Incentive compatible elicitation procedures

Duncan James[1]

[1]*Economics Department, Fordham University, NY, USA*
Email: dujames@fordham.edu

Abstract: This paper presents new elicitation procedures that can be used to obtain data on subjects' choices in probabilistic settings. These procedures are incentive compatible and interesting in their own right. They can also be viewed as either complements to or alternatives to Becker-DeGroot-Marschak, depending on the aims of the investigator using these procedures. Additionally, results obtained in demonstrating the use of these procedures provides further evidence (corroborating that in Plott and Zeiler (2005)) on the transience of valuation gaps across institutions and further evidence (corroborating that in James (2007)) that deviation from risk neutrality and subject error are correlated.

1. INTRODUCTION

This paper introduces, and proves the incentive compatibility of, two incentive compatible elicitation procedures. These two procedures are "dual" procedures, but not to each other. That is, one of the new procedures can be thought of as the dual to the selling-version of BDM, and the other new procedure can be thought of as the dual to the buying-version of BDM. These procedures allow for an "out of sample" test of conjectures previously made with regard to subject error and non-risk-neutral behavior, including the possibility of a relationship between playing dominated strategies and exhibiting non-risk-neutral behavior. Data from these procedures, from experiments run at Florida State University and Fordham University during 2006 and 2007, are presented which speak to this topic and to the issue of valuation gaps between procedures.

Risk preference parameter estimates that are unstable over time, or inconsistent across institutions/elicitation procedures, force us to think carefully about what data on subject choice in probabilistic settings might reflect. Isaac and James (2000) and Berg, Dickhaut, and McCabe (2005) both document inconsistency in risk preference parameters estimated within subjects, across the first price sealed bid (FPSB) auction and the selling version of BDM. Specifically, subjects in both studies most often appear risk averse in the FPSB, but risk seeking in BDM. (Isaac and James note also a subset of subjects who do the opposite: they appear risk seeking in the FPSB, but risk averse in BDM.)

There is already some suggestion in the existing literature that subject error may play an important role in shaping risk preference parameter estimates based on choice data. Frederick (2005) surveys a number of papers examining the possibility of connections between problem solving ability on the one hand, and risk preferences or time preferences on the other. Notably Benjamin and Shapiro (2005) find that students with high math SAT scores are more likely to make expected value maximizing (i.e. risk neutral) choices. James (2007) finds that with the buying and selling versions of BDM, subjects who violate the dominant strategy of truthful revelation of value (which is independent of risk preferences in lotteries with degenerate end-state probabilities, and observable as the dollar stake in the 100% probability end-state) are likely to generate risk preference parameter estimates relatively far from risk neutrality. Subjects in the buying version of BDM who played dominated strategies tended to appear risk averse; subjects in the selling version of BDM who played dominated strategies tended to appear risk-seeking; subjects who correctly played the dominant strategy were more likely to have risk preference parameter estimates indistinguishable from risk neutrality.

If one is to entertain subject error as at least a partial explanation of non-risk-neutral risk preference parameter estimates, then experiments allowing for an "out-of-sample" test of this conjecture would likely be helpful. It could be argued that an "out-of-institution" test would be a particularly compelling form of out-of-sample test; an out-of-institution test using a new institution for which expected, canonical results do not yet even exist might be even more so. It is in that spirit that the institutions and design put forth in this paper should be taken. Much as Dorsey and Razzolini (2003) allowed for a fresh look at choice and estimated risk preferences in the FPSB, by re-designing the subject's operation of the FPSB relative to the traditional version (see Cox, Smith and Walker (1988)), the institutions in this paper allow for a re-examination in a fresh setting of conjectures developed with the traditional versions of BDM.

2. DUAL INCENTIVE COMPATIBLE ELICITATION PROCEDURES

2.1 The dual to the selling-version of BDM

The dual to the selling-version of BDM functions in the following manner:
1) The subject receives a cash endowment (at the beginning of each round).
2) Given a particular pair of possible cash lottery payoffs (high and low), and the subject's cash endowment for that round, all announced at the beginning of the round, the subject is asked: what probability of the high lottery payoff occurring would make the subject willing to exchange their cash endowment for ownership of the lottery proceeds? (Elicitation requesting the probability of the low lottery payoff is also possible, merely requiring a consistent re-working of the remaining steps.)

