Quadratic Voting *

Steven P. Lalley† E. Glen Weyl‡

February 2014

While democracy often leads to the tyranny of the majority, alternatives that have been proposed by economists suffer even more severe problems related to multiple inefficient equilibria and budget balance. A simple mechanism, *Quadratic Voting* (QV), resolves all of these concerns, offering a practical and efficient alternative to one-man-one-vote. Voters making a binary decision purchase votes from a centralized clearing house paying the square of the number of votes purchased. Funds raised are returned to participants is an essentially arbitrary manner. We show that, under standard conditions, QV achieves full efficiency in large populations in any Bayesian equilibrium. Even when these conditions are relaxed in a variety of ways, QV achieves very high efficiency, especially compared to majority rule. The quadratic form is essentially unique because it is the only function with a linear derivative.

JEL Classifications: C78, D61, D71, H41

Keywords: social choice, market design, large markets, Bayesian mechanism design, vote buying

THIS MANUSCRIPT IS HIGHLY PRELIMINARY AND INCOMPLETE. DO YOU NOT CIRCULATE IT OR CITE IT WITHOUT THE AUTHORS' PERMISSION. ANY RESULTS NOT FORMALLY PROVEN IN THIS DRAFT SHOULD BE CONSIDERED AS CONJECTURES, SOMETIMES BASED ON APPROXIMATIONS, NOT FINAL PROVEN RESULTS. RESULTS MISSING EVEN FORMAL STATEMENTS ARE NOT YET EVEN FORMAL CONJECTURES.

---

*This paper replaces a previous working paper titled "Quadratic Vote Buying" sole-authored by Weyl. We are grateful to David Ahn, Drew Fudenberg, Jacob Goeree, Jerry Green, Alisha Holland, Scott Duke Kominers, Steven Levitt, Paul Milgrom, Stephen Morris, David Myatt, Roger Myerson, Michael Ostrovsky, Eric Posner, Azeem Shaikh, Sang-Seung Yi and Richard Zeckhauser, as well as seminar and conference participants at the École Polytechnique, Seoul National University, the Toulouse School of Economics, the Universidad de los Andes, the University of Chicago, the University of Washington at St. Louis, the Wisdom and Public Policy Research Conference for helpful comments and for the financial support of the Institut D'Économie Industrielle and the Social Sciences Division at the University of Chicago received by Weyl. Conor Mahany, Kevin Qian, Tim Rudnicki, Daichi Ueda, Joshua Bosshardt and especially Albert Tan and Arvid Viaene supplied excellent research assistance. All errors are our own.

†Department of Statistics, University of Chicago: 5734 S. University Avenue, Chicago, IL 60637; lalley@galton.uchicago.edu, http://www.stat.uchicago.edu/~lalley/.

‡Department of Economics, University of Chicago: 1126 E. 59th Street, Chicago, IL 60637 and Toulouse School of Economics, Manufacture de Tabacs, 21 allées de Brienne, 31000 Toulouse, France; weyl@uchicago.edu, http://www.glenweyl.com.

(D)emocracy is the worst form of government, except all the others that have been tried.

– Sir Winston S. Churchill

1 Introduction

Prohibitions on gay marriage seem destined to be remembered as a classic example of the "tyranny of the majority" that has plauged democracy since the ancient world. While in many countries an increasingly narrow majority of voters oppose the practice, the value it brings to those directly affected seems likely to be an order of magnitude larger than the costs accruing to those opposed, as we discuss in greater detail in Section 2. Majority rule, however, offers no opportunity to express intensity of preference, allowing such inefficient policies to persist. While most developed countries have institutions, such as independent judiciaries and log-rolling, that help protect minorities, these are often slow, insufficient and plauged with their own inefficiencies. In this paper we propose and analyze a simple, efficient and robust mechanism, *Quadratic Voting* (QV), that could break Churchill's pessimistic conclusion about collective decision-making.

Voters use a quasi-linear currency to purchase votes for or against a proposed changed from the status quo from a planner (with arbitrarily many votes to supply) and pay the square of the number of votes they purchase. Funds raised are distributed back to the voters in essentially any way desired.[1] In Section 4 we show that, under independent private values the drawn from any distribution with a finite mean and values on both sides of the origin, a Bayesian equilibrium exists regardless of the number of voters. More importantly, *in any such equilibrium* the inefficiency of the mechanism relative to the welfare-maximizing choice dies off at a rate that depends on the tail properties of the distribution. If distribution has thin tails (all moments finite), then inefficiency dies at rate $1/N$, where N is the number of voters; if it has a tail matching that of the US income distribution, inefficiency dies at rate $1/\sqrt{N}$. In either case inefficiency is quantitatively small for reasonable population sizes FILL IN.

QV builds on the "Quadratic Mechanism" of Groves and Ledyard (1977*a*) for the provision of continuous public goods under complete information in Walrasian econoimes.[2] However, to our knowledge, QV for discrete collective decisions and our game theoretic, incomplete information, finite population analysis are novel to this paper.[3] In particular while under complete information essentially any convex cost function results in efficiency (Maskin, 1999), under imperfect information

[1]The only constraint is that each voter must receive back the same fraction of the funds paid in *directly by her*.

