Public Choice (2017) 172:195–222
DOI 10.1007/s11127-017-0413-4

Ethical considerations on quadratic voting

Ben Laurence[1] · **Itai Sher**[2]

Received: 31 August 2016 / Accepted: 30 September 2016 / Published online: 27 February 2017

Abstract This paper explores ethical issues raised by quadratic voting. We compare quadratic voting to majority voting from two ethical perspectives: the perspective of utilitarianism and that of democratic theory. From a utilitarian standpoint, the comparison is ambiguous: if voter preferences are independent of wealth, then quadratic voting outperforms majority voting, but if voter preferences are polarized by wealth, then majority voting may be superior. From the standpoint of democratic theory, we argue that assessments in terms of efficiency are too narrow. Voting institutions and political institutions more generally face a legitimacy requirement. We argue that in the presence of inequalities of wealth, any vote buying mechanism, including quadratic voting, will have a difficult time meeting this requirement.

Keywords Quadratic voting · Majority voting · Utilitarianism · Democratic legitimacy

1 Introduction

Quadratic voting (QV) is a voting mechanism in which voters purchase votes. The price of votes is proportional to the square of the number of votes purchased. The alternative for which the most votes are purchased wins. QV applies to binary decisions. Lalley and Weyl (2016) present a model in which QV is approximately efficient in large populations. Posner and Weyl (2015) propose a political application of QV as a method for the democratic

✉ Itai Sher
itaisher@gmail.com

Ben Laurence
benlaurence@uchicago.edu

[1] Department of Philosophy, University of Chicago, 1115 E. 58th St., Chicago, IL 60637, USA

[2] Department of Economics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA

selection of legislation and candidates, and argue that it is superior to majority rule. In this paper, we explore some ethical considerations on voting with a view to evaluating the normative case for QV over majority rule in the political context.

The most pressing and intuitive concern with vote buying schemes like QV involves the way they transform economic inequality into political inequality. For, on the face of it, vote buying seems to give greater weight to the preferences of the wealthy who can more easily afford to buy votes than the less wealthy. This is the first objection that will come to mind for most people upon an initial encounter with QV. In order to judge the relative normative credentials of QV and majority rule, we seek to carefully distinguish and explore two distinct concerns in this neighborhood. In Sect. 2, we explore the interaction of QV with economic inequality from a utilitarian point of view. We find that if voter preferences are independent of wealth, then QV is utilitarian preferred to majority voting, but if voter preferences are polarized by wealth, then majority voting may be superior. In Sect. 3, we move beyond the utilitarian perspective to consider questions that concern democratic legitimacy. This latter part of our paper addresses the question: why should we treat the fact that wealthier citizens can buy more votes differently from the fact that they can buy more of everything else? We argue that in the presence of inequalities of wealth, any vote buying mechanism, including QV, will have a difficult time meeting a democratic legitimacy requirement. We end by mentioning a variant of QV that would not be subject to these worries.

2 Utilitarian considerations

This section explores the utilitarian assessment of QV. Our argument, like that of Lalley and Weyl (2016), is spelled out in the context of a formal economic model. However, the argument is, in essence, philosophical. So we begin by sketching our argument informally. For a reader who is primarily interested in the broad philosophical argument, we recommend reading only this introductory part of Sect. 2, skipping the technical sections—Sects. 2.1–2.4—and moving directly to Sect. 3, where we assess QV in terms of democratic legitimacy. Within the utilitarian section, Sects. 2.3.5 and 2.4 are relatively less technical and contain material of substantive and philosophical interest. So some readers less interested in technical aspects may also want to skim the preceding sections for essential details and focus on these. Readers who are skeptical that the broad philosophical argument really addresses the economic arguments for QV are invited to read Sects. 2.1–2.4.

Suppose that we must make a public decision between two alternatives. Under majority voting, voters can express support for the alternative they prefer, but not the intensity of their preference. So, for example, an alternative slightly preferred by many people might beat an alternative that is passionately preferred by a smaller number. Intuitively, from a utilitarian perspective, this is bad because the losers may forgo more utility than the winners gain.

The thought animating QV is that if we put a price on votes, then voters could express their preference intensity—more passionate voters are willing to pay more—and we solve the problem. Specifically, in QV, it costs v^2 dollars to purchase v votes, which can then be used to vote for either alternative. The alternative with more votes wins. The pricing rule v^2 induces voters to purchase votes in proportion to their utility for the outcome, it is claimed, and so the outcome that maximizes aggregate utility wins. Why v^2 is the right pricing rule is explained in Sects. 2.1–2.2.

We don't think this argument is correct. We think that the arguments in favor of QV show only that under QV, voters will purchase votes in proportion to their willingness to pay. Willingness to pay is not the same as utility in an ethical sense. It is not utility in any of the senses that have been proposed by philosophers. This point has also been long understood by economists. In Sect. 2.3.1 we lay out our understanding of ethical utility.

There is a systematic divergence between utility and willingness to pay: if a rich person and a poor person care about a decision equally—the decision has the same impact on their utility—the rich person will be willing to pay more than the poor person for that decision to be made. This is because the rich person will have to give up only certain luxuries if she spends \$$y$, whereas the poor person will have to give up more basic wants or necessities for the same expenditure. So under QV, the preferences of the rich will be overrepresented relative to their true ethical weight.

This *could* mean that majority voting is better than QV from a utilitarian point of view. If preferences are independent of wealth, in the sense that overall poor voters have the same distribution of utilities for outcomes as rich voters, then QV will be optimal from a utilitarian perspective. The divergence between willingness to pay and utility will not cause any distortion in the outcome. If, on the other hand, preferences are polarized by wealth, so that the poor prefer one thing and the rich another, then the louder voice of the rich under QV may win out even when the poor care more. If there are more poor then rich, then under majority voting the poor win, as the utilitarian would have liked. It is fundamentally ambiguous whether QV or majority voting is better from a utilitarian standpoint.

The rest of the section is devoted to spelling out this argument in a formal way. We focus on the well known point that when "utility" is used to mean a numerical representation of preference, as is typical in economic models, such utility does not coincide with what philosophers have meant by utility. There is also an efficiency argument for QV, which does not rest on a conflation of different senses of utility. We argue that utilitarian considerations shed doubt on this efficiency argument as well, that such efficiency arguments are especially problematic when applied to controlling political institutions, and that QV specifically is not a good mechanism for settling distributive questions.

The outline of the remainder of Sect. 2 is as follows. Section 2.1 provides formal preliminaries. Section 2.2 presents quadratic voting formally. Sections 2.3 makes our basic utilitarian argument, and Sect. 2.4 addresses the efficiency argument.

2.1 Preliminaries

2.1.1 Environment and voter preferences

Consider a public decision $x \in \{0, 1\}$, where 0 represents decision 0 and 1 represents decision 1. Decision 1 may be the decision to undertake a public project, and decision 0 may be the decision not to do so.

Let $N = \{1, \ldots, n\}$ be the collection of agents. Each agent i has a wealth endowment $w_i^e \in \mathbb{R}_+$, representing i's initial wealth. An outcome for i is a pair (x, w_i), where x is the public decision and w_i is i's *final* wealth holdings. Final wealth w_i can differ from w_i^e because i has made or received a payment. For each agent i, there is a quantity of money $\hat{u}_i$ such that i is indifferent between $(0, w_i^e)$ and $(1, w_i^e - \hat{u}_i)$. Assuming that i likes money, $\hat{u}_i > 0$ if i prefers decision 1 and $\hat{u}_i < 0$ if i prefers decision 0. So $|\hat{u}_i|$ is i's willingness to pay for decision 1 to replace decision 0 if i prefers 1, and is the minimum payment i is willing to accept for 1 to replace 0 if i prefers 0. In determining $\hat{u}_i$, voters take into account all consequences of the decisions. For example, if decision 1 but not decision 0 requires costs that must ultimately be paid by citizens, each voter will factor her portion of these costs into her valuation.

Let us assume that i's preferences over outcomes (x, w_i) are represented by the quasilinear utility function

$$\hat{U}_i(x, w_i) = \hat{u}_i x + w_i \tag{1}$$

If $x = 0$, then $\hat{u}_i x = 0$ and if $x = 1$, then $\hat{u}_i x = \hat{u}_i$. So we can think of i's utility for decision 0 as being normalized to zero, so that i's utility for decision 1 is $\hat{u}_i$. In addition, let us assume that $\hat{U}_i$ is a von Neumann–Morgenstern utility function so that preferences over lotteries over outcomes can be represented by expectations of $\hat{U}_i$. This implies that the agent is risk neutral.

Let us think of both quasilinearity and risk neutrality as only *approximations*. That is, for the range of payments that might be made in the voting mechanism, preferences are very close to quasilinear and the agent is approximately risk neutral, but for larger payments, deviations from these assumptions may become apparent.[1]

Observe finally that we have not made the assumption that the different utility functions $\hat{U}_i$ can be used for ethically meaningful interpersonal comparisons.

2.1.2 A model of costly voting

This section presents a market model of voting taken from Lalley and Weyl (2015), which is an earlier draft of Lalley and Weyl (2016). Some of the ideas in Lalley and Weyl's model derive originally from Hylland and Zeckhauser (1980). The model provides the basis for our formal analysis of QV.

Let $v \in \mathbb{R}$ be a quantity of votes. If v is positive, then v represents $|v|$ votes for decision 1 and if v is negative, then v represents $|v|$ votes for decision 0. Let $c : \mathbb{R} \to \mathbb{R}_+$ be a **costly voting rule** that maps votes into dollars. Assume that c is even, which means that $c(v) = c(-v)$. So we can think of c as a function of the absolute value $|v|$. Assume that c is differentiable, convex, and strictly increasing in $|v|$, and that $c(0) = 0$. The interpretation is that it requires \$$c(|v|)$ to purchase $|v|$ votes, which can then be cast for either alternative, decision 0 or decision 1. The alternative that receives the most votes wins.

Election proceeds are refunded (approximately) equally to all citizens. Let v_j be the votes purchased by voter j. Then voter i is refunded $\frac{\sum_{j \in N \setminus i} c(v_j)}{n-1}$ dollars, which is the average payment made by voters other than i. Voter i's refund is independent of how i votes, and even whether i votes at all, so as not to affect i's incentive to vote.

