Transferable Votes for Fair-Share Spending

Introduction Dec 15, 2014

- Participatory Budgeting
- other questions of resource allocation
- This paper describes our efforts at applying STV to questions of resource allocation.

How it works

The system as used in real elections

Evaluation
- changing the quota
  - even over 50% quota still gives fair shares
- elimination order
  - delete last loser
- later no harm
- ease of use for voters, motivation for voting
- fairness, how to evaluate fairness;
  - fairness and utility

Future work

Conclusion
- why you should care

- how fair share representation does not in and of itself lead to fair share resource allocation

ntroduction

STV, and voting systems generally, have mainly been considered in the context of election of candidates. This paper describes our efforts to apply STV to questions of resource allocation.

A timely and interesting example of applying voting to resource allocation is the participatory budgeting movement. Participatory Budgeting, PB generally includes the use of voting as the final step in allocating part of a municipal or other public budget. It has been used since 1989 in Porto Alegre, Brazil to allocate part of the municipal budget---on average $71.5 million. [http://www.participatorybudgeting.org/about-participatory-budgeting/examples-of-participatory-budgeting/] Participatory budgeting has since spread to hundreds of cities across Latin America and dozens of others around the world. In recent years, it has taken root in North America, starting with a ward in Chicago in 2009, with citizens allocating one million dollars; the 2014 (?) process in 10 council districts of New York City will use citizen involvement and, in the end, voting, to allocate $14 million dollars.

Participatory budgeting is not necessarily the only way in which voting could be used for resource allocation. Most notable is Elinor Ostrom's Nobel Prize winning work on allocating common pool resources. http://en.wikipedia.org/wiki/Elinor_Ostrom We speculate that voting could potentially be useful to communities allocating forests, fisheries, oil, grazing land, or water. For simplicity, we will use money in our examples.

Generally, participatory budgeting processes have used some variant of bloc vote; some may have used Borda count or cumulative vote. [http://techpresident.com/blog-entry/australia-e-participatory-budgeting-experiment] Readers will be well aware of the issues with these voting rules in candidate elections; it seems likely to us that those issues are even worse applied to resource allocation. We are naturally interested in applying STV in some form instead.

In addition to its practical interest in participatory budgeting, we also believe that applying STV to resource allocation may provide interesting insights simply from the point of view of studying the technical issues of STV. From a technical point of view, STV for resource allocation is a direct generalization of ordinary STV, and as such may provide a new and interesting perspective on certain technical issues. Such issues include the Hare quota vs the Droop quota, or the Meek method vs the Warren method. We will call out such issues over the course of this paper.

We will begin with a detailed discussion of how we believe STV could best be applied to resource allocation. We will then proceed to consider the use of our method in real elections; we have ranked-choice ballots from real participatory budgeting elections in a small community, and we evaluate the results of applying our method.

Our new method builds on earlier methods and retains their three essential principles.(1)

Projects with set funding levels

An election is specified by a total amount of money to be spent, together with a list of n candidate projects P_1 to P_n, each specifying its amount of money a_1 to a_n.

Each voter's ballot is a ranking of some or all the projects. Each valid ballot will allocate that voter's fair share of the total money, total/voters.

Any STV method can then be used to perform the allocation and determine which projects are funded. For concreteness we will use Meek's method in our discussion and examples.

We begin with a simple STV structure. In the first iteration, each voter's full share is distributed to the first project on their ballot. If any project P_i has received its full amount a_i, that project is "elected". If there is a surplus, it will be distributed to ballots' later preferences in later iterations, in the standard way which we will not discuss here.

In Meek's method this is done by assigning a keep value k_i to project P_i at each step of the iteration, where k_i is initially one. At each step of the iteration, a ballot supporting P_i gives it k_i of the remaining share allocated by that ballot. If at some step of the iteration P_i receives funding f_i > a_i, then the new keep value will be k_i * a_i / f_i. Since the keep value will be lower in the next iteration, each ballot supporting P_i has more funds available for later preferences.

Eliminations

When it is no longer possible to make progress by transferring surplus, then it is necessary to choose a candidate project for elimination. In a standard election one would eliminate the candidate with the lowest vote total. In a budgeting election, where different projects require different amounts to be elected, we must consider other criteria. Two initially reasonable ideas are (1) eliminate the project with the lowest proportion of current funding relative to its needed amount, or (2) eliminate the project with the greatest absolute amount remaining before it is funded. It is likely that the proportion criteria is better because it does not create an incentive to split a costly project into small pieces, each of which would have a small absolute deficit, so as to avoid elimination.