[2]As a matter of intellectual history, Weyl arrived at the idea of QV independently of Groves and Ledyard (1977*a*), much as Clarke (1971) and Groves (1973) arrived at the mechanism of Vickrey (1961) independently; Weyl became aware of the Groves and Ledyard (1977*a*) paper after circulating a first draft of this paper.

[3]Goeree and Zhang (2013) simultaneously and independently proposed a similar, but distinct, mechanism that we discuss in greater detail in Subsection 7.3. However like Groves and Ledyard (1977*a*) and Hylland and Zeckhauser (1980) their analysis only considers the Walrasian limit and not convergence towards it.

the quadratic form is crucial, as Hylland and Zeckhauser (1980) showed in unpublished work.[4] Adapting their continuous logic to our binary setting, in a large market, where no voter has a significant impact on the chance of a tie occurring, the marginal benefit of a vote is proportional to a voter's value. For voters to purchase votes proportional to their values, thereby insuring efficiency, their marginal cost of purchasing votes must therefore be proportional to the number of votes purchased, as voters equate the marginal benefit of a vote to its marginal cost. The square is the unique cost function with this property. As a result, we show in Section 5 that the only vote-buying rule that are robustly limit-efficient in the way that QV is have zero first derivative at the origin and a strictly positive, but finite, second derivative. That is they are equivalent to QV about the origin, which is the only relevant point in the limit.

While this essentially unique optimality of QV is a compelling motivation for considering it, for it to be a practical mechanism it must be robust to a much broader range of environments than the simple one in which our main results are established. Thus in Section 6 we relax a number of the assumptions underlying our basic analysis. We allow for various forms of collusion and other manipulations, imperfect knowledge of the value distribution by voters, imperfect voter rationality, common values and small populations. In all cases the efficiency of QV remains extremely high under reasonable conditions, sometimes converging even faster to perfect efficiency than in our baseline analysis. Most importantly, QV consistently outperforms standard majority rule.

This contrasts sharply with the other mechanisms economists have proposed for making binary collective decisions, as we discuss in Section 7. Perhaps the widely known proposal is the application proposed by Tideman and Tullock (1976) of the Vickrey (1961)-Clarke (1971)-Groves (1973) (VCG) mechanism. In this mechanism, voters report their valuations and, if and only if their report changes the outcome, pay the amount by which all other voters valued the outcome prior to their input. In addition to its efficient equilibrium, however, this mechanism has an infinite number of inefficient equilibria in which, for example, any two voters may receive their desired outcome at zero cost by both reporting values extreme enough to ensure that neither is unilaterally pivotal. In addition the mechanism often requires resources to be destroyed in order to operate and suffers from a number of other crippling flaws highlighted by Ausubel and Milgrom (2005) and Rothkopf (2007) that make VCG, in Rothkopf's words, "not practical". The Expected Externality mechanism of Arrow (1979) and d'Aspremont and Gérard-Varet (1979), and a recent variant proposed by Goeree and Zhang (2013), eliminate the budget surplus problem and somewhat mitigate the collusion challenge but require the planner, and all agents, to know the distribution of valuations. Other mechanisms(Hurwicz, 1977; Maskin, 1999; Walker, 1981; Crémer and McLean, 1988) rely even more heavily on details of the information structure and/or have a large multiplicity of equilibria.

[4]Again, Weyl was unaware of of the Hylland and Zeckhauser work, whose intuition is more closely related to that behind our work than is Groves and Ledyard's, until the first version of this paper was circulated. Weyl is now collaborating with Hylland and Zeckhauser to finally publish their paper; more details on this collaboration are discussed below, especially in our concluding Section 9.

In addition to our formal analyses, we have more informally investigated the robustness of QV along numerous other dimensions. The surprising resilience of this simple mechanism, which, unlike the others discussed above, is not "fitted" to any specific context, has persuaded us that it has great promise for practical applications. We therefore, in Section 8, briefly discuss some of the applications that we and others are developing in other work. From the most easily-implemented to the most ambitious these are committee decision-making with an artificial currency, corporate governance, the assembly of complementary goods subject to holdout problems, international governance institutions, representative bodies, referenda (like those on gay marriage) and elections of representatives.

In Section 9 we conclude by discussing directions we are taking and that others might take to extend and apply the analysis here. Most proofs and some formal development of the robustness of QV are contained in appendices that follow the main text. While most of the paper is intended to be accessible to a broad audience of economists, Section 4-6 (with the exception of Subsection 5.1) are more technical in nature and can be skipped by readers with a more applied interest. Conversely Sections 2 and 8 contain little formal content and are primarily intended for motivation; thus more theoretically-inclined readers may wish to skip these sections.

2 Motivation

According to the Census, in 2010 lesbian, gay, bisexual and transgender (LGBT) voters constituted approximately 4% of the population of California and voters in same-sex couple households constitute approximately .7% of California's population. Given that, according to a survey by The Wedding Report, the average wedding alone costs more than $25,000 and LGBT couples are on average wealthier than non-LGBT couples[5], it seems reasonable to suppose that the benefit of marriage to the same-sex household couples is at least $100,000 per voter and given both the option value and dignity concerns it seems likely that the option to marry is worth at least $20,000 to other LGBT voters. This implies a per-capita (across the whole population) willingness-to-pay to see Proposition 8 defeated by LGBT voters of $1360.