Following Lalley and Weyl (2015), we model voting as a market in which voters purchase influence. This is the **price-taking model**. A **collective decision problem** is a tuple $\{N, S, \bar{u}\}$, where N is the collection of agents, S, a positive real number, is the supply of influence, and $\bar{u} = (\hat{u}_i : i \in N) \in \mathbb{R}^N$ is the profile of utilities for decision 1 (see Sect. 2.1.1). A **price-taking equilibrium** of a collective decision problem under voting

[1] The range of payments i will potentially make in the voting mechanism is likely to include only payments substantially smaller than $\hat{u}_i$. For simplicity, it is natural to assume that $\hat{U}_i$ is approximately quasilinear for payments as large as $\hat{u}_i$. This would validate our assumption that i is indifferent between $(0, w_i^e)$ and $(1, w_i^e - \hat{u}_i)$. However, we do not want to assume quasilinearity for arbitrarily large payments.

rule c is an **influence vector** $\bar{I}^* = (I_i^* : i \in N) \in \mathbb{R}^N$, a **price** $p^* \in \mathbb{R}_{++}$ and a decision $x^* \in \{0, 1\}$ such that

$$\textit{Price-taking}: \quad I_i^* \text{ maximizes } \hat{u}_i I_i - c(p^* I_i) \quad \text{over all } I_i \in \mathbb{R}, \tag{2}$$

$$\textit{Market clearing}: \quad \sum_{i \in N} \left| I_i^* \right| = S, \tag{3}$$

$$\textit{Vote majority}: \quad x^* = 1 \Leftrightarrow \sum_{i \in N} p^* I_i^* \geq 0. \tag{4}$$

The **equilibrium vote profile** is the vector $\bar{v}^* = (v_i^* : i \in N) \in \mathbb{R}^n$ whose i-component is $v_i^* = p^* I_i^*$.

We now interpret this model. While the price-taking model does not explicitly represent any randomness in valuations, imagine, for the moment, a **Bayesian game** in which the utility profile $\bar{u}$ is drawn randomly, such that voter valuations are mutually statistically independent.[2,3] Each voter is informed only of her own utility $\hat{u}_i$, and then purchases votes. Let Q_i^* be the probability that decision 1 would win the election if i did not purchase any votes.

By buying votes, voter i purchases influence. Let I_i be the additional probability that i's favorite alternative wins given the number of votes that i purchases. Given the influence that i obtains by purchasing votes, the probability that the outcome is decision 1 is $Q_i^* + I_i$. $I_i > 0$ if i votes for decision 1, and $I_i < 0$ if i votes for decision 0. Acquiring I_i units of influence requires $v_i = p^* I_i$ votes. Thus, p^* is the "price of influence". Equivalently, we can think of $\frac{1}{p^*}$ as a voter's *marginal pivotality*, that is, the additional probability of being pivotal she purchases with an additional vote. For what follows, it is not essential that the is conversion from votes to pivotality probabilities be linear; it matters only that each voter perceives (approximately) the same marginal pivotality.

Given the assumptions made on voter utility in Sect. 2.1.1—recall in particular that the utility of decision 0 is zero—the voter's problem is

$$\max_{I_i} \left[\underbrace{\hat{u}_i \times (I_i + Q_i^*)}_{\text{expected value of decision}} - \underbrace{c(p^* I_i)}_{\text{vote cost}} \right]. \tag{5}$$

Aligning this with the voter's utility function in (1), if voter i purchases influence I_i, i's final wealth w_i is $w_i^e - c(p^* I_i) + \frac{\sum_{j \in N \setminus i} c(v_j)}{n-1}$, but since $w_i^e + \frac{\sum_{j \in N \setminus i} c(v_j)}{n-1}$ enters the objective function additively and is independent of the voter's choice of influence I_i, this term can be omitted from (5). Moreover, (5) differs from the price-taking condition (2) only in that (5) contains a term $\hat{u}_i Q_i^*$. However, $\hat{u}_i Q_i^*$ is a constant that enters the objective additively and so does not affect the optimal choice of influence. So, we can eliminate $\hat{u}_i Q_i^*$, and arrive at (2).

The market-clearing condition (3) posits a supply of influence S, which is necessary to give influence a price p^*. The supply S is just a formal device for setting this equilibrium price. By varying S, while holding $\{N, \bar{u}\}$ fixed, p^* can be made to equal any positive real number. So S can be chosen to make $\frac{1}{p^*}$ equal to the (commonly perceived) marginal

[2] Independence is the most favorable assumption for QV to achieve its efficiency benefits.

[3] Of course, this does not mean that valuations are independent of wealth.

pivotality in the Bayesian game. Then the market-clearing condition (2) would lead to the same purchase of votes $\bar{v}^*$ as would occur conditional on realization $\bar{u}$ in the game. Finally, (4) enforces the condition that the outcome that receives the most votes wins.

In sum, we can think of the price taking equilibrium as providing a snapshot of an ex post realization in a Bayesian voting game. However, the relation between the price taking equilibrium and the Bayesian voting game is only heuristic. For a rigorous analysis of a Bayesian game model of QV, see Lalley and Weyl (2016). The price-taking model is much simpler to analyze than the Bayesian game model. That is why we build on the price-taking model in this paper.

Observe finally that the profile of wealth endowments $\bar{w}^e = (w_i^e : i \in N)$ is not part of the price-taking model. This is because, once we know the utility profile $\bar{u}$, we no longer need to know $\bar{w}^e$ to calculate the equilibrium. This does not mean, however, that the utilities are independent of wealth. Given that utility is only *locally* quasilinear (see Sect. 2.1.1), a large change in initial wealth w_i^e may change willingness to pay $\hat{u}_i$.

2.2 Quadratic voting

QV is the costly voting rule $c(v) = kv^2$ for some $k > 0$. For simplicity, assume that $k = 1$, so that $c(v) = v^2$.

2.2.1 Why QV is advocated

Under QV, the voter's optimization problem—the problem in the price-taking condition (2)—becomes

$$\max_{I_i} \quad \hat{u}_i I_i - (p^* I_i)^2.$$

The first-order condition for an optimum is $\hat{u}_i = 2(p^*)^2 I_i^*$. Equivalently, i's optimal influence is $I_i^* = \frac{\hat{u}_i}{2(p^*)^2}$. Converting influence into votes (using $v_i^* = p^* I_i^*$), i purchases

$$v_i^* = \frac{\hat{u}_i}{2p^*} \tag{6}$$

votes. Equation (6) shows that voters purchase votes in proportion to their value $\hat{u}_i$ for decision 1. So, using (6) and (4), under QV,

$$x^* = 1 \Leftrightarrow \sum_{i=1}^{n} \hat{u}_i \geq 0. \tag{7}$$

(7) says that decision 1 wins if and only if aggregate willingness to pay for decision 1 to replace decision 0 is positive. Lalley and Weyl (2015) show that each costly voting rule c and collective decision problem $\{N, S, \bar{u}\}$ determine a unique price-taking equilibrium. They refer to c as **robustly efficient** if for all collective decision problems, this unique equilibrium satisfies (7). The above argument shows that QV is robustly efficient. QV is the *unique* rule with this property.

Proposition 1 (Lalley and Weyl 2015) *A costly voting rule c is robustly efficient if and only if $c(v) = kv^2$* for some $k > 0$.

2.2.2 *The argument for QV over majority voting*

In majority voting, each voter is entitled to one vote and the alternative that receives the most votes wins.

Suppose there are three voters: Ann, Bob, and Carol. Ann and Bob each prefer decision 0, and each would be willing to pay \$1 to cause decision 0 to replace decision 1. Carol prefers decision 1 and willing to pay \$3 for her preference. Under majority voting, there are two votes for decision 0, and one vote for decision 1. So decision 0 wins. Yet $\hat{u}_{\text{Ann}} + \hat{u}_{\text{Bob}} + \hat{u}_{\text{Carol}} = -1 - 1 + 3 = 1 > 0$. So decision 0 wins while the aggregate willingness to pay for decision 1 to replace decision 0 is positive. As we have seen above, under QV, decision 1 wins.

This observation is the basis for the argument that QV is superior to majority voting. But what is the argument precisely? We can think of two candidates.

1. From a utilitarian point of view, decision 1 (the outcome of QV) is better than decision 0 (the outcome of majority voting) because the sum of utilities for decision 1 is greater.
2. If we enact decision 0, then everyone could be made better off: we could instead enact decision 1, and Carol could make a payment of $\$\frac{4}{3}$ to Ann and a payment of $\$\frac{4}{3}$ to Bob. The utility of each of the three agents would rise by $\frac{1}{3}$. So we should not enact decision 1. Majority voting, unlike QV, enacts an outcome against which there is a decisive argument.

We consider the first argument in Sect. 2.3, and the second in Sect. 2.4. The second argument suggests that we consider voting mechanisms not in isolation, but rather as bundled with compensating transfers. Section 2.3 analyzes voting mechanisms as stand-alone mechanisms, and Sect. 2.4 considers mechanisms bundled with compensating transfers.

2.3 A utilitarian analysis

2.3.1 *Utilitarianism*

Utilitarianism is a major ethical theory with a long tradition in philosophy. It says that the best outcome is the one that maximizes the sum of agents' utilities. Utility has been interpreted in different ways. It might represent the experiences of pleasure and pain or it might accommodate broader aspects of wellbeing, desire, preference, or a person's good.

We will not take a stand on philosophical debates on what precisely utility is. However, there is one thing that ethical utility cannot be. Ethical utility cannot merely be a representation of preference, which was the meaning given to the utility functions $\hat{U}_i$ in Sect. 2.1.1. Indeed, if $\hat{U}_i$ represents i's preferences over outcomes (or over lotteries), then for any positive real number α_i, so does $\alpha_i \hat{U}_i$, but how the sum $\sum_i \alpha_i \hat{U}_i$ ranks outcomes depends on the choice of the numbers α_i.[4]

[4] If the utility functions $\hat{U}_i$ are quasilinear—that is, of the form $\hat{U}_i(x, w) = \hat{u}_i x + \beta_i w_i$—and if in addition $\beta_i = 1$ for all i, then one might think that the only ethically reasonable way to calibrate utilities would be to set $\alpha_i = \alpha$ for some α and all i, since if $\alpha_i > \alpha_j$, then the utilitarian would want to make a boundless transfer from j to i. However, by the same token, it is not reasonable to believe to begin with that ethical utility—or decision utility for that matter—is quasilinear for arbitrarily large wealth transfers. Indeed, in Sect. 2.1.1, we only assumed that the $\hat{U}_i$ are (approximately) quasilinear when wealth transfers are not too large.

