''Third down with a yard to go'': recursive expected utility and the Dixit–Skeath conundrum

Simon Grant[a,b,*], Atsushi Kajii[c], Ben Polak[d]

[a]*Department of Econometrics and Operations Research, Tilburg University, Tilburg, The Netherlands*
[b]*School of Economics, Faculty of Economics and Commerce, Australian National University, Canberra, ACT 0200, Australia*
[c]*Institute of Policy and Planning Sciences, University of Tsukuba, Tsukuba, Japan*
[d]*Economics Department, Yale University, New Haven, CT, USA*

Abstract

In two-outcome strictly competitive games, equilibrium mixed strategies do not depend on ultimate prizes. Dixit and Skeath [Games of Strategy (1999) Norton, New York] find this 'counter-intuitive'. We show this invariance comes from reduction, not independence; and provide conditions for 'intuitive' comparative statics under recursive expected utility. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Recursive expected utility; Strictly competitive games; Comparative statics; Non-expected utility

JEL classification: C72; D81

1. Introduction

This paper illustrates a new use of the recursive expected utility model applied to game theory.[1] The application allows us to demonstrate some comparative static properties of the model. We derive analogs of the Arrow–Pratt coefficients of absolute and relative risk aversion and show, in the context of an example, what restrictions on these coefficients are sufficient to yield desired comparative static results. The example we use to illustrate the model, is drawn from Dixit and Skeath's *Games of Strategy* (1999).

*Corresponding author. Tel.: +61-2-6125-4602; fax: +61-2-6125-0383.

E-mail address: simon.grant@anu.edu.au (S. Grant).

[1]There is a growing literature applying non-expected utility models to game theory. But as far as we know, the only paper that applies any form of the recursive expected utility model is Dekel et al. (1991).

Dixit and Skeath (1999), pp. 219–223) pose a 'conundrum'. They consider the game between the offense and the defense in an American football game. It is third down and the offense has a yard to go. Suppose the offense can only choose to run or to pass. Similarly, the defense can only chose optimally to defend the run or the pass. The probability that the offense goes on to win the game depends on whether or not its action is matched by the defense. Run and pass, however, are not equivalent. If the offense attempts a pass against a run defense, then the offense wins the game with very high probability, but if meets a pass defense, offense loses with very high probability. On the other hand, if the offense runs against a pass defense the probability of the offense winning the game is not very much higher than if it meets a run defense. As Dixit and Skeath put it, run is the 'percentage' or safe play for the offense, whereas pass is the 'risky' alternative. The equilibrium of this game involves mixed strategies.

Dixit and Skeath observe:

> " . . . people often say that if the occasion is really important, in the sense that winning versus losing is a big difference in payoffs, then one should use the percentage play more often. Thus they argue that the offense may throw a long pass on third down with a yard to go in an ordinary season game, but the risk of this play is so large that it should not be used in the superbowl."

Contrary to this intuition, however, the equilibrium mixing probabilities between the two plays are completely independent of the ultimate prizes.

> "The theory says that you should mix the percentage play and the risky play in exactly the same proportions on a big occasion as you would on a minor occasion. So which is right: theory or intuition?"

Dixit and Skeath suggest that perhaps it is the theory that is too rigid, and specifically that the problem arises from the use of expected utility in constructing the payoffs. They call for attempts to provide game and strategy with alternative foundations to bring the theory and intuition back into line: "here is an interesting research opportunity".

We show that Dixit and Skeath's conundrum can be resolved, in a natural way, using the recursive expected utility model. Contrary to Dixit and Skeath's claim, it is not "linearity in the probabilities" (that is, the independence axiom of expected utility) that is the problem. Any non-expected utility model that satisfies both the reduction of compound lotteries axiom and first-order stochastic dominance will produce the invariance result that Dixit and Skeath find unintuitive: that is, the equilibrium mixing probabilities in two-outcome strictly competitive games will still be invariant to changes in the ultimate prizes. On the other hand, if we retain the independence axiom but instead relax the reduction of compound lotteries axiom (that is, if we use Kreps and Porteus's (1978) recursive expected utility model), then this invariance no longer holds.

The equilibria in games such as Dixit and Skeath's football example involve two sets of probabilities: those representing each agent's beliefs about the other's actions, and those coming from the residual uncertainty of the game. The reduction axiom requires agents to treat these two sources of uncertainty in the same way. Agents are assumed simply 'to multiply through' the probabilities to form one-stage lotteries on the final prizes. Dixit and Skeath's intuition is that people do not conflate multiple sources of uncertainties in this way. For the offense, the pass play seems 'risky' in the sense that most of the total uncertainty is resolved in the first stage, hence the 'utilities' of the second-stage

uncertainty (that part, left to nature) are 'spread out'. The running play seems 'safe' in that most of the total uncertainty is resolved after the first stage, hence the 'utilities' of the second-stage uncertainty are 'close together'. The recursive expected utility model, by dropping reduction, allows the agent to distinguish the immediate uncertainty concerning the choices of her opponent from the later uncertainty involving the whims of nature.

