

Yootopia!

Daniel M. Reeves Bethany M. Soule Tejaswi Kasturi
Yahoo! Research Columbia University Yahoo! Research

Ver. 305: 2007-12-23 20:47, checked in at 20:35:07 by dreeves

Abstract

The nascent Yootopia Project at Yahoo! Research brings together three related research agendas: mechanisms for group decision making, prediction, and payment infrastructure. The currency (scrip system) is called Yootles and underlies (but is orthogonal to) the group decision and prediction mechanisms. We present an array of currency-agnostic decision mechanisms for small groups, describing new and existing mechanisms for (1) choosing among a short list of options, (2) choosing among an effectively innumerable list of options, as in meeting scheduling, (3) allocating shared goods and responsibilities, (4) public good provision, and (5) bilateral trade. We list desirable mechanism properties and describe the tradeoffs that the mechanisms make among them. Finally, we describe a small step towards synthesis of group decisions and group prediction: an interface for friendly wagers.

1 Introduction

There are plenty of established mechanisms for making group decisions, from arguing to sophisticated voting systems. Often choosing an outcome involves explicit negotiation, compromising, quid pro quoing, etc. Yet these mechanisms often fail to make the choice that would have maximized group welfare. There are many classic pitfalls in group decision making (like various forms of groupthink) that can be addressed with more structure in the decision-making process.

Group prediction is a related problem and suffers similarly from a lack of structure. Experts are hard to identify, they have biases, and are notoriously bad judges of their own degree of certainty. Prediction markets are an increasingly popular approach to incentivizing information revelation and intelligently aggregating it to predict the future [Surowiecki, 2005].

To achieve greater social welfare¹ and fairness in everyday decision making we suggest there is a need for a common metric for strength of preferences (utility), and the ability to directly transfer utility from one person to another. Of course, these are the roles that money plays in society but money is commonly

¹I.e., sum of the participants' utilities.

eschewed for these valuable purposes in settings such as social groups and within organizations. For this reason, as well as to provide a currency for prediction markets, we introduce *yootles*.² The mechanisms we describe here, however, are currency-agnostic.

Key to adoption of this approach to group decision making is accessibility and ease of use. We have implemented an SMS (Short Message Service) phone interface to the mechanisms presented here, with a web interface to them at <http://yootopia.org>. We document this interface in Appendix B.

Consider first the value of a common metric for utility. People often have no way to compare preferences. “I want” vs. “I really want” may do a poor job of comparing how much we value different options—a prerequisite for making socially optimal choices. Second, groups often have at best extremely rudimentary means for compensating the people on the losing end of decisions. Decision mechanisms with yootles payments (i.e. auctions) address both of these shortcomings through explicit compensation and preference elicitation.³

Admittedly, when one is not in the habit of quantifying one’s preferences, it can seem a daunting task. Nonetheless, it is quite possible to do and with practice can even become simple. When you are unsure of your own value for a particular outcome, we find a simple binary search is extremely effective at pinning down your *indifference point*. “Would I take that plane ticket if it were free?” “Definitely.” “Would I buy it if it cost \$1000?” “No way.” “\$500?” “Uhh, No.” “\$250?” “Definitely.” “\$375?” “Yeah, I guess.” “\$438?” “Uhh, well...” The point where you are truly torn is your true utility.⁴

But why not just use an existing currency (like dollars) for these purposes? In fact, in some settings an exchange rate will quickly emerge and then yootles would be money in every sense. But some groups may want to level the playing field among a group of people of varying financial means. This is difficult to enforce but a like-minded group can achieve this property by agreeing, for example, to only use yootles for decision-making and prediction tasks. Yootles may also simply provide a more explicit framework for casting currency as a descriptor of a person’s utility—a concept that most people outside certain fields like economics and artificial intelligence are not used to. Finally, there may be various social impediments to the adoption of government currency for the purposes we propose. Paying for influence in group decisions is not palatable to many. And spending one’s own money in the context of one’s job is avoided in corporate culture.

Yootles, just like with government currency, work because of mutual agreement. And just like with with money, the whole economy is nothing but a scorekeeping system for owed favors (broadly defined). Participants start with a balance of zero and simply owe or are owed yootles. All yootles transac-

²The name is a variant on *utils* or *utiles*—a hypothetical unit of utility or happiness.

³Quantifying one’s preferences can be awkward and difficult when not in the habit of doing it, however. When unsure of your own value for a particular outcome, we find a quick mental binary search is extremely effective at pinning down your *indifference point*.

⁴One advantage of iterative mechanisms, which we will discuss later, is to spare the participants this sometimes painful introspection [Parkes, 2005].

tions, debts, and balances are tracked in a ledger system. This paper presents group decision making as a key application of the yootles infrastructure, being launched by Yahoo! Research at www.yootles.com. We largely defer discussion of this payment infrastructure and prediction markets in Yootopia, though these are all interrelated.

Key to adoption of this approach to group decision making is accessibility and ease of use. We have implemented an SMS (Short Message Service) phone interface to the mechanisms presented here, with a web interface to them at <http://yootopia.org>. We document this interface in Appendix B.

2 Desirable Mechanism Properties

Following are commonly desirable mechanism properties for group decision-making, often treated as constraints on a mechanism designer [Mas-Colell et al., 1995]. Famous impossibility results [Green and Laffont, 1979] preclude achieving all of these properties in any mechanism simultaneously and in the next section we describe various tradeoffs that can be made among these properties.

Fundamental to the study of mechanisms is the notion of an agent's preferences. We refer to an agent's *type* synonymously with preferences to refer to the private information an agent has that affects its utility for the different outcomes. Typically, this is simply the utility values it would get from each possible outcome of the decision mechanism.

EFF (Social Efficiency) A socially efficient (a.k.a. socially optimal or welfare optimal) mechanism always chooses the outcome that maximizes total utility (a.k.a. social welfare). The degree of social efficiency of a mechanism is the sum of everyone's utility for the choice the mechanism made divided by the greatest social welfare over all possible choices.

PAR (Pareto Efficiency) A Pareto efficient mechanism is one which guarantees an outcome that cannot be improved upon without making some participant worse off. This is commonly viewed as a minimum requirement for a mechanism. It may be forgivable to find a suboptimal solution but to offer solution A when every single participant at least weakly prefers B will be viewed as a failure on the part of the mechanism designer. Note that $\text{EFF} \implies \text{PAR}$.⁵

(DS)IC (Dominant Strategy Incentive Compatibility) In a dominant strategy incentive compatible (a.k.a. strategy-proof or non-manipulable) mechanism it is in every participant's best interest to truthfully report their preferences, no matter how anyone else is playing the game. We will use the terms IC and DSIC interchangeably. Note that DSIC does not imply EFF or PAR, the Prisoners' Dilemma being a famous counterexample where the dominant strategy

⁵Proof: Consider a mechanism that is socially efficient but not Pareto efficient. Then someone's utility can be increased without decreasing anyone else's. This yields greater social efficiency, which is a contradiction.

is defection despite mutual cooperation yielding strictly greater utility for both players.

BNIC (Bayes-Nash Incentive Compatibility) In a Bayes-Nash incentive compatible mechanism it is in every participant's best interest to truthfully report their preferences, so long as everyone else is also using the strategy of truthful bidding.⁶ This assumes, however, that there is a common-knowledge distribution from which agents' preferences are drawn (and that they have no other information about these preferences). Since this assumption tends not to hold in practice, it will often mean little to claim the BNIC property for practical mechanisms. Note that since any dominant strategy equilibrium is also a Nash equilibrium, $IC \implies BNIC$. By $*IC$ we will mean the set of both properties, $\{BNIC, DSIC\}$.

BB (Budget Balance) In a budget balanced mechanism, the sum of all payments made is zero. In other words, the only payments are transfers to other participants. A BB mechanism does not run a surplus or require a subsidy. (Weak budget balance means that the mechanism never requires a subsidy but may run a surplus.) We will typically view budget balanced mechanisms as advantageous but if a for-profit entity is designing the mechanism, they may prefer to maximize the surplus (revenue) generated. In this case the mechanism would be said to be aiming for Revenue Maximization (REV). The literature on "optimal auctions" [Myerson, 1981] refers to this design goal, especially in the context of mechanisms for selling goods to multiple buyers.

