Ethical Considerations on Quadratic Voting[*]

Ben Laurence[†]   Itai Sher[‡]

March 22, 2016

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 welfare or 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.

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 (2015a) 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 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,

[†]Department of Philosophy, University of Chicago; email: benlaurence@uchicago.edu
[‡]Department of Economics, University of California, San Diego; email: itaisher@gmail.com

vote buying seems to give greater weight to the preferences of the wealthy who can better 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 Section 2, we explore the interaction of QV with economic inequality in the utilitarian context that is QV's natural home. We find that if voter preferences are independent of wealth, then quadratic voting is utilitarian preferred to majority voting, but if voter preferences are polarized by wealth, then majority voting may be superior. In Section 3, we move beyond the utilitarian perspective to consider questions that concern democratic legitimacy. This latter part 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 quadratic voting, 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 evaluates quadratic voting from a utilitarian perspective. The next section will look to ethical considerations beyond utilitarianism. At times, it appears that the advocates of quadratic voting support it on utilitarian grounds. For instance, Posner and Weyl (2015) write, "... quadratic voting solves the preference aggregation problem by giving proper weight to preferences of varying intensity ... ." While they do use the phrase "utilitarian efficiency", Lalley and Weyl (2015a) predominantly emphasize efficiency rather than preference intensity. Recall that an outcome is Pareto efficient if it is impossible to make one person better off without making someone else worse off. It is well known that efficiency and utilitarian optimality are distinct properties. One may believe, for example, that interpersonal comparisons of utility are meaningless and still hold that efficiency is a desirable characteristic. In some settings, such as that studied by Lalley and Weyl (2015a), the efficient outcomes coincide with those that maximize the sum of agents' willingnesses to pay for the outcome.[1] This formula is reminiscent of the utilitarian formula of maximizing the sum of utilities, but, of course, as is well known, these two formulas are not the same.

The term "utility" has many meanings and it is easy to slide back and forth between

---

[1]This coincidence occurs if one takes the model literally. We think, however, that it is better to interpret the quasilinear utilities as a local approximation, in which case, the set of outcomes that maximize the sum willingness's to pay in Lalley and Weyl's model is better interpreted as the subset of efficient outcomes that can be achieved with transfers sufficiently small so as not to significantly change voters' willingness to pay for outcomes.

different meanings. Utility may mean "decision utility"; that is, it may simply encode preference or choice. Alternatively, "utility" may have a richer ethical meaning as in the tradition of utilitarianism.[2]

The purpose of this paper is to assess quadratic voting on ethical grounds. This requires that we appeal to the richer ethical meaning of "utility". The formal results of Lalley and Weyl (2015a) are really best interpreted as showing that QV is efficient rather than as showing that it is utilitarian optimal.[3] This leaves open the question of how QV fares from a utilitarian perspective. That is the question we take up here.

Even within the utilitarian tradition, utility has been interpreted in many ways. For our purposes, we do not need to take a stand on the best interpretation of utility. Below, we write about utility as if it can be represented by the preferences of an impartial utilitarian ethical observer, following an economic tradition pioneered by Harsanyi (1953). However, this is inessential to our argument. We only make a single ethical assumption on utility: namely, that the marginal utility of a dollar is decreasing in wealth, where "utility" is used in an interpersonally comparable ethically significant sense. In other words, 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 because it would make more of a difference to the poor person. This assumption is widespread in the utilitarian tradition.

2.1 QV as a Stand-Alone Policy

Let us clarify the nature of the analysis considered here. We now consider QV as a *stand alone policy*. This means that a wealth distribution is given and there is no complimentary policy to deal with distributive concerns. Indeed, in our analysis, we even neutralize the refund generated by QV, invoking what we call a *no redistribution constraint*. We explain the motivation for this last feature of our analysis below. So, this section amounts to discussing QV as an isolated innovation without other supporting policies. In Section 2.2, we discuss bundling QV with a redistributive policy and related issues. We also say a few words about the nature of the inefficiency that is generated when majority voting outperforms QV from a utilitarian perspective.

---

[2]It is not that preferences are irrelevant from an ethical perspectives. Indeed, some versions of utilitarianism are founded on satisfying agent's preferences. But in this case, utility must do more than just "encode" preferences in the ordinal sense that a more preferred outcome is given a higher utility. Utility must also capture preference intensity, understood in an ethically significant and interpersonally comparable sense.

[3]Even assuming quasilinear utility, efficiency and utilitarian optimality should be distinguished. We can see this by observing that quasi linearity is just a functional form assumption on the utility function representing preferences and so it cannot tell us how, ethically, benefits of different agents should be traded off. So, while an allocation that maximizes the sum of quasilinear utilities is efficient, it may not be utilitarian optimal. The utilitarian may prefer a specific efficient allocation.