In the following, we show what specific form of the recursive expected utility model give the comparative statics that Dixit and Skeath want; namely, that players are inclined to play more cautiously (that is, run) in the superbowl than in the regular season.[2] We believe that the comparative static techniques illustrated here, in particular the analogies to Arrow–Pratt coefficients, can be applied more generally.

Section 2 formally sets up the problem. Section 3 discusses dropping independence but retaining reduction. Section 4 does the opposite. Section 5 shows how to generate the Dixit Skeath comparative statics using recursive expected utility, and provides an example.

2. The setting

Consider a strictly competitive contest between two players, row and column. There are two ultimate outcomes, row wins (and column loses) or row loses (and column wins). If row wins, he gets $\bar{x}$; and if he loses, he gets $\underline{x}$, where $\bar{x} > \underline{x}$. Similarly, if column wins, she gets $\bar{y}$; and if she loses, she gets $\underline{y}$, where $\bar{y} > \underline{y}$. Which outcome occurs depends not only on the players' actions but also on chance. The players' choices are simultaneous. Row chooses between S or R and column chooses between s or r. Later, we will identify S with Dixit and Skeath's 'safe' (or 'percentage') strategy, and identify R with their 'risky' strategy. Each action-pair results in a lottery. For example, given the action-pair (R,s), let π_{Rs} be the probability that row wins $\bar{x}$ (in which case column gets $\underline{y}$); and let $1 - \pi_{Rs}$ be the probability that column wins $\bar{y}$ (in which case row gets $\underline{x}$). Writing $[\pi;\bar{x},\underline{x}]$ for the lottery which yields $\bar{x}$ with probability π and $\underline{x}$ with probability $1 - \pi$, this contest can be summarized in the matrix below.

	s	r
S	$[\pi_{Ss};\bar{x},\underline{x}],[1 - \pi_{Ss};\bar{y},\underline{y}]$	$[\pi_{Sr};\bar{x},\underline{x}],[1 - \pi_{Sr};\bar{y},\underline{y}]$
R	$[\pi_{Rs};\bar{x},\underline{x}],[1 - \pi_{Rs};\bar{y},\underline{y}]$	$[\pi_{Rr};\bar{x},\underline{x}],[1 - \pi_{Rr};\bar{y},\underline{y}]$

We are interested in mixed-strategy equilibria so we assume: $\pi_{Rs} > \pi_{Ss}$, $\pi_{Rr} < \pi_{Sr}$, $1 - \pi_{Rs} < 1 - \pi_{Rr}$ and $1 - \pi_{Ss} > 1 - \pi_{Sr}$. We use the terminology of a mixed-strategy equilibrium being an

---

[2]In fact, in our Yale classroom experiments we also found that invariance failed — there was too much switching — but (unlike Dixit's Princeton classroom experiments) we found that many people switched the other way. By reversing the inequalities in Proposition 3, this behavior can also be incorporated in the recursive expected utility model.

equilibrium in beliefs.[3] Let q^* denote row's equilibrium belief that column will play s, and let p^* denote column's equilibrium belief that row chooses S.

As Dixit and Skeath note, under the standard assumption of expected utility, the equilibrium mixed-strategy profile in such two-outcome strictly competitive games is invariant to the sizes of the final prizes, $\bar{x}$, $\underline{x}$, $\bar{y}$ and $\underline{y}$. In particular, row's equilibrium belief q^* is given by

$$q^*\pi_{Ss}[u(\bar{x}) - u(\underline{x})] + (1 - q^*)\pi_{Sr}[u(\bar{x}) - u(\underline{x})] = q^*\pi_{Rs}[u(\bar{x}) - u(\underline{x})] + (1 - q^*)\pi_{Rr}[u(\bar{x}) - u(\underline{x})],$$
$$(1)$$

where u denotes row's von Neumann–Morgenstern utility index. The q^* that solves this equation clearly does not depend on $\bar{x}$ or $\underline{x}$. The equilibrium belief q^* is simply that which makes the total probability of 'winning' from S equal to that from R, regardless of exactly what it is you win or lose. Similarly, p^* does not depend on $\bar{y}$ or $\underline{y}$.[4]

Definition. We say that the invariance property of mixed equilibria in two-outcome strictly competitive games is satisfied if, for all such games with $\pi_{Rs} > \pi_{Ss}$, $\pi_{Rr} < \pi_{Sr}$, $1 - \pi_{Rs} < 1 - \pi_{Rr}$ and $1 - \pi_{Ss} > 1 - \pi_{Sr}$, the equilibrium beliefs (p^*, q^*) do not depend on the ultimate prizes, $\bar{x}$, $\underline{x}$, $\bar{y}$, and $\underline{y}$.

