



Rejecting small gambles under expected utility

Ignacio Palacios-Huerta*, Roberto Serrano

Department of Economics, Brown University, Box B, Providence RI 02912, United States

Received 24 April 2005; accepted 7 September 2005

Available online 24 March 2006

Abstract

This paper contributes to an important recent debate around expected utility and risk aversion. Rejecting a gamble over a given range of wealth levels imposes a lower bound on risk aversion. Using this lower bound and empirical evidence on the range of the risk aversion coefficient, we calibrate the relationship between risk attitudes over small-stakes and large-stakes gambles. We find that rejecting small gambles is consistent with expected utility, contrary to a recent literature that concludes that expected utility is fundamentally unfit to explain decisions under uncertainty. Paradoxical behavior is only obtained when calibrations are made in a region of the parameter space that is not empirically relevant.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Risk aversion; Expected utility

JEL classification: D00; D80; D81

“The report of my death was an exaggeration.”

Mark Twain (May 1897)

1. Introduction

Within expected utility theory, risk aversion is identified with the concavity of the Bernoulli utility function u on wealth w (Arrow (1971), Pratt (1964); see Mas-Colell et al. (1995, chapter 6)). The expected

* Corresponding author.

E-mail addresses: ipalacios@brown.edu (I. Palacios-Huerta), roberto_serrano@brown.edu (R. Serrano).

URLs: <http://www.econ.brown.edu/~iph> (I. Palacios-Huerta), <http://www.econ.brown.edu/faculty/serrano> (R. Serrano).

utility framework has been severely criticized in a recent literature that concludes that diminishing marginal utility is an utterly implausible explanation for appreciable risk aversion.¹ The basis of the criticism can be best illustrated in Rabin (2000), who calibrates the relationship between risk attitudes over small and large stakes gambles under expected utility. Using his results, it is possible to present striking statements of the following kind: if a decision maker is a risk-averse expected utility maximizer and if he rejects gambles involving small stakes over a large range of wealth levels, then he will also reject gambles involving large stakes, sometimes with infinite positive returns. For instance, “suppose that, from any initial wealth level, a person turns down gambles where she loses \$100 or gains \$110, each with 50% probability. Then she will turn down 50–50 bets of losing \$1000 or gaining *any* sum of money,” or “suppose we knew a risk averse person turns down 50–50 lose \$100/gain \$105 bets for any lifetime wealth level less than \$350,000 . . . Then we know that from an initial wealth level of \$340,000 the person will turn down a 50–50 bet of losing \$4000 and gaining \$635,670.” (Rabin (2000, p. 1282)).

From this paradoxical, even absurd, behavior towards large-stakes gambles, Rabin (2000) and other authors conclude that expected utility is fundamentally unfit to explain decisions under uncertainty. This paper challenges this conclusion. We show that the flaw identified in this literature has little empirical support.² In particular, we show that it is the assumption of rejecting small gambles over a large range of wealth levels, and *not* expected utility, that does not typically match real-world behavior. In articulating our response, it is more useful not to argue whether expected utility is literally true (we know that it is not, since many violations of its underpinning axioms have been exhibited). Rather, one should insist on the identification of a useful range of empirical applications where expected utility is a useful model to approximate, explain, and predict behavior.

To be more specific on Rabin’s criticism, let p be the proposition “agent a is a risk averse expected utility maximizer,” let q be “agent a turns down the modest-stakes gamble X for a given range of wealth levels,” and let r stand for “agent a turns down the large-stakes gamble Y .” Then, Rabin’s statements can be expressed as: if p and q hold, then so does r . He then convincingly argues that most individuals do not turn down Y . This amounts to saying that r does not hold. From here, Rabin (2000) jumps to the conclusion that p does not hold. But this conclusion is not warranted unless q is either tautological or at least empirically compelling.

The plausibility of q as an assumption is argued in the literature purely by appealing to the reader’s introspection. Introspection, however, may sometimes be misleading: what people think they *would do* in a thought experiment may turn out to be quite different from what they *actually do* when confronted with a similar real-life situation. Indeed, we shall see that q is far from being empirically relevant. We report experimental evidence that supports our claim, but most of the evidence we wish to bring to bear is empirical. In doing so, we stress a methodological point. In assessing levels of risk aversion, it is important to study what real economic agents actually do when facing very different problems involving uncertainty. These data sets have been carefully analyzed by empirical economists, who have estimated the levels of the risk aversion coefficient over a wide range of problems.

