Consequences of Social Risk in Small and Moderate Sized Deliberative Democracies

Lilia Chang and Bora Uyumazturk

Stanford University, Stanford, CA 94305 USA
lilia.chang@stanford.edu, buyumazturk@stanford.edu

Abstract. The voting function that a group adopts is to some extent a reflection of its priorities. Many small to moderate sized communities use consensus voting to aggregate individual preferences in order to ensure that minority voices are heard and valued. In one such group on Stanford University's campus, unanimous approval is easier achieved than one might expect: votes almost always pass without any argument. In this paper, we gather empirical data to assess the true preferences of voters in this community and find that members often vote against their true preferences. We then present a game-theoretic model which captures the dishonest behavior of voters by taking into account social pressure. Finally, we explore the possibility of using a different voting scheme, quadratic voting, instead of consensus in these small to moderate group contexts. Using data from votes in Columbae, we provide evidence which suggests that QV may be more efficient than consensus voting from a maximin perspective in practice, and therefore may better achieve the goals of communities which aim to value minority voices. In doing so, we contribute to the growing body of work which attempts to bridge normative social choice theory, which considers just the aggregation of preferences, and on deliberative democracy, which formalizes deliberation as a key component of decision-making.

1 Introduction

Collective decision-making in small groups typically involves a discussion prior to reaching a conclusion. Whether it be amongst governing politicians, in the board room, or at the dinner table, deliberation and the justification of ones preferences are considered essential to most decision-making processes in these small to moderate sized group settings. The explanation is simple: the members want to consider many perspectives before coming to a decision, and furthermore the small to moderate group sizes make this feasible through direct deliberation. These sorts of democratic models of decision-making fall under the term *deliberative democracy.*

Columbae, a co-operative house at Stanford University, is an example of a community that seeks to integrate the ideals of deliberative democracy into its way of life. Columbae is one of many such communities: the *New York Times* in 2018 reported that cooperative houses are booming in popularity – ones that are cross-generational, and consist of not just college students [14]. By their own ethics, these houses are motivated to give everyone equal say in the house and to especially defer to subsets that perhaps deserve special attention: groups that are in the minority in thought or advantage. As such these houses have chosen a scheme of deliberation with a unanimous vote – which we refer to as *consensus* or CV – as opposed to other aggregation mechanisms such as majority rule.

Consensus, a procedure adopted from the Quaker tradition, aims to ensure that minority voices are taken into account; something that most other voting systems (such as majority rule, for example) cannot guarantee. In the case of Columbae, the house's fundamental policies regarding "meat, smoke, and noise" in the house are discussed and then voted upon. One would expect these talks to be contentious, however most policies pass without much debate. Even if the vote was preceded by long discussion, residents almost always approve of the policy come voting time.

Thus, it is reasonable to question whether consensus voting protects minority opinions as well as its practitioners might hope. There are reasons to suspect otherwise: debating policies in front of the entire house takes considerable time and effort. Moreover, in these small to moderate sized communities there may be a social cost to dissenting from the majority opinion. Minority voices – or those who perceive themselves as so – may not end up speaking up, or very strong minority voices may drown them out.

Columbae is a community with specific ideals and a simple voting mechanism, however it is unclear whether their voting scheme works as intended. This scenario highlights a gap in normative social choice theory: while the literature has various justifications for different aggregation rules on the basis of specific theoretical desiderata, it doesn't have suggestions for when groups have a particular ideology that it wants to enforce.

This is the setting which we explore in our paper. We examine whether group members participating in consensus are honest in their revelation of preferences; furthermore, we present a model to better understand the conditions under which dishonest voting occurs. In answering these questions, our contributions are three-fold:

- First, we have empirical data from surveying Columbae community members which approximate the true preferences of consensus participants. Our findings show that people do not vote truthfully even if they are in the majority, which we argue is due to the magnified effect of perceived social risk in the scheme.
- Secondly, we contribute a game theoretic model which explicitly models perceived social risk. While the game theoretic literature on deliberative democracy and, specifically, on this variant of consensus, recognize that individuals are incentivized to withhold their true preferences, our model shows the concerning degree to which individuals will keep their true beliefs private, supporting the data from our first contribution.
- Thirdly, in search of an alternative mechanism that encourages true-preference revelation and protects minorities, we explore how a new scheme, *quadratic*

voting, would fare in this community by surveying the same members. QV has created excitement in the social choice community since its founding in 2013 but empirical support for it is sparse. Analysis of our data implies that QV may better achieve Columbae's desire to protect and give voice to minorities than consensus, which should encourage future research in QV and its implementation.

We now outline the rest of the paper. In section 2 we present the relevant literature on deliberative democracy, game theoretic models of consensus voting schemes, and quadratic voting. Section 3 contains our empirical findings on the truthfulness of voters who participated in consensus voting. In section 4, we present a game theoretic model of consensus voting which takes into account social factors and characterize its equilibria. Section 6 presents the results of a simulated vote using QV in the same community. Finally, in sections 7 and 8 we compare QV and CV using a maximin social welfare function and discuss the possible benefits of applying QV in a deliberative context.

2 Related Work

In this section we given an overview of the relevant literature from *social choice theory* – the study of aggregating individual preferences – and on *deliberative democracy* – the study of sharing of preferences between individuals. We then introduce *quadratic voting* as a new alternative mechanism to normative voting schemes.

The most commonly used procedure for binary choice decisions is majority rule (MR), and social choice theory provides extensive, formal justification for this. When assuming a setting with rational agents who, given binary options, aim to maximize their utility, MR is the only rule that uniquely satisfy some desiderata (May's Theorem) [9]. It is also clearly strategyproof: if one votes dishonestly, she only increases her chance of a less-preferred outcome. There exists, too, a consequentialist argument for MR: assuming an informed and rational populace, the probability of a correct majority decision converges to 1 as the number of voters increases (Condorcet's Jury Theorem) [9]. Finally, by the Rae-Taylor theorem, MR is the rule that maximizes an individual voters expected utility [13].

