

Responsive Lotteries

Uriel Feige^{*} and Moshe Tennenholtz[†]

March 31, 2011

Abstract

Given a set of alternatives and a single player, we introduce the notion of a *responsive lottery*. These mechanisms receive as input from the player a reported utility function, specifying a value for each one of the alternatives, and use a lottery to produce as output a probability distribution over the alternatives. Thereafter, exactly one alternative wins (is given to the player) with the respective probability. Assuming that the player is not indifferent to which of the alternatives wins, a lottery rule is called *truthful dominant* if reporting his true utility function (up to affine transformations) is the unique report that maximizes the expected payoff for the player. We design truthful dominant responsive lotteries. We also discuss their relations with scoring rules and with VCG mechanisms.

1 Introduction

We consider a setting where there are n alternatives A_1, \dots, A_n and a single player. We assume that the player has a cardinal utility function over the alternatives, in the sense of Von-Neumann and Morgenstern. Namely, the player has a utility vector $U = (u_1, \dots, u_n)$, with utility value u_i associated with the respective alternative A_i , and this utility vector determines the preference of the player over different lotteries. Formally, given two lotteries, one that associates probabilities p_i with the respective alternative A_i , and the other associates probabilities q_i with the respective alternative A_i , the player prefers the former lottery if $\sum p_i u_i > \sum q_i u_i$, the latter lottery if $\sum p_i u_i < \sum q_i u_i$, and is indifferent over the choice of lotteries if $\sum p_i u_i = \sum q_i u_i$. Recall that Von-Neumann and Morgenstern show that if all that the player knows is his preferences over every conceivable pair of lotteries, and that these preferences are consistent in the sense that they satisfy a certain set of axioms (these axioms are natural, though there is a well known debate whether they actually reflect human behavior), then this in fact defines a utility function that is unique up to positive affine transformations (shift by a scalar and multiplication by a positive scalar – game theory literature often calls these linear transformations). We shall assume throughout that the player is not *indifferent* to the alternatives, namely, that there are at least two alternatives A_i and A_j with $u_i \neq u_j$. As utility functions are defined only up to affine transformations, we shall often represent utility functions in one of two canonical forms: either as *unit range*, meaning that $\min_i u_i = 0$ and $\max_i u_i = 1$, or as *unit sum*, meaning that $\min_i u_i = 0$ and $\sum_i u_i = 1$.

We introduce here a concept that we call a *responsive lottery*.

Definition 1 *Given a set of alternatives A_1, \dots, A_n and a single player, a responsive lottery is a mechanism that operates as follows:*

^{*}Department of Computer Science and Applied Mathematics, Weizmann Institute, Rehovot, Israel. uriel.feige@weizmann.ac.il. Work done at Microsoft R&D Center, Herzelia, Israel.

[†]Microsoft R&D Center, Herzelia, Israel, and Faculty of Industrial Engineering and Management, Technion, Haifa, Israel. moshet@ie.technion.ac.il.

1. The player provides a report $X = (x_1, \dots, x_n)$, where $X \in R^n$.
2. Using a function f from R^n to R^n , which is called the lottery rule, one computes a probability vector $f(X) = P = (p_1, \dots, p_n)$, with $p_i \geq 0$ and $\sum_i p_i = 1$.
3. A lottery is held and alternative A_i wins with probability p_i .

The lottery is *responsive* in the sense that the corresponding probabilities are not given in advance, but rather determined in response to the report of the player. The notion of an alternative winning the lottery should be aligned with what the utility function of the player refers to. For example, if the alternatives are the choice of seat in a certain flight (say, a window seat, an aisle seat, or a middle seat) and the utility function refers to the value the player associates with sitting in such a seat, then following the lottery the player should be seated in a seat corresponding to the winning alternative.

Given a responsive lottery and a utility vector U for the player, we say that the report X of the player is *honest* if $X = U$. Note that since utility functions are defined only up to affine transformations, we assume here that both X and U are given in the same canonical form (say, unit-sum). We say that the report X of the player is *rational* if X is such that $f(X)U = \sum p_i u_i$ is maximized. Namely, the player chooses a report that maximizes his expected payoff.

Definition 2 A lottery rule for a responsive lottery is truthful dominant if it has the property that for every utility function of the player, the honest report is rational, and every rational report is honest (or equivalently, given the first condition, the second condition is that the rational report is unique).

The truthful dominance property can be seen to combine three properties.

1. *Rational invertibility.* For every report X there is at most one utility function U for which X is a rational report.
2. *Rational uniqueness.* For every utility function U , there is a unique rational report X .
3. *Incentive compatibility.* For every utility function U , the report $X = U$ is rational.

Observe that rational invertibility and rational uniqueness are properties of the range of the lottery rule. Given a lottery rule f that is rational invertible, obtaining a truthful dominant lottery out of it involves only appending in front of it an appropriate permutation mapping π , that maps a report X to the report $Y = \pi(X)$ such that $Z = f(Y)$ maximize ZX . By rational invertibility, this now implies that given a utility function U , the unique rational report is $X = U$. This idea is similar to the revelation principle in mechanism design (see [5], for example). Note on the other hand that given a lottery rule that is incentive compatible, there does not seem to be a straightforward way to turn it into a truthful dominant rule. For example, the lottery rule that assigns $p_i = 1$ for the index i for which x_i is largest is incentive compatible, but there are many rational reports that give out no information beyond which is the preferred alternative for the player. Truthful dominance requires much more – that the rational report reveals the whole utility function.