We are interested in whether this invariance extends beyond the standard expected-utility model. For our purposes, there are two key assumptions in the standard model. First, is the reduction of compound lotteries axiom (hereafter, 'reduction'). Without reduction, we can model agents as having preferences over two-stage lotteries, the first stage representing each player's uncertainty about the other's action, and the second the 'π-lottery' over outcomes. The second assumption is the independence axiom. Recall that Dixit and Skeath suggest that resolving this axiom will resolve the conundrum. We refer to models that relax independence but maintain reduction as atemporal nonexpected-utility models. We refer to models that relax reduction but maintain independence as recursive expected-utility models.

3. Atemporal non-expected utility

Let $\mathscr{L}$ be the set of one-stage (real-valued) lotteries. Let the function $V:\mathscr{L} \to \mathbb{R}$ represent a general (complete, transitive and continuous) preference relation over $\mathscr{L}$. Under atemporal nonexpected utility, row's equilibrium belief q^* is given by

$$V([q^*\pi_{Ss} + (1 - q^*)\pi_{Sr}; \bar{x}, \underline{x}]) = V([q^*\pi_{Rs} + (1 - q^*)\pi_{Rr}; \bar{x}, \underline{x}]). \tag{2}$$

That is, given q^*, the (reduced) lottery over outcomes induced by S is indifferent to that induced by R. Just as with expected utility, the equilibrium belief q^* is simply that which makes the total probability

---

[3]We adopt the 'belief' rather than the 'randomizing' terminology simply because we find it helps our intuitions, but nothing in our formal results depends on this. Since there are only two ultimate outcomes, there is no existence problem when we relax independence.

[4]For these equilibrium beliefs to be well-defined, we are assuming that $u(\bar{x}) - u(\underline{x})$ is finite. We retain this or the equivalent 'boundedness' assumptions below.

of 'winning' from S equal to that from R. Thus, as long as winning is strictly better than losing, this equilibrium q^* is invariant to the exact value of the prizes. Strictly speaking, this argument requires that V be strictly increasing in the probability of winning, but most widely used atemporal nonexpected-utility models retain such monotonicity: they assume that preferences respect first-order stochastic dominance. This discussion is summarized in the following proposition.

Proposition 1. *Provided the players' preferences respect first-order stochastic dominance, the invariance property of mixed equilibria in two-outcome strictly competitive games is maintained if we relax independence but retain reduction.*

4. Recursive expected utility

Recall that each action by a player results in that player facing a two-stage lottery. For example, if row believes column is choosing s with probability q, then row's choosing S yields the two-stage lottery where the second-stage lottery $[\pi_{Ss};\bar{x},\underline{x}]$ occurs with first-stage probability q, and the second-stage lottery $[\pi_{Sr};\bar{x},\underline{x}]$ with first-stage probability $1 - q$. Let us write such two-stage (real-valued) lotteries in the form $X = [q;[\pi_{Ss};\bar{x},\underline{x}],[\pi_{Sr};\bar{x},\underline{x}]]$, and let L^2 be the set of such two-stage lotteries. Let the function $W:L^2 \to \mathbb{R}$ represent a complete, transitive and smooth preference relation over L^2.

The function W satisfies recursive expected utility if it can be written in the following form:

$$W([q;[\pi;\bar{x},\underline{x}],[\pi';\bar{x},\underline{x}]]) = \Phi([q;U([\pi;\bar{x},\underline{x}]),U([\pi';\bar{x},\underline{x}])]), \tag{3}$$

where both $\Phi:\mathscr{L} \to \mathbb{R}$ and $U:\mathscr{L} \to \mathbb{R}$ are expected utility functions representing preference relations over one-stage lotteries. To see why this is called recursive expected utility, notice that $U([\pi;\bar{x},\underline{x}])$ and $U([\pi';\bar{x},\underline{x}])$ are the evaluations of the two second-stage lotteries using the expected utility function U. Thus, $[q;U([\pi;\bar{x},\underline{x}]),U([\pi';\bar{x},\underline{x}])]$ is the one-stage lottery that assigns probability q to getting the second-stage utility $U([\pi;\bar{x},\underline{x}])$, and probability $(1 - q)$ to getting the second-stage utility $U([\pi';\bar{x},\underline{x}])$. The right side of expression (3) is the evaluation of this one-stage lottery using the expected utility function Φ. Notice that the functions Φ and U are not necessarily the same: Φ measures attitudes towards the uncertainty about the opponent's actions; U measures attitudes towards the residual uncertainty of nature.

