Anton Arbman Hansing*

Quadratic Voting and Heterogeneous Beliefs

Bachelor's essay in Economics

Supervisor: Tommy Andersson
Department of Economics
Lund University
August 19, 2020

*anton.arbman.hansing@hhs.se

Abstract

Quadratic voting has been proposed as a voting procedure that aggregates preferences in a near-utilitarian way. A voter's optimal strategy is, however, not solely a function of her preferences, but also of her beliefs about how others will vote. This thesis presents a formal model of quadratic voting that, unlike previous models, allows for heterogeneity in the beliefs of voters. The specification is then used to answer two questions about the relationship between beliefs and voting outcomes. Firstly, how does heterogeneity in beliefs impact election outcomes? Secondly, what incentives do individual voters have to acquire information about how others will vote? Results suggest that heterogeneity in beliefs leads election outcomes to represent preferences only noisily. As long as preferences and beliefs are independently distributed, the detrimental effects on efficiency decline as the number of voters grows large. If beliefs and preferences are correlated, however, heterogeneous beliefs may lead to outcomes that are systematically inefficient.

Contents

1 Introduction

In firms, organizations, democratic states, and everyday life, collective decisions are often made through voting procedures. These procedures allow each eligible member of the group to express her preference by placing a vote, and determines the collective choice of the group as a function of the votes cast. The structure of a vote, as well as the decision rule determining the relationship between votes and the collective choice, can be designed in numerous ways. Despite extensive study, political scientists and economists differ widely in their opinions about the relative merits of different such systems.

One recurring criticism of common voting systems, such as majority voting, is that they fail to capture preference intensity. In an election between two alternatives, the vote of a nearly-indifferent voter is typically weighted equally to that of a voter with an intense preference for a particular alternative. This feature can lead to a *tyranny of the majority*, where the weak preference of a majority trumps the strong preference of a (potentially almost as numerous) minority. In a referendum on the legality of same-sex marriage, for example, a nearly-indifferent opposing majority might win even if a large minority in favor cares deeply about the issue.

Quadratic Voting (QV) has been proposed as a voting system that aims to achieve (approximately) utilitarian-efficient outcomes in the presence of heterogeneous preference intensities. Under QV, voters choose freely how many votes to cast in each election, paying a monetary price proportional to the square of the number of votes cast. The alternative for which the highest number of votes is purchased is implemented. Since the marginal cost of each incremental vote is increasing, voters with stronger preferences are likely to purchase higher numbers of votes, yielding them a higher influence on the outcome of the election. Several formal models have been put forward to illustrate the benefits of such a system (Lalley & Weyl, 2018, 2019; Laurence & Sher, 2017). QV has also been used with tentatively positive results in both experimental settings (Casella & Sanchez, 2019) and for real-world purposes, most notably in the Democratic Party caucus of the Colorado House of Representatives (Eason, 2019).

In the existing, theoretical literature on quadratic voting, participants decide upon a number of votes to cast by solving an expected utility maximization problem. Placing a higher number of votes increases the probability that the election outcome is favorable, but incurs a private cost that increases on the margin. The two determinants of a voter's behavior are hence (i) her preference, including its intensity, and (ii) her perceived probability of changing the outcome of the election, conditional on placing a certain number of votes. While preferences have been assumed to vary exogenously, all voters have been assumed identical with respect to their beliefs. That is, all voters share the same estimate of how likely they are to swing the election when placing a certain number of votes. As a consequence, variability in election outcomes is in these models fully attributable to variability in preferences, effectively masking the role of strategy-relevant beliefs.

Several contributions have offered critiques of these formal models. Laurence and Sher (2017) explore the consequences of wealth inequality and differences in the marginal utility of money, and find that QV may perform poorly for questions on which opinions are polarized by wealth. They also argue that QV fails an important democratic legitimacy requirement: that citizens should have equal opportunity to affect election outcomes. Weyl (2017) models, among other problems, the risk that voters may collude to buy more votes at lower prices. To my knowledge, however, the possibility of voters differing in beliefs about the closeness of the election (and hence, the value of an additional vote) remains unexplored in the existing literature.

The purpose of this thesis is to study how differences in beliefs – in particular, beliefs about the voting behaviors of other voters – affect the outcome of QV elections. The rest of the thesis is structured as follows. Section 2 provides a formal introduction to the study of voting systems and discusses the strategic role of voter beliefs. In Section 3, quadratic voting is introduced through two previous specifications. I then present my own specification, incorporating the possibility of heterogeneous beliefs, and demonstrate the conditions under which it attains efficient outcomes. In Section 4, I explore two applications of this novel specification. First, I analyze how differences in beliefs may lead to inefficient outcomes. Second, I model how an individual voter will value information about the voting behavior of others. Section 5 concludes the thesis.

2 Voting, Beliefs and Strategy

This section provides an introduction to the formal study of voting systems. A set of individuals $\mathcal{N}$ are set to collectively decide on an outcome from a set of possible alternatives $\mathcal{X}$. Each individual i chooses a vote v_i from some set of possible votes $\mathcal{V}$. The structure of a vote depends on the voting systems. In binary elections, for example, it may be the case that $v_i \in \{\text{Yes, No}\}$, while in elections featuring more options, each vote may constitute a complete ordering of all alternatives, e.g. $v_i = (1 : \text{Red Party}, 2 : \text{Blue Party}, 3 : \text{Green Party})$.

The outcome of the election is determined by a *decision rule*, i.e. a function $\Phi : \mathcal{V}^{\mathcal{N}} \to \mathcal{X}$ that, for each possible set of votes, outputs a uniquely identified collective choice. Voters will be assumed to choose their voting behavior in order to maximize their expected utility. A voter i whose preferences are described by a utility function U_i will hence choose to vote as

$$v_i^* = \underset{v_i \in \mathcal{V}}{\arg\max}\ \mathbb{E}\big[U_i(v_i)\big] \tag{1}$$

where $\mathbb{E}[U_i(v_i)]$ denotes the utility that she expects to attain when placing a vote v_i, subject to her relevant probabilistic beliefs about the world. For our purposes, we can define a *voting system* as a tuple $(\mathcal{V}, \Phi)$ describing (i) the structure of admissible votes and (ii) the decision rule determining the outcome for every possible vector $\mathbf{v} = [v_i : i \in \mathcal{N}]$. A common voting system is *plurality rule*.

Example 1 (Plurality rule)**.** Under plurality rule, each voter places a vote for exactly one of the available alternatives, and the outcome of the election is determined as the alternative that receives the highest total number of votes. Formally,

- $\mathcal{V} = \mathcal{X}$; the set of possible votes is the set of alternatives.
- $\Phi(\mathbf{v}) = \arg\max_{x \in \mathcal{X}} |\{v_i = x; i \in \mathcal{N}\}|$; the chosen outcome is whichever alternative receives the highest number of votes.[1]

ꝏ

Under plurality rule between two competing alternatives, the optimal voting behavior of a voter is simple: whatever her beliefs about how others will vote, it is optimal to vote for her preferred alternative. In the terminology of game theory, voting for one's preferred option is a *dominant strategy*. In these situations, May's theorem (May, 1952) shows that majority rule is the single voting system that aggregates ordinal preferences to unique collective choices while (i) treating all voters, and all outcomes, equally, and (ii) responding positively to preferences.[2]

If, however, the election is between three or more alternatives, it is no longer necessarily optimal to vote for one's preferred alternative. As an example, consider an election between three alternatives A, B, and C. A voter preferring alternative A to B, and B to C, may nevertheless choose to vote for B if she does not believe that A will enjoy sufficient support from other voters to have a chance at winning. Her rational strategy is then determined partly by her beliefs of how others will vote, as opposed to being uniquely defined by her preferences. We will refer to this behavior of belief-informed voting as *strategic voting.* [3]

[1] Here, $|\cdot|$ denotes the cardinality of a set. For the sake of conciseness, we disregard the possibility of ties.

[2] *Positive responsiveness* means, slightly simplified, that if the outcome is a draw before an additional individual is added to the group of voters, then the outcome must change to whichever alternative is preferred by that individual.