The work done in social choice theory may make MR seem like an obvious choice for an aggregation rule. But reality poses other problems: MR is vulnerable to the "tyranny of the majority," in which a ruling majority stronghold the votes or policies for their own interests, leaving the minority party with few devices to overturn the vote.

In small groups, this may be mitigated by what organically occurs in practice prior to or during a vote: voters are given a chance to justify their preferences. A session of deliberation may be held to "promote substantive, balanced, and civil discussion," and allows these minorities to convince others to change their minds [3].

Forms of democratic decision-making that involve deliberation are studied under the umbrella of *deliberative democracy* (DD). The study of DD takes into account the deliberative aspect of decision-making in contrast to social choice theory, which ignores this component and focuses on just the aggregation rule. The variants of DD take two general forms: a *mixed deliberative model*, where there is time for discussion prior to an aggregation of preferences, or a *pure deliberative model*, where group members talk until they naturally arrive at a resolution.

Though it is just one form of DD, we focus on the mixed model that uses consensus voting as its aggregation rule due to its prominence in practice among small to moderate-sized communities. Groups like Columbae, that operate as *co-operative houses* or *cooperative communities*, are one key use case for this model, which we shall refer to as *consensus*. The Foundation for Intentional Community, a 501(c)(3) nonprofit that advises and funds cooperative communities, has published extensive literature on consensus. In their recommended handbook on implementing consensus from the Center for Conflict Resolution, the authors write that consensus is “a powerful tool for building group unity and strength” [2]. The *New York Times* reported on the recent, multi-generational rise in co-operative communities across the nation [14]. In other words, the specific case of the Columbae community is broadly applicable and relevant to other real-life settings.

2.1 Previous Work on Modeling Consensus

Again for shorthand, we refer to the model of consensus with deliberation as just consensus or CV. To preface our own model of consensus, we give an overview of the literature on game theoretic models of deliberative democracy.

Consensus theoretically should protect minority votes: only one person needs to vote against the policy in order for it to be overturned. It is also strategyproof: voting against ones honest preference is almost surely harmful, and so in its abstract form, it’s a sensical choice for communities like Columbae.

But when accounting for the payoff associated with revealing one’s true preferences in deliberation, consensus should no longer be as appealing of an option. Previous work focuses on the truthfulness of voters under consensus, or lack there-of. Rarely does it encourage honest revelation of preferences: Feddersen and Pesendorfer show that truthfulness is only a dominant strategy in the case where voters share underlying preferences [6]. Austen-Smith and Feddersen show that full truthfulness in consensus amounts to it being common knowledge that everyone already shares identical preferences [1]. It then follows that this equilibrium exists under a deliberative model with any other aggregation rule. Meanwhile, in the most usual case where others preferences are unknown, majority rule encourages deliberation while consensus does not. Put simply, “unanimity rule seems uniquely bad at incentivizing truth-telling in deliberation” [3], and it’s rather arbitrary for communities to herald it as the best option.

Other undesirable effects of the consensus model include the fact that Condorcet's Jury Theorem no longer holds and so the probability of "incorrect" decisions may actually increase with the number of participants [5][6].

The procedure itself may be time and energy-consuming: it's intuitive that the time taken to search for a policy deserving unanimous approval may grow – more members means possibly more deliberation. CV may also encourage strategic bargaining when searching for an optimal policy: that is, if voter A knows what voter B wants in a policy, she may hold off on voting until she gets B to comply with her [13]. But contrary to these efficiency concerns, consensus procedures in Columbae, by observation, take little time. Most votes on a policy receive all fives and the issue of strategic bargaining does not occur. To corroborate this observation, another empirical study found that "if anything, efficiency improved with larger groups" [7].

In summary, while many small to moderate communities may choose consensus as its social choice function to enforce their ideals, the scheme is highly unreliable and incentivizes undesirable behavior from its participants. The extant literature however leaves room for us to explore under what specific conditions we can expect to observe undesirable, dishonest behavior.

2.2 Quadratic Voting

Despite its noble aim of allowing a community to freely discuss the policies at hand, we will see that CV in practice has various drawbacks. We explore an alternative mechanism, namely *quadratic voting*, proposed by Weyl in 2013 [15].

QV, like CV, aims to address the issues posed by tyranny of the majority. While CV relies on the community members to regulate the procedure (hence why social factors are embedded in the system), QV treats a referendum like an auction. So while CV places additional burden of social pressure on minorities[1], under QV, people are expected to *pay* the price they put on their desired outcome, in addition to an externality. Lalley and Weyl prove that the price a person should pay is equivalent to the externality she imposes on others for her desired outcome. This externality grows quadratically, therefore the cost of casting votes should grow quadratically as well [16]. QV then avoids the "tyranny of the majority" problem by enabling persons to have more than just a single vote, yet it does not give complete veto power since the cost of voting multiple times grows quickly.

2.3 QV In Practice

The theoretical literature on QV implies that it should protect minorities against disenfranchisement as opposed to other rules [11]. Previous simulations of QV,

[1] We clarify the use of the vague term "minority." What we imply is a set of persons at unique disadvantage compared to the rest of the population, outside of the fact that they constitute a minority in population.

using data on votes from past referendums, make minority-victories more frequent and result in higher average welfare [4].

Additionally, QV allows voters to vote more "expressively" than in other traditional voting mechanisms, where voters are allocated just one vote per policy ("one-person, one-vote" or "1p1v"). This, in conjunction with QV's policy of scarcity, forces voters to think more carefully about their actions than in other mechanisms. For these reasons, the literature on QV suggests that QV is also a useful survey tool, as it is more accurate in soliciting valuation of true preferences than traditional Likert surveys [8][10][12] .

While the theoretical properties of QV make it a promising candidate for communities that aim to protect minorities, it is a recent invention and so the research on how it fares in practice is scarce: the work on its implementation mostly relates to its use as a survey tool. This leaves room to examine how QV may perform in cooperative communities like Columbae.

