Nonmanipulable collective decision-making for games

Rob LeGrand
Angelo State University, San Angelo, Texas, U.S.A.
Timothy Roden
Angelo State University, San Angelo, Texas, U.S.A.
Ron K. Cytron
Washington University, St. Louis, Missouri, U.S.A.

ABSTRACT

This chapter explores a new approach that may be used in game development to help human players and/or non-player characters make collective decisions. We describe how previous work can be applied to allow game players to form a consensus from a simple range of possible outcomes in such a way that no player can manipulate it at the expense of the other players. We then extend that result and show how nonmanipulable consensus can be found in higher-dimensional outcome spaces. The results may be useful when developing artificial intelligence for non-player characters or constructing frameworks to aid cooperation among human players.

INTRODUCTION

Teamwork is important in many games. Whether they are human players or non-player characters (NPCs) or both, game entities must often work together to achieve goals. However, those goals do not always coincide perfectly, and, even when they do, players will not always agree on the best next course of action to take.

Much research (see especially Rabin 2002, sec. 7) has explored effective decision-making for individual game agents, even in a multiagent context. By contrast, in this work, we assume that all agents have already individually decided which of the available outcomes (which are usually actions) they prefer over others, and we assume the agents desire to use those preferences to reach consensus for the group.

For an example game situation, imagine a team of wargame players with a common goal: band together to attack the western coast of a continent held by a common enemy. They could attack the coast's northernmost point, the southernmost point or anywhere in between, and each player has a different favorite attack point. If the players can be trusted to express sincere preferences, their preferred points could simply be averaged to give the consensus point. Averaging, however, can allow some players to gain a better outcome from their point of view by exaggerating their preferences, whereas other aggregation mechanisms may never reward such insincerity.

When a group of players aims to benefit all members of the group by coordinating their actions, a method of combining their preferences into a single outcome is useful, but the usefulness may disappear if individual players can manipulate the outcome by expressing insincere preferences. Here we present a set of nonmanipulable collective-decision-making methods that apply to a wide range of game situations.

In the sections below we review previous work that informs ours, look at several game situations that motivate our approach and present the ideas that provide an innovative solution.

BACKGROUND IDEAS

The core ideas of this chapter, while new, are based in extant work from fields such as computer science, mathematics, political science and economics.

Mechanism design

Returning to the above wargame example, if a team of players is trying to agree on a coastal attack point, their preferred points could simply be averaged to give the consensus point, but doing so sometimes rewards insincerity on the parts of the players. The field of *mechanism design* (Nisan 2007) has evolved to find decision-making mechanisms that satisfy particular properties, often some kind of immunity or resistance to strategic manipulation.

Strategic manipulation is a common problem in collective decision-making. It is well known that voters can gain advantage under most voting systems by voting insincerely (Gibbard 1973, Satterthwaite 1975). Examples include voting for an alternative that is not a voter's first choice and ranking alternatives untruthfully. Traditionally, this problem is discussed in political science, but more recently the techniques of computer science have been applied with success (Bartholdi, Tovey & Trick 1989, Conitzer & Sandholm 2003, Elkind & Lipmaa 2005, Procaccia & Rosenschein 2006). In this chapter we explore a particular approach to creating manipulation-resistant mechanisms.

The Declared-Strategy Voting framework

Declared-Strategy Voting (DSV) is a computationally-based response to manipulable voting systems (Cranor & Cytron 1996, Cranor 1996). Under DSV, each voter submits preferences over the available outcomes. The DSV system then uses those preferences to vote optimally (and, perhaps, insincerely) on each voter's behalf in a simulated election using some underlying voting method. It continues to cast optimal ballots on behalf of each voter until an equilibrium is found or some other stopping criterion is reached. The outcome at equilibrium is then taken as the DSV outcome, or the results of the voting rounds could be used by a policy-maker to reach a justifiable decision.

The hope is that, since the DSV system is attempting to vote strategically on each voter's behalf, no voter will have a reason to mislead the system by expressing insincere preferences—in fact, an attempt to mislead the system may easily backfire. DSV has been shown to be effective in transforming some manipulable voting systems into manipulation-immune ones with the same available outcomes.

A previous result: AAR DSV

In past work, we have successfully applied DSV to an average-approval-rating (AAR) system, which takes voters' ratings between 0% and 100% as input and outputs their average (LeGrand 2008, LeGrand & Cytron 2008). Such systems are widely used for rating movies, music, and buyer/seller reliability on the Internet on websites such as Amazon, Metacritic and Rotten Tomatoes; similar systems are used for many diverse applications. We assume only that each voter has an ideal outcome between 0% and 100% and prefers the outcome to be as close to that ideal as possible. So, for example, a voter whose ideal outcome is 20% cannot prefer 40% to 30%. We found that the resulting system had some surprising and attractive properties: Given reasonable assumptions, optimal voting strategy is unique and can be fully characterized, and, given any input ratings, the outcome at equilibrium is unique. We discuss these results in more detail below.

