Agent-Based Models of Quadratic Voting

Jacob Kelter, Andreas Bugler, and Uri Wilensky

1 Introduction and Research Questions

A democracy must have a voting system, and there are multiple to choose from. None is perfect, each with strengths and limitations, depending on prevailing conditions. In this paper, we compare a recently proposed voting system, called quadratic voting (QV), with the more common majority rule system. There are many metrics that can be used to evaluate a voting system. A common one is to measure the "population welfare" resulting from a vote, defined as the sum of each individual's welfare. An individual who benefited from the outcome has a positive welfare and one who was harmed has a negative welfare, with the magnitude determined by the amount of benefit or harm.

The most common voting mechanism, familiar in democracies around the globe, is one person one vote (1p1v). In binary decisions, 1p1v results in majority rule, which is equivalent to choosing the preference of the median voter since voters can only express the direction of their preference but not the intensity. If the majority of the population is only weakly opposed to an outcome and a minority is strongly in favor, 1p1v will fail to maximize welfare by siding with the majority. In the extreme, this can result in "tyranny of the majority," in which the minority's fundamental rights are violated. Attempts to avoid such tyranny are as old as modern democracy itself. John Adams, recognizing this problem would occur in a government consisting of a single elected body, argued for a three-branch government to protect against it [1].

J. Kelter (✉) · A. Bugler · U. Wilensky
Northwestern University, 633 Clark St, Evanston 60208, USA
e-mail: jacobkelter@u.northwestern.edu

A. Bugler
e-mail: andreas@u.northwestern.edu

U. Wilensky
e-mail: uri@northwestern.edu

Z. Yang and E. von Briesen (eds.), *Proceedings of the 2020 Conference of The Computational Social Science Society of the Americas*, Springer Proceedings in Complexity, https://doi.org/10.1007/978-3-030-83418-0_8

Recently, an alternative voting mechanism known as quadratic voting (QV) has been proposed, aiming to maximize the aggregate welfare of a collective binary (yes/no) decision by allowing voters to express the intensity of their preferences, not just the direction [2, 3].

In the original formulation of QV [10], votes are purchased for money at a cost of v^2 dollars to cast v votes. The money raised is then divided evenly among voters and returned to them. On average, this results in the losing side of the vote gaining some money as consolation. Voting with money has the advantage of allowing citizens to make trade-offs between private goods and public goods. However, granting unequal voting power based on wealth is both politically and legally infeasible, even if it were desirable [4]. As an alternative, all voters can instead be allotted a certain number of voting credits to purchase votes. Upon reaching voting age, a citizen would be allotted a stock of voting credits and would subsequently receive an additional allotment at set intervals of time [4]. Each election, the voter can "purchase" v votes for a given referendum (or candidate) at a cost of v^2 voting credits. Assuming a voter, i, has a quasi-linear utility for retaining voting credits for future referenda, then the voter's expected payoff is:

$$\text{expected payoff} = pv_i u_i - v_i^2 \ , \tag{1}$$

where p is the probability of a single vote flipping the outcome (known as the marginal pivotality), v_i is the number of votes purchased, and u_i is the voter's utility for the referendum (negative if in opposition). Differentiating with respect to votes shows that the voter maximizes payoff by buying votes proportionally to her utility. Voting rationally under QV thus requires voters to correctly determine their utility for a given outcome and to estimate the marginal pivotality of a vote.

Under ideal conditions, QV has been shown to optimize welfare. If all voters estimate the same marginal pivotality and vote rationally, then the sum of votes, $\sum v_i$, will have the same sign as the sum of utilities, $\sum u_i$, and the mechanism will maximize aggregate welfare (assuming all voters also correctly determined their utility). The question then becomes: how do we model the process of voters estimating marginal pivotality? One formulation of QV adopted the price-taking assumption [5, 6] that all voters simply agree on p and proved that QV is then the unique optimal vote buying mechanism [2]. The price-taking assumption is equivalent to assuming that all voters use exactly the same (or equivalent) processes for estimating marginal pivotality, regardless of what the process was. Although this is clearly unrealistic, it serves as a useful bound of optimality. Later work assumed that all voters know the distribution from which the population's utilities are drawn and use this information to estimate marginal pivotality. Under this particular shared process for determining marginal pivotality, although QV is not perfect, it is still superior to 1p1v in both infinite and finite populations in terms of welfare loss—the probability that a mechanism does not maximize welfare [7, 8]. These are valuable results, but it is difficult to extend the analytic methods used to obtain them for modeling more realistic voter knowledge and behaviors. Additionally, no results have been published for models of modified QV (mQV) in which voters have a fixed number of voting credits, despite this being the most likely scenario for real-world implementations of QV [9].