Ethical utility must be *interpersonally comparable*, and utilities must be calibrated in an ethically reasonable way. For example, we cannot calibrate utilities in such a way that it is more important that Ann not suffer a minor itch than that Bob not starve. Harsanyi (1953) proposed a formal model in which the ethical sum $\sum_i U_i$ represents the preferences of an impartial observer who might turn out to be any citizen with equal probability.[5] This is one way that economists have thought about ethical utility.

We make an ethical assumption on utility that is widespread in the utilitarian tradition, namely, the assumption that the marginal utility of a dollar is decreasing in wealth. Under this assumption, if we had one dollar, and could either give it to a rich person or a poor person, then, other things being equal, it would be ethically better to give it to the poor person. Intuitively, the dollar makes more of a difference to the poor person.

Henceforth, we use the term *utility* to refer to ethical utility. In contrast, we refer to $\hat{U}_i$ from Sect. 2.1.1 as the *value function* and to $\hat{u}_i$ as *willingness to pay*. We assume that ethical utility is given by

$$U_i(x, w_i) = u_i x + g(w_i). \tag{8}$$

As above, $x \in \{0, 1\}$ is the public decision and w_i is final wealth. Here u_i, which can be positive or negative, is i's utility of the public decision 1 and the utility of public decision 0 is zero. The utility of wealth function $g(w_i)$ measures the utility an agent receives from having wealth w_i. For simplicity, g is the same for all agents i. We assume that g is concave to capture the diminishing marginal utility of wealth. For simplicity, we assume that u_i is not a function of wealth w_i.

We rename the set of all agents N°—the reason will become apparent below. The utilitarian objective is to maximize the sum of utilities:

$$\sum_{i \in N^\circ} U_i(x, w_i). \tag{9}$$

We introduce utilitarianism as a way of ethically evaluating QV. We do not wish to bundle this ethical criterion with a change in the assumption about how agents behave. For this reason, we assume that g is piecewise linear.[6] More precisely, we assume that there exists a positive number M and for all i, there exists a positive number b_i such that

$$\forall t \in \mathbb{R}, \quad |t| < M \Rightarrow g\left(w_i^e - t\right) = g\left(w_i^e\right) - b_i t. \tag{10}$$

Recall that w_i^e is i's wealth endowment. Figure 1 displays a utility of wealth function g satisfying our assumptions.

The number $b_i = g'(w_i^e)$ represents i's marginal utility of wealth at w_i^e. We assume that M is large enough that no transfers that occur through the voting mechanism exceed M. We also assume that U_i not only measures ethical utility, but also that U_i is a von Neumann–Morgenstern utility function that represents i's preferences. Under these assumptions, over the relevant range of transfers, each agent both has quasilinear utility and is risk neutral.

Note finally that while above we assumed for simplicity that u_i does not depend on wealth, nothing essential in our analysis below would change if we were to write $u_i(w_i)$ instead of u_i in (8) for the utility of the public decision and assume that $u_i(w_i)$ is constant over the range of transfers that actually occur in the voting mechanism. Then u_i could

[5] Vickrey (1945) was a precursor to this analysis.

[6] If, instead, we assumed that g is strictly concave, but the slope of g changes sufficiently slowly, then nothing of substance would change.

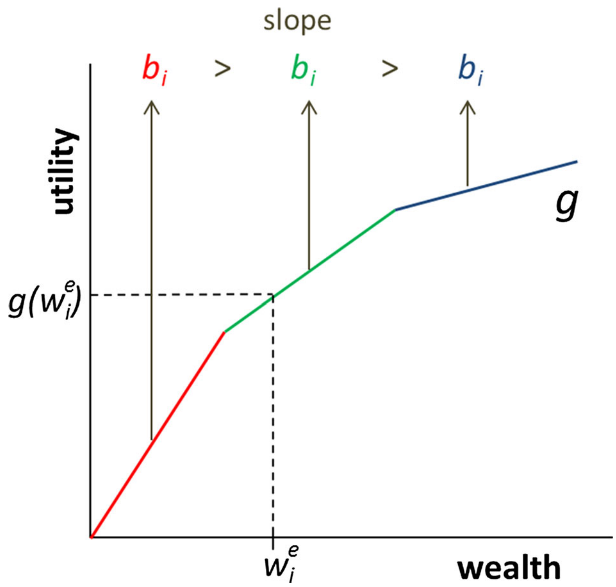

Fig. 1 The utility of wealth function g. The function g is piecewise linear and concave. w_i^e is the wealth endowment. b_i is the marginal utility of wealth and corresponds to the slope of the function. The *color* of the label b_i matches the *color* of the linear segment whose slope it represents. Since g is concave, b_i diminishes as wealth increases

simply be understood as meaning $u_i(w_i)$ everywhere below. This may be desirable because, when utility is understood in an ethically significant sense, for various reasons the utility consequences for an agent of a given public decision may depend on her wealth.[7]

2.3.2 Neutralizing the refund

Now that utility is understood in an ethical sense, we can see that QV's refund is ethically significant because it amounts to a transfer among voters with potentially different marginal utilities of wealth. But is the refund's effect on total utility positive or negative? One can easily see that the refund is ethically ambiguous. Imagine that voters are split into two groups A and B. Voters in A care passionately about the election and voters in B are relatively indifferent. So voters in A purchase a large number of votes while voters in B purchase a small number of votes. As proceeds are refunded equally to all voters independently of voting behavior, the refund then amounts to a net transfer from A to B. B could be large or small in comparison to A and voters in B could be wealthy or poor in comparison to voters in A. Depending on how these parameters are resolved, the transfer

[7] For example, suppose that the question is whether to create a new public park. Suppose that the enjoyment that voter i would get from the park h_i does not vary with i's wealth, but each citizen would have to make a payment of T to finance the park independently of her wealth. Then i's net utility from the park is $h_i - b_i T$, where $b_i = g'(w_i^e)$. Since the marginal utility of wealth b_i does not change over the range of transfers that might occur in the voting mechanism, this example satisfies our assumptions with $u_i(w_i) = h_i - g'(w_i)T$. As another example, the public decision may involve building a facility that the agent may only want to use if she is poor.

could be from rich to poor or from poor to rich and its effect may be larger or smaller. So, the refund could have a positive or negative effect on aggregate utility.[8]

To neutralize the effect of the refund, imagine that we live in a society in which redistribution via the refund is prohibited. Call this the **no redistribution constraint**. Imagine that society is split into two groups: *citizens* and *non-citizens*. Only citizens may vote. However, non-citizens do not care about the public decision, so that for each non-citizen i, $u_i = 0$. Non-citizens do value wealth. Every dollar raised through the election from citizen i is transferred to some non-citizen j with the same marginal utility of wealth as i.[9] Citizens receive no refund. The utilitarian cares equally about the utility of citizens and non-citizens. So utility is neither lost nor gained in the transfer from citizens to non-citizens. N°, the set of agents whose utilities the utilitarian sums in (9), now refers to the set of all agents—citizens and non-citizens. Henceforth N will refer to the set of citizens only.

Under the no redistribution constraint, elections have no effect on the wealth component of utility g and non-citizens do not care about the public decision. So the utilitarian objective (9) simplifies to

$$\sum_{i \in N} u_i x, \tag{11}$$

and given the profile of utilities $\bar{u} = (u_i : i \in N)$, if two voting mechanisms select different public decisions, the utilitarian prefers the mechanism that selects decision 1 if and only if $\sum_{i \in N} u_i \geq 0$. Structurally, this looks very similar to the efficiency criterion (7).

The motivation for the no redistribution constraint is twofold: (i) it makes utilitarian evaluation as parallel as possible to the efficiency evaluation of Lalley and Weyl (2015), in which transfers are irrelevant, and, more importantly, (ii) it focuses the utilitarian evaluation on the substantive evaluation of merits of the public decision to be made by the election, rather than on incidental features of the rebate, which is merely a byproduct of the election.

2.3.3 A utilitarian analysis of quadratic voting

Expressing the voter's optimization problem in terms of the ethically calibrated utility representation (8) rather than the "dollar denominated" utility representation (1), (5) becomes

$$\max_{I_i} \left[\underbrace{u_i \times (I_i + Q_i^*)}_{\text{expected utility of decision}} + \underbrace{g\left(w_i^e - c(p^* I_i)\right)}_{\text{wealth utility given vote cost}} \right]. \tag{12}$$

[8] In large elections, the election proceeds should be small relative to aggregate willingness to pay. This follows from results of Lalley and Weyl (2016) and Weyl (2017). Taking this into account would itself limit the effect of the refund, and so makes it less important to neutralize its effect, as we do below, although it is still useful to do so, for analytical precision.

[9] Positing non-citizens simplifies some mathematical expressions below, but nothing of substance would change if we assumed that every dollar raised from citizen i is transferred to some other *citizen* j with the same marginal utility of wealth as i.