Our initial examination of real-world results suggests that elimination order, known to be at least a theoretical issue generally in STV, may be especially problematic in budgeting elections with projects of varying amounts. Thus it may be practically necessary to consider methods such as Sequential STV or CPO-STV which attempt to work around problems of elimination order.

We have tested a method which we call Delete Last Loser (DLL). The last candidate eliminated holds on to the weights of some voters but is not elected and so may cause wasted votes. More importantly, eliminating that Last Loser before another tally can give candidate(s) with similar appeal a better chance to win. At the end of the previous tally, the Last Loser had the best chance it could reasonably get. Sequential STV is similar and requires less calculation, but DLL probably works better at fitting winners into a budget where candidates vary in “cost”.

Eventually every project is either elected or eliminated, and the election is complete. If the election results include the final keep values, each voter can take the election results and determine exactly how their share of the total funding was allocated.

Exhausted ballots

If we allow voters to leave some projects unranked, then there is an issue of exhausted ballots. What happens to what is left of a voter's share of the money when they have no more non-eliminated preferences to allocate it to? In standard Meek STV elections of candidates, exhausted ballots are considered to fractionally reduce the number of valid votes, and thus the quota. In a budgeting election, the corresponding effect is to increase each voter's share; and the voter whose ballot was exhausted will allocate less money than other voters.

An alternative is not to allocate the exhausted funds at all. They can go back to the funding source for use in future years. This option, which is not available in a standard election of representatives, is probably most in line with the voter's intentions: none of the remaining projects are worth funding.

Quota

We have managed to discuss an entire STV election without referring to the notion of quota. In a budgeting election, the notion of quota—how many votes a candidate needs to be elected—is subsumed by funding: each project needs a certain amount of money in order to be funded.

As it turns out, this reduces to the Hare quota. To see his, consider a budgeting election where there are **s** dollars to be spent and each project costs $1. If there are v voters, then to be funded a project needs v/s votes. For example, if v is 50 and s is $500. Then 50/$500 = $0.10 is the value of 1 vote. So each $1 project needs 10 votes.

[If each project costs $1 and needs 10 voters, then to spend $500 funding 500 projects, 50 voters need to cast 5,000 votes; that's 100 votes each. *Isn't a decision with 500 projects is too big and rare*?]

[If each project costs $100 and needs 10 voters, then to spend $500, 50 voters need to cast 50 votes; that's 10 winning votes each.]

[The ratio s/v determines how much each voter is allowed to allocate, which we call a voter's share. RL maintains that the ratio v/s never determines the number of voters an item must get. That requirement is exogenous.

[Is s the traditional notation for this total amount? We might use s for the share amount.]

[As it turns out, this reduces to the Hare quota. To see this, consider a budgeting election where there are s dollars to be spent and each project costs $100. If there are v voters, then to be funded a project needs $100/(s/v) voters' shares. For example, if v is 50 and s is $500, then each voter's weight or share is $500/50=$10. So each $100 project needs 10 voters' shares.

=[For example, v is 50, s is $500 and there are 5 seats to fill. Then each candidate needs to receive

$500/5 seats = $100 to win a seat, and each voter starts with $500/50 voters = $10. That is the "weight" of each voter's share.]

Or, start by saying each voter's weight is $1 and it goes to 1 candidate? Then change it to an election in which each voter helps several winners.]

We find it very interesting that the clear and intuitive notion, that a project needs to be allocated the dollars that it will spend, corresponds to the Hare quota, which is almost universally passed over in favor of the Droop quota in applications of STV for electing candidates. Moreover there is no immediately obvious way to base a budgeting election, where projects might have widely varying amounts, on the Droop quota.

There is another notion of quota which needs to come into play to use this method for allocation of public funds. The funding-based quota only involves number of dollars, not number of voters. Thus it would be possible for a small project to be funded by a small number of voters, or even a single voter. This is not consistent with the idea that the election is about spending public funds on public goods. Thus, we suggest an additional requirement, that there be a minimum number of votes (as opposed to dollars) that must go to a project before it can be funded. This would be expressed as a portion of the valid votes: perhaps 10% or 20% of all the voters must support a project before it can be funded. The amount would be set before the election and, technically, is part of the specification of the election. Although this “quota” has in common with the standard STV quota that it represents a fraction of the ballots, it is the funding-based

quota that is more mathematically analogous to the standard STV quota, and so to avoid confusion we will call this percentage of voters the support requirement.

Practically, the support requirement needs to be tuned to the political needs of any given budgeting election. A smaller support requirement gives more power to minority groups to pool their funding to get projects they want funded. A larger support requirement ensures that all funded projects have broad support, even if only at lower preferences.