A responsive lottery with a truthful dominant lottery rule may be viewed as a mechanism for elicitation of the utility function of the player. Recall that the work of Von-Neumann and Morgenstern already implies that utility functions can be inferred by observing preferences over lotteries. However, the procedure implicit in [7] involves a (potentially infinite) sequence of comparisons between pairs of lotteries (or a comparison among infinitely many lotteries, which is not feasible in practice). The mere fact that lottery comparisons are performed more than once is problematic for elicitation of utility functions. If the winning alternative is not actually given to the player after each lottery, the player

might not have incentives to report the truth. And if the winning alternative is given to the player after each lottery (assuming that this can be practically done), then the issue of complementarities among the alternatives might distort the original utility function of the player. We circumvent these difficulties by having only a single lottery. The aspect of this lottery that allows the elicitation of the utility function (if the player is rational) is its responsive nature. In a sense, the player is choosing among infinitely many lotteries. The rational invertibility property implies that the choice of the player allows one to infer his utility function (assuming that the player is rational). The incentive compatibility property makes it easy for the rational player to choose one lottery out of the infinite set of lotteries.

We assume infinite precision in the values of the utility function, in the reports and in the probabilities assigned by the lottery rule to the alternatives. Namely, they are real numbers. Employing our lottery rules with finite precision will obviously introduce rounding errors. We ignore this issue in this paper.

1.1 Related work

This manuscript refers only to utility functions as defined by Von-Neumann and Morgenstern [7]. It may be interesting to extend this work (if possible) to other notions of utility function (for the need for other notions, see for example [5] or [4]), but this is beyond the scope of the current work.

As far as we know, our notion of truthful dominant responsive lotteries is new. However, it is related to some other mechanisms for eliciting information from players. Most often, unlike our responsive lotteries, these mechanisms involve incentives in the form of transfer of money. *Strictly proper scoring rules* provide a mechanism for eliciting the belief of a player regarding the probabilities of future events. This is done by giving monetary rewards that depend on the predictions of the player, and on the actual realization of the future events. The relation between scoring rules and responsive lotteries will be discussed more extensively in Section 4. VCG mechanisms are a method for eliciting the true value that a bidder has for items that are sold in an auction. The incentives are built into the monetary payments that the bidder makes if he wins the item. The relation between VCG mechanisms and responsive lotteries will be discussed in Section A.2.

1.2 Our results

We view the introduction of the concept of truthful dominant responsive lotteries as one of the contributions of this work. Our main results are as follows:

1. We present a geometric approach for designing truthful dominant responsive lotteries, and use it to design what we call the *spherical lottery rule*. This lottery rule is continuous – a small change in the reports results in a small change in the probabilities of the alternatives. See Section 2
2. For three alternatives we present an algebraic approach for designing truthful dominant lottery rules. These rules are continuous. See Section 3.
3. We present methodologies for transforming any truthful dominant lottery rule over three alternatives to a truthful dominant lottery rule over $n > 3$ alternatives. The resulting lottery rule is not continuous. See Section 3.1.
4. We show a transformation from *bounded* proper scoring rules for n events to truthful dominant lottery rules for n alternatives. The resulting voting rule is not continuous. The transformation does not apply if the scoring rule is unbounded (such as the *logarithmic score*). See Section 4.

5. We show how the VCG mechanism (which involves money and multiple agents) can be used to design truthful dominant lottery rules (that involves only one player and no money). Also here, the resulting lottery rules are not continuous. See Section A.2 in the appendix.
6. We show a transformation from truthful dominant lottery rules for $n + 2$ alternatives to proper scoring rules for n events. By way of example, we use this transformation to derive a well known proper scoring rule, the *quadratic score*. We also show a transformation from truthful dominant lottery rules for $n+1$ alternatives to proper scoring rules for n events. Combining this with item (5) implies a methodology for deriving strictly proper scoring rules from the VCG mechanism. See Sections A.1, B.2 and B.3 in the appendix.

1.3 Some remarks

In our lotteries exactly one alternative wins. Our lottery mechanisms do not assume that if no alternative wins then the player gets 0 utility. If we wish to encompass situations in which valid outcomes include the possibility that no alternative wins, or that more than one alternative wins, the lottery mechanism needs to add these possible outcomes as additional alternatives.

The way of providing incentives to the player is by the choice of the winning alternative. There is no transfer of money involved in our mechanisms. Money can be introduced into our mechanisms by specifying alternatives that involve receiving or paying money.

The incentives in a truthful dominant lottery rule only refer to reporting the exact true utility function. In our mechanisms, there will also be some correlation between how close a report is to the true utility function and the expected value of the report. However, we make no formal claims regarding the nature of this correlation, and do not exclude the possibility that among two different reports, the one “further away” from the true utility function (according to some metric to be chosen by the reader) results in higher expected payoffs.

Some of the lottery rules that we design are continuous – a small change in the reports results in a small change in the probabilities of the alternatives. We view continuity as a desirable property for lottery rules, if one wishes them to be used in practice. Discontinuity of the lottery rule might have negative psychological effects on players who are not sure about their utility functions. They might spend too much time deliberating among reports that are almost identical but that lead to very different probability vectors. We note that for all our lottery rules, even the discontinuous ones, the value of their expected payoff is continuous (even though the probability vector might not be continuous), provided that the reports are honest.

An important aspect of a lottery mechanism is its *economic efficiency*. Namely, we want the winning alternative to be the one that is actually preferred by the player. For more than two alternatives, there are no economically efficient lottery rules that are truthful dominant. However, we remark that there are truthful dominant lottery rules that achieve almost perfect economic efficiency, though we do not advocate using them (see Section 5).

Our lottery rules provide ex-ante incentives to reveal the true utility function. However, this does not exclude the possibility that the player will experience ex-post regret. This issue too will be discussed in Section 5.

1.4 Ordinal utilities

Though our work is concerned with cardinal utilities, it may be instructive to consider first the case of ordinal utilities. In this case a responsive lottery is truthful dominant if the unique optimal report for the player is to report the alternatives in his order of preference (and specify ties, if there are any).

If $n = 1$, the problem is not interesting. The winning alternative is determined regardless of what the voter reports. If $n = 2$, the voter may report his preferred alternative, and the winning alternative is the reported alternative. If $n \geq 3$, there is no deterministic lottery with the rational uniqueness property. There are $n!$ different ranking orders (in fact more, if one allows ties), and only n possible winners. For any deterministic mechanism with $n \geq 3$, there are different orders that result in the same winner. Even if with respect to both orders reporting the truth gives the best payoff to the player, there is no incentive to the player in distinguishing between these two orders in the report.

