HANDBOOK OF EXPERIMENTAL ECONOMICS RESULTS

Edited by

CHARLES R. PLOTT
California Institute of Technology

and

VERNON L. SMITH
Chapman University

Notice
No responsibility is assumed by the publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made

For information on all North-Holland publications visit our website at books.elsevier.com

Working together to grow
libraries in developing countries

www.elsevier.com | www.bookaid.org | www.sabre.org

ELSEVIER   BOOK AID International   Sabre Foundation

Chapter 115

ON THE PERFORMANCE OF THE LOTTERY PROCEDURE FOR CONTROLLING RISK PREFERENCES

JOYCE E. BERG and THOMAS A. RIETZ

Henry B. Tippie College of Business, University of Iowa, Iowa City, IA 52242-1000, USA

JOHN W. DICKHAUT

Carlson School of Management, University of Minnesota, 321 19th Avenue South, Minneapolis, MN 55455, USA

1. Introduction

In theory, the lottery-based induction procedure for controlling risk preferences (Roth and Malouf, 1979; Berg et al., 1986) allows experimenters to induce subjects to display pre-specified risk preferences. Thus, like induced value theory (Smith, 1976), it extends experimenters' ability to perform controlled laboratory tests by controlling parameters, such as preferences, that are exogenous to the theory being tested. Theoretically, the procedure is incontrovertible. However, the empirical performance of the procedure is an open issue. In this paper, we present the theoretical basis for the induction procedure and examine how it works in practice, using two example sets of data.

2. Inducing Risk Preferences in Theory

Often, experiments are conducted with explicit dollar payoffs or with units of exchange (e.g., "francs") that later convert to dollars at a fixed rate. Subjects' preferences over wealth can be concave (risk averse), convex (risk seeking), linear (risk neutral) or a combination of these in different regions. Thus, such payment mechanisms leave risk preferences uncontrolled. To see how this can affect choice, consider the utility function depicted by the heavy, curved line in the left panel of Figure 1, where the horizontal axis is denominated in dollars (or francs converted directly into dollars). The utility function depicted represents a person who is risk averse in his choices among monetary gambles: the person strictly prefers the expected value of the gamble to the gamble itself. If the graph were convex instead, the person would be risk loving and strictly prefer the gamble to the expected value of the gamble. That is, depending on their risk preferences, subjects may make different choices among risky alternatives. This is problematic in an experiment when we wish to determine whether behavior is in accordance with a particular theory and the predictions of the theory depend on risk preferences. Deviations can occur because risk preferences differ from those assumed
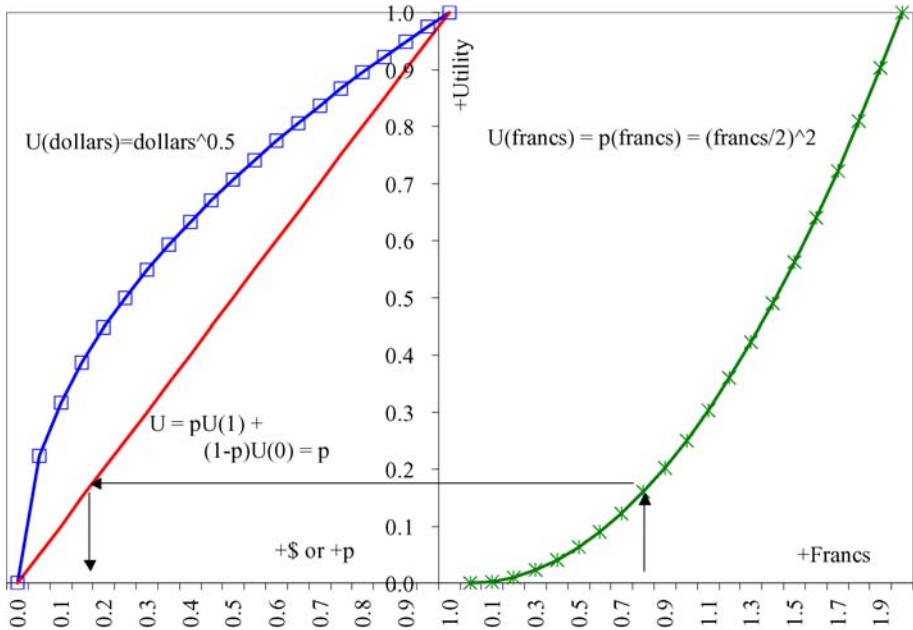
Figure 1. A graphical depiction of inducing a risk averse subject to have risk seeking preferences. The left panel shows the utility function for a subject with the risk averse utility function: $U(\text{dollars}) = \text{dollars}^{0.5}$. The straight line gives the expected utility function for a gamble with a \$1 payoff with probability p. The utility function is normalized so that $U(1) = 1$ and $U(0) = 0$. (This can be done with any utility function since expected utility is unique up to a positive affine transformation.) Then, the expected utility equals p. The right panel shows the desired, risk seeking utility function $U(\text{francs}) = (\text{francs}/2)^2$. This function maps francs into the probability of winning the \$1 prize. Since the expected utility is p, the subject's utility for francs is given by the transformation from francs into p. In this case, $U(\text{francs}) = (\text{francs}/2)^2$.