Most importantly, we were able to prove that no AAR DSV voter can achieve an outcome closer to ideal than by voting sincerely. As an example, imagine a group of three game players who are deciding how much of a collective resource to use towards an immediate goal, such as how much of their magic store to employ against the next major enemy, or how much of a healing pack to use before a fight. Say that the three players have sincere preferences [25%, 40%, 70%] and express them sincerely to an AAR DSV system. After DSV iteratively applies optimal strategies on behalf of each player, the equilibrium becomes [0%, 20%, 100%], giving the outcome 40%. Now, if the first player, somewhat dissatisfied with this outcome, had expressed the preference 0% instead of 20% in an effort to pull the outcome closer to 20%, the DSV equilibrium and outcome would be the same as above; in essence, the DSV system is already manipulating on behalf of each player, so misleading the DSV system is unnecessary and fruitless. For this AAR voting system, the DSV framework perfectly internalizes voting strategy—players can never gain advantage by exaggerating their position.

If this meta-voting system is to be used to find consensus in real games, it will be important to be able to calculate the outcome quickly. Fortunately, it turns out that there is an efficient algorithm to calculate the AAR DSV outcome, one that is not significantly slower than simply sorting the numerical inputs of the players.

AAR DSV offers a way for agents to find consensus within a numerical range (a line segment) without the possibility of successful manipulation by selfish agents. Many game situations, however, require finding consensus within other outcome spaces. We will see that the DSV framework is flexible enough to work well with other useful outcome spaces.

A NEED: FINDING COLLECTIVE CONSENSUS IN GAMES

Some games are purely competitive and have no room for cooperation and thus no need to find consensus among a group of players. But more common are games that feature multiple agents who are, at least under some circumstances, motivated to cooperate to the advantage of all, and the rise of online and social gaming is making these situations more common. Agents in games include both human players and computer-controlled NPCs, and opportunities for cooperation can emerge among human players, among NPCs, and even between humans and NPCs. Frequently, to cooperate most efficiently, these groups of agents need to come to some sort of explicit consensus. We will give three example situations from real games that illustrate benefits of our approach, each with a different outcome space.

Finding consensus in one dimension

Often game agents need to come to a numerical consensus within the space of a line segment, the endpoints of which may be thought of as 0% and 100%. Many times, agents may need to agree on how much, from 0% to 100%, of a common resource to use at a given time. As another, more concrete, example, imagine a game like Sid Meier’s Colonization, in which a single colony may have many cities, each individually controlled by either a human player or an NPC. One of the most important facets of the game is deciding when your colony should revolt against its parent country, and a colony’s players may disagree on the best time to revolt. Triggering a successful revolt is based on the average level of discontent for the entire colony, so it is sometimes desirable to build up a measure of discontent in individual cities. At any given time, then, each player may have a different ideal target for the colony-wide average level of discontent, and the player controlling a given city has exclusive control over the level of discontent associated with that player’s city.

As an example, consider a colony with three cities, each controlled by a different player. The three players prefer that the colony-wide average level of discontent be [25%, 40%, 70%]. The players could manipulate their own cities’ discontent levels to move the colony-wide average towards their ideals, struggling against one another, but doing so would require sacrifices to be made in other ways. Alternatively, they could “vote” their ideal colony-wide averages and agree to abide by the results, thus requiring smaller changes to each city's discontent level. If the system that resolves input ratings into an outcome is chosen carefully, it will never reward insincerity and the players will be able to find consensus with confidence. The AAR DSV system, introduced above and explored in more detail below, fits the bill: It takes a vector of ratings between 0% and 100% as input, outputs a consensus rating, and is impossible for a voter to manipulate through insincere voting. If the input vector were [25%, 40%, 70%] as in the example above, the DSV equilibrium would be [0%, 20%, 100%], giving the outcome 40%, which is the ideal outcome for one of the players.

Finding consensus within a hypercube

In other game situations, agents need to find consensus inside a hypercube of some dimensionality; each dimension can be seen as ranging between 0% and 100%. One motivating example comes from Age of Empires, a real-time strategy game that can accommodate many human players and many NPCs. Several human players working together may need to decide collectively the location for their newest base. Every point in the game's square map has two coordinates, one ranging between the west edge, 0%, and the east edge, 100%, and one between the south edge, 0%, and the north edge, 100%. Each player has a most-preferred point for the base and would like to see the base built as close to that ideal point as possible. Players may be assumed to prefer points closer by Euclidean distance to their ideal point to points farther from it, so a player with the ideal point (60%, 10%) can be assumed to prefer that the base be built at (65%, 20%) than at (50%, 30%). When a nonmanipulable system is used to aggregate input points into an outcome point, it will not be possible for the player with ideal (60%, 10%) to change the outcome from (50%, 30%) to (65%, 20%) by changing their vote point from the sincere (60% 10%) to any other point.

Finding consensus within a simplex

Another common kind of game decision to make collectively involves the allocation of a finite resource among several uses. Essentially, the agents must agree on a point within a simplex of some dimensionality, where each dimension ranges between 0% and 100% and the values chosen for each dimension sum to exactly 100%. For example, in a space war game such as Star Trek Online, the officers on a starship must often decide how to deploy its limited energy resources; shields, weapons and maneuvering all require energy. The ship's officers may be controlled by separate AI programs, which may sincerely disagree on the policy best for the ship. At a given time, one AI officer may estimate (30%, 60%, 10%) as the best policy, giving 30% of the available energy to shield systems, 60% to weapons and 10% to the helm, while another may prefer (50%, 50%, 0%). Each officer may be assumed to prefer points nearer to ideal in each dimension to points farther in each dimension; for example, we may assume that an officer that most prefers (30%, 60%, 10%) will prefer the outcome (40%, 40%, 20%) to (50%, 30%, 20%). Even when all relevant agents are controlled by a computer, it may be that a nonmanipulable consensus-finding system is preferred; for the sake of realism, AI agents may be programmed to take advantage of manipulable systems whenever possible (unless they are part of a hive collective such as the Borg).