To make our points precise, we construct a formal model, parallel to the price-taking equilibrium model of Lalley and Weyl (2015a).[4] The price-taking model is less rigorous than the Bayesian Nash equilibrium model that Lalley and Weyl (2015a) also present, but the price-taking model is simpler, and it is sufficient to make our points. Our model modifies the Lalley and Weyl's price-taking model to accommodate the utilitarian perspective.

Suppose that there is a public decision $x \in \{0, 1\}$. For example, $x = 1$ may be the decision to undertake a public project, and $x = 0$ may be the decision not to undertake the project. There is a collection N of agents such that each agent i has utility function:

$$U_i(x, w) = u_i x + g_i(w_i)$$

where u_i is i's utility of the public decision $x = 1$, w_i is i's wealth, and $g_i(w_i)$ is the utility that i gets from wealth w_i. All agents receive a utility of 0 from the decision $x = 0$, so whether an agent prefers $x = 0$ or $x = 1$ depends on whether u_i is negative or positive.

Assume that there is a utilitarian ethical observer who would like to maximize:

$$\sum_{i \in N} U_i(x, w_i) \tag{1}$$

That is, the utility functions U_i have been put on an interpersonally comparable scale from the standpoint of the ethical observer. Next assume for simplicity that all agents convert wealth into utility in the same way. Formally,

$$g_i = g_j =: g, \quad \forall i, j \in N,$$

so that g is the common utility function of wealth. g is increasing in w_i. Assume diminishing marginal utility of wealth. Formally, this means that g is concave.

A voting mechanism is described as follows. Knowing her own value for an electoral outcome, a voter chooses how many votes to purchase. The voter then uses her purchased votes to vote for either $x = 0$ or $x = 1$ in the election. The alternative that receives the most votes wins. A voting institution is therefore determined by a cost function c for purchasing votes. For simplicity, we can think in terms of purchasing votes *for* $x = 1$, so that purchasing v votes for $x = 0$ is equivalent to purchasing $-v$ votes for $x = 1$. We allow that votes may be purchased in non-integer quantities. We assume that $c(0) = 0$.

One might not initially see how majority voting can be fitted into this framework, but it can. Specifically, consider the cost function:

$$c(v) = \begin{cases} 0, & \text{if } v \in [-1, 1] \\ +\infty, & \text{otherwise.} \end{cases} \quad \text{(majority voting)}.$$

---

[4]This version of the model has been removed from the latest draft of Lalley and Weyl (2015b).

Thus, one can costlessly buy up to one vote for $x = 1$ or up to one vote for $x = 0$, but purchasing more votes is prohibitively expensive, and so effectively impossible.

The cost function for quadratic voting is:

$$c(v) = kv^2 \quad \text{(quadratic voting)}, \tag{2}$$

where k is some positive constant.

In QV, the proceeds of the election are refunded to the population. The key property of this refund is that it is independent of how one votes, and even whether one votes. Otherwise, the refund would have to be included in the cost function c. The most natural way to refund the proceeds would be to give every member i of the population a refund of $\frac{\sum_{j \in N \setminus i} c(v)}{|N| - 1}$. Under this scheme, everyone receives an (approximately) equal refund from the election. One can easily show that this version of the refund has an ambiguous effect on utility from the utilitarian perspective.[5] To see that there is no strong logic dictating the effect of the refund, imagine that the population of voters is split into two sets A and B. Voters in A care very much about the election and those in B care very little. So voters in A purchase a large number of votes while voters in B purchase a small number of votes. Still voters in B will receive their share of the refund, because a critical property of the refund is that it is independent of voting behavior. 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, and depending on how these parameters are resolved, the refund could have any conceivable effect on utility.

If there were not diminishing marginal utility of wealth, this would not be an issue, because the utilitarian observer would be indifferent about transfers. But it is a crucial part of our story that there is diminishing marginal utility of wealth. So, in order to come to an unambiguous conclusions, we would like to neutralize the effect of the refund. There is a more fundamental reason for doing this. We really want to look at the utilitarian assessment of the *outcome* of the vote, and thinking about the redistributive transfers generated by the voting mechanism detracts from a clear assessment of this issue.

To address the above issues and neutralize the refund, we introduce the following device. Let us suppose that in the society we are talking about, redistribution of wealth through the election refund is forbidden. Imagine that the population is divided into two parts, citizens and non-citizens. The citizens vote and non-citizens are not allowed to vote. Each dollar raised through the election from citizen i is transferred to a non-citizen j whose wealth prior to the transfer is such that the marginal utility of the dollar is the same for i and j. $C \subseteq N$ is the collection of citizens and $N - C$ is the collection of non-citizens. Assume that for all $i \in N - C$, $u_i = 0$, so non-citizens do not care about the public decision. Therefore we don't have to think about the effect of the public decision on non-citizens.