Using (10), (12) becomes

$$\max_{I_i}\left[u_i \times (I_i + Q_i^*) + g\left(w_i^e\right) - b_i c(p^* I_i)\right] \tag{13}$$

Eliminating constants $u_i Q_i^*$ and $g\left(w_i^e\right)$, which do not affect the optimization, and dividing by b_i yields an analog of the price taking condition (2) in which basic ethically significant primitives, u_i and b_i, replace willingness to pay $\hat{u}_i$:

$$\textit{Price-taking}: \quad I_i^* \text{ maximizes } \frac{u_i}{b_i} I_i - c(p^* I_i) \quad \text{over all } I_i \in \mathbb{R}. \tag{14}$$

Here $\frac{u_i}{b_i}$ takes the place of $\hat{u}_i$. Indeed, expressing utility in terms of ethical primitives, $\frac{u_i}{b_i}$ is i's marginal willingness to pay for a higher probability of decision 1, and if M in (10) is large enough, i is indifferent between outcomes $(0, w_i^e)$ and $(1, w_i^e - \frac{u_i}{b_i})$,[10] so that $\frac{u_i}{b_i}$ is i's willingness to pay for decision 1 to replace decision 0 as well. Relating the two utility representations of i's preferences, (1) and (8), we have

$$\hat{u}_i = \frac{u_i}{b_i}. \tag{15}$$

Define an **ethically specified collective decision problem** to be a tuple $\{N, S, \bar{u}^\circ, \bar{b}\}$, where $\bar{u}^\circ = (u_i : i \in N)$ is a profile of public decision utilities and $\bar{b} = (b_i : i \in N)$ is a profile of marginal wealth utilities. Parallel to Sect. 2.1.2, an *ethically specified* price taking equilibrium consists of an influence vector $\bar{I}^*$, a price p^*, and a decision x^* that satisfy (14), (3), and (4). Mathematically, this is equivalent to the equilibrium definition of Sect. 2.1.2, but the utility primitives are now expressed in ethically significant units.

For QV (with constant $k = 1$), using the individual vote demands (6) and (15), we have

$$v_i^* = \frac{1}{2p^*}\frac{u_i}{b_i}. \tag{16}$$

So when utility is ethically specified, in QV, different voters i do *not* purchase votes in proportion to their utility u_i, but rather in proportion to willingness to pay $\frac{u_i}{b_i}$. Since the wealthy have a lower marginal utility of wealth b_i, the wealthy buy votes disproportionately to their utility. QV fails to be optimal, for the utilitarian, because it is biased toward the wealthy.

Say that a costly voting rule c is **robustly utilitarian** if for all ethically specified collective decision problems, the (unique) price-taking equilibrium under rule c satisfies:

$$x^* = 1 \Leftrightarrow \sum_{i=1}^{n} u_i \geq 0. \tag{17}$$

Voting rule c is robustly utilitarian if it always selects the best outcome from a utilitarian perspective. Contrast the following result with Proposition 1.

Proposition 2 *There does not exist a costly voting rule that is robustly utilitarian.*

Proposition 1 implies that, restricting attention to collective decision problems $\{N, S, \bar{u}^\circ, \bar{b}\}$ in which all agents have the same marginal utility of wealth—$b_i = b, \forall i$—QV

[10] In particular, $g(w_i^e) = u_i + g(w_i^e - \hat{u}_i) \Leftrightarrow g(w_i^e) = u_i + g(w_i^e) - b_i\hat{u}_i \Leftrightarrow \hat{u}_i = \frac{u_i}{b_i}$.

is the unique voting rule that always satisfies (17). So if any voting rule is robustly utilitarian, it must be QV. But since, as we have just seen, when there exist i and j such that $b_i \neq b_j$, QV does not lead to votes proportional to utilities, QV sometimes violates (17).[11] So QV is not robustly utilitarian.

2.3.4 Wealth-weighted quadratic voting

The preceding analysis suggests that to achieve the ideal of QV relative to ethically specified utility rather than willingness to pay, voting rules must be made conditional on wealth. Define a **wealth sensitive voting rule** to be a function $c : \mathbb{R} \times \mathbb{R}_{++} \rightarrow \mathbb{R}_+$ that maps votes purchased v and the marginal utility of wealth b into a vote cost $c(v, b)$. We assume that for all b, $c(\cdot, b) : \mathbb{R} \rightarrow \mathbb{R}_+$ is even, differentiable, convex, and strictly increasing in $|v|$, and that $c(0, b) = 0$. If the utility of wealth function g is known, b_i can be derived from w_i^e (specifically, $b_i = g'(w_i^e)$), so that c can be thought of as a function of wealth rather than of the marginal utility of wealth. In practice, the formula for determining the cost of purchasing votes would take voter wealth into account. Under this generalization, a price taking equilibrium for an ethically specified collective decision problem $\{N, S, \bar{u}^\circ, \bar{b}\}$ is a tuple $\{\bar{I}^*, p^*, x^*\}$ satisfying

$$\textit{Price-taking}: \quad I_i^* \text{ maximizes } \frac{u_i}{b_i} I_i - c(p^* I_i, b_i) \quad \text{over all } \mathrm{I_i} \in \mathbb{R}. \tag{18}$$

as well as (3) and (4). To restore the attractive properties of QV once utility is ethically specified, we must "undo" the effect of dividing by b_i. This motivates the following definition: **wealth-weighted QV** is a rule of the form $c(v, b) = \frac{k}{b} v^2$ for some constant $k > 0$. Since poorer voters have a higher marginal utility of wealth, wealth-weighted QV gives poorer voters a discount on votes. Suppose, for example, that $g(w) = \log(w)$,[12] so that marginal utility is inversely proportional to wealth, $g'(w_i^e) = \frac{1}{w_i^e}$. Observe that in this case wealth-weighted QV yields $c(v, \frac{1}{w_i^e}) = w_i^e k v^2$, so that if $k' = k \times 100$, buying one vote would cost $(k' \times 1)\%$ of your wealth, buying two votes would cost $(k' \times 4)\%$ of your wealth, and so on.

The analysis of the preceding sections implies that:

Proposition 3 *A wealth sensitive voting rule is robustly utilitarian if and only if it is an instance of wealth-weighted QV.*

This might seem to make wealth-weighted QV attractive. Indeed, if we take our model as a literal and complete representation of reality, then wealth-weighted QV seems to be "the" utilitarian solution.

Yet, taking a broader perspective, the advantages of wealth-weighted QV are balanced by significant disadvantages. An attractive feature of ordinary—as opposed to wealth-weighted—QV is that it is "detail free": it requires no knowledge about the agent's utility function. In contrast, wealth-weighted QV requires knowing a voter's marginal utility of wealth. We have assumed for simplicity that all voters have the same marginal utility of wealth function g. This assumption, as we have made it, has both empirical and normative

[11] To see this, assume wlog that $b_i < b_j$, consider the case where $u_i < 0 < u_j$, $|u_j| > |u_i|$, $|\frac{u_i}{b_i}| > |\frac{u_j}{b_j}|$, and $u_\ell = 0, \forall \ell \in N \backslash \{i, j\}$.

[12] This departs from the assumption above that g is piecewise linear; to be consistent with this, we may assume that g is a piecewise linear approximation to log.

content.[13] If different voters have different marginal utility of wealth functions g_i, the marginal utility of wealth will not be a function of wealth alone, and so the optimal voting rule would require more than just information about wealth. To solve this, one might attempt to posit a single representative utility of wealth function g that one might hope to provide a reasonable approximation. The more fundamental issue is that the question of what is the marginal utility of wealth is not a purely empirical question; it is largely a normative question. Some of the simplifying assumptions we made above might suggest that one could learn the marginal utility of wealth from the agent's degree of risk aversion over large gambles, but that conclusion is just an artifact of our model, produced by our simplifying assumptions.[14] At a more basic level, determining the marginal utility of wealth requires normative value judgments about the ethical tradeoffs involved in providing resources to people at different wealth levels. Of course, advocacy of any policy requires similar value judgments, and it is much better not to pretend otherwise. However, the basic voting rules that underlie our political system need to be durable and robust. That is why it is desirable for those voting rules to be detail free, and not to build in very specific, and, hence, inevitably speculative value judgments into those rules.

Second, to begin with, we should expect adoption of even ordinary QV to be politically problematic. Putting a price on votes is unlikely to be perceived by citizens as just another tax, such as, e.g., an income tax. Whether justified or not—we discuss the merit of related charges in Sect. 3—at least some people are likely to perceive QV as in some ways akin to a partial form of disenfranchisement. It is difficult to predict counterfactual sentiment, but making QV wealth-weighted would add a layer of complexity to what may already be a fraught issue. One might think that giving poor people a discount, as wealth-weighted QV would do, would mitigate the perception of disenfranchisement. On the other hand, creating any kind "wealth test" for voting, whether it favors rich or poor, may be offensive to many.

2.3.5 *Quadratic voting versus majority voting*

Putting wealth weighted QV aside, we now return to our central concern: we compare (ordinary) QV and majority voting. We simplify the problem and consider two extreme cases that we take to epitomize the cases in which QV would perform well and in which it would perform poorly.

We split the citizens into two classes, rich R and poor P. All rich citizens have wealth endowment w_R^e and all poor citizens have wealth endowment w_P^e, where $w_P^e < w_R^e$. Let α be

[13] See footnote 14.

[14] In Sect. 2.3, we assumed, *first*, that U_i measures ethical utility. Second, we assumed that U_i also represents i's von Neumann–Morgenstern preferences over lotteries over outcomes. The second assumption (combined with the assumption of a common utility of wealth function g) is what might suggest that risk preferences determine ethical utility. We made the second assumption primarily to make the utilitarian analysis of Sect. 2.3.1 as completely parallel to the preceding analysis of QV in Sects. 2.1–2.2. However, the claim that there should exist a common utility function that simultaneously satisfies the first and second assumptions above, a claim that is integral to Harsanyi's version of utilitarianism, is a philosophically contentious one. Even if we were to grant that U_i can play both roles, in reality, different agents will have different risk preferences. If $\hat{g}_i$ and $\hat{g}_j$ represent i and j's preferences over wealth gambles, then so do $\alpha_i\hat{g}_i$ and $\alpha_j\hat{g}_j$ for any positive numbers α_i and α_j, and it will require an ethical judgment to calibrate the utilities—that is, to determine the ethically correct ratio α_i/α_j—before these utilities can be added up to generate an ethically significant quantity in the way required by utilitarianism. These considerations show that the notion that we can empirically infer the marginal utility of wealth from observation without making ethical value judgments is mistaken.

the proportion of rich citizens in the population. Assume that $\alpha \in (0, \frac{1}{2})$. We evaluate QV and majority voting conditional on a realization $\bar{u}^\circ = (u_i : i \in N)$ of voter utilities. We consider two cases:

1. *Issues independent of wealth* The distribution of utilities for the public decision is independent of wealth. Formally, for every utility u_i, ℓ poor citizens have utility u_i for decision 1 if and only if $\frac{\alpha}{1-\alpha}\ell$ rich citizens have utility u_i for decision 1.
2. *Issues polarized by wealth* There are two utility levels u_P and u_R such that all poor citizens have utility u_P for decision 1 and all rich citizens have utility u_R for decision 1, where $u_R < 0 < u_P$.

