire attitudes.

A Primer on the Design of Economic Experiments

Advantages of Experiments

What are the advantages of doing laboratory experiments in economics? Plott lists five potential advantages. First, laboratory experiments generate data in a controlled environment. This allows experiments to be replicated independently by anyone skeptical of the original results. Second, most experiments can be executed in a few hours and can be used to simulate transactions that might take days or even years to observe in the real world. Third, experiments in laboratory settings are generally much less expensive than those in field settings. Fourth, the environment is flexible and can be changed readily to simulate a variety of conditions, some of which cannot be controlled for or even observed in the field. This allows the investigation of a wider range parameters. Finally, to subject a theory to experimental validation, proponents or opponents of the theory are required to "operationalize" it: does the theory specify certain conditions required for it to work, such as certain characteristics of agents, maintained functional forms, econometric regularities, stopping rules, equilibrium conditions, etc.? Quite often, these conditions are defined abstractly or not at all in the theory, but they must be dealt with concretely in even the simplest of experimental settings.

Features of Experiments

Experimental economics can be defined as the study of economic behavior under controlled and replicable conditions. *Control* is necessary to assure that the variable under examination (the one that constitutes the "treatment") is the only parameter being varied. In this context it is similar to the ceteris paribus conditions so prevalent in economic theory. *Replicability* is necessary to ensure that later researchers who wish to verify or extend a study will have reasonable success in imitating the conditions under which the original study was conducted. As pointed out by Hudson in this session, this interpretation differs from that of econometrics, whereby replicability simply refers to the ability to reproduce the same results with the same data set, an exercise that should be—but surprisingly often contrives not to be—as deterministic as the sum of two plus two.

Every experiment uses what is called a "protocol." The instructions made available to the participants are an important part of the protocol, because these describe the rules and rewards. The protocol also involves a detailed description of the experimental treatment and control conditions. Designing the protocol is one of the most creative aspects in developing a new experiment in a novel research area. Later experimenters usually use the same instructions to control for replicability of the rules while varying the treatment.

The general protocol used in laboratory experiments in economics has become fairly standard. Subjects are typically recruited from convenient undergraduate classes. This is acceptable when the hypothesis being tested does not apply only to a specific population, as is usually the case in economic theory. A set of instructions that describes how to earn money in the experiment is read to the subjects at the start of the experiment. The role of these cash payments is described by Smith (1976) in his theory of "induced value."

After reading the instructions, subjects then engage in some game-playing or role-playing situation that embodies the essential features of the hypothesis being tested. Although these situations sometimes appear to be gross oversimplifications of reality, they go much farther in approximating authentic conditions than do most economic theories. It is in this way that laboratory experiments bridge the gap between theory and observations of the real world.

During the experiment, the behavior and responses of the subjects are observed and recorded. Typically, these data are then analyzed with simple but robust statistics such as t-tests, F-tests, ANOVA, or their nonparametric counterparts. Because much behavior involves dynamic situations and learning, simple graphs and time charts often yield the most revealing patterns of behavior.

Laboratory sessions usually last only 2–4 hours, to avoid complications from fatigue or boredom on the part of subjects. Sometimes it is necessary to train subjects beforehand in the mechanics of the game or to preselect those who exhibit characteristics essential to the research question, such as subjects who are risk-neutral in the domain of task rewards, or those who have engaged in cooperative behavior in a previous game.

Sufficient Conditions for Experiments

What makes laboratory experiments a valid source of data about the real world? Smith (1982b) proposed five sufficient conditions that constitute a *valid, controlled microeconomic experiment:*

(1) *nonsatiation:* subjects prefer more money to less;
(2) *salience:* the experiment has motivational relevance that clearly links the reward to the task;
(3) *dominance:* the anticipated rewards in the experimental setting dominate any other costs or benefits that might affect performance of the task;
(4) *privacy:* subjects are informed only about their own payoffs, to control for interpersonal utility; and
(5) *parallelism:* propositions derived from laboratory experiments will apply wherever similar ceteris paribus conditions hold.