We aim to investigate the performance of QV with more realistic voting behavior and variations of the QV mechanism. These can be difficult to model with the analytical tools used in previous QV research. Therefore, we employ agent-based methods for the modeling. Previous research on QV assumed highly idealized processes for voters determining marginal pivotality, and no models have been published for analyzing mQV. So, we investigate the following research questions: (1) How does variation in perceived marginal pivotality (PMP) affect the efficiency of QV in terms of welfare loss (WL)—the fraction of the time that the result of voting does not maximize welfare and (2) what are the patterns of WL under mQV in which people vote on multiple referenda at once with a fixed number of voting credits?

2 QV with Variance in Percieved Marginal Pivotality (PMP)

2.1 Agent-Based Model of QV

The model reported in this section builds on previous work investigating the efficiency of unmodified QV [7, 8]. The critical assumption of QV compared to mQV is that voters are assumed to be able to purchase as many votes as they need to fully express the intensity of their preferences. At the same time, voters are assumed to have a quasi-linear value for retaining voting credits for future elections.

The purpose of this model is to investigate the performance of QV when voters do not all estimate the same marginal pivotality. Voters are the only agents in the agent-based model (ABM), and they are assigned utilities for the referendum under consideration. In one condition, utilities are normally distributed around zero (normal distribution condition). Different standard deviations for the normal distribution condition were tested and did not affect results. In the second condition (Prop 8 condition), utilities are calibrated following previous work [7, 10] to model a "tyranny of the majority" situation based California's Proposition 8.

In a 2008 election in California, a referendum was held on Proposition 8, which proposed to ban same sex marriage. The resulting votes were 52% in favor to 48% opposed. Since the California election used the 1p1v system, the proposition passed. However, there is reason to doubt that this decision maximized the welfare of the voting population. Since lesbian, gay, bisexual or transgender (LGBT) people would be greatly affected by the decision, it is likely that they cared about the outcome significantly more than other voters. Thus, our intuitions might be that the outcome of the vote failed to maximize welfare. In the calibration to represent this scenario, the population is divided into four groups: Approximately 3.3% of the voting population is LGBT and single, and approximately 0.7% are members of LGBT couple

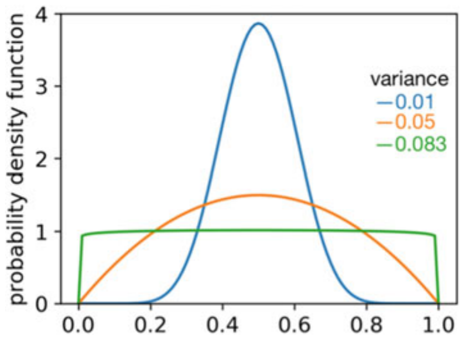

Fig. 1 The beta distribution with three different variances. At low variance, it is approximately normal and at variance 0.083 it is approximately a uniform distribution on the interval 0 to 1

households. These two groups are assumed to have average utilities equal to $20,000 and $100,000 respectively in opposition to Proposition 8 (see [7] for justifications of these numbers). Proposition 8 passed 52–48%, implying that an additional 44% of population beyond the LGBT community opposed it. This 44% and the remaining 52% who supported Proposition 8 are assumed to both have utilities of $5000 on average, either in favor or in opposition to the proposition. More specifically, non-LGBT voters draw a utility from a uniform distribution between 0 and $10,000, either in favor or against the proposition. LGBT singles draw a utility from a uniform distribution between $5,000 and $35,000, and members of LGBT couples draw a utility from a uniform distribution between $20,000 and $180,000. This results in a mean utility of $960 against Proposition 8, while the median voter is in favor.