For heuristic purposes, it is useful to rewrite the recursive expected utility function W in terms of two von Neuman–Morgenstern utility indices, v and u, the former used when evaluating the expected utility of first-stage lotteries and the latter used when evaluating the expected utility of second-stage lotteries. We can then rewrite expression (3) as:

$$W([q;[\pi;\bar{x},\underline{x}],[\pi';\bar{x},\underline{x}]]) =$$
$$qv \circ u^{-1}(u(\underline{x}) + \pi(u(\bar{x}) - u(\underline{x}))) + (1 - q)v \circ u^{-1}(u(\underline{x}) + \pi'(u(\bar{x}) - u(\underline{x}))), \tag{4}$$

To help understand this form, notice that $u(\underline{x}) + \pi(u(\bar{x}) - u(\underline{x}))$ is the expected u-utility of the second-stage lottery $[\pi;\bar{x},\underline{x}]$. Thus, $u^{-1}(u(\underline{x}) + \pi(u(\bar{x}) - u(\underline{x})))$ is just the certainty equivalent of this second-stage lottery according to the expected u-utility preferences. Similarly, $u^{-1}(u(\underline{x}) + \pi'(u(\bar{x}) - u(\underline{x})))$ is the corresponding certainly equivalent of the second-stage lottery $[\pi';\bar{x},\underline{x}]$. Thus, the function

W is evaluated by taking the expected v-utility of the first-stage lottery that assigns probability q to the (second-stage) certainty equivalent of $[\pi;\bar{x},\underline{x}]$ and $1 - q$ to the (second-stage) certainty equivalent of $[\pi';\bar{x},\underline{x}]$.

We can now give an expression for row's equilibrium belief q^* assuming that the agent has recursive expected utility preferences. Let u in expression (4) be the utility index corresponding to the function U in expression (3). In this case, $\varphi := v \circ u^{-1}$ is the utility index corresponding the function Φ. Row's equilibrium belief q^* is then given by

$$
\begin{aligned}
&q^*\varphi(u(\underline{x}) + \pi_{Ss}(u(\bar{x}) - u(\underline{x}))) + (1 - q^*)\varphi(u(\underline{x}) + \pi_{Sr}(u(\bar{x}) - u(\underline{x}))) = \\
&q^*\varphi(u(\underline{x}) + \pi_{Rs}(u(\bar{x}) - u(\underline{x}))) + (1 - q^*)\varphi(u(\underline{x}) + \pi_{Rr}(u(\bar{x}) - u(\underline{x}))).
\end{aligned}
\tag{5}
$$

That is, the equilibrium q^* is such that the two-stage lottery induced by S is indifferent to that induced by R.

Unlike the standard expected utility or atemporal non-expected utility cases (Eqs. (1) and (2)), the q^* that solves Eq. (5) is not necessarily that which makes the total probability of 'winning' from S equal to that from R. In the recursive expected utility case, the first-stage q-probabilities do not multiply the second-stage π-probabilities directly. Instead, they multiply the function φ evaluated at the utilities (the 'U's') of the second-stage π-lotteries. Suppose that this φ function is non-linear; for example, if the utility index v is a concave transformation of the utility index u, then φ is concave. Concavity of φ is like risk aversion in a utility index, except that risk here is not over outcomes but over second-stage utilities. Thus, if φ is concave, all other things being equal, the agent would prefer the U's to be less spread out. For fixed $\bar{x}$ and $\underline{x}$, this is equivalent to saying that the agent would prefer the π's to be less spread out since (since $U([\pi;\bar{x},\underline{x}]) = u(\underline{x}) + \pi(u(\bar{x}) - u(\underline{x}))$). That is, all other things being equal, she would prefer the 'safe' to the 'risky' strategy. If the agent is to be indifferent between her two strategies, the equilibrium q^* must compensate her for taking the 'risky' strategy.

Suppose now that — as in the Dixit–Skeath story — the prizes, $\bar{x}$ and $\underline{x}$, are no longer fixed. Recall that, in the standard model, an agent's attitude toward risky outcomes depends on both the level of the base outcome and the magnitude of the spread of outcomes. Similarly, in the recursive expected utility model, an agent's attitude toward risky second-stage utilities depends on both her base level of second-stage utility and the magnitude of the spread of these utilities. Thus the risk premium required in the equilibrium q^*, will not in general be invariant to the prizes, $\bar{x}$ and $\underline{x}$.

The following proposition makes this intuition more precise. For recursive expected utility, the invariance property only holds if φ is linear; that is, if the first-stage utility index v is a positive affine transformation of the second-stage utility index u. This condition is equivalent to imposing the reduction of compound lotteries axiom. Proposition 1 showed that, regardless of whether or not preferences satisfy independence (given first-order stochastic dominance) reduction is sufficient for the invariance property. Proposition 2 shows that, given independence, reduction is necessary.

Proposition 2. *The invariance property of mixed equilibria in two-outcome strictly competitive games is not maintained if we relax reduction but retain independence. Indeed, for recursive expected utility, the invariance property implies reduction.*

Proof. See Appendix A.