Technically, a support requirement is implemented by capping the percentage that any voter is allowed to give to a project. This can be done, for instance, by specifying an initial keep value for each project less than 1 in Meek's method. It should be noted that this means that even the initial allocation from a single ballot might support multiple projects. This removes the "Single" from STV and makes satisfying the later-no-harm criterion difficult(*), exacerbating the previously mentioned problems of elimination order.

[(*) Each voter could be allowed to try to save a threatened favorite by briefly withholding his or her votes from lower-choice items. Or we could give 2 votes to the first dollar of a voter's top project and step down to 0 votes for the last dollar in the voter's share, thus reducing the chance it will hurt the voter's current top project.]

In those budgeting elections where the nomination of projects is a separate process involving some amount of public scrutiny and approval before each project is put on the ballot, the support requirement

may not be necessary. If, on the other hand, the election is open to any proposed project, a support requirement is clearly necessary to avoid absurdity.

A quota of votes can prove enough voters think the item is a public good, so it is authorized to get public money. But a quota of serious contributors can prove enough voters think the item is also a top priority for public funding compared with other proposals. A quota of money can prove enough voters think it is worth its cost at a certain level of funding.

The new quota is made by combining two requirements: a minimum number of serious supporters and the item's cost. The item must meet both requirements. The vote quota must be filled to prove the item's breadth of support. The cost quota must be filled to prove their intensity of support.

We must combine the latter two support requirements into one that will decide which item to eliminate. We combine them by capping the amount a voter may give an item at its cost divided by the quota. The amount an item collects during a tally then shows both how many votes support it and how close they are to paying its cost.

Why not hold two separate votes? The first would approve proposals for public funding then the second would elect the set of winners.

A) Staging two votes takes longer to reach a decision. Staging just one vote is quicker.

B) Staging two votes demands more time and effort from voters. Staging just one vote is easier.

C) When evaluating/appraising an item's public good, people who put their share of the money on the line are likely to be more serious and thoughtful than people who give approval votes with no money.

D) A group of conspirators can manipulate the two step method more reliably. A simple example: Even strangers could agree to approve each others' proposals for items to go in front of their homes: a bike rack at my house, a new street tree for you, a new driveway apron for someone else. If everyone keeps the agreement to give cost-free approval votes, then each of them could then use his or her share to buy a gift for oneself.

The conspirators in this simple example probably could not gather enough individuals to meet a support requirement of 10% or 20%. A more likely plan would use agreements by residents of several blocks or apartment buildings.

In the two step method, the approval vote is tallied first; then each voter knows before casting the second ballot whether or not all others have kept the agreement. But in the 1 step method with secret ballots, the agreement is unenforceable and therefore less inviting. In the case of elected representatives whose votes are public, they will not want to be seen spending public money on a "bridge to nowhere" or other costly projects with narrow appeal.

Projects allowing a range of funding levels

Now consider a more general scenario where each project has not a specified amount, but a range of possible amounts to be determined by the voters. In

this case an election is specified by a total amount of money to be spent, together with a list of n candidate projects P_1 to P_n, each specifying a minimum funding level min_i and a maximum funding level max_i.

Each voter's ballot is a ranking of some or all of the projects together with a proposed funding level for each, where the proposed level must be between the project's minimum and maximum.

A voter putting project P and amount A at a given preference level is interpreted as "I wish to put any money I have left after funding higher preferences, as much as one is allowed to give, toward funding project P, up to a total funding for P of A dollars".

Ranking multiple funding levels for a single project

We might allow a voter to rank a project at different preferences at several different funding levels. For instance a voter might have preferences

P $500
Q $500
P $1000

This means the voter's share will (as much as possible given the constraints of a support requirement) go to funding P up to $500. If the voter has money left over, what's left will go as much as possible to funding Q up to $500. Finally, if there is still money left over in the voter's share, what's left will go to raising P from $500 to $1000.

It is unclear whether the additional expressivity of this allowance is worth the extra complexity. In particular, it might be confusing to voters. For

example, one voter expressed to us the idea that it seemed sensible that multiple funding levels for a project should go down, rather than up, in later preferences. This reflects the idea that the voter might like to express “As my first preference, I would like to get P funded at $1000; if that is not possible, as a second preference I would like to get P funded at $500”. In actual fact, a preference of “P at $1000” is interpreted as the desire to fund P up to a maximum of $1000, so all lower funding levels are already part of it.