in theory or because the theory does not explain behavior. Without a reliable method of controlling for risk preferences, we cannot untangle these two possibilities.[1]

In theory, the lottery-based induction procedure allows the experimenter to control risk preferences. By using a unit of reward tied to the probability of winning the higher prize in a two-prize gamble, the lottery procedure induces expected utility maximizing subjects to behave as if they have pre-specified risk preferences relative to this unit of reward regardless of their native preferences over monetary gambles. Thus, subjects can be induced to act as if they are risk averse, risk loving, or risk neutral.

Risk preference induction depends on expected utility maximization as depicted by the lighter straight line in the left panel of Figure 1. This line shows the expected utility of a two-prize gamble with payoffs of \$0 and \$1 (following Varian, 1984, p. 159). We

---

[1] Alternatively, the experimenter could choose to measure native risk preferences and use that information in analyzing experimental results. Which choice is better depends on the experiment and its design.

have normalized utility so that the utility of $0 is 0 and the utility of $1 is 1. Thus, the bottom axis can also be interpreted as the probability of winning the $1 prize. (We can always normalize in this way since expected utility functions are unique up to a positive affine transformation.)

Preference induction relies on the result that, independent of the shape of the utility function or the size of the prizes, expected utility is linear in the probability of winning the higher of two prizes in a two-prize lottery.[2] Graphically, expected utility as a function of probability is a straight line (as shown in Figure 1) independent of the original utility function. That is,

$$E(U) = pU(M_h, \mathbf{X}) + (1 - p)U(M_l, \mathbf{X}),$$

where $p = $ probability of winning the higher valued prize, $M_h = $ higher valued prize, $M_l = $ lower valued prize, and $\mathbf{X} = $ is a vector of all other components in the utility function.

When $U(M_h, \mathbf{X})$ and $U(M_l, \mathbf{X})$ are normalized to 1 and 0 respectively, we have:

$$E(U) = pU(M_h, \mathbf{X}) + (1 - p)U(M_l, \mathbf{X}) = p.$$

Preferences are induced by using an experimental unit of exchange (say, "francs") that is later converted into the probability of winning the higher of two monetary prizes (instead of converting directly into dollars).[3] The conversion function determines how subjects should behave relative to francs. If the conversion function is $p = V(\text{francs})$ then, expected utility maximizing subjects will maximize $V(\text{francs})$ because they maximize the probability of winning the higher valued prize. Thus, they act as if they each have the utility function $V(\text{francs})$ regardless of their preferences over dollars!

Figure 1 shows how this works graphically. Suppose you would like to investigate the effect of risk-seeking behavior on market prices and want to induce the utility function, $V(\text{francs}) = (\text{francs}/2)^2$, shown in the right panel. To do this, undertake the following procedures:

(1) Have subjects trade in francs in a market with a maximum possible payoff (normalized here to 1) and a minimum possible payoff (normalized here to 0).
(2) Translate francs into the probability of winning the higher of a two-prize lottery according to the function $p = V(\text{francs}) = (\text{francs}/2)^2$.
(3) At the end of trading, run an induction lottery to determine the prize won, thus determining the ultimate payoffs to subjects.

