Revealed Preferences for Risk and Ambiguity

Donald Brown, Chandra Erdman, Kirsten Ling and Laurie Santos

Paris - Dauphine

November 2010

Choice under Uncertainty: Risk and Ambiguity

(1) *Risk, Uncertainty and Profit* (1921) by Knight — see his chapter on "The meaning of risk and uncertainty."

(2) *The General Theory of Unemployment, Interest and Money* (1936) by Keynes — see his chapter on "The state of long-run expectation."

(3) A central contribution of Knight and Keynes is the distinction between risk and ambiguity.

A Quotation from Keynes (1937)

"By uncertain knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know."

This distinction between risk and ambiguity is absent in the expected utility (EU) model of decision-making under risk, due to Von Neumann and Morgenstern (1944), and Savage's (1954) model of decision-making under uncertainty, but it is the genesis of Ellsberg's (1961) seminal critique of Savage's theory of subjective expected utility (SEU).

The Ellsberg Paradox

(1) In Ellsberg's two-color thought experiment, subjects make pairwise choices between a risky urn, where the relative frequencies of the two outcomes are 1/2, and an ambiguous urn, where the relative frequencies are unknown.

(2) In the first trial, if the subject chooses an urn and draws a black ball then she receives $100, the "good" outcome, but if she draws a white ball then she receives $0.00 dollars, the "bad outcome."

(3) In the second trial the payoffs are reversed.

Preferences for Ambiguity

(1) Subjects that choose the ambiguous urn on both trials are said to be ambiguity-seeking and subjects that choose the risky urn on both trials are said to be ambiguity-averse.

(2) Ambiguity-seeking subjects in the Ellsberg experiment "act as if," the perceived probability of the "good" outcome is greater than the relative frequency of the "good" outcome.

Preferences for Ambiguity (Continued)

(3) Ambiguity-averse subjects in the Ellsberg experiment "act as if" the perceived probability of the "bad" outcome is greater than the relative frequency of the "bad" outcome.

(4) Ellsberg's (1961) explanation of the two color Ellsberg paradox: "... we would have to regard the subject's subjective probabilities as being dependent upon his payoffs, his evaluation of the outcomes ... it is impossible to infer from the resulting behavior a set of probabilities for events independent of his payoffs."

(5) That is, agents choose actions and beliefs.

The Duke Model

(1) Huettel et al. (2006) in The Center for Cognitive Neuroscience at Duke, proposed a new model of decision-making under uncertainty with proxies for risk-aversion and ambiguity-aversion, consistent with Ellsberg's explanation of the two color Ellsberg paradox.

(2) The proxies are β for risk-aversion, where β is the coefficient of relative risk-aversion for the utility function $u(w) = w^{\beta}$, and α for ambiguity-aversion in the α-max min expected utility model.

(3) If $\beta = 1$ then u is linear and the subject is risk-neutral. If $\beta > 1$ then u is convex and the subject is risk-seeking. If $\beta < 1$ then u is concave and the subject is risk-averse.

The Duke Model (Continued)

(4) The α-max min expected utility of an ambiguous lottery, $x = (x_1, x_2)$ is $(1 - \alpha)u(x_1 \vee x_2) + \alpha u(x_1 \wedge x_2)$.

(5) Huettel et al. interpret α as a measure of ambiguity-aversion, where $\alpha \in [0, 0.5)$ denotes ambiguity-seeking, $\alpha = 0.5$ is ambiguity-neutral, and $\alpha \in (0.5, 1]$ denotes ambiguity-averse.

Resolving the Ellsberg Paradox with the Duke Model

(1) The utility of an ambiguous lottery in the Duke model is the α-max min expected utility and the utility of a risky lottery is the expected utility. The Duke model assumes that subjects maximize utility in choosing between a pair of lotteries.