Calculating with multiple funding levels

Consider a Meek-STV based budgeting election with multiple funding levels. At each iteration, each project P_i has a list f_1, …, f_{n_i} of funding levels which are being considered. (Initially, think of this as all the funding levels which have been proposed by voters; but see “Final funded preference”, below...?) Each funding level f_j is associated with a keep factor k_j. The initial keep factor is determined from the support requirement. The keep factor will be zero for eliminated funding levels (which are necessarily all funding levels above a least such).

(Trying with traditional def of keep factor)

We can adapt the usual notion of keep factor in a straightforward manner to be the fraction of a voter’s remaining share which goes toward funding project P from f_{j-1} to f_j. It needs to be kept in mind that, except for the minimal funding level f_1, the voter will already have spent some share raising P up to f_{j-1}, and so the “remaining share” will never be the full share in those cases.

B1 Now consider a voter's ballot with a preference "P from 0 to f_j", where the voter's total share is s, of which r remains after offering funds to higher preferences. Then the voter will offer a_1 = k_1 * r to get the project funded to f_1, and a_2 = k_2 * (r – a_1) to raise the project from f_1 to f_2, and so on up to a_j = k_j * (r – sum a_1 to a_{j-1}) to raise the project to f_j .

Note that the voter's ballot is not affected by splitting a single preference "P from 0 to f_j" into two consecutive preferences "P from 0 to f_i", "P from f_i to f_j ".

C1 A split preference might express "Raise P from f_i to f_j ", in which case the voter will offer a_{i+1} = k_{i+1} * r to raise P from f_i to f_{i+1}, and so on up to a_j = k_j * (r – sum a_{i+1} to a_{j-1}) to raise the project to f_j .

??? If we know the keep factor k for P from 0 to 100, what are the keep factors for P from 0 to 50 and P from 50 to 100? Answer: P 0 → 50 gets half of what P 0 → 100 got, with the same remaining share to start, so gets k/2. P 50 → 100 gets the other half, but with a remaining share of (1- k/2)*s; so it gets (k/2) / (1-k/2). In particular this means if k = 1, then the keep factor of P 50 → 100 is 1: P keeps everything the voter has left at that point.

… More generally: if k is the keep factor for P from a to c, then the keep factor for P from a to b is k_1 = k*(b-a)/(c-a); and the keep factor for P from b to c is k_2 = k / (1 – k_1) * (c-b)/(c-a) = k*(c-b)/(c-a – k*(b-a)).

(trying with idiosyncratic keep factor)

For technical reasons we find it beneficial to use an unusual notion of "keep factor" which is the portion of the needed funds which will be allocated by a single voter supporting this project at this funding level. The traditional notion of keep factor is the portion of one voter's full share which will be allocated by a single voter. The two definitions are clearly intercalculable; but it is easier to work with the former for at least two reasons: (1) the initial keep factor, which is determined by the support requirement, will be the same for all projects and funding levels; (2) we can more easily start considering a new funding level partway through the calculation. What this second point means is the following. Suppose at some iteration each voter supporting project P at $200 as a first preference is giving $10; and suppose each voter's share is $20. The keep factor is 50% of the share, or 5% of the funding amount. Now suppose we need to consider how much a voter raising P from $0 to $100 should give, and then separately how much a voter raising P from $100 to $200 should give. The amount each voter will give is $5 in each case; that is 25% of the share, but, still in each case, 5% of the funding amount. Thus, considered as a portion of the funding amount, we can introduce an intermediate funding level $100 without having to change the keep factor for $200.

B2 Now consider a voter's ballot with a preference "P at f_j", where the voter's total share is s, of which r remains after offering funds to higher preferences. Then the voter will offer k_1 * f_1 to get the project funded to f_1, and k_2 * (f_2 – f_1) to raise the project

from f_1 to f_2, and so on up to k_j * (f_j – f_{j-1}) to raise the project to f_j.

C2 A split preference might express “Raise P from f_i to f_j”, in which case the voter will offer k_{i+1} * (f_{i+1} – f_i) to raise P from f_i to f_{i+1}, and so on up to k_j * (f_j – f_{j-1}) to raise the project to f_j.

Warren and Meek methods

The calculation given so far is appropriate for Warren’s method. Meek requires the portion of the vote offered at non-first preferences to be proportional to the amount of share remaining. So an extra step is required to reduce the offered funding for non-first preferences; this turns out to be more complicated in the presence of multiple funding levels. It is necessary to ensure that a ballot which raises P from $0 to $200 in a single preference, behaves identically to a ballot which raises P from $0 to $100 in one preference and then from $100 to $200 in a second preference.