Our methodology is the following. Rather than relying on introspection to validate q , we first investigate the theoretical implications of q on the preferences of the decision maker. We shall

¹ See, for example, Hansson (1988), Kandel and Stambaugh (1991), Rabin (2001), Rabin and Thaler (2001) and other references therein. Samuelson (1963), Machina (1982), Segal and Spivak (1990) and Epstein (1992) also study various issues that are related to this literature.

² See Watt (2002), who has independently obtained results related to ours.

demonstrate that it implies a specific positive lower bound on the coefficient of absolute risk aversion, and show how it can be calculated exactly. Indeed, something beyond strict concavity is being assumed.³

After obtaining this lower bound on risk aversion implied by q , the relevant issue is empirical, that is, the question is whether this bound is broadly consistent with the shape of the utility functions supported by empirical evidence. We argue that this bound is often unreasonably high. For instance, the assumption that a person turns down gambles where she loses \$100 or gains \$110 for *any* initial wealth level implies that the coefficient of relative risk aversion must go to infinity when wealth goes to infinity, while the assumption that a 50–50 lose \$100/gain \$105 bet is turned down for any lifetime wealth level less than \$350,000 implies a value of the same coefficient no less than 166.6 at \$350,000. In contrast, a vast body of empirical evidence consistently indicates that the coefficient of relative risk aversion is estimated to be in the single-digit range.

With the single-digit range of empirically plausible coefficient values in hand, we calibrate the relationship between risk attitudes over small-stakes and large-stakes gambles. The important question here, far from obvious a priori, is whether one could still find paradoxes similar to the ones found in [Rabin \(2000\)](#). Our answer is in the negative: we do not find paradoxical rejections of large-stakes gambles. Thus, it follows that paradoxical behavior is only obtained when calibrations are made in a region of the parameter space that is not empirically relevant.

In conclusion, within expected utility, the relevance of the criticism when comparing behavior towards small-stakes and large-stakes gambles relies on the validation of assuming q . We take the debate to the empirical arena, and we argue that much empirical evidence contradicts the plausibility of Rabin's q in a variety of settings. We conclude that expected utility is often a useful model to explain decision making under uncertainty.

2. Rejecting small gambles: evidence

As explained for example in [Segal and Spivak \(1990\)](#), any expected utility agent with a differentiable utility function must accept infinitesimal gambles of positive expected value, because locally these preferences amount to risk neutrality. However, as soon as the gambles are no longer infinitesimal, but of any finite size, expected utility is compatible with both accepting and rejecting small gambles. It is then necessary to perform more powerful tests to examine the question of the appropriateness of expected utility preferences.

There is much work, both experimental and empirical, concerning the behavior of agents towards small-stakes gambles under expected utility. [Camerer \(1995, pp. 638–642\)](#) provides perhaps the most comprehensive survey of experimental evidence on individual decision making, virtually gathered in all cases from what may be considered small-stakes gambles. He finds substantial support for expected utility. See also [Camerer and Ho \(1994\)](#), [Loomes and Segal \(1994\)](#), [Hey and Orme \(1994\)](#) and [Harless and Camerer \(1994\)](#).

As to the range of risk aversion estimates obtained in experimental studies of one-person decision problems and games, the evidence is quite robust. The estimated values of the coefficient of relative risk

³ This is contrary to what is claimed in the literature. For instance, “the calibration theorem is entirely nonparametric, assuming nothing about the utility function except concavity” ([Rabin \(2000, p. 1282\)](#)). However, we show that there are large families of concave functions that are being ruled out by q .

aversion r_R are quite low, certainly within the single-digit range. See, for instance, Holt and Laury (2002), Goeree et al. (2002, 2003).