IR (Individual Rationality) In an individually rational mechanism, no risk-neutral agent will strictly prefer to opt out of the mechanism rather than participate. Any mechanism in which participants may place bids of zero and guarantee that they do no worse than opting out altogether is IR. Note that any mechanism can be made IR by including "opt out" in the strategy space.

EQ (Fairness or Equitability) The degree of fairness of an outcome for n people is n times the utility of the least happy person divided by the total utility (equivalently, the ratio of the minimum utility to the average utility). In other words, perfect fairness means everyone is equally happy. We call a mechanism fair, or *equitable*, if it guarantees perfectly fair outcomes.

ENV (Envy-freeness) An envy-free mechanism yields outcomes in which no one wants to trade places with anyone else.

SMP (Simplicity) For many applications, the complexity (either computational or cognitive) of a mechanism may need to be minimized. Like many of

⁶The Revelation Principle actually guarantees that, with some caveats, we can turn any mechanism into one that is Bayes-Nash incentive compatible.

the above, this goal may need to be traded off against other mechanism properties [Nisan and Ronen, 1999]. For the purposes of small group decision making, we will consider a mechanism to have the property SMP if the mechanism can be conducted easily without a computer.

Some of the above mechanism properties may be further specialized. A mechanism has a property *ex ante* if it has the property in expectation before any of the participants' preferences are known. A mechanism has a property *ex interim* if, from the perspective of any particular participant, it has that property in expectation with only the participant's preferences known. Finally, a mechanism has a property *ex post* if it has the property regardless of any of the participants' preferences.

If the mechanism includes a randomization device we model that by having Nature be one of the participants and the result of the randomization constitutes Nature's "preferences". Flipping a coin, then, to decide who gets the last piece of cake is *ex ante* fair and *ex interim* fair from the perspective of the participants other than Nature. It is *ex post* quite unfair.

3 Auctions for Group Decisions

We now describe a potpourri of decision mechanisms for various group decisions. We start with some trivial, degenerate, and mundane mechanisms before continuing on to some widely applicable and well-known mechanisms, as well as some practical original mechanisms. We do this to highlight the improvements that can be made over some of the mechanisms traditionally applied to the domain of group decision making.

Table 1 summarizes the properties of the mechanisms we describe in this section.

Most of the mechanisms we consider here take a set of reported agent preferences and return one of a set of possible outcomes along with a payment (positive or negative) to or from each agent. This makes them direct revelation mechanisms (DRMs) or, simply, direct mechanisms.⁷ A DRM is a sealed-bid mechanism where the bids are in fact agents' full preference functions, mapping every possible outcome to a utility value. Mechanisms of this form are very general. A voting mechanism, for example, ignores the magnitudes of the reported nonzero valuations for options and sets all the payments to zero. A bilateral bargaining mechanism takes the buyer's and seller's valuation and reservation prices, chooses an outcome from the set {TRADE, NO TRADE}, and sets the payments to be a sale price and the negative of the sale price (in the case of outcome TRADE).

The Revelation Principle [Mas-Colell et al., 1995] in fact guarantees that for any mechanism, involving arbitrarily complicated sequences of messages be-

⁷The Revelation Principle [Mas-Colell et al., 1995] in fact guarantees that for any arbitrarily complicated mechanism, there exists a DRM that is equivalent in terms of how it maps preferences to outcomes and is BNIC.

	EFF	PAR	IC	BNIC	BB	EQ	ENV	SMP
Coin Flip	0	0	1	1	1	0	0	1
Dictatorship	0	1	1	1	1	0	0	1
Voting	0	1	0	0	1	0	0	1
VCG	1	1	1	1	0	0	0	0
SGA(1/3, 0)	1	1	0	1	1	> 0	0	1
SGA(1/2, 0)	1	1	0	0	1	$2/3^*$	0	1
SGA(0, 1)	1	1	0	0	1	> 0	0^*	1
DAUC	0^*	1	0	0	1	> 0	0	0
NFA	0^*	1	0	0	1	0^*	0^*	1

Table 1: Properties of decision mechanisms. Starred entries indicate that the property holds ($= 1$) when players bid truthfully despite it not being an equilibrium (i.e., for altruistic agents). SGA is the two-player shared-good auction with uniform types. The entry for voting applies to many common voting systems including approval voting, Borda count, and instant runoff.

tween participants, there exists a direct revelation mechanism that is equivalent in terms of how it maps preferences to outcomes and is BNIC. Consider a meta-mechanism that consists of the original mechanism with proxies inserted that take reported preferences from the agents and play a Nash equilibrium on the agents’ behalf in the original game (every game with a finite number of possible actions has a Nash equilibrium [Nash, 1951]). Similarly, the revelation principle means that if a mechanism has a dominant strategy equilibrium then we can construct one that is DSIC.

Sometimes reporting preference functions is straightforward—as in the case of specifying a utility for each of a short list of options—while for other social choice functions it is prohibitive. And even when reporting preference functions is straightforward, the mechanism may not be able to find a Nash equilibrium of the game. So although it is theoretically without loss of generality to consider only direct revelation mechanisms, indirect mechanisms are important in practice [Conitzer and Sandholm, 2004] (see Section 3.9).

3.1 Coin Flip

We start with perhaps the simplest of all decision procedures. The coin flip mechanism refers more generally to picking an option for a group of participants by choosing among the options randomly (picking the i th of n outcomes with probability $1/n$). This mechanism has the advantage of being ex interim fair (EQ). It is also trivially IC in that the outcome is independent of the participants’ bids (no bids in fact are asked for). The biggest downside of randomization as a mechanism is efficiency—an option may be chosen that makes everyone miserable, with no possible bias toward options that are better for the group. Thus, Coin Flip fails to achieve PAR or EFF.

Coin Flip does achieve two properties that are harder to come by in more sophisticated mechanisms. It trivially balances the budget (BB), and with reasonable constraints on the possible outcomes (excluding, for example, the “rob agent i blind” option), Coin Flip is individually rational (IR). Any mechanism with no payments is of course trivially budget balanced and typically satisfies individual rationality as well.

Like many decision mechanisms, Coin Flip will not in general be envy-free (ENV). For example, if the randomly chosen option is to give agent i the last piece of cake, everyone but i will be envious. Envy-freeness is typically achieved with payments—if agent i paid more for the cake than any of the others thought it was worth then the ENV property would be satisfied.

Of course, for all its shortcomings, the coin flip mechanism has extremely low cognitive and computational complexity (SMP).

3.2 Dictatorship

Dictatorship means throwing out every participant’s preferences except for one (the dictator) and optimizing with respect to only the dictator’s preferences. This has greater expected efficiency than Coin Flip and does achieve Pareto efficiency (as long as the dictator breaks ties magnanimously) but otherwise is no better than Coin Flip.

Incentive compatibility is trivially satisfied since either your preferences are disregarded or followed completely. Not surprisingly, for many decision problems Dictatorship is the least fair mechanism, though if a dictator is chosen randomly, then this mechanism, like Coin Flip, is *ex ante* fair. For all of the other properties (BB, IR, ENV, SMP) the story for Dictatorship is identical to that for Coin Flip.

3.3 Voting

We define voting very generally as any decision mechanism (i.e., mapping from preferences to outcomes) that never involves payments. Thus, Coin Flip and Dictatorship both qualify as degenerate forms of voting. It turns out in fact that for choosing among more than two options, Dictatorship and mechanisms like Coin Flip that may rule out options even when unanimously preferred are the only voting mechanisms that are dominant strategy incentive compatible. This is the famous Gibbard-Satterthwaite impossibility theorem [Gibbard, 1973, Satterthwaite, 1975, Gibbard, 1977].

Just as Coin Flip and Dictatorship are degenerate decision mechanisms that take full preferences, largely ignore them, and return an outcome, many voting mechanisms take full preferences and ignore all but the ranking information. For Borda count, for example, the n options are sorted for each agent by utility and reassigned a number from one to n . In the case of approval voting all positive utilities are mapped to one. In fact, in our implementation of approval voting positive utilities are mapped to one and the rest to zero unless there are no positive utilities reported in which case negative utilities are mapped to zero

and the zero utilities (the default when unspecified) are mapped to one. This provides a convenient way to approve of *all except* a set of options.