[3] An alternative definition, used in the context of ordinal voting systems, is that a voter votes *strategically* if and only if her vote misrepresents her true preference ordering (Kawai and Watanabe, 2013, Myatt, 2007, Gibbard, 1973). This definition is less useful in the context of cardinal voting systems such as QV, since these rarely provide an objective way to derive a "sincere" vote from a set of (even cardinal) preferences.

2.1 Strategic Voting

In the example of plurality rule, we saw that the optimal voting strategy may depend on a voter's beliefs if the election is between more than two possible outcomes. This problem is not restricted to plurality rule. Instead, Gibbard (1973) and Satterthwaite (1975) famously showed that no deterministic, non-dictatorial[4] voting system allows for dominant strategies in the presence of more than two alternatives.[5] Strategic voting is hence an inescapable component of multi-option elections.

Voter beliefs have also been shown to empirically affect voting behaviors. Forsythe et al. (1993) demonstrate in a laboratory setting that subjects vote differently in plurality rule situations if shown polls of their peer's preferences in advance. Cain (1978) studies data from the 1970 British general election, and finds that third-party voting (i.e., the proportion of votes for other parties than the Labour and Conservative parties) was lower in constituencies where the margin between the top parties was small. This pattern is interpreted as evidence that voters who prefer alternative parties nonetheless have a tendency to vote for either the Labour party or the Conservative party, if the election is believed to be tightly contested.

The fact that voting behaviors reflect not only preferences, but also strategic beliefs, has notable detrimental effects on the outcomes and interpretability of elections. One consequence is that voters that are well-informed about how others will vote may systematically attain more preferable election outcomes than voters without such knowledge. This is not an intended feature of voting systems, that are typically motivated as means of aggregating preferences. While there are domains, such as financial markets, that are designed to grant larger influence to better-informed participants, voting systems are not such a domain – in particular when the information in question is not about the outcomes of policies, but about the voting behaviors of others.

Furthermore, when voters' actions are affected by their beliefs, preferences can no longer be inferred from election results. In Swedish parliamentary elections, for example, seats are distributed proportionally between all parties receiving at least 4% of the national vote. Voters whose preferred party is expected to receive significantly less support than this threshold will likely hesitate to vote for their first choice. The public support of the ideas put forward by this party will hence likely be underestimated by their election result. In the US, self-identified "independent" voters have outnumbered both Democrats and Republicans throughout the last decade (Gallup Inc., 2020). If self-identifying as independent is interpreted as a preference for third parties, then the fact that almost all votes in presidential elections are for

[4] A voting system is *dictatorial* if its outcome depends only on the vote of a single voter, i.e. if $\Phi(\mathbf{v}) = \Phi(v_i)$ for a single $i \in \mathcal{N}$.

[5] This result was further generalized to every stochastic voting system except the random dictatorship by Hylland (1980).

Democrat or Republican candidates suggests that strategic voting is ubiquitous. This, in turn, makes it challenging to determine the nature of independent voters' disagreement with the major parties.[6]

We have now seen that strategic voting has an inescapable role in elections between three or more alternatives, to potentially detrimental effects on both the efficiency of election outcomes and the interpretability of voting results. As illustrated in Example 1, common voting systems avoid these adverse consequences in binary elections, since the same behavior is then optimal regardless of one's beliefs. As we will see, however, QV may suffer from strategic voting even when only two options are under consideration. The remainder of this thesis will be aimed at understanding the relationships between strategic beliefs, voting decisions, and election outcomes in quadratic voting.

3 Quadratic Voting

This section introduces the concept of quadratic voting (QV). The first part sets up the mechanics of the voting system, specifies the assumed preferences of voters, and defines the primary efficiency criterion used to evaluate the outcomes of the election. The second part briefly discusses two solution concepts previously proposed by Lalley and Weyl (2018, 2019). The third part introduces the belief-centered model put forward in this thesis, and specifies assumptions under which this model coincides with the previously proposed models.

3.1 General Setup

Consider a set of individuals $\mathcal{N}$ tasked to make a collective decision between two alternatives: maintaining the current state of affairs, and implementing some proposed change. The outcome of the decision is described by the binary variable X:

$$X \in \{0 : \text{Maintain status quo}, 1 : \text{Implement some change}\}. \tag{2}$$

Each voter i chooses a number of votes to cast $v_i \in \mathbb{R}$, where $v_i > 0$ represents $|v_i|$ votes in favor of the proposed change, and $v_i < 0$ represents $|v_i|$ votes in favor of the status quo. Each voter pays a monetary cost $c(v_i) = v_i^2$ to cast the votes.

The utility of each voter is quasilinear in wealth. In addition, let $u_i \in \mathbb{R}$ denote voter i's willingness to pay to implement the proposed change with certainty, instead of maintaining the status quo with certainty. Hence, $u_i > 0$ if and only if i strictly prefers the outcome $X = 1$

[6] Of course, an alternative interpretation is that "independent" voters are simply open to vote for either major party, in which case the informational problem of strategic voting is likely smaller.

over $X = 0$. Since utility is assumed to be quasilinear in wealth, the willingness to pay u_i can equivalently be described as i's preference over election outcomes. Furthermore, we will sometimes refer to $|u_i|$ as the *preference intensity* of i. This can be thought to capture the extent to which the voter cares about the particular issue at hand. In sum, the preferences of each individual i are fully described by the utility function

$$U_i(X, v_i) = X \cdot u_i - v_i^2. \tag{3}$$

The outcome of the election is determined as the alternative which receives the highest number of votes: $X = 1$ if and only if $\sum_{i \in \mathcal{N}} v_i \geq 0$. X can therefore be thought of as a function partly determined by v_i. A utility-maximizing voter will hence take two aspects into consideration when determining how to vote. Firstly, placing a larger number of votes will increase the probability of a favorable election outcome. Secondly, however, it will also incur a higher private cost. Furthermore, due to the quadratic cost of votes, the marginal cost of placing an additional vote in increasing in v_i.

Let v_i^* denote the number of votes purchased by i in some (later defined) equilibrium. The motivation of quadratic voting is that it claims to achieve a particular type of efficiency, defined in Lalley and Weyl (2018).

Definition 1 (adapted from Lalley and Weyl (2018))**.** A voting system is *robustly efficient* if, for any $\mathbf{u}$, the equilibrium votes $\mathbf{v}^*$ satisfy:

$$\text{sign}\left(\sum_{i \in \mathcal{N}} v_i^*\right) = \text{sign}\left(\sum_{i \in \mathcal{N}} u_i\right) \tag{4}$$

or, equivalently, if the alternative for which the aggregate preference is the highest, is always implemented in equilibrium. ꝏ

We may notice that efficiency, in this sense, is a utilitarian criterion; a voting system is efficient if and only if it consistently implements the alternative for which the aggregate utility is maximized. Plurality rule fails this criterion, since it may well be the case that a majority prefers one alternative, while the aggregate preference is higher for the other option. For example, it may be the case the $u_i < 0$ for a majority of a voters i, but that the sum $\sum_i u_i$ is nonetheless greater than 0. In this case, plurality voting implements the choice $X = 0$, while the efficient choice is $X = 1$.

This efficiency criterion can be argued for on two distinct grounds. Firstly, it can be interpreted as a utilitarian criterion – it requires the implementation of whatever alternative maximizes aggregate utility. This may seem like a desirable characteristic for those who view voting as a means to implement outcomes that are collectively optimal in a consequentialist sense. It will, however, likely fail to impress those that view voting from rights-based[7] or contractualist[8] perspectives.