The translation of francs to probability can be read from Figure 1. Start with the level of francs earned by the subject in the right panel. The desired level of utility for this level of francs is $(\text{francs}/2)^2$. Taking this desired level of utility into the left panel to the

---

[2] Technically, induction does not depend specifically on expected utility maximization. It works for any preference function that gives preferences that are linear in probabilities.

[3] Note that the prize does not need to be monetary. However, we discuss monetary prizes so that the lottery technique is more easily compared to induced value theory.

expected utility function (the straight line) shows the probability that must correspond to this level of francs to induce the desired preferences in francs (here, $p = (\text{francs}/2)^2$).

This procedure has been implemented in a number of different ways. Berg et al. (1986), Berg, Dickhaut, and Rietz (2003) use "spinners" where the probability of winning (chances the spinner lands in the "win area") is determined by the number of points the subject earns in a choice or pricing task and the desired induced utility function. If the spinner stops in the win area, the subject wins the higher monetary prize. Rietz (1992) uses a box of lottery tickets numbered 1 to 1000. A ticket is drawn randomly from the box. If the ticket number is less than or equal to the number of points earned by the subject, the subject wins the high monetary prize.

3. Evidence

3.1. Inducing Risk Neutrality: Evidence from Sealed Bid Auctions

We will begin with evidence from attempts to induce risk neutral preferences in sealed bid auctions. Harrison (1989), Walker, Smith, and Cox (1990), Rietz (1992) all attempt to induce risk neutral preferences in similar sealed bid auction experiments. All run series of four-person, private-value, first-price sealed bid auctions with values drawn from a uniform distribution (over a range which we will normalize to 0 to 1000). Some use dollar payoffs and some use a lottery procedure designed to induce risk neutral preferences.[4] All compare the dollar payoff results to the induced results.

Vickrey (1961) derives the symmetric Nash equilibrium bid functions for traders with risk neutral preferences as:

$$\text{Bid} = \big[(n-1)/n\big] \times \text{Value},$$

where n is the number of bidders in the auction. Thus, in these 4 person auctions, bids should be 3/4 of value for risk neutral traders. Cox, Smith, and Walker (1984) show that risk averse bidders will use a bid function with a higher slope than that of risk neutral traders. Intuitively, bidders trade off expected value for a higher probability of winning the auction. Thus, risk aversion is one possible explanation for the commonly observed "over-bidding" relative to the risk neutral bid function in sealed bid auctions (see Cox, Roberson, and Smith, 1982, 1984, 1985, 1988; Harrison, 1989; Walker, Smith, and Cox, 1990; Rietz, 1992). Alternatively, over-bidding could result from a positive intercept. Intuitively, this results from a utility of winning the auction that is independent of the profit received.

The red line in Figure 2 shows the average level of overbidding (bidding greater than the predicted risk neutral bid) as a function of value. This is measured as the difference

---

[4] Rietz (1992) also runs second-price sealed bid auctions and attempts to induce risk averse preferences in some treatments.
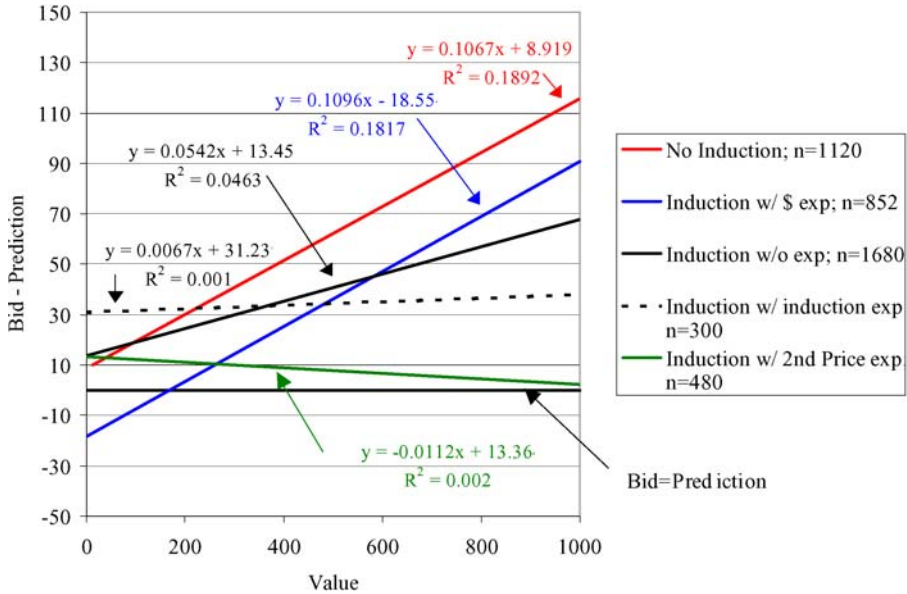
Figure 2. Least-squares trend lines for deviations in bids from the risk neutral prediction using data from Harrison (1989), Walker, Cox, and Smith (1990), and Rietz (1993). A trend line slope of "0" would indicate on-average, risk-neutral bidding behavior. Within each treatment, data is aggregated across sources. The treatments are as follows. "No Induction" includes data from dollar-valued auctions. "Induction w/$ exp" includes data from auctions in which risk neutral preferences were induced on subjects who had previously participated in dollar-valued auctions. "Induction w/o exp" includes data from auctions in which risk neutral preferences were induced on subjects without previous auction experience. "Induction w/induction exp" includes data from auctions in which risk neutral preferences were induced on subjects who had previously participated in auctions with risk neutral, induced preferences. "Induction w/2nd Price exp" includes data from auctions in which risk neutral preferences were induced on subjects who had previously participated in second-price auctions under risk neutral, induced preferences.