3 Evaluation of True Preferences

We begin with an attempt to evaluate the behavior of voters under the consensus model, and specifically to test the literature's hypothesis that voters are more likely to vote dishonestly, as outlined in §2.

3.1 Data and Methods

We surveyed residents of Columbae on the policies that were voted upon in a meeting held in September 2018. Members at this particular, quarterly meeting, called "Meat, Smoke, and Noise" (MSN), decided on the house's most fundamental policies regarding consumption of meat, where/when people are allowed to smoke, and what the quiet hours of the house are. Of the 55 existing members, 42 responded, of which 34 were present.

The survey, which was put forth six weeks after the initial meeting, consisted of 10 policy proposals, enumerated from 0 to 9. Users were asked to rate the strength of agreement or disagreement on a Likert scale: "1" denoting strong disagreement and "5" being strong agreement, with "3" as indifference. We consider a vote below 3 to be a "negative" vote and a vote above as a "positive" vote. At the actual MSN meeting, all policies passed.[2]

3.2 Results

In our discussion going forward we shall use only the results of participants who are residents and were present at MSN. We are assuming that the opinions

[2] We can say that at the actual meeting, everyone expressed "fours and fives". This is because the specific implementation of consensus in Columbae allows for the revelation of preference on a scale of 0-5. Fours and fives mean approval of the policy.

expressed of the participants were held when the meeting occurred and acknowledge that perhaps not all of the participants felt as negatively about certain policies at the time, but use it as a proxy.

With this assumption, the results of the Likert survey show that under CV, Columbae members at times keep private their true opinions. Table 1 presents the median vote per policy and a simulated MR vote according to the Likert preferences.

Policy	**Median**	**Positive votes - negative votes**
0	5	24
1	**2**	**-20**
2	4	18
3	**2**	**-13**
4	3.5	13
5	4	17
6	3	4
7	3	3
8	4	24
9	4	11

Table 1. Likert survey results. A vote above 3 is "positive," and below is "negative." Bolded rows indicate results that would have overturned CV outcomes.

The frequencies with which people voted negatively on policies 1 and 3 is especially significant. More than half meeting attendees voted for a policy they otherwise strongly disagreed with. We show the histogram of Likert responses for policy 1 in Figure 1 (other results in Appendix B).

If we count one person's vote on a policy as one vote, there were 90 votes that were against true preferences compared with 171 votes that were in line with true preferences – that is, one third of the time, people were motivated to vote strictly against their true preferences. There were 79 votes that truthfully indifferent, but voted in favor anyways. So, 26% of the time, participants voted strictly against their true preferences (e.g. they voted for a policy when they disagreed with it), and only 50% of participants were strictly in favor of policies that they voted for. These results indicate that voters are highly incentivized to withhold their true preferences.

4 Modeling CV in the Presence of Social Factors

We incorporate social factors into a game theoretic model of consensus, motivated by the results in §3. Consider the individuals participating in CV as players with the aim of maximizing their individual payoff given the behavior of the other players. We hypothesize that the observed dishonest voting is due to the mechanism's component of public deliberation: in other words, the payoff of

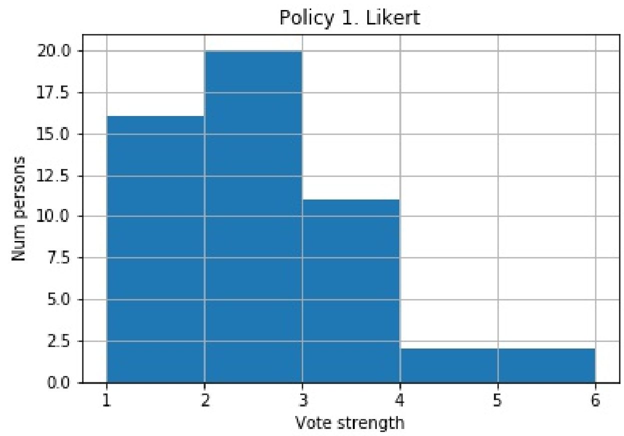

Fig. 1. Histogram of Likert responses for Policy 1.

an individual may also depend on how their public votes will be perceived by their community members. With this in mind, we propose a model that embeds social factors into an individual's welfare function in CV: one that describes how a voter may maximize their utility by harboring their dissenting views.

Definition 1. *Let the welfare of individual i in CV be written as,*

$$W_i = \sum_{j=1}^{m} W_{ij}(v), \tag{1}$$

for W_{ij} being the welfare function for the individual for the jth policy, expressed as

$$W_{ij}(v) = (h_{ij} - \frac{1}{2})r_j - \frac{k_i}{n-1} \sum_{b \neq i}^{n} (v_{bj} - v)^2 \tag{2}$$

with

- $v \in \{0, 1\}$ *wheres* $v = 0$ *denotes a negative vote, and* $v = 1$ *a positive vote;*
- $h_{ij} \in [0, 1]$, *uniformly distributed, denoting i's benefit ($h_{ij} > \frac{1}{2}$) or harm ($h_{ij} < \frac{1}{2}$) imposed by policy j;*
- r_j *being whether policy j passes (1 if it does, -1 else), and*
- $k_i \geq 0$ *being i's "self-consciousness" coefficient, between 0 and 1.*

We now focus on the case where the k_i are constant among all participants. We look for symmetric Nash equilibrium in players use a particularly simple strategy: we say that a player i uses a *cutoff strategy* with threshold c if they vote 1 if $h_i > c$ and 0 otherwise.

We will refer to this c as an "indifference threshold," since in such an equilibrium a player i must be indifferent to voting 1 or 0 when $h_i = c$. The following result characterizes these types of equilibria:

Theorem 1. *In this model of consensus voting there exists a symmetric Nash equilibrium consisting of cutoff strategies with indifference threshold* $c = 1 - k^{\frac{1}{n-1}}$.

Proof. See Appendix A.

We should note that there also exist equilibrium with threshold $c = \frac{1}{2}$ (in which everyone votes honestly) and $c = 1$ (in which everyone votes against the policy).