5. A resolution of the conundrum

It remains to show how the equilibrium beliefs change in response to changes in the prizes. Consider a setting similar to that of Dixit and Skeath, in which the action S results in almost the same lottery for row regardless of what action column chooses. The action R on the other hand results in a very good lottery for row if column chooses s and a very bad lottery for row if column chooses r. In particular, suppose that $\pi_{Rs} > \pi_{Sr} > \pi_{Ss} > \pi_{Rr}$. In this case, it is as if S is a 'safer' action for row than is R.

Following Dixit and Skeath, suppose we raise the stakes by increasing the value of winning, $\bar{x}$, or decreasing the value of losing, $\underline{x}$, or some combination of the two. For example, compare a regular season game with the superbowl. Introspection might suggest that, in the superbowl, row might be less willing to incur risk and hence more inclined to choose S. Therefore to keep row indifferent, his equilibrium belief q^* that column will play s will have to increase.

The following proposition gives conditions on the composite function $\varphi = v \circ u^{-1}$ associated with recursive expected utility that yield this comparative static.

Proposition 3. *Suppose that $\pi_{Rs} > \pi_{Sr} > \pi_{Ss} > \pi_{Rr}$; and that the row player's preferences over two-stage lotteries satisfy recursive expected utility with associated composite function $\varphi: \mathbb{R}_+ \to \mathbb{R}_+$ with $\varphi(0) = 0$. Then a sufficient condition for row's equilibrium belief q^* that column will play s to increase as $\underline{x}$ decreases and/or $\bar{x}$ increases is that φ is concave, $-\varphi''(w)/\varphi'(w)$ is decreasing and $-w\varphi''(w)/\varphi'(w)$ is increasing.*

Proof. See Appendix A.

One way to view this result is that, if row holds the same beliefs about column's actions in the superbowl as he does in the regular season, then the safe alternative now looks more attractive. If commentators hold this naïve belief, then they would predict that row is more likely to play S. Of course, it does not follow that, in equilibrium, row will in fact play S more often. If we interpret mixed strategies as randomizations then, as usual, row's equilibrium mix is that which makes column indifferent. Under this interpretation, the result above says that column is more likely to defend the safe alternative in the superbowl than in the regular season. It is tempting to think that, for this same ordering of the π's, a similar comparative static result applies to column's beliefs (and hence to row's randomization). Since $1 - \pi_{Rr} > 1 - \pi_{Ss} > 1 - \pi_{Sr} > 1 - \pi_{Rs}$, however, it is *not* the case that s is a 'safer' action for column than is r in Dixit and Skeath's contest. In fact, it can be shown that the change in p^* for the analogous changes in $\bar{y}$ and $\underline{y}$ cannot be signed for the general case.[5]

Proposition 3 resembles standard results in comparative statics for expected-utility theory originally due to Pratt (1964).[6] To require $-\varphi''(w)/\varphi'(w)$ to be decreasing is analogous to assuming decreasing absolute risk aversion. Similarly, to require $-w\varphi''(w)/\varphi'(w)$ to be increasing is analogous to assuming increasing relative risk aversion. This analogy should not, however, be taken literally. The function φ is not a utility index. Loosely speaking, rather than mapping outcomes to utils, it maps

---

[5]The problem is that the difference in the relevant 'mean values', $\bar{\pi}(\pi_{Rr}, \pi_{Rs}) - \bar{\pi}(\pi_{Ss}, \pi_{Sr})$, cannot be uniquely signed.
[6]For other examples of comparative statics under expected utility, see Gollier (1995).

utils to utils. Concavity of φ does not imply risk aversion in the usual sense, but rather 'risk-aversion' in second-stage utilities. In our setting, risk-aversion in second-stage utilities means that, all other things being equal, the agent dislikes the uncertainty being resolved by actions in the first stage rather than by nature in the second stage.[7]

Pushing the analogy further, decreasing absolute (respectively, increasing relative) risk aversion is analogous to decreasing absolute (respectively, increasing relative) aversion to first-stage resolution of uncertainty. In our setting, as the outcome from losing gets worse or/and the outcome from winning gets better, the agent becomes more averse to making almost everything determined by the players' actions rather than by nature, and so becomes less inclined to pass. Other comparative statics can be obtained by similar analogies.[8]

To make clear this is not an issue of risk aversion per se, the following example satisfies the conditions of Proposition 3, but does not require either of the underlying utility indices, u or v, to be risk averse.

Example. Let $u(x) = (x + 1)^a - 1$ and $v(x) = (x + 1)^b - 1$, where $a > b > 0$. Then $\varphi := v \circ u^{-1}$ is given by $\varphi(w) = (w + 1)^{b/a} - 1$, which is concave and has $\varphi(0) = 0$. Moreover,

$$\frac{-\varphi''(w)}{\varphi'(w)} = \left(\frac{a - b}{a}\right)\frac{1}{w + 1}$$

and

$$\frac{-w\varphi''(w)}{\varphi'(w)} = \left(\frac{a - b}{a}\right)\frac{w}{w + 1}.$$

So the former is decreasing in w, and the latter is increasing in w, as required. But, if a and b are greater than one, then the underlying utility indices are risk loving.