This motivates considering randomized mechanisms (lotteries). Given a report that ranks the alternatives from 1 to n , we may let the j th alternative win with probability $\frac{2(n-j)}{(n-1)(n-2)}$ (or any other probability distribution that decreases with rank). If the player possesses a complete order over the alternatives, then the dominant strategy for the player is to report his true ranking.

Note however what happens if the player views two of the alternatives as being equivalent (a tie). Then asking the player to report a total order (with no ties) forces the player not to be truthful. Hence we should allow the player to report a tie among some alternatives. It is natural in this case to redistribute the probability of winning equally among the tied alternatives. Note however that by now we lost both the rational uniqueness property and the rational invertibility property. In case of a tie in the ranking, reporting a tie is not the unique best strategy. Reporting an arbitrary order among the tied alternatives gives the same expected payoff to the player as reporting a tie. This illustrates part of the challenges in designing truthful dominant mechanisms.

In this work we shall be interested not only in learning ordinal utilities, but cardinal utilities. Hence we wish to learn more than just the ranking, but we also have greater control of the rewards. The player can be incentivized not only by the choice of order among winning alternatives (which alternatives have higher probability of winning than others), but also by the choice of the actual values of these probabilities.

2 A geometric view of truthful dominant lottery rules

We present here a geometric view of truthful dominant lottery rules. We use it to design what we call the *spherical lottery rule* (which shares some common principles with the *spherical scoring rule*). More generally, the approach presented here leads to a geometric characterization of a wide class of truthful dominant lottery rules.

We shall use the following notation. The number of alternatives is n , the utility vector of the player is $U = (u_1, \dots, u_n)$, the report vector is $X = (x_1, \dots, x_n)$, and the probability vector is $P = (p_1, \dots, p_n)$, taken from an infinite set \mathcal{P} of *feasible* probability vectors (\mathcal{P} is the range of f for the lottery rule). Observe that all vectors $P \in \mathcal{P}$ are nonnegative and lie on the hyperplane $\sum p_i = 1$. Let $\bar{1} = \frac{1}{\sqrt{n}}(1, \dots, 1)$ denote the unit vector in the direction of the n -dimensional all 1 vector. Given an arbitrary vector Y , it can be decomposed into $Y = \alpha_Y \bar{1} + \beta_Y Y^\perp$, where Y^\perp is a unit vector orthogonal to $\bar{1}$, $\alpha_Y = \langle Y, \bar{1} \rangle$ and $\beta_Y = \langle Y, Y^\perp \rangle$. We assume w.l.o.g. that the sign of Y^\perp is chosen so that $\beta_Y \geq 0$. Observe that for all $P \in \mathcal{P}$ we have that $\alpha_P = \frac{1}{\sqrt{n}}$. Given a utility function $U = \alpha_U \bar{1} + \beta_U U^\perp$ and a probability vector $P = \frac{1}{\sqrt{n}} \bar{1} + \beta_P P^\perp$ for the responsive lottery, the payoff to the player is $\langle U, P \rangle = \alpha_U / \sqrt{n} + \beta_U \beta_P \langle U^\perp, P^\perp \rangle$. The rational report X for the player is the one for which $P = f(X)$ maximizes $\langle U, P \rangle$, and hence the maximum expected payoff attainable by the player is $\max_{P \in \mathcal{P}} \{ \langle U, P \rangle \}$. Observe that the optimal choice of $P \in \mathcal{P}$ is the one maximizing $\beta_P \langle U^\perp, P^\perp \rangle$ (since β_U is positive). The optimal P is preserved under positive affine transformations to U , because these transformations only change α_U and β_U (without flipping its sign) but not U^\perp .

We now describe a methodology for deriving truthful dominant lottery rules. (Presumably this methodology characterizes all continuous truthful dominant rules. Proving this appears to be an exercise

in formalities that does not add interesting insights, and hence will not be pursued here.) A compact convex body K will be called *nice* if (1) for every point z on its boundary ∂K there is a unique hyperplane H such that $H \cap K = z$, and (2) for every two points on ∂K , the line joining them lies entirely within K . For example, balls, ellipsoids and eggs are nice convex bodies, whereas polyhedrons are not. For nice convex bodies, for every vector v , there is a unique value $t(v)$ such that the closed halfspace $H(v) = \{x | \langle x, v \rangle \leq t\}$ (whose defining hyperplane is orthogonal to v) contains K and $\partial H \cap \partial K \neq \emptyset$. Moreover, ∂H and ∂K intersect in exactly one point.

Consider the $(n - 1)$ -dimensional subspace of R^n defined by the hyperplane $\sum p_i = 1$. Within its nonnegative orthant (satisfying $p_i \geq 0$ for every i) consider an arbitrary nice convex body K . Let \mathcal{P} (the set of feasible probability vectors for a responsive lottery) be precisely ∂K . Given a report X , consider the halfspace $H(X^\perp)$ as described above, and choose $P = f(X)$ to be the unique point $z \in \partial K$ intersecting $\partial H(X^\perp)$. This maximizes the projection of P on X^\perp , and hence maximizes $\langle P, X \rangle$. For this choice of lottery rule f , given a utility vector U , reporting $X = U$ maximizes the expected payoff. Moreover, for any report $X \neq U$, the probability vector $P = f(X)$ will be one that is strictly inferior to $f(U)$ in terms of the expected payoff.