In Figure 2 we plot the indifference threshold specified in Theorem 1 with different n and k and observe that as n increases, c decays drastically for a fixed k-coefficient. Therefore despite having the same "self-consciousness factor" for both small and large n, the effect on voting behavior is much more pronounced when n is large.

The most basic case of $n = 2$ yields a Nash equilibrium where the cutoff decreases linearly, namely $c = 1 - k$. Because the payoffs of the players become highly coupled as k increases, they are able to sustain an equilibrium where the policy passes with high probability, even if both players actually would prefer to privately vote no. Intuitively, their fear of disagreeing with each other in effect prevents them from being completely truthful and overturning the vote in their favor, a phenomenon which we observe in our empiricla data.

To summarize, this model shows that for moderate n (e.g. 16, 32), even a very small k coefficient can sustain a cut-off significantly below 1/2, indicating that people may vote more enthusiastically than they really feel due to social pressures. In particular, this matches the behavior we observe in the Columbae voting results.

5 Evaluating QV in Columbae

We have given theoretical and empirical evidence which suggests that consensus fails to enforce the ideals of cooperative-communities. In this section we move on to simulate quadratic voting to explore how this alternate mechanism fares in comparison to consensus in small groups.

5.1 Methods

We asked participants of the first survey to fill out a second survey via a platform provided by the engineers at weDesign. The platform brings the user to a video which instructs them on how to use the platform to vote on the same 10 policies as in the first survey. There are two buttons next to each policy, one with a "thumbs-up" and another with a "thumbs-down" sign. The user can click on these buttons, with one click indicating one vote, with each sequential vote for a given policy being more expensive (the first vote being worth one "credit,"

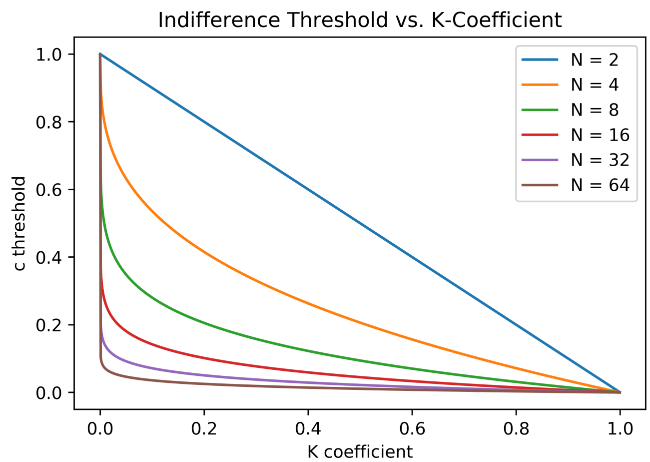

Fig. 2. As k or N grows, the threshold of indifference drops sharply.

the second worth four, the third worth nine, and so on). The user can see the number of "credits" they have left at the top so that they don't have to calculate how much they have left themselves. Each person is initially given 100 credits to spend.

5.2 Results

The results of the QV simulation are in Table 2. For each policy, we list the sum of votes and the median vote. We also provide a simulated MR vote based on the QV data, which scales each person's vote to -1, 0, 1 based on the sign of their votes. In the last row is the sum and median number of leftover credits.

As shown in Table 2, the two policies that would have been negated by the simulated MR did not pass in QV. Additionally, one of the more contentious policies, policy 6, would have been overturned. We provide the histogram of votes in Appendix B in Figures 3-6.

In the next section we analyze these results under a specific welfare function to compare with the consensus vote.

6 Comparing CV and QV

Given the results of simulated MR, QV, and CV, how do we measure which was most effective in protecting minorities? In this section we argue for a specific outcome, that produced by a maximin function, and compare it to the results of CV and QV.

Policy	**Sum of QV Votes**	**Median**	**Sim. MR**
0	108	3	30
1	**-51**	**-2**	**-21**
2	38	1	20
3	**-40**	**-1**	**-15**
4	27	1	14
5	35	1	15
6	**-14**	**0**	**0**
7	28	1	7
8	83	2.5	22
9	32	1.0	14
Credits left	700	8.5	–

Table 2. Results of QV survey. The Simulated MR column denotes the number of positively-signed allocations subtracted by number of negatively-signed allocations. The bolded rows indicate the policies that would have overturned CV results.

Like voting functins, measures of community welfare also depends on the community's chosen priorities. For example in a utilitarian community – where individuals are treated equally and the goal is to simply maximize overall welfare – it's easy to see how the needs of a majority may supercede those of a more-needy minority. This ideal of overall welfare would clearly not match the priorities of a community like Columbae's. Instead, we argue that an analysis of overall welfare by a maximin function best suits Columbae's purposes.

Definition 2. *Consider a profile $\langle W_1, \ldots, W_n \rangle$ for $W_i(x)$ denoting i's welfare as a result of some outcome x. x is a **maximin** outcome if for every other alternative y, x is (weakly) preferred to y if $\min_{i \in N} W_i(x) \geq \min_{i \in N} W_i(y)$ [9].*

As put by Economist M.E. Yaari and Psychologist Maya Hillel in their empirical study of various welfare functions, an outcome is maximin when "[dividing] a given bundle of goods in such a way that, after the division, the position of the least advantaged individual shall be as high as possible" [17]. That is, a voting aggregation function is deemed maximin efficient if its outcome is the maximin outcome.

CV's priorities are in line with maximin's deference to the neediest voter, as Yaari and Hillel's study show empirically. They surveyed a population to gauge the most intuitively "fair" way of dividing goods between two people with their own interests. When the players had different "needs", e.g. necessary nutritional quota, the consensus was that splitting the goods "maximin" was the most fair [17]. As CV aims to protect those with greatest needs, maximin is a logical choice for optimizing welfare.

6.1 Approximating Need

We use the results of the QV survey to approximate what each person's "need" is. As each person started with the same number of credits, we claim that the number of credits spent are a reasonable proxy for the reasons stated in §2.3.