The above examples are inspired by currently available games, but we believe social cooperative games are potentially the best-suited games for our method. Most of these games are still in development as designers and game programmers figure out how to implement effective cooperative strategies among players that play non-concurrently; our techniques do not require players to indicate their preferences simultaneously, and so they may prove especially useful for this kind of game.

USING DSV TO FIND STABLE CONSENSUS IN ONE DIMENSION

First we consider the simplest outcome space that players may need, which is a line segment; we will assume that the available outcomes range between 0 (0%) and 1 (100%). We further assume:

- There are n players that want to find a consensus in the inclusive interval from 0 to 1.
- Each player i submits a value v_i in the inclusive interval between 0 and 1; the resultant vector $\vec{v}$ is used to determine the consensus outcome.
- Each player i has an ideal outcome r_i and prefers that the outcome be as near to r_i as possible.

This last assumption, that of *single-peakedness* (Moulin 1980), is required for the conclusions reached below, but it is important to point out that it rules out certain preference orderings. We are assuming, for example, that an agent that prefers an outcome of 0.2 to 0.3 cannot also prefer 0.4 to 0.3. Still, it is an imminently reasonable assumption for many game situations.

Average aggregation

Perhaps the most natural outcome function to use is the average of the inputs, which minimizes the sum of squared distances between the outcome and the inputs. While the Average aggregation function is sensitive to each voter's input, it has an important disadvantage: Voters can often gain by voting insincerely. For example, if $n = 3$, $\vec{r} = [0.25, 0.4, 0.7]$ and all three players express their sincere preferences, then $\vec{v} = [0.25, 0.4, 0.7]$ and the Average outcome is 0.45. Consider player 3, whose ideal outcome is $r_3 = 0.7$. That player could achieve a better outcome by not expressing the sincere preference $v_3 = 0.7$ and instead choosing $v_3 = 1$. The resulting Average aggregation yields the outcome 0.55, which, being closer to 0.7, is preferred by player 3 to 0.45.

Rationally optimal strategy for Average aggregation

Using the Average outcome opens the door for manipulation, but investigating the nature of that manipulation further will prove fruitful. If we assume that all players want only to optimize the outcome from their own points of view, we can characterize rational voting. If all other players have expressed their preferences and player i is deciding how to choose v_i, the ideal outcome r_i could be achieved by choosing $v_i = r_i n - \sum_{j \neq i} v_j$, but this choice for v_i is allowed only if it is between 0 and 1. In general, player i can move the Average outcome as near to r_i as is possible by choosing

$v_i = \min\left(\max\left(r_i n - \sum_{j \neq i} v_j, 0 \right), 1 \right)$, which is the rationally optimal strategy for any player i.

An example will illustrate that, if all players use this strategy iteratively, an equilibrium may be reached from which no player would change. Imagine that the three players from above with sincere preferences $\vec{r} = [0.25, 0.4, 0.7]$ begin by expressing $\vec{v} = [0, 0, 0]$. If they then rationally adjust their strategies in order of descending ideal preferences, they would calculate as follows.

$$\text{Calculate: } v_3 = \min\left(\max\left(r_3 n - \sum_{j \neq 3} v_j, 0 \right), 1 \right) = \min(\max(2.1 - 0, 0), 1) = 1$$

$$\text{Update: } \vec{v} = [0, 0, 1]$$

$$\text{New outcome: } \bar{v} \approx 0.333$$

$$\text{Calculate: } v_2 = \min\left(\max\left(r_2 n - \sum_{j \neq 2} v_j, 0 \right), 1 \right) = \min(\max(1.2 - 1, 0), 1) = 0.2$$

$$\text{Update: } \vec{v} = [0, 0.2, 1]$$

$$\text{New outcome: } \bar{v} = 0.4$$

$$\text{Calculate: } v_1 = \min\left(\max\left(r_1 n - \sum_{j\neq 1} v_j, 0\right), 1\right) = \min(\max(0.75 - 1.2, 0), 1) = 0$$

$$\text{Update: } \vec{v} = [0, 0.2, 1]$$

$$\text{New outcome: } \bar{v} = 0.4$$

After one pass through the players, an equilibrium from which no player would change has been found, and the Average outcome at this equilibrium is $\bar{v} = 0.4$. Player 1 would prefer a smaller outcome, but v_1 is already as small as is allowed; player 3 would prefer a larger outcome, but v_3 is already as large as is allowed; player 2 has set v_2 exactly where it must be to achieve the outcome $\bar{v} = r_2 = 0.4$. Thus Average aggregation is manipulable by strategic players willing to submit insincere preferences.