(2) The Huettel et al. model is consistent with Ellsberg's explanation of the two-color paradox. The utility of the risky urn is $[u(0) + u(100)]/2$ and the utility of the ambiguous urn is $(1-\alpha)u(100) + \alpha u(0)$, where $u(0) = 0$. If the agent is ambiguity-averse then $(1-\alpha) \in [0, 0.5)$. Hence $u(100)/2 > (1-\alpha)u(100)$ and the agent chooses the risky urn on both trials. If the agent is ambiguity-seeking then $(1-\alpha) \in (0.5, 1]$. Hence $u(100)/2 < (1-\alpha)u(100)$ and the agent chooses the ambiguous urn on both trials.

The Duke Experiment

(1) Huettel et al. asked 13 subjects to make pairwise choices between lotteries with different degrees of uncertainty, i.e., certain, risky and ambiguous, and used *fMRI* data to identify regions in the brain that are activated during the choice process.

(2) For each subject, β is chosen to maximize the number of correct predictions in the risky-risky and risky-certain trials, using the expected utility model. The *fMRI* data identified a region of the brain that is activated during the choice process, call it region R.

The Duke Experiment (Continued)

(3) Given the estimated β, α for each subject is chosen to maximize the number of correct predictions in the ambiguous-risky and ambiguous-certain trials, using the α-max min expected utility model. The *fMRI* data identified a different region of the brain that is activated during this choice process, call it region A.

(4) A is inactive when R is active and R is inactive when A is inactive. As is common in the neural science literature, this double dissociation *fMRI* study is interpreted as independence of the two choice behaviors.

Limitations of the Duke Experiment

(1) Unfortunately, the estimation procedure in the Duke study is not identified, i.e., there are several values of α and β that maximize the number of correct predictions.

(2) Recently, Levy et al. (2010) offered an alternative explanation of the hypothesized finding of differential activation in parts of the brain as a consequence of choice under risk and ambiguity.

Limitations of the Duke Experiment (Continued)

(3) In the Duke study subjects were told, ex post, the probabilities used to resolve ambiguous lotteries, possibly allowing subjects to learn the ambiguous probabilities.

(4) In the Levy et al. study, where the ambiguous probabilities were not resolved, the levels of neural activation in regions R and A resulting from choice under risk and ambiguity were comparable.

The Mixed Logit Model

(1) Given these limitations of the Duke experiment, we recast the Duke model as a random utility model, more specifically a parametric, mixed logit model. The mixed logit model allows us to estimate a parametric, bivariate distribution over α and β from pairwise choices between risky and ambiguous lotteries made by subjects randomly selected from the population.

(2) The random utility model was first proposed in psychology by Thurstone (1927) in a form now called the binomial probit model, and subsequently introduced in economics by Marschak (1960) who investigated the properties of choice probabilities for utility functions subject to random perturbations.

(3) McFadden (1974) introduced the conditional logit model. In the binomial case, this is the well-studied logistic model in biostatistics. See McFadden's Nobel Lecture for a brief history of the origins of the random utility model.

Random Effects

(1) In the mixed logit model presented in this paper,the proxies for ambiguity-aversion and risk-aversion, α and β, are treated as random effects, i.e., random variables uncorrelated with the explanatory variables.

(2) There are two criteria for using a random effects model in lieu of a fixed effects model. First, the data is generated by taking a random sample of subjects from some fixed population.

(3) Our sample is randomly selected from the population of Yale students, matriculating in the summer session and fall term of 2009. Second, the explanatory variables — the payoffs and probabilities defining the lotteries — must be uncorrelated with the random effects, α and β.

(4) This is certainly true in our experiment in which the payoffs and probabilities defining the lotteries in the pairwise comparisons are generated randomly and independently for each subject.

The Bernstein–von Mises Theorem

(1) We interpret α and β as random effects with a bivariate log-normal distribution, parameterized by unknown hyper-parameters Ψ, using the Bayesian perspective we can estimate the posterior means.

(2) The posterior means are consistent estimates of the individual-level random effects, α_j and β_j.