Although it may seem self-referential to use the spending under QV to evaluate QV's own performance, it is not necessarily the case that QV should still be maximin-efficient for every policy: one can easily imagine the case where the person with the most need is outweighed by a larger group of people with less intense preferences.

6.2 Results

In Table 3 we have the $\max_j$ and $\min_j$ votes cast for each policy j by any person in the QV survey. We determine that if $\max_j + \min_j > 0$, the policy's passage is in agreement with the maximin outcome. If $\max_j + \min_j = 0$, as is the case for Policy 9, we also list the second to largest and second smallest values.

Policy	**Max vote**	**Min vote**	**Max + Min**
0	7	-1	6
1	2	-5	-3
2	6	-5	1
3	2	-6	-4
4	3	-2	1
5	4	-4	0
6	4	-7	-3
7	8	-6	2
8	7	-5	2
9	5,4	-5,-3	0,1

Table 3. Max and min votes allocated per policy in QV survey.

As predicted by our model, CV fares poorly by maximin. CV's outcome did not compensate the "neediest-person" on three occasions: on policies 1, 3, and 6. Interestingly, policies 1 and 3 also happened to have majority-negative opinion, as shown in Table 1. That is, the majority disagreed with the these policies and yet they passed. This observation supports our model's hypothesis that the probability of truthful-voting falls to 0 for slightly self-conscious voters in moderate-sized communities.

In contrast, QV performed better according to maximin. It makes sense that QV did so: needy voters under QV are willing and able to purchase more votes.

7 Suggestions for Columbae, and Similar Communities

In this section we use our results to inform our suggestions for communities like Columbae, that have the goal of enforcing their particular ideologies through a chosen voting mechanism.

QV's maximin efficiency relative to consensus makes it an optimistic candidate for cooperative communities. However QV relies on the scarcity of voice-credits to force voters to be discriminating in their decision-making. The trade-off to this is that a voter is unable to express enthusiasm (or strong disapproval) for many separate policies. And if credits are to be used over the course of a year, it may be difficult for voters to presently know what they should be saving up for.

However the simulation of QV in Columbae shows that house members are more scarce with their credits than they need to be. While members could have spent all their credits – after all, there was no followup referendum – only eleven people out of the thirty-four had less than five credits leftover (5% of original amount). Some users, when asked, reported that they did not spend all their credits because it did not feel "right" to swing a vote so much in their preferred directions, since they did not actually care that much. This shows that for at least a community sized as Columbae's, having an excess of credits may not pose the issue of freely spending on votes, since members are sensitive to how their votes impact others.

Given this qualitative evidence, perhaps an allocation of credits that expires at the end of each referendum makes the most sense for the Columbae community, so to not limit expression of enthusiasm between referenda.[3] Having credits expire and be newly allocated could mitigate the possibility of collusion with other house members outside of meetings and so the voting function could more accurately gauge what members' "true" preferences are.

7.1 Opportunity for a New Model: Deliberative QV

We note that the QV model that was tested was a private vote and lacked consensus' aspect of deliberation. Given that the process of discussion is essential to these groups, it would seem that communities may be resistant to forgoing the tradition of deliberation and would prefer still consensus regardless of QV's other beneficial traits. However we contend that QV is not irreconcilable with deliberative democracy: perhaps a deliberative form of QV can be implemented. Namely, voters may participate in a deliberation session prior to the vote and then use QV to aggregate their preferences.

Of course, one could implement a deliberation session prior to any aggregation rule. But we argue that QV as the chosen rule would actually encourage

[3] One could argue that this is equivalent to a Borda count – however, allowing voters to decide how many policies to support, to what extent, and to alternatively abstain, allows them to express their opinions with a granularity that Borda and other systems cannot offer.

deliberation more than, say, MR. Imagine the scenario where a soft majority, say 60% are in favor of a given policy. They care about the policy initially, say with $h = 0.6$. Let the minority have $h = 0.2$, as in they are strongly against the policy. Because minority people have soft veto power given by QV (by choosing to spend large amount of votes), they can still incur a large cost on the majority. This is in contrast to the majority vote case, where no matter what the minority does, the majority can win regardless. However in QV, the majority has a large incentive to try to rationalize their position to minority to reduce the cost. The minority has incentive to deliberate for similar reasons: to reduce the amount they have to spend to overturn policy.

Thus it seems worthwhile for communities like Columbae to consider implementing a new form of deliberative democracy that instead uses QV as its aggregation rule, to both improve the quality of deliberation and incentivize true preference revelation in the end.

8 Conclusion

Practitioners of deliberative democracy often choose consensus with the assumption that participants will be truthful. We examine this assumption by surveying current Columbae members on their "true" opinions of policies. We conclude that a model for CV should incorporate social factors like the cost of disagreeing with others. The presence of a dishonest equilibrium showed that social pressures can motivate voters to hide their true preferences, even when they are actually in the majority.

By using a maximin overall welfare function, we conclude that CV does not meet its aims of protecting minorities or deferring to the most disadvantaged.

We tested how an alternative mechanism called Quadratic Voting would fare in comparison to CV and are optimistic about its being potentially more maximin efficient than CV. While social pressures in CV in effect squash the preferences of voters, QV seems to better reflect voter preferences and compensate those with greater need. Additionally, we argue that a deliberative form of QV would satisfy the need for both deliberation and truthfulness in such communities.

Borne out of the desire to mitigate inequality, compensate for people's needs, and protect minorities, deliberation with a consensus vote has instead become an institution tends to obscure individual preferences. By removing the assumptions of honest voting and explicitly modeling social pressure, we have shown both empirically and theoretically that other methods like QV may better carry out CV's aims.

9 Acknowledgements

Thank you to Douglas von Kohorn at weDesign and Glen Weyl at Microsoft Research for their generous help with implementing the QV platform, in addition to Tim Roughgarden (Columbia University) and Alex Slivkins (Microsoft Research New York) for their guidance and mentorship.