In every decision mechanism discussed in this paper, including the degenerate mechanisms, all voting systems, and the auction-based mechanisms in the next sections, the social choice is made by applying any filters to the preferences and then choosing the option with the highest sum. Payments (generally made from the winners to the losers) are calculated as a function of the chosen option and the reported utilities but, again, for voting mechanisms the payments are always set to zero. As described below, our principle motivation for preferring auction mechanisms over voting is to achieve greater social efficiency and fairness.

Voting, of any form, is not socially efficient (EFF) in that it does not guarantee the outcome maximizing total utility. This is because it does not fully take into account the strength of participants' preferences. For example, consider a vote to choose a restaurant that includes a steakhouse and a Thai restaurant. If Alice is a vegetarian and abhors the steakhouse option, the most influence she can exert is to rank it last or otherwise put the full weight of her vote elsewhere. If the steakhouse is otherwise popular, any voting mechanism will choose it even if the others prefer it only mildly over Thai. Given the others' mild preferences and Alice's strong preference, Thai is clearly the socially optimal choice, contrary to the result of the vote. Note that this is a problem with all possible variants of voting (approval voting, Borda count, instant runoff, etc.) because fundamental to any voting system is the requirement that each participant have equal influence on the outcome. As we will see in the subsequent subsections, auction-based mechanisms remove this constraint.

In the restaurant example above, the steakhouse option is Pareto efficient since some participants prefer it to Thai (Thai is Pareto efficient as well). Voting mechanisms in fact are Pareto efficient (PAR) in general since they will always pick an option that at least some people prefer.

We know (Gibbard-Satterthwaite) that no reasonable voting system achieves IC and many voting mechanisms are famously susceptible to strategizing. In many political elections, for example, expressing your truthful preference for a third party candidate is tantamount to wasting your vote if you are convinced that, given your expectation of others' strategies, your first choice is unlikely to win. In principle (the revelation principle) there exist BNIC voting mechanisms but this is of little practical value due to the reliance on a common knowledge distribution of voter preferences. Thus we characterize voting systems generally as not BNIC in Table 1.

Voting does not achieve fairness (EQ), even *ex ante*, since preferences may be polarized such that certain participants can expect to be unhappy with the outcome. Voting is of course commonly thought of as fair in the colloquial sense of giving every participant an equal opportunity to influence the decision.

For budget balance (BB), individual rationality (IR), and envy-freeness (ENV), the story for voting systems is the same as that for Coin Flip and Dictatorship. Finally, many forms of voting have the advantage of being computationally simple (SMP), requiring only a show of hands or paper ballots. This property is

lost, however, for more sophisticated voting mechanisms.

3.4 Vickrey-Clarke-Groves (VCG) Mechanism

None of the decision mechanisms discussed up to this point have involved payments. As discussed in Section 1, the use of payments allows the explicit comparison of strength of preferences as well as the ability to redistribute utility to achieve greater fairness. We start with a classic mechanism that achieves what none of the previous mechanisms can: social efficiency.

The VCG [Vickrey, 1961, Clarke, 1971, Groves, 1973] mechanism is impressively versatile (though many criticisms have been levied against it, for example, that it is vulnerable to collusion). It can be applied whenever agents can specify their full preferences and when finding the socially optimal choice given those preferences is computationally feasible. For the group decisions we are addressing, these requirements are largely met. As we will see, VCG achieves most of the key mechanism properties at the expense of BB.

The mechanism proceeds by collecting sealed bids from each agent, which the mechanism takes to be their (possibly strategically falsified) preferences. It then selects the outcome O^* that maximizes total reported utility. IC is then achieved thanks to a simple rule for establishing the payments: each agent pays the externality⁸ it imposes on the rest of the group. In other words, for each agent i , the mechanism solves the subproblem with all agents but i , yielding outcome O_{-i}^* . If $O^* = O_{-i}^*$, i.e., i is not *pivotal*, then i 's participation cost the group nothing and i pays nothing. But if the social optimum changes without i ($O^* \neq O_{-i}^*$) then i was pivotal in the final decision and pays precisely what it cost the rest of the group: $\text{pay}_i = \sum_{j \neq i} U_j(O_{-i}^*) - \sum_{j \neq i} U_j(O^*)$. In words, a pivotal agent i pays the difference between the total utility of everyone else without i , and their utility with i .

Intuitively this means that if an agent's preferences do not impact the rest of the group, the auction costs it nothing. If an agent does pay, its payment is entirely insensitive to the amount of its bids (conditional on being asked to pay at all—i.e., influencing the outcome). For this reason, it is a dominant strategy for agents to report their preferences truthfully in VCG.

Incentive compatibility in VCG comes at the cost of budget balance (BB). In fact, when the externalities are negative (as is the case for choosing among contentious options) VCG maximizes revenue among all efficient mechanisms. When externalities are positive (as is the case for chore division) VCG requires a subsidy. Another potential complaint about VCG is a low degree of fairness (EQ) and lack of envy-freeness (ENV). Next we consider mechanisms that address these shortcomings.

⁸An agent's externality is the aggregate utility (often negative) that the other participants derive from its participation. For example, telephone usage has positive externalities (the more others use the technology the more useful it is); for road usage, the externalities are negative due to congestion.

3.5 Shared-Good Auction (SGA): Two-Person “Voting”

Consider the problem of two people trying to decide between two options. Unless both players prefer the same option, no standard voting mechanism (with either straight votes or a ranking of the alternatives) can help with this problem. VCG is a nonstarter with no one to provide a subsidy and third-party payments tantamount to a pure efficiency loss (burning money).

Instead we propose a simple auction: each player submits a bid and the player with the higher bid wins, paying some function of the bids to the loser in compensation. Reeves [2005] considered a special case of this auction and gave the example of two roommates using it to decide who should get the bigger bedroom and for how much more rent. We also find it a most practical way to allocate tasks for which two people have joint responsibility—e.g., deciding who books flights for a joint trip.

We sometimes refer to this mechanism as an *un-sharing auction*—it allows one agent to sell its half of a good to the other joint owner (or pay the other to take on its half of a “bad”).

We define a space of mechanisms for this problem that satisfy BB, IR, and (with a minor caveat noted below) EFF. The following is a payoff function defining a space of games parameterized by the function f .

$$u(t, a, t', a') = \begin{cases} t - f(a, a') & \text{if } a > a' \\ \frac{t - f(a, a') + f(a', a)}{2} & \text{if } a = a' \\ f(a', a) & \text{if } a < a', \end{cases} \quad (1)$$

where $u()$ gives the utility for an agent who has a value t for winning (its type) and chooses to bid a (its action) against an agent who has value t' and bids a' . Finally, $f()$ is some function of the two bids.⁹ In the tie-breaking case the payoff is the average of the two other cases, i.e., the winner is chosen by the flip of a fair coin.

We next consider a restriction of the class of mechanisms defined by Equation 1:

Definition 1. $SGA(h, k)$ is the mechanism defined by Equation 1 with $f(a, a') = ha + ka'$.

These are games in which the winner pays a linear function of the two bids. For example, in $SGA(1/2, 0)$ the winner pays half its own bid to the loser; in $SGA(0, 1)$ the winner pays the loser’s bid to the loser. We now give Bayes-Nash equilibria for such games when agents’ values (types) are uniformly distributed on an interval.

Theorem 1. *That is, for $h, k \geq 0$ and types $U[A, B]$ with $B \geq A + 1$ the following is a symmetric Bayes-Nash equilibrium of $SGA(h, k)$:*

$$a(t) = \frac{t}{3(h+k)} + \frac{hA + kB}{6(h+k)^2}.$$

⁹ Reeves [2005] considered the case $f(a, a') = a/2$.

We provide the proofs of theorems in Appendix A.

We can now characterize the truthful mechanisms in this space. According to Theorem 1, SGA(1/3, 0) is BNIC for $U[0, B]$ types. The following theorem says this is the *only* truthful mechanism for uniform types.