between the bid and the predicted bid, aggregating the data from all the dollar payoff auctions in Harrison (1989), Walker, Smith, and Cox (1990), and Rietz (1992) using a least-squares trend line. Subjects with low values bid higher than the risk neutral prediction and the amount of "over-bidding" increases with value.[5] Thus, the slope (0.1067) is greater than the risk neutral bid function prediction (0), as risk-averse preferences would predict. Inducing risk neutral preferences should flatten the slope of the bidding line. We classify the value of winning the auction as one of the "other" factors in the utility function, a factor unaffected by induction, so we do not predict that the positive intercept will decrease with induction.

---

[5] We note that, in addition to being true in aggregate, this in nearly universally true for each individual subject.

The rest of the lines in Figure 2 show, the average level of overbidding for various risk-neutral preference-induction treatments, aggregating the data from all similar-treatment auctions in Harrison (1989), Walker, Smith, and Cox (1990), and Rietz (1992) using least-squares trend lines.

When induction is attempted on subjects who have already been in dollar payoff auctions (shown as the blue line), the intercept drops but the slope (0.1096) changes little. Apparently it is difficult to break the over-bidding behavior in subjects who have had experience in dollar-payoff auctions. Rietz (1992) refers to this as hysteresis. Induction on inexperienced subjects or subjects experienced with induction in the same or similar environments meets with more success.

When induction is attempted on subjects who have no previous experience in sealed bid auctions (shown as the solid black line), there is a significant reduction in the slope of the bid function (down to 0.0542). The slope falls further (to 0.0067) when subjects return for a second set of induced-preference auctions, as shown by the dashed black line. In fact, this treatment results in a slope closest to the risk neutral prediction of 0.

Finally, when subjects are given the opportunity to learn about the induction mechanism in second-price sealed bid auctions before using it in first-price sealed bid auctions (the green line), bids conform quite closely to the risk-neutral predictions. Rietz (1992) suggests that, because there is a dominant strategy in second-price sealed-bid auctions, subjects are able to learn about the induction mechanism without having to learn about optimal strategies at the same time. Note also that, as values increase, the slight negative slope (-0.0114) results in bids becoming even closer to predictions. This is consistent with the importance of saliency in experimental payoffs. The chances of winning the auction increase and the rewards become more salient as the value increases.

Overall, the evidence from sealed bid auctions suggests that:

(1) It is more difficult to induce preferences when subjects have already formed strategies under dollar payoffs.
(2) Under induction, the behavior of inexperienced subjects conforms more closely to the risk-neutral predictions than the behavior of inexperienced subjects under dollar payoffs.
(3) Experience with the induction mechanism, especially in a similar, but less complex context, increases the correspondence between the actual outcomes and the risk neutral prediction.

Finally, Rietz (1992) also shows that risk averse preference induction results in bid functions that closely track the appropriate risk averse predictions. We will address the ability to induce risk seeking and risk averse preferences in more detail in the next two sections.