Assuming that LGBT voters voted on California's 2008 Proposition 8, which banned gay marriage, at the same rates as other voters in the population and that all LGBT voters opposed the measure, the measure's passage by 52% to 48% implies that, among the 96% of non-LGBT voters, the measure was supported 52% to 44%. Assuming that, on average, these not-directly-affected voters had on average the same willingness to pay (for ideological or ethical reasons) to see the initiative go their way and that this was no greater than $5000 on average (which seems quite high) the average willingness-to-pay resulting from non-LGBT voters is $400 in favor of Proposition 8.

[5]See Experian Marketing's 2013 Lesbian, Gay, Bisexual, Transgender Demographic Report.

Thus, unless this calculation is significantly off, Proposition 8 seems a clear example of Pareto-inefficient tyranny of the majority. If a proposition involving appropriate transfers could have been arranged that committed the state never to ban gay marriage, it likely would have received overwhelming support. However, arranging such transfers to achieve Pareto-improvements, beginning from Pareto-inefficient allocations, is typically infeasible in large-scale political contexts both because of the incentives they create for rent-seeking (Coate and Morris, 1995) (viz. passing oppressive measures just to be paid off) and because of incomplete information (Mailath and Postelwaite, 1990) (viz. individuals claiming to oppose gay marriage just to receive a payment). Nonetheless, Proposition 8 seems likely to have not just been unjust in the sense many of have claimed but also inefficient in the standard utilitarian sense and a system that would systematically avoid such tyranny of the majority would likely be Pareto-improving, or nearly so, by a substantial amount.

Of course, the fact that majority rule and other social institutions fail to accurately incorporate the intensity of preferences means that a calculation like that above is inherently guesswork. Even if its conclusions are incorrect, it makes clear that replacing majority rule with a system *that is capable of accurately reflecting the intensity of voter preferences* has the potential to bring large aggregate welfare gains. In fact Proposition 8 is merely one small, if salient, example of problem of tyranny of the majority about which political theorists have been concerned at least since the time of Aristotle (c. 350 B.C.E.*b*). Other examples arise in nearly every walk of life, from the trivial to the epochal:

- Voters voting on a meeting time end up picking a date slightly preferred by a majority that excludes most members of the minority from being able to attend at all.
- A promising recruitment candidate with a relatively narrow focus is rejected because they cannot muster a majority in a diverse economics department.
- Latin American polities elect redistributive populist governments that wreck their economies and Middle Eastern polities elect divisive sectarian governments that lead to coups and civil wars.
- International organizations ride roughshod over the sovereignty of small member states on issues central to their national interest, while swamping great powers in endless coalition building to pass needed measures.

The best-functioning organizations typically have mechanisms in place designed to make such inefficient outcomes less likely, such as log-rolling, favor-trading, lobbying, absolute protections of minority rights, etc. These checks and balances, however, are both often insufficient and carry with them inefficiencies in the form of governmental paralysis and corruption that are all too familiar. A practical formal mechanism that can, as the market economy does for private goods, facilitate efficient trade on collective decisions, is therefore badly needed.

3 Mechanism, Model and Intuition

We now formally define and analyze our proposed mechanism. To emphasize the independence of the mechanism from the details of the modeling environment, we define it prior to most of the details of the model we use to study it.

3.1 The Quadratic Voting mechanism

N voters $i = 1, \ldots, N$ must make a binary collective decision that impacts all of their welfare of whether to stick with the status quo or adopt and alternative. Every voter has some amount of currency (possibly unlimited). Under the Quadratic Voting (QV) mechanism, each voter may purchase votes $v_i \in \mathbb{R}$ for which she pays the square of the votes she purchases v_i^2 out of her currency. The alternative is adopted if and only if $V \equiv \sum_i v_i > 0$.

For definiteness, we assume that each voter receives $\frac{1}{N-1}$ of the revenues paid in by all other voters and none of the revenue collected directly from him; however none of the results below depend on this particular redistributive scheme. Any rule in which all revenues are returned and each voter receives the same share of the revenues she herself pays suffices to establish essentially all results that follow. Thus voters pay out, on net, $v_i^2 - \frac{1}{N-1} \sum_{j \neq i} v_j^2$.

Remark 1. QV is budget balanced as

$$\sum_i \left(v_i^2 - \frac{1}{N-1} \sum_{j \neq i} v_j^2 \right) = \sum_i v_i^2 - \frac{N-1}{N-1} \sum_j v_j^2 = 0.$$

3.2 Baseline model

To analyze QV we now make the standard assumptions of the expected utility-maximizing, independent private values environment with quasilinear utility that is used so frequently in mechanism design theory. Some of these assumptions are relaxed in Section 6.