3) To establish incentive compatibility, the subject is pitted against a random-number generator (uniform distribution[0,1]) such that if the probability of occurrence for the high state that they write down is lower than the analogous value drawn from the random number generator, they will be held to have taken ownership of the proceeds to the lottery (in exchange for their endowment for that round); otherwise they retain their cash endowment for that round.

4) The conditions of the lottery are then finalized as follows: the possible payoffs are still as previously fixed and disclosed to the subject; the *probability* of the high payoff occurring in the lottery is now set equal to the number drawn from the random-number generator (and the probability of the low state is then set equal to $(1-p_{high})$). (As noted in (3), in exchange for taking ownership of the lottery proceeds, the subject would give up their cash endowment for that round, alternatively, in the event that the subject does not write down a lower value for the probability of the high payoff than that drawn from the random number generator, no exchange would take place, and the subject would keep their cash endowment for that round.)

2.1.1 Incentive compatibility of the dual to the selling-version of BDM

Suppose the agent decides to report a higher value (call it p_{over})for that probability than the probability that would be sufficient to make her indifferent (hereafter, p_{true}). Then one of three circumstances must occur:

      (1) draw $< p_{true} < p_{over}$

      (2) $p_{true} < p_{over} <$ draw

      (3) $p_{true} <$ draw $< p_{over}$

In (1) and (2), her reporting a probability higher than that which would establish indifference does not change the outcome relative to reporting p_{true}. In (1) she would still not receive the lottery; in (2) she would still receive the lottery, and with the probability of the high payoff still set equal to the draw. However, in (3) the outcome would be changed. Moreover, the outcome would change for the worse, as she would *not* receive the lottery, and the lottery would by construction have better odds than would have made her indifferent between it and her original cash holdings.

Suppose instead that the agent decides to report a lower value (call it p_{under}) to the experimenter than the probability that would be sufficient to make her indifferent (hereafter, p_{true}). Then one of three circumstances must occur:

      (1) draw $< p_{under} < p_{true}$

      (2) $p_{under} < p_{true} <$ draw

      (3) $p_{under} <$ draw $< p_{true}$

In (1) and (2), her reporting a probability lower than that which would establish indifference does not change the outcome relative to reporting p_{true}. In (1) she would still not receive the lottery; in (2) she would still receive the lottery, and with the probability of the high payoff still set equal to the draw. However, in (3) the outcome would be changed. Moreover, the outcome would change for the worse, as she *would* receive the lottery, and the lottery would by construction have worse odds than would have made her indifferent between it and her original cash holdings.

Hence either over- or under-stating the probability of a high payoff to the lottery which would be sufficient to establish indifference between that lottery and the agent's original cash holdings cannot help the agent, but could in some circumstances hurt the agent. Hence it is a dominant strategy to truthfully report that probability.

2.2 The dual to the buying-version of BDM

The dual to the buying-version of BDM functions in the following manner:

1) As in the buying-version of BDM, the subject's endowment consists solely of whatever the experimenter has decided to give her in cash prior to her commencing the valuation procedure; that is, the subject's endowment takes the form of a lump sum of working capital (and not a cash endowment per round).

2) Given a particular pair of possible cash lottery payoffs (high and low), and a third cash figure, which for convenience purposes we shall call the premium, the subject is asked: what probability of the low lottery payoff occurring would make the subject willing to take on the burden of *paying the lottery's outcome to the experimenter* in exchange for *receiving the premium with certainty*? (Elicitation requesting the

probability of the high lottery payoff is also possible, merely requiring a consistent re-working of the remaining steps.)

3) To establish incentive compatibility, the subject is pitted against a random-number generator (uniform distribution [0,1]) and is told that if the probability of occurrence for the low state that they write down is lower than the analogous value drawn from the random number generator, they will be held to have taken on the burden of paying to the experimenter the proceeds of the lottery; otherwise no exchange takes place.

4) The conditions of the lottery are then finalized as follows: the possible payoffs are still as previously fixed and disclosed to the subject; the *probability* of the low payoff occurring in the lottery is now set equal to the number drawn from the random-number generator (and the probability of the high state is then set equal to $(1-p_{low})$). (As noted in (3), in exchange for taking on the burden of paying out the lottery's eventual payoff, the subject will receive the premium with certainty; alternatively, in the event that the subject does not write down a lower value for the probability of the low payoff than that drawn from the random number generator, nothing is exchanged, and the subject's payment for that round is zero.)