Each voter has a perceived marginal pivotality (PMP) drawn from a beta distribution bounded 0 to 1 (this value can be divided by a large number, e.g. 1000, to obtain more realistic values of p, but this is optional because it only scales vote numbers but not their proportions). We vary the variance of the beta distribution from 0 (all voters have the same PMP) to 0.083, at which point a beta distribution is essentially a uniform distribution from 0 to 1, as shown in Fig. 1. This distribution does not have deep theoretical or empirical justification. We use it because its lower bound is zero (marginal pivotality cannot be negative, as it is a probability) and it is approximately a normal distribution at low variances. Our model is agnostic as to how voters determine PMP, instead investigating the impact of variance in PMP on WL, regardless of how this variance came about. Voters cast votes rationally to maximize their utility by maximizing Eq. 1 (with p replaced by p_i, the voter's PMP). Note that we assume voters perfectly estimate their own utility for the referendum.

We compare QV with rational voters to both 1p1v and to a QV scenario in which voters vote in the right direction according to their own utilities but with a random magnitude between 0 and 3. We call this QV-RDRM for "right direction random magnitude." QV-RDRM models a scenario in which voters can correctly determine whether or not a referendum will benefit them but have no ability to determine by how much. This is unlikely in practice, especially because a study of real people using QV to express opinions on a questionnaire found that votes were often approximately normally distributed [11]. Due to its implausibility, QV-RDRM serves as a useful

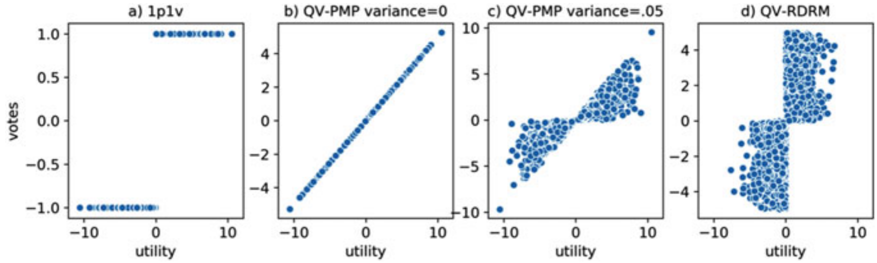

Fig. 2 Votes cast versus utility for (from left to right) 1p1v, QV with 0 variance in PMP, QV with variance in PMP equal to 0.05, and QV-RDRM

lower bound on possible outcomes from QV. The choice of a random magnitude between 0 and 3 is arbitrary. Different bounds for the random magnitude will change the results quantitatively, but not qualitatively.

We also vary population size from ten to 10,000. Each parameter combination was run 1000 times using models written in NetLogo [12].[1] Welfare loss (WL) is calculated as the fraction of runs with negative aggregate payoff, P_a, which is the sum of each voter's utility multiplied by the sign of the vote outcome: $P_a = \sum_i(\text{sign}(\sum_j v_j)u_i)$.

To aid in interpreting the results of the next section, Fig. 2 shows voting behavior versus utility for a population of agents. Figure 2a shows voting behavior under 1p1v: those with positive utility vote +1 and those with negative utility -1 regardless of the magnitude of the utility. Figure 2b shows votes versus utility for QV when all voters have the exact same PMP; votes are linearly proportional to utility, ensuring that welfare is optimized. Figure 2c shows votes versus utility for QV when there is variance in voters' PMP. The best fit line is still a linear relationship, but there is considerable variance in votes cast for a given utility, especially for high-magnitude utilities. Figure 2d shows votes versus utility for QV-RDRM, in which all voters vote in the right direction given their utilities, but with a random magnitude.

2.2 Results

2.2.1 Normal Distribution Condition

In the normal distribution condition, WL does not depend on the population size for 1p1v, QV, or QV-RDRM. Figure 3 plots the WL of QV versus variance of PMP for this condition along with the WL of 1p1v and QV-RDRM for reference. Unsurprisingly, WL of QV increases with increased variance in PMP, but even when PMP is a flat distribution between 0 and 1 (variance of 0.083), WL is significantly lower than 1p1v.

[1] Source code for the models can be found at https://github.com/jzkelter/qv.

