Eliciting risk attitudes - How to avoid mean and variance bias in Holt-and-Laury lotteries

Norbert Hirschauer Oliver Musshoff Syster C. Maart-Noelck Sven Gruener (preprint version: June 2013)

The final publication is available at http://www.tandfonline.com (2014, Applied Economics Letters, 2014, 21(1): 35-38; DOI: 10.1080/13504851.2013.835474).

Abstract

This article shows that including inconsistent subjects in a Holt-and-Laury analysis will bias the mean, as well as the variance of the risk attitudes of the subject group of interest to an extent that cannot be determined a priori and that must not be neglected. One might be tempted to simply drop inconsistent subjects from the analysis to avoid such biases in a population-level analysis. Unfortunately, however, this is not a solution: first, the sample size may fall to an unacceptably low level. Second – and even more important – simply dropping inconsistent subjects from the analysis may introduce another unknown bias since systematic differences may exist in the risk preferences of those who answer consistently and those who do not. One must thus conclude that, if the group of interest contains a large proportion of inconsistent subjects, the whole set-up of the Holt-and-Laury lottery (HLL) experiment must be critically reconsidered and the experiment eventually repeated.

Keywords: Holt-and-Laury lottery; random-choice bias; population-level analysis; risk attitude

JEL Classification: D81; C90

1 Eliciting subjective risk attitudes – where are we?

Starting with Binswanger (1981) as an early precursor, over the last decades economists have increasingly used incentivized laboratory experiments to elicit subjective risk attitudes. Holt and Laury (2002) proposed a specific discrete choice experiment in which subjects are presented with a menu of lottery choices. The procedure has become known as Holt-and-Laury lottery (HLL). Conventionally the menu comprises ten consecutive choices between paired lotteries. Along this sequence of choices, the transition from the less risky ("safe") lottery A to the more risky lottery B is rewarded by an increasing risk premium. While subsequently being transformed into a risk aversion coefficient, risk attitudes are initially measured by an individual's "number of safe choices" (HLL-value) before crossing over to the riskier lottery B. In the last decade the HLL has virtually become *the* standard method for eliciting subjective risk attitudes.

The information obtained from a HLL can be used in two different ways: first, the heterogeneity of the obtained *individual* risk attitudes may be used – very generally speaking – as one of several variables in econometric modeling. Second, the analysis may be directly concerned with *population-level* risk attitudes. This boils down to the question whether different groups of people differ with regard to the mean and the variance of their risk preferences. Our note is concerned with the latter type of analysis. In population-level estimations a bias may occur if standard HLL procedures are replicated without considering the possibility of random responding by those whose risk attitudes are to be assessed. Making random choices between lottery A and B represents an inconsistent response behavior because the risk premium offered in the HLL increases monotonically along the sequence of the ten paired lottery choices.

The problem of inconsistent behavior in population-level analysis has been noted in many studies. Many authors (including Holt and Laury themselves) argue, however, that the bias regarding the average number of safe choices is negligible because only few subjects behave inconsistently and switch back and forth (e.g., Abdellaoui et al., 2011; Holt and Laury, 2002; Houser et al., 2010). Quite in contrast to that other authors find high inconsistency levels (e.g., Charness and Viceisza, 2011; Jacobsen and Petrie, 2009; Galarza, 2009).

Unfortunately, a generally recognized standard to deal with inconsistent responding in a HLL has not been established as yet. On the contrary! Some population-level studies, even though they find high inconsistency rates, do not correct for the resulting mean and variance bias (e.g., Charness and Viceisza, 2011; Galarza, 2009; Jacobsen and Petrie, 2009). Other studies propose a variety of differing approaches to deal with inconsistent subjects in population-level-analysis (e.g., Deck et al., 2008; Holm et al., 2012; Masclet et al., 2009).

2 Mean and variance bias in Holt-and-Laury lotteries

Some seemingly inconspicuous subtleties determine the type and the extent of the random-choice bias. These subtleties are associated with the question of how inconsistent choices, such as moving *back* to the "safe" lottery A after having previously crossed over to the riskier lottery B, are addressed.

Three procedural HLL variants have been proposed to deal with inconsistent choices:

(1) Inconsistent subjects are dropped from the analysis on the grounds that they have not understood the game. Their apparently random choice between lotteries is not seen as an indication of their risk attitude. For consistent subjects, who do not switch back from B to A, the number of safe choices is determined by their transition from A to B (e.g., Holm et al., 2012).

(2) Inconsistent subjects are included in the analysis, but the number of safe choices is determined by totaling an individual's *overall* choices for lottery A (e.g., Holt and Laury (2002); Deck et al., 2008).

(3) Inconsistent subjects are included in the analysis and the number of safe choices is determined by an individual's *initial* transition to the riskier lottery B (e.g., Masclet et al., 2009).