The criterion from Definition 1 can, however, also be justified as a way of guaranteeing Kaldor-Hicks efficiency. In the case of voting outcomes that are inefficient in the above sense, a Pareto improvement can always be achieved by changing the election outcome and letting the winners pay compensation to the losers. To see this, let $\mathcal{N}_-$ and $\mathcal{N}_+$ denote the set of voters for which $u_i < 0$ and $u_i > 0$, respectively. Similarly, let $U_- = \sum_{i\in\mathcal{N}_-} |u_i|$ and $U_+ = \sum_{i\in\mathcal{N}_+} |u_i|$ denote the aggregate preference of each group. Then, a voting system is efficient if any only if $X = 1$ exactly when $U_+ > U_-$. If an inefficient system proposes $X = 0$ in this case, then there exists a wealth transfer of T from $\mathcal{N}_+$ to $\mathcal{N}_-$ (where $U_- < T < U_+$) that leads to a Pareto improvement. This transfer would increase the aggregate utility attained by $\mathcal{N}_+$ by $(U_+ - T) > 0$ and that of $\mathcal{N}_-$ by $(T - U_-) > 0$. In an efficient voting system, transfers of these kinds can never constitute such improvements.

Unlike major voting systems, QV involves monetary transactions. It is, however, not intended to be source of revenue for the organizer of the election. The amount paid for votes is typically assumed to be repaid in full to the voters, such that for an election with N voters, each voter i will receive a refund of $\frac{1}{N-1}\sum_{j\neq i} v_j^2$. Since this amount is received independently of i's voting decision it does not impact her behavior, and is hence overlooked in the remainder of this thesis. Notably, it is also possible to implement versions of QV based on artificial currencies. One such system, proposed by e.g. Laurence and Sher (2017), is to provide every voter with a fixed amount of *vote credits* that are then used to purchase votes (at a quadratic price) across a pool of elections. While this is often how QV has been implemented in practice, the theoretical literature has primarily covered monetary QV. This thesis will adhere to this norm.

3.2 Two Previous Solution Concepts

We have now defined the general setup of the quadratic voting model. Yet, to understand the outcomes it produces, we must further discuss how voters decide on a number of votes to place. Recall that every voter perceives voting as a trade-off between two goals: firstly, that of attaining her preferred election outcome with as high probability as possible, and secondly, that of minimizing her cost of voting. Lalley and Weyl (2018, 2019) propose two distinct models for how voters solve this problem, that will be succinctly laid out in this

[7]e.g. arguing that individuals have an intrinsic right to influence decisions that affect them.

[8]e.g. believing that individuals voluntarily agree to take part in collective decision-making because they expect to benefit personally.

section. Slight adaptations from the original presentations have been made for pedagogical and consistency purposes.

3.2.1 Price-taking equilibrium

The first solution concept, from Lalley and Weyl (2018)[9], involves voters perceiving themselves as purchasers of *influence* $I_i \in \mathbb{R}$, rather than votes v_i. Assume that in each election, the supply of such influence is finite and denoted by S. The *vote-price of influence*, denoted $p > 0$, is the cost, in votes, of one unit of influence. Purchasing a certain level of influence I_i is hence associated with a cost of $(pI_i)^2$. Since voters will demand higher levels of influence if the price is low, there exists a market-clearing equilibrium price p^* and corresponding influence levels $\{I_i^*\,;\, i \in \mathcal{N}\}$ for which

(i) $I_i^* = \arg\max_{I_i} \left\{u_i \cdot I_i - (p^* I_i)^2\right\}$ for every $i \in \mathcal{N}$, and

(ii) $\sum_i I_i^* = S$.

The (implicit) price of influence p^* is set in such a way that the supply is exerted, i.e. that $\sum_i I_i^* = S$. The key assumption in this solution concept is that voters will, without necessarily communicating, converge on a particular perceived price p^*. This is the number of votes that they expect is required to attain one unit of influence.

From (i), we find that voters in equilibrium will choose to buy a level of influence

$$I_i^* = \frac{u_i}{2(p^*)^2}, \tag{5}$$

corresponding to equilibrium votes of

$$v_i^* \equiv p^* I_i^* = \frac{2}{p^*} \cdot u_i. \tag{6}$$

Since v_i^* is proportional to u_i with a coefficient $\frac{2}{p^*} > 0$ that is shared among voters, we find that

$$\operatorname{sign}\left(\sum_{i \in \mathcal{N}} v_i^*\right) = \operatorname{sign}\left(\sum_{i \in \mathcal{N}} \frac{2}{p^*} \cdot u_i\right) = \operatorname{sign}\left(\sum_{i \in \mathcal{N}} u_i\right). \tag{7}$$

In this price-taking equilibrium model, quadratic voting is hence robustly efficient according to the specification in Definition 1.

[9]The notation used in this section is based on a previous version of the paper, which is also found in Laurence and Sher (2017).

3.2.2 Symmetric, Bayesian model

In the price-taking model described above, voters were assumed to perceive an identical price of influence p^*, but the emergence of this agreement was not specified. This section presents a possible game-theoretic microfoundation for such an agreement, also specified in Lalley and Weyl (2019).

Assume that each voter i is aware of her own preference u_i and the probability distribution f_u from which the preferences of other voters are drawn. She does not, however, know the realizations of u_j for $j \neq i$. A voter i purchasing v_i votes attains a utility of

$$U_i(v_i) = \Psi(v_i + V_{\text{-}i}) \cdot u_i - v_i^2 \tag{8}$$

where $V_{\text{-}i} \equiv \sum_{j \neq i} v_j$ and $\Psi : \mathbb{R} \to [0, 1]$ is a function of class $\mathbb{C}^\infty$ such that $\Psi(V) - \frac{1}{2}$ is odd and $\lim_{V\to\infty} \Psi(V) = 1$. In other words, $\Psi(v_i+V_{\text{-}i})$ is a continuous function approximating the discrete step function $X = X(v_i + V_{\text{-}i})$, that takes the value 0 if $v_i + V_{\text{-}i} < 0$ and 1 otherwise. The reason for this slight complication is proof-theoretical, rather than substantive. Voters will aim to maximize

$$\mathbb{E}\big[\Psi(v_i + V_{\text{-}i}) \cdot u_i - v_i^2\big] \tag{9}$$

subject to their beliefs about $V_{\text{-}i}$. A necessary condition for optimality is hence

$$v_i^* = \frac{u_i}{p(v_i^*)}, \text{ where } p(v_i) = \frac{1}{2 \cdot \mathbb{E}[\Psi'(v_i + V_{\text{-}i})]}. \tag{10}$$

The authors go on to show that this game allows a Nash equilibrium, and if the number of voters is sufficiently large, $p(v_i)$ is approximately constant in v_i. This entails that the optimal voting strategy v_i^* is approximately linear in u_i and that

$$\mathbb{P}\left(\text{sign}\Big(\sum_{i\in\mathcal{N}} v_i^*\Big) = \text{sign}\Big(\sum_{i\in\mathcal{N}} \frac{u_i}{p(v_i^*)}\Big) = \text{sign}\Big(\sum_{i\in\mathcal{N}} u_i\Big)\right) \tag{11}$$

is high. The inefficiency, defined as the probability of an election outcome that is suboptimal in light of the preferences u_i, approaches 0 as the number of voters grows large.

3.3 A Belief-Centered Model

The previous section laid out two solution concepts to the QV model – one based on voters perceiving themselves as purchasers of influence (sharing identical estimates of the price thereof), the other based on voters knowing the distribution of preferences f_u and voting on the assumption that this distribution is common knowledge, and that other voters are rational expected utility maximizers.

While these solution concepts elegantly convey the idea that QV has beneficial efficiency qualities in the abstract, they rely on strong symmetry assumptions. Notably, all voters share a common belief in the distribution of preferences, as well as how voters translate said preferences to voting strategies. What happen if voters, for whatever reason, do not adhere to these assumptions?

In this section, I will lay out a model of QV that allows each voter to have arbitrary beliefs about the voting behavior of others. I will specify how voters choose a strategy in this setting, and exemplify one set of assumptions for which QV is still robustly efficient in this framework.

Like before, let the preference of each voter i be characterized by a scalar u_i, describing the utility that i attains if the election results in a proposed change, compared to the continuation of the status quo. Let v_i denote the number of votes purchased by i, and $V_{\text{-}i} = \sum_{j \neq i} v_j$ denote the sum of votes purchased by other voters. Then, the utility of i is determined as

$$U_i(v_i|V_{\text{-}i}) = \begin{cases} u_i - v_i^2, & \text{if } v_i + V_{\text{-}i} \geq 0 \\ -v_i^2, & \text{if } v_i + V_{\text{-}i} < 0, \end{cases} \tag{12}$$

where the first case corresponds to the enactment of the proposed change, and the second to the continuation of the status quo.