Nonsatiation and salience are required to create a *microeconomic environment.* Nonsatiation is a powerful axiom of preference theory that enables us to make predictions about a person's preferences among bundles solely from the observable and measurable quantities that make up the bundles. Salience requires that subjects understand the task to be performed, the rewards to be earned under various conditions, and that these should be obviously related to the performance of that task.

Dominance and privacy are needed for *experimental control.* If subjects incur opportunity costs of time or effort that are much greater than their expected reward from the experiment they may become distracted or bored or may hurry through the task to leave sooner, and their responses will be much more variable. When specific experience or expert knowledge is not a condition of the hypothesis being tested, the concept of dominance is the principal motivation for using students in experiments: the opportunity costs of students are less than those of the general public. In addition, it is essential to provide privacy, because an experiment in which subjects are expected to maximize their *absolute* payment can become uninterpretable if they are really just trying to outperform others in their *relative* payment, performance or ranking (sometimes called a *tournament setting*).

The first four conditions are required to achieve *internal validity,* "the basic minimum without which any experiment is uninterpretable: did in fact the experimental treatments make a difference in this specific experimental instance?" (Campbell and Stanley, p. 5). Internal validity is a necessary condition for any test of hypotheses to be a valid source of data about the real world. Again, this is achieved in the experimental sciences by control—the ability to link a specific response to a specific stimulus because all other variables are being held constant or are randomly assigned to the replicates. The achievement of control is what allows the experimenter to select the best explanation of the results from among several competing explanations, and this is what enables science to progress.

Parallelism, the last condition, allows results from the laboratory to be transferred to the many other situations that constitute the "real world." It is a condition of *external validity* or generalizability: "to what populations, settings, treatment variables, and measurement variables can this effect be generalized?" (Campbell and Stanley, p. 5). Often, this condition is difficult to satisfy in practice. For example, a straightforward test of the hypothesis that a "nuclear winter" would follow a nuclear war would be unthinkable. In cases where a critical field validation experiment is impractical, it may be possible to expand the number of laboratory tests to include an increasing variety of conditions so as to define the range over which the hypothesis could be expected to hold. This is how the inverse square law of attraction was extended from celestial mechanics to the motions of atoms.

Note that a direct field test of a hypothesis that seems to satisfy the condition of parallelism but has not secured control cannot use inferential statistics to eliminate alternative hypotheses. Sources of error from uncontrolled variables can often produce confidence regions that are large enough to include the results predicted by other hypotheses. Thus, internal and external validity are both required to provide the necessary *and* sufficient conditions for a laboratory experiment to be a valid source data about the real world. Even though each is necessary, neither by itself is sufficient.

Challenges of Experiments

Having extolled the virtues of laboratory experiments, let us explore some of the ways that experiments can go wrong, such as when one or more of the conditions mentioned above are not met. A common source of problems is that subjects do not understand the task because the instructions are unclear or misleading or simply because the task is too complicated. Consider an experiment in which subjects are required to estimate the subjective expected utility of 20 choices in five categories in under 30 seconds. Obviously, the complexity of the task is far out of proportion to the resources available to the subjects for the performance of the task. Violations of the conditions of salience, dominance, and parallelism could all be present in this situation. Most experiments engender, to some degree, a test of "rationality," such as maximizing earnings or behaving as theory predicts. However, when subjects are not given a reasonable chance of behaving optimally, generalization to the real world may be questioned. If, in the real world, the 20 choices above were in fact the potential oil reserves located in five regions, then a team of geologists and engineers given several months could probably reach a decision close to the optimum predicted by theory. Some of the controversy relating to subjects' use of "heuristics and biases" may arise from unreasonable demands on computational ability in experimental settings.

Another problem is that the structure of the reward mechanism may suggest to subjects a strategy that was not anticipated by the experimenters—another problem of salience. For example, in situations involving the elicitation of subjective probabilities, certain payment functions can induce some subjects to give predictions that differ from their true beliefs, because their payoffs are higher that way (Nelson and Bessler). In this session, the discussion by Lusk on incentive compatibility in elicitation of willingness to pay is intimately concerned with issues of salience, dominance, and even nonsatiation.