Proposition 4 *If issues are independent of wealth, then QV is the utilitarian optimal voting rule.*

The argument is straightforward: in this case, within each wealth class, the marginal utility of wealth is constant and, hence, within the class, utility is proportional to willingness to pay. So under QV, among the votes from a given class, the utilitarian best alternative for that class wins. But both wealth classes have the same distribution of utilities. So the alternative that maximizes utility overall wins. So when valuations are independent of wealth, not only is QV better than majority voting, but it is optimal among all voting rules.

Proposition 5 *Assume that issues are polarized by wealth and let $u_R = -1$. There exist thresholds t_0 and t_1* with $0 < t_0 < t_1$ *such that for all $u_P \in (0, t_0)$, QV is strictly utilitarian preferred to majority voting, for all $u_P \in (t_0, t_1)$, majority voting is strictly utilitarian preferred to QV, and for $u_P \in (t_1, +\infty)$, majority voting and QV choose the same outcome, so they are equally good.*

Remark 1

- $(0, t_0)$ is the **tyranny of the majority region** for majority voting, where a majority that does not feel very strongly gets its way over a minority that feels more strongly. This is the region where counting votes fails to be a good guide to social utility.
- (t_0, t_1) is the **corrupting influence of money region** for QV, where a wealthy minority gets its way at the expense of the less wealthy majority who cares more. This is the region where willingness to pay fails to be a good guide to social utility.
- $(t_1, +\infty)$ is the **doesn't matter how you count region** where counting votes and adding willingnesses to pay give the same answer.

The proof of Proposition 5 is in the Appendix. Intuitively, in region $(0, t_0)$, the poor win under majority voting, despite not caring much about the decision, because there are more of them. In region (t_0, t_1), the poor care more in aggregate than the rich about the public decision, but not enough to win under QV given their wealth disadvantage. In region $(t_1, +\infty)$, the poor care so much that they win regardless of the voting method.

We have assumed that the refund is handled as in Sect. 2.3.2 so that the utilitarian objective is (11). However, if election proceeds are only a small fraction of aggregate willingness to pay for public decisions, as one would expect in large elections,[15] then vote payments and refunds will have only a small effect on utility, and Proposition 5 will be robust to the precise assumption one makes about the refund rule.

[15] This follows from the analysis of Lalley and Weyl (2016) and Weyl (2017).

2.4 QV, compensation, and distribution

We have seen that the utilitarian will sometimes prefer the public decision selected by majority voting to that selected by QV. Still, one might think that the efficiency of QV would imply that when QV makes the inferior decision, the losers can be compensated, making everyone better off. Alternatively, one might claim that if we start off from an optimal (or at least a sufficiently good) position with regard to distribution subject to incentive constraints and other constraints, then efficiency objectives and utilitarian objectives will tend to align. We now consider these arguments.

2.4.1 The distributive inefficacy of QV

Let us start by considering the question of whether QV would be favorable to compensation for losses induced by policies. So let us ask, if QV is the mechanism for making public decisions, how would it be decided that the losers—those who are harmed by a public decision—are to be compensated? Could we hold an election using QV to decide whether to compensate the losers? As Posner and Weyl (2015) point out in arguing that QV would not lead to unjust expropriation, a dollar is worth a dollar to everyone. So if the ballot question were “should we compensate the losers in the amount of \$$y$?”, the aggregate willingness to pay for compensation by the losers would be \$$y$, and the willingness to pay by the winners *not* to compensate would also be \$$y$. So QV would be at a deadlock. So we should not be confident that, if QV is used to decide, transfers compensating the losers for any given public decision would actually be implemented.

Going beyond compensation for losses, a similar argument suggests that QV would be powerless with respect to purely distributive questions, and would be suspect for questions with a large distributive component. Consider a question with both a distributive and efficiency component: suppose that decision 1 is more efficient than decision 0, but also favors the rich at the expense of the poor. We might think of decision 1 as effectively a policy that bundles an efficiency gain with a transfer of \$$y$ from poor to rich. QV detects the efficiency gain, and supports it, but, as in the purely distributive case, it ignores the \$$y$ transfer. In other words, it takes only part of what is at stake into account. This suggests that QV is not a general purpose mechanism; for those important public decisions that have a significant distributive component, QV would have to be supplemented by some other decision mechanism.

Posner and Weyl (2015) present an argument aimed to address distributive concerns. They suggest that people often do not know what position in society they will ultimately occupy and redistributive policy is like a public good that helps insure everyone against bad outcomes. At one extreme, they imagine that all individuals face the same uncertain prospects, not knowing their future social position or wealth. At this extreme, we are effectively behind a veil of ignorance, where all individuals share a common view, and, under QV, distributive concerns would be handled by the common desire to insure against a bad social realization. Of course, we are not at the extreme case. Posner and Weyl (2015) write, “In reality, some investments are sunk, and some uncertainty is realized, but likely about roughly equal amounts of each, at least when averaged over the population. In this setting, QV would produce the optimal social-insurance plan covering the residual level of uncertainty. Given the balance, it is likely the optimum would resemble that in the first case, where agents choose the social-insurance system behind the veil of ignorance.”

Distributive policy is reduced to efficient provision of social insurance, a public good, and as QV favors efficient decisions, it will handle such decisions well.

Observe that in the extreme case, if we are behind a veil of ignorance with interests completely aligned, then all voting methods, including QV and majority voting, will produce the same result. We do not need a special voting institution when voters are unanimous. If, in the more realistic case, we are far from unanimity on distributive questions, we should have little confidence that the electorate faces essentially the same unanimity inducing problem it would face behind the veil of ignorance. Different groups will face different lotteries over future prospects, and so different segments of the population will be concerned with different risks. For example, risks associated with starting positions at birth will already be realized. Different groups will also have different preferences with regard to distribution in view of their relative positions in society. Posner and Weyl acknowledge such differences but they don't provide a reason for thinking that in conditions under which there is substantial disagreement about desirable distributive policy, aggregate willingness to pay for insurance against residual risks will reflect the preferences of an agent who behind the veil of ignorance might become anyone in today's society with equal probability.

In sum, the argument is completely compelling only in narrow circumstances: it shows that if when people vote on distributive policy, they view the issue as one of insuring themselves against future risks, and to a reasonable approximation all people view themselves as facing the same risks from the same starting positions, QV will support utilitarian optimal policies. But this is not really a victory for QV, because in these circumstances all reasonable voting methods, including majority voting, will work well. In the more realistic case in which different groups face different risks, we have not been given reason to believe that the weight QV will place on different groups is proportional to their interests and, in particular, QV may not put sufficient weight on risks faced by disadvantaged groups. To the extent that distributive policy concerns addressing disadvantage rather than providing insurance, QV will tend to be indifferent.

Contrary to picture painted by Posner and Weyl, we do not think that distributive questions can be reduced to questions of efficiency. Once this is admitted, there is no remaining justification for selecting voting institutions on efficiency grounds alone. The conceptual scheme that separates questions of efficiency from distributive questions, as if they can be undertaken by two different branches of the government, is misplaced when the task is to select controlling political institutions. We cannot say, we will select our basic political institutions as those that are most efficient or that will select the most efficient outcomes, and then designate some other institution to take care of distributive questions. What other institution is this going to be?

2.4.2 Utility versus efficiency

It is important to realize that the efficiency of a public decision is not independent of the wealth distribution. Suppose that group A prefers decision 0 and group B prefers decision 1. Perhaps the decision is a cultural decision, such as the decision of whether to permit gay marriage. Let us suppose that the preferences, as well as the intensity of preference for the decision, would not change if we were to make a wealth transfer from group A to group B. Then it may be that at present, willingness to pay for decision 0 exceeds willingness to pay for decision 1, because voters in A are willing to pay a great deal for decision 0. If there were a large transfer of wealth from voters in A to voters in B, then voters in A may no longer be willing to pay so much for decision 0, and willingness to pay for decision 1 may be in excess of willingness to pay for decision 0. The transfer of wealth may have little impact on the intensity of preference of voters for the issue. However, when voters become poorer, they become less willing to pay for many of the things that they desire. If the issue is gay marriage, the utilitarian merit of the issue may have little to do with whether the voters who support it or those who oppose it happen to be wealthier. We conclude from this that the *substantive merit* of a policy is not, in general, determined by voters' willingness to pay for it.

In our simple model, in which utility is of the form $U_i(x, w_i) = u_i x + g(w_i)$, so that the utility the voter gets from the decision is additively separable from her utility of wealth, decision 1 is the substantively superior decision if and only if $\sum_i u_i > 0$. Conditional on the realization of utilities u_i, whether decision 1 satisfies this criterion is independent of the wealth endowment. On the other hand, QV selects decision 1 if and only if decision 1 is efficient, that is, if and only if $\sum_i \frac{u_i}{g_i'(w_i^e)} > 0$, which does depend on the wealth endowment. Because which public decision is efficient depends on the wealth distribution, we will refer to a public decision as being *locally* efficient at a given distribution.

It is clearly possible that $\sum_i u_i > 0$ and $\sum_i \frac{u_i}{g_i'(w_i^e)} < 0$. So the utilitarian criterion and efficiency may diverge. One doesn't need the utility of the public decision to be additively separable from wealth utility for such a divergence to be possible; far from it. One needs only that the willingness to pay for an increment of utility varies with wealth, which is a very easy condition to satisfy.

It is true that when decision 0 is locally more efficient than decision 1, then *in principle* there is some transfer from citizens who support decision 0 that would make everyone better off than if decision 1 were taken and no transfers were made. By the same token, *in principle*, the utilitarian may prefer to take decision 1, which may be the decision that has the most merit, and couple that with such a large transfer from the supporters of decision 0 to the supporters of decision 1 that decision 1 actually becomes efficient. Or perhaps the best thing to do would be to separate the distributional question from the question of substantive merit, take the substantively better decision, decision 1, and, if a transfer is to be made, select a transfer that is distributively good, rather than a transfer aimed at supporters of one decision or the other.