Voters are assumed to choose v_i to maximize expected utility, subject to their beliefs about the voting behavior of others. Let v_i^* denote this optimal voting strategy. In the absence of uncertainty about how others will vote, this behavior is trivially computed.

Proposition 1. *Assume that a voter i is certain of the value of $V_{\text{-}i}$. Then,*

$$v_i^* = \begin{cases} 0, & \textit{if } \operatorname{sign}(u_i) = \operatorname{sign}(V_{\text{-}i}) \text{ or } |u_i| < V_{\text{-}i}^2, \\ |V_{\text{-}i}| \cdot \operatorname{sign}(u_i), & \text{otherwise.} \end{cases} \tag{13}$$

ꝏ

This result follows directly from the maximization of Equation 12 with respect to v_i for fixed values of $V_{\text{-}i}$. The intuition behind the result is that a voter will find it worthwhile to purchase a non-zero number of votes only if (i) abstaining from voting would lead to an undesirable outcome and (ii) the number of votes required to swing the election is sufficiently low to motivate the cost of the votes. As we move toward the case where the voter is uncertain about the value of $V_{\text{-}i}$, the same considerations will apply; the optimal voting behavior will be determined by the probability that the voter assigns to being able to swing the election without purchasing an excessively costly number of votes.

To study the voting choice of an individual that is uncertain of how others will vote, we will introduce the subjective probability operator $\mathbb{P}_i$ and the belief function f. For any event A, $\mathbb{P}_i(A)$ denotes the probability that i assigns to the realization of A.

Definition 2. Let i be a voter whose beliefs are described by the subjective probability operator $\mathbb{P}_i$. Then, the *belief function* f of i is a function such that for every $V_1, V_2 \in \mathbb{R}$,

$$\mathbb{P}_i(V_1 \leq V_{\text{-}i} \leq V_2) = \int_{V_1}^{V_2} f(V)\mathrm{d}V. \tag{14}$$

ↂ

The belief function f is a probability density function, satisfying $f(V) \geq 0$ for all V and $\int_{-\infty}^{\infty} f(V)\mathrm{d}V = 1$. Having defined an uncertainty measure over $V_{\text{-}i}$, we can now compute the *subjective expected utility* of i when placing v_i votes as

$$\begin{aligned} \mathbb{E}\Big[U_i(v_i)\Big] &\equiv \int_V U_i(v_i|V)f(V)\mathrm{d}V && (15)\\ &= -v_i^2 + u_i \int_{-v_i}^{\infty} f(V)\mathrm{d}V && (16)\\ &= -v_i^2 + u_i \left(\int_{-v_i}^{0} f(V)\mathrm{d}V + \int_0^{\infty} f(V)\mathrm{d}V \right) && (17) \end{aligned}$$

In this expression, $-v_i^2$ is the deterministic cost associated with placing v_i votes, $\int_0^{\infty} f(V)\mathrm{d}V$ is the probability that i assigns to the election resulting in a change if she abstains from voting, and $\int_{-v_i}^{0} f(V)\mathrm{d}V$ is the probability of changing the outcome of the election, conditionally on placing v_i votes. These probabilities are illustrated by the size of the blue and orange areas in Figure 1.

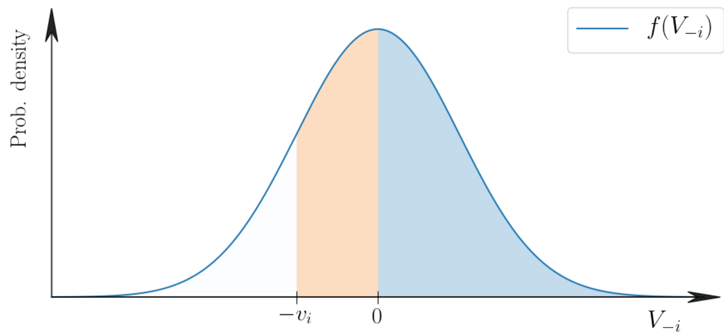

Figure 1: Graph of a possible belief function f. The size of the blue area represents the probability of winning the election without voting, i.e. the probability assigned to $V_{-i} \geq 0$. The orange area marks the additional probability of winning when placing v_i votes.

Having a closed-form expression for $\mathbb{E}[U_i(v_i)]$, we can compute the optimal vote v_i^* as

$$v_i^* = \arg\max_{v_i} \left\{ \mathbb{E}\Big[U_i(v_i)\Big] \right\} \tag{18}$$

$$= \arg\max_{v_i} \left\{ -v_i^2 + u_i \int_{-v_i}^{\infty} f(V) dV \right\}. \tag{19}$$

Without further assumptions on f, we can say little about the characteristics of v_i^*. Two conclusions may, however, be drawn. Firstly, refraining from voting ($v_i^* = 0$) may be optimal only for indifferent voters (for whom $u_i = 0$) and voters that perceive it to be impossible to swing the election (for whom $f(V) = 0$ in a neighborhood of 0). All other voters will place at least some small number of votes. Secondly, if f is continuous, then a necessary condition for optimality is that $2v_i^* = u_i \cdot f(v_i^*)$.

3.3.1 Sufficient conditions for efficiency

The above specification makes no assumptions on the shape of the belief function f of each voter apart from it constituting a probability density function. From the expression for the optimal voting strategy, Equation 19, it is clear that the beliefs are a crucial determinant of how a rational voter will act. In fact, the only positive effect of voting on utility is driven by the belief that V_{-i} might be close enough to 0 for a voter to plausibly swing the election. Any claim regarding the efficiency of outcomes from QV must hence be based on assumptions on these beliefs. Here, I will provide two conditions on the beliefs of different voters that jointly imply robustly efficient outcomes.

Firstly, assume the beliefs f of each voter to be identical. In the case of an election with 100 voters, this means that each voter shares the same subjective probability distribution over how the other 99 voters will choose to vote. This holds, for example, in the model presented in Section 3.2.2. If beliefs are identical, voters with identical preferences will vote identically; $u_i = u_j$ implies $v_i^* = v_j^*$ for every pair of voters i, j.

This condition, while strong, is not sufficient for efficiency, because it does not describe how the optimal voting scheme v_i^* changes with the value of u_i. For this, an additional assumption is necessary, namely that f is constant in a sufficiently large neighborhood of 0. By *sufficiently large* we mean that for any possible preference u_i, the belief function $f(V)$ is constant for $|V| < v_i^*$. Hence, each voter i must consider the probability of swinging the election to be linear in v_i, for the values of v_i that she realistically considers to place. Since v_i^* is itself a function of f (see Equation 19), this condition has a recursive nature, but the intuition is straight-forward: a voter must consider the possibility of swinging the election to be twice as high when buying two votes instead of one, but not necessarily when buying two billion votes instead of one billion votes.

If the above condition holds, then $f(V) = c$ for some $c > 0$ on this neighborhood of 0. The optimal voting strategy can then be computed as

$$\begin{aligned} v_i^* &= \arg\max_{v_i}\{-v_i^2 + u_i \int_{-v_i}^{\infty} f(V)\mathrm{d}V\} && (20)\\ &= \arg\max_{v_i}\{-v_i^2 + u_i \int_{-v_i}^{\infty} c \cdot \mathrm{d}V\} && (21)\\ &= \arg\max_{v_i}\{-v_i^2 + u_i \cdot v_i \cdot c\} && (22)\\ &= \frac{1}{2} c u_i. && (23) \end{aligned}$$

The number of votes purchased by each voter is hence linear in her preference u_i. Recall that robust efficiency requires that

$$\operatorname{sign}\left(\sum_{i \in \mathcal{N}} v_i^*\right) = \operatorname{sign}\left(\sum_{i \in \mathcal{N}} u_i\right), \tag{24}$$

which holds in this case, since

$$\operatorname{sign}\left(\sum_{i \in \mathcal{N}} v_i^*\right) = \operatorname{sign}\left(\sum_{i \in \mathcal{N}} \frac{1}{2} c u_i\right) = \operatorname{sign}\left(\sum_{i \in \mathcal{N}} u_i\right). \tag{25}$$

How plausible is the assumption that $f(V)$ is approximately flat around 0? Intuitively, it appears defensible when the number of voters is large. In this case, the range of plausible values of $V_{\text{-}i}$ would likely be much wider than for the plausible values of v_i^*, analogously to how the variance of a sum of N independent random variables increases with N. Voters would then believe that $|V_{\text{-}i}| >> |v_i^*|$ with high probability, and likely assign similar probabilities to (for example) $V_{\text{-}i} \approx \frac{1}{2} \cdot v_i$, $V_{\text{-}i} \approx v_i$ and $V_{\text{-}i} \approx 2 \cdot v_i$, for the values v_i that they are considering. If the number of voters is low, however, the same argument need not apply, and the assumption of f being flat around 0 appears less well-founded.