The empirical evidence on decisions under uncertainty in natural environments is vast, and not easy to summarize. However, the relevant conclusion for the purposes of our analysis is again noticeably robust: values of r_R in the single-digit range appear to rationalize the reactions to risk in virtually all circumstances examined in the literature, including data sets involving small, medium and large gambles. See evidence on small gambles in Cicchetti and Dubin (1994), Evans and Viscusi (1991), Miravete (2001); see also Friedman and Savage (1948). For larger gambles, see for example Metrick (1995), Mehra and Prescott (1985), Kocherlakotta (1996), Luttmer (1996), Barsky et al. (1997), Rosenzweig and Wolpin (1993), Binswanger (1980).

We thus consider that the experimental and empirical evidence conclusively indicates the range of parameter values where it would be empirically relevant to calibrate attitudes towards risk for gambles of different size. We can safely take that range of preferences to be those for which r_R is in the single-digit range.⁴

3. Rejecting small gambles: testable implication

Rabin (2000) shows that if an individual is a risk averse expected utility maximizer and rejects a given gamble of equally likely gain g and loss l , $g > l$, over a given range of wealth levels, then he will reject correspondingly larger gambles of gain G and loss L . We investigate now the implication of the “rejecting small gambles” assumption.

For a decision maker with wealth level w and twice continuously differentiable Bernoulli utility function u , denote the Arrow–Pratt coefficient of absolute risk aversion by $r_A(w, u) = -\frac{u''(w)}{u'(w)}$, with the coefficient of relative risk aversion being $r_R(w, u) = w \cdot r_A(w, u)$. We next show that the assumption that an expected utility maximizer turns down a given gamble with gain g and loss l for a given range of wealth levels implies that there exists a positive lower bound on $r_A(w, u)$. In fact, this positive lower bound can be calculated exactly and, therefore, provides a testable implication of the assumption.

Proposition. Let u satisfy non-increasing absolute risk aversion. Let I be an interval in the positive real line. If for every $w \in I$,

$$\frac{1}{2}u(w + g) + \frac{1}{2}u(w - l) < u(w),$$

there exists $a^* > 0$ such that the absolute risk aversion coefficient $r_A(w, u)$ is greater than a^* for all $w \in I$. Moreover, the highest such a^* is the solution to the equation

$$f(a) = e^{al} + e^{-ag} - 2 = 0.$$

⁴ A final observation concerning the evidence is in order. While the theory usually focuses on life-time wealth, empirical studies typically use per period data (e.g., monthly income, yearly consumption, and so on). Under the standard assumption of time stationarity of preferences, one can easily show that estimates of the preference parameters for the per period utility function correspond to those of the utility function over life-time wealth. Moreover, as Rubinstein (2001) notes, “nothing in the von Neumann-Morgenstern (vNM) axioms dictates use of final wealth levels . . . vNM are silent about the definition of prizes . . . The definition of prizes as final wealth levels is no less crucial to Rabin’s argument than the expected utility assumption.” It is certainly possible that subjects may view lab choice situations and “short-term” problems in isolation, rather than in conjunction with other sources of income. This is Rubinstein’s (2001) response to Rabin (2000). See Cox and Sadiraj (2001) for a related point.

Proof. Suppose not. Then, for every $a > 0$ there exists $w \in I$ such that $r_A(w, u) < a$. In particular, this holds for the unique $a^* > 0$ solving the equation $f(a) = 0$. (By the intermediate value theorem, note that a^* is well defined because f is continuous, $f(0) = 0$, $f(\infty) = \infty$, $f'(0) < 0$, $f'(a) = le^{al} - ge^{-ag}$ is first negative, vanishes at a single point and is positive thereafter).

Consider the constant absolute risk aversion utility function $v(w) = -e^{-aw}$ for $a < a^*$. For such a choice of a , $f(a) < 0$, i.e.,

$$e^{al} + e^{-ag} < 2,$$

or

$$\frac{1}{2} \left(-e^{-a(w-l)} \right) + \frac{1}{2} \left(-e^{-a(w+g)} \right) > -e^{-aw}.$$

Thus, an individual with utility function v would agree to play the small-stakes lottery with gain g and loss l starting from any wealth level w .