But strategic manipulation may not be so undesirable if an equilibrium can always be found as rapidly as in the above example whatever the players' sincere preferences. Given a set of *n* players and their sincere preferences $\vec{r}$, LeGrand (2008) shows, by counting the number of players that must be strategizing at each of the two extremes, that any average $\bar{v}$ at equilibrium must satisfy two inequalities:

$$\left|\{i : \bar{v} < r_i\}\right| \leq \bar{v}n$$

$$\bar{v}n \leq \left|\{i : \bar{v} \leq r_i\}\right|$$

LeGrand (2008) then proves that:

- At least one equilibrium $\vec{v}$, at which $(\forall i)\, v_i = \min\left(\max\left(r_i n - \sum_{j\neq i} v_j, 0\right), 1\right)$ (so that no player *i* would be motivated to change v_i unilaterally), must exist. (The proof shows that the same algorithmic approach as seen in the example above will always find an equilibrium in one pass, which requires showing that no player *i*, after choosing v_i, would want to change its value later in the pass.)
- Multiple different such equilibria may exist, but all such equilibria must have the same Average outcome. (The proof shows by contradiction that any two averages at equilibrium, which must satisfy the above two inequalities, must be equal.)

It follows that, given the vector of sincere preferences $\vec{r}$ as input, the equilibrium Average outcome is unique and can be defined as a mathematical function.

Average-Approval-Rating DSV

The DSV framework allows players to express their ideal outcomes, then "votes" for them iteratively until a stable outcome emerges. So, applied to the Average system discussed above, players would input their preferences and the DSV system would simulate the same rationally strategic voting illustrated above, reliably giving a unique outcome. Doing so explicitly provides an effective AAR DSV algorithm, but LeGrand (2008) proves that, given a vector of expressed preferences $\vec{v}$, the AAR DSV outcome can be calculated yet more efficiently:

sort $\vec{v}$ so that $(\forall i \leq j)\, v_i \geq v_j$
$w \leftarrow 0$
for *i* = 1 to *n* do
    $w \leftarrow w + \min(\max(v_i n - w, 0), 1)$

return $\frac{w}{n}$

This algorithm runs in $O(n \log n)$ time if a $O(n \log n)$-time sort is used.

The most important property of this new AAR DSV system is that it cannot be manipulated by strategic players who are willing to submit insincere preferences, as the Average system can. LeGrand (2008) proves that AAR DSV never rewards insincerity: No player *i* can move the AAR DSV outcome closer to the ideal r_i by expressing a preference other than $v_i = r_i$; the proof entails first proving that:

- An AAR DSV outcome cannot be increased (respectively, decreased) without increasing (decreasing) at least one of the inputs. (The proof is by contradiction and uses the two inequalities that any average at equilibrium must satisfy.)
- If an AAR DSV outcome is smaller (respectively, larger) than one of the inputs, the outcome does not change when that input is increased (decreased). (The proof relies on the uniqueness of the average at equilibrium.)

Once these two points are proved, it is straightforward to show that no player *i* can gain (but may lose) from moving v_i away from sincerity. This nonmanipulability result is satisfying because it shows that the DSV framework, by using the inputs to "vote" on each player's behalf, perfectly internalizes all required strategy, allowing players to focus on more important matters than attempts to manipulate the consensus.

Therefore, if a collective decision is to be made inside a line segment, the AAR DSV approach provides a way to find a consensus in $O(n \log n)$ time without the possibility of any advantage gained by insincere players. Above we imagined a team of wargame players banding together to attack the western coast of a continent held by a common enemy. If they used AAR DSV to decide how far north or south to attack, none of the players would have any reason to express an insincere preference, and the players would be able to forget about outfoxing each other and concentrate on the attack itself.

USING DSV TO FIND STABLE CONSENSUS IN A HYPERCUBE

As discussed above, another potentially important outcome space for game players is a hypercube, essentially a cross product of line segments. For example, players might want to agree on a point inside the two- or three-dimensional game map to attack or at which to meet, or they may be making several independent 0%-to-100% decisions at once. This problem is essentially equivalent to making two, three or more 0%-to-100% decisions as a package; 0% might mean the west edge of a map and 100% might mean the east edge, etc. We will assume that the available outcomes in each of the *d* dimensions range between 0 (0%) and 1 (100%). We further assume:

- There are *n* players that want to find a consensus inside the hypercube of dimension *d*.
- Each player *i* submits a *d*-dimensional vector v_i, each scalar coordinate of which is in the inclusive interval between 0 and 1; the resultant vector $\vec{v}$ is used to determine the consensus outcome.
- Each player *i* has an ideal outcome r_i and prefers that, in each of the *d* dimensions individually, the outcome be as near to r_i as possible.

Instead of this last assumption, that of intradimensional single-peakedness, we could assume simply that each player *i* would like to minimize the Euclidean distance between the outcome and r_i, but that is actually a stronger assumption than we require. Fortunately, if each player aims simply to minimize the

Euclidean distance between the actual outcome and their ideal outcome, as arguably would often be the case, then this notion of dimension-independence holds.

However, our assumption does rule out certain preference orderings that might conceivably occur in a real game. For example, consider a game map over which a group of RPG players is deciding collectively where to meet. It may be that one player, player *i*, prefers to meet near a river that flows northeast to southwest, so if the consensus is to meet somewhere in the west, player *i* would prefer meeting points farther to the south, whereas if the consensus is to meet somewhere in the east, player *i* would prefer meeting points farther to the north. In other words, player *i* might prefer both (0, 0) to (1, 0) and (1, 1) to (0, 1). Such a player would not be able to isolate one ideal point r_i that satisfies our assumptions; only a richer input space would allow expressing such preferences. But our assumption is reasonable for many common game situations, including any in which the dimensions are completely independent.