4 Application

We have now specified a belief-centered model of QV that allows for arbitrary, asymmetrical beliefs among voters. This model highlights that the optimal voting behavior is crucially determined by a voter's beliefs, and suggests that the robust efficiency of voting outcomes is achieved only in particular circumstances – for example when beliefs are identical and flat around 0.

In this section, we will explore two ways in which this novel specification can help us gain insight into voting behaviors and likely outcomes of quadratic voting. First, we will consider how the efficiency of outcomes is impacted when voters differ in beliefs. To this aim, we will simulate elections where the preference u_i and belief c_i of each voter is sampled randomly from some joint distribution. We can thereby compute how each such simulated participant would have voted, and hence the outcome of the election. By repeating this type of simulation for different distributions of u_i and c_i, we can analyze the case when preferences and beliefs are independent as well as the consequences of correlations between them.

Second, we will analyze the incentives of an individual voter to acquire information, such as polling data, about how others will vote. As voters with more accurate beliefs tend to achieve better outcomes in strategic voting situations, we expect them to be willing to pay non-zero amounts for this kind of information. This is modeled by considering a voter adhering to Bayesian rationalism, i.e., having some prior belief about how other will vote that she updates (in accordance with Bayes' rule) upon observing new evidence. I provide an expression for the expected value of an arbitrary piece of information, as long as the recipient of the information has a priori knowledge of its evidential accuracy.

This section can be thought of as an exploratory analysis of the relationships between beliefs and voting behaviors in quadratic voting. To the extent that these relationships are modeled quantitatively, the parameter values used are intended to illustrate certain mechanisms, rather than to realistically describe real-world situations. The purpose is to add to the list of potential ways that QV outcome can be less efficient in practice than in theory, rather than to generate sharp predictions about the relative magnitude of different problems.

4.1 Efficiency Effects of Heterogeneous Beliefs

The fundamental argument in favor of QV is its alleged property of implementing *efficient* outcomes, i.e., outcomes that maximize the aggregate utility of participating voters. To what extent can we expect this property to be attained when voters differ in beliefs? This section will explore this key question through a simple Monte Carlo simulation study.

Let the belief of each voter i be characterized by a probability density function f_i. Assume, as in Section 3.3.1, that each f_i is constant in a neighborhood 0, i.e. that $f_i(V) = c_i > 0$ for sufficiently small magnitudes of V. We will refer to the constant c_i as the *perceived pivotality* of i, since it describes the likelihood that i assigns to swinging the election, conditionally on placing a certain number of votes. From Equation 19, we find that a voter with preference u_i will choose to vote as

$$
\begin{aligned}
v_i^* &= \arg\max_{v_i} \Big\{ - v_i^2 + u_i \int_{-v_i}^{\infty} f(V) \mathrm{d}V \Big\} && (26)\\
&= \arg\max_{v_i} \Big\{ - v_i^2 + u_i \cdot v_i \cdot c_i \Big\} && (27)\\
&= \frac{1}{2} \cdot c_i \cdot u_i. && (28)
\end{aligned}
$$

An election between such voters will be determined by the sign of $\sum_i v_i^*$, which is identical to the sign of $\sum_i c_i u_i$. The outcome is *efficient* insofar as this sign is equal to the sign of $\sum_i u_i$.

Assume that for each voter i, u_i and c_i are sampled randomly from some joint, two-dimensional distribution f_{uc}. How likely is it that the outcome of the election coincides with the preferred outcome, i.e. what is

$$
\mathbb{P}_{-} \equiv \mathbb{P}\left(\operatorname{sign}\Big(\sum_i c_i u_i \Big) = \operatorname{sign}\Big(\sum_i u_i \Big) \right) \tag{29}
$$

for different generating functions f_{uc}, and numbers of voters? Similarly, we might be interested in the expected welfare loss of an election where voters differ in beliefs. To this end, we will define the *utility loss* L of an election as

$$
L = \begin{cases} \frac{1}{N} \Big| \sum_i u_i \Big|, & \text{if} \operatorname{sign}\left(\sum_i c_i u_i\right) \neq \operatorname{sign}\left(\sum_i u_i\right) \\ 0, & \text{otherwise,} \end{cases} \tag{30}
$$

where N denotes the number of voters. This quantity takes a value of 0 if the election outcome is optimal, and $\frac{1}{N} | \sum_i u_i |$ otherwise. Similarly, the *expected loss* $\mathbb{E}[L]$ is equal to 0 if any only if the optimal outcome is implemented with probability 1.

4.1.1 Independent sampling

As a base case to illustrate the detrimental efficiency effects of heterogeneity in beliefs, consider the scenario where preferences and beliefs are independently distributed. Let $u_i \sim U(-1,1)$ and $c_i \sim U(0,1)$. Figure 2 shows simple point estimates, i.e. observed average values, of $\mathbb{P}(X=1)$, $\mathbb{P}_{-}$, and $\mathbb{E}[L]$, for $N = 1, ..., 10^4$ voters. As expected in this symmetric case, about half of all elections result in $X = 1$ regardless of N. In addition, the number of inefficient outcomes is fairly constant around 16% for any $N \geq 3$. The expected loss per voter, however, decreases with N. This follows directly from the central limit theorem – the distribution of $\frac{1}{N}\sum_i u_i$ will asymptotically approach a zero-mean Gaussian distribution with variance decreasing in N. The average voter will hence grow "asymptotically indifferent" to the election outcome as the number of voters grows large.

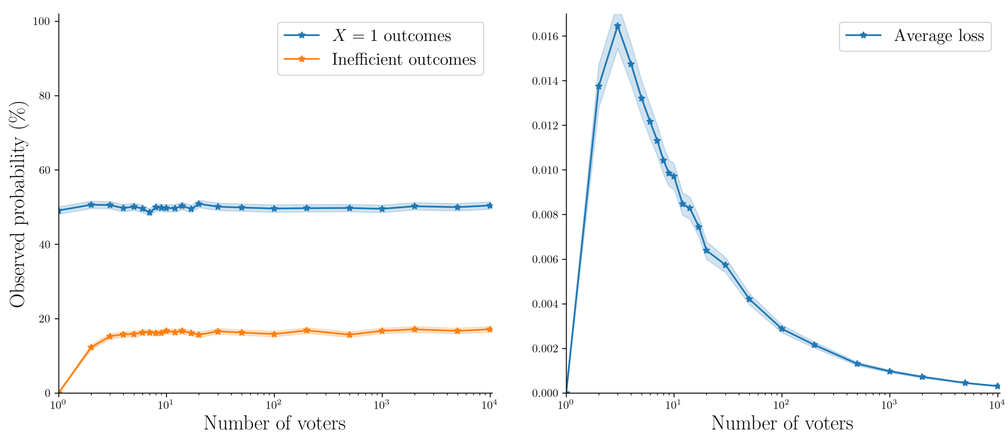

Figure 2: Monte Carlo estimated probabilities of $X = 1$ and inefficient election outcomes (left) and average loss (right) as a function of the number of participating voters. Shaded areas represent confidence intervals at significance level $\alpha = 0.05$ (exact binomial intervals are used for the probability estimates, while the average loss intervals are computed from asymptotic normality). Based on 10^4 simulations for each number of voters.