Denote by $w' \in I$ the wealth level for which $r_A(w', u) = a$. By non-increasing absolute risk aversion, for $w \in I$, $w \geq w'$, the individual with utility function v is at least as risk averse as the one with utility function u . Therefore, using the well-known characterization of comparisons of risk attitudes across individuals, found for example in Mas-Colell et al. (1995, chapter 6), it follows that

$$\frac{1}{2}u(w+g) + \frac{1}{2}u(w-1) > u(w)$$

for every $w \in I$, $w \geq w'$, which is a contradiction.

The hypothesis of non-increasing absolute risk aversion is generally accepted. Note, however, that it is not essential to the argument in the above proof. We use it only in the last step to assert that the range of wealth levels over which the “rejecting the small-stakes lottery” assumption is violated constitutes an interval of the form $[w', \infty)$, something stronger than we need. In the absence of the non-increasing absolute risk aversion assumption, continuity of the utility function would suffice to obtain the same result over some arbitrary interval. This would also be enough for our purposes.

Hence, contrary to several statements in Rabin (2000, 2001) and Rabin and Thaler (2001), the conclusion to be drawn from this proposition is that the assumption of rejecting the small-stakes gamble does go beyond the assumption of concavity of the Bernoulli utility function. A positive lower bound on risk aversion is also assumed, and this bound is independent of the interval I over which the assumption is made. This lower bound on the coefficient of relative risk aversion clearly increases when for a given small-stakes gamble we enlarge the interval I over which it should be rejected. This means for example that the assumption that a given gamble is rejected for all wealth levels is incompatible with the agent becoming risk neutral at some sufficiently high level of wealth, feature shared by a large class of concave utility functions. The proposition implies that $r_R(w, u)$ must go to infinity as wealth goes to infinity. \square

4. Rejecting small gambles: some new calibrations

In this section we calibrate attitudes towards risk for gambles of different size in the region where the combination of prizes (which we will continue referring to as wealth) and the curvature of the utility function yields empirically plausible parameter values. That is, having turned the assumption of rejecting

Table 1

Wealth levels at which an individual with CRRA (γ) utility function stops rejecting a 50–50 lose \$100/gain \$ g lottery

γ	g	
	125	110
2	400	1000
3	501.3	3300.5
4	2003.1	4401.2
5	2504.9	5502
6	3006.9	6602.7
7	3508.8	7703.5
8	4010.8	8804.3
9	4512.8	9905.1
10	5014.9	11,006
11	5516.9	12,106.8
12	6018.9	13,207.6
20	10,035.4	22,014.2
30	15,056	33,022.5
40	20,076.7	44,030.8
50	25,097.3	55,039.2

small gambles into an empirically testable proposition we examine whether or not rejecting small gambles, when the coefficient $r_R(w, u)$ is in the single-digit range, induces paradoxical behavior. Although the point we are raising is general, only for computational simplicity, we shall work with the class of CRRA (constant relative risk aversion) Bernoulli utility functions $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$ for $\gamma \geq 0$. For this utility function, $r_R(w, u) = \gamma$. In order to facilitate comparison with the literature we next consider calibrations based on gambles similar to Rabin (2000). This is the only reason to use these gambles, in the absence of a clear definition of what it is a “small gamble.” One can certainly argue that there is a great deal of subjectivity in evaluating the magnitude of a gamble.⁵

The assumption of rejecting small-stakes gambles is generally made over a given range of wealth levels. In Table 1 we calculate, for two small-stakes lotteries and for different values of γ , the largest wealth level at which an individual rejects the lotteries.

Note that the values of these wealth levels are extremely small. Therefore, the empirical relevance of the assumption, for decision makers with γ in the single digits, would seem to be quite limited.

Continuing with the specification of CRRA utility, the next question we examine is how high is the bound a^* associated with the given small-stakes lottery of gain g and loss l . On the basis of the same lotteries used in Rabin’s (2000) Tables 1 and 2, we calculate in Table 2 their corresponding values of a^* , as defined in the proposition above. The table also shows, for wealth levels of \$300,000 and \$30,000, the induced values of γ .