Average aggregation is easily generalized to a higher-dimension hypercube by taking the average of each coordinate separately, effectively calculating the *centroid*, the center of mass given a set of unit masses. Alternatively, one can imagine attaching Hookean springs of equal spring constants to each fixed input point, then gluing the other ends of the springs together; the glue point will come to rest at the centroid. This generalization is equivalent to finding the point that minimizes the sum of squared distances between that point and all of the input points. The resulting system is rotationally invariant and is equivalent to conducting *d* separate and independent Average elections (LeGrand & Cytron 2008). Thus, the results above for strategic behavior under the one-dimensional Average system apply to the "election" for each coordinate. In particular, if each voter has *separable* preferences (Border & Jordan, 1983), so that preferences in one dimension are independent of preferences in all other dimensions, conducting a *d*-dimensional AAR DSV election is equivalent to conducting *d* parallel one-dimensional AAR DSV elections, and so gives a nonmanipulable system. (Such a preference-function space is not *abundant* by Zhou's (1991) definition.)

To illustrate briefly, consider two players with ideal preferences $\vec{r} = [(0.2,0.3),(0.6,0.1)]$. If they used Average aggregation and applied rational strategy iteratively, they would soon reach the equilibrium $\vec{v} = [(0,0.6),(1,0)]$, giving the outcome $(0.5,0.3)$. At this equilibrium they are pulling each coordinate of the outcome toward their ideal outcomes to the greatest extent possible, which in this case requires insincerity. If the players decided to use AAR DSV instead of Average, they could submit their sincere preferences, $\vec{v} = \vec{r} = [(0.2,0.3),(0.6,0.1)]$, giving the outcome $(0.5,0.3)$ directly, and the players would have no incentive to be insincere.

Therefore, if a collective decision is to be made inside a hypercube, the AAR DSV approach provides a way to find a consensus without the possibility of any advantage gained by insincere players: Simply use the one-dimensional AAR DSV in each of the *d* dimensions. Using the AAR DSV algorithm given above, a nonmanipulable consensus can thus be found in $O(dn \log n)$ time.

USING DSV TO FIND STABLE CONSENSUS IN A SIMPLEX

The AAR DSV approach gives satisfying results for the line-segment and hypercube outcome spaces, but there are other outcome spaces that may be useful in game situations. Another useful way to generalize the one-dimensional space between 0 and 1 into higher dimensions is into a simplex, as discussed above. For example, consider another decision situation, in which players decide how a fixed amount of a resource should be allocated among several uses, such as (30%, 60%, 10%). An AAR-style method would have each player suggest an allocation and then average them to give the outcome. We will assume that the available outcomes in each of the *d* dimensions range between 0 (0%) and 1 (100%); these coordinates of any one input point or outcome must sum to 1 (100%). We further assume:

- There are *n* players that want to find a consensus inside the simplex of *d* dimensions. (Because the coordinates must sum to 1, this simplex is mathematically a $(d-1)$-dimensional space, but we will refer to *d* dimensions for convenience.)
- Each player *i* submits a *d*-dimensional vector v_i, each scalar coordinate of which is in the inclusive interval between 0 and 1 and all of which sum to 1; the resultant vector $\vec{v}$ is used to determine the consensus outcome.
- Each player *i* has an ideal outcome r_i and prefers outcome *a* to outcome *b* whenever *a* is nearer to r_i than *b* is in at least one of the *d* dimensions and *a* is farther from r_i than *b* is in none of them.

Again, instead of this last assumption, we could assume simply that each player *i* would like to minimize the Euclidean distance between the outcome and r_i, but it would be a stronger assumption than we require.

In the hypercube space, it was possible to use one-dimensional AAR DSV in each dimension to arrive at a multidimensional consensus. Effectively, each dimension was completely independent of the others. But in the simplex space, the coordinates of a point restrict the allowed values of the other coordinates; in a sense, the dimensions fail to be independent by virtue of the shape of the space. For example, if one coordinate of a point has the value 0.4, then no other coordinate can have a value higher than 0.6.

Still, as before, we can characterize optimally strategic "voting" for Average aggregation within a simplex. As in the one-dimensional space, a rational player would express $v_i = r_i n - \sum_{j \neq i} v_j$ if it were inside the allowed simplex space. Otherwise, it must be projected to the (always uniquely) nearest point in the simplex, which may be on a vertex, an edge, a face, etc. For example, if $d = 3$, then the simplex is an equilateral triangle with vertices $(1,0,0)$, $(0,1,0)$ and $(0,0,1)$, and rational strategy can be characterized as follows: If we calculate $r_i n - \sum_{j \neq i} v_j$ and call its three coordinates *x*, *y* and *z*, then player *i* should express

$$v_i = \begin{cases} (1,0,0) & if\ x \geq y+1\ and\ x \geq z+1 \\ (0,1,0) & if\ y \geq x+1\ and\ y \geq z+1 \\ (0,0,1) & if\ z \geq x+1\ and\ z \geq y+1 \\ \left(x+\frac{z}{2}, y+\frac{z}{2}, 0\right) & if\ x \leq y+1\ and\ y \leq x+1\ and\ z \leq 0 \\ \left(x+\frac{y}{2}, 0, z+\frac{y}{2}\right) & if\ x \leq z+1\ and\ y \leq 0\ and\ z \leq x+1 \\ \left(0, y+\frac{x}{2}, z+\frac{x}{2}\right) & if\ x \leq 0\ and\ y \leq z+1\ and\ z \leq y+1 \\ (x,y,z) & if\ x \geq 0\ and\ y \geq 0\ and\ z \geq 0 \end{cases}$$

The rational strategy functions for higher-dimensional simplex spaces will have more cases but follow a similar pattern.