Appendix A

Proof of Proposition 2

Let $\underline{u} := u(\underline{x})$, and let $\Delta u := u(\bar{x}) - u(\underline{x})$. It is enough to show that q^* cannot be invariant both to changes in $\underline{u}$ (holding Δu fixed) and to changes in Δu (holding $\underline{u}$ fixed) unless the preferences satisfy reduction. For q^* to be invariant to changes in $\underline{u}$ (holding Δu fixed), differentiating Eq. (5) with respect to $\underline{u}$, we require

$$q^*\varphi'(\underline{u} + \pi_{S_s}\Delta u) + (1 - q^*)\varphi'(\underline{u} + \pi_{S_r}\Delta u)$$
$$= q^*\varphi'(\underline{u} + \pi_{R_s}\Delta u) + (1 - q^*)\varphi'(\underline{u} + \pi_{R_r}\Delta u).$$

---

[7]More generally, Grant et al. (1998) show that the curvature of φ captures the agent's attitudes towards the timing of resolution of uncertainty.

[8]For example, sufficient conditions for q^* to decrease as the value of both prizes, $\underline{x}$ and $\bar{x}$ are increased by the same amount as those of Proposition 3 plus risk aversion with respect to lotteries that are degenerate in the first stage.

Evaluating this expression at $\underline{u} = 0$ and $\Delta u = 1$, rearranging and substituting for q^* from expression (5), we get

$$\text{sign}\left[\frac{\mathrm{d}q^*}{\mathrm{d}\underline{u}}\right] = -\text{sign}[q^*(\varphi'(\pi_{Rs}) - \varphi'(\pi_{Ss})) - (1 - q^*)(\varphi'(\pi_{Sr}) - \varphi'(\pi_{Rr}))]$$

$$= \text{sign}\left[\frac{(\varphi'(\pi_{Sr}) - \varphi'(\pi_{Rr}))}{(\varphi(\pi_{Sr}) - \varphi(\pi_{Rr}))} - \frac{(\varphi'(\pi_{Rs}) - \varphi'(\pi_{Ss}))}{(\varphi(\pi_{Rs}) - \varphi(\pi_{Ss}))}\right]. \tag{A.1}$$

So, for $\mathrm{d}q^*/\mathrm{d}\underline{u}$ to be zero, we require the difference in the last square bracket to be zero for all π's such that a mixed-strategy equilibrium applies. This in turn implies that

$$\frac{(\varphi'(\pi_1) - \varphi'(\pi_2))}{(\varphi(\pi_1) - \varphi(\pi_2))} = \alpha$$

for all $\pi_1 > \pi_2$, where α is a constant. Allowing $\pi_2 = \pi_1 + \Delta\pi$, and taking $\Delta\pi$ arbitrarily small, by l'Hôpital's rule, we obtain

$$\frac{\varphi''(\pi_1)}{\varphi'(\pi_1)} = \alpha$$

for all π_1 (where $\varphi'(\pi_1) > 0$). It is as if the composite function φ must exhibit 'constant absolute risk-aversion' (although, we do not require the function to be concave).

Similarly, for q^* to be invariant to changes in Δu (holding $\underline{u}$ fixed) we require

$$q^*\pi_{Ss}\varphi'(\underline{u} + \pi_{Ss}\Delta u) + (1 - q^*)\pi_{Sr}\varphi'(\underline{u} + \pi_{Sr}\Delta u) =$$

$$q^*\pi_{Rs}\varphi'(\underline{u} + \pi_{Rs}\Delta u) + (1 - q^*)\pi_{Rr}\varphi'(\underline{u} + \pi_{Rr}\Delta u).$$

Again, evaluating this expression at $\underline{u} = 0$ and $\Delta u = 1$, rearranging and substituting for q^* from expression (5), we get

$$\text{sign}\left[\frac{\mathrm{d}q^*}{\mathrm{d}(\Delta u)}\right] = \text{sign}\left[\frac{(\pi_{Sr}\varphi'(\pi_{Sr}) - \pi_{Rr}\varphi'(\pi_{Rr}))}{(\varphi(\pi_{Sr}) - \varphi(\pi_{Rr}))} - \frac{(\pi_{Rs}\varphi'(\pi_{Rs}) - \pi_{Ss}\varphi'(\pi_{Ss}))}{(\varphi(\pi_{Rs}) - \varphi(\pi_{Ss}))}\right]. \tag{A.2}$$

So, for $\mathrm{d}q^*/\mathrm{d}(\Delta u)$ to be zero, we require the difference in the last square bracket to be zero for all π's such that a mixed-strategy equilibrium applies. This in turn implies that

$$\frac{(\pi_1\varphi'(\pi_1) - \pi_2\varphi'(\pi_2))}{(\varphi(\pi_1) - \varphi(\pi_2))} = \rho$$

for all $\pi_1 > \pi_2$, where ρ is a constant. Again, allowing $\pi_2 = \pi_1 + \Delta\pi$, and applying l'Hôpital's rule, we obtain

$$\frac{\pi_1\varphi''(\pi_1)}{\varphi'(\pi_1)} = \rho - 1$$

for all π_1. It is as if the composite function φ must also exhibit 'constant relative risk-aversion'.