Let W_i be the calculation C2 above of the funding offered to take project P from f_i to f_{i+1}. To find the fun 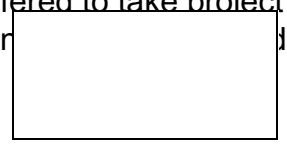 der Meek’s method, we calculate:

M_0 = W_0 * r_0 / s

M_i = W_i * r_i / (s – sum M_0 to M_{i-1})

where s is the voter’s whole share, and r_i is the amount of the share remaining before offering any funding beyond f_i. For a first preference, r_0 = s and r_i = s – M{i-1}, and so M_i and W_i coincide.

For funding levels contained within a single preference, r_i = r_{i-1} – M_{i-1}. In the case of a split preference, it is possible that r_i < r_{i-1} – M_{i-1}; for example in a ballot with the three preferences: first choice raise P from 0 to f_i, second Q from 0 to g, and third P from f_i to f_j. So the split choice for P would give less money than a ballot with a single first choice to raise P from 0 to f_j. ?

It is noteworthy that Warren's method has a certain technical naturalness in this context which Meek's method lacks. This is not purely a result of our redefinition of keep factor. In fact, if keep factor is fraction of a full share, consider again the case of a voter raising P from $0 to $100 and then raising P from $100 to $200. If the voter is to give $5 for each step, that may be 25% of the share, but in the second case it is 33% of the remainder after funding P up to $100. So in order to make the Meek calculation a simple multiple of keep factor and remaining share, we have to define keep factor, not as the portion of a whole share given, but as the portion of a whole share minus what was already given to smaller funding levels for the same project.

… but splitting the preference could cause things to behave differently for the last funded preference, for Meek or Warren...

Final funded preference

The final funded preference on a ballot requires special handling. This is the preference which exhausts the voter's share. Again we must ensure that

a ballot behaves the same way when an amount of funding for P is expressed as a single preference or multiple consecutive preferences. To this end the final preference should be reduced to the funding level which would exactly exhaust the voter's share. Only if such a funding level would be below the project's minimum …

An example. Suppose the support requirement is 100 people, and each voter's share is $20. A voter who supports P at $4000 can't actually afford to give $4000/100 = $40. So we consider the voter to support P at $2000 instead; raising P from $2000 to $4000 is a second preference.

If $4000 is the minimum funding level for P, it does not allow this split, so each [this?] voter just gives half as much – and counts as only a half vote for P.

The effect of this is to ensure that the funds given are maximal. Giving $20 to P at $2000 is 1 vote and is more likely to be effective than giving $20 to P at $4000.

??? Maybe we should consider it as only supporting up to the minimum, so that the minimum gets more support; or maybe that is automatically taken care of by elimination.

The system as used in real elections

Twin Oaks Community is an income-sharing community of about 100 adults in rural Virginia. It was founded in 1967 and in recent years has had an

operating budget of about $2 million.

- changing the quota
- even over 50% quota still gives fair shares
- ease of use

Conclusion

Some ‘why you should care’ bits:

As noted above, from a technical point of view, STV for resource allocation is a direct generalization of ordinary STV, and as such may provide a new and interesting perspective on certain technical issues.

Fair-share representative democracy does not in and of itself lead to fair-share resource allocation. It is commonplace in modern democracies for an elected council to allocate a budget; and generally this means that a bare majority of the council has the power to allocate 100% of the budget. In terms of power to influence the budget, votes which elect non-majority council members are effectively wasted votes.

This __ [failure, limiting fair-share democracy], leads to the bases of many serious conflicts.

Electoral reform is very difficult because relatively unfair voting methods have become entrenched by centuries of use. Because participatory budgeting is still young, we have a unique opportunity to introduce better voting methods now — voting methods that are more expressive and more fair.

Footnotes

This version of Movable Money Votes, MMV, is Meek's STV if each voter gets $1 and if each candidate costs $ [number of voters / (seats + 1)].

This fact has let us test his program using all of the real elections in the ERS data set. The results are identical to those using OpenSTV Meek. The only hiccup was where a random tiebreaker was needed.

Note on Warren's method

Warren's method has an interesting interpretation in a budgeting election. Under Warren's method, each ballot supporting a project gives the same amount of money as every other ballot, with one exception. The exception is for ballots where the project is the voter's last-funded preference, in which case it uses what is left of the ballot's available funds, and thus possibly less than the equal share – so it gets less than 1 vote. Under Meek's method, each ballot supporting a project gives the same proportion of the ballot's total share left after funding higher preferences; thus in general, ballots which support a project at a higher preference (among funded projects) will give it more money.