For any group whose members show no (or very few) inconsistencies, procedures (1), (2) and (3) coincide. Holt and Laury (2002) found little inconsistency in their group of students. Consequently, they could resort to procedure (2) – or could have resorted to (3), for that matter – without being confronted with a significant bias problem. However, poor education, different cultural backgrounds, miscommunication and a general unfamiliarity or distaste of lotteries may cause inconsistencies in HLL choices. Table 1 illustrates the problem by contrasting HLL results for a pool of German students with those for a group of Kazakh farmers.

Table 1: Effects of HLL variants on the number of group members (n), and the group-mean (μ) and -variance (σ^2) of the number of safe choices

	German students	Kazakh farmers
Procedure (1): inconsistent subjects dropped from the analysis; number of safe choices determined by an individual's transition from A to B	$n = 99$ $\mu = 5.9$ $\sigma^2 = 3.0$	$n = 43$ $\mu = 5.8$ $\sigma^2 = 14.9$
Procedure (2): inconsistent subjects included in the analysis; number of safe choices determined by totalling an individual's choices for lottery A	$n = 104$ $\mu = 5.9$ $(p = .859)$ a) $\sigma^2 = 2.9$ $(p = .864)$ b)	$n = 99$ $\mu = 5.5$ $(p = .337)$ a) $\sigma^2 = 7.4$ $(p = .001)$ b)
Procedure (3): inconsistent subjects included in the analysis; number of safe choices determined by an individual's initial transition from A to B	$n = 104$ $\mu = 5.8$ $(p = .581)$ a) $\sigma^2 = 3.6$ $(p = .300)$ b)	$n = 99$ $\mu = 3.1$ $(p < .001)$ a) $\sigma^2 = 12.5$ $(p = .296)$ b)

a) p-value of a Mann-Whitney-U-test in comparison to procedure (1).

b) p-value of a Levene-test in comparison to procedure (1).

Similar to the results reported by Holt and Laury (2002), little inconsistency is found in the pool of students. Hence, procedures (1), (2) and (3) generate near-identical results without a significant bias in mean or variance. According to a Mann-Whitney-U-test, the null hypotheses of no difference in mean between procedures (2) and (1) and between procedures (3) and (1) *cannot* be rejected. The same applies to the variance according to Levene-test statistics. The picture is different for Kazakh farmers. Within this group, 57% of subjects show inconsistent behavior in that they switch forth and back. How inclusion of these subjects in the analysis affects the group-mean and -variance depends on how inconsistent behavior is dealt with procedurally when determining the number of safe choices. For the sake of easy demonstration of biasing mechanisms that occur within the different procedural designs, we assume that subjects who show inconsistent behavior have not understood the game and hence make random choices when asked to choose either lottery A or B ("random-choice subjects").

Mean-biasing mechanism in procedure (2): Including "random-choice subjects" in the analysis and equalizing the total number of A-choices with the number of safe choices boils down to adding noise in the form of a binomial distribution $B_{m,p}(k)$, with $m = 10$ (= number of repeated lottery choices), $p = 0.5$ (= probability of choosing one or the other lottery) and $k \in \{0, \dots, 10\}$ representing the HLL scores that may result from random choice. The inclusion of this probability mass distribution with its mean of 5 shifts the group-mean of the number of safe choices – except for one chance constellation: no relevant shift in mean is caused *if* the mean number of safe choices also coincidentally amounts to approximately 5 within the consistent subgroup.

Variance-biasing mechanism in procedure (2): At first sight, one might be surprised that the inclusion of noise, according to procedure (2), reduces the group-variance of the number of safe choices. This observation, however, is easily explained. Given the binomial distribution $B_{10,0.5}(k)$, the numbers k of A-choices (and thus HLL scores) cluster around 5, with nearly two thirds of scores being between 4 and 6. The inclusion of these scores reduces the variance of the entire group if the consistent subgroup is more heterogeneous than the inconsistent group and has, by coincidence, a mean number of safe choices of approximately 5. The variance reduction, as observed in our example, does not constitute a general effect, however. Quite on the contrary: depending on how far the mean number of safe choices within the consistent subgroup deviates from 5, the inclusion of a bulk of scores around 5 may reduce, or not change, or increase the variance of the entire group.