Theorem 2. *With $U[0, B]$ types ($B > 0$), SGA(h, k) is BNIC if and only if $h = 1/3$ and $k = 0$. Furthermore, for $U[A, B]$ types with $A > 0$ there is no setting of h and k such that SGA(h, k) is BNIC.*

Of course, by the revelation principle, it is straightforward to construct a mechanism that is BNIC for any $U[A, B]$ types. However, to be a proper auction [Krishna, 2002] the mechanism should not depend on the types of the participants. In other words, the mechanism should not be parameterized by A and B . With this restriction, the revelation principle fails to yield a BNIC mechanism for arbitrary uniform types.

By construction the game is budget balanced (BB) regardless of the strategies since there is a single transfer payment from the winner to the loser. Likewise, individual rationality (IR) is guaranteed by construction since either agent can opt out by bidding zero and guaranteeing itself zero payoff.

The caveat to the claim of efficiency for SGA is that the strategies must be monotone increasing in type, i.e., greater utility for winning can never lead you to bid less. This is the case for the Nash equilibrium in Theorem 1, for truthful bidding, and arguably for any other sane bidding strategy. If additionally both agents play the same monotone strategy then the agent with the higher type will always win. Since the transfer payments, being budget-balanced, do not contribute to the efficiency, SGA will achieve the efficient outcome (EFF).¹⁰

The shared-good auction is ex post fair when the winning agent pays half its surplus to the loser in compensation. For example, SGA(1/2, 0) would be ex post fair (EQ) if played altruistically. When played strategically with $U[A, B]$ types, we can calculate the degree of fairness from the Nash equilibrium in Theorem 1, which for $h = 1/2$ and $k = 0$ reduces to $(2t + A)/3$. The utility of the least happy agent is $\min(\frac{2t_w + A}{6}, t_w - \frac{2t_w + A}{6})$, where t_w is the winner's type, which reduces to $(2t_w + A)/6$ since $t_w \geq A$. Since the total utility is t_w , the degree of ex post fairness is $2/3 + A/(3t_w)$ and (taking the expectation with respect to t_w) the degree of ex ante fairness is $2/3 + A/(A + 2B)$.

Theorem 3. *For strategic agents with $U[A, B]$ types ($B \geq A + 1$) playing the symmetric equilibrium given by Theorem 1, there is no setting of h and k such that SGA(h, k) guarantees ex post fairness (EQ).*

Finally, we consider the remaining properties of SGA. The shared-good auction is Pareto efficient (PAR) since someone ends up with the good. When h and k are such that the winner pays at least the loser's bid to the loser—i.e., $k \geq 1$ —then the mechanism is envy-free (ENV). Finally, it is straightforward

¹⁰The total social welfare, then, is the type of the player with the highest type, i.e., the first order statistic of the type distribution. The expected social welfare for $U[A, B]$ types is $A/3 + 2B/3$.

to collect bids, pick the largest, and pay a simple function of the bids, with no computational assistance (SMP). This is especially true for $h = 0$ and $k = 1$ —i.e., when the winner pays the loser the loser’s bid. It is this last case, SGA(0, 1), that we generalize in the next section.

3.6 A General Decision Auction (DAUC)

The SGA family comprise versatile decision mechanisms for two players. We now generalize to the n -player case with the Decision Auction (DAUC)—a budget balanced mechanism that modifies VCG by redistributing the VCG surplus so as to increase fairness with limited impact on incentive compatibility.

Others have proposed similarly motivated modifications of VCG. Cavallo [2006] has characterized preference distributions for which the surplus can be partially redistributed with no loss of IC or EFF. Faltings [2004] proposes the simple method of achieving BB by arbitrarily excluding some participants from the decision and funneling the surplus to them. This retains IC at the cost of some EFF. Parkes et al. [2001] apply a similarly motivated VCG payment adjustment approach in the setting of combinatorial exchanges.

DAUC uses a related method for redistributing the surplus. After making the VCG payments each participant receives back a piece of the surplus in proportion to the difference between their highest bid and their bid for the winning option (participants whose highest bid was placed for the winning option will thus not receive any payout). This, in effect, means that losers get compensated in proportion to how badly they wanted some other result. If someone is neutral among the options and bids zero across the board, they will not receive any payment, whereas someone who really wanted one of the losing options will receive a higher payout—provided that at least one person was pivotal in the auction. If no one is pivotal then, like straight VCG, the mechanism does not require any payments. Note that in the case of two agents and two options this mechanism reduces to SGA(0, 1). Thus we know from above that at least in the two-player case with uniform types, the Decision Auction is efficient even when played strategically.

The redistribution of payments amongst the losers means that this mechanism is not incentive compatible (IC). For example, placing your highest bid for something you know will lose will garner you more of the redistributed surplus. Likewise, if you suspect you will be pivotal and may need to pay close to your full bid to influence the outcome, you may prefer to reduce your bid in order to lose (have a less preferred option chosen) and get compensated. In general, however, you can’t have greater influence on the chosen outcome by changing your bids, you can only capitalize on losing. This can be considered a form of partial incentive compatibility in the sense that, among a small group of people who care foremost about finding the social optimum, reasonable participants are unlikely to try to capitalize off the rest of the group. This is in contrast to straight voting, for example, where strategizing may be required to keep your vote from being wasted.

DAUC also retains incentive compatibility in a different sense: a sufficiently

risk-averse agent, assuming a minimally diffuse distribution of others' preferences, will not inflate its bids since doing so entails a risk of a negative utility (paying more for an outcome than it's worth). Truthful bidding, in fact, is a minimax strategy¹¹ in DAUC.

In general, when dealing with mechanisms that are not incentive compatible, it is useful to introduce some notion of the degree of incentive compatibility of the mechanism. To these ends we would like to measure how much an agent stands to gain by bidding its best response to truthful bidding. This can be done for the two-player version of this game, as it reduces to SGA(0, 1), and we get the following result:

Theorem 4. *For two players with types $U[0, B]$ truthful bidding is an ε -equilibrium in DAUC, where $\varepsilon = \frac{B^3}{72}$.*

For example, for the case of $U[0, 1]$ types, an agent can improve its utility $1/72$ by deviating from truth-telling—a gain of 4%. Proof of this theorem is provided in the full version of this paper.

3.7 Nominally Fair Auction (NFA)

What if, hypothetically, we trusted participants to just play nicely (altruistically)? What mechanism properties besides *IC could we achieve? The answer is all of them. NFA—for Nominally (or perhaps Naively) Fair Auction—is a simple direct revelation mechanism achieving, with altruistic agents, every property in Section 3 (except *IC). Like VCG, it chooses the social optimum from reported preferences (which in the case of VCG are the true preferences due to IC) but allocates payments to maximize fairness, i.e., to make everyone equally happy. Specifically, with n agents, let b_i be the i th agent's bid for the winning option, and U be the total utility for the winning option. Then each agent i receives (or pays if negative) $U/n - b_i$. Since the sum of these payments is zero, regardless of the strategies, we achieve BB regardless of agent strategies.

NFA is not generally BNIC. Consider the two-player case where this mechanism reduces to SGA(1/2, 0). With types drawn from $U[A, B]$, we see from the equilibrium strategy that the game is not BNIC except for agents with the lowest possible value A (the Nash equilibrium is $(2t + A)/3$). As for DAUC, however, we can find the degree of incentive compatibility in the two-player case of NFA:

Theorem 5. *For two players with types $U[0, B]$ truthful bidding is an ε -equilibrium in NFA, where $\varepsilon = \frac{B^3}{48}$.*

There are cases where NFA is BNIC, if not DSIC. Consider the case of allocating a shared good with a known common utility (perhaps no one wants

¹¹A minimax strategy is found by considering for each strategy the lowest possible utility if that strategy is followed (as if the other players cared for nothing but to harm you) and picking the one for which this minimum is maximized. In other words, it is the strategy which maximizes worst-case utility.

the good itself but it has a known resale value). It is in equilibrium to bid the full value of the good, which amounts to splitting the good evenly among all participants. In fact, any time there is complete information reasonable people tend to pick NFA [Schelling, 1960].¹²

NFA is also envy-free (ENV) with altruistic bidders. Suppose it weren't and for some set of preferences an agent wanted to trade places with another. But by social efficiency, the one who wanted it most got it. By fairness, the two are receiving equal utility after the payments. Thus, the agent would receive less utility by switching, which is a contradiction. Finally, NFA is cognitively, computationally, and mathematically simple (SMP). (At least, it is no more complicated than splitting a dinner bill.)