It is natural to require (though not necessary) that the nice convex body K has geometric symmetries that reflect the intention that a-priori, all alternatives are treated symmetrically. In particular, in this case the center of mass of K will be at $\frac{1}{\sqrt{n}}\bar{1}$. Of all convex bodies, the most symmetric one is the ball, and its boundary is a sphere. Our *spherical lottery rule* uses a sphere centered at $\frac{1}{\sqrt{n}}\bar{1}$. To maximize the the variability in expected payoffs, this sphere has maximum possible radius. This radius is governed by the need to stay in the nonnegative orthant. A closest point P on the boundary of this orthant to the center of the sphere is $(0, 1/(n - 1), \dots, 1/(n - 1))$ for which $P^\perp = (-1/n, 1/n(n - 1), \dots, 1/n(n - 1))$. Hence the radius of the sphere is $1/\sqrt{n(n - 1)}$ (implying among other things that no entry in P is larger than $2/n$). Observe that using the spherical lottery rule, given a report $X = \alpha_X \bar{1} + \beta_X X^\perp$, the probability vector P is derived simply by projecting X^\perp on the sphere (along the line connecting X^\perp to the center of the sphere). We can assume that $X^\perp \neq 0$, by the assumption that the player is not indifferent. Hence

$$f(X) = \frac{1}{\sqrt{n}}\bar{1} + \frac{1}{\sqrt{n(n - 1)}}X^\perp = \left(\frac{1}{n}, \dots, \frac{1}{n}\right) + \frac{1}{\sqrt{n(n - 1)}}X^\perp$$

One readily observes that $P (= f(X))$ is an affine transformation of the report X . Hence the spherical lottery rule can be viewed as a normalization of the utility vector U with $\alpha_U = 1/\sqrt{n}$ and $\beta_U = 1/\sqrt{n(n - 1)}$, and following this normalization one simply takes $P = U$.

Theorem 3 *The spherical lottery rule described above is truthful dominant.*

The proof of Theorem 3 is implicit in the discussion preceding it. But let us sketch here yet another proof. Observe that for the spherical lottery rule, all vectors in \mathcal{P} have the same norm $1/\sqrt{n - 1}$. Hence the inner product $\langle U, P \rangle$ is maximized by the vector in \mathcal{P} that minimizes the angle with U . This vector is precisely the projection of U on the sphere, and $f(X)$ is this projection if and only if X is a positive affine transformation of U .

3 Three alternatives

In this section we design truthful dominant responsive lottery mechanisms for three alternatives. We assume that the utility function of the player is normalized to be unit range $0 = u_1 \leq u_2 \leq u_3 = 1$, and so are his reports.

If all reports are identical, then we set $p_i = 1/3$ for every alternative. If the reports are not identical, let the reports (after normalization) be $0 = x_1 \leq x_2 \leq x_3 = 1$. For simplicity of notation, let $x = x_2$.

Theorem 4 *Any responsive lottery over three alternatives satisfying all the following conditions is truthful dominant.*

1. $p_i \geq 0$ for $i \in \{1, 2, 3\}$.
2. $\sum p_i = 1$.
3. $p_1 \leq p_2 \leq p_3$, with $p_1(x) = p_2(x)$ iff $x = 0$ and $p_2(x) = p_3(x)$ iff $x = 1$.
4. p_2 is strictly increasing in x and p_1 and p_3 are strictly decreasing in x .
5. The derivatives satisfy $xp'_2(x) + p'_3(x) = 0$ for every $0 \leq x \leq 1$.

Proof. Conditions 1 and 2 are satisfied by every responsive lottery. Condition 3 ensures that in the optimal reports, the alternatives are ranked in their true order of preference (satisfy the ordinal aspect of the utility function). Specifically, $u_1 \leq u_2 \leq u_3$, with $u_1 = u_2$ iff $x_1 = x_2$ and $u_2 = u_3$ iff $x_2 = x_3$. Note that we assume that the player is not indifferent, and hence $u_1 < u_3$. After normalization to unit range, we have $0 = u_1 \leq u_2 \leq u_3 = 1$. For simplicity of notation, let $u = u_2$. We need to prove that the optimal report x is $x = u$.

The payoff to the voter is $v = u_1p_1(x) + u_2p_2(x) + u_3p_3(x) = up_2(x) + p_3(x)$. The derivative of v with respect to x is $up'_2(x) + p'_3(x)$ which equals 0 if $x = u$, by Condition 5, and only if $x = u$, by Condition 4 that implies that the derivatives are nonzero. This is the unique extremum for v . Condition 4 implies that this is a maximum rather than a minimum. ■

We remark that replacing condition 5 in Theorem 4 by the weaker condition that $-p'_3(x)/p'_2(x)$ is strictly increasing in x , from 0 to 1, will result in a rational invertible mechanism, though not necessarily incentive compatible.

Theorem 4 allows for many truthful dominant mechanisms. It is natural to limit the possible choices by normalizing the rewards such that there is some report for which $p_1 = 0$. It is easy to see that in conjunction with Theorem 4 this amounts to postulating that when $x = 1$ we have $p_1 = 0$ and $p_2 = p_3 = 1/2$. By way of example, we present two mechanisms that satisfy Theorem 4 and this additional requirement. They are named after the largest degree in the polynomials that are involved.

- The 3-alternative quadratic lottery rule: $p_1 = \frac{1-2x+x^2}{6}$, $p_2 = \frac{1+2x}{6}$, $p_3 = \frac{4-x^2}{6}$.
- The 3-alternative cubic lottery rule: $p_1 = \frac{1-3x^2+2x^3}{8}$, $p_2 = \frac{1+3x^2}{8}$, $p_3 = \frac{6-2x^3}{8}$.

3.1 Extension to more than three alternatives

Here we present two approaches for extending truthful dominant responsive lotteries over three alternatives to truthful dominant responsive lotteries with $n > 3$ alternatives.

The first of these approaches is as follows. Given the report of the player, pick uniformly at random three alternatives. If all three have the same reported utility, let each one of them win with probability $1/3$. Otherwise, normalize the part of the report of the player that refers to these three alternatives so that it becomes unit sum, and apply a 3-alternative truthful dominant responsive lottery on these three alternatives. It is not hard to see that this gives a truthful dominant responsive lottery.