Each voter has a value, $u_i \in \mathbb{R}$, measured in units of the currency, for the alternative realizing. Voters are expected wealth maximizers and have unlimited units of the currency. Voter i thus chooses her v_i to maximize

$$\mathbb{E}_i \left[u_i 1_{V_{-i}+v_i>0} - v_i^2 \right].$$

where $V_{-i} \equiv \sum_{j \neq i} v_i$. Values are drawn independently and identically across voters from an atomless distribution F with a bounded density f and support on $(\underline{u}, \overline{u}) \subseteq \mathbb{R}$, where $\underline{u} < 0 < \overline{u}$. We assume f has either all moments finite or finite first and second moments and a regularly varying (viz. Pareto) tails. This covers essentially every distribution we are aware of that has been used in economic analysis of preferences or income. We denote the mean of f as μ and its standard deviation as σ^2. If it has Pareto tails we refer to its lower Pareto tail as α_- and the constant on the Pareto tail as k_- and its upper Pareto tail as α_+ and constant k_+; if it has all moments finite we say $\alpha_- = \alpha_+ = \infty$. Our assumption of finite first two moments implies $\alpha_-, \alpha_+ > 2$. If it is has exponentially dying tails on one or both sides we denote α on that side as ∞.

We define the welfare achieved as $\frac{U \cdot 1_{V>0} - U \cdot 1_{V \leq 0}}{2\left(U \cdot 1_{U>0} - U \cdot 1_{U \leq 0}\right)} + \frac{1}{2}$, where $U \equiv \sum_i u_i$ and the expected welfare, $\mathbb{E}[W]$, as the ex-ante expectation of this. Note that (expected) welfare is always between 0 and 1, assuming U is not (identically) 0; 0 arises from (always) making the wrong decision and 1 arises from (always) making the right decision. Expected inefficiency (EI), $\mathbb{E}[I]$, is $1 - \mathbb{E}[W]$.

3.3 "Perfectly competitive" analysis

Consider a "type-symmetric" equilibrium in which each voter uses the same function v to map from her type u to the number of votes she buys. Then, at this equilibrium, the distribution, call it P, of the sum of votes of $N-1$ voters' votes, from the perspective of the remaining voter, is the same for all voters. Thus a voter with utility u maximizes, over her choice of v, $u\left[1 - P(-v)\right] - v^2$ as $1 - P(-v)$ is the chance of the alternative being adopted. Assuming P's density exists, $p(-v)$ is the density of an individual being pivotal by buying an additional vote, which we refer to as the *density of pivotality*.

As we argue extensively in the next section, in any equilibrium for a large population the dependence of the density of pivotality on v, at least over the range of values of v that all but a tiny measure of voters would consider buying, is very small. Intuitively, in a large population no individual has a significant influence at equilibrium on the chance of an election being tied as she is a small part of the aggregate votes being purchased. Otherwise she would have an incentive to act cheaply as a dictator, which could not be an equilibrium as other individuals would have an incentive to do the same. Effectively the difference between Proposition 8 passing or failing by 400 votes is very small and few voters would consider buying more than 400 votes. Thus, as with a price in a market for private goods, the density of pivotality is something which any small individual in a large society may take as fixed and given by the invisible hand of the market. Thus, for the purposes of our intuitive analysis in this subsection we take p as constant.

By taking the derivative of her utility, we obtain that an individual's marginal benefit from buying a small additional unit of vote is pu, the density of pivotality multiplied by her value. Her marginal cost of buying an additional vote is evidently $2v$. Equating these and noting that the cost of a vote purchase is convex while its benefit is linear, her optimal vote is given by

$$v(u) = \frac{p}{2}u. \tag{3.1}$$

The decision is thus made based on the sign of $V = \sum_i \frac{p}{2} u_i$ which is evidently the same as the sign of U. Thus, because each individual buys votes proportional to her value, the decision is always made based on the sign of the sum of values, maximizing social welfare. The reason is that the derivative of the square is linear in the number of votes and the marginal benefit of voting is linear in value because p is constant. As a result, under QV, a voter who intends only her own gain is led by an invisible hand to promote an end which was no part of her intention.

4 Main Results

The argument of the previous section relied on a number of unproven suppositions, namely the type-symmetry of equilibrium, the existence of a density of pivotality and its constancy just as the efficiency of the market economy for the provision of private goods depends on all participants facing the same prices, the existence of prices that clear the market and the constancy of prices as a function of the quantities chosen by participants. Neither of these sets of assumptions are satisfied exactly when the number of individuals is finite. However, recent work has shown that these assumptions are approximately satisfied in large, but finite, market mechanisms such as the double auction (Satterthwaite and Williams, 1989; Rustichini, Satterthwaite and Williams, 1994; Cripps and Swinkels, 2006; Azevedo, Weyl and White, 2013). In this section we show the same is true of QV, completing the analogy.

4.1 Existence

We begin by establishing the existence of a monotone, type-symmetric equilibrium in finite populations, before moving on in the next subsection to show that all equilibria in large markets behave approximately as desired and then using the approximations to calculate the rate at which efficiency obtains.

Lemma 1. *For any $N > 1$ there exists a type-symmetric Bayesian Nash Equilibrium v that is monotone increasing.*

While purely technical, this result reassures us that the substantive results in the following subsections are not vacuous. For the most part, the proof is just a standard combination of techniques from Reny (1999) and Reny (2011). However, two small innovations are worth noting. First, the existence of different types, rather than players, with conflicting preferences (with values of different signs) is used to ensure reciprocal upper semi-continuity. Second, Reny (2011)'s result is extended to unbounded value distributions whose first moments exist.