Bilateral Trade Using NFA, we can also cast Bilateral Trade (a.k.a. bargaining) as a group decision problem. Myerson and Satterthwaite [1983] show that it is impossible for any mechanism to achieve EFF, IR, and BB in bilateral trade for any reasonable (essentially, overlapping) distribution of preferences. In order to use NFA for bilateral trade, the seller submits a negative bid for how much it will cost to give up the good in question—option TRADE—with an implicit zero bid for option NO TRADE. The buyer submits a positive bid for how much they would be willing to pay for option TRADE, with zero for option NO TRADE. (Of course, either agent may strategically misreport these numbers.) If the buyer's bid and seller's bid sum to greater than zero then, appropriately, TRADE will be chosen, otherwise NO TRADE. Since NFA splits the reported surplus evenly this means that the selling price is fixed halfway between the (magnitude of) the seller's and buyer's values. Chatterjee and Samuelson [1983] give the Nash equilibrium of this game for $U[0, 1]$ types for a seller (1) and a buyer (2):

$$\begin{aligned} a_1(t_1) &= 2t_1/3 + 1/4 \\ a_2(t_2) &= 2t_2/3 + 1/12. \end{aligned}$$

Note that NFA does not apply to the case of multilateral trade as the mechanism shares surplus over *all* participants. This is clearly not correct if the final outcome is a transaction that involves only some subset of the bidders—i.e. in the case of general trade we don't want to reward bidders just for participating in the auction.

Joint Purchase Joint Purchase (a.k.a. public good provision) refers to the problem of n agents deciding if a good to be shared is worth the collective cost. In order to apply NFA to Joint Purchase, we introduce the special player, Nature (see Section 2). Since the seller is not participating in the auction, but rather the good has an exogenous price, we need a way to include the actual price of the good in the auction. We have Nature submit a negative bid equal to the price of the good. The participants will submit their (presumably positive)

¹²Consider the example of Alice selling her car to her friend Bob when they know the dealership would buy it for \$1000 and sell it for \$2000. They'll naturally choose a sale price of \$1500, per NFA. We see how NFA achieves this result below.

bids. The set of possible choices consists of {BUY, NO BUY}. The mechanism proceeds as described above, selecting the option that maximizes total utility, and calculating payments, with the one additional caveat that Nature makes and receives no payments. Since Nature is really just a placeholder for the cost of the good in question, we don't want to share the surplus with it. The two-player version of this game was suggested by Reeves [2005].

Here we give the Nash equilibrium of NFA for two players both reporting nonnegative values for a single option, with an exogenous cost (negative bid from Nature) of $c \geq 2/3$.

Theorem 6. *For the game defined by the payoff function*

$$u(t, a, t', a') = \begin{cases} t - a + \frac{a+a'-c}{2} & \text{if } a + a' \geq c \\ 0 & \text{otherwise,} \end{cases}$$

and types distributed $U[0, 1]$, the following is a symmetric Bayes-Nash equilibrium:

$$a(t) = 2t/3 + c/4 - 1/6.$$

3.8 Sequential Pie Auction

Using auction mechanisms for voting can be controversial. A common objection is captured by the mantra “One person, One vote.” The Sequential Pie Auction mechanism addresses this objection by explicitly giving each participant the same number of votes (one or two, as we will see below) and then giving them the opportunity to buy and sell their votes from each other in a series of auctions.

While vote buying is taboo in political elections it is common for small group voting (“I’ll support your proposal if you make these concessions”). And although such mechanisms will still meet resistance on the grounds that vote buying should not be explicit, we argue that the distinction is rather meaningless.

The Sequential Pie Auction mechanism is as follows:

1. Pick a number of votes, k , to give to each of the n participants such that the total number of votes is an even number (i.e., $k = 1$ if n is even, else $k = 2$).
2. While votes are not all owned by the same person, repeat:
 - (a) Pick the two biggest equal size blocks of votes such that each of the two blocks are owned by different people, breaking ties arbitrarily.
 - (b) Hold a shared-good auction to reallocate all the votes in both blocks to one of the two people.

Note that the special case of $n = 2$ is simply a single shared-good auction—both players have an equal stake in the decision and the first and only iteration of the auction gives the choice to one or the other.

3.9 Iterative Decision Auction

We now consider iterative mechanisms [Parkes, 2006]. The above mechanisms handle the case of voting on a small list of options. The iterative decision auction addresses the case when there are so many options that it is prohibitive for participants to explicitly assign values to each of them—for example, picking a time for a meeting. In this mechanism, a default option is chosen arbitrarily and broadcast to the participants who then increase their bids (from the default of zero on all options) on a subset of the options (of their choosing) which they prefer to the current default. Based on the current bids, the socially optimal choice is re-broadcast as the new default option and the process repeats. When none of the participants respond to the default by increasing any of their bids, the auction ends and the current default is chosen, with payments made according to any of the above mechanisms (DAUC, NFA, VCG, etc).

The iterative decision auction retains the key properties of the mechanism that is used to decide the payments based on the final bids and last default option broadcast (we call this the base mechanism). For example, its degree of budget balance is naturally the same as the base mechanism since no payments are made until the final bids are fed to the base mechanism. Similarly, the social efficiency, fairness, and individual rationality are inherited from the base mechanism. Disregarding the cognitive/computational burden of constructing and transmitting one’s full preferences, the iterative decision auction also inherits the incentive compatibility of the base mechanism.

4 A Simple Prediction Mechanism

As described in the next section, we have implemented a simple interface for friendly wagers as the first step toward integrating group decision-making with group prediction. The wagering system allows participants to enter probabilities of various outcomes, providing the average of the probabilities as an aggregate measure of the consensus judgment. To motivate participation, users also specify a worst-case amount of yootles they are willing to lose if they are wrong. The system then computes fair odds based on the reported probabilities of pairs of participants. Fair odds are taken to be $\frac{1-p}{p}$ to 1 where p is the average of the reported probabilities. In this way, the participants have equal assessments of the expected payout of the wager. Finally, the payouts are scaled by the minimum of the amounts the participants were willing to risk. When the actual outcome resolves, one of the participants specifies this to the mechanism, which then effects the appropriate yootles payouts (from the loser of the bet to the winner).

Our next goal is to integrate full-fledged prediction markets in Yootopia. In particular, we would like to implement conditional prediction markets that depend on policy (group decision) questions in the spirit of Robin Hanson’s Futarchy¹³ to further marry group prediction and group decisions.

¹³<http://hanson.gmu.edu/futarchy.html>

5 Implementation

We have created an auction/voting/wagering system implementing the mechanisms described in this paper. This system needs to have the properties that it be easy to use, (nearly) universally available, and easily modifiable so that new mechanism types can be implemented quickly without disturbing existing auctions. To that end, the system uses SMS (Short Message Service) text messages from participants for all tasks required to create and run an auction or wager. Alternatively, participants without SMS access can also use a web browser to access the same capabilities, using the same commands as SMS participants. Users can even interact with the yootles system via a command line tool, written using a provided Perl API. The bulk of the interface is self-explanatory, and can be read by either visiting the website (currently at <http://yootopia.org>) or texting the keyword YOOTLES to the phone code 4INFO (44636). The interface allows the user to create an auction or wager, designate the mechanism used to run the auction, bid or bet, check results, and transfer payments according to the mechanism results. Each auction is designated by a short keyword at creation time which is used for all subsequent commands. Any user can create an auction, which can also be created simply by bidding on an unknown auction name. Any other user can bid on that auction if they know the keyword, and they will receive the current results of the auction. Bidding without specifying an amount will also return the current auction results. At the end of the auction any participant can initiate a yootle payout process to automatically enter payments (if any) into the ledger system. Using the same command users can also make arbitrary payments that are not associated with an auction keyword.

Wagering uses the same system, except that users specify their expected probability of a given outcome and the maximum amount of yootles they are willing to lose. After the wager is resolved, any participant specifies whether the specified outcome occurred, and payments are given out accordingly. We use a ledger system to keep track of yootle balances for all users in the system.