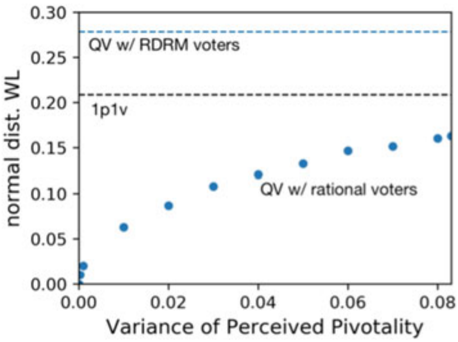

Fig. 3 Welfare loss versus variance in perceived marginal pivotality (PMP) when utilities are drawn from a normal distribution. Increased variation in PMP decreases the performance of QV, but it still outperforms 1p1v

QV-RDRM performs worse than 1p1v, because it is essentially 1p1v with random noise that sometimes flips the outcome from what it would have been under 1p1v. If 1p1v had a WL of 0.5, then the noise would not matter on average, but 1p1v only has a WL of approximately 0.2 in the normal distribution condition. So, the random noise in QV-RDRM worsens WL about 80% of the time and improves it only about 20% of the time. The exact amount that QV-RDRM underperforms 1p1v will depend on the distribution that random votes are drawn from, but the logic holds that QV-RDRM will, at best, be equivalent to 1p1v in the normal distribution condition.

2.2.2 Prop 8 Condition

In the Prop 8 calibration, WL depends on both the number of voters and the variance in PMP. Figure 4 shows WL of 1p1v, QV-RDRM and QV versus number of voters. As expected, WL of 1p1v in this condition approaches one as the number of voters increases, because with a large number of voters, the median and mean utilities of the population match those of the sampling distribution with a very high probability, which have opposite signs for the Prop 8 calibration. With smaller numbers of voters, there is a significant chance that the sample drawn from the Prop 8 distribution will have a median and mean with the same sign, in which case, 1p1v maximizes utility. QV-RDRM performs similarly to 1p1v. Since QV-RDRM is essentially 1p1v with noise, as long as 1p1v has WL less than 0.5, QV-RDRM performs slightly worse than 1p1v. When the WL of 1p1v surpasses 0.5 at around 100 voters, then QV-RDRM performs slightly better than 1p1v.

In contrast, WL of QV (with rational voters) decreases with population size, reaching essentially zero with 1000 voters regardless of the variance in PMP. All WL for QV in this model is due to voters under/overestimating PMP. Figure 4 also makes clear that even the worst performance of QV with rational voters is better than the best of 1p1v under these conditions.

These results confirm that QV is robust to variation of PMP (e.g., from imperfect information) under the assumptions that voters are still rational utility maximizing agents. This model still assumes that PMP does not depend on utility of the agents, for

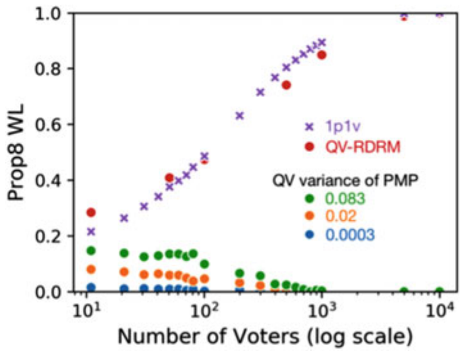

Fig. 4 Welfare loss (WL) versus number of voters using the Prop 8 calibration. WL of QV approaches 0 regardless of the variance in perceived marginal pivotality while WL of 1p1v approaches 1

example through heterogeneous social factors across preference groups. In essence, this model assumes all voters are homogeneous except for differences in utility. Even the variation in PMP in this model is only heterogeneous in the sense that agents have different values assigned randomly, but all agents still have PMP drawn from the same distribution.

3 Modified QV with Multiple Simultaneous Issues

3.1 Agent-Based Model of mQV

There are multiple potential versions of modified QV (mQV) [4]. In the version modeled here, each voter in the ABM is given a fixed number of voting credits to allocate among multiple referenda on a single ballot. This contrasts with other mQV proposals in which voters are given a fixed number of voting credits at set intervals (e.g., every 10 years) to use across elections as they see fit. We chose to model a version of mQV with a fixed number of voting credits for a single election for two reasons. First, this is the easiest to implement in the real world since it doesn't require a system to keep track of voters' voting credits over time, and, as it turns out, it is the version that has been tested "in the wild" to elicit preferences [11]. Second, it simplifies the model because voters don't have to make predictions about how much they will care about referenda on future ballots. Instead, voters only have to compare referenda on the current ballot.