---

[5]In Section 2.2, we return to the possibility that the redistributive transfers generated by the refund are made conditional on wealth.

Given the *no redistribution constraint* and the above provision for maintaining it, every dollar taken from one agent i is transferred to another agent j with the same marginal utility of wealth, the redistribution of wealth through the refund will have no effect on utility. Therefore the utilitarian ethical observer would like to choose $x \in \{0, 1\}$ so as to maximize

$$\sum_{i \in C} u_i x \tag{3}$$

Thus, the utilitarian ethical observer judges the outcome solely on the basis of whether the utility maximizing public decision is made and ignores the (utility-neutral) redistributive transfers.

Following Lalley and Weyl, we now define the notion of a price taking equilibrium. A **collective decision problem** is a tuple $\{C, S, \bar{u}, g, \bar{w}^e\}$, where C is the collection of citizens, S, a positive real number, is the supply of influence, $\bar{u} \in \mathbb{R}^N$ is the profile of utilities for the public decision, g is the (common) utility of wealth function, and $\bar{w}^e \in \mathbb{R}^N_{++}$ is the profile of wealth endowments, so that the i-component of $\bar{w}^e$, w_i^e is the wealth endowment of voter i.

A **price-taking equilibrium** of a collective decision problem under voting rule c is an **influence vector** $I^* \in \mathbb{R}^N$, a **price** $p^* \in \mathbb{R}_{++}$ and a decision $x^* \in \{0, 1\}$ such that:

$$I_i^* \text{ maximizes } u_i I_i + g(w_i^e - c(p^* I_i)) \text{ over all } I_i \in \mathbb{R} \quad \text{(Price-taking)} \tag{4}$$

$$\sum_{i \in C} |I_i^*| = S \quad \text{(Market clearing)} \tag{5}$$

$$x^* = 1 \Leftrightarrow \sum_{i \in C} p^* I_i^* \geq 0 \quad \text{(Majority rule)} \tag{6}$$

The **equilibrium vote profile** is the vector $v^* \in \mathbb{R}^n$ whose i-component is $v_i^* = p^* I_i^*$. That is the price p^* converts a vote into a certain quantity of influence. We now comment on each of the equilibrium conditions.

The price-taking condition says that voters maximize their utility taking the price of influence as given. The voter's objective trades of a benefit $u_i I_i$ against a cost $g(w_i^e) - g(w_i^e - c(pI_i))$. To understand the benefit term, consider (informally) that if voter i did not vote, there would be some lottery over the public decision. This lottery is outside of the voter's control. However, if the voter purchases influence I_i she can increase the probability of her preferred outcome by I_i. In other words, I_i should be interpreted as i's probability of being pivotal. $u_i I_i$ is then the difference in utility between the situation where the voter purchases I_i units of influence and that in which she purchases zero influence. The voter is attempting to trade-off this difference against the cost of making it. $g(w_i^e) - g(w_i^e - c(pI_i))$ is the utility cost of purchasing influence I_i.

The most intuitive way of thinking about the difference between our price-taking equilibrium model and that of Lalley and Weyl ($2015a$) is that Lalley and Weyl assume a constant marginal utility of wealth so that $g(w_i) \equiv w_i$, so that the objective function in the price taking condition becomes $u_i I_i - c(p^* I_i)$.[6] However, this is probably not the best way of thinking of it. In our view, a better interpretation is to think of Lalley and Weyl ($2015a$) as agreeing that there is diminishing marginal utility of wealth, but maintaining that utility is approximately quasilinear in the relevant range. This means that Lalley and Weyl's results concern (local) efficiency rather than utilitarian optimization. In any event, it is important to note once again that, in our paper, the marginal utility of wealth has ethical significance. That is, the utility of wealth is measured precisely in the units of value that the utilitarian ethical observer wishes to maximize.

The market clearing condition simply says that the total influence demanded is equal to the supply of influence. The majority rule condition says that the outcome for which more votes are purchased wins the election.

The price-taking equilibrium model is only heuristic; it is not completely coherent. There are several problems with it. First, the supply of influence is treated as exogenous whereas it should be endogenous; the total supply of influence will depend on how voters vote. Second, in this model, the price of influence is determined by the actual realization of voter utilities for the policy whereas it should be determined by the ex ante probability distribution over voter utilities. This is related to a third problem (which is perhaps not distinct from the second): Conditional on a realization of the vector $\bar{u}$ (and equilibrium behavior), the probability of being pivotal induced by purchasing any definite number of votes v_i is either one or zero; it is only ex ante (or at the interim stage) that a purchase of a definite number of votes leads to a probability of being pivotal strictly between 0 and 1.