(3) The Bernstein–von Mises theorem provides an alternative classical method of estimating the individual-level random effects. That is, maximum likelihood estimation of α_j and β_j.

(4) The Bernstein–von Mises theorem shows that the Bayesian and classical estimates of the individual-level random effects α_j and β_j are asymptotically equivalent.

The Yale Experiment

(1) To replicate the essentials of the Duke. experiment, we consider pairwise choices in 200 monetary lotteries made by 30 randomly chosen Yale undergraduates in 2009.

(2) As in the Duke experiment, each lottery involves choices between a known payoff, payoffs with known probabilities, and payoffs with unknown probabilities.

(3) We refer to these lotteries as certain, risky and ambiguous lotteries, respectively. In our experiment, each subject chooses between 40 risky-certain pairs, 40 risky-risky pairs, 40 ambiguous-certain pairs, 40 ambiguous-ambiguous pairs and 40 risky-ambiguous pairs.

The Yale Experiment (Continued)

(1) All ambiguous lotteries have two positive outcomes and all certain lotteries have one positive outcome.

(2) In the risky-certain pairs and the risky-risky pairs, all risky lotteries have one zero outcome and one positive outcome, but in the risky-ambiguous pairs both ambiguous and risky lotteries have two positive outcomes.

(3) At the start of each trial, subjects are given a pairwise choice between lotteries, represented by two pie charts. Subjects are instructed to choose the lottery on the left or right by typing "*L*" or "*R*." Once a choice is made, a box appears around the chosen lottery and the other lottery disappears.

The Yale Experiment (Continued)

(4) Expected values of lotteries are chosen at random, whole-dollar amounts between \$5 and \$25, and expected values of pairs of lotteries are matched within 20 percent. The probability of winning the amount presented in a certain lottery is always 1, and the probabilities of winning amounts presented in risky and ambiguous lotteries are chosen randomly between 0.25 and 0.75, and varied across gambles.

(5) Finally, the payoff of the lottery is displayed at the bottom of the screen. The figure on the next slide displays pairs of risky, certain and ambiguous lotteries. After completion of 200 trials, subjects are paid winnings from 4 randomly selected trials. Winnings ranged from \$0 to \$93 in a single trial, and \$35 to \$99 overall. Ex post, subjects are not told the probability of payoffs in an ambiguous lottery.

The Yale Model

(1) In the Duke model, α and β are interpreted as parameters and the choice probability, p_A, for x_A in the pairwise comparison between lotteries x_A and x_B is defined as the percent correctly predicted. This is not the case for the mixed logit model that we present.

(2) In our model, the choice probability, p_A, for x_A in the pairwise comparison between lotteries x_A and x_B is interpreted as the proportion of individuals in the population, with the same preferences for risk and ambiguity, that choose x_A or is interpreted as the proportion of times that a single individual chooses x_A in repeated pairwise comparisons between options x_A and x_B.

Choice Probability

(1) To derive estimates of risk-aversity, β, and ambiguity aversity, α, we follow the Huettel et al. protocol and first estimate $\hat{\beta}$, using pairwise comparisons between risky lotteries. That is, we maximize the log-likelihood defined by the logistic cdf

$$\Lambda[\eta] \equiv \frac{\exp \eta}{1 + \exp \eta}.$$

(2) If

$$X \equiv (x_1, x_2; \pi_1, \pi_2) \text{ and } Y \equiv (y_1, y_2; \eta_1, \eta_2),$$

is a pair of risky lotteries, where p_X is the probability of choosing X over Y and the expected utility of the chosen risky option X is

$$E(U(X)) \equiv \pi_1 u(x_1) + \pi_2 u(x_2)$$

then, the choice probability $p_X(U(X))$ is

$$p_X(U(X)) \equiv \Lambda[\ln E(U(X)) - \ln E(U(Y))].$$

MLE of the Yale Model

(1) We estimate α_j and β_j by maximizing the log-likelihood of each subject's pairwise choices in risky and ambiguous lotteries. Following Huettel et al., we use a two-step procedure to estimate α_j and β_j for each subject $j = 1, 2, ..., 30$.