3.2. Inducing Risk Aversion and Risk Seeking: Evidence from Paired Choice Tasks

Berg et al. (1986) attempt to induce both risk averse and risk seeking preferences. Across these treatments, they compare the choices that subjects make over paired bets. The bets in a pair differ only in variance. Each bet has the same expected value, but one
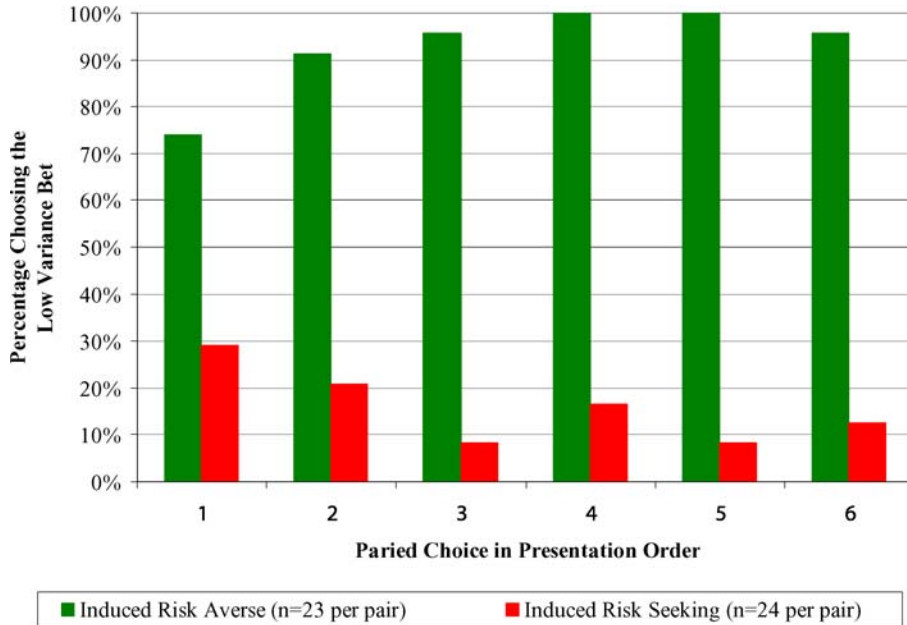
Figure 3. Percentage of subjects choosing the low-variance bet in paired choice tasks in Berg et al. (1986).
The green bars are choices made by subjects with induced risk averse preferences. The red bars are choices
made by subjects with induced risk seeking preferences.

is a relatively high variance bet while the other is a relatively low variance bet. Figure 3
shows the percentage of subjects who chose the low variance bet of each pair. Subjects
with induced risk aversion chose the low variance bet the majority of the time (100% in
some cases) and they chose it significantly more often than subjects with induced risk
seeking preferences. The evidence here suggests that inducing different risk preferences
results in a significant change in behavior as predicted.[6]

3.3. *Inducing Risk Aversion and Risk Seeking: Evidence from the*
Becker–DeGroot–Marshak Procedure

Berg et al. (1986) also study induced risk averse and risk seeking preferences using
a pricing task. Valuations for gambles are elicited as prices for the gambles using
the Becker, DeGroot, and Marschak (1964) procedure. In this incentive-compatible
procedure, subjects submit a minimum acceptable sales price for each gamble. Then,