In the appendix (Section B.1) we present another approach in more detail. When there are $n + 2$ alternatives A_0, \dots, A_{n+1} and the reports are $0 = x_0 \leq x_1 \leq \dots \leq x_n \leq x_{n+1} = 1$, it will give the $((n + 2)$ -alternative) quadratic lottery rule:

•

$$p_0 = \frac{n - \sum_{i=1}^n 2x_i + \sum_{i=1}^n (x_i)^2}{n^2 + 3n + 2}$$

•

$$p_i = \frac{n + 2x_i}{n^2 + 3n + 2}$$

for $1 \leq i \leq n$

•

$$p_{n+1} = \frac{2n + 2 - \sum_{i=1}^n (x_i)^2}{n^2 + 3n + 2}$$

4 Comparison with scoring rules

Our notion of truthful dominant lottery rules can be seen to be related to a notion known as *strictly proper scoring rules for categorical variables*. (For information on scoring rules, see for example [6] and [9].) Unlike our lottery rules, scoring rules involve transfer of money. In their basic form, scoring rules assume that the utility derived from money is linear in the amount of money, though this assumption can be avoided (for bounded scoring rules) by replacing transfer of a given amount of money by a lottery over a larger amount [1].

Let us provide a dictionary that we will use for both settings, that we call *lottery rules* and *scoring rules* respectively. To make the connection between lottery rules and scoring rules more apparent, it is convenient to normalize the utility function (and report) of the player to be unit sum.

n - number of mutually exclusive alternatives (for lottery rules) or mutually exclusive events (for scoring rules). Eventually, exactly one alternative will win/ one event will happen.

y_i - true utility of alternative i (for lottery rules) or true probability of event i happening (for scoring rules).

x_i - reported utility of alternative i (for lottery rules) or predicted probability of event i (for scoring rules).

$R_i(x_1, \dots, x_n)$ - probability of alternative i winning (for lottery rules) or reward to predictor (score of predictor) if event i happens (for scoring rules).

$\sum_i y_i R_i(x_1, \dots, x_n)$ - expected reward to player (for lottery rules) or to predictor (for scoring rule).

A dominant truthful lottery rule is one in which (unless the voter is indifferent) the unique report maximizing the expected reward is $x_i = y_i$. A strictly proper scoring rule is one in which the unique prediction maximizing the expected reward is $x_i = y_i$.

Note the similarity in the formula for expected reward for lottery rules and scoring rules. However, there are some noticeable differences between the two scenarios. (1) For lottery rules, there are no requirements when the voter is indifferent (all utilities are equal). For scoring rules, we wish the predictor to express true predictions even if all probabilities are equal. (2) For lottery rules, reports that differ by a positive affine transformation are equivalent. For scoring rules, this is not the case. (3) For lottery rules, the rewards (probabilities) are nonnegative and sum up to 1. For scoring rules, there is no such restriction. In fact, there is no strictly proper scoring rule in which rewards always sum up to exactly 1, since it must be the case that the unique maximum for sum of rewards is attained when all predictions are $1/n$.

Because of the above differences, there is no immediate equivalence between lottery rules and scoring rules. However, there are certain algebraic transformations between these two concepts. See Section A.1 in the appendix for more details. Moreover, there are geometric characterizations of scoring rules in a spirit similar to that of our geometric approach of Section 2. See [3]. In particular, our spherical lottery

rule is based on a high dimensional sphere, and so is the spherical scoring rule (though these are two different spheres).

5 Convex combinations

In general, given one truthful dominant mechanism, one can generate others by the method of taking *convex combinations*. We say that a mechanism is a convex combination of two mechanisms M_1 and M_2 if there is some probability $0 < q < 1$ such that with probability q the mechanism employs M_1 and with probability $1 - q$ it employs M_2 . Equivalently, given the reports, the probability of a given alternative to win is the convex combination (with weights q and $1 - q$) of the respective probabilities in M_1 and M_2 . The following proposition is self evident.

Proposition 5 *A convex combination of a truthful dominant responsive lottery with an incentive compatible responsive lottery is truthful dominant.*

As an example for the use of Proposition 5, the following responsive lottery is truthful dominant. Choose the alternative with highest reported value as a winner with probability q , and with the remaining probability employ the spherical lottery rule. Observe that as q approaches 1, this voting rule converges to optimal economic efficiency, but at the cost of weakening the incentives for the player to distinguish in his report between the less desirable alternatives.

As another example, consider the issue of ex-post regret involved in responsive lotteries. Even though reporting the true utility function is optimal for the player ex-ante, the player may suffer ex-post regret after the lottery is held. For the 3-player quadratic lottery rule, given any report other than $(0, 1, 1)$ the least desirable alternative might win the lottery, and then the player may regret not having reported $(0, 1, 1)$ which would have avoided this possibility. To prevent this ex-post regret, one may take a convex combination of the quadratic rule with the uniform rule (each alternative equally likely to win), which ensures that regardless of the report of the player, the least desirable alternative has some probability of winning. This decrease in ex-post regret comes at the cost of economic efficiency (among other things).

The availability of several (infinitely many) mechanisms that are truthful dominant allows one to introduce some additional objective function, and select the mechanism that optimizes this additional property. For example, among all truthful dominant mechanisms one may want to select the mechanism minimizing the maximum probability with which the least desirable alternative wins. (This probability is $1/6$ for the 3-player quadratic lottery rule.) However, for this particular objective function, there is no truthful dominant mechanism that minimizes it. For every truthful dominant mechanism, the value of this objective function is strictly positive. But then it can be lowered by taking a convex combination with an incentive compatible mechanism. This relates to the fact that the notion of truthful dominant mechanisms is defined using strict inequalities, and its closure is the incentive compatible mechanisms.