It is first worth noting that for the wealth level of \$300,000 very few values of γ are in the single-digit range, or even in the teens. No single-digit value arises when the gamble involves losing \$100 or \$1000. Only when the rejected gamble involves losing $l = \$10,000$, which would not appear to be a small-stakes gamble, such low values start to arise consistently. In an attempt to generate more γ coefficients in the

⁵ If one is to assume the rejection of an initial small gamble, this should hold with independence of how it is presented (e.g., as the choice between telephone calling plans, as a minor health hazard, or as a choice in a first-price auction or a matching pennies experiment). The fact that it is difficult to say what “small” or “large” stakes are suggests a possible role for reference points.

Table 2

Lower bounds on the coefficient of absolute risk aversion for an individual that rejects a 50–50 lottery lose l /gain g (a^*) for any range of wealth levels, and associated lower bound on the coefficient of relative risk aversion for wealth levels \$300,000 and \$30,000

l/g	a^*	$\gamma=300,000a^*$	$\gamma=30,000a^*$
100/101	.0000990	29.7	2.9
100/105	.0004760	142.8	14.2
100/110	.0009084	272.5	27.2
100/125	.0019917	597.5	59.7
100/150	.0032886	986.5	98.6
1000/1050	.0000476	14.2	1.4
1000/1100	.0000908	27.2	2.7
1000/1200	.0001662	49.8	4.9
1000/1500	.0003288	98.6	9.8
1000/2000	.0004812	144.3	14.4
10,000/11,000	.0000090	2.7	0.2
10,000/12,000	.0000166	4.9	0.4
10,000/15,000	.0000328	9.8	0.9
10,000/20,000	.0000481	14.4	1.4

single-digit range, we examine the same lotteries for a wealth level as extremely low as \$30,000. In this case, single-digit coefficients arise for some gambles where l =\$1000, and for all gambles where l =\$10,000, which are hardly small-stake gambles for an individual with that wealth level. For the lowest stakes gambles involving l =\$100 a single-digit γ is only found when g =101. We conclude from Table 2 that empirically plausible, single-digit values of γ are compatible with the assumption only when the loss l in the gamble is a significant proportion of the individual's wealth. We thus learn that the assumption of rejecting truly small gambles does not hold, when applied to all the decision makers that are behind the experimental and empirical evidence mentioned above.

Table 3

If averse to 50–50 lose l /gain g for wealth levels \$300,000 and \$30,000 with CRRA utility and coefficient of relative risk aversion γ , will also turn down a 50–50 lose L /gain G bet; G 's entered in table

Wealth: \$30,000				Wealth: \$300,000		
l/g	100/101	1000/1050	1000/1100	l/g	10,000/11,000	10,000/12,000
γ	2.9	1.4	2.7	γ	2.7	4.9
L				L		
400	416	–	–	15,000	17,341	19,887
600	636	–	–	17,000	20,072	23,572
800	867	–	–	20,000	24,393	29,792
1000	1107	–	–	22,000	27,435	34,490
2000	2479	2205	2439	25,000	32,266	42,574
4000	6538	4917	6259	30,000	41,116	59,870
6000	14,538	8329	13,168	40,000	62,594	126,890
8000	40,489	12,749	30,239	50,000	91,268	∞
10,000	∞	18,686	495,319	75,000	239,089	∞

Finally, for various lotteries in Table 2 that yield values of γ in the single-digit range, Table 3 displays the best large-stakes lottery with gain G and loss L that the decision maker would reject.

It is apparent that these rejections are no longer paradoxical. For instance, for a wealth of \$300,000 the agent turns down gambles involving losses L ranging from 5% to 15% of his wealth and gains G that appear reasonable. The same can be said for a wealth of \$30,000. In this case, note that relative to wealth these values of L are ten times greater than those in Rabin (2000). Thus, not even for these much larger gambles paradoxical behavior is obtained. Finally, it is worth stressing that gambles with $G=8$ are turned down only when potential losses L represent a significantly great proportion of the individual's wealth.

The reasonable behavior described in these large-stakes gambles contrasts with the paradoxes in Rabin (2000) and in other authors in the literature, and indeed may be viewed as a further confirmation of the empirical soundness of single-digit values for $r_R(w, u)$.