As an example of rationally strategic players finding consensus in a simplex, consider $n = 2$ players trying to decide on the best way to allocate an amount of magic points among $d = 3$ uses, such as to attack the current enemy, to heal the players and to save for future situations. The two players have sincere preferences $\vec{r} = [(0.5,0.5,0),(0.25,0.25,0.5)]$ and use Average aggregation. After repeatedly applying the rational strategy detailed above, they would soon reach the unique equilibrium $\vec{v} = [(0.5,0.5,0),(0,0,1)]$, giving the outcome $(0.25,0.25,0.5)$, which is player 2's ideal outcome. From this equilibrium, neither player can move the Average outcome nearer to ideal by Euclidean distance.

Using Average aggregation in the one-dimensional space, it was possible to prove the following properties (LeGrand 2008):

- A strategic equilibrium always exists.
- The equilibrium outcome is unique.
- The resulting DSV function is nonmanipulable by insincere agents.

The equivalent properties may yet be proved in the simplex space; until then, it will be useful to look for counterexamples via Monte Carlo simulations.

Experiments and results in a simplex

We wrote software using the C programming language to simulate many random "elections", checking each one to see whether the properties held. The setup was as follows:

- For each combination of *d* and *n* from $d \in \{3,4,5,6\}$ and $n \in \{2,3,5,8,13\}$, run one million simulated elections in *d* dimensions with *n* agents.
- For each election, do the following:
    - Randomly generate a vector $\vec{r}$ of *n* *d*-dimensional points, each coordinate of which has exactly p_g decimal places, from a uniform distribution in the simplex.
    - Taking the points in $\vec{r}$ as the agents' sincere preferences, iteratively vote strategically on behalf of each agent, using a randomized ordering of agents, until an equilibrium $\vec{v}$ is reached. Stop only when all coordinates of the points in $\vec{v}$ are stable to p_e decimal places.
    - Repeat the iterative strategy from scratch, again using a randomized ordering of agents, until an equilibrium is reached. Take note of an *outcome-uniqueness violation* if the Average outcomes at the two found equilibria differ to p_d decimal places in any coordinate.
    - Construct another vector $\vec{r}'$ of points such that r_1' is regenerated randomly as before and $r_i' = r_i$ for all $i > 1$, representing an attempt on the part of agent 1 to deceive the DSV system with insincere strategy, and find an equilibrium with the new $\vec{r}'$. Take note of a *successful Euclidean-distance manipulation* if the equilibrium Average outcome using $\vec{r}'$ is nearer, to p_d decimal places, in Euclidean distance to r_1 than is the equilibrium Average outcome using $\vec{r}$. Also take note of a *successful dimension-domination manipulation* if the equilibrium Average outcome using $\vec{r}'$ is nearer, to p_d decimal places, to r_1 than is the equilibrium Average outcome using $\vec{r}$ in at least one of the *d* dimensions and farther, to p_d decimal places, in none of them.

We ran the above simulations with the values $p_g = 4$, $p_e = 13$ and $p_d = 7$. We chose p_d larger than p_g so as to detect much smaller differences among outcomes than the differences among input points, and we chose p_e larger than p_d so as to ensure that we did not detect a difference among outcomes that was due only to the inaccuracy of the DSV outcomes; the choice of $p_e = 13$ was small enough for the DSV function to converge to an equilibrium quickly in practice. None of the simulated elections failed to find an equilibrium, and none found two equilibria with different Average outcomes. While by no means a proof, this result leads us to suspect simplex AAR DSV to be a mathematical function of its inputs, meaning that, given the same inputs, it will always give the same outcome.

On the other hand, we found very quickly that simplex AAR DSV *is* Euclidean-distance-manipulable by insincere agents; Table 1 shows how often the attempted manipulations were successful in our simulations. In fact, a manipulation opportunity exists for the simplex example given above. If both players submit sincere preferences, the DSV system will apply rational strategy on the part of both players, find the unique equilibrium $[(0.5,0.5,0),(0,0,1)]$ and give the outcome $(0.25,0.25,0.5)$, player 2's ideal outcome. But if player 1, instead of the sincere $(0.5,0.5,0)$, insincerely expresses the preference $(0.2,0.6,0.2)$, the DSV system would instead find the unique equilibrium $[(0,1,0),(0.25,0,0.75)]$. The resulting outcome, $(0.125,0.5,0.375)$, is only 0.28125 away from player 1's ideal in Euclidean distance, an improvement from the sincere outcome, which was 0.375 away. Player 1 has thus successfully manipulated simplex AAR DSV, at least by Euclidean distance.

However, while this manipulated outcome is closer to player 1's ideal outcome in the second and third dimensions (0 *vs.* 0.25 and 0.375 *vs.* 0.5, respectively), it is farther from ideal in the first dimension (0.375 *vs.* 0.25). In fact, our simulated elections found no example of a successful dimension-domination manipulation (Table 2). This result, while again no proof, suggests that the simplex AAR DSV system may be immune to strategic manipulation at least in a weak sense: It may be impossible for an agent to use insincerity to move a simplex AAR DSV outcome to one which is strictly better, *i.e.*, closer to ideal on at least one dimension and farther on no dimension. In other words, from any one agent's point of view, insincerity that improves the outcome in one dimension must make it worse in another dimension.