References

1. D. AUSTEN-SMITH and T. J. FEDDERSEN. Deliberation, preference uncertainty, and voting rules. *American Political Science Review*, 100(2):209217, 2006.
2. M. Avery. *Building united judgment: a handbook for consensus decision making.* Center for Conflict Resolution, 1991.
3. A. Bchtiger, J. S. Dryzek, J. Mansbridge, M. Warren, and C. List. Democratic deliberation and social choice: A review. 10 2018.
4. A. Casella and L. Sanchez. Storable Votes and Quadratic Voting. An Experiment on Four California Propositions. NBER Working Papers 25510, National Bureau of Economic Research, Inc, Jan. 2019.
5. F. Dietrich and K. Spiekermann. Epistemic democracy with defensible premises. Technical report, 2013.
6. T. Feddersen and W. Pesendorfer. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *The American Political Science Review*, 92(1):23–35, 1998.
7. E. Hoffman and M. L. Spitzer. The coase theorem: Some experimental tests. *The Journal of Law Economics*, 25(1):73–98, 1982.
8. L. Kaplow and S. D. Kominers. Who will vote quadratically? voter turnout and votes cast under quadratic voting. *Public Choice*, 172(1-2):125 – 149, 2017.
9. C. List. Social choice theory. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2013 edition, 2013.
10. J. S. Masur. Quadratic voting as an input to cost-benefit analysis. *Public Choice*, 172(1-2):177 – 193, 2017.
11. E. A. Posner and E. G. Weyl. Quadratic voting as efficient corporate governance. *The University of Chicago Law Review*, 81(1):251–272, 2014.
12. D. Quarfoot, D. von Kohorn, K. Slavin, R. Sutherland, D. Goldstein, and E. Konar. Quadratic voting in the wild: Real people, real votes. *Public Choice*, 172(1-2):283 – 303, 2017.
13. W. Shughart II. Dennis c. mueller, public choice iii. *Public Choice*, 118:469–473, 02 2004.
14. T. Verde. Theres community and consensus. but its no commune. *The New York Times*, Jan 2018.
15. G. Weyl and S. Lalley. Quadratic voting: How mechanism design can radicalize democracy. *American Economic Association Papers and Proceedings*, 1(1), 12 2017.
16. G. Weyl and E. Posner. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society.* 2018.
17. M. E. Yaari and M. Bar-Hillel. On dividing justly. *Social Choice and Welfare*, 1(1):1–24, 1984.

A Original Proofs

A.1 Proof of Theorem 1

In the following, let the number of players be n. Again, for simplicity we assume that a fixed k coefficient is shared between all players. We drop the subscript j and just consider a single policy throughout.

Consider a player i. The indifference condition at c implies that when $h_i = c$, we necessarily have

$$W_i(0) = W_i(1) \tag{3}$$

$$\begin{aligned}\Rightarrow W_i(0) &= W_i(1|r_{-i} = 1)P(r_{-i} = 1)\\ &\quad + W_i(1|r_{-i} = -1)P(r_{-i} = -1)\end{aligned} \tag{4}$$

(here r_{-i} denotes the outcome of the vote among all the players except i, i.e. $r_{-i} = 1$ if all other players vote 1 and -1 otherwise). We now compute these utilities.

With $h_i = c$, we have:

$$W_i(0) = 1/2 - c - \frac{k}{n-1}E\big[\sum_{i' \neq i}^{n} v_{i'}^2\big] \tag{5}$$

$$= 1/2 - c - \frac{k}{n-1}\big(\sum_{i' \neq i}^{n} E[v_{i'}^2]\big) \tag{6}$$

$$= 1/2 - c - \frac{k}{n-1}\big(\sum_{i' \neq i}^{n} E[v_{i'}]\big) \tag{7}$$

$$= 1/2 - c - k(1-c), \tag{8}$$

where (8) follows from the fact that $v_j^2 = v_j$, since it is an indicator variable. Note that the expectation here is taken over the randomness in the preferences of the other players, which are distributed uniformly in $[0, 1]$.

As for the utility of voting 1 given $r_{-i} = 1$, we have:

$$W_i(1|r_{-i} = 1) = c - 1/2 - \frac{k}{n-1}E\big[\sum_{i' \neq i}^{n}(v_{i'} - 1)^2 | r_{-i} = 1\big] \tag{9}$$

$$= c - 1/2 - \frac{k}{n-1}\sum_{i' \neq i}^{n}(1-1)^2 \tag{10}$$

$$= c - 1/2, \tag{11}$$

where (11) follows from the fact that all votes must be 1 if $r_{-i} = 1$.

To find the probability of the event that $r_{-i} = 1$, since all h_i are independent and uniformly distributed between 0 and 1, it follows that $P(r_{-i} = 1) = (1 - c)^{n-1.}$

Now for the other conditional payoff:

$$
\begin{aligned}
W_i(1|r_{-i}=-1) =& 1/2 - c \\
& - \frac{k}{n-1} E\big[\sum_{i' \neq i}^{n} (v_{i'} - 1)^2 | r_{-i} = -1\big] \qquad (12) \\
=& 1/2 - c \\
& - \frac{k}{n-1} \big(\sum_{i' \neq i}^{n} E[(v_{i'} - 1)^2 | r_{-i} = -1]\big) \qquad (13) \\
=& 1/2 - c \\
& - \frac{k}{n-1} \big(\sum_{i' \neq i}^{n} E[\mathbb{I}\{v_{i'} = 0\} | r_{-i} = -1]\big) \qquad (14) \\
=& 1/2 - c \\
& - \frac{k}{n-1} \big(\sum_{i' \neq i}^{n} P(v_{i'} = 0 | r_{-i} = -1)\big). \qquad (15)
\end{aligned}
$$

Now we compute the conditional probability:

$$
\begin{aligned}
P(v_{i'} = 0 | r_{-i} = -1) &= \frac{P(r_{-i} = -1 | v_{i'} = 0) P(v_{i'} = 0)}{P(r_{-i} = -1)} \qquad (16) \\
&= \frac{c}{1 - (1-c)^{n-1}}. \qquad (17)
\end{aligned}
$$