These results refute assertions such as “paradoxical implications are not restricted to particular contexts or particular utility functions,” or “within the expected-utility framework, for *any* concave utility function, even very little risk aversion over modest stakes implies an absurd degree of risk aversion over large stakes” (Rabin (2001, p. 203)). That is, much more than “very little risk aversion over modest stakes” is needed to generate paradoxical behavior. Indeed, this is only obtained when calibrations are made in a region of the parameter space that is not empirically relevant.

Lastly, it is important to note that a rather straightforward empirical implication of the calibrations made in Rabin (2000), under his belief that the assumption of rejecting small gambles holds, is that “when measuring risk attitudes maintaining the expected-utility hypothesis . . . data sets dominated by modest-risk investment opportunities are likely to yield much higher estimates of risk aversion than data sets dominated by larger investment opportunities” (Rabin (2000, p. 1287)). But as already discussed in Section 2, contrary to this implication, the empirical evidence gathered from many different studies consistently obtains estimates of $r_R(w, u)$ that vary very little across a wide heterogeneity of scales of risk, as these estimates are narrowly confined to the single-digit range.

5. Concluding remarks

Using a problem posed to one of his colleagues as a starting point, Samuelson (1963) argues that, under expected utility, the rejection of a given single gamble for all wealth levels implies the rejection of the compound lottery consisting of the single gamble being repeated an arbitrary number of times. Samuelson's exercise sheds light on the fact that some decision makers may be misapplying the law of large numbers when accepting a compound lottery (the colleague's response was that he would reject the single lottery, but accept its compound version). However, Samuelson was clearly aware of the crucial importance of the assumption of rejecting the single lottery for *all* wealth levels or a large range thereof: “I should warn against undue extrapolation of my theorem. It does not say that one must always refuse a sequence if one refuses a single venture: if, at higher income levels the single losses become acceptable, and at lower levels the penalty of losses does not become infinite, there might well be a long sequence that it is optimal” (p. 112). Indeed, it may very well be the case that Samuelson's colleague was not fooled by any fallacy of large numbers. He simply violated the assumption of rejecting the given small-stakes lottery for all wealth levels or large range thereof.

The main advantage of expected utility is its simplicity and its usefulness in the analysis of economic problems involving uncertainty. As often argued in the literature, its predictions sometimes conflict with

people's behavior. This has led economists to develop various non-expected utility models which can often accommodate actual behavior better. The non-expected utility research agenda is an important one, and there is no question that we should continue to pursue it. However, expected utility should not be accused when it is not at fault. The analysis in this paper shows how certain paradoxical examples in the literature are many times counterfactuals. Paradoxical behavior is only obtained when calibrations are made in a region of the parameter space that is not empirically relevant. In a more recent paper, [Rabin and Thaler \(2001\)](#) continue to drive home the theme of the demise of expected utility and compare expected utility to a dead parrot from a Monty Python show. To the extent that all their arguments are based on the calibrations in [Rabin \(2000\)](#), the expected utility parrot may well be saying that "the report of my death was an exaggeration."

Acknowledgment

We gratefully acknowledge Salomon research awards from Brown University. Serrano also acknowledges support from NSF grant SES-0133113, the Alfred P. Sloan Foundation and Deutsche Bank. We are indebted to Oscar Volij, who collaborated in the first stages of the project. We thank Robert J. Aumann, Juan Dubra, Jerry Green, Faruk Gul, Peter Hansen, Mark Machina, Casey Mulligan, Matthew Rabin, Amy Salsbury, Lars Stole, Richard Thaler, Rajiv Vohra, Peter Wakker, and participants in the Luncheon Seminar of the School of Social Science at IAS in Princeton, the "Rationality on Fridays" Workshop in Jerusalem, and the 2001 CEPR Conference on Psychology and Economics in Brussels for helpful comments.