The real question is not what policy would be best if it were accompanied by complementary transfers. The real question is, which policy would be best given whatever transfers, if any, are actually likely to accompany the policy. What we and the voters should be concerned with is what *will* happen, not what *could* happen. If transfers are likely to accompany a policy or are in fact an explicit part of the policy, if, for example, they are bundled with its financing through the tax system, voters will factor this into their votes, and we may want to view these transfers as part of the policy. If more efficient alternatives were generally accompanied by transfers that made them Pareto improving relative to their less efficient alternatives, then we would generally see unanimity in elections, which is not the case.[16] As the alternatives before the electorate will not be Pareto ranked, there is a strong argument for evaluating voting institutions in terms of their propensity to bring about outcomes that would be substantively best rather than outcomes that are merely locally efficient.[17] We have seen in Sect. 2.3.5 that from a utilitarian point of view QV

[16] For a discussion from a very different perspective, but that echoes the themes raised here, see Buchanan (1959).

[17] With respect to the financing of public decisions, Kaplow (2004) writes, "of the many (consistent) ways that one could adjust the income tax system to achieve budget balance, ideally the intrinsic features of providing a public good or correcting an externality would not become entangled with concerns about the

selects the substantively best policies when issues are independent of wealth, but it may systematically fail to do so when issues are polarized by wealth.

2.4.3 Arguments that we should expect efficiency and the utilitarian criterion to align

We now explore one other kind of argument. This is not a compensation argument, but rather an argument to the effect that, in general, there should be a strong tendency for efficiency and the utilitarian criterion to align. While we are no longer concerned with compensation arguments, let us take as our starting point another common criticism of such arguments having to do with information problems. Concerning lump-sum distributive transfers, Weyl (2015) writes:

> Vickrey (1945), among others, objected that such transfers were typically infeasible as granting them would require information unavailable to the state. Atkinson and Stiglitz (1976) responded by providing an alternative defense of the Kaldor-Hicks criterion. Hylland and Zeckhauser (1979) and Kaplow (2004) elaborated this into a comprehensive foundation for the modern application of the criterion. They argued that in the presence of an optimal income tax that "takes care" of redistribution, the envelope theorem for society implies that a dollar in the hands of any individual is equally valuable. While a dollar may be worth less directly in the hands of a rich individual it encourages individuals to become rich generating tax revenue that eventually benefits the poor.

The most straightforward reading of, e.g., Hylland and Zeckhauser (1979) is as a compensation argument. It says that in the presence of information problems and the incentive constraints to which they give rise, compensation associated with a more efficient policy can be carried out through the tax system in an incentive compatible way. More precisely, under certain assumptions, any policy may be accompanied by a tax adjustment that helps to channel efficiency gains into distributively desirable outcomes. However, that is not the lesson that Weyl draws here. Weyl's argument echoes formal results of Christiansen (1981) and Boadway and Keen (1993). Weyl is claiming that when we are at a utilitarian optimal tax system, then, taking incentive constraints into account, it is equally good from a utilitarian point of view to give a dollar to anyone, and hence a policy should be evaluated by aggregate willingness to pay. On Weyl's view, this is so without making a *further* adjustment to the tax system to accommodate the new policy.[18] In the paragraph following the one we cite, Weyl lists various qualifications that must hold for his argument to be valid.

In personal communication, Weyl has argued that in the presence of a utilitarian optimal tax system, by selecting the more efficient policy, QV will always select the utilitarian preferred policy. Specifically, he argued that, while we are of course not at the optimal tax system, in rich countries, we are close enough that the above logic provides a reasonable approximation. We stress again that Weyl is not claiming that policies voted on would be

Footnote 17 continued
proper extent of income redistribution. Therefore, a distribution-neutral approach to policy analysis is warranted." Our paper deals with a somewhat different topic, voting institutions, but we agree with the basic sentiment that public decisions should not constantly be entangled with unrelated distributive questions; rather public decisions should be decided on their merits. However, putting aside broader moral considerations that go beyond the utilitarian framework, we think that the utilitarian criterion is superior to the efficiency criterion as a measure of the intrinsic merit of a policy.

[18] We have verified with Weyl that this is the view he was articulating here.

accompanied by corresponding tax adjustments to make them distributively desirable; that would just be a more intricate version of the compensation argument that we have already addressed. Rather, he is arguing that in the current circumstances, more efficient policies will have a strong tendency to be utilitarian superior without compensation.

We have several objections to this line of argument. First, and most obviously, we are skeptical that we are close to a utilitarian optimal tax system. In the absence of sufficient evidence to this effect, the argument does not get off the ground. Second, how close to an optimal tax system must we be for this to be a reasonable approximation? How close is it reasonable to expect that we are? These claims must be made more precise. Third, even if the argument holds for policies that have small effects, it seems unlikely that it would be valid for policies that have large effects. The democratic system has to process issues that have large effects. Finally, the basic voting system underlying a democracy must work well in a variety of circumstances. So even if we are transiently close to an optimal distribution factoring in incentive constraints, that fortunate occurrence is unlikely to last indefinitely. The need for a good voting system may be greater when circumstances are bad than when they are good.

Consider the broader argument that for some reason, not necessarily because we are at an optimal tax system, efficiency and the utilitarian objectives will coincide. The idea, mentioned in Sect. 2.4.1, that agents have common interests with respect to social insurance also had this flavor.

For any public decision x, let u_i^x be i's utility under x. If, for example, a public project generates utility benefits because it generates incentives to work, as Weyl suggests in the quotation above, then for any beneficiary, those benefits will be included in u_i^x. This assumes that voters are rational and farsighted. In this setting, the condition that efficiency and the utilitarian criterion rank decisions x and y in the same way amounts to

$$\underbrace{\sum_i u_i^x \geq \sum_i u_i^y}_{\text{utilitarian superiority of } x} \Leftrightarrow \underbrace{\sum_i \frac{u_i^x}{b_i} \geq \sum_i \frac{u_i^y}{b_i}}_{\text{greater efficiency of } x}, \tag{19}$$

where, recall, b_i is voter i's marginal utility of wealth. Condition (19) holds trivially if wealth is distributed equally, since then $b_i = b_j, \forall i, j$, but we are not to imagine that that is the case. Another case in which (19) holds is when issues are independent of wealth in the sense introduced in Sect. 2.3.5.

Let X be the set of all potential public decisions. It would be very restrictive to assume that (19) holds for all pairs of policies in X. Let us start with the easiest case to analyze, simply because it is two-dimensional, which is when issues are polarized by wealth as in Sect. 2.3.5, so that there are two types of agents, rich and poor, with $b_R \neq b_P$, and all agents within a wealth class share the same utility for all projects. Each policy x corresponds to the utility pair (u_P, u_R) it generates in $\mathbb{R}^2$, in which case X can be represented as a subset of $\mathbb{R}^2$. Then the linear indifference curves of the utilitarian objective $u_P + u_R$ and the efficiency objective $\frac{u_P}{b_P} + \frac{u_R}{b_R}$ in u_P-u_R space will have different slopes, and so if X is a two-dimensional convex set, many pairs of policies in X will violate (19). Higher-dimensional cases, with many different wealth levels, will be analogous.

Of course what matters is not whether there will ever be divergence between efficiency and the utilitarian objective, but how frequent it will be, and how serious it will be when it occurs. The answer to this question, especially when we consider the broad range of issues that can come up for vote, including not just economic issues, but also cultural issues such as gun rights and gay marriage, and the very uncertain agenda-setting process for selecting pairs of issues to come up for vote, strikes us as highly non-obvious. We do not see how one can support the claim that the efficiency and utilitarian objectives will in general align.

These considerations suggest that the utilitarian assessment of QV does not primarily concern the inherent properties of QV as a voting institution. Let us grant that QV will select the more efficient option in pairwise comparisons. The question becomes: how well does the ranking of policies in terms of efficiency or willingness to pay mirror the utilitarian ranking? Note that the question is not whether utilitarian optimal policies are efficient, but rather whether the pairs of policies that actually come up for vote, given whatever agenda-setting process selects them, will be ranked similarly by the utilitarian and efficiency criteria. Pareto efficiency (subject to whatever constraints are operative) is highly desirable property—a necessary condition for optimality—if a utilitarian is facing a constrained optimization problem, since if it fails to be satisfied, everyone can be made better off.[19] It is very different to say that measures of efficiency, such as aggregate willingness to pay, point in the same direction as the utilitarian criterion for pairwise comparisons. Yet it is really only the latter claim, a claim we believe to be very difficult to establish, that is relevant to the utilitarian assessment of QV. It is best to admit what should be fairly obvious anyway, that efficiency is one thing and utilitarian merit is another.

3 Democratic legitimacy

The previous section evaluated QV from the standpoint of utilitarianism. This section sets aside these utilitarian concerns to consider QV through the lens of a broader range of social values relevant to public decisions, especially the value of democratic legitimacy. Most of our discussion assumes that society faces a fixed binary decision. The formal results on QV are for this case. Sections 3.6 and 3.7 briefly discuss agenda-setting and multiple alternatives.

3.1 Political institutions and public goods

We now revisit the motivation for QV with an emphasis on the role that it implicitly assigns to political institutions. Lalley and Weyl (2015) write, "Economists have typically been skeptical of the possibility of public decisions being taken as efficiently as private goods are allocated, as reflected in the formal results of Arrow (1951), Samuelson (1954), Gibbard (1973) and Satterthwaite (1975) and in informal attitudes in work such as Friedman (1962). In this paper we have argued that this attitude may be an artifact of particular institutions. Public goods do not appear to pose a fundamentally harder mechanism design problem than that posed by private goods." This quotation suggests the following picture. Markets efficiently provide private goods, but not public goods. Public institutions may also fail to efficiently provide public goods, but an ideally designed public political institution—such as QV—may be able to solve the problem. One might go further and say that the distinction between market institutions and (ideally designed) political institutions is the distinction between institutions for efficiently allocating private goods and institutions for efficiently making public decisions. We have challenged this

[19] Pareto efficiency is a desirable property when we are near an optimum; if for example, we are far from a utilitarian optimum (e.g., one person has all of the resources), then Pareto efficiency in and of itself has little merit.

conception on utilitarian grounds above, but another challenge arises from considerations of the role that political institutions are designed to play.