It is well-known, however, that the only functions that satisfy constant relative and constant

absolute risk aversion are linear. By the definition of φ, this implies that the utility index v is an affine transformation of the utility index u. In this case, we can write $v(x) = a + bu(x)$ (where $b > 0$) and

$$W([q;[\pi;\bar{x},\underline{x}],[\pi';\bar{x},\underline{x}]])$$

$$= q[a + bu \circ u^{-1}(U([\pi;\bar{x},\underline{x}]))] + (1 - q)[a + bu \circ u^{-1}(U([\pi';\bar{x},\underline{x}]))]$$

$$= a + b(qU([\pi;\bar{x},\underline{x}]) + (1 - q)U([\pi';\bar{x},\underline{x}]))$$

$$= a + bU([q\pi + (1 - q)\pi';\bar{x},\underline{x}]).$$

But $[q\pi + (1 - q)\pi';\bar{x},\underline{x}]$ is simply the one-stage lottery that is the reduction of the two-stage lottery $[q;[\pi;\bar{x},\underline{x}],[\pi';\bar{x},\underline{x}]]$. That is, if the invariance property of mixed equilibria in two-outcome strictly competitive games applies under the recursive expected-utility preferences W then the preferences also satisfy reduction.

Proof of Proposition 3

Set $\underline{u}: = u(\underline{x})$ and $\Delta u: = u(\bar{x}) - u(\underline{x})$. As $\underline{x}$ decreases, $\underline{u}$ decreases. As $\underline{x}$ decreases and/or $\bar{x}$ increases, Δu increases. So, it is enough to show that q^* is decreasing in $\underline{u}$ (that is, expression (6) is negative), and increasing in Δu (that is, expression (A.2) is positive).

We deal first with changes in $\underline{u}$. Set $\pi_1 > \pi_2$. Let $w_i: = \underline{u} + \pi_i \Delta u$ and let $\zeta_i: = \varphi(w_i)$ for $i = 1,2$. Then, by the mean value theorem,

$$\varphi' \circ \varphi^{-1}(\varphi(w_1)) - \varphi' \circ \varphi^{-1}(\varphi(w_2))$$

$$= \varphi' \circ \varphi^{-1}(\zeta_1) - \varphi' \circ \varphi^{-1}(\zeta_2)$$

$$= (\zeta_1 - \zeta_2)\frac{\mathrm{d}}{\mathrm{d}\zeta}\varphi' \circ \varphi^{-1}(\bar{\zeta}(\zeta_1,\zeta_2))$$

for some $\bar{\zeta}(\zeta_1,\zeta_2)$ in the interval (ζ_2,ζ_1). Applying the inverse function rule, we get

$$\varphi'(w_1) - \varphi'(w_2) = (\varphi(w_1) - \varphi(w_2))\frac{\varphi''[\bar{w}(w_1,w_2)]}{\varphi'[\bar{w}(w_1,w_2)]}$$

for some $\bar{w}(w_1,w_2)$ in the interval (w_2,w_1) (since $\varphi'(w) > 0$).

Let $A(w): = -\varphi''(w)/\varphi'(w)$. We next show that the mean value is increasing in its arguments. By the implicit function theorem,

$$\frac{\partial}{\partial w_1}\bar{w}(w_1,w_2) = \frac{[A(w_1) - A(\bar{w}(w_1,w_2))]\varphi'(w_1)}{A'(\bar{w}(w_1,w_2))[\varphi(w_1) - \varphi(w_2)]},$$

which is greater than zero, since (by our assumptions) A is strictly monotone. Similarly, $\bar{w}(w_1,w_2)$ is also increasing in w_2.

Putting w_{Sr} for $\underline{u} + \pi_{Sr}\Delta u$ etc., and using this 'mean-value' method, we can rewrite the bracketed term in expression (A.1) as

$$\frac{\varphi''(\bar{w}(w_{Sr},w_{Rr}))}{\varphi'(\bar{w}(w_{Sr},w_{Rr}))} - \frac{\varphi''(\bar{w}(w_{Rs},w_{Ss}))}{\varphi'(\bar{w}(w_{Rs},w_{Ss}))}.$$

Since $\pi_{Rs} > \pi_{Sr} > \pi_{Ss} > \pi_{Rr}$, we have $\bar{w}(w_{Sr},w_{Rr}) < \bar{w}(w_{Rs},w_{Ss})$. Hence the expression is negative. That is, q^* is decreasing in $\underline{u}$.