The reason for referring to the cash figure which the subject stands to gain as the premium should by now be clear: this procedure is essentially incentive compatible insurance underwriting. If any exchange takes place (and as just discussed, in some cases nothing will be exchanged), the exchange will consist of the subject *giving* the rights to the lottery proceeds to the experimenter and *receiving* a cash payment.

2.2.1 Incentive compatibility of the dual to the buying-version of BDM

The proof is analogous to that used for the dual to the selling version of BDM (in section 3.1.1).

3. DESIGN

Each experimental session consisted of 52 periods, divided into 4 segments of 13 periods each; the first and last segments implemented either the dual-to-selling or dual-to-buying (but never both for a single subject); the middle two segments implemented the selling version and the buying version as controls (all subjects participated in both, but the order in which these institutions were used was reversed for half the subjects). For both the dual-to-buying and dual-to-selling institutions, elicitation was split across (but not within) subjects between elicitation of the high state probability and elicitation of the low state probability (half the subjects in each elicitation mode). The lottery end-states were $0 and $2 in every institution, in every round. The cash amount initially owned by the subject (in dual-to-selling) or offered to the subject in exchange for taking on the lottery payout to the experimenter (in dual-to-buying) was varied from round to round, from $0 up to $2 via amounts in between. (In the control segments using the selling and buying versions of BDM, probabilities across the end states were varied instead.)

The experiments each took around ninety minutes to run. Subjects were recruited from students in upper division economics classes at Fordham University and Florida State University during 2006 and 2007. 28 subjects took part in each of the dual-to-selling and dual-to-buying (56 total subjects). Prior to commencing use of an institution, the subjects read the instructions for that institution and completed two practice rounds (without monetary payment) during which the operation of the institution was reviewed and questions were fielded. The instructions were written with the objective of giving each institution its best shot at success. That is to say, the instructions expressly inform the subjects that it is in their interest to truthfully report their value, and this point is then illustrated in the instructions with cases showing the harmful effects of either over-statement or under-statement of their response probability. Upon completion of the experiment, subjects were paid their earnings from the experiment.

4. DATA

A first way to get a sense for how these institutions influence subject responses is by examining behavior over time in each institution. In each case we note, on average across subjects, a damping over time of subject deviations from the probability equivalent (P.E.) that would be submitted by a risk neutral agent. The average absolute deviation from the risk neutral benchmark is 10.2% (per subject, per round) in the early segment of dual-to-selling, but 7.2% in the late segment. Similarly average absolute deviation drops

from 9.8% to 5.3 % in the dual-to-buying institution, from the early segment to the late segment. In each institution the data become less volatile over time.

One can also examine the data in terms to signed (+/-) deviations. To this end, note that in reporting such results probability equivalents for the dual-to-selling are represented as probability equivalent of the low state (i.e. low state lottery payout from the experimenter to the subject), while probability equivalents for the dual-to-buying as represented as probability equivalent of the high state (i.e. high state lottery payout from the subject to the experimenter). This provides for a sign convention and an orientation to graphs such that above (resp., below) the zero axis corresponds to a greater (resp., lesser) willingness to take on the proposed exchange, and thus observations above (below) the zero axis correspond to risk-seeking (risk aversion). In the dual-to-selling, the average signed deviation from the risk neutral benchmark is -5.2% (per subject, per round) in the early segment, and -4.8% in the late segment. In the dual-to-buying, average signed deviation from the risk neutral benchmark is +0.6% in the early segment, and -1.6% in the late segment. The dual-to-buying yields data closer to risk neutrality.