3.2 The broader purpose of political institutions

Voting institutions, and political institutions more generally, can be assessed on grounds other than efficiency (or utilitarian grounds as in Sect. 2). To make the relevance of other considerations clear, it is useful to reflect on some public decisions made through political institutions. These include decisions about whether to allow abortion, euthanasia, gambling, prostitution, and sale of organs, whether to go to war, whether to increase or decrease the government's supply of nuclear arms, decisions about the progressivity of the tax system and redistributive transfers, about public provision of healthcare and education, about how we ought to deal with climate change, decisions concerning the balance between privacy, freedom and security, and about what sort of criminal justice system we ought to have, what sort of antidiscrimination and affirmative action laws there ought to be, whether to make reparations to the victims of past wrongs, and so on.

Considerations of efficiency play a crucial role in many of these decisions, but clearly many of these decisions depend on social values beyond efficiency. Take the first issue on the list: abortion. At least part of the question is whether abortion is murder. Even issues such as climate change, in which efficiency figures centrally, contain other important moral dimensions, such as what we owe to future generations.

3.3 The legitimacy requirement

Important public decisions share some characteristics. They concern basic social values. Citizens regularly disagree about these decisions, either because they hold different social values, or because they draw different conclusions about what policies those values support. Such decisions often concern the structure and operation of institutions that affect large numbers of people, and, once made, are backed by the coercive powers of the state. To summarize, these decisions have high stakes, involve social values, and are subject to disagreement.

Disagreement is a crucial feature. There is no feasible system that can choose policy A over B if and only if A is more just, or if and only if A produces more welfare. Not only is this impossible, but there will be disagreement about what is just, or what will enhance welfare. This disagreement motivates a focus on the procedure whereby the decision is made. Does the procedure manage the disagreement well and fairly? Does it give everyone a say who ought to have one?

These considerations suggest that public decisions face a democratic legitimacy requirement. This is a requirement on the *process* of making public decisions rather than their outcomes. We formulate it as follows.

Democratic legitimacy requirement *For major public decisions*, (*1*) *citizens should be consulted*, *and* (*2*) *should have an equal opportunity to influence the outcome*.

The consultation condition (1) rules out, for example, a lottery that decides the outcome independently of anyone's opinion. Under the lottery, everyone is treated equally but no one has any influence. Equality is not sufficient; citizens must collectively control the decision. Condition (2) refers to *opportunity* because it is sufficient to provide citizens the opportunity to influence the outcome. If some citizens choose not to exercise this opportunity, this does not undermine legitimacy. The legitimacy requirement applies to public decisions and not private decisions, such as those within an individual's personal sphere.[20]

Our formulation is only provisional; it should be viewed as a sketch that points to the general sort of consideration we think is important; we cannot hope to elaborate a complete theory of democratic legitimacy here. We are not advocating that the legitimacy requirement should take absolute priority over efficiency or any other value. What we do maintain is that democratic legitimacy is a serious value, which is critically important in the evaluation of alternative voting arrangements.

To illustrate the legitimacy requirement, let us compare two nations who, under different political regimes, decide to go to war. The two nations find themselves in the same circumstances, and they make the same decision. In the first nation, the decision is made democratically. In the second, the decision is made by an authoritarian government. The reasons for the first and second nations to go to war have equal merit. Nonetheless, the second nation fails to meet a democratic legitimacy requirement that the first nation meets. The difference concerns who makes the decision and how it is made. This difference has justificatory consequences: citizens of the second nation have a complaint against the decision that citizens of the first nation lack. Citizens in the second nation can say, "we were not consulted," or "we did not have a say." A citizen of the first nation, even one who rightly opposed the war, cannot issue such a complaint.

As the same decision is made in both cases, it is not efficiency that separates the two decisions, but rather democratic legitimacy.[21] Now consider instead the case where the two nations make different decisions. For at least some issues, the difference between the ways the decisions were made—and, hence, whether the decisions were legitimate—may be more important than the question of which decision was made. This too cannot be accounted for by a conception of political institutions that evaluates such institutions exclusively in terms of the efficiency of their outcomes.

3.4 Specific egalitarianism for decision-making authority

An authoritarian government choosing to go to war is an extreme case of inequality in the distribution of authority. But a problem of legitimacy arises also for milder forms of inequality in decision-making authority. For example, unequal voting rights raise a concern about democratic legitimacy. Democratic legitimacy thus appears to require some kind of equality in decision-making power, and this is what our formulation of the legitimacy requirement attempts to capture.

Tobin (1970) introduced the term "specific egalitarianism" to describe a demand of equality (or more equality) with respect to certain specific goods (such as healthcare and education). Economists often resist such specific egalitarianism and seek equality—to whatever degree they do—on a global level, that is, with respect to goods considered in

[20] A characterization of the boundary between public and private decisions is beyond the scope of this paper.

[21] One may argue that viewed from a broader perspective, the decisions may not be equally efficient. If one considers the entire system of democratic government in contrast with the authoritarian system, with all of their consequences, the democratic system may be more efficient. This is efficiency at the level of the entire political system, not efficiency at the level of the binary decision under consideration. We do not think that this broader notion of efficiency can capture the motivation behind the legitimacy requirement. In any event, the formal efficiency results for QV establish that QV is efficient at the level of binary decisions, and so we focus on this narrower notion of efficiency.

aggregate. Weyl echoes this sentiment and extends it to argue against specifically political equality, writing, in personal correspondence, "Many have advocated specific egalitarianism for some private goods, like health care, educations, organs, etc. But in general people seem to feel that it is ok to have general egalitarianism for most private goods, but want to wall off public goods entirely. Is this a coherent position? What, other than the technological structure, makes the categorical difference between public and private goods?"

We do think that there is an argument for specific equality with respect to the allocation of decision-making authority in a democratic society. However, it is important to clarify that this is not an argument that some goods, such as public goods, should be allocated more equally than others. Rather, we think that decision making authority should be specifically equally allocated. This is a requirement of equality at the higher order level of the collective decision mechanism rather than the lower order level of the specific allocation or outcome of the decision.

While legitimacy concerns higher order equality, it also bears on equality at the level of the allocation. Systems for allocating goods that result in inequality are sometimes justified on grounds of efficiency, but nevertheless raise pressing questions about justice. In particular, these inequalities should be justifiable to those on their lower end. When these inequalities have been imposed by an illegitimate procedure, we think that it is difficult to meet this justificatory burden. When inequalities can be justified, it is, we think, a crucial part of the story that the system for allocating goods is at least contestable through a reasonable public decision procedure in which the affected parties are consulted on terms of equality. To the extent that citizens can decide as democratic equals that the benefits of allowing some degree of inequality, in terms of welfare, freedom, and efficiency, outweigh the costs, the justification for material inequalities may be strengthened. Democratic equality is thus a critical ingredient in the legitimation of material inequality.[22]

3.5 QV and democratic legitimacy

QV is designed to solve the problem of making efficient public decisions, but how does it fare with respect to the criterion of democratic legitimacy?

Let us separate the analysis into two cases: (1) an artificial case in which all citizens have equal wealth, and (2) a more realistic case in which different citizens have different wealth.

In the equal wealth case, QV will likely involve unequal *exercise* of authority in any particular decision, because different people value voting on the issue differently, and so purchase different numbers of votes. For example, if Ann buys more votes than Bob because Ann cares more than Bob about the issue under vote, then Ann will have more influence on that decision. However, since both Ann and Bob could have bought the same number of votes (roughly) at the same sacrifice of material well-being, Ann and Bob had (roughly) equal opportunity to exercise influence. We think that it is really equality of *opportunity* for influence that matters for legitimacy. If citizens have equal opportunity to

[22] In his contribution to this volume, Ober (2016) argues that QV has trouble with a requirement of democratic legitimacy. While we are friendly to his conclusion, we do not base the legitimacy requirement on the claim that we have equal common interests in the distribution of public goods. On our account, democratic legitimacy is a higher order property that holds when people are consulted and have equal opportunities to influence the outcome of public decisions. On our account, QV has problems with democratic legitimacy even in cases where legislation concerns the distribution of private goods, as well as in cases where individuals have different levels of interest in public goods.

vote but they do not exercise that opportunity since the issue matters to them less, this does not appear to present a problem for legitimacy. In light of the choice to forgo exercising the opportunity, they have no basis for a complaint that their views were not adequately consulted in reaching the decision.

But now consider the unequal wealth case. QV allocates the opportunity to exercise political decision-making authority according to wealth. This directly raises the legitimacy problem. To see this, put QV aside for the moment, and imagine that there are two groups of citizens, R and P. Citizens in group R are allowed to vote under a majority rule system, and citizens in group P are not allowed to vote. This clearly raises issues of legitimacy, and citizens in P have a cause for complaint. Now consider QV instead of majority voting, letting R stand for "rich" and P stand for "poor". Consider what happens in the limit as the rich become very rich and the poor become very poor, so poor that it is prohibitively expensive for them to vote. In the limit the situation is essentially the same as that in which the citizens in P were simply excluded from voting while those in R were allowed to vote. But the problem does not just occur in the extreme case; the problem is continuous, and the more unequal the wealth distribution the more severe the problem becomes.

Majority voting will not face this problem, at least not in the sharp form presented above. Each voter is granted a single vote. In most circumstances, this should mean that voters have an equal opportunity for influence, at least within the formal voting institution. This equality of opportunity will generally not depend on wealth or other extraneous factors.

The case for majority voting vis-à-vis equality of opportunity for influence seems unimpeachable with sufficient symmetry. However, a voting rule in itself may not be able to guarantee that the situation will be symmetric in the requisite respects. For example, certain individuals may belong to a persistent minority. How this is to be analyzed in terms of opportunity for influence is a complicated matter, but this certainly raises a doubt about the degree to which even majority voting can realize this ideal.