This simple simulation contrasts with the case of shared beliefs, $c_i = c$ for every voter i, where $\mathbb{P}_{-} = \mathbb{E}[L] = 0$ with certainty. The mere introduction of noise into beliefs is hence detrimental to the efficiency of voting outcomes, even if this noise is independent of the preferences of voters. In large elections, however, the expected efficiency loss of belief heterogeneity tends to 0 as long as preferences and beliefs are independent.

4.1.2 Correlated sampling

In the real world, however, few variables that describe individuals are independent or even uncorrelated. It is not difficult to hypothesize that there may be patterns in which individuals perceive themselves as pivotal in QV elections. Perhaps highly educated voters (with preferences predictably different from the population as a whole) estimate their pivotality more accurately than the average person. Or perhaps some particular ideological group tend to overestimate their impact on the election results. To model belief-preference correlation, consider the case where we sample c_i and u_i jointly, allowing for the following relationship:

$$u_i \sim U(-1, 1) \tag{31}$$

$$c_i = \rho\left(\frac{u_i + 1}{2}\right) + (1 - \rho)\hat{c}_i \tag{32}$$

$$\text{where } \hat{c}_i \sim U(0, 1) \text{ and } 0 \leq \rho \leq 1. \tag{33}$$

From this generation, it is still the case that $u_i \in [-1, 1]$ and $c_i \in [0, 1]$, but higher u_i have a tendency to come together with higher c_i – a relationship that grows with ρ. In the edge cases, $u_i = 1$ leads to $c_i \sim U(\rho, 1)$ and $u_i = -1$ leads to $c_i \sim U(0, 1 - \rho)$.

Even a relatively low correlation ρ, like $\rho = 0.1$ in Figure 3 below, leads to $X = 1$ becoming the dominant outcome as the number of voters grows large. It does, however, require about 10000 voters for the probability of $X = 1$ to approach 100%. Interestingly, the average loss (right) approaches 0 for large N even as the probability of an inefficient outcome increases. The reason for this is similar to that of the independent sampling; as long as $\mathbb{E}[u_i] = 0$, the sum $\frac{1}{N}\sum_i u_i$ will converge towards 0. In words, our assumption that the mean person is indifferent about election outcomes leads large populations to be nearly indifferent on average.

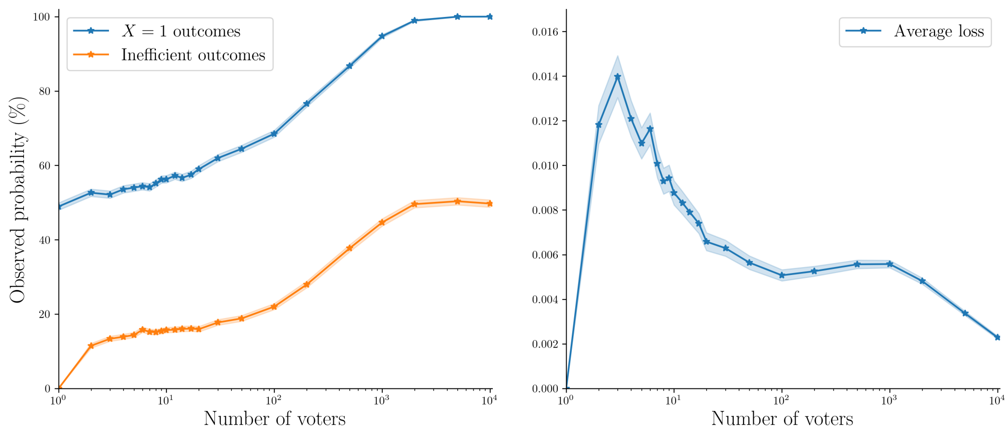

Figure 3: Simulation results for correlated sampling.

4.1.3 Non-zero mean preferences

So far, we have found that heterogeneous beliefs induce noise in elections. This lowers the average efficiency of chosen outcomes even when the average voter is indifferent, but the effects diminishes as the number of voters grows large. Now, consider the case where voters on average prefer one outcome over the other, but those who prefer the other outcome are more confident in their ability to swing the election. Let $\bar{u} \in \mathbb{R}$ denote the average preference and ρ the preference-belief correlation. Sample u_i and c_i as:

$$u_i = \bar{u} + \hat{u}_i \tag{34}$$

$$c_i = \rho \left(\frac{\hat{u}_i + 1}{2} \right) + (1 - \rho)\hat{c}_i \tag{35}$$

$$\text{where } \hat{c}_i, \sim U(0,1) \text{ and } \hat{u}_i \sim U(-1,1) \tag{36}$$

The interesting case is when $\bar{u} < 0$, when the average voter is skeptical to the proposed change, but enthusiastic voters are more likely to view themselves as pivotal. This makes preferences and beliefs "pull in different directions". Figure 4 shows the simulation results for $\rho = 0.1$, $\bar{u} = -0.01$:

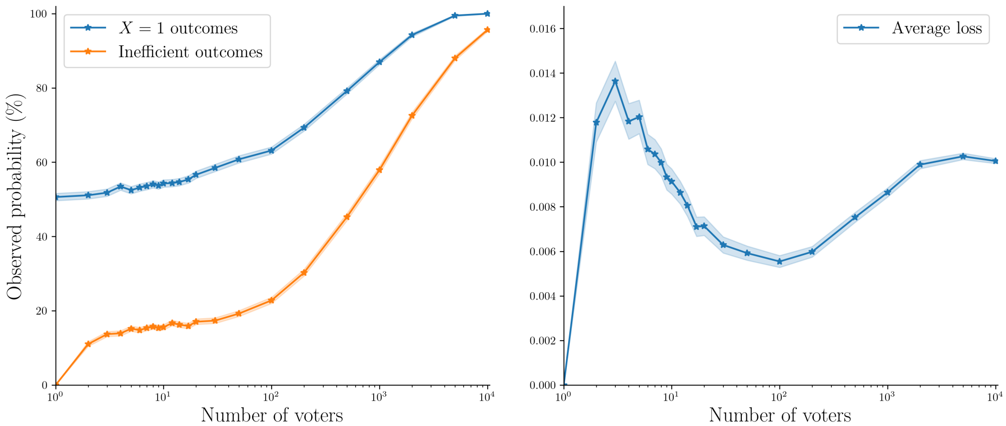

Figure 4: Simulation results for non-zero mean preferences with correlated beliefs.

With these parameter values, beliefs trump preferences – as the number of voters increases, the number of elections with a positive outcome ($X = 1$) approaches 100%. The average utility loss (right) exhibits a non-monotone pattern – it is high for < 10 voters, low for about 100 voters, and high again for > 1000 voters. I conjecture that the high loss when the number of voters is low is due to statistical noise (it's relatively more common e.g. that one outlier voter, perceiving herself as highly pivotal, swings the election in an inefficient direction) while the high loss for high N is due to the systematic relationship between c_i and u_i.

In this series of simulations, we have seen that when voters differ in their perceived pivotality, the ability of QV to implement efficient outcomes deteriorates. As long as beliefs and preferences are uncorrelated, this issue primarily arises when the number of voters is low. If certain preferences go together with certain beliefs, however, efficiency is seriously harmed regardless of the number of participating voters.

The next application will discuss one possible way in which differences in beliefs may arise: different voters may have access to different information.

4.2 Value of Information

We have found that the optimal voting strategy, and hence the expected utility, of each voter depends on her beliefs. Does this imply that voters with more accurate beliefs systematically attain higher levels of utility, than peer with less accurate beliefs? If so, how much would an individual voters be willing to pay for accurate information about the voting behavior of others?

This section presents a framework for how Bayesian voters estimate the value of information about the voting behavior of others. The first part provides a primer on the central concepts of Bayesian rationality. The second part specifies how these relate to the situation of a QV participant observing some piece of information, and the third formalizes how highly the participant values said information.