Further problems may arise when rewards fail the dominance criterion. This is most often the case when payments are too low and subjects, perceiving inadequate compensation for their time or effort, behave inconsistently or erratically. In this session, Hudson discusses "hypothetical bias" and "endowment effects," both of which are examples of problems that can arise from violations of the dominance criterion. Also in this session, Barnett and Kriesel present objections to rewarding students for their participation in teaching experiments, which is another example of dominance considerations.

As a final comment on sources of errors in experiments, the results of experiments are sometimes questioned on the grounds that the subjects used were not "representative," which can be viewed as a violation of parallelism. This is certainly a valid criticism if the theory being tested is conditional on a specific subject pool (e.g., stock brokers); however, most theories in economics do not specify that the economic agents in question must have certain characteristics—physical, behavioral, cultural, socioeconomic, or otherwise. Nevertheless, even where the theory does not specify subject pool characteristics, it can happen that the task is so designed that it prevents some or all subjects from being able to express the full range of behavior. Thus, cases arise in which sociologists accuse economists of training people to be free riders, because the only subject pools that displayed such behavior were populated by economics students (Marwell and Ames). Careful examination of the conditions under which the experiment was conducted usually reveals some element that, although originally overlooked, turns out to make a considerable difference when the experiment is replicated independently. For example, in the free-rider experiments, all of the subjects eventually began to free ride when the experiment was allowed to continue for multiple periods (Isaac, McCue, and Plott). In a study of probabilistic stock price forecasting (Stael von Holstein), statisticians understood the rules of the probability calculus that were needed to perform "correctly," whereas stock brokers did not. In any case, good experimental design and protocol should obviate problems that might arise concerning the representativeness of subjects.

Incentives

Let us return now to the question of rewarding subjects, particularly students in a classroom setting. What purpose do rewards serve? A number of studies have examined payment versus nonpayment as a treatment variable. Siegel found that, when he held the complexity of the task constant and increased the reward, the number of reward-maximizing (salient) choices made by subjects increased. Then, when he held the reward constant and increased the complexity of the task, the number of reward-maximizing choices decreased. Smith (1976) found that, when there were no rewards or when rewards were chosen randomly, responses were much less consistent than when there was a known reward. Phillips and Edwards investigated subjects' ability to incorporate new information into their decisions, which they defined as "learning." They found that more learning occurred under payoff conditions than under nonpayoff conditions and that there was less variation between subjects' responses in the payoff group than in the nonpayoff group. In a study that used hypothetical (no pay) and real (pay) gambles, Jiranyakul examined several theories that were devised to explain results that were inconsistent with the classical theory of expected utility (EU). In almost every case, subjects who gave responses inconsistent with classical EU theory when the reward was hypothetical gave EU-consistent responses when the gamble was associated with real payoffs.

Because significant differences between treatments are more likely to be found when variance is smaller, these results indicate that experimenters have a better chance obtaining unambiguous results when they pay subjects. Moreover, to the extent that experiments in the classroom are directed toward improving students' ability to incorporate new information into their decisions ("learning"), then rewarding them is evidently preferred to not rewarding them. Even when the purpose of a classroom experiment is simply to demonstrate an economic concept, the likelihood of being able to do so increases when students are properly motivated to give consistent responses.

It seems obvious from the examples above that, unless one has strong prior evidence that the experiment will meet the conditions of salience and dominance without payments, one should plan to reward subjects. What is surprising to us is not that so many studies have been satisfactorily conducted without payments but that so many critics would prefer to compromise the statistical significance and interpretation of valuable economic data rather than deal systematically with the questions of salience and dominance.

Teaching

Barnett and Kriesel's paper in this session advocates the use of experimental economics in the classroom. In part, the authors base their support of classroom experiments on the work of psychologist Carl Jung and educational specialist David Kolb, who identified different types of learners and styles of learning. Barnett and Kriesel suggest that, although agricultural economics students are largely *concrete* learners who would prefer more "hands-on" approaches to learning, most undergraduate teaching is conducted in a lecture format that provides little opportunity for concrete examples or experiences. On the basis of work by Dobbins et al., (1995), Barnett and Kreisel argue that experiential learning techniques are best for providing students with concrete examples of abstract concepts, which in turn helps them learn those concepts better.