---

[6] Prasnikar (1998) demonstrates that the comparative static results hold for a much larger set of gambles.
She also builds a method of calibrating the degree of error in induction enabling her to demonstrate more
precisely the relationship between saliency and the performance of the lottery method.
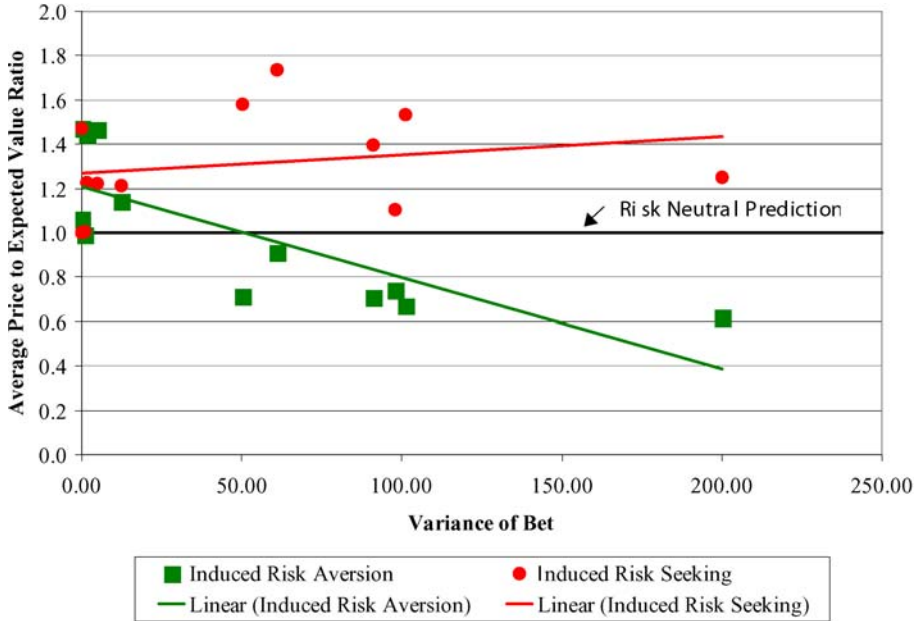
Figure 4. Ratio of average prices to expected values of gambles in Berg et al. (1986). Prices are elicited using the incentive compatible mechanism of Becker, DeGroot, and Marschak (1964). The lines "Linear (Induced Risk Aversion)" and "Linear (Induced Risk Seeking)" represent linear least-squares trend lines.

a random draw from a known distribution determines an "offer" price. If the offer price exceeds the minimum acceptable sales price, the subject sells the bet at the offer price. If not, the subject plays the bet. The dominant strategy for this pricing task is to reveal one's true value as the minimum acceptable sales price.

Figure 4 shows the ratios of average stated prices to expected values for the gambles as a function of variance. The risk neutral prediction is that prices will equal expected values, making the ratio one. Risk averse subjects should price gambles at less than expected values, with the discount increasing with risk. Risk seeking subjects should price gambles at more than expected values, with the premium increasing with risk. This pattern is clearly shown in Figure 4. The evidence suggests that inducing different risk preferences results in shifts in valuations as predicted.

4. Summary

In this chapter, we describe the lottery procedure for inducing preferences over units of experimental exchange and show how it is supported by several experiments on behavior in simple contexts.

We consider the evidence from several papers by different researchers that attempt to induce risk neutral preferences in first-price sealed bid auctions. The evidence is quite clear in these auction experiments: The experiences of subjects affect how the inducing technique performs. Experience with monetary payoffs appears to dampen the effect of the induction technique so much that results differ little from those observed under monetary payoffs. This appears to be a hysteresis effect resulting from the prior monetary payoff auctions because the results come significantly closer to the risk neutral prediction when subjects have no previous auction experience. Results come even closer to the risk neutral prediction as subjects gain experience in auctions run with the induction mechanism. Finally, the results point to the importance of simple settings as learning tasks. Convergence toward the risk neutral prediction appears to be accelerated by experience with the induction mechanism in second-price sealed bid auctions (where there is a dominant strategy for bidding).

We also review evidence from a set of paired choice and pricing tasks designed to determine whether subjects' revealed preferences over gambles are consistent with attempted risk preference induction. There is strong support for the performance of inducing when subjects choose between paired gambles. Subjects induced to be risk seeking nearly always choose the riskier gamble, while those induced to be risk averse choose the less risky one. There is similar support for pricing gambles, but the strength of the effect is a function of the variance of the gambles. This is consistent with other experimental evidence about the importance of saliency. Risk preferences matter little when there is little risk! As risk increases, risk preferences should become more important and, in fact, we see this in the experiments. Subjects appear to price gambles more consistently with their induced risk preferences as variance increases.

The lottery technique can be a powerful experimental tool. Theoretically it depends on very few assumptions and is therefore robust to many conditions. We note several of interest to experimenters:

(1) Preferences can be induced in single person or multiple person settings.

(2) The ability to induce preferences is independent of an equilibrium concept.