Planned enhancements include adding restrictions for who can bid in a given auction and resolving auction name conflicts by changing to a local rather than global namespace, using the auction creator's realm of trusted users.¹⁴ Also, additional security restrictions so that only designated owners of an auction can change its mechanism or payout proceeds are necessary. The documentation of this system is included in Appendix B and (more up to date) at <http://yootopia.org>.

6 Conclusion

We have introduced a currency (yootles) for facilitating group decision making and prediction. While in many settings yootles will become equivalent to any regular, government-issued currency their use can be kept intentionally separate, allowing for them to function in situations where there may be reasons to eschew

¹⁴See <http://ripplepay.com/about>.

the use of money. We propose use of an online ledger system to track balances, and peer-to-peer credit networks to mitigate issues of default.

Our primary focus in this paper is on an array of decision mechanisms for common group decision problems. In particular we introduce a new mechanism (DAUC) for redistributing the VCG surplus to achieve greater fairness while retaining a semblance of incentive compatibility. We also introduce a mechanism (NFA) for achieving all desirable mechanism properties in the case that agents play altruistically. For many applications this is less useful but we describe applications of NFA to domains where it is a sensible mechanism when played strategically (bilateral trade and joint purchase). Additionally we are able to give limited results on the degree of incentive compatibility of these two mechanisms (in the restricted case of two-players), finding that DAUC is closer to incentive compatible than NFA.

For domains such as meeting scheduling where the submission of full preferences is costly we propose iterative versions of the above mechanisms. Finally, we describe our implementations of these mechanisms in a new service from Yahoo! Research.

Acknowledgments

David Pennock, Seung-Taek Park, David Yang, Yiling Chen, and Daniel Raffel have provided invaluable and ongoing assistance in this work.

References

- Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of vcg payments. In *Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 2006.
- Kaylan Chatterjee and W. F. Samuelson. Bargaining under incomplete information. *Operations Research*, 31:835–851, 1983.
- Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- Vincent Conitzer and Tuomas Sandholm. Computational criticisms of the revelation principle. In *Fifth ACM Conference on Electronic Commerce*, pages 262–263, New York, 2004.
- Boi Faltings. A budget-balanced incentive-compatible scheme for social choice. In *AAMAS-04 Workshop on Agent-Mediated Electronic Commerce*, 2004.
- Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.

- Jerry R. Green and Jean-Jacques Laffont. *Incentives in Public Decision-Making*. North-Holland Publishing Company, 1979.
- Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- Vijay Krishna. *Auction Theory*. Academic Press, 2002.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- John Nash. Non-cooperative games. *Annals of Mathematics*, 2(54):286–295, 1951.
- Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *Thirty-first ACM Symposium on Theory of Computing*, pages 129–140, 1999.
- David C. Parkes. Iterative combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 2. MIT Press, 2006.
- David C. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and AI*, 44:269–302, 2005. Special Issue on the Foundations of Electronic Commerce.
- David C. Parkes, Jayant R Kalagnanam, and Marta Eso. Achieving budget-balance with Vickrey-based paymentschemes in exchanges. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 1161–1168, 2001.
- Daniel M. Reeves. *Generating Trading Agent Strategies*. PhD thesis, University of Michigan, 2005.
- Mark A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–216, 1975.
- Thomas C. Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA, March 1960. ISBN 0674840313. URL <http://www.amazon.co.uk/exec/obidos/ASIN/0674840313/citeulike-21>.
- James Surowiecki. *The Wisdom of Crowds*. Anchor, August 2005. ISBN 0385721706. URL <http://www.amazon.co.uk/exec/obidos/ASIN/0385721706/citeulike-21>.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

Appendix

A Proofs

A.1 Proof of Theorem 1

We show that for the two-player game with types $U[A, B]$ and payoff function

$$u(t, a, t', a') = \begin{cases} t - ha - ka' & \text{if } a > a' \\ \frac{t - ha - ka' + ha' + ka}{2} & \text{if } a = a' \\ ha' + ka & \text{if } a < a', \end{cases}$$

with $h, k \geq 0$ and $B \geq A + 1$ that the following is a symmetric Bayes-Nash equilibrium strategy:

$$\frac{t}{3(h+k)} + \frac{hA+kB}{6(h+k)^2}. \quad (2)$$

Consider first the special case that $h = k = 0$. Equation 2 prescribes a strategy of bidding ∞ and it is clear that this is a dominant strategy in a game where the winner is the high bidder with no payments required.¹⁵ We will now assume that $h + k > 0$.

Define $m \equiv \frac{1}{3(h+k)}$ and $c \equiv \frac{hA+kB}{6(h+k)^2}$ and let T be a random $U[A, B]$ variable giving the opponent's type. Noting that the tie-breaking case ($a = a'$) happens with zero probability¹⁶ given that (2) is a continuous function of a uniform random variable, we write the expected utility for an agent of type t playing action a as

$$\begin{aligned} \text{EU}(t, a) &= E_T[u(t, a, T, mT + c)] \\ &= E[t - ha - k(mT + c) \mid a > mT + c] \Pr(a > mT + c) \\ &\quad + E[h(mT + c) + ka \mid a < mT + c] \Pr(a < mT + c) \\ &= E \left[t - ha - kmT - kc \mid T < \frac{a-c}{m} \right] \Pr \left(T < \frac{a-c}{m} \right) \\ &\quad + E \left[hmT + hc + ka \mid T > \frac{a-c}{m} \right] \Pr \left(T > \frac{a-c}{m} \right) \end{aligned} \quad (3)$$

We consider three cases on the range of a and find the optimal action a_i^* for each case i .

Case 1: $a \leq Am + c$. ($\implies \frac{a-c}{m} \leq A$)

The probabilities in (3) are zero and one, respectively, and so the expected utility is:

$$\text{EU}(t, a) = hm \frac{A+B}{2} + hc + ka.$$

¹⁵This assumes that the space of possible bids includes ∞ . More generally, the dominant strategy is the supremum of the bid space but if this is not itself a member of the bid space (as is the case if the bid space is \mathbb{R}) then there is in fact no Nash equilibrium of the game.

¹⁶In the same sense that a real random variable with support of measure greater than zero has zero probability of occurring at a prespecified value in its support.

This is an increasing function in a , implying an optimal action at the right boundary: $a_1^* = Am + c$. Thus the best expected utility for case 1 is

$$\text{EU}(t, a_1^*) = \frac{2A + B}{6}.$$

Case 2: $a \geq Bm + c$. ($\implies \frac{a-c}{m} \leq B$)

The probabilities in (3) are one and zero, respectively, and so the expected utility is:

$$\text{EU}(t, a) = t - ha - km \frac{A + B}{2} - kc.$$

This is a decreasing function in a , implying an optimal action at the left boundary: $a_2^* = Bm + c$. Thus the best expected utility for case 2 is

$$\text{EU}(t, a_2^*) = t - \frac{A + 2B}{6}.$$

Case 3: $Am + c < a < Bm + c$.

Knowing that $\frac{a-c}{m}$ is between A and B it is straightforward to compute the probabilities in (3) and the conditional expectation of T . So we write $\text{EU}(t, a)$ as:

$$\begin{aligned} & \left(t - ha - km \frac{A + \frac{a-c}{m}}{2} - kc \right) \left(\frac{a-c}{m} - A \right) \\ & + \left(hm \frac{B + \frac{a-c}{m}}{2} + hc + ka \right) \left(B - \frac{a-c}{m} \right) \\ = & (-108a^2h^4 - 432a^2kh^3 - 648a^2k^2h^2 - 432a^2k^3h - 108a^2k^4 + 36aAh^3 \\ & + 72ath^3 + A^2h^2 + 4B^2h^2 + 4ABh^2 + 72aAkh^2 + 36aBkh^2 - 36Ath^2 \\ & + 216akth^2 + 36aAk^2h + 72aBk^2h + 8A^2kh + 8B^2kh + 2ABkh \\ & + 216ak^2th - 60Akh - 12Bkh + 36aBk^3 + 4A^2k^2 + B^2k^2 \\ & + 4ABk^2 + 72ak^3t - 24Ak^2t - 12Bk^2t) / (24(h+k)^2). \end{aligned}$$

Since this is a concave function of a the maximum is where the derivative with respect to a is zero, that is (skipping the tedious algebra for which we used Mathematica):

$$\begin{aligned} \frac{\partial \text{EU}(t, a)}{\partial a} &= 0 \\ \implies a_3^* &= \frac{t}{3(h+k)} + \frac{hA + kB}{6(h+k)^2}. \end{aligned}$$

Since $A \leq t \leq B \implies Am + c \leq a_3^* \leq Bm + c$, a_3^* is in fact in the allowable range for case 3. The expected utility for case 3 is then

$$\text{EU}(t, a_3^*) = \frac{3t^2 + A^2 + B^2 + A(B - 6t)}{6}.$$

It now remains to show that neither $\text{EU}(t, a_1^*)$ nor $\text{EU}(t, a_2^*)$ is greater than $\text{EU}(t, a_3^*)$ for any t .