4.2 Characterization

While our proof is more general, we focus exposition on the case when the distribution has a bounded distribution; that is $\underline{u}$ and $\overline{u}$ are both finite. Equilibrium differs greatly depending on

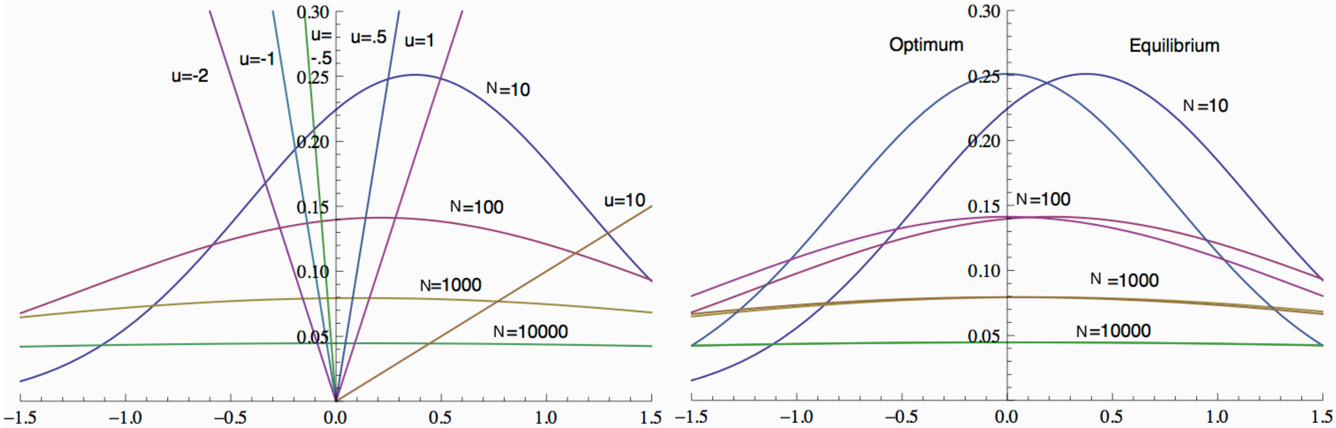

Figure 1: Approximate equilibrium behavior when $\mu = 0$, $\sigma = 1$ and $\mathbb{E}\left[u^3\right] = 0$. On the left curves represent $p(-v)/2$ for various values of N. Rays represent $\frac{v}{u}$ for various values u. Intersections are optima for individuals with the relevant values

whether $\mu = 0$ or $\mu \neq 0$ and we now characterize the nature of equilibrium in each case.

While non-generic, the case when $\mu = 0$ corresponds most closely to the simplest intuition for why p is approximately constant in the limit. Furthermore, despite its non-genericity, it may arise frequently in the equilibrium of a broader political game where candidates or candidate initiatives converge toward efficiency (Ledyard, 1984). For both reasons we begin by discussing this case. In particular, we show that any equilibrium has all individuals buying vote approximately (for large N) in proportion to their values, as we conjectured. As a result, p, the distribution of the sum of any $N-1$ votes, has a mean of approximately 0. As the number of voters grows larger, its standard deviation grows as well, though on the order of $\sqrt[4]{N}$ rather than $\sqrt{N}$ because, as it grows, the p declines and thus individuals buy fewer votes. Thus the standard $\sqrt{N}$ growth of the standard deviation is "split" (geometrically) between a decline in p and thus votes and a growth in the standard deviation. As a result each individual purchases votes that die at a rate $\frac{1}{\sqrt[4]{N}}$, for a cost that dies at a rate $\frac{1}{\sqrt{N}}$, leading to aggregate votes that grow at a rate $N^{\frac{3}{4}}$ and aggregate revenue that grows at a rate $\sqrt{N}$.

By the central limit theorem, p is approximately normal and as its standard deviation rises it becomes flatter and wider about its peak near 0. Furthermore individuals buy fewer votes as p diminishes and thus move less far along the distribution. Thus all individuals perceive approximately equal densities of pivotality, an approximation that grows increasingly accurate as N increases.

This behavior is pictured in Figure 1. This shows our computational solutions for our asymptotic approximate equilibria values of $p(-v)$ in a case when $\mu = 0$ for various values of N. Optimal votes-purchased-to-utility ratios for individuals with various values u are pictured based on the first-order condition, $p(-v) = \frac{v}{2u}$. The rays emanating from the origin are the value of the right-hand side of this first-order condition for various values of u. If p were identical across individuals, then all rays would intersect p at the same level of p. This is not exactly true, as we see: in particular, individuals with extreme values intersect further down the slope of the distribution than do individuals with modest values. As a result, they will be under-weighted, causing some inefficiency in finite populations, as we discuss in the next subsection. However, as N grows large every individual converges to purchasing according to approximately the same linear proportion (even in proportionate terms) of their value as p grows wide and flat, as formally stated in the following lemma.