(3) The technique is immune to wealth changes during the experiment.[7]

(4) There is no limitation on the form of the induced preference function, $V(.)$, with the caveat that the range of V must be mapped in a 0 to 1 probability range.[8]

(5) There is no limitation on the dimensionality of the induced preference function, $V(.)$, so that $V(.)$ could be used to induce a multi-period utility function. Thus, if

---

[7] Suppose we had the subject make two choices and between choices we used the lottery technique to pay the subject. Using the technique after choice 1 would in no way alter our ability to induce using exactly the same procedure on choice two. Preferences are still linear in probability even after the wealth change and the function used to transform units of experimental exchange to probability will determine the utility function that is induced.

[8] Frequently, given the structure of economic theory (e.g., portfolio and agency theory) monotonic functions (e.g., linear or risk averse utility functions) are necessary to test the predictions of theory. However, V (francs) could be much more general and in fact non-monotonic or non-differentiable.

francs$_1$ and francs$_2$ represented the amount of francs received at the end of each of two periods then a multi-period utility function can be defined by:

$$V(\text{francs}_1, \text{francs}_2) = p.$$

(6) Preferences can be induced even when subjects are not expected utility maximizers, provided that (i) it is possible to reduce the payoffs in the setting to be one of two prizes, and (ii) preferences are linear in probability. Thus, induction should "work" for some of the proposed replacements of expected utility theory such as rank dependent utility theory and regret theory.[9]

Finally, because the lottery technique of inducing risk preferences relies on a strict subset of the axioms of expected utility theory, to reject induction is to reject expected utility theory.

Acknowledgements

References

Becker, G.M., DeGroot, M.H., Marschak, J. (1964). "Measuring utility by a single-response sequential method". Behavioral Science 9, 226–232.

Berg, J.E., Dickhaut, J.W., Rietz, T.A. (2003). "Preference reversals and induced risk preferences: Evidence for noisy maximization". Journal of Risk and Uncertainty 27, 139–170.

Berg, J.E., Daley, L.A., Dickhaut, J.W., O'Brien, J.R. (1986). "Controlling preferences for lotteries on units of experimental exchange". Quarterly Journal of Economics 101, 281–306.

Cox, J.C., Roberson, B., Smith, V.L. (1982). "Theory and behavior of single object auctions". Research in Experimental Economics 2, 1–43.

Cox, J.C., Smith, V.L., Walker, J.M. (1984). "Theory and behavior of multiple unit discriminative auctions". Journal of Finance 39, 983–1010.

Cox, J.C., Smith, V.L., Walker, J.M. (1985). "Experimental development of sealed-bid auction theory: Calibrating controls for risk aversion". AEA Papers and Proceedings 75, 160–165.

Cox, J.C., Smith, V.L., Walker, J.M. (1988). "Theory and behavior of first price auctions". Journal of Risk and Uncertainty 1, 61–99.

Harrison, G.W. (1989). "Theory and misbehavior of first-price auctions". The American Economic Review 79, 749–762.

Prasnikar, V. (1998). "How well does utility maximization approximate subjects' behavior? An experimental study". Mimeo.

---

[9] Even for prospect theory when probabilities are bounded away from the endpoints, the valuations of outcomes are weighted by a monotonic function, $\varphi(p)$, of the probability of the preferred outcome. Thus, in principle, if we could determine $\varphi(p)$, we would be able to induce an arbitrary function under prospect theory by mapping francs into the probability of winning the larger prize using $p = \varphi^{-1}(V(\text{francs}))$. Then, subjects would act as if maximizing $V(\text{francs})$.

Rietz, T.A. (1993). "Implementing and testing risk preference induction mechanisms in experimental sealed bid auctions". Journal of Risk and Uncertainty 7, 199–213.

Roth, A.E., Malouf, M.W.K. (1979). "Game-theoretic models and the role of bargaining". Psychological Review 86, 574–594.

Smith, V.L. (1976). "Experimental economics: Induced value theory". American Economic Review 66, 274–279.

Vickrey, W. (1961). "Counterspeculation, auctions, and competitive sealed tenders". The Journal of Finance 16, 8–27.

Walker, J.M., Smith, V.L., Cox, J.C. (1990). "Inducing risk neutral preferences: An examination in a controlled market environment". Journal of Risk and Uncertainty 3, 5–24.