References

- Arrow, K.J., 1971. *Essays in the Theory of Risk Bearing*. Markham Publishing Company, Chicago, IL.
- Barsky, R.B., Juster, F.J., Kimball, M.S., Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement study. *Quarterly Journal of Economics* 112, 537–579.
- Binswanger, H.P., 1980. Attitude toward risk: experimental measurement in rural India. *American Journal of Agricultural Economics* 62, 395–407.
- Camerer, C., 1995. Individual decision making. In: Kagel, J., Roth, A. (Eds.), *Handbook of Experimental Economics*. Princeton University Press.
- Camerer, C., Ho, T.-H., 1994. Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty* 8, 167–196.
- Cicchetti, C.J., Dubin, J.A., 1994. A microeconomic analysis of risk aversion and the decision to self-insure. *Journal of Political Economy* 102, 169–186.
- Cox, J.C., Sadiraj, V., 2001. Risk aversion and expected-utility theory: coherence for small-and large-stakes gambles. Department of Economics University of Arizona, Working Paper, 01–03.
- Epstein, L.G., 1992. Behavior under risk: recent developments in theory and applications. In: Laffont, J.J. (Ed.), *Advances in Economic Theory, VI World Congress of the Econometric Society*, vol. II. Cambridge University Press.
- Evans, W.N., Viscusi, W.K., 1991. Estimation of state-dependent utility functions using survey data. *Review of Economics and Statistics* 73, 94–104.
- Friedman, M., Savage, L.J., 1948. The utility analysis of choices involving risk. *Journal of Political Economy* 56, 270–304.
- Goeree, J.K., Holt, C.A., Pfaffrey, T.R., 2002. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* 104, 247–272.
- Goeree, J.K., Holt, C.A., Pfaffrey, T.R., 2003. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior* 45, 97–113.

- Hansson, B., 1988. Risk aversion as a problem of conjoint measurement. In: Gardenfors, P., Sahlin, N.-E. (Eds.), *In Decision, Probability, and Utility*. Cambridge University Press, Cambridge.
- Harless, D.W., Camerer, C.F., 1994. The predictive utility of generalized expected utility theories. *Econometrica* 62, 1251–1289.
- Hey, J.D., Orme, C., 1994. Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62, 1291–1326.
- Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. *American Economic Review* 92, 1644–1655.
- Kandel, S., Stambaugh, R.F., 1991. Asset returns, investment horizon, and intertemporal preferences. *Journal of Monetary Economics* 27, 39–71.
- Kocherlakotta, N.R., 1996. The equity premium: it's still a puzzle. *Journal of Economic Literature* 34, 42–71.
- Loomes, G., Segal, U., 1994. Observing different levels of risk aversion. *Journal of Risk and Uncertainty* 9, 239–256.
- Luttmer, E.G.J., 1996. Asset pricing in economies with frictions. *Econometrica* 64, 1439–1467.
- Machina, M., 1982. Expected utility analysis without the independence axiom. *Econometrica* 50, 277–323.
- Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. *Microeconomic Theory*. Oxford University Press.
- Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. *Journal of Monetary Economics* 15, 145–161.
- Metrick, A., 1995. A natural experiment in “Jeopardy!”. *American Economic Review* 85, 240–253.
- Miravete, E.J., 2001. Choosing the wrong calling plan? Ignorance, Learning, and Risk Aversion. Mimeo, Dept. of Economics. University of Pennsylvania.
- Pratt, J., 1964. Risk aversion in the small and in the large. *Econometrica* 32, 122–136.
- Rabin, M., 2000. Risk aversion and expected-utility: a calibration theorem. *Econometrica* 68, 1281–1292.
- Rabin, M., 2001. Diminishing marginal utility of wealth cannot explain risk aversion. In: Kahneman, D., Tversky, A. (Eds.), *In Choices, Values, and Frames*. Cambridge University Press.
- Rabin, M., Thaler, R.H., 2001. Anomalies. Risk aversion. *Journal of Economic Perspectives* 15, 219–232.
- Rosenzweig, M.R., Wolpin, K.I., 1993. Credit market constraints, consumption smoothing, and the accumulation of durable assets in low-income countries: investments in bullocks in India. *Journal of Political Economy* 101, 223–244.
- Rubinstein, A., 2001. Comments on the Risk and Time Preferences in Economics. Mimeo Tel-Aviv University.
- Samuelson, P.A., 1963. Risk and uncertainty: a fallacy of large numbers. *Scientia* 98, 108–113.
- Segal, U., Spivak, A., 1990. First order versus second order risk aversion. *Journal of Economic Theory* 51, 111–125.
- Watt, R., 2002. Defending expected utility theory. *Journal of Economic Perspectives* 16, 227–229.