Lalley and Weyl recognize this, and only put forward the price-taking equilibrium model as a heuristic preliminary to their more rigorous Bayesian Nash equilibrium analysis. Indeed, the price-taking equilibrium model has been removed from the latest draft of the paper. Nevertheless, we prefer to use this model here. First, the price-taking equilibrium analysis is much simpler than that of the Bayesian Nash equilibrium analysis, and we do not wish to entangle the philosophical issues that are our focus with the complexities of the Bayesian Nash equilibrium analysis. At the same time, we wish to show that our philosophical points can easily be formalized and made rigorous. Thus the price taking equilibrium analysis seems a good compromise. The price taking equilibrium model captures the essential point, which is that quadratic voting will induce vote buying proportional to voters' willingness to pay for the outcome, and will thus generate the outcome that maximizes the sum of willingnesses to pay. We now address the question of how a utilitarian would compare such a rule to majority voting under the no redistribution constraint.

---

[6] w_i^e is a constant that can be ignored.

In what follows, we assume that:

$$\forall i, \exists b_i \in \mathbb{R}_+, \forall t \in [-M, M], \quad g(w_i^e - t) \approx g(w_i^e) - b_i t \tag{7}$$

where M is some number sufficiently large that payments in the election will not exceed M. In other words, g is approximately linear over the relevant range. We will treat the approximate equality (7) as a literal equality for our analysis. The reader may assume, if she wishes, that g is a piecewise linear concave function, which would make the approximation exact. In any event, (7) only simplifies the analysis without changing it in an essential way. b_i is i's marginal utility of wealth.

Let us relax the voter's problem for a moment and assume that different agents may face different cost functions c_i in the election. Now suppose that each agent solves:

$$\max_{I_i} u_i I_i + g(w_i^e - c_i(p^* I_i)) \tag{8}$$

instead of (4). Again p^* is the market clearing price. Define:

$$h_i(t) := g(w_i^e) - g(w_i^e - t),$$
$$h_i^{c_i}(v) := h_i(c_i(v)).$$

Then (8) becomes:

$$\max_{I_i} u_i I_i + g(w_i^e) - h_i^{c_i}(p^* I_i)$$

Using an analysis similar to the proof of Theorem 1 of Lalley and Weyl ($2015a$), one can show that the price-taking model with (8) instead of (4), equilibrium will lead to the outcome that maximizes the utilitarian objective (3) for all collective decision problems if and only if there exists $k > 0$ such that:

$$h_i^{c_i}(v) = kv^2, \quad \forall i \in C. \tag{9}$$

Using (7), it follows that:

$$c_i(v) \approx \frac{k}{b_i} v^2 \tag{10}$$

The problem is that we have been thinking about a relaxation in which different agents may face different cost functions. However, in reality, we are operating under the constraint that:

$$c_i = c_j := c, \quad \forall i, j \in C. \tag{11}$$

So (10) is infeasible. It follows that:

Proposition 1 *There is no voting rule c that induces an equilibrium that is utilitarian optimal (i.e., maximizes (3)) for all collective decision problems.*

In particular, quadratic voting does not achieve this. This result highlights the difference between maximizing the sum of willingnesses to pay and maximizing the sum of utilities in an ethically meaningful sense. Quadratic voting maximizes the former but not the latter; no voting rule *such that the price of votes does not depend on one's wealth* can maximize the latter.

So what voting rule should we choose? One approach would be the following. For simplicity, hold fixed all aspects of the collective decision problem except the utility profile $\bar{u} \in \mathbb{R}^C$. Let $x(\bar{u}, c)$ be the outcome that would result given utility profile $\bar{u} \in \mathbb{R}^C$ and let $\mathbb{C}$ the set of possible cost functions. Let $\mu \in \Delta(\mathbb{R}^C)$ be the probability measure over utility profiles. To find the optimal voting rule, ideally we would like to solve:

$$\max_{c \in \mathbb{C}} \int_{\mathbb{R}^C} \sum_{i \in C} u_i x(\bar{u}, c) \mu(d\bar{u}). \tag{12}$$

This would be a purist approach, and who knows what outcome it would produce.

Instead, let us compare quadratic voting and majority voting. That is the task we have set for ourselves. 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_R^e, where $w_P^e < w_R^e$. Let α be the proportion of rich citizens in the population. Assume that $\alpha \in (0, \frac{1}{2})$. 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 $x = 1$ if and only if $\frac{\alpha}{1-\alpha}\ell$ rich citizens have utility u_i for $x = 1$.

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 $x = 1$ and all rich citizens have utility u_R for $x = 1$, where $u_R < 0 < u_P$.