(2) That is, our estimator is two-step maximum likelihood estimator. To estimate β_j, we ask each subject to choose between 40 risky-certain pairs and 40 risky-risky pairs of lotteries. Here we assume that each subject is maximizing expected utility, which only depends on β_j.

(3) To estimate α_j, we ask each subject to choose between 40 ambiguous-certain pairs and 40 pairs of ambiguous-ambiguous lotteries. Here we assume that each subject is maximizing α-max min expected utility, which depends on both α_j and β_j, where we use the previously estimated value $\hat{\beta}_j$ and need only estimate α_j.

(1) To estimate $\hat{\beta}$, we use the multiplicative random utility model, where the expected utility of the risky option X is given by

$$EU(X) = \pi_1(x_1)^\beta + \pi_2(x_2)^\beta$$

and the choice probability for X is

$$p_X(\beta) \equiv \mathrm{Prob}[\ln EU(X) + \varepsilon_X > \ln EU(Y) + \varepsilon_Y]$$

where ε_X and ε_Y are i.i.d. extreme value random variables.

Estimates of Risk-Aversion (Continued)

(2) The choice probability for choosing X is

$$p_X(\beta) \equiv \frac{\exp[\ln EU(X) - \ln EU(Y)]}{(1 + \exp[\ln EU(X) - \ln EU(Y)])}.$$

(3) For pairs of risky-certain and risky-risky lotteries, $x_2 = y_2 = 0$. Hence the choice probability for choosing X as a function of β is

$$p_X(\beta) = \frac{\exp[\beta \ln(x_1) + \ln \pi_1 - \beta \ln(y_1) - \ln \eta_1]}{(1 + \exp[\beta \ln(x_1) + \ln \pi_1 - \beta \ln(y_1) - \ln \eta_1])}.$$

The MLE for Each Subject's Risk-Aversion

(1) We denote the chosen lotteries as X^j in each pair of 40 risky-certain and 40 risky-risky lotteries. The likelihood of the observed risky choices in the 80 pairwise comparisons$\{(X^j, Y^j)\}_{j=1}^{j=80}$ as a function of β, is

$$\prod_{j=1}^{j=80} p_{X^j}(\beta).$$

(2) The log-likelihood

$$\frac{1}{80}\sum_{j=1}^{j=80} \ln p_{X^j}(\beta) = \frac{1}{80}\frac{\sum_{j=1}^{j=80} \ln(\exp[\beta \ln(x_1^j) + \ln \pi_1^j - \beta \ln(y_1^j) - \ln \eta_1^j]}{(1 + \exp[\beta \ln(x_1^j) + \ln \pi_1^j - \beta \ln(y_1^j) - \ln \eta_1^j]))}$$

is strictly concave in β, hence the *MLE* of $\hat{\beta}$ is identified and consistent.

The Null Hypothesis

(1) Our null hypothesis is that economic preferences for risk and ambiguity are independent, where β is a measure of the subject's tolerance for risk and α is a measure of the subject's attitude towards ambiguity.

(2) The alternative hypothesis is that economic preferences for risk and ambiguity are correlated.

(3) Under the null hypothesis, every function of α and every function of β are independent.

(4) In particular, the LL for β and risky-certain or risky-risky data is independent of the LL for α and ambiguous-certain or ambiguous-ambiguous data, for every fixed value of β, e.g., $\hat{\beta}$, the estimate of β.

Estimates of Ambiguity-Aversion

(1) To estimate α, we use the additive random utility model, where subjects evaluate ambiguous lotteries, using α-max min expected utility.