The fact that voting credits cannot be saved for future elections breaks one of the assumptions of the original QV formulation: that voters have a quasi-linear value for retaining voting credits for future elections [2]. However, when voting on any given referendum, voters still want to retain voice credits for the remaining referenda on the ballot. More importantly, the fixed number of voting credits limits the ability of voters with very strong preferences to fully express them. It has already been

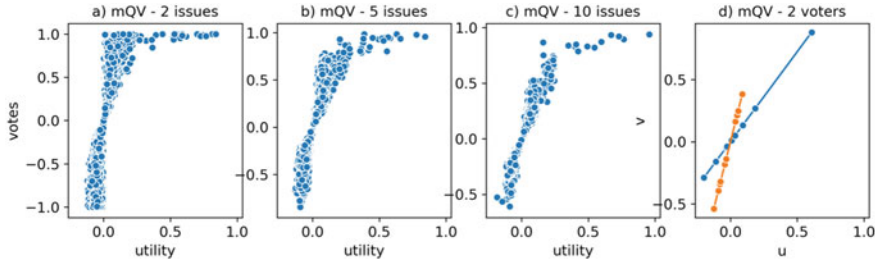

Fig. 5 Votes versus utility for mQV with **a** two, **b** five and **c** ten referenda on the ballot. **d** Shows Votes versus utility for two specific voters with ten referenda on the ballot, one of whom has a very high utility for one of the referenda

recognized that this will have some impact on the efficiency of mQV [4], but no quantitative models of the impact have been published.

Once again, voters are the only agents in the model. They are assigned a utility for each of the referenda under consideration. In the normal distribution condition, all utilities are drawn from a normal distribution centered at zero. In the Prop 8 condition, utilities for one of the referenda are drawn from the Prop 8 calibration described in Sect. 2.1 while utilities for all other referenda are drawn from normal distributions.

The voters split their votes across referenda proportionally to their utilities without estimating marginal pivotality. This is equivalent to voters assuming they have the same marginal pivotality for each issue. We investigate the effect of the number of voters and the number of referenda on WL of mQV. WL is calculated as the fraction of runs with negative aggregate payoff for each issue. Each parameter combination was run 1000 times.

To aid in understanding the results, Fig. 5 shows plots of votes versus utilities for mQV with various numbers of referenda using the Prop 8 calibration. If only one referendum is on the ballot, mQV is equivalent to 1p1v, because voters will use all of their voting credits for that single referendum. With only two referenda (shown in Fig. 5a), mQV looks like QV with variation in PMP for small utilities, but like 1p1v for large utilities. As the number of referenda increases (Fig. 5b, c), the behavior of mQV approaches that of normal QV. However, even with 10 issues, the voters with the highest utilities are not able to vote proportionally to their utilities because of the limited number of voting credits as seen in Fig. 5c. This is only compared to voters with lower utilities on all issues. High-utility voters still vote linearly with their utilities when looking at only their own votes, but since they have at least one high utility issue, the slope of their votes versus utility graph is smaller than that of a voter without any high utilities as shown in Fig. 5d (the blue line represents the high utility voter). When looking at the entire population in aggregate, it appears that high-utility voters are unable to vote linearly with utilities.

3.2 *Results*

3.2.1 Normal Distribution Condition

When utilities for all issues are from the normal distribution condition, the WL once again does not depend on the population size for either 1p1v or mQV. Figure 6 shows the average WL of all referenda as the number of referenda varies for mQV, along with the WL of 1p1v for reference (which does not depend on the number of referenda). The WL of mQV decreases with the number of referenda, because as voters spread their votes among more referenda, their vote buying behavior on any single referendum approaches that of single-issue QV with retainable voting credits. Even with just two referenda, the WL of mQV is significantly lower than that of 1p1v. Of course, with just one referendum, mQV is equivalent to 1p1v, because each voter will just put all of their credits towards their preferred outcome on the single issue.