Since $t \geq A$ there exists a $\delta \geq 0$ such that $t = A + \delta$. And since $B \geq A + 1$ there exists an $\varepsilon \geq 0$ such that $B = A + 1 + \varepsilon$. First, $\text{EU}(t, a_3^*) \geq \text{EU}(t, a_2^*)$ because

$$\begin{aligned}
& (\delta - 1)^2 \geq 0 \\
\implies & \delta^2 - 2\delta + 1 \geq 0 \\
\implies & \delta^2 + 1 \geq 2\delta \\
\implies & (A + \delta - A)^2 + 2A + 1 \geq 2A + 2\delta \\
\implies & (t - A)^2 + 2A + 1 \geq 2t \\
\implies & t^2 + A^2 + 2A + 1 \geq 2At + 2t \\
\implies & 3t^2 + 3A^2 + 6A + 3 + (3A\varepsilon + \varepsilon^2 + 4\varepsilon) \geq 6At + 6t \\
\implies & 3t^2 + A^2 + (A^2 + 2A + 2A\varepsilon + \varepsilon^2 + 2\varepsilon + 1) + (A^2 + A + A\varepsilon) - 6At \\
& \geq 6t - A - 2A - 2 - 2\varepsilon \\
\implies & 3t^2 + A^2 + (A + 1 + \varepsilon)^2 + A(A + 1 + \varepsilon) - 6At \geq 6t - A - 2(A + 1 + \varepsilon) \\
\implies & 3t^2 + A^2 + B^2 + AB - 6At \geq 6t - A - 2B.
\end{aligned}$$

Finally, $\text{EU}(t, a_3^*) \geq \text{EU}(t, a_1^*)$ because

$$\begin{aligned}
& (t - A)^2 \geq 0 \\
\implies & t^2 - 2At + A^2 \geq 0 \\
\implies & t^2 + A^2 \geq 2At \\
\implies & 3t^2 + 3A^2 \geq 6At \\
\implies & 3t^2 + 3A^2 + (3A\varepsilon + \varepsilon^2 + \varepsilon) \geq 6At \\
\implies & 3t^2 + 3A^2 + 3A + 3A\varepsilon + \varepsilon^2 + \varepsilon - 6At \geq 3A \\
\implies & 3t^2 + (A^2 + A + \varepsilon) - 6At + (A^2 + 2A + 2A\varepsilon + \varepsilon^2 + 2\varepsilon \\
& + 1) + A^2 \geq 3A + \varepsilon + 1 \\
\implies & 3t^2 + A(A + 1 + \varepsilon) - 6At + A^2 + (A + 1 + \varepsilon)^2 \\
& \geq 2A + (A + \varepsilon + 1) \\
\implies & 3t^2 + AB - 6At + A^2 + B^2 \geq 2A + B. \quad \square
\end{aligned}$$

A.2 Proof of Theorem 3

First, note that to prove any fairness results for SGA, we must take $k = 0$. When $k > 0$ the payment made by the winner is a function of both a and a' . Since a' is independent of the utility associated with the actual outcome, we can't make any guarantees on fairness when the payment is in part a function of a' . In other words, to achieve fairness we want to only consider the utility of the actual outcome, namely t —the type of the winner.

Let's look at the degree of fairness of an SGA($h, 0$) auction. Recall that degree of fairness is $\frac{n \cdot u_{\min}}{\sum_i u_i}$ where n is the number of participants and u_i the utility of each participant for the final outcome. In SGA(h, k) the degree of fairness is $2(ha + ka')/t$, in our case $k = 0$, so $2ha/t$. We know that agents are bidding strategically according to the equilibrium given in Theorem 1, so we know that a is a function of the winning agent's type, t , and h . Substituting this into the equation for degree of fairness, we get:

$$\frac{2h(t/3h + A/6h)}{t} = \frac{2t + A}{3} \cdot \frac{1}{t} = \frac{2}{3} + \frac{A}{3t}$$

We see that the degree of fairness is independent of h and is roughly $2/3$ with some additional additive function of A and t . Note, the only time the auction is completely fair, is when $t = A$. Recall that A is the lowest possible type for any agent. For an agent with lowest possible type to win, the other agent must also have had type A , since in general $t \geq A$, and if the losing agent had even a slightly higher value, it would have won the auction. \square

A.3 Proof of Theorem 4

In order to show that truthful bidding is an ε -equilibrium, we want to show that an agent stands to gain no more than ε by playing their best response to truthful bidding.

Since the two-player case of DAUC reduces to SGA($0, 1$) we can use Equation 3 from the proof of Theorem 1 to find the expected utility of each strategy. First note that the theorem gives three cases over the range of a , however, we notice also that cases one and two tell us that a will fall into the range examined in case three. Therefore we can take the simplified version of the equation given in Case 3:

$$\begin{aligned} EU(t, a) = & \left(t - ha - km \frac{A + \frac{a-c}{m}}{2} - kc \right) \left(\frac{a-c}{m} - A \right) \\ & + \left(hm \frac{B + \frac{a-c}{m}}{2} + hc + ka \right) \left(B - \frac{a-c}{m} \right). \end{aligned} \quad (4)$$

For an opponent strategy of truthful bidding, $m = 1$ and $c = 0$, and since $h = 0$ and $k = 1$ the above becomes:

$$EU(t, a) = \left(t - \frac{A+a}{2} \right) (a - A) + a(B - a). \quad (5)$$

We want to compare the expectation over all t for two specific actions, namely truthful bidding ($a = t$), and the best-response strategy ($a = 5t/6$) given by Reeves [2005].

We integrate with respect to t over the interval $[0, B]$ to find that $EU(t, t) = \frac{B^3}{3}$, and that $EU(t, 5t/6) = \frac{25B^3}{72}$. You lose $\frac{B^3}{72}$ by bidding truthfully. \square

A.4 Proof of Theorem 5

As above we use case 3 version of Equation 3 from Theorem 1 with an opponent strategy of truthful bidding. Since $h = 1/2$ and $k = 0$ in the two-person case of NFA, Equation 3 becomes:

$$EU(t, a) = (t - a/2)(a - A) + \left(\frac{B + a}{4}\right)(B - a). \quad (6)$$

We compare the expectation over all t for truthful bidding ($a = t$), and the best-response strategy given by Reeves [2005] ($a = t/2$) by integrating with respect to t over the interval $[0, B]$. We find that $EU(t, t) = \frac{B^3}{3}$, and that $EU(t, t/2) = \frac{17B^3}{48}$. You lose $\frac{B^3}{48}$ by bidding truthfully. \square

B API and SMS Interface to Decision Mechanisms

We have implemented the mechanisms described in this paper as an SMS application, which is a direct mapping to the underlying yootles API. The current interface can be used for mechanisms in their sealed-bid or iterative forms. In the latter case participants simply resubmit bids in light of partial results. See <http://yootopia.org> for the most up-to-date documentation of the system.