Figures 1-3 present the data as time series of average (signed) deviations. As seen in Figure 3, the difference between the responses generated by the two institutions diminishes over time; this is consistent with the findings of Plott and Zeiler (2005). (Note also that the break in the series in each graph corresponds to the middle segments employing buying and selling versions of BDM as controls; these institutions did indeed return results typical of their use elsewhere (e.g. James (2007)), with average sell offers above expected value in the selling version, and buy offers below expected value in the buying version. Further discussion of them is precluded by concerns of focus and space.)
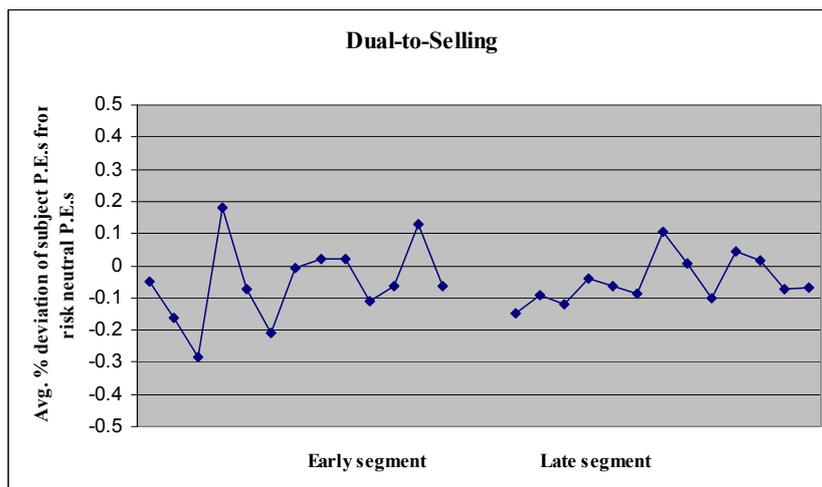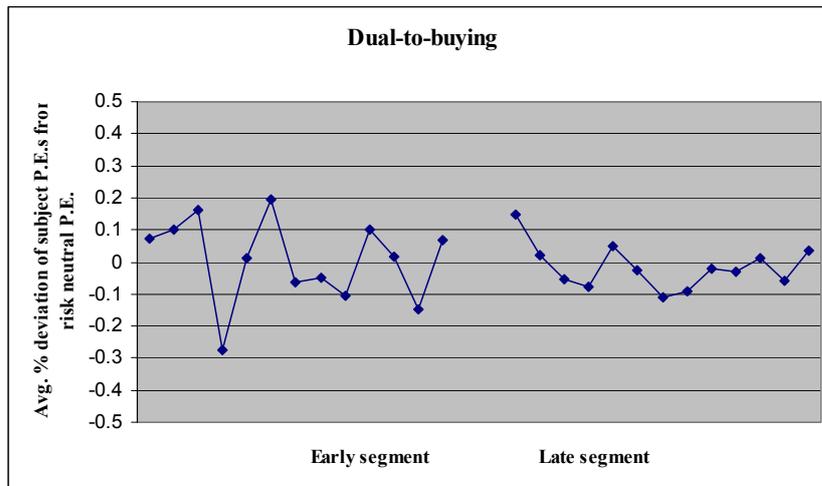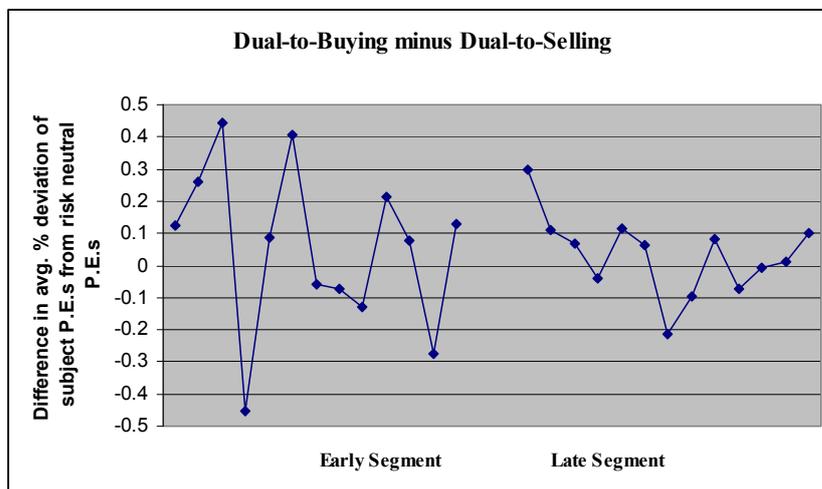

Figure 1

Figure 2

Figure 3

A second way to assess influences on the subject responses observed is made possible by the incentive compatible nature of these institutions. The inclusion of rounds where the lottery is degenerate allows for an appraisal of subject apprehension/utilization of the dominant strategy of truthful revelation. Since any subject, regardless of risk preferences, should respond with a probability equivalent for the low state of 1 (resp. 0) in dual-to-selling if their endowment in that round is equal to the low state (resp. high state) lottery payoff, we can detect failure to apprehend and play the dominant strategy on the part of a given subject on the basis of her responses in such rounds. (An analogous construction holds for the dual-to-buying.) Doing so we find that that exact revelation goes from 32% to 64% from the first segment to the second segment of the dual-to-selling, while approximate revelation, where subjects come within .01 of the correct response, goes from 50% to 96%. For dual-to-buying, we find that exact revelation stays at 42% in both segments, while approximate revelation increases from 68% to 75%.