Proposition 2 *If issues are independent of wealth, then QV is the utilitarian optimal voting rule.*

The argument is straightforward. $x = 1$ will win among the rich voters if and only if the sum of willingnesses to pay among rich voters for $x = 1$ rather than $x = 0$ is positive. However because rich voters all have the same marginal utility of wealth, it follows that $x = 1$ will

9

win if and only if the sum of utilities for $x = 1$ is greater than the sum of utilities for $x = 0$. The same facts hold for the poor voters. Since the distribution of preferences is the same in both groups, $x = 1$ maximizes the sum of utilities among rich voters if and only if $x = 1$ maximizes the sum of utilities among poor voters. So $x = 1$ will win if and only if $x = 1$ maximizes utility among all voters and $x = 0$ will win if and only if $x = 0$ maximizes utility among all voters.

Proposition 2 shows not just that QV is better than majority voting, but that it is optimal among all voting rules in this case.

Proposition 3 *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 their 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 their 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 willingness to pay give the same answer.*

The proof of Proposition 3 is in the appendix.

To illustrate the ideas on an informal intuitive level, suppose that we have a set of individuals of equal number who are identical except for their preferences between two policies, 1 and 2. Then one would expect those with greater preference intensity to be willing to pay more for votes than those with less intense preferences. What makes the QV payment rule special is that it will induce voters to buy votes in proportion to their willingness to pay for the outcome they prefer. So QV will plausibly lead to the utilitarian optimum. This result holds if all voters have the same wealth, and it also holds if voters have different wealth but wealth is independent of preferences.

But now take the other extreme. Suppose that there are two groups R and P of equal size, all voters are identical in terms of their preference intensity for whichever alternative they prefer, but they have unequal wealth. Moreover, suppose that preferences are polarized by wealth: The wealthier voters in R want one outcome and the poorer voters in P want

another. Then, the wealthier group will win. So in this case there is no connection between the utilitarian merit of the alternatives and outcome of the election. Even if P grows in size, so that it makes a majority and the intensity of preference of those in P increases, R may win under QV, despite the fact that P's preferred policy is utilitarian preferred and would win under majority rule.

2.2   QV and Redistribution

Above we have assumed that the refund is utility neutral, so that we could assess the *election outcome* in utilitarian terms. For example, in the case of rich and poor voters, there may be a third group of indifferent voters who absorb the refund. We now come back to the possibility of bundling QV with some redistributive transfers, as well as discussing related issues.

The efficiency of QV suggests that even when QV does not select the utilitarian optimum while majority voting does, if QV is bundled with a set or redistributive transfers, then this bundle may be utilitarian preferred to the majority voting outcome. Let us briefly consider this.

Let $W = \sum_{i \in N} w_i^e$ be the total wealth endowment of citizens and non-citizens . An **allocation** is a pair $(x, \bar{w})$ such that x is a public decision and $\bar{w}$ is a wealth profile satisfying $\sum_{i \in N} w_i = W$. We cannot really say of a public decision x that it is or is not Pareto efficient. It depends on how wealth is distributed. If even one person prefers public decision x, it is efficient to select that public decision and give that person all the wealth. So it is only an allocation that can be said to be efficient or inefficient.

This contrasts with utilitarian optimality. At least in our simple model, the public outcome x that maximizes the utilitarian objective (3) can truly be said to be utilitarian optimal independently of the allocation. Suppose, for example, that $x = 1$ is utilitarian optimal in the sense that it maximizes (3). Then consider any allocation $(0, \bar{w})$ in which $x = 0$. The allocation $(1, \bar{w})$ with the same wealth distribution but with $x = 1$ is utilitarian preferred. This may well be an artifact of our simple model. In reality, it will often be the case that a utilitarian will not be able to assess the public decision in isolation.

Consider the following "canonical example":

Canonical example Let $x = 0$ be the QV outcome, let $x = 1$ be the majority voting outcome, and suppose that $x = 1$ is the utilitarian optimum.

Let $(1, \bar{w})$ be the allocation generated by majority voting. Since the sum of willingnesses to pay for 0 rather than 1 is positive (under the initial wealth distribution), it is true that there is some wealth distribution $\bar{w}'$ such that $(0, \bar{w}')$ is Pareto preferred to $(1, \bar{w})$. On the other hand, $(1, \bar{w}')$ is *utilitarian preferred* to $(0, \bar{w}')$. So why not choose $(1, \bar{w}')$ instead?

Moreover, if $\bar{w}''$ is the egalitarian wealth distribution according to which $\bar{w}_i'' = \frac{W}{|N|}$ for all $i \in N$, then $(1, \bar{w}'')$ is utilitarian preferred to all allocations involving $x = 0$; if $x = 0$, there are no transfers that will compensate everyone relative to $(1, \bar{w}'')$ from a utilitarian point of view. Now, this model is terribly over-simplified; for one thing, $(1, \bar{w}'')$ may well not be truly feasible. But it does suggest that there is no simple general purpose *utilitarian* argument that the outcome *should be* whatever QV selects along with appropriate compensating transfers.