3.2.2 Prop 8 Distribution Condition

When utilities for one of the referenda are drawn from the Prop 8 calibration, WL depends on both the number of voters and the number of referenda. Figure 7 shows WL for the Prop 8 calibrated referendum when utilities for all other referenda are drawn from the normal distribution condition. Having only one referendum is equivalent to 1p1v since voting credits are limited to this election. As expected under 1p1v, WL for the Prop 8 calibrated referendum approaches one with increasing population, because, with very high probability, the mean and median of the population will match those of the sampling distribution which have opposite signs. Since only 4% of the sampling distribution has strong utilities in the Prop 8 calibration, the smaller the population is, the more likely it is that there will not be many people with strong utilities, resulting in the mean and median utilities being of the same sign. For this reason, smaller populations lead to lower WL of 1p1v under the Prop 8 calibration.

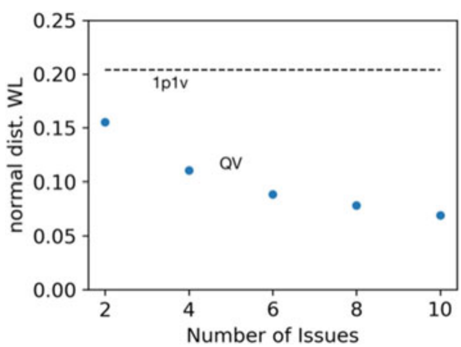

Fig. 6 Welfare loss (WL) of mQV versus number of referenda with all utilities drawn from zero-centered normal distributions

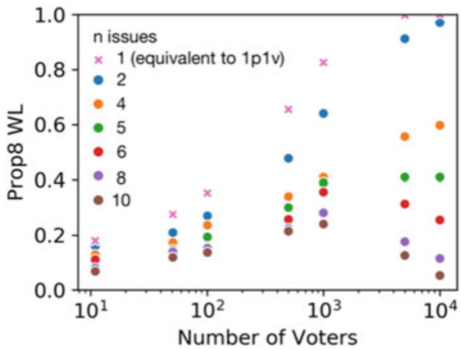

Fig. 7 Welfare loss on the Prop 8 referendum versus number of voters using multi-issue QV

The performance of mQV improves as the number of referenda increases, because voting behavior of the population approaches the ideal of being linear with utility, as seen earlier in Fig. 5. With ten referenda and 10,000 voters, WL is only about 5%. Surprisingly, even with six or more referenda, WL increases up until around 1000 voters before starting to decrease with the number of voters. This is due to the probability of there being enough high-utility voters to flip the sign of the mean utility to be opposite that of the median without there being enough to flip the sign of the total votes. If the number of voting credits were not restricted, then a voter who flipped the mean utility would be able to cast enough votes to flip the outcome as well, but with restricted voting credits this is not always the case. As the number of voters increases past 1,000, the probability that there will be enough high utility voters to flip the vote outcome increases and WL decreases. These results are of course specific to the Prop 8 calibration used, but they should be qualitatively similar to any situation in which the mean and median utilities have opposite signs.

4 Discussion and Future Work

Our results confirm the superiority of single-issue QV over 1p1v for maximizing aggregate welfare even with large variations in perceived marginal pivotality. Additionally, we offer the first (as far as we know) formal model of mQV and confirm its superiority over 1p1v on all multi-issue referenda under both normal distribution and Prop 8 conditions.

We also revealed a potentially surprising relationship between WL and number of voters in Prop 8 calibrations of mQV, in which WL increases with the number of voters up until about 1,000 before decreasing with the number of voters. In some ways this can be considered an artifact of the model construction. With small populations, it is likely that not enough high utility voters are drawn from the utility distribution to make the mean and median utilities different signs, in which case both mQV and 1p1v maximize utility.

These results further strengthen the case for some form of QV over 1p1v for collective binary decision making. Even with large amounts of noise in perceived marginal pivotality, QV produced lower welfare loss than 1p1v in our models, regardless of the number of voters and the utility distribution used. The only case in which QV performed worse than 1p1v was under the assumption that voters cast votes in the right direction according to their utilities, but with completely random magnitudes. Even with this unrealistic assumption, QV only slightly underperforms 1p1v. Barring other factors, this suggests QV will perform similarly to 1p1v in the worst case and significantly better than 1p1v in the best case.