To think about this, consider that voting institutions must be robust over the long sweep of time. Imagine an initial "constitutional" stage in which we select our basic voting institutions. If we choose majority voting, we may be able to predict that there will sometimes be persistent minorities. Nevertheless, majority voting is quite a neutral rule that doesn't particularly disfavor any ex ante identifiable group. Perhaps there is superior neutral mechanism that will mitigate the problem of persistent minorities. But it cannot be an improvement, from the standpoint of equal opportunity, to implement a system that will permanently provide the wealthy with a voting advantage. That would be to substitute for transient disadvantages that arise from many sources a permanent disadvantage based an easily identifiable characteristic. Moreover, this characteristic, namely wealth, is itself one of the most important sources of social power. So if we institute QV or some similar institution at the initial stage, we will be announcing to citizens ex ante that if you lack economic power, your opportunity to influence political outcomes through the formal voting mechanism will thereby be diminished as well. This shows clearly, to our mind, the relative advantage of majority voting along this dimension.

3.6 QV and the current system of majority rule

Looking at formal voting institutions in their broader social context, one may argue more forcefully that under the current system of majority rule, different voters have different influence. For example, wealthy voters will find it easier to make campaign contributions, and to hire lobbyists. They thereby have the opportunity to exercise outsized influence on the outcome of the electoral process, or, once candidates are elected, on legislative agenda-setting, and on the drafting of legislation. Posner and Weyl (2015) argue that buying votes through QV could be a substitute for other forms of influence buying, and indeed, that the net effect of moving from our system of majority voting to QV would be that money would have less influence on politics. If true, this would be relevant to assessing the comparative democratic legitimacy of their proposed reform and the current status quo. As we are sure they would agree, more would need to be done to both flesh out the proposed reforms and to empirically support the contention that they would lessen the influence of the wealthy.

We do not wish to pretend that our current political system is close to ideal. Our strategy has been to pursue a simple apples to apples comparison of voting systems in order to highlight a normative consideration that we think should be emphasized. We put aside many highly speculative questions about how the introduction of QV would indirectly change the political environment as a whole, and focused on what we believe would be the relatively direct and unavoidable effects of its introduction. In the treatment of broader issues, the relevant comparison would not be between QV and majority voting as isolated voting systems, but rather between QV and majority voting in the context of a much broader account of the surrounding political environment and complementary political institutions.

Even if there turn out to be gains of democratic legitimacy elsewhere in the system, in terms of superior agenda-setting or in terms of reduction of the influence of wealth through other means—a point that we are not ready to concede—they come at the cost of legitimacy that comes with adopting the voting procedure of QV. This cost is significant, because from the standpoint of legitimacy, inequality of opportunity for influence that is formally incorporated into the voting system is likely to be worse than similar informal inequality generated by the social context in which voting occurs. This is because formal voting rules that give advantage to wealthier voters signal more strongly than the availability of outside opportunities for influence that we as a society do not care about providing everyone with an equal say. For this reason, we think the offsetting gains to democratic legitimacy in the reduction of informal influence face a high bar, and not a bar that is likely to be met.

3.7 Multiple alternatives

Our analysis has been conducted in the context of a fixed binary decision. How would the analysis be affected if there were multiple alternatives?[23] It is well known that majority (or plurality) voting have problems in the presence of multiple alternatives (Condorcet 1785). Indeed, it has been claimed that no voting rule can be satisfactory in the presence of multiple alternatives (Arrow 1951).

We have a few comments. First, the existing formal results on QV only concern binary public decisions, and in the presence of a fixed binary decision, majority voting does not face commonly discussed problems associated with Arrow's theorem, agenda-setting, or strategic voting. Second, the democratic legitimacy requirement is a procedural requirement that has to do with shared joint control of decisions. While such issues have been discussed by social choice theorists, the legitimacy requirement does not map straightforwardly onto the requirements most commonly discussed, such as, e.g., those that feature

[23] This is related to the issues raised in Sect. 3.6. Agenda-setting is only important against a background of multiple alternatives. That drafting of legislation, not just voting on legislation is important implies that there are more than two alternatives.

in Arrow's theorem or the Gibbard–Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975). It would be interesting think about legitimacy in the context of multiple alternatives, but it would take us too far afield. The comparison of QV and majority voting with respect to the legitimacy requirement is relatively straightforward in the case of a fixed binary decision that we have considered.

4 Conclusion

This paper evaluated QV from two ethical perspectives, a utilitarian perspective and the perspective of democratic legitimacy. Both perspectives revealed problems for QV when wealth is unequally distributed. Of the two, the utilitarian perspective issued in a finer diagnosis: even if wealth is unequally distributed, QV will outperform majority voting if voters' interests are independent of their wealth. However when interests are polarized by wealth, so that the interests of the rich and poor are opposed, then majority voting may be superior to QV. The comparison between QV and majority voting is thus ambiguous from a utilitarian perspective.

The perspective of democratic legitimacy was not similarly ambiguous. Given the inevitability of economic inequality, vote-buying mechanisms unlike majority rule, come at a cost to the democratic legitimacy of voting institutions. This comparison is straightforward.

Let us contrast the two perspectives. The utilitarian perspective evaluates voting institutions on the basis of the utility of their outcomes. In contrast, our arguments about democratic legitimacy stem from the intuition that all citizens have equal *claims* on society's major public decisions. It is useful to compare such "claim rights" to property rights. To make the comparison to utility vivid, imagine that you and I own an object together. You come to me and say, "I care about the object more. So you should let me do with it what I wish." I reply, "that may be so, but it is irrelevant. We own the object equally, and I have just as much of a right to it as you." We think that something similar is involved in our equal rights to have a say in decisions that concern us all.

While our common claims on social decisions do not take absolute precedence over considerations of utility, it is very important to honor these claims. Moreover, we are not in a position to say that considerations of utility conflict with these claims because it is very difficult to know which voting system would be best from a utilitarian point of view.

We close with a variant of QV that we think is promising, and merits further exploration: QV with an artificial currency, a variant raised by Posner and Weyl. For example, a voter may have a budget of votes that she can spend over several elections, or within an election over several issues. Returning to the common ownership analogy, suppose that we jointly own a collection of objects. Then there is no violation of our ownership claims if we negotiate a division, which is what QV with an artificial currency effectively allows us to do.

Acknowledgements We thank Glen Weyl for many illuminating and enjoyable conversations on this topic. We are grateful to Emilee Chapman for detailed comments and to participants at the Conference on Quadratic Voting and the Public Good.

Appendix: Proof of Proposition 5

Let b_P and b_R be the marginal utilities of wealth for the poor and rich respectively. Then $b_R < b_P$. Using (16) and the fact that $u_R < 0 < u_P$, QV selects decision 1 if and only if

$$(1-\alpha)\frac{u_P}{b_P} \geq \alpha\frac{|u_R|}{b_R},$$

or equivalently, if and only if

$$u_P \geq \frac{\alpha}{1-\alpha}\frac{b_P}{b_R}|u_R|.$$

On the other hand, for a utilitarian, it is optimal to select decision 1 if and only if

$$u_P \geq \frac{\alpha}{1-\alpha}|u_R|.$$

Since $b_P > b_R$, the threshold on u_P that QV demands for selecting decision is too high.

Majority voting always selects decision 1, since there are more poor than rich voters.

Select t_0 as $t_0 = \frac{\alpha}{1-\alpha}|u_R|$. When $u_P < t_0$, then it is optimal to select decision 0, and indeed QV selects decision 0, while majority voting will select decision 1. Select $t_1 = \frac{\alpha}{1-\alpha}\frac{b_P}{b_R}|u_R|$. Then on (t_0, t_1), it is optimal to select decision 1, majority voting selects decision 1, and QV selects decision 0. Finally, when $u_P > t_1$, then both majority voting and QV select decision 1. This completes the proof. □

References

Arrow, K. J. (1951). *Social choice and individual values*. Hoboken: Wiley.

Atkinson, A. B., & Stiglitz, J. E. (1976). The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics, 6*(1–2), 55–75.

Boadway, R., & Keen, M. (1993). Public goods, self-selection and optimal income taxation. *International Economic Review*, 463–478.

Buchanan, J. M. (1959). Positive economics, welfare economics, and political economy. *The Journal of Law and Economics*, *2*, 124–138.

Christiansen, V. (1981). Evaluation of public projects under optimal taxation. *The Review of Economic Studies*, *48*(3), 447–457.

Condorcet, M. (1785). *Essay on the application of analysis to the probability of majority decisions*. Paris: Imprimerie Royale.

Friedman, M. (1962). *Capitalism and Freedom*. Chicago: University of Chicago Press.

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

Harsanyi, J. C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy*, *61*, 434–435.

Hylland, A., & Zeckhauser, R. (1979). Distributional objectives should affect taxes but not program choice or design. *Scandinavian Journal of Economics*, *81*(2), 264–284.

Hylland, A., & Zeckhauser, R. (1980). *A mechanism for selecting public goods when preferences must be elicited*. Unpublished manuscript

Kaplow, L. (2004). On the (ir)relevance of distribution and labor supply distortion to government policy. *The Journal of Economic Perspectives*, *18*(4), 159–175.

Lalley, S. P., & Weyl, E. G. (2015). *Quadratic voting*. Unpublished manuscript.

Lalley, S. P., & Weyl, E. G. (2016). *Quadratic voting*. Unpublished manuscript, June 2016.

Ober, J. (2017). Equality, legitimacy, interests, and preferences: Historical notes on quadratic voting in a political context. *Public choice*. doi:10.1007/s11127-017-0409-0.

Posner, E. A., & Weyl, E. G. (2015). Voting squared: Quadratic qoting in democratic politics. *Vanderbilt Law Review*, *68*, 441–500.

Samuelson, P. A. (1954). The pure theory of public expenditure. *Review of Economics and Statistics, 36*(4), 387–389.
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.
Tobin, J. (1970). On limiting the domain of inequality. *The Journal of Law and Economics*, *13*(2), 263–277.
Vickrey, W. (1945). Measuring marginal utility by reactions to risk. *Econometrica*, 319–333.
Weyl, E. G. (2015). *Price theory*. Unpublished manuscript.
Weyl, E. G. (2017). The robustness of quadratic voting. *Public choice*. doi:10.1007/s11127-017-0405-4.