(1) N. Tideman and D. Richardson. Better voting methods through technology: The refinement-manageability trade-off in the single transferable vote. Public Choice, 103(1):13–34, April 2000.

"STV is not a specific voting method, but rather a family of voting methods that share certain principles:

"1. Every voter should be allowed to allocate all of his vote to the candidate[s] of his choice.

"2. If a candidate has more than enough votes to be elected, then surplus votes should be transferred to the next available candidates in the rankings of those who voted for the candidate with the surplus.

"3. If the candidate to whom a vote is presently allocated is excluded, then that vote should be transferred to the next available candidate in that voter's ranking."

Note that "later no harm" is not one of Tideman's three principles — although it is a top priority for FairVote's Executive Director, Rob Richie, and others.

Sequential STV — a new version. I.D. Hill and Simon Gazeley, Voting matters, 15, 13–16. 2002.

Sequential STV — a further modification I.D. Hill and Simon Gazeley; Voting Matters - Issue 20, June 2005

==========

History of Voting in Participatory Budgeting

Since 1976, Twin Oaks Community has used several voting methods to select "One-Time Resource Allocations" (OTRAs) and set ongoing budgets in their annual "Trade-Off Game" (TOG). To invigorate the "game," the planners occasionally change the rules a little, but it is always a kind of Participatory Budgeting.

(Twin Oaks is a diverse intentional community, founded in 1967 in rural Virginia, with around 100 members. They do not have a central leader but govern themselves democratically, with responsibility shared among various managers, planners, and committees.)

Participatory Budgeting (PB) in big city budgets dates from 1989 in Porto Alegre, Brazil. It has spread to over a thousand municipalities, most in South America and Western Europe. PB elections usually use a plurality rule such as Bloc Vote.(1) A few elections have used simple averaging of votes.(2)

1) The 2011 PB vote in Chicago was like Bloc Voting: Voters got 6 votes and as it happened, 6 projects won. Often the vote is like Limited Voting: each voter is allowed to cast a few votes less than the number of projects likely to win funding. Bloc Voting is a plurality rule, not proportional. (U.S. courts ruled it unconstitutional or in violation of the Voting Rights Acts where it was adopted to prevent minority representation.)

2) A 2011 tally averaged the votes for the street-resurfacing fund in Chicago's 49th Ward.

"1. Streets Vote – A separate question will appear on the ballot asking voters to vote on a specific percentage of the 2011 aldermanic menu budget that should be devoted exclusively to street resurfacing. The votes on this separate question will be tallied and an average calculated. That average will determine the percentage of the 2011 aldermanic menu budget that will be devoted exclusively to street resurfacing." Emphasis in original.

http://participatorybudgeting49.wordpress.com/about/2011-ballot/

Fair-share Spending does no averaging. There is no incentive to exaggerate one's allocation to a project. As mentioned above, there are no incentives to break a project into smaller pieces nor to combine several projects into larger a one. [The latter appears to happen in PB with bloc voting, e.g. consolidating proposals from different schools or adding amenities to larger proposals...]

Advantages of FS over previous PB voting

A more optimal allocation of resources:

Cost-awareness

In particular, when considering voting for resource allocation we want to consider cost-awareness. In the 2009 Chicago participatory budgeting process, a vote for a park was worth $501; but if given to fund bike racks, it was worth only $31.

From the Chicago 2009 participatory budgeting process:

Bike Rack
$12000 / 388 = $31 per vote.

New Playground at Touhy Park,
$175000 / 394 = $501 per vote.

From the Chicago 2010 participatory budgeting process:

Intersection safety at Clark and Farwell,
$2600 / 334 voters = $8 per vote.

Traffic Signal on Clark at Chase,
$230000 / 494 voters = $466 per vote.

The bike racks pleased 388 voters, for only $31 per

voter. The intersection safety project cost only $8 per voter. Recent PB elections in Chicago try to conceal, but not fix, this flaw in their voting method by banning projects under $25000; some of those may be the most cost-effective proposals. Elections have also cut the total value of proposals relative to funding, so there are fewer losers with wasted votes. (In 2009, 63% of the votes went to losers, 7% went to winner surplus.)

Cambridge Massachusetts first voted for PB projects in 2014. Their Bloc Vote tally was not cost aware: A winning vote might be worth $339, or just $7. A vote for the new restroom was worth at least $339. But a ballot that voted for *all* five other winners was worth about *half* as much, $175. More than half the votes cast were for losers and so were worth nothing. The whole PB fund could win more votes: 9,216 votes, 73.48% for winners, instead of only 6,258 votes, 49.9%. The PB money could build fourteen proposals instead of only six.