Let us now briefly consider some simple ways of realizing additional utility gains using standard QV and some variants. Above, we have neutralized the refund, because, as we have argued, its effect is on utility is ambiguous. Whatever its effect on aggregate utility, the refund will partially compensate the losers. The refund will also transfer wealth from passionate voters to indifferent voters, whatever are their policy preferences and wealth levels. A utilitarian would prefer the following modification: Rather than using a refund in which everyone receives an equal share of the election proceeds, the refund could be conditioned on wealth, so as to amount to a transfer from rich to poor. Call this **QV with wealth adjusted transfers**. Note this would generally not change the outcome of the election. We mention this version primarily because it would undermine our argument that the effect of the refund is ambiguous. The refund would always be positive from the utilitarian's perspective. We believe that, even in this version, majority voting would sometimes be utilitarian preferred to QV, but this is not something we have explored.

Let us consider another more interesting variant of QV. Proposition 1 shows that there is no costly voting rule that does not depend on wealth that is always utilitarian optimal. However, our analysis, and specifically (10), suggests that the costly voting rule $c(v, w) = \frac{k}{g'(w)}v^2$ is utilitarian optimal. Here v is the number of votes purchased, w is the voter's (initial) wealth, $g'(w)$ is the marginal utility of wealth and k is a constant. Call this **wealth weighted Quadratic Voting**. Under this scheme, poorer voters would pay less than richer voters for the same number of votes. Unlike standard unweighted QV, wealth weighted QV would always make the utilitarian optimal decision. In the canonical example, QV with wealth adjusted transfers selects $x = 0$ whereas wealth weighted QV selects $x = 1$. Wealth weighted QV must face the difficult issue of measuring the marginal utility of wealth. This marginal utility is not a purely behavioral quantity. It requires an ethical judgment that might itself be the subject of democratic disagreement.

While, of course, no one is talking about replacing majority voting from today to tomorrow, it is an understatement to say that the task of eventually persuading the public to adopt QV instead of majority voting even for some political decisions would be a difficult one. Our guess is that either version of wealth adjusted QV would be even more politically unpalatable than standard QV. Standard QV is also much simpler than QV with these adjustments. It is worth noting that Posner and Weyl (2015) proposed a refund that would be equally split among all voters.

A final point is in order about bundling QV with compensating redistribution by whatever means and about QV and distributive policy more generally. Imagine a democratically elected government enacting a policy that passes the cost benefit test. The government in its role as the representative of the people can bundle this policy with a redistributive scheme that will compensate the losers. Matters are quite different when the policy in question is itself the democratic process for selecting representatives and controlling policy. If QV is really to be taken as controlling policy generally, then questions of redistribution should be decided by QV. As Posner and Weyl themselves argue, a dollar is worth a dollar to everyone, and so QV could not be used to implement pure transfers. For the same reason, QV would be indifferent between bundling QV with redistributive transfers and not doing so, and QV would choose QV bundled with transfers over majority voting if and only if QV would choose QV without additional transfers over majority voting; indeed, QV would choose either version of itself over majority voting.

Since QV is apparently impotent when it comes to distributive policy and such decisions are an important part of democratic decision making, presumably, QV would have to be combined with some other decision method to provide a foundation for political institutions with universal scope. We cannot elaborate on this further here, but, in our view, a system in which QV controls all non distributive questions and distributive questions are settled by some other decision method would be very complicated.

2.3 Summary

Our conclusion is that theoretical utilitarian considerations leave the comparison between QV and majority voting ambiguous. Perhaps a utilitarian could bring empirical evidence to bear, but we view the utilitarian question as a difficult one to decide with any precision.

3 Democratic Legitimacy

The previous section addressed QV from the standpoint of utilitarianism. We now put our utilitarian qualms aside, and move on to discuss normative considerations that go beyond utilitarianism.

We return to the position of considering QV as a stand-alone policy, but we briefly consider some complimentary policies in Section 3.5. For most of this section, we abstract away from agenda setting and assume that a fixed *binary* decision is given. Some of our conclusions would have to be qualified if we were to consider agenda setting and multiple alternatives, and we will qualify them in Section 3.6.

To address broader normative concerns, we must directly address the purpose of QV.

3.1 What Problem is QV Designed to Solve?

Lalley and Weyl (2015a) 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 even go further and say that the distinction between market institutions on the one hand, and (ideally designed) political institutions on the other is the distinction between institutions for efficiently allocating private goods and institutions for efficiently making public decisions. If one does not wish to take this last step, then it is important to consider what other roles public 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 Section 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 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, whether to provide public education and healthcare, in what quantities, and by what means, 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.

Economic considerations, and efficiency in particular, play a crucial role in many of these decisions, but clearly the resolution of many of these decisions depends 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.