The argument for changes in Δu is similar but more involved. Using the same notation as before, by the mean value theorem,

$$\pi_1 \varphi'(w_1) - \pi_2 \varphi'(w_2)$$

$$= \left(\frac{\varphi^{-1}(\zeta_1) - \underline{u}}{\Delta u} \right) \varphi'\left(\varphi^{-1}(\zeta_1)\right) - \left(\frac{\varphi^{-1}(\zeta_2) - \underline{u}}{\Delta u} \right) \varphi'\left(\varphi^{-1}(\zeta_2)\right)$$

$$= (\zeta_1 - \zeta_2) \frac{\mathrm{d}}{\mathrm{d}\zeta} \left[\left(\frac{\varphi^{-1}(\bar{\zeta}) - \underline{u}}{\Delta u} \right) \varphi'\left(\varphi^{-1}(\bar{\zeta})\right) \right]$$

where $\bar{\zeta}$ is in the interval (ζ_2, ζ_1). Applying the inverse function rule, the last expression becomes:

$$\frac{(\zeta_1 - \zeta_2)}{\Delta u} \left[\frac{\varphi'(\bar{w})}{\varphi'(\bar{w})} + \bar{\pi}\Delta u \frac{\varphi''(\bar{w})}{\varphi'(\bar{w})} \right],$$

where $\bar{\pi}$ is in the interval (π_2, π_1), and $\bar{w} = \underline{u} + \bar{\pi}\Delta u$. That is,

$$\pi_1 \varphi'(w_1) - \pi_2 \varphi'(w_2) = \frac{(\varphi(w_1) - \varphi(w_2))}{\Delta u} \left[1 - \frac{\bar{\pi}\Delta u}{\bar{w}} R(\bar{w}) \right]$$

where $R(w) = -w\varphi''(w)/\varphi'(w)$.

We next show that the mean value is increasing in π_1 and π_2. By the implicit function theorem

$$- \frac{\partial \bar{\pi}}{\partial \pi_1} \frac{1}{\bar{w}^2} [\underline{u}R(\bar{w}) + \bar{\pi}\Delta u\bar{w}R'(\bar{w})] = \left[\frac{\bar{\pi}\Delta u}{\bar{w}} R(\bar{w}) - \frac{\pi_1\Delta u}{w_1} R(w_1) \right].$$

Given our assumptions, R is positive and increasing, so $\bar{\pi}$ is increasing in π_1. Similarly, it is increasing in π_2.

Let $\bar{\pi}(\pi_{Sr},\pi_{Rr})$ be the relevant mean-value for π_{Sr} and π_{Rr}, and let $\bar{\pi}(\pi_{Rs},\pi_{Ss})$ be the relevant mean-value for π_{Rs} and π_{Ss}. Since $\pi_{Rs} > \pi_{Sr} > \pi > \pi_{Rr}$, we have $\bar{\pi}(\pi_{Rs},\pi_{Ss}) > \bar{\pi}(\pi_{Sr},\pi_{Rr})$. Applying the mean-value method, we can rewrite the bracketed term in expression (A.2) as:

$$\frac{\bar{\pi}(\pi_{Rs},\pi_{Ss})\Delta u}{\underline{u} + \bar{\pi}(\pi_{Rs},\pi_{Ss})\Delta u} R(\underline{u} + \bar{\pi}(\pi_{Rs},\pi_{Ss})\Delta u) - \frac{\bar{\pi}(\pi_{Sr},\pi_{Rr})\Delta u}{\underline{u} + \bar{\pi}(\pi_{Sr},\pi_{Rr})\Delta u} R(\underline{u} + \bar{\pi}(\pi_{Sr},\pi_{Rr})\Delta u)$$

which is positive since R is positive and increasing. Hence q^* is increasing in Δu.

References

Dekel, E., Safra, Z., Segal, U., 1991. Existence and dynamic consistency of Nash equilibrium with non-expected utility preferences. Journal of Economic Theory 55, 229–249.

Dixit, A., Skeath, S., 1999. In: Games of Strategy. Norton, New York.

Gollier, C., 1995. The comparative statics of changes in risk revisited. Journal of Economic Theory 66, 522–535.

Grant, S., Kajii, A., Polak, B., 1998. Intrinsic preference for information. Journal of Economic Theory 83, 233–259.

Kreps, D., Porteus, E., 1978. Temporal resolution of uncertainty and dynamic choice theory. Econometrica 46, 185–200.

Pratt, J., 1964. Risk aversion in the small and the large. Econometrica 32, 122–136.