This narrower notion of immunity to strategic manipulation, which essentially makes strong assumptions about preferences voters may have, avoids previous impossibility results (Moulin 1980, Zhou 1991). Even if this assumption about players' preferences would not always hold in real game situations, this simplex AAR DSV method does seem to be at least quite resistant to strategic manipulation.

Table 1. Incidence of successful Euclidean-distance manipulation in simplex simulation

d \ *n*	2	3	5	8	13
3	6.4109%	4.2370%	4.0829%	3.6739%	3.4174%
4	4.7949%	7.0336%	5.7254%	5.3783%	5.0183%
5	3.1722%	6.3082%	6.6661%	6.1147%	5.7822%
6	2.0793%	4.9974%	7.0774%	6.5012%	6.2127%

Table 2. Incidence of successful dimension-domination manipulation in simplex simulation

d \ *n*	2	3	5	8	13
3	0.0000%	0.0000%	0.0000%	0.0000%	0.0000%
4	0.0000%	0.0000%	0.0000%	0.0000%	0.0000%
5	0.0000%	0.0000%	0.0000%	0.0000%	0.0000%
6	0.0000%	0.0000%	0.0000%	0.0000%	0.0000%

MAKING ASYNCHRONOUS DECISIONS

In many real games, collective decisions must be made quickly and players may need to express their preferences simultaneously. In such situations, while a nonmanipulable method of aggregating preferences removes any chance of successful manipulation, a manipulable method like Average aggregation may work well enough in practice if the players have too little insight into the others' preferences to take advantage of them. A player may not know enough to inform a strategically insincere "vote".

On the other hand, in many social games, players are often playing at different times and on different schedules; Sid Meier's Civilization World (under development; originally titled Civilization Network) may prove to be an excellent example. If a cooperating group of players needs to make a collective decision, the players may often need to register their preferences asynchronously. Our AAR-DSV-based methods adapt to this environment gracefully. In fact, even if expressed preferences are made public before all of them have been expressed, players cannot take advantage of them. Players can never do better (in the senses discussed above) than to express sincere preferences, and so no unfair advantage is generally gained by expressing preferences sooner or later than other players.

Besides Civilization World, there are many other games currently being developed for Facebook and other social networking websites. Our approach could potentially find a use in many games in which players cooperate by casting "votes" for collective actions or other kinds of outcomes.

FUTURE RESEARCH DIRECTIONS

There are several potentially interesting and useful directions for future research. First, in light of the results of our simulations, it may be possible to prove for the simplex outcome space the equivalent results already proved for the line-segment and hypercube spaces. Also worth consideration may be collective-decision-making methods with other outcome spaces, whether continuous spaces of different shapes or discrete spaces.

Discrete outcome spaces certainly occur in real games. They occur whenever players are faced with a collective decision among finitely many mutually exclusive choices, such as whether to deploy phasers or photon torpedoes against an enemy, or which of several computed paths to follow. A continuous outcome space may even be effectively discretized when players' preferences are multiple-peaked, such as when all players would prefer to attack either of the two flanks of an opposing army than to attack the center. Previous impossibility results (Arrow 1951, Gibbard 1973, Satterthwaite 1975) may preclude the possibility of finding perfectly nonmanipulable decision-making protocols with (effectively) discrete outcome spaces, depending on assumptions made about players' preferences, but the DSV framework may prove successful in minimizing opportunities for insincere strategy.

Perhaps more immediately important than future theoretical directions, however, is deployment of these techniques in real games. It will be instructive to note how useful and intuitive human players find them when cooperating with other human players and NPCs. These techniques could also be used when groups of NPCs make collective decisions, and such use may make a game's AI become more effective and/or seem more realistic—they may prove to increase the apparent independence of NPC agents. Whether used with human players or NPCs, these techniques offer a flexible alternative to hierarchical group organization. Eventually, they may inform future designs for multiagent AI frameworks in real games.

CONCLUSION

In this chapter we have presented new methods for collective decision-making in games and explored their immunity and/or resistance to manipulation by insincere players. We specifically applied our approach to the line-segment, hypercube and simplex outcome spaces, giving examples of each.

We believe that our AAR DSV approach has many advantages:

- It can contribute to both the AI and the player-to-player frameworks.
- It allows agents to concentrate on determining optimal policies instead of on deceiving other agents with whom they are ostensibly cooperating.
- It places a relatively small burden on players: They need only indicate their ideal outcome; no complex ranking or rating of outcomes is needed.
- It allows cooperating agents to find a compromise immediately, which is painless, rather than by fighting it out through the game, which may be costly to all.
- It allows important decisions to be made whether players indicate their preferences simultaneously in real time or at different times.
- Its outcome functions are decisive, efficient and easily implemented.
- It is sufficiently general to be applied to decision situations in almost any kind of game.

Certain assumptions, concerning the players' preferences and continuity of the outcome space, are required for AAR DSV to work perfectly, but these assumptions are at least approximately true in many game situations. We conclude that our approach is likely to be useful when designing real games.