6 Applications

The notion of truthful dominant responsive lotteries is a mathematical construct that may have practical applications. We believe that in choosing a truthful dominant lottery rule for a practical application, one would need to strike a careful balance among several considerations (such as ex-post regret, economic efficiency and the strength of the incentives, see Section 5), and this will be possible only if the number of alternatives is fairly small (four alternatives appears to be a good number). Also, we believe that a carefully designed user interface may help the players understand the concept of responsive lotteries and the effect of the choice of report on the expected reward. For example, one may imagine an interface which includes sliding bar controls for each alternative, and a screen showing a pie chart for the relative

probability of each alternative winning. As the player moves the sliding bars to indicate to which extent he values each alternative, the pie chart changes dynamically. Having such a user interface in mind is one of the reasons why we wish lottery rules to be continuous (and among our lottery rules for more than three alternative, only the spherical one is continuous).

Though the purpose of this manuscript is mainly to develop the mathematical theory of truthful dominant responsive lotteries, we briefly and informally discuss some potential applications. A formal treatment of these and other potential applications will hopefully be undertaken elsewhere. In all cases below we assume that the winning alternative can actually be given to the player after the lottery is held (or at least, that the player believes that this is what will happen).

Experimental psychology. If one wishes to gain a quantitative understanding of preferences of people over a set of alternatives, one may in principle use truthful dominant responsive lotteries. For example, a psychophysical experiment may study the relative sensation of pleasure or pain associated with various temperatures, and the alternatives may be those of putting one's hand in containers of water of various (possibly unpleasant but not harmful) temperatures. One may not want to repeat such an experiment many times with the same subject, due to effects of adaptation, and responsive lotteries may serve as a way of eliciting more information in fewer experiments. As always in experimental settings, caution is needed in performing experiments and in interpreting the results (which at best indicate what were the true preferences of the subject under the conditions of the experiment).

Market research. A company may use responsive lotteries to gain understanding of the preferences of its potential costumers. For example, an airline company may offer some passengers on a flight a responsive lottery over the choice of seat (say, a window seat, an aisle seat, or a middle seat) so as to get a sense of what the true preferences of costumers are.

Multiple-agent mechanism design. In many settings one is interested in designing a mechanism in which agents report their utilities, and then some global decision is taken so as to optimize some objective function that depends on the true utilities of the agents. The difficulty is often in incentivizing the agents to reveal their true utilities. Mechanisms based on statistical approaches often take a small sample of agents, ask them for their utility function, and use this output so as to reach a global decision that effects the agents not in the sample. In such a setting (and assuming no externalities), the agents in the sample have no incentive to not tell the truth, and hence are sometimes assumed to be truthful. See for example [2] for the use of a statistical approach in the design of combinatorial auctions. If this statistical approach is combined with truthful dominant responsive lotteries then there is more justification in assuming that the agents in the sample are truthful. Let us provide a hypothetical example.

Suppose a company wants to reward each one of its employees with a \$100 gift certificate to some store chain. There are two store chains that are being considered (say one specializes in electronics, one in sports). However, (almost) all gift certificates should be to the same store chain, as then the company gets a big discount from the store chain. How can the company decide among the store chains? One option is to sample at random a small number of employees and offer each one of them a truthful dominant responsive lottery over four alternatives, where two of them are the gift certificates to the two chains, and the other two alternatives are \$50 and \$100 in cash (for calibration). Each employee in the sample actually does get the alternative that wins the respective lottery. All remaining employees get gift certificates to just one store chain, and this store chain is determined based on the information elicited by the responsive lotteries (and on the objective of the company, which may be for example to maximize welfare). Arguably, employees in the sample will actually reveal their true references. If the total number of employees is large, then with high probability this mechanism leads to almost optimal economic efficiency: the sample size may be chosen to be large enough to be representative, yet small enough to make the marginal inefficiencies small (inefficiencies resulting from giving sampled employees more expensive rewards, and from occasionally giving sampled employees less favorable rewards, due to

the random nature of responsive lotteries).

Acknowledgements

We thank Ran Smorodinsky and Aviv Zohar for many helpful discussions.

References

- [1] Franklin Allen. Discovering personal probabilities when utility functions are unknown. *Management Science*, Vol. 33, No. 4, April 1989, 542–544.
- [2] Shahar Dobzinski, Noam Nisan, Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. *STOC 2006*: 644–652.
- [3] Tilmann Gneiting and Adrian Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 359–378, 2007.
- [4] Daniel Kahneman and Amos Tversky. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, XLVII (1979), 263-291.
- [5] Andreu Mas-Colell, Michael Whinston, Jerry Green. *Microeconomic Theory*. Oxford University Press 1995.
- [6] J. McCarthy. Measures of the value of information. *Proceedings of the National Academy of Sciences*, 42, 654–655, 1956.
- [7] John von Neumann, Oscar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, Princeton, third edition 1954.
- [8] Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani. *Algorithmic Game Theory*, Cambridge, 2007.
- [9] L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66, 783–801, 1971.

A Main appendix

A.1 Transformation between lottery rules and scoring rules

We first design scoring rules based on lottery rules (in which utilities and reports are normalized to be unit sum).

Given n events and predicted probabilities x_i , one may attempt to associate an alternative with each event, and use the probabilities of the corresponding lottery rule as rewards for the scoring rule. However, this fails due to differences 1 and 2 listed above. Hence we shall design scoring rules for n alternatives based on lottery rules for $n + 1$ alternatives. To the scoring rule, add the alternative A_0 with $y_0 = x_0 = 0$. Given reports $x_1 \leq \dots \leq x_n$ for the scoring rule, append to them the report $x_0 = 0$ and then use the lottery rule. Let $y_1 \leq \dots \leq y_n$ be the true probabilities. The properties of the lottery rule easily imply that it is optimal for the predictor to set $x_n \geq x_i$ for all i . Thereafter, since we add A_0 , the lottery rule implies that to maximize $\sum_{i=0}^n y_i R_i(x_1, \dots, x_n) = \sum_{i=1}^n y_i R_i(x_1, \dots, x_n)$ one needs to have $x_i/x_n = y_i/y_n$ for every i . The requirement that $\sum_{i=1}^n x_i = 1$ then implies that the unique optimal prediction is indeed the true probabilities.