Our cost-aware FS method does not call for such limitations on the choices offered to voters. It also gives voters the proper incentive to favor cost-effective proposals. Thus it tends toward a more optimal allocation of resources.

A more equitable distribution of resources: proportional, fair share

[Notice we say "equitable" rather than "egalitarian." That is because no voting method can give those who merely vote as much influence as those who discuss and develop proposals. Perhaps it is right that there should be some reward for those who devote more

time and effort to the decision-making process.]

More incentive to vote: Because fewer votes are wasted and more are effective and perhaps decisive.

(Currently, all forms of Proportional Representation stop after an election. Then all of the power to set policies and budgets goes to the most cohesive majority of representatives. So we might consider the votes cast to elect the minority representatives were wasted votes. The use of FS can change that, by making votes for an opposition party or faction effect allocation of the discretionary fraction of the budget. This can give minority voters more incentive to vote for their representatives.)

Previous Related Work

In February 1996, at the suggestion of second author Robert Loring, the River Road Unitarian Church in Bethesda, Maryland used Droop STV to select which 3 social justice projects would be included on every member's annual pledge form.

In May 1996, Helen Popenoe, chair of the social justice council of the Greater Washington Association of Unitarian Universalist Congregations, asked Loring to design and lead a fair-share, multi-winner vote to allocate a special fund.

Each project's chairperson stated their request and answered questions. The requested amounts varied by an order of magnitude. (The largest was cut in half due to discussion feedback.) Each member wrote which proposals he/she would fund while staying within the total budget. A show of hands tallied the votes for each

proposal. The one with fewest votes was dropped; those who had helped it transferred their money to other proposals. This process was repeated until the remaining proposals fit within the total budget. Notice the method was cost aware and had no preset support requirement.

By 1998 Loring's DemocracyEvolves website on AOL offered a more sophisticated algorithm. This method averaged the amounts voters offered a project, so there was an incentive to exaggerate one's allocation to one's favorite project.

"Loring Allocation Rule (LAR) uses Condorcet's rule to fund central winners and a transfer process similar to STV to distribute some winners. Each rep's "weight" is her share of the budget -- a personal account that rewards efficient use. She ranks items and gives them $ votes, the amounts she thinks each should get. Her $ vote / a quota of ballots (set in the by-laws) is the amount her ballot "offers" to an item.

"If her $ vote for A is $300 and the quota is 10, her offer is $30. If A gets its quota of 10 offers, their total is the average $ vote (as if the $ votes had been totaled and then divided by the quota). The item's budget equals its total offers from contributors.

"Offers combine a quota of supporters (with a minimum $ vote) and the item's budget. (its budget = its total offers = its average $ vote.) These 2 criteria prove the intensity and breadth of support, to pay for the item and qualify it as a public good.

"Her ballot offers money to its top-ranked items, as many as her weight can afford. LAR suspends and

transfers offers from the weakest items 1 at a time. It also transfers surplus votes until all funds go to items with exactly a quota of supporters."

https://web.archive.org/web/20011122092332/http://accuratedemocracy.com/elect.htm#LAR

Also in the late 1990s, Loring presented workshops at the annual Twin Oaks Communities Conference to teach Movable Money Votes, MMV, a version of FS that let each voter place cards on a tally board. http://accuratedemocracy.com/a_workshop.htm#money

In 2003, he distributed the workshop handouts at the Claim Democracy Conference in Washington, DC. The handouts were also being distributed on the Web:

https://web.archive.org/web/20040218133314/http://www.accuratedemocracy.com/download/a_handout.pdf

A section of Loring's Accurate Democracy website explains FS for lay readers, plus some budget rules we think are less useful. Also online, you can try the easy and powerful ballot used by Twin Oaks and created by first author, Robert Tupelo-Schneck. http://tupelo-schneck.org:8080/tag/

At least two other researchers have independently investigated how to apply STV to resource allocation or budget setting. Both of these researchers deserve credit for publication, although our work was independent. In addition to offering a more detailed exposition in this paper, we also believe that certain details of our proposal lead to a superior system in practice.

In 2004 James Green-Armytage presented a proposal on the election-methods mailing list (http://lists.electorama.com/pipermail/election-methods-electorama.com/2004-May/013125.html or http://web.archive.org/web/20090515035023/http://fc.antioch.edu/~james_green-armytage/vm/budget.htm). In his proposal each voter controls an equal share of a total budget. Each preference item on my ballot specifies a project, the maximum amount of my share I am willing to allocate to that project, and also a "project cap" which is my desired maximum funding level for the project. I will not give the project more than the maximum amount specified, nor will I contribute to getting its funding above my specified project cap.