B.1 Keywords

yootles

Responds with a welcome message and list of available commands.

yhelp

Responds with a menu for more detailed help.

yreg NAME

Allows a user to register their phone. This is optional and only has the effect of referring to the user as NAME instead of their phone number in subsequent messages from the server.

ydel TAG

Deletes the auction or wager called TAG. If TAG is “phone” this command has a special meaning...

ydel phone

Disconnects a user from their current phone number. This allows the user to connect their account to a different phone number, and allows others to use the old phone number.

ydel TAG OPTION

Deletes the given option in the auction called TAG.

ybid TAG (OPTION AMOUNT)*

Allows a participant to cast a vote or place a bid. The TAG parameter refers to a previously created auction, or an auction called TAG is created if it doesn't exist. The server replies to each bid with intermediate auction results. Any number of OPTION AMOUNT pairs may be specified, giving the participant's bids for each of the named options. Options specified in a participant's bid are automatically appended to the list of options for the vote. Any unspecified option defaults to a zero bid and all bids are renormalized by a translation constant so that the minimum bid is zero (i.e., from each bid subtract the minimum bid, assuming more than one option). Bid commands may be reissued in which case they override the previous amount (and thus erroneous options can be voided by repeating them with a zero bid).

If no OPTION AMOUNT pairs are specified, the null bid serves simply to ping the server for the latest results.

Finally, if no option is specified and only one amount, the option defaults to the sender's own name (if registered) or phone number. This simplifies the interface in the common case that the list of options and the list of participants coincide. Note also that for the case of "drawing straws" each participant can place a negative bid for themselves which is equivalent to a positive bid of the same magnitude for everyone else.

ywager TAG PROBABILITY AMOUNT

Places a bet in wager TAG (which has possible outcomes of "yes" and "no"), with a specified probability of "yes" between 0 and 1, and a worst-case amount that the participant is willing to lose. (The YWAGER command is now also generalized to support any number of possible outcomes.)

yreset TAG

Purges all the options and bids/bets from auction or wager TAG.

ypay TAG

Confirm the payments specified in auction TAG's results and transfer them on the yootles ledger.

ypay TAG OUTCOME

Resolves a wager called TAG with the specified outcome (must be either "yes" or "no"), and credits or debits participants the appropriate number of yootles based on their bets and probabilities.

ypay AMOUNT RECIPIENT COMMENT*

An alternative use of the **ypay** keyword simply transfers the given amount of yootles to the recipient on the ledger. (If the amount is negative this is a

payment in the opposite direction.) Anything typed after `RECIPIENT` goes in the comments field for the yootles transaction.

`ycred` `USER` [`AMOUNT`]

Extend an amount of credit to a user (by default 0.02).

`ymech` `TAG` `MECH` `PARAMETERS*`

The `ymech` keyword sets the decision mechanism to be used for auction `TAG`. The default option for `YMECH` is `dauc` for the decision auction. Following is a fuller list of mechanisms coming soon. In all of them, the winning option is the one with the greatest sum after the bids (votes) are interpreted.

`dauc` Decision Auction—either the sealed-bid version (Section 3.6) or iterative version (Section 3.9) depending on whether participants resubmit bids in response to partial results.

`nfa` Nominally Fair Auction (Section 3.7).

`apr` Approval Voting—For each participant their positive bids are mapped to one and their nonpositive bids are mapped to zero.

`bor` Borda Count—With n options, each participant's bids are mapped to a number of points with n points to their most preferred option and 1 to their lowest (ties broken arbitrarily). The option with the highest sum of points wins.

`rng` Range voting—In Range Voting, voters' bids are renormalized to range from 0 to 1. As in all mechanisms, the default bid for all options is 0. The option with the highest total is chosen.

`irv` Instant Runoff Voting—In Instant Runoff Voting, if no option gets a simple majority (more than half the participants rated it highest) then the option which the least number of people ranked highest is eliminated (if your ranking was A-B-C and A is eliminated then your new ranking is B-C) and the process repeats until a simple first-choice majority winner is found.

`wei` Weighted Voting—Each participant's bids are adjusted translationally so the lowest bid is zero and then rescaled so they sum to one. This is a generalization of plurality voting in which an agent may divide its one vote among any number of options.

`jpa` Joint-Purchase Auction (Section 3.7).

`fav` Favor Auction—A generalization of Bilateral Trade, described in Section 3.7. This is a front-end for NFA. The buyer places a positive self-bid and the seller a negative self-bid. The sale price, if the bids overlap, is the average of the two bids.

`exch` An exchange, or call market mechanism.

C Yootles Questions, with Answers

1. What prevents someone from bidding a million yootles to get their way?

You can only bid as many yootles as you have plus whatever credit line other users have extended to you.

2. Will rich participants be able to walk all over the poor, by buying yootles with money?

Different groups may evolve different conventions. One such convention may be a general understanding that it's not kosher to buy yootles with dollars, either directly or indirectly. This of course is not strictly enforceable which is why it could only be loosely enforced in the form of a social taboo. If yootles are purchasable with money, yootling serves as a means for wealth redistribution.

3. What if I'm generous with yootles and my friend is stingy?

Being generous with yootles violates the spirit of yootles. Yootles are a measure of your utility. Be generous in your actions, not your yootle transfers. Remember that a yootle that you give away is in essence an IOU for a favor done, so being generous with your yootles is like being generous in accepting favors from someone else. A person who is stingy with yootles is someone who is generous with their time, or with giving way to others' desires.

4. Are yootles susceptible to inflation? And how do I judge the value of a yootle?

Yes, if an entity goes more negative than it can make good on, this causes inflation for the same reason that the government printing excess money causes inflation. And, like money, yootles are susceptible to inflation in the sense that if people come to expect, say, an hour of cooking to fetch a million yootles then bids will inflate accordingly. In other words, price bubbles can form. Any number of things can ground the value of a yootle, including another currency. This is possible even if money and yootles are never exchanged. The expectation (along with some adjustments by the central bank) could keep the value of a yootle grounded (insomuch as your local currency is grounded). Another way to ground the value of a yootle is to make it worth an hour of unskilled labor. This is the rationale for a local currency in Ithaca, New York: Ithaca Hours (although in reality, Hours are grounded strictly in dollars).

5. Should there be an interest rate?

Yes, humans, being mortal, have a fundamental discount factor—utility now is better than utility later. The current incarnation of the yootles ledger sets a universal (but modifiable) interest rate of 5.375%. The upcoming version of the ledger will allow any pair of participants to choose a custom interest rate (including zero or even negative), and in fact multiple interest rates for different chunks of credit they extend to each other.

6. How about letting people get out of yootle debt by donating to charity?

That’s a very nice idea (thanks to Dave Pennock), but it’s not a perfect solution. The charity would have to itself be a yootles account which would sell yootles to anyone who’s negative and use the cash for charity. The charity account would be exempt from the requirement that it not stay negative indefinitely. However many negative yootles the charity account has is how many positive dollars the group has donated. The downside is that the further negative the charity account goes, the more inflated the yootles currency is. Say I’m feeling philanthropic and pay 1000 yootles in some auction which I later replace from the charity account. That’s 1000 yootles I flooded the group with and a yootle will be devalued accordingly. It also may undermine some good incentive properties: you now want to make me pay more, even if I’m not paying you, because if I go negative I’ll probably end up paying a charity that you presumably care about.

7. Why is yootle inflation bad?

- You lose the property that a yootle is worth a dollar, which may make it harder to quantify your preferences.
- It’s harder to learn about your preferences over time when the value of a yootle is changing.
- It destroys a nice application of the yootles system: tracking debts (small or large) of real money between members of the group.
- And, most generally, it hurts social efficiency.

8. How can you ensure that a yootle will always be worth a dollar?

If yootles become more (or less) valuable, the central bank can inject yootles into the system (or tax them out).

9. I’m an altruist. I get utility (“yootility”) from increasing others’ utility. What happens with me in the system?

We would try to change your worldview and get you to trust in the Invisible Hand. Presumably what you really want is to maximize social welfare—not simply to maximize everyone’s utility intentionally at your

own expense. Participating sincerely in a yootles-based system is a better way to achieve that than you operating on your assumptions about others' utility functions. Once people are accustomed to expressing their utility functions it's amazing how much more effectively they can increase each other's utility.

10. What is the connection to RipplePay.com?

Our new ledger system was inspired by RipplePay. RipplePay begins with the insight that logging an IOU on a ledger is in a very real sense the same as paying someone. They have run with this idea and have implemented an alternative to the current yootles ledger system in which all debts are tracked pairwise between individuals who trust each other.