In Section B.2 in the appendix we show how one can derive an n -alternative scoring rule from an $(n+2)$ -alternative lottery rule with utilities in unit range form. We use this approach to derive the well known quadratic scoring rule.

We now show how to design lottery rules (in unit sum representation) based on scoring rules. We assume here that the rewards in the scoring rule are bounded. Some strictly proper scoring rules (such as the *logarithmic score*) are not bounded: if $x_i = 0$ and event i does happen, the reward might be $-\infty$. However, some other strictly proper scoring rules (such as the *quadratic score*, and the *spherical score*) are bounded.

Consider an arbitrary strictly proper scoring rule with nonnegative rewards for which for any report of predicted probabilities, the sum of all rewards is at most 1. Every strictly proper scoring rule with bounded rewards can be transformed into such a scoring rule by adding a large enough constant to all scores so that they are nonnegative, and thereafter multiplying all scores by a properly chosen constant so that they sum up to at most 1. Now we can use the rewards of the scoring rule (which are nonnegative and sum up to at most 1) as the probabilities of each of the alternatives winning in the lottery rule. This almost works, except for the fact that the probabilities might sum up to less than 1, and then with some remaining probability no alternative wins. This is not allowed in lottery rules. (To see why this is not allowed, think of the case that one of the original alternative was itself the null alternative, and the player reported nonzero utility for it.) Hence we need to decide which alternative wins with the remaining probability. Here are two options that do not work:

1. A random alternative is chosen with uniform probability. This can easily be seen to fail, by considering the quadratic scoring rule. After scaling this rule gives $p_i = (2x_i + 1 - \sum_j (x_j)^2)/(n+1) \geq 0$ with $\sum p_i = (n+2 - n \sum (x_j)^2)/(n+1) \leq 1$. Updating p_i to $p_i - (1 - \sum p_j)/n$ we get the linear functions $(2x_i + 1 - 1/n)/(n+1)$. This is not even incentive compatible. Regardless of the true utility function of the player, the report maximizing the expected payoff for linear rules is to report only the best alternative ($x_n = 1$ and $x_i = 0$ for all $i \neq n$).
2. A random alternative is chosen with probability proportional to the reward that the scoring rule gives for it. This is equivalent of scaling all the p_i by the same multiplicative factor so that they all sum up to 1. This can easily be seen to fail for the spherical scoring rule, which gives $p_i = x_i/\sqrt{\sum (x_j)^2}$. After this scaling we just get $p_i = x_i$ which again is linear and hence not incentive compatible.

Given that the above options do not work, we present an approach that does work. Given the reported utilities, we first give every alternative except for the one with smallest report (breaking ties arbitrarily, and redistributing probabilities at the end if needed) probability $1/n$ of winning. Only on the remaining $1/n$ probability, we make the probability of winning be the normalized value of the rewards of the scoring rule. If the sum of rewards of the scoring rule is less than 1, with the remaining probability we let the lowest alternative win. Since $y_1 = 0$, this does not change the expected payoff. Also, this cannot make the probability of winning the least desirable alternative higher than the probability of winning any other alternative. Hence we do inherit from the scoring rule the property that to maximize $\sum_{i=1}^n y_i R_i(x_1, \dots, x_n)$, one needs to have $x_i = y_i$.

We note that the approach described above of creating lottery rules out of scoring rules gives lottery rules whose only discontinuity is when the utility of an alternative becomes equal to that of the least favorable alternative.

For scoring rules with unbounded rewards (as in the logarithmic score, that would give $p_i = \log x_i$), the transformation above does not apply, and it is not clear whether there is a simply way of deriving lottery rules out of them.

A.2 Lottery rules from VCG

The VCG mechanism (see for example [5] or [8], and references therein) uses money to elicit the true utility function of bidders. It is always incentive compatible, and when there is the right kind of competition for items, being truthful may become the unique optimal strategy. Here we show how VCG mechanisms can be used in order to derive truthful dominant lottery rules (without money).

Given a utility function $0 = u_1 \leq u_2 \leq \dots u_n = 1$ for the player, let the reports be $0 = x_1 \leq \dots x_n = 1$. (To simplify notation, we assume here that the order among x_i is the same as the order among u_i . The justification for this assumption will become apparent soon.) Let π be some arbitrary probability distribution with full support in the range $[0, 1]$. For simplicity, one may think of π as the uniform probability distribution. Every such π will result in a lottery rule (that depends on π). Choose y_1, \dots, y_n independently at random, each from the probability distribution π . Think of the y_i values as the reported bids of a second bidder in a VCG mechanism. The winning alternative is the one maximizing the sum $x_i + y_i$. (We assume for simplicity that ties will not occur, as indeed is the case for uniform π .) In addition, the player pays $\max_j y_j - y_i$. Being truthful is the unique optimal strategy for the player if the payment is in units that correspond to the units of his utility function. Those units are u_n , which by our simplifying assumption (to be justified soon) are the same as x_n . This payment can be simulated by transferring $\max_j y_j - y_i$ from the probability that alternative n wins to the probability that alternative 1 wins.

However, we need to maintain the property that the probability that alternative 1 wins is lowest among all alternatives, and the probability that alternative n wins is highest. To achieve this, we play the VCG game only with probability $1/(n+1)$. In addition, let each of the alternatives $2 \leq i \leq n-1$ win with probability $1/(n+1)$, and let alternative n win with probability $2/(n+1)$. Now we can indeed shift probability from alternative n to alternative 1 while maintaining the above property.

In summary, the truthful dominant responsive lottery is as follows.