Mean-biasing mechanism in procedure (3): When including inconsistent subjects in the analysis and determining the number of safe choices by counting A-choices until the *initial* transition to the riskier lottery B equates to including a positively skewed discrete distribution. Its probability mass function and thus the probabilities P of inconsistent HLL scores are quickly derived by using the multiplication rule: $P(HLL = 1) = 0.5^1$, $P(HLL = 2) = 0.5^2 = 0.25$, …, and $P(HLL = 10) = P(HLL = 0) = 0.5^{10} = 0.00098$. That is, the bulk of inconsistent scores are clustered towards the lower end of the zero-to-ten-scale, with approximately 94% of scores being below 5 and only 3% scores being above 5. By including these inconsistent scores with their mean of 2, a downward bias of the group-mean is generated,

Variance-biasing mechanism in procedure (3): The variance of the skewed distribution resulting for inconsistent HLL scores according to procedure (3) is relatively small, with 75 of scores being either one or two. Its impact on group-variance depends on the specific context under consideration. If the consistent subgroup were clustered around the same mean and if it were more heterogeneous than the inconsistent subgroup, group-variance would be reduced. In a more realistic constellation, in which the mean of the consistent subgroup deviates from the mean of 2 of the inconsistent subgroup, the inclusion of the inconsistent group, even though it has little variance in itself, will increase group-variance. While, as a consequence of these opposite effects, no significant variance bias was found in our example, we can neither exclude upward nor downward variance bias for procedure (3) in general.

The risk of obtaining distorted HLL means and variances increases if experiments are carried out inappropriately, i.e. without providing subjects with proper information on the "rules of the game" or without guaranteeing incentive compatibility. One should note, however, that, despite all efforts, the inclusion of inconsistent subjects cannot be completely avoided when carrying out HLL.

3 Conclusion

As shown above, the inclusion of inconsistent subjects, who make random choices between HLL lottery pairs, may distort the mean as well as the variance of the risk attitudes of the subject group of interest. We have termed this problem "random-choice bias" in population level analysis. Type, extent and direction of the random-choice bias depend on how inconsistent choices are dealt with methodically.

According to Holt and Laury (2002) it does not matter how to handle little inconsistency in the subject group of interest. In the case of "little inconsistency" all three above-mentioned procedures could be used to deal with inconsistent subjects in population-level analysis. However, it is neither known a priori how much inconsistency occurs nor how much bias it produces. Studies on the risk attitudes of groups which ignore the problem of inconsistency, discredit HLL because they introduce both mean and variance bias to an unknown degree.

Even though our analysis has clarified that the "random-choice bias" must not be neglected, the methodological conclusion that is to be drawn if multiple-switching behavior is found is not obvious. While one might ostensibly argue that no information regarding a population's risk attitude is gained by including "random-choice subjects" into HLL, we cannot solve the problem by simply dropping inconsistent subjects from the analysis: first, the sample size may fall to an unacceptably low level that prevents a meaningful analysis. Second – and even more important – systematic differences may exist in the risk preferences of those who answer consistently and those who do not. Simply dropping inconsistent subjects from the analysis may thus introduce a new bias in population-level analysis the magnitude of which cannot be assessed either. Hence, if the group of interest contains a large proportion of inconsistent subjects, the whole setup of the HLL experiment should be critically reconsidered. This may mean to better adjust the general experimental set-up to suit the exigencies of the context under consideration and to repeat the experiment. It may also include the verification of the trustworthiness of HLL data by triangulating methods. Triangulation may simply imply to consider alternative framings: in addition to using the wording "lottery", the HLL choice could be framed as an investment decision, and instead of changing probabilities, changing payoffs could be used in the consecutive lottery choices (cf. Brick et al. 2012). Triangulation may also mean to include non-experimental methods such as psychometric scales.

References

Abdellaoui, M., Driouchi, A., L'Haridon, O., 2011. Risk aversion elicitation: reconciling tractability and bias

Binswanger, H.P., 1981. Attitudes toward risk: theoretical implications of an experiment in rural India. The Economic Journal 91(364), 867-890.

Brick, K., Visser, M., Burns, J. (2012). Risk Aversion: Experimental Evidence from South African Fishing Communities. American Journal of Agricultural Economics 94(1), 133–152.

Charness, G., Viceisza, A., 2011. Comprehension and risk elicitation in the field: evidence from rural Senegal. IFPRI discussion papers 1135, International Food Policy Research Institute (IFPRI).

Deck, C., Lee, J., Reyes, J. et al. (2008) Measuring risk attitudes controlling for personality traits. Working Paper No. 0801, Department of Economics, Florida International University, Miami.

Galarza, F., 2009. Choices under risk in rural Peru. MPRA Paper 17708, University Library of Munich.

Holm, H.J., Opper, S., Nee, V., 2012. Entrepreneurs under uncertainty: an economic experiment. Working Papers 2012:4, Lund University, Department of Economics.

Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. American Economic Review 92(5), 1644-1655.

Houser, D., Schunk, D., Winter, J., 2010. Distinguishing trust from risk: an anatomy of the investment game. Journal of Economic Behavior and Organization 74(1-2), 72-81.

Jacobson, S., Petrie, R., 2009. Learning from mistakes: what do inconsistent choices over risk tell us? Journal of Risk and Uncertainty 38(2), 143-158.

Masclet, D., Colombier, N., Denant-Boemont, L., Lohéac, Y., 2009. Group and individual risk preferences: a lottery-choice experiment with self-employed and salaried workers. Journal of Economic Behavior and Organization 70(3), 470-484.