(2) Given the pair of ambiguous lotteries $W \equiv (w_1, w_2)$ and $Z \equiv (z_1, z_2)$, the logit choice probability for choosing W as a function of α, for fixed $\hat{\beta}$, is

$$p_W(\alpha, \hat{\beta}) = \frac{\exp\{[\alpha(w_1 \wedge w_2)^{\hat{\beta}} + (1-\alpha)(w_1 \vee w_2)^{\hat{\beta}}] - [\alpha(z_1 \wedge z_2)^{\hat{\beta}} + (1-\alpha)(z_1 \vee z_2)^{\hat{\beta}}]\}}{[1 + \exp\{[\alpha(w_1 \wedge w_2)^{\hat{\beta}} + (1-\alpha)(w_1 \vee w_2)^{\hat{\beta}}] - [\alpha(z_1 \wedge z_2)^{\hat{\beta}} + (1-\alpha)(z_1 \vee z_2)^{\hat{\beta}}]\}}.$$

The MLE for Each Subject's Ambiguity-Aversion

(1) We denote the chosen lotteries as W^j in each pair of 40 ambiguous-certain and 40 ambiguous-ambiguous lotteries.

The likelihood of the observed ambiguous choices in the 80 pairwise comparisons$\{(W^j, Z^j)\}_{j=1}^{j=80}$ as a function of α, for fixed $\hat{\beta}$, is

$$\prod_{j=1}^{j=80} p_{W^j}(\alpha, \hat{\beta}).$$

(2) The log-likelihood

$$\frac{1}{80}\sum_{j=1}^{j=80} \ln p_{W^j}(\alpha, \hat{\beta})$$

$$= \frac{1}{80}\sum_{j=1}^{j=80} \ln \frac{\exp\{[\alpha(w_1\wedge w_2)^{\hat{\beta}}+(1-\alpha)(w_1\vee w_2)^{\hat{\beta}}]-[\alpha(z_1\wedge z_2)^{\hat{\beta}}+(1-\alpha)(z_1\vee z_2)^{\hat{\beta}}]\}}{[1+\exp\{[\alpha(w_1\wedge w_2)^{\hat{\beta}}+(1-\alpha)(w_1\vee w_2)^{\hat{\beta}}]-[\alpha(z_1\wedge z_2)^{\hat{\beta}}+(1-\alpha)(z_1\vee z_2)^{\hat{\beta}}]\}}.$$

The MLE for Ambiguity-Aversion(Continued)

(3) If $\hat{\beta} \neq 0$, then the log-likelihood function is globally concave in α.

(4) If $\hat{\beta} = 0$, then $\hat{\alpha}$ is not identified. That is, if $\hat{\beta} = 0$, then for all $\hat{\alpha} \in [0, 1] : p_{W^J}(\hat{\alpha}, \hat{\beta}) = 1/2$.

(5) Hence $\hat{\alpha}$ is indeterminate and the six subjects with indeterminate $\hat{\alpha}$ "act as if" they flip a fair coin to choose between any pair of ambiguous lotteries. These 6 subjects were excluded from our statistical analysis.

Hypothesis Test for Correlation

To estimate the correlation between risk and ambiguity, we estimate $\hat{\alpha}$ and $\hat{\beta}$ for each subject and regress $\hat{\alpha}$ against $\hat{\beta}$, where we exclude the six subjects with unidentified $\hat{\alpha}$. The slope coefficient of the regression is not significantly different from 0 at the 0.05 level, indicating linear independence between risk-aversion and ambiguity-aversion. Hence we cannot reject the null hypothesis of independence of revealed preferences for risk and ambiguity.

Hypothesis Test for Independence

We construct a 2×2 contingency table, where the columns are labeled AA, for ambiguity aversion and AS, for ambiguity seeking,and the rows are labeled RA, for risk averse, and RS , for risk seeking, where we have omitted the six subjects with unidentified α:

5	0
19	0

Figure 2

We see that the cells in the second column are both zero. Consequently, $\text{Prob}(AA|RA) = \text{Prob}(AA|RS) = 1$. Hence in our choice experiment it follows from Fisher's exact test that the revealed preferences for risk and ambiguity are independent.