Lemma 2. *When $\mu = 0$, in any equilibrium $v_i(u) = \frac{p_N}{2}\left(1 + \frac{\epsilon_i(u,N)}{N}\right) u$ where $p_N \equiv \frac{\sqrt{2}}{\sqrt{\sigma}\sqrt[4]{2\pi(N-1)}}$ and for all i, $|\epsilon_i(u,N)| < \epsilon(u)$ and $\epsilon(u)$ FILL IN.*

While this sort of equilibrium is perhaps the simplest way in which p can come to be approximately constant for all individuals, it cannot constitute an equilibrium when $\mu \neq 0$. To see this, suppose that $\mu > 0$ and all individuals purchased votes approximately in proportion to their. Then the mean of the distribution of votes would grow relative to its standard deviation as the former is proportional to N while the latter is proportional to $\sqrt{N}$. By standard large deviation arguments, this would imply an exponential decay of $p(0)$ in N and thus of the number of votes bought by any individual. Because exponential decay is much faster than linear, the total number of votes bought by all individuals would similarly decay essentially exponentially to 0, along with the the chance of the alternative not being adopted. But then any individual with negative value would have an incentive to buy enough votes to guarantee the alternative would not be adopted, as she could do this at very small cost in exchange for a change of probability of the outcome near unity. But clearly this would not involve vote purchases proportionate to value and thus would break the proposed equilibrium.

The actual equilibrium incorporates precisely this feature that breaks the other proposed equilibrium. Let us focus, without loss of generality, on the case when $\mu > 0$. Equilibrium strategies exhibit a discontinuity, with most types pursuing the "moderate" strategy of buying votes approximately proportional to their values and a small group of "extremists" with u near $\underline{u}$ buying enough votes to single-handedly bring the election close to a tie. These extremists are necessary to the equilibrium, as it is their behavior that increases the chance of a tie back to a non-trivial level, encouraging the moderates to buy a small, but not exponentially small, number of votes. In particular, in order for the extremists to be willing to buy "the whole vote", but for there not to be a growing number of such extremists, the mean of the number of votes purchased by $N-1$ individuals must asymptote to a constant and so the number of votes bought by any moderate must die at a rate $\frac{1}{N}$, thus her expenditure must die at rate $\frac{1}{N^2}$ and total expenditures by all moderates at rate $\frac{1}{N}$. In order for this to occur, the chance of a tie, which again is all driven by the event of an extremist existing, must die at rate $\frac{1}{N}$ and thus the probability that any given individual is an extremist, the measure of types that pursue this strategy, must die at rate $\frac{1}{N^2}$. Thus equilibrium occurs because

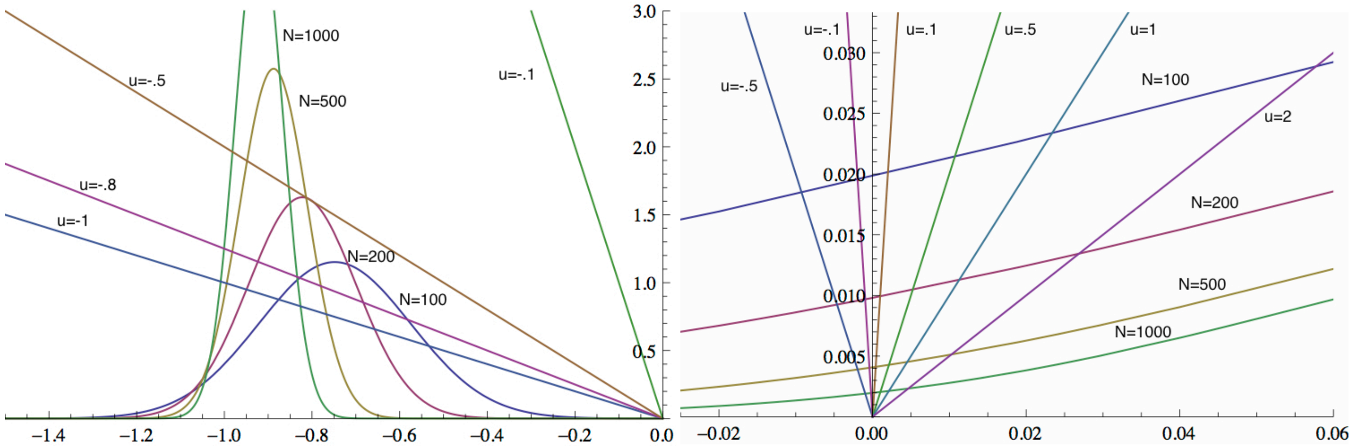

Figure 2: Approximate equilibrium behavior when u is uniform on $[-1, 2]$. On the left is the behavior of the extremists. The bell curves represent the approximate density of pivotality $p(-v)$ for various values of N. The rays represent, as in Figure 1, v/u for various various values of u. On the right is the behavior of the"moderates. The curves and rays are as on the left, but in a region local to the second, smaller peak resulting from the event that an extremist exists. On the right all intersections are optimal, but on the left there are often three intersections; the furthest left intersection is an optimum if and only if the area below the corresponding curve and above the ray is at least the area above the curve and below the ray. Otherwise the optimum is the intersection near the origin, as is the focus in the right panel

the moderates' fear of the extremists keeps them buying sufficient votes to deter all but the most diehard extremists from usurping the vote.