4.2.1 Bayesian rationality

Bayesian rationality is a method of probabilistic reasoning specifying how to update beliefs in light of new information. It owes its name to Bayes' rule, that states that for any hypothesis H and evidence E (such that $\mathbb{P}(E) > 0$), it holds that

$$\mathbb{P}(H|E) = \frac{\mathbb{P}(E|H) \cdot \mathbb{P}(H)}{\mathbb{P}(E)}. \tag{37}$$

The formula tells us that the correct probability to assign to a hypothesis H after observing evidence E, i.e. $\mathbb{P}(H|E)$, is proportional to the product of the likelihood of the evidence given the hypothesis, $\mathbb{P}(E|H)$, and the *prior probability* $\mathbb{P}(H)$ assigned to the hypothesis before observing evidence. $\mathbb{P}(E)$ denotes the prior probability assigned to observing the evidence E, rather than some distinct piece of evidence $\tilde{E}$.

In its original form, Bayes' rule is specified for discrete events. The formula can, however, be trivially translated to continuous hypotheses and evidence. Letting f denote probability density functions, the equivalent update rule is

$$f(h|e) = \frac{f(e|h) \cdot f(h)}{f(e)} \tag{38}$$

where $f(h)$ is often referred to as the *prior belief*, $f(e|h)$ as the *likelihood function*, and $f(h|e)$ as the *posterior belief*.

4.2.2 A Bayesian voter

For some voter i, let

- $f_0(V_{\text{-}i})$ denote her prior belief over $V_{\text{-}i}$ before observing the evidence in question,
- e represent the evidence: a noisy observation of the true value of $V_{\text{-}i}$ such that $f(e|V_{\text{-}i})$ is known and symmetric around $V_{\text{-}i}$,
- $f_e(e)$ denote i's belief over what evidence will be observed, and
- $f_p(V_{\text{-}i}|e)$ describe i's posterior belief over $V_{\text{-}i}$ after observing a particular piece of evidence e.

From the continuous version of Bayes' rule, we find that these quantities must be related as

$$f_p(V_{\text{-}i}|e) = \frac{f(e|V_{\text{-}i}) \cdot f_0(V_{\text{-}i})}{f_e(e)} \tag{39}$$

where the denominator

$$f_e(e) = \int_V f_0(V) f(e|V) \mathrm{d}V \tag{40}$$

can be interpreted as a normalizing constant to ensure that the posterior is indeed a probability density for any observed e. We can train our intuition about Equation 39 by considering two extreme cases.

First, assume that the observed evidence e is perfect, that is, equal to the true value $V_{\text{-}i}$ with certainty. If this is the case, then $f(e|V_{\text{-}i})$ is a degenerate distribution – $f(e|V_{\text{-}i}) = 0$ for any $e \neq V_{\text{-}i}$. In this case, the posterior $f_p(V_{\text{-}i}|e)$ will also take on a non-zero value at a single point only. Hence, any voter observing perfect evidence will (intuitively) be certain about the value of $V_{\text{-}i}$ afterwards.[10]

As an alternative example, let the voter be certain about $V_{\text{-}i}$ *before* observing evidence, i.e., let $f_0(V^*_{\text{-}i}) > 0$ for a single value $V^*_{\text{-}i}$. Then, the voter will still be sure that this is the correct value after observing evidence, as long as $f(e|V^*_{\text{-}i}) > 0$ for the observed e.[11] This captures the intuition that additional information has no (ex-ante) value for someone who is already certain about the value of a parameter.

4.2.3 Value of information

Having formalized how voters update their beliefs in the light of new information, we can now determine how that information might induce better voting decisions. Without evidence, Equation 16 shows that a voter i will expect to attain a utility of

$$\mathbb{E}_0[U_i|v^*_i] = -(v^*_i)^2 + u_i \int_{-v_i}^{\infty} f_0(V) \mathrm{d}V \tag{41}$$

$$= -(v^*_i)^2 + u_i \cdot \mathbb{P}_0(v^*_i + V \geq 0) \tag{42}$$

[10]This assumes that, for the true value $V^*_{\text{-}i}$, $f_0(V^*_{\text{-}i}) > 0$.

[11]If, for some observed e, $f(e|V_{\text{-}i}) \cdot f_0(V_{\text{-}i}) = 0$ for every $V_{\text{-}i}$, then the posterior belief is not well-defined. We will disregard this problem by assuming that each voter assigns non-zero prior probability to the true value $V^*_{\text{-}i}$.

where $\mathbb{P}_0$ refers to the subjective probability with respect to the prior f_0. If she expects to observe evidence, she will instead expect to attain a utility of

$$\mathbb{E}_e\Big[\mathbb{E}_p\big[U_i|v_i^e\big]\Big] = -\mathbb{E}_e\big[(v_i^e)^2\big] + u_i \cdot \mathbb{E}_e\Big[\mathbb{P}_p(v_i^e + V \geq 0)\Big] \tag{43}$$

where $\mathbb{E}_e$ is the expectation over what evidence will be observed, v_i^e is the optimal vote after observing a particular evidence e and $\mathbb{P}_p$ is the probability with respect to the posterior distribution f_p (that depends on the evidence).

The difference between equations 41 and 43 shine light on the value of the evidence. Evidence can be beneficial in two distinct ways:

1. the evidence can lead i to vote less in expectation, if $\mathbb{E}_e\big[(v_i^e)^2\big] < (v_i^*)^2$, and
2. the evidence can lead to a higher probability of winning the election, if it is the case that $\mathbb{E}_e\Big[\mathbb{P}_p(v_i^e + V \geq 0)\Big] > \mathbb{P}_0(v_i^* + V \geq 0)$.

Notably, both of these positive effects can occur simultaneously. Consider the following example.

Example 2 (Perfect evidence)**.** Let i be a voter with $u_i = 2$ and a prior belief $f_0(V) = \frac{1}{2}$ for $-1 \leq V \leq 1$. From Equation 19 we find that, without information, she will place $v_i^* = \frac{1}{2}$ votes at a cost of $(v_i^*)^2 = \frac{1}{4}$. Furthermore, she will expect to win with a probability of $\frac{3}{4}$, that is, the probability that she assigns to $V_{\text{-}i} + \frac{1}{2} \geq 0$.

Now, assume that she will obtain perfect evidence, i.e. observe some $e = V_{\text{-}i}^*$. Her optimal vote is then given by Proposition 1. If she observes $e \geq 0$ (which, from her prior, she expects will happen with 50% probability), she will infer that $V_{\text{-}i} \geq 0$ and place no votes, since she expects to win without voting. If she observes $-1 < e < 0$, she will place $v_i^e = |e|$ votes – just enough to win the election. Since her prior is that the true value $V_{\text{-}i}^*$ is necessarily greater than or equal to -1, she considers any smaller value of e impossible. Consequently, prior to observing evidence, i expects to pay $\frac{1}{6}$ for her votes[12], and win with probability 1.

The evidence has hence both reduced the expected cost of her votes from $\frac{1}{4}$ to $\frac{1}{6}$, and increased her perceived probability of winning from $\frac{3}{4}$ to 1.[13] ꝏ

Seeing that voters expects to attain a higher utility from observing some evidence, it is natural to introduce a term for the value of this evidence.

[12] $\mathbb{E}_e[(v_i^e)^2] = \int_{e=-1}^{0} |e|^2 f_e(e)de = \frac{1}{6}$, where $f_e(e) = f_0(e)$, since her "prior of what evidence she will observe" is the same as her prior over the value of $V_{\text{-}i}$.

[13] One may note that if instead $u_i = 1$, then the evidence *increases* the expected cost of voting. The value of the information is still positive, however, since this increase in cost is compensated by a higher probability of a favorable election outcome.

Definition 3. Let e be some noisy observation of $V^*_{\text{-}i}$. Then, the (ex-ante) *value of information* VoI of e for a voter i is

$$\text{VoI} = \mathbb{E}_e\Big[\mathbb{E}_p\big[U_i|v_i^e\big]\Big] - \mathbb{E}_0\big[U_i|v_i^*\big]. \tag{44}$$

In Example 2, the voter i expects a utility of $\frac{5}{4}$ without evidence[14], and $\frac{11}{6}$ with evidence[15]. The value of information can be computed as

$$\text{VoI} = \frac{11}{6} - \frac{5}{4} = \frac{7}{12}. \tag{45}$$

The voter will hence be willing to pay $\frac{7}{12}$ monetary units to attain the information in question. With lower-quality evidence (i.e. when $f(e|V)$ is not a degenerate distribution), the value of information must typically be computed through numerical methods.