1. Obtain the report $0 = x_1 \leq \dots x_n = 1$.
2. With probability $(n-1)/(n+1)$: each of the alternatives $2, \dots, n$ wins with conditional probability $1/(n-1)$ (or true probability $1/(n+1)$).
3. With probability $2/(n+1)$: play the VCG mechanism with random y_j and let i be the index maximizing $x_i + y_i$, and let the largest y_j value be \hat{y} .
 - (a) With conditional probability $1/2$ (true probability $1/(n+1)$), let alternative i win.
 - (b) With conditional probability $1/2$ (true probability $1/(n+1)$), let alternative 1 win with conditional probability $\hat{y} - y_i$, and alternative n win with probability $1 + y_i - \hat{y}$.

Remark. In Section 4 we have seen that truthful dominant lottery rules can be used in order to derive strictly proper scoring rules. This implies that our approach can also be used in order to derive scoring rules from VCG mechanisms. See Section B.3 in the appendix for a short discussion on this.

B Secondary appendix

B.1 From three alternatives to more

We present here an approach for using 3-alternative truthful dominant lottery rules to derive n -alternative truthful dominant lottery rules with $n > 3$.

Choose values $0 < q_2 \leq q_3 \leq \dots \leq q_n$. (The requirement that $q_n \geq q_{n-1}$ may be relaxed as the example below shows.) Choose an arbitrary truthful dominant three alternative lottery rule M_3 , and let m denote the maximum probability with which the third alternative wins in M_3 .

Given a unit range report $0 = x_1 \leq \dots \leq x_n = 1$ of the player for the n -alternatives mechanism, for every $1 < i < n$ consider M_3 with alternatives A_1, A_i, A_n and reports $0, x_i, 1$. Let p_1^i, p_i^i, p_n^i be the probabilities assigned by M_3 to the three alternatives in this case. Scale these probabilities by $\frac{q_2}{m(n-3)}$, thus obtaining q_1^i, q_i^i, q_n^i .

The tentative probability assigned to alternative A_i is p_i , where $p_1 = \sum_{i=2}^{n-1} q_1^i$, $p_i = q_i + q_i^i$ for $1 < i < n$, and $p_n = q_n + \sum_{i=2}^{n-1} q_n^i$. To get the actual probabilities, scale by $1/(\sum p_i)$.

Note that the scaling of p_1^i ensures that $p_1 \leq p_2$. It is not hard to see that it is optimal for the player to report the true ordinal preference over alternatives, and that given that the player does so, reporting the true x_i is optimal, a property inherited from M_3 .

As an example, let us apply the second approach together with the quadratic lottery rule when $n = 4$. We choose $q_2 = q_3 = 1/6$ and $q_4 = -1/3$ (the minimum value that will ensure $p_4 \geq p_3$). After normalization we get that given a report $0 = x_1 \leq x_2 = x \leq x_3 = y \leq x_4 = 1$ we have the 4-alternative quadratic lottery rule:

- $p_1 = \frac{2-2x-2y+x^2+y^2}{12}$, $p_2 = \frac{2+2x}{12}$, $p_3 = \frac{2+2y}{12}$, $p_4 = \frac{6-x^2-y^2}{12}$.

Note that the 4-alternative quadratic lottery rule is not continuous in the following sense: as x_2 approaches 0 (and x_3 does not), p_2 does not converge to p_1 . Likewise, as x_3 approaches 1 (and x_2 does not), p_3 need not converge to p_4 .

More generally, when there are $n + 2$ alternatives A_0, \dots, A_{n+1} we obtain the $((n + 2)$ -alternative) quadratic lottery rule referred to in Section 3.1.

B.2 Deriving the quadratic scoring rule from the quadratic lottery rule

Here we show how one can derive an n -alternative scoring rule from an $(n + 2)$ -alternative lottery rule with utilities in unit range form. Add two auxiliary alternatives A_0 and A_{n+1} with reported utilities 0 and 1 respectively. Note that adding the auxiliary alternatives ensures that the player is not indifferent, and that the x_i values are now treated on an absolute scale rather than a relative one. Now the lottery rule for $n + 2$ alternatives gives a scoring rule for the n events. Note that the lottery rule also has some probability of $R_{n+1}(x_1, \dots, x_n)$ of alternative A_{n+1} winning. For the expected score to reflect this, add $R_{n+1}(x_1, \dots, x_n)$ to the reward of the predictor.

As an example, let us design a scoring rule for n events based on the quadratic scoring rule. Given predicted probabilities $0 \leq x_1 \leq \dots \leq x_n$ with $\sum x_i = 1$, we need to add two alternatives A_0 and A_{n+1} , assume reports $x_0 = 0$ and $x_{n+1} = 1$ and employ the quadratic scoring rule on $n + 2$ alternatives. We then see that the reward for event i happening is $\frac{2x_i + 3n + 2 - \sum_{j=1}^n (x_j)^2}{n^2 + 3n + 2}$. Note that by scaling all rewards by the constant $(n^2 + 3n + 2)$ and then subtracting the constant $(3n + 3)$ the rewards become $2x_i - 1 - \sum_{j=1}^n (x_j)^2$. This is a well known scoring rule, called the *Brier score* or *Quadratic score*.

B.3 Deriving scoring rules from VCG mechanisms

The n -alternative lottery rule of Section A.2 derived from the VCG mechanism naturally gives a scoring rule for alternatives $2, \dots, n - 1$. The reward for event i happening is the expected probability for alternative i winning, plus the expected probability of alternative n winning. Computing these expectations might not be easy (they depend both on the predictions x_i and the probability distribution π). This computational difficulty may be overcome if we allow probabilistic scoring rules in which the

reward is not a deterministic function of the event that happens, but rather a probabilistic function, and the predictor wishes to maximize the expected reward (this is consistent with the philosophy of scoring rules). In this case, do the following. Let k be the event that actually happened. Draw variables y_1, \dots, y_n at random. Let j be the index maximizing $x_j + y_j$, and let \hat{y} be the largest of the y_i values. Give a reward of $1 + y_j - \hat{y}$. In addition, if either $j = k$ or $j = n$ then give an additional reward of 1. Linearity of expectation shows that the expected reward for this probabilistic reward function given that k happened is exactly equal to the reward of the VCG based mechanism.