However, there are other factors to consider, which we plan to investigate in future work. First, our mQV results are based on the assumption that voters do not estimate marginal pivotality for each referendum. A psychologically and socially plausible mechanism for voters to estimate their marginal pivotality, for example from information about peers and polls, could lead to significantly different outcomes. Second, we assumed that voters have perfect information about their own utilities (except in QV-RDRM), but it is perfectly possible for voters to think a policy will benefit them when it will not. A model in which voters have perceived utility for a referendum which can differ from the actual value they will experience if the referendum is implemented is another important future research direction. Third, we assumed in both of our models that utility and PMP are uncorrelated. In the real world it is likely that both would be influenced by social processes of peer influence and signaling. In essence, the models presented here (and all previous models of QV) are models of socially isolated voters, uninfluenced by any social milieu. Social processes leading to even a small correlation between utility and PMP could drastically change the outcomes of QV. Finally, it has been suggested that real world voting behavior under QV would be mostly determined by social and psychological factors rather than pure instrumental rationality [13]. This warrants an investigation of the results for various heuristic voting rules beyond the unrealistic "right direction random magnitude" heuristic analyzed here.

In future work, we intend to address the above issues by modelling the effects on QV of various psychological and social factors such as (a) having voters calculate PMP from peers and polling information for mQV elections, (b) including a model of voter turnout which can depend on both individual utility and social influence, (c) giving voters perceived utilities that are affected by social factors such as advertising and peer influence and which may not align with their true utilities, and (d) giving voters realistic heuristic voting rules that are not based on strict utility-maximizing calculations.

References

1. Adams, J. *A Defence of the Constitutions of Government of the United States of America*, Google-Books-ID: 0JuNQoYVKMC (C. Dilly, 1788)
2. Lalley, S. P. & Weyl, E. G. Quadratic Voting: How Mechanism Design Can Radicalize Democracy. *AEA Papers and Proceedings* **108**, 33–37. ISSN: 2574-0768. https://www.aeaweb.org/doi/10.1257/pandp.20181002 (2018)
3. Weyl, E. G. Quadratic Vote Buying. *SSRN Electronic Journal* (2013)
4. Posner, E. A. & Stephanopoulos, N. O. Quadratic election law. *Public Choice* **172**, 265–282. ISSN: 0048-5829, 1573-7101. http://link.springer.com/10.1007/s11127-017-0415-2 (2017)
5. Laine, C. R. Strategy in Point Voting: A Note*. The Quarterly Journal of Economics 91, 505–507. ISSN: 0033-5533. https://www.jstor.org/stable/1885981 (1977)
6. Mueller, D. C. Constitutional Democracy and Social Welfare. *Q J Econ* **87**, 60–80. ISSN: 0033-5533. https://academic.oup.com/qje/article/87/1/60/1907208 (1973)
7. Chandar, B. & Weyl, E. G. Quadratic Voting in Finite Populations. https://papers.ssrn.com/abstract=2571026 (2019)
8. Lalley, S. & Weyl, E. G. *Nash Equilibria for Quadratic Voting* SSRN Scholarly Paper ID 2488763 (Social Science Research Network, Rochester, NY, 2019). https://papers.ssrn.com/abstract=2488763
9. Posner, E. A. & Weyl, E. G. Quadratic voting and the public good: introduction. *Public Choice* **172**, 1–22. ISSN: 1573-7101. https://doi.org/10.1007/s11127-017-0404-5 (2017)
10. Weyl, E. G. The robustness of quadratic voting. *Public Choice* **172**, 75–107. ISSN: 1573-7101. https://doi.org/10.1007/s11127-017-0405-4 (2017)
11. Quarfoot, D. et al. Quadratic voting in the wild: real people, real votes. *Public Choice* **172**, 283–303. ISSN: 1573-7101. https://doi.org/10.1007/s11127-017-0416-1 (2017)
12. Wilensky, U. NetLogo 1999. http://ccl.northwestern.edu/netlogo/.
13. Kaplow, L. & Kominers, S. D. Who will vote quadratically? Voter turnout and votes cast under quadratic voting. *Public Choice* **172**, 125–149. ISSN: 0048-5829, 1573-7101. http://link.springer.com/10.1007/s11127-017-0412-5 (2017)