The behavior of the two groups is illustrated in Figure 2 in the case when u is uniform on $[-1, 2]$. The left panel shows $p(-v)$ over a relatively large range of v values. This looks normal with the peak growing sharper and a bit more extreme as N grows large. Individuals with positive or small negative values have rays that do not intersect the peak. However, the rays for larger negative values intersect the peak, and thus have three intersection points with p. The middle intersection is a local minimum, but the other two intersections are local maxima. Which is the global maximum depends on whether the area below the peak and above the ray is greater or less than the area beneath the ray and above the v-axis. If former area is greater, the optimal action is to act as an extremist, buying enough votes to ensure defeat of the alternative with high probability. This is true only for values of u very close to $\underline{u} = -1$, ensuring that only the most extreme individuals follow the extremist strategy.

The right panel shows matters from the perspective of moderate voters who do not "go over the peak". Their global optimum is the first intersection with the distribution, near a smaller local peak created by the possibility of the extremist existing near the origin. The small probability of the extremist existing pushes this peak down toward zero rapidly so that moderates span over a small part of the range. This reduces the effective slope for them and causes them to choose to buy votes approximately in proportion to their values. The slight, and asymptotically vanishing, bias favors positive votes as in the event of the extremist existing moderates who buy positive votes make a tie more likely.

In the case of an unbounded distribution, and particularly the case of Pareto tails, the basic structure of equilibrium is similar but the relevant computations are more involved. As the population grows larger, the threshold for extremist behavior gets pushed increasingly far into the tails of the distribution and, as a result, the probability of an extremist dies more slowly as a larger probability is required to maintain sufficient voting by moderates. The following lemma summarizes these results.

Lemma 3. *When $\mu > (<) 0$, in any equilibrium, if $u > (<) - sign(\mu)\sqrt{\frac{p_N}{2} N |\mu|}\left(1 + \frac{\nu_i(N)}{\sqrt{N}}\right)$, then $v_i(u) = \frac{p_N}{2}\left(1 + \frac{\sqrt{\log(N)}\epsilon_i(u,N)}{N^{\frac{\alpha+3}{2(\alpha+1)}}}\right) u$ where $p_N = 2\left(\frac{\alpha}{(\alpha+1)N^{2\alpha}(k\mu)^{2\alpha+1}}\right)^{\frac{1}{\alpha+1}}$, α and k are are the values for the sign opposite to that of μ, for all i, $|\epsilon_i(u,N)| < \epsilon(u)$, $\epsilon(u)$ FILL IN, $|\nu_i(u,N)| <$ FILL IN. Furthermore in any equilibrium if $u_i < (>) - sign(\mu)\sqrt{\frac{p_N}{2} N |\mu|}\left(1 + \frac{\nu_i(N)}{\sqrt{N}}\right)$ then $v_i(u) = -\frac{p_N}{2} N\mu\left(1 + \frac{\zeta i(N,u)}{\sqrt{N}}\right)$ where $|\zeta_i(u,N)| < \zeta(u)$ and $\zeta(u)$ FILL IN.*

Our logic above emphasizes why an equilibrium of the form we describe exists rather than directly establishing that *every* equilibrium must take this form. Much of the proofs in the appendices are concerned with this uniqueness. Our basic argument is simply that the distribution of sums of independent and not necessarily identically distributed random variables converge to a distribution that is, at least near its mean, approximately normal and approximately identical when the behavior of any individual is removed. This implies that optimal behavior for all agents must be approximately identical and that the set of such behaviors possible is of low dimension, determined only by the small set of parameters of such normal distributions. This small number of behaviors can be whittled, by contradictions, down to the cases we highlight above. These contradictions roughly take the form of our argument that no equilibrium in which all individuals buy votes approximately in proportion to their values is possible when $\mu \neq 0$.

Furthermore, within these behaviors, equilibrating forces imply uniqueness. In the $\mu = 0$ case as p_N increases votes increase, raising the standard deviation of aggregate votes and thus reducing p_N. In the $\mu \neq 0$ case if the number of extremists increases this raises the p_N, there by raising the number of votes purchased and reducing the number of individuals who wish to be extremists by making this behavior more costly.

4.3 Efficiency

The characterization of equilibrium in the previous subsection implies that limiting behavior corresponds to our basic argument in Subsection 3.3 and thus that in the limit full efficiency is obtained. However, it characterizes equilibrium behavior much more tightly than this and allows us to compute rates and constants of convergence, which constitute our main theorem.

Theorem 1. *If $\mu = 0$ then $\mathbb{E}[I] \leq \frac{A}{N}$. If $\mu \neq 0$ then $\mathbb{E}[I] \leq \frac{B}{N^{\frac{\alpha-1}{\alpha+1}}}$ where α is from the side of the distribution opposite to the sign of μ.*

The rate of convergence to efficiency thus depends, except in the case when $\mu = 0$, on the thickness of the tails of the distribution of values. When the distribution has thin or no tails ($\alpha = \infty$) then convergence is at the same $\frac{1}{N}$ rate when $\mu \neq 0$ as when $\mu = 0$. However, when the distribution has Pareto tails, convergence is slower. For example, when $\alpha = 3$, a high estimate for the top-end of the US income distribution (Diamond and Saez, 2011), convergence is at a rate of $\frac{1}{\sqrt{N}}$. If the distribution is extremely fat tailed, like the most extreme estimates of the top end of the income distribution in the United States of $\alpha = 1.5$, convergence is at a rate of $\frac{1}{\sqrt[5]{N}}$. Intuitively, the fatter the tails of the distribution of valuations the more extreme pool of individuals there exists as candidate extremists. Deterring these extremists requires individuals to buy a large number of votes, which can only occur when the extremists triumph and ruin efficiency with sufficiently large probability.