Plugging this back into (15), we see that

$$
\begin{aligned}
W_i(1|r_{-i} = -1) &= \frac{1}{2} - c - \frac{k}{n-1} (\sum_{i' \neq i}^{n} P(v_{i'} = 0 | r_{-i} = -1)) \qquad (18) \\
&= \frac{1}{2} - c - \frac{k}{n-1} (\sum_{i' \neq i}^{n} \frac{c}{1 - (1-c)^{n-1}}) \qquad (19) \\
&= \frac{1}{2} - c - \frac{kc}{1 - (1-c)^{n-1}}. \qquad (20)
\end{aligned}
$$

Also note that

$$
P(r_{-i} = -1) = 1 - P(r_{-i} = 1) = 1 - (1-c)^{n-1}. \qquad (21)
$$

After inserting these expressions into (4) and simplifying, we get the condition:

$$
2(c - 1/2)((1-c)^{n-1} - k) = 0. \qquad (22)
$$

Therefore the cutoff symmetric Nash equilibria are given by the roots of this polynomial. We see that $c = 1/2$, the "honest" solution, always exists. However, we also have the solution

$$
c = 1 - k^{\frac{1}{n-1}}, \qquad (23)
$$

which is our "dishonest" equilibrium. This completes the proof. ■

B Supplementary figures

In this auxiliary section we provide the histograms for the Likert survey responses (on the left) and the corresponding histograms from the QV survey responses (on the right). We also provide at the end the histogram of number of people who have a given number of leftover credits after completing the QV survey.