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. In sum, they 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 manner in which public decisions are made. The requirement does not primarily concern an external assessment of the *outcome* of decisions – e.g., the specific level of public goods such as sanitation and national defense; rather, it concerns the *process* of decision-making and who gets to have a say in it. The legitimacy requirement concerns the proper distribution of authority to decide. The legitimacy requirement does not apply to genuinely private decisions such as those within an individual's personal sphere. The precise boundary of the decisions to which it should apply is a difficult issue that we will not address here, but it clearly applies to the sorts of important public decisions that we have discussed above.

Let us compare two nations who, under different political regimes, decide to go to war. The two nations find themselves in essentially 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. Citizen 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.[7] Now consider instead the case where the two

---

[7]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, but even if

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 that 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.

3.3 Democratic Legitimacy as Consultation and Equality of 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.

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 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.

It is important to distinguish these two levels, which interact with one another in interesting ways. Although we cannot explore their interactions here at any length, the following observation is in order. 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 who have a

---

it could, the formal efficiency results for QV establish that QV is efficient at the level of binary decisions. Much more analysis is needed to analyze efficiency at the broader level of the entire political system.

complaint against the arrangement if anyone does. 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 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 a critical ingredient in the legitimation of material inequality.

A few clarifying remarks are in order. We do not wish to say that equality in the allocation of decision-making authority is *sufficient* for democratic legitimacy. Indeed, a lottery which decides the outcome independently of anyone's opinion allocates decision-making authority equally, but is not a democratically legitimate way of making important decisions. It is also necessary that citizens *are* consulted, which is not the case in the lottery. But we think that some form of equality is necessary for democratic legitimacy. We say a little more about the nature of the required equality in the next section and in the conclusion. A full theory of democratic legitimacy is clearly beyond the scope of this paper. It should also be clear from the way we have presented our argument that we are not advocating here that democratic legitimacy take absolute priority over efficiency or any other value. We do not intend to stake out such a stark position. What we do maintain is that democratic legitimacy is a serious value, which is critically important in the evaluation of alternative voting arrangements.

3.4   QV and Democratic Legitimacy

QV is designed to solve the problem of making efficient public decisions. It is not designed to solve the democratic legitimacy problem. It would call for explanation if QV – which was designed to solve the efficient decision-making problem – also solved the legitimacy problem. Indeed, there is reason to worry that it cannot. Since democratic legitimacy requires political equality, and QV involves vote buying, it seems likely that QV will have trouble in this area.

To be more precise, let us separate the analysis into two cases: (i) an artificial case in which all citizens have equal wealth, and (ii) 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 will value voting on the issue differently, and so will 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.[8] If citizens have equal opportunity to 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 R are allowed to vote under a majority rule system, and citizens in 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 acute the problem becomes.

Majority rule does not face this problem, at least in the context of a fixed binary decision, as we are considering here for both QV and majority voting. In securing to each person the right to cast a vote, it allows each the same opportunity to exercise influence over the decision. Since this opportunity is not a function of wealth, majority rule operates the same way under conditions of equality and inequality in wealth distribution. For, let the inequalities be what they may, still the decision procedure accords every citizen an equal opportunity to exercise influence over the outcome by casting her single vote.

3.5   QV and the Current System of Majority Rule

However, one may argue 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 who doggedly pursue their legislative agendas. 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. These are real problems, and that they raise disturbing questions of democratic legitimacy for our polity. Posner and Weyl (2015) argue that buying votes through QV could

---

[8]Again, equality, in any form, is not sufficient for democratic legitimacy; it is necessary. There may be other aspects of the democratic legitimacy requirement of which even QV with equal wealth would run afoul, but we do not explore any such additional requirements here.

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 context of fixed binary decisions and without considering other indirect influences of money on politics. 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 may be worse than similar informal inequality generated by the social context in which voting occurs. This is because formal voting rule 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 will have to meet a high bar.

3.6 Multiple Alternatives

Another important issue to consider is that of multiple alternatives. It is well known that majority (or plurality) rule faces problems when dealing multiple (more than two) alternatives, and it could be that these problems are relevant to the legitimacy that we have claimed for majority voting above. This will also raise the issue of agenda setting again, to which we alluded in the previous section.

Problems with majority voting were raised long ago, and have been formalized in the voluminous literature on social choice (Condorcet 1785, Arrow 1951). These problems can

be separated into various categories: (i) The problems raise the question of whether there is any coherent notion of the "will of the people" to be aggregated by a democratic procedure. (ii) They raise the possibility that an agenda setter could exert excessive influence on the outcome of elections. (iii) They raise the possibility that voters will vote strategically rather than sincerely and so the voting system will not operate on the true preferences of voters (Gibbard 1973, Satterthwaite 1975).