A third approach is to check whether there is a relationship between deviation from risk neutrality and adoption of the dominant strategy. In the first segment, greater absolute deviation from risk neutrality and failure to adopt the dominant strategy are positively correlated in cross-section (for the Wilcoxon rank-sum test, calculated $z = -2.384$, critical $z = -1.6$). In the second segment, after adoption of the dominant strategy has increased and response variability has dropped, the correlation goes away. These results are similar to those in James (2007), and thus establish "parallelism" by extending the prior result to a new setting.

James, D. Incentive compatible elicitation procedures

5. Discussion

The dual-to-selling and dual-to-buying institutions represent new tools available for use by empirical economists. In addition, the initial results from their use complement existing results on several topics. As noted in the section 5, the results here corroborate those in James (2007), both in terms of noting mutability of subject response over the course of an experiment and in terms of detecting a possible relationship between subject skill in an institution and the type of responses they give. In addition, the results here can be viewed as corroborating the results of Plott and Zeiler (2005). That is, the difference in responses across the institutions discussed in this paper can be thought of as an example of a new kind of "willingness to pay-willingness to accept gap". But as in Plott and Zeiler, that gap appears to be unreliable, and decreasing over time (see Figure 3).

Finally, it has been suggested in the context of the selling version and buying version of BDM (Kachelmeier and Mehata (1992)) that responses from the selling version are to be discounted more than responses from the buying version, on the grounds that the subjects in the selling version incorrectly try to "negotiate" with the experimenter over sale of the lottery each subject owns (as of the beginning of a given round); applying this argument leads to discarding risk-seeking results from the selling version, and retaining risk averse results from the buying version, as accepted empirical results. But what then are we to make of the results from the dual institutions? The dual-to-selling institution produces results which are risk averse; might not these risk averse data be subject to the same criticism that elsewhere leads to dismissal of risk seeking data, as the dual-to-selling is essentially the selling version run in reverse? Conversely, the dual-to-buying institution produces data closer to risk neutrality than does dual-to-selling; given that dual-to-buying is essentially the (elsewhere favored) buying version run in reverse, does that recommend the data closer to risk neutrality?

REFERENCES

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, 9, 226-232.

Benjamin, Daniel J. and Jesse M. Shapiro. "Who is Behavioral? Cognitive Ability and Anomalous Preferences." Harvard University working paper 2005.

Berg, Joyce, John Dickhaut, and Kevin McCabe. (2005), Risk Preference Instability Across Institutions: A Dilemma."*Proceedings of the National Academy of Sciences*, 102, 4209-4214.

Cox, James C., Vernon L. Smith, and James M. Walker. "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1988, 1, 61-99.

Dorsey, Robert and Laura Razzolini (2003), "Explaining Overbidding in First Price Auctions Using Controlled Lotteries," *Experimental Economics*, 6, 123-140.

Frederick, Shane. "Cognitive Reflection and Decision Making", *Journal of Economic Perspectives*, 2005, 19, 25-42.

Isaac, R. Mark and Duncan James (2000), Just Who Are You Calling Risk Averse? *Journal of Risk and Uncertainty*, 20, 177-187.

James, Duncan (2007), Stability of Risk Preference Parameter Estimates within the Becker-DeGroot-Marschak Procedure, *Experimental Economics*, 10(2), pp.123-141.

Kachelmeier, Steven J. and Mohamed Shehata. "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People's Republic of China." *American Economic Review*, 1992, 82, 1120-1141.

Plott, Charles and Kathryn Zeiler. "The Willingness to pay-Willingness to Accept Gap, the "Endowment Effect," Subject Misconceptions and Experimental Procedures for Eliciting Valuations." *American Economic Review*, 2005, 95, 530-545.