In the system we have presented in this paper, the voter only specifies the "project cap"; the maximum amount of share allocated is either unlimited or implicitly derived from the "project cap" by the "support requirement" factor (see section REF). We suspect that in a real election a single number is the most that should be expected from a voter; additionally the size of a voter's share may not be determined until the number of votes is determined. Green-Armytage also supposes that a project approval stage precedes the project funding stage, rather than allowing any project to be funded that has sufficient support as we suggest in this paper; as such he avoids having the consider how to formulate an elimination rule for projects of varying size, even though he does suggest that "severely underfunded" projects might transfer funds.

In 2006 Frank McGonigal presented a proposal at the TED 06 ("Towards Electronic Democracy")

conference (http://www.mi.imati.cnr.it/conferences/ted06/TED/TED06_mcgonigal.pdf). McGonigal applies STV to projects of various sizes, but does not suggest that voters can specify their own funding levels for projects. McGonigal observes three “problems” with STV for budgets: (1) the notion of quota reduces to the Hare quota rather than the Droop quota (see section REF); (2) low-cost projects may get funded even if relatively unpopular; and (3) more work is required to determine an appropriate and cost-aware elimination rule when projects vary in size (McGonigal does not explictly give such a rule). We share the first observations, but do not consider it to be necessarily problematic. We think FS solves the two other problems.

Motivation

A ruling majority commonly controls all government policies and spending. Such all-or-nothing political systems motivate political and even military wars. If a country uses FS, voters and leaders will feel much less motivation to fight seesaw battles.

Mosul is the focus of more intense warfare than most Iraqi cities because whoever controls it gets all the profits from nearby oilfields. Factions have used ethnic cleansing to force other groups to leave the city, making the victors the majority.

Future Work

Evaluating Results

We have run elections and simulations using several methods for Fair-Share Spending. It would be good to find ways of comparing their results:

Meek versus Warren,

Sequential STV versus Delete Last Loser,

protect your threatened favorite(s) by briefly withholding your cards from lower-choice items versus weight your favorite (with, say, one and one half votes) balanced giving only a half vote to lower-choice items versus neither.

[transferable votes versus ~ “voting cards”]

Measuring Utility

[Above “toward a more optimal use of resources”]

The utility value of the set of winners may relate to fairness: If a large group does not get its fair share, then the utility value is probably low. We do not advocate an interpersonal comparison of utility scores for voting or measuring fairness of resource distributions.

Measuring Fairness

To measure fairness, we can calculate how much of the spending can be attributed to each ballot. Then we might contrast the ballot which won the most with the ballot which won the least, or contrast the average of the bottom 20% with the average of the top 20%. It is

easy to calculate a Lorenz curve and a Gini coefficient to contrast the most winning and least winning voters. In this context, the Lorenz curve is a concave graph representing the cumulative distribution of winnings, starting with the least successful voter, then adding the second worst and proceeding to the most successful voter.

But should we include winning proposals from the whole ballot or just its top choices or just the items it helped to fund? All three are easy to calculate for each voter and compare within the electorate.

The Twin Oaks OTRA vote in 1995 used a plurality method. The ratio of most winning to least winning was almost 3: 1. In the 2009 vote using FS, the ratio was even lower. Virtually all ballots spent their full share unless the ballot was exhausted.

We can get more details: How far down in rankings did each ballot go?

We hypothesize that the new voting method will distribute “effective votes” more evenly among voters than the previous methods of PB voting did, just as STV does compared to multi-winner plurality voting.

* Footnote: The winner/loser ratio in the open market is far greater than in any of these PB cases. The ratio of pay to CEOs versus average workers in the Fortune 50 is 213:1 but that does not include total compensation. In the USA as a whole it is ___ . In Japan the ratio is just ___ . 147-to-1, the reported ratio in Germany

There is no ‘blueprint’ for organizations to copy; only lessons from the experience of repeated testing and evaluation in a variety of situations.

"Brazil does not offer a ‘blueprint’ for other countries to copy; only lessons from experience." http://ras.sagepub.com/content/78/4/733.abstract

Moses N Kiggundu "Anti-poverty and progressive social change in Brazil: lessons for other emerging economies" International Review of Administrative Sciences, December 2012 vol. 78 no. 4 733-756,

Sprott School of Business, Carleton University, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada Email: moses_kiggundu@carleton.ca