In response, we point out first that the existing formal results on QV that we are aware of concern binary public decisions, and in the presence of a fixed binary decision, majority voting does not face any of the problems associated with Arrow's theorem, agenda setting, or strategic voting.

Still it is important to address these concerns, not in general, but specifically as they bear on the *democratic legitimacy* of majority voting. That is, we do not attempt to defend (some variant of) majority voting in all respects, but its credentials with respect to democratic legitimacy, and our defense is very partial. While full consideration of this issue would require a great deal of space, we must be very brief.

First we address (i). If one were to conceive of the legitimacy requirement as that of public decisions reflecting the collective will of the people, then legitimacy may be hopeless when preferences and attitudes are such that there is no collective will. We have conceived of democratic legitimacy as requiring that everyone be consulted and have an equal opportunity for influence. This requirement makes sense even if there is no collective will. With respect to (iii), it may be that the possibility of or incentive for strategic voting undermines equality of opportunity for influence or one's ability to be consulted, but the case is not straightforward. If others have misrepresented their opinions, that does not prevent me from expressing mine, and the fact that I may have an incentive to misrepresent my opinion or even that I choose to do so does not mean that I am not consulted on equal terms. Moreover, even in the presence of incentives to misrepresent, in majority voting and its cousins, ex ante, no one individual seems to have more influence than any other.

We now come to agenda setting, which is a real concern. If the agenda is controlled by a powerful group to further its own interests, this really could undermine democratic legitimacy from our perspective. However to address this point, and to compare QV and majority voting in the presence of powerful agenda setters, we come back to a point raised in the previous section: We really do need to appeal to a much broader account of the social and political environment in which QV or majority voting would function.

4   Conclusion

In this paper, we have investigated the ethical merits of Quadratic Voting from two philosophical perspectives. First, we looked at QV from a utilitarian perspective. Next we looked

at QV from the perspective of democratic legitimacy. We came to similar conclusions in both cases, but there are also important differences. Both the utilitarian and the legitimacy perspectives suggest that QV is problematic when wealth is unequally distributed. The utilitarian perspective adds to this 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 from a utilitarian perspective. The comparison between QV and majority voting is ambiguous from a utilitarian perspective.

When assessing the democratic legitimacy of QV under conditions of inequality, the comparison was not ambiguous in the same way. Vote-buying mechanisms, unlike majority rule, come at a cost to the democratic legitimacy of voting institutions. At least, this comparison is straightforward and unambiguous subject to a given binary agenda and against a fixed background of other forms of political influence.

Which considerations do we think are most relevant to the assessment of QV against majority voting: the utilitarian considerations or those of democratic legitimacy? Both are important, but we tend to side with the latter. We believe that it is a central consideration that we all have equal *claims* on our society's major public decisions. In some sense these claims are analogous to property rights, although perhaps the analogy should not be taken too far. 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.

We close with a variant of QV that we think is promising, and merits further exploration: That is, 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.

A Appendix

Proof of Proposition 3

Using (7), a poor and rich agent's utility function can be written respectively as

$$U_P(x, w) \approx u_P x + g_P(w_P^e) + b_P(w - w_P^e) \tag{13}$$

$$U_R(x, w) \approx u_R x + g_R(w_R^e) + b_R(w - w_R^e), \tag{14}$$

where $b_R < b_P$, and in (13) w is appropriately close to w_P^e, while in (14) w is appropriately close to w_R^e. It follows that QV selects $x = 1$ (approximately) if and only if[9]

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

or equivalently, (approximately) if and only if

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

On the other hand, from a utilitarian point of view it is optimal to select $x = 1$ 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 choosing $x = 1$ is too high.

In contrast to QV, majority voting always selects $x = 1$.

Let us select t_0 in the proposition so that $t_0 = \frac{\alpha}{1-\alpha}|u_R|$. When $u_P < t_0$, then it is optimal to select $x = 0$, and indeed QV will select $x = 0$, while majority voting will select $x = 1$. Let us 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 $x = 1$. That is what majority voting does, while QV selects $x = 0$. Finally, when $u_P > t_1$, then both majority voting and QV select $x = 1$. □

References

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

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

---

[9]Recall that $u_R < 0 < u_P$.

Gibbard, A. (1973), 'Manipulation of voting schemes: a general result', *Econometrica* pp. 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.

Lalley, S. P. & Weyl, E. G. (2015*a*), 'Quadratic voting', *Unpublished Manscript, February 2015* .

Lalley, S. P. & Weyl, E. G. (2015*b*), 'Quadratic voting', *Unpublished Manscript, December 2015* .

Posner, E. A. & Weyl, E. G. (2015), 'Voting squared: quadratic voting in democratic politics', *Vand. L. Rev.* **68**, 441–500.

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 & Economics* **13**(2), 263–277.