REFERENCES

Arrow, K. J. (1951). *Social Choice and Individual Values.* New York, New York: Wiley.

Bartholdi III, J. J.; Tovey, C. A.; & Trick, M. A. (1989). The computational difficulty of manipulating an election. *Social Choice and Welfare, 6*, 227–241.

Border, K. C., & Jordan, J. S. (1983, January). Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, *50*(1), 153–170.

Conitzer, V., & Sandholm, T. (2003, August). *Universal voting protocol tweaks to make manipulation hard.* Paper presented at the 18th International Joint Conference on Artificial Intelligence, Acapulco, Mexico.

Cranor, L. F. (1996, December). *Declared-Strategy Voting: An Instrument for Group Decision-Making*. Doctoral dissertation, Washington University, St. Louis.

Cranor, L. F., & Cytron, R. K. (1996, April). *Towards an information-neutral voting scheme that does not leave too much to chance.* Paper presented at the Midwest Political Science Association Annual Meeting, Chicago, Illinois.

Elkind, E., & Lipmaa, H. (2005, December). *Hybrid voting protocols and hardness of manipulation.* Paper presented at the 16th Annual International Symposium on Algorithms and Computation (ISAAC), Sanya, Hainan, China.

Gibbard, A. (1973, May). Manipulation of voting schemes: A general result. *Econometrica*, *41*(3), 587–601.

LeGrand, R. (2008, April). *Computational Aspects of Approval Voting and Declared-Strategy Voting*. Doctoral dissertation, Washington University, St. Louis.

LeGrand, R., & Cytron, R. K. (2008, September). *Approval-rating systems that never reward insincerity*. Paper presented at the 2nd International Workshop on Computational Social Choice (COMSOC), Liverpool, England.

Moulin, H. (1980, January). On strategy-proofness and single peakedness. *Public Choice*, *35*(4), 437–455.

Nisan, N. (2007). Introduction to Mechanism Design (for Computer Scientists). In N. Nisan, T. Roughgarden, E. Tardos & V. V. Vazirani (Ed.), *Algorithmic Game Theory* (pp. 209–241). Cambridge, England: Cambridge University Press.

Procaccia, A. D., & Rosenschein, J. S. (2007). Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, *28*, 157–181.

Rabin, S. (Ed.). (2002). *AI Game Programming Wisdom*. Hingham, Massachusetts: Charles River Media.

Satterthwaite, M. A. (1975, April). Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, *10*(2), 187–217.

Zhou, L. (1991, January). Impossibility of strategy-proof mechanisms in economies with pure public goods. *The Review of Economic Studies*, *58*(1), 107–119.

ADDITIONAL READING SECTION

Black, D. (1948, February). On the rationale of group decision-making. *Journal of Political Economy*, *56*(1), 23–34.

Conitzer, V., Lang, J., & Sandholm, T. (2003). How many candidates are needed to make elections hard to manipulate? In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-03)* (pp. 201–214). Bloomington, Indiana.

Conitzer, V., & Sandholm, T. (2002). Complexity of manipulating elections with few candidates. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI '02)* (pp. 314–319). Edmonton, Canada.

Downs, A. (1957). *An Economic Theory of Democracy*. New York, New York: Harper.

Ephrati, E., & Rosenschein, J. S. (1993). Multi-agent planning as a dynamic search for social consensus. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI '93)* (pp. 423–429). Chambéry, France.

Merrill III, S. (1988). *Making Multicandidate Elections More Democratic*. Princeton, New Jersey: Princeton University Press.

Pennock, D. M., Horvitz, E., & Giles, C. L. (2000). Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI '00)* (pp. 729–734). Austin, Texas.

Procaccia, A. D., Rosenschein, J. S., & Zohar, A. (2007, January). Multi-winner elections: Complexity of manipulation, control and winner-determination. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI '07)* (pp. 1476–1481). Hyderabad, India.

Vazirani, V. V., Nisan, N., Roughgarden, T., & Tardos, É. (2007). *Algorithmic Game Theory*. Cambridge, England: Cambridge University Press.

Von Neumann, J., & Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Princeton, New Jersey: Princeton University Press.

KEY TERMS & DEFINITIONS

Average-Approval-Rating (AAR) DSV: A specific application of the DSV framework to a continuous outcome space that uses the Average aggregation procedure (simply averaging the inputs) as the internal voting method for the simulated election.

Declared-Strategy Voting (DSV): A framework for collective decision-making that uses agents' preferences to vote on their behalf in a simulated election, the result of which becomes the DSV outcome.

Hypercube: A generalization of a cube into any dimension. For example, $\{(x, y, z): 0 \le x \le 1 \; and \; 0 \le y \le 1 \; and \; 0 \le z \le 1\}$ is a 3-hypercube (cube).

Mechanism design: A subfield of game theory that aims to design mechanisms that output a decisive outcome given agents' input and are relatively robust to rationally strategic agents.

Outcome: The output of a collective-decision-making mechanism.

Outcome space: The set of allowed outcomes over which agents have preferences.

Simplex: A generalization of a triangle into any dimension. For example, $\{(x, y, z): 0 \le x \le 1 \; and \; 0 \le y \le 1 \; and \; 0 \le z \le 1 \; and \; x + y + z = 1\}$ is a 2-simplex (triangle).

Single-peaked preferences: Preferences for which it is never true that *a* and *c* are each preferred to *b* when $a \le b \le c$.