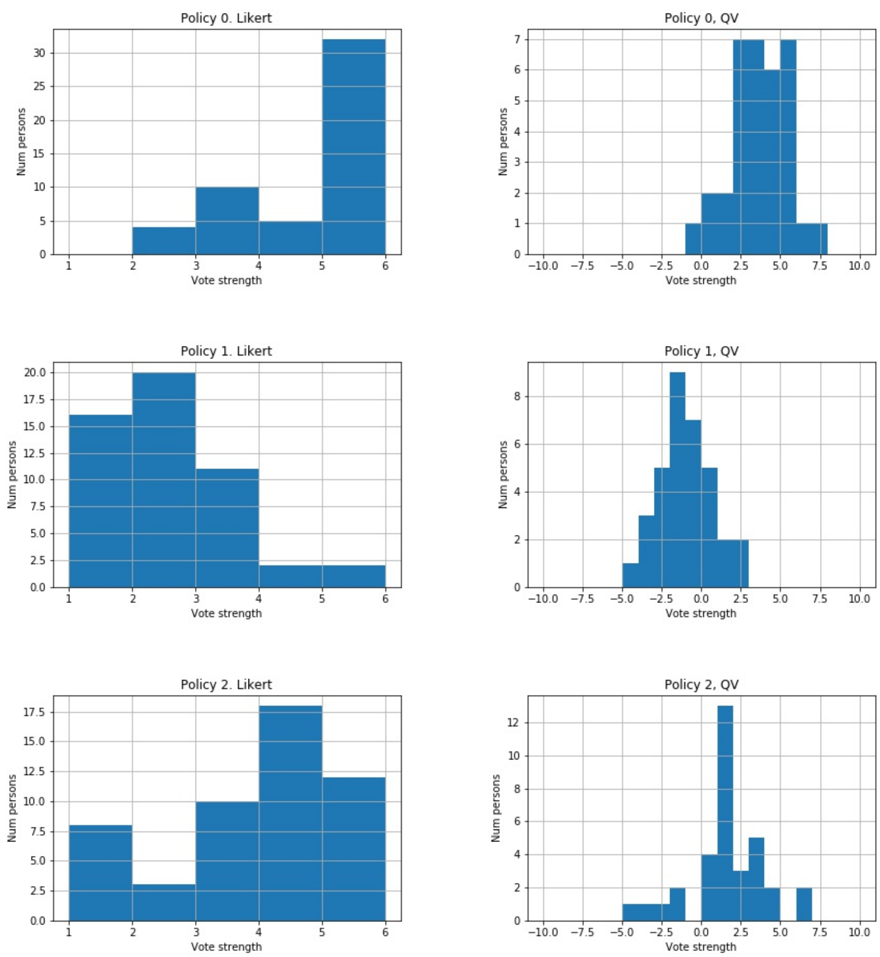

Fig. 3. Vote results for policies 0-2, Likert and QV

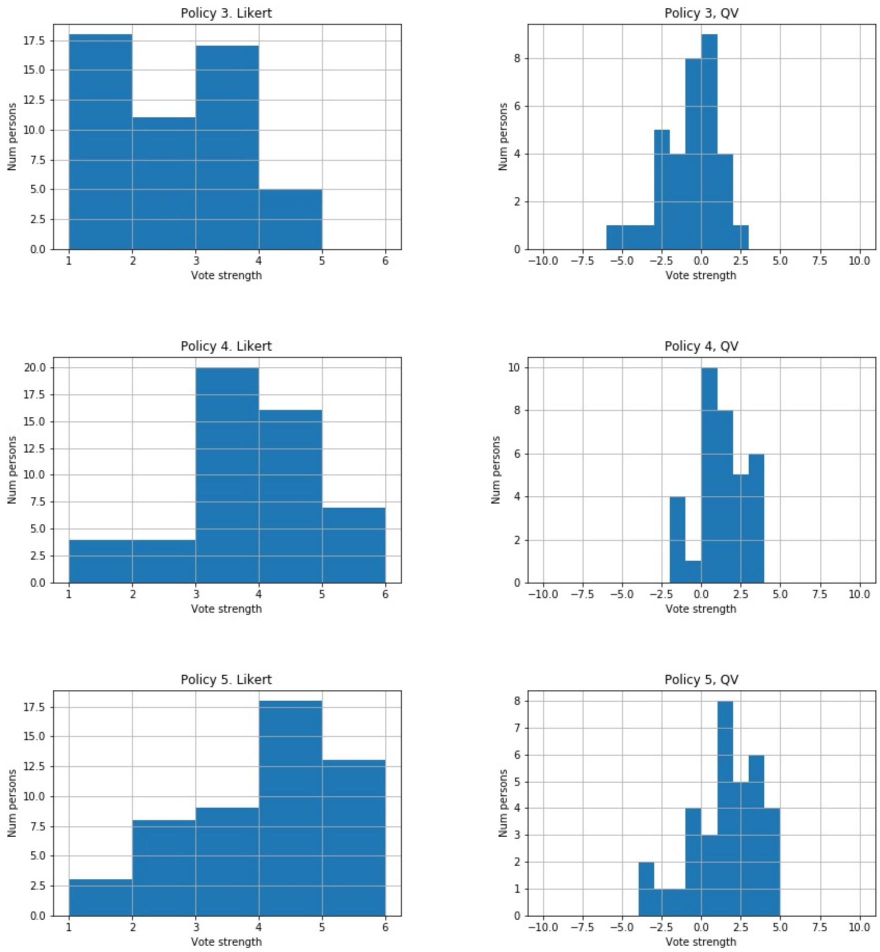

Fig. 4. Vote results for policies 3-5, Likert and QV

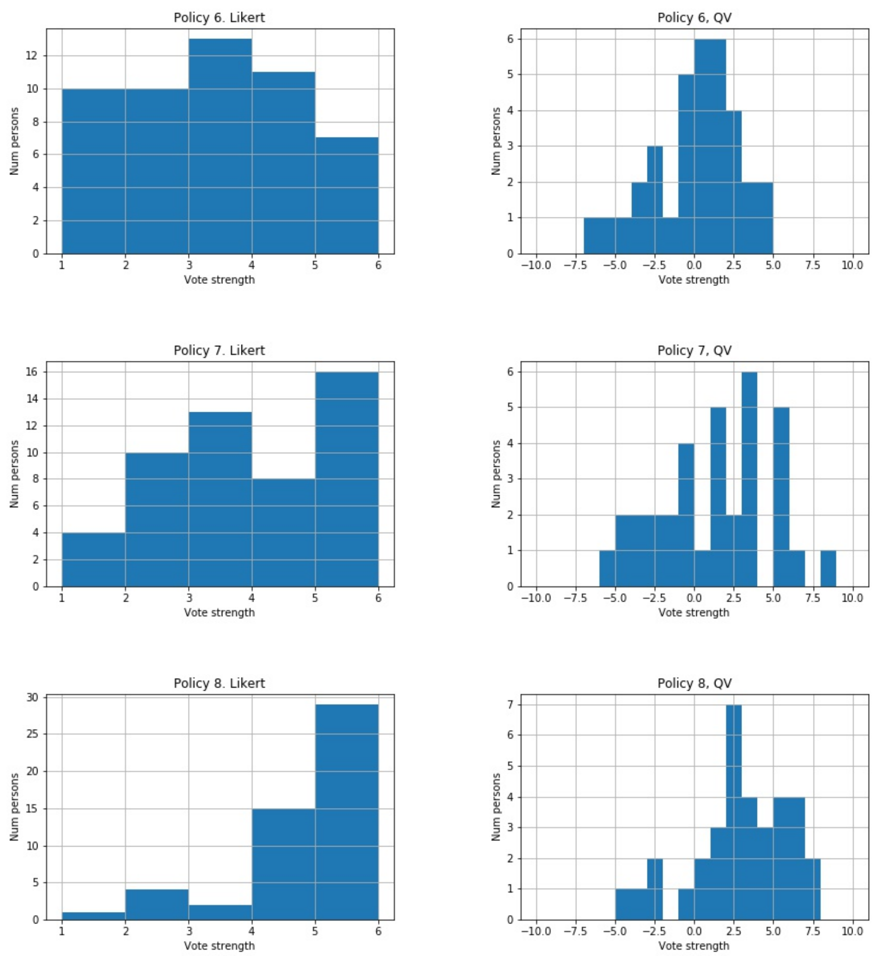

Fig. 5. Vote results for policies 6-8, Likert and QV

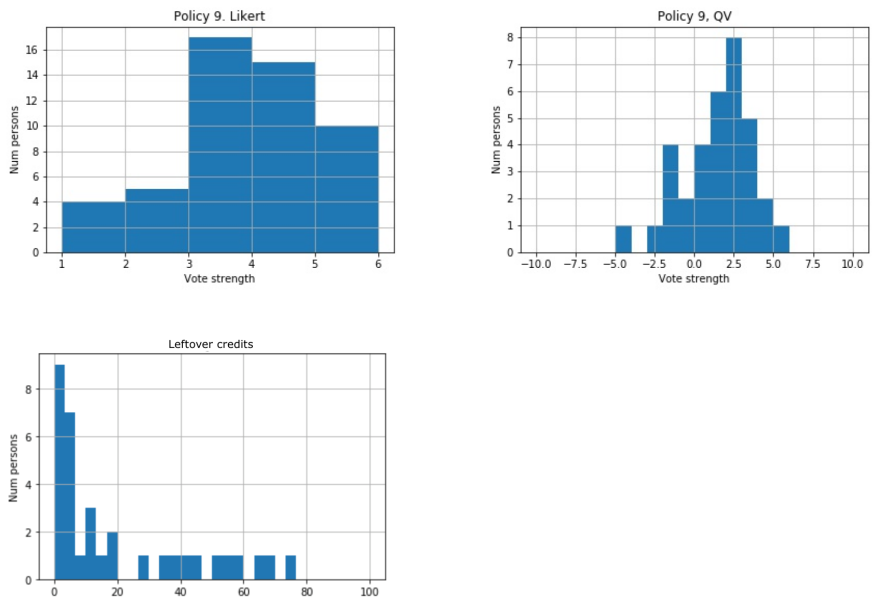

Fig. 6. Vote results for policy 9 and leftover credits from the QV vote