In their concluding remarks, Barnett and Kriesel propose that grant-giving agencies should provide more funds for experiments in *research,* arguing from their regression results that the greater use of experiments in research leads to the greater use of experiments in the classroom. But why should the greater use of experiments in the classroom be advocated? Why should experiments be used relative to other teaching methods?

We do not know. This response is not a statement against the use or usefulness of experiments in the classroom but rather is an acknowledgment of Barnett and Kriesel's assessment that the empirical evidence suggesting that experiments facilitate learning is limited and inconclusive. We simply do not know whether experiments are unequivocally better than other experiential learning tools nor, for that matter, whether experiential learning tools are better than the standard lecture format. Some critics argue that experimental economists should use experiments to determine the value of experiments in the classroom. Designing such an experiment could in itself be a useful exercise, but if we require this scrutiny for experiments, would it not be consistent to require the same scrutiny of all teaching methods, even lectures?

A 1993 article in the *Journal of Economic Education* by Rendigs Fels entitled "This is What I Do, and I Like It" has probably caused more mischief on this issue than any other. Fels used the title phrase to derisively characterize the meager empirical support provided by early proponents of classroom experiments. Furthermore, his statement (p. 365) "It is ironic that those who use controlled experiments in their research on economics do not use controlled experiments to evaluate their teaching" has undoubtedly been used by countless reviewers as a justification for rejecting manuscripts on experiments for the classroom.

This myopic view is lamentable for a number of reasons. In the first place, Fels himself outlined a research agenda for testing the efficacy of classroom games in teaching "one important idea" (which he suggested should be allocative efficiency), but because, by his own admission, he had never tried experiments, his protocol was not very explicit. In fact, it is only five lines long and ends with the statement: "I would cajole, browbeat, or bribe colleagues who are expert in such matters to evaluate the results, *holding other things constant.*" (p. 369, italics ours). As we described in the previous section, holding other things constant is the sine qua non of a controlled experiment, and the achievement of control is one of the most demanding and creative aspects of designing an experiment. The fact that such a protocol is not standardized in the pedagogy of our profession suggests that it may be exceedingly difficult to achieve satisfactorily.

In the second place, Fels (p. 369) conceded that "It may safely be predicted that a single experiment using one or two class hours would not make an important difference in how much students learn. (The same could be said of any use of one class hour, including canceling it altogether.)" Indeed, the same could also be said of case studies, simulations, PowerPoint slides, Internet access, distance learning, econometrics, graphs, field trips, guest lectures, classroom discussions, reading and writing assignments, homework, exams, and lectures. Yet none of these uses of class time is held to the same standards of evaluation as experiments. By analogy, we should also demand that anyone wishing to publish a case study should be required to evaluate its effectiveness with case study methods or otherwise be condemned by the "irony" of the situation.

None of this recrimination is productive. Teachers should recognize the following premises as axioms, which we hold to be self-evident: students have different learning styles, instructors have different teaching styles, and a bigger toolbox of teaching methods is preferred to a smaller one, even if we do not use all the tools. In the words of a Palestinian bar owner (National Public Radio), on the advantages of living in America: "It is good to have everything. Then you can choose."

References

Battalio, R.C., J.H. Kagel, R.C. Winkler, and R.A. Winett. "Residential Electricity Demand: An Experimental Study." *Review of Economics and Statistics* 61(1979):180–89.

Bohm, P. "Revealing Demand for an Actual Public Good." *Journal of Public Economics* 24(1984): 135–51.

Campbell, D.T., and J.C. Stanley. *Experimental and Quasi-Experimental Designs for Research.* Chicago: Rand McNally, 1963.

Cason, T.N., and A.W. Williams. "Competitive Equilibrium Convergence in an Posted Offer Market with Extreme Earnings Inequities." *Journal of Economic Behavior and Organization* 14(1990):331–52.

Davis, D.D., and C.A. Holt. *Experimental Economics.* Princeton, NJ: Princeton University Press, 1993.