The above analysis provides support for the claim that participants in QV elections have incentives to inform themselves about the voting behaviors of others. Since this incentive is likely to be larger for voters with intense preferences, there is a risk that such voters would go to greater lengths than the average voter to estimate their pivotality accurately. This, in turn, is a potential source of belief heterogeneity.

5 Conclusions and Discussion

The topic of this thesis has been to explore the role of strategic beliefs in quadratic voting: how does the optimal voting behavior depend on a voter's beliefs about other voters, and how does heterogeneity in said beliefs impact election outcomes? To this end, I have formalized a novel, belief-centered model of voting behaviors. I have then explored two applications of this model: one analyzing how heterogeneity in beliefs affects voting outcomes, and one formalizing an individual voter's incentive to acquire information about how others will vote.

The first application finds that belief heterogeneity has detrimental effects on the outcomes of QV elections, even if beliefs and preferences are independently distributed. If the purpose of the voting system is to aggregate preferences, then belief heterogeneity can be seen as a noise introduced into the aggregation process – the higher the levels of heterogeneity and the lower the number of voters, the smaller the likelihood that the optimal alternative will be chosen. If beliefs and preferences are correlated, the heterogeneity in beliefs is no longer just a noise in the aggregation process, but also a bias. Outcomes can then be systematically inefficient even as the number of voters grows large.

[14] Paying $\frac{1}{4}$ for 75% probability of winning (and gaining $u_i = 2$ utility).
[15] Paying in expectation $\frac{1}{6}$ for certain win.

The second application specifies how a rational voter may have a willingness to pay for information about the voting behavior of others. This highlights that QV may give disproportionate influence to well-informed individuals and interest groups. Although this tendency exists in any voting system when more than two alternatives are under consideration, we have at least three reasons to believe if might have a particularly large impact under QV. Firstly, QV encourages participants to vote strategically even in binary elections, such as the final stages of many countries' presidential elections.

Secondly, the set of possible votes is infinite in QV (corresponding to the entire real axis). Placing an optimal number of votes therefore requires nuanced estimation of the minuscule probability of swinging the election. Under plurality rule, in contrast, voters typically consider no more than a few different vote options. The threshold for information to impact a voter's strategy is therefore higher. Thirdly, the traditionally most common method of acquiring voting-relevant information – polls – is likely to be less useful under QV. This will be the case at least as long as said polls do not reliably capture preference intensities.

In summary, voting under QV appears to require higher levels of cognitive ability and information than under most other common voting systems. This may in and of itself generate heterogeneity in the voting strategies employed by different voters, which would in turn lead to higher inefficiency than predicted by the "belief-symmetric" models in the existing literature.

Throughout this thesis, voters have been assumed to follow the rational choice model of voting: they vote with the sole motivation of potentially affecting the election outcome in a direction that aligns with their preference. In recent decades, this framing of voter motivation has been criticized. For example, it has been argued that an individual voter typically has a relatively modest stake in the outcome of each election (rarely above the order of a few thousand dollars), and should in many voting situations hardly expect to change the election outcome with a higher probability than perhaps one in a million. The expected utility of voting – the product of the probability of swinging the election and the benefit of a favorable outcome – would then likely be in the magnitude of cents. This is, it is argued, far from sufficient to compensate for the inconvenience of making a trip to a polling station. This observation, coined the *paradox of voting* in the seminal work of Downs (1957), challenges the rational choice model of voting.

Three solutions are commonly proposed to the alleged paradox. Firstly, it may simply be the case that voters vastly overestimate their probability of swinging an election. They may hence, subjectively but erroneously, estimate voting to be worthwhile from a selfish, expected utility maximizing perspective (Acevedo & Krueger, 2004). This hypothesis would, if true, give further weight to the concerns put forward in this thesis, since it implies that voters have systematically biased beliefs about their probability of impacting election outcomes. Secondly, individuals may choose to vote for prosocial reasons; they may value not only the potential benefits of a new policy to themselves, but to other individuals, too (Edlin et al., 2007). This greatly increases the total expected benefits of voting, to the point of making it rational for sufficiently altruistic voters. Thirdly, a voter may gain utility not primarily from the outcome of an election, but from the act of voting itself – perhaps by being perceived as a responsible citizen, signaling some tribal membership, or experiencing a sense of civic duty (Riker & Ordeshook, 1968). If this were true, beliefs about pivotality are likely to be of little relevance for voting decisions, and the formal method of analysis deployed in this thesis largely misguided.

Of course, voter motivations may differ substantially between contexts – a citizen voting for presidential candidates may have different decision criteria from a board member voting about budget priorities for a corporation. Although many examples in this essay have been drawn from the political sphere, QV can be used equally well for collective decision-making in private organizations or informal situations. If motivations of voters are context-dependent, then the same is likely to hold for the strategic role of beliefs.

This thesis is an exploratory analysis of the strategic and informational aspects of quadratic voting. The arguments put forward are generally driven by simplified examples, and the extent to which the problems apply in real-world elections is far from determined. One particularly promising avenue of future research is therefore more empirical and qualitative. In the real world, what considerations do QV participants state as influential to their voting decisions? How salient are arguments about the expected closeness of the elections? Do voting strategies change significantly with experience? This class of research questions seem fitting for experimental study.

References

Acevedo, M., & Krueger, J. I. (2004). Two egocentric sources of the decision to vote: The voter's illusion and the belief in personal relevance. *Political Psychology*, *25*(1), 115–134.

Cain, B. E. (1978). Strategic voting in Britain. *American Journal of Political Science*, *22*, 639–655.

Casella, A., & Sanchez, L. (2019). *Storable votes and quadratic voting. An experiment on four California propositions* (Working Paper No. 25510). National Bureau of Economic Research.

Downs, A. (1957). *An economic theory of democracy*. Harper.

Eason, B. (2019). Colorado's experiment in quadratic voting tried to remove politics from budget decisions [Accessed: Jul 20, 2020]. *Colorado Sun*. https://coloradosun.com/2019/05/28/quadratic-voting-colorado-house-budget

Edlin, A., Gelman, A., & Kaplan, N. (2007). Voting as a rational choice: Why and how people vote to improve the well-being of others. *Rationality and Society*, *19*(3), 293–314.

Forsythe, R., Myerson, R. B., Rietz, T. A., & Weber, R. J. (1993). An experiment on coordination in multi-candidate elections: The importance of polls and election histories. *Social Choice and Welfare*, *10*(3), 223–247.

Gallup Inc. (2020). Party affiliation [Accessed: Aug 14, 2020]. https://news.gallup.com/poll/15370/party-affiliation.aspx/

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

Hylland, A. (1980). *Strategy proofness of voting procedures with lotteries as outcomes and infinite sets of strategies* [Unpublished].

Kawai, K., & Watanabe, Y. (2013). Inferring strategic voting. *American Economic Review*, *103*(2), 624–62.

Lalley, S. P., & Weyl, E. G. (2018). Quadratic voting: How mechanism design can radicalize democracy. *AEA Papers and Proceedings*, *108*, 33–37.

Lalley, S. P., & Weyl, E. G. (2019). *Nash equilibria for quadratic voting* [Available at SSRN at 2488763].

Laurence, B., & Sher, I. (2017). Ethical considerations on quadratic voting. *Public Choice*, *172*(1-2), 195–222.

May, K. O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, *20*, 680–684.

Myatt, D. P. (2007). On the theory of strategic voting. *The Review of Economic Studies*, *74*(1), 255–281.

Riker, W. H., & Ordeshook, P. C. (1968). A theory of the calculus of voting. *The American Political Science Review*, *62*(1), 25–42.

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

Weyl, E. G. (2017). The robustness of quadratic voting. *Public Choice*, *172*(1-2), 75–107.