Dobbins, C.L., M. Boehlje, S. Erickson, and R. Taylor. "Using Games to Teach Farm and Agribusiness Management." *Review of Agricultural Economics* 17(1995):247–55.

Hayek, F.A. "The Use of Knowledge in Society." *American Economic Review* 35(1945):519–30.

Holt, C.A., L. Langan, and A.P. Villamil. "Market Power in Oral Double Auctions." *Economic Inquiry* 24(1986):107–23.

Isaac, R.M., K.F. McCue, and C.R. Plott. "Public Goods Provision in an Experimental Environment." *Journal of Public Economics* 26(1985): 51–74.

Jiranyakul, K. "Utility Function in the Domains of Gains and Losses: An Experimental Study." Ph.D. thesis, Texas A&M University, 1986.

Jung, C.G. *Psychological Types.* Trans. by H. G. Baynes, rev. by R.F.C. Hull. Princeton, NJ: Princeton University Press, 1976 (originally published in 1921).

Kagel, J.H. "Token Economies and Experimental

Economics." *Journal of Political Economics* 80(1972):779–85.

Kershaw, D., and J. Fair. *The New Jersey Income Maintenance Experiment.* New York: Academic Press, 1976.

Kolb, D.A. *Experiential Learning.* Englewood Cliffs, NJ: Prentice Hall, 1984.

Lynch, M., and N. Gillespie. "The Experimental Economist." Interview of Vernon Smith dated December 2002. Internet site: http://www.reason.com/0212/fe.ml.the.shtml (accessed January 5, 2003).

Marwell, G., and R.E. Ames. "Economists Free Ride: Does Anyone Else? Experiments on the Provision of Pubic Goods. IV." *Journal of Public Economics* 15(1981):295–310.

Menkhaus, D.J., C.T. Bastian, O.R. Phillips, and P.D. O'Neill. "Supply and Demand Risks in Laboratory Forward and Spot Markets: Implications for Agriculture." *Journal of Agricultural and Applied Economics* 32(2000):159–73.

National Public Radio. Program broadcast March 1, 1999.

Nelson, R.G., and R.O. Beil, Jr. "A Classroom Experiment on Oligopolies." *Journal of Agricultural and Applied Economics* 27(1995):263–75.

Nelson, R.G., and R.O. Beil, Jr. "Pricing Strategy Under Monopoly Conditions: An Experiment for the Classroom." *Journal of Agricultural and Applied Economics* 26(1994):287–98.

Nelson, R.G., and D.A. Bessler. "Subjective Probabilities and Scoring Rules: Experimental Evidence. *American Journal of Agricultural Economics* 71(1989):363–69.

Pechman, J.A., and P.M. Timpane, eds. *Work Incentives and Income Guarantees: The New Jersey Negative Income Tax Experiment.* Washington, DC: The Brookings Institution, 1975.

Phillips, L., and W. Edwards. "Conservatism in a Simple Probability Inference Task." *Journal of Experimental Psychology* 72(1966):346–54.

Plott, C.R. "Experimental Methods in Political Economy: A Tool for Regulatory Research." *Attacking Regulatory Problems: An Agenda for Research in the 1980's.* A.R. Ferguson, ed. Cambridge, MA: Ballinger Publishing, 1981.

Siegel, S. "Decision Making and Learning Under Varying Conditions of Reinforcement." *Annals of the New York Academy of Sciences* 89(1961): 766–83.

Smith, V.L. "Experimental Economics: Induced Value Theory." *American Economic Review* 66(1976):274–79.

———. "Markets as Economizers of Information: Experimental Examination of the Hayek Hypothesis." *Economic Inquiry* 20(1982a):165–79.

———. "Microeconomic Systems as an Experimental Science." *American Economic Review* 72(1982b):923–55.

———. "Experimental Methods in Economics." *The New Palgrave: Allocation, Information, and Markets.* J. Eatwell, M. Milgate, and P. Newman, eds., pp. 94–111. New York: W.W. Norton and Company, 1989.

Stael von Holstein, C-A.S. *Assessment and Evaluation of Subjective Probability Distributions.* Stockholm: The Economic Research Institute, 1970.