OTHER COMPARATIVE STATICS ON CONSTANTS

NUMERICAL RESULTS

5 Uniqueness

At first blush the square rule appears somewhat arbitrary. Wouldn't other convex rules work as well? It is well-known, for example, that in complete information environments where Groves and Ledyard (1977*a*) first proposed a Quadratic Mechanism (for continuous public goods), many other rules do work (Maskin, 1999).

We address these queries in two ways in this section. First, we consider a broader class of parametric rules that nests QV along with democracy and dictatorship and provide a simple analysis showing why QV is the only efficient rule. Second, we provide a more fundamental, axiomatic characterization showing that QV is essentially the unique optimal rule satisfying certain properties; unfortunately the proof and interpretation of this characterization is somewhat more technical.

5.1 Convex power voting

The most natural class of mechanisms that nests QV are those where votes may be bought at the cost of some power greater than unity of the number of votes purchased. With a few complications and limitations, the approach taken in the previous section can be used to validate the intuitive "perfectly competitive" analysis of Subsection 3.3 applied to these mechanisms as well. Thus in this section we analyze this broader class of mechanisms in this intuitive manner.

Under these broader regimes, the cost of vote purchase are $|v|^x$, where $x > 1$. The marginal benefit of buying an additional vote is again pu while its cost is now $\text{sign(v)}x|v|^{x-1}$. Again solving for the first-order condition for optimal vote purchases we obtain

$$v^{\star}(u) = \left(\frac{p}{x}\right)^{\frac{1}{x-1}} \operatorname{sign}(u)|u|^{\frac{1}{x-1}} = k|u|^{\frac{2-x}{x-1}}u,$$

where k is a constant across voters. Thus the decision is made based on the sign of $\sum_i |u|^{\frac{2-x}{x-1}}u$ rather than the sign of U. When $x = 2$ these are identical, but otherwise different individuals are weighted differently depending on the magnitude of their values. If $x < 2$ then $2 - x > 0$ and greater weight is placed on individuals with larger values: intensity of preference is over-weighted. If $x > 2$ then $2 - x < 0$ and greater weight is placed on individuals with smaller values: intensity of preference is under-weighted.

Two limits of this mechanism are particularly salient. In the limit as $x \to \infty$, $\frac{2-x}{x-1} \to -1$ and each individual has a weight inversely proportional to the magnitude of her value. Thus every individual has just one vote effectively and the decision is made by simple majority rule. That is, as $x \to \infty$, convex power vote buying converges to standard majority rule.

On the other hand, in the limit as $x \to 1$ the weight placed on an individual with a slightly higher magnitude value becomes arbitrarily larger than that placed on an individual with a slightly lower magnitude of value. Thus the weight placed on the single individual (assuming there is one) with the largest magnitude of value, the greatest intensity, becomes larger than the weight placed on all other individuals and that individual acts as dictator. Thus linear vote-buying yields dictatorship of the most intense (likely wealthiest), not efficiency. Similar result have been obtained through a variety of other concepts of equilibrium in the market for votes (Dekel, Jackson and Wolinsky, 2008; Casella, Llorente-Saguer and Palfrey, 2012). It is also consistent with a classical tradition dating at least to Olson (1965) that in resource-driven politics the most concentrated interests win. This may be an important reason for the hostility among the public to vote buying.

Intermediate values of $x \neq 2$ will resemble various forms of government discussed in political theory. Low values of x will have most weight placed on a small group of voters with very intense preferences; given that intensity is measured in dollars and the wealth distribution is likely to be more dispersed, especially at the top, than is the distribution of idiosyncratic preferences this will resemble an oligarchy. High values of x will over-weight individuals with weak preferences without being purely democratic and will thus resemble some form of republican government. Because Aristotle (c. 350 B.C.E.*b*) did not have a continuous spectrum along which these various forms of government he compared could be placed, it was not clear to him as it is through this analysis where his "golden mean" (Aristotle, c. 350 B.C.E.*a*) lies, though he clearly understood why other forms do not perform as well.[6]

[6]In fact, his logic in adjudicating between democracy and oligarchy in Book Six, Part III of *Politics* is similar to our logic that interests of all individuals should be linearly aggregated in a utilitarian fashion, and that incentives for information revelation must be correctly designed to ensure this is possible:

> *Democrats say that justice is that to which the majority agree, oligarchs that to which the wealthier class;*

5.2 Axiomatic characterization

While the set of convex power voting schemes is which to embed QV, clearly it is not exhaustive of mechanisms for binary collective decisions. As we know from a large literature briefly surveyed in Section 7 below, many other mechanisms exist that are efficient in the classical non-cooperative, indepdent private values model set up in Subsection 3.2. As we discuss in the next section, the most important thing that differentiates QV from these other mechanisms is its robustness to a range of environments that deviate from these simple assumptions by allowing features like collusion. There is, however, no guarantee that QV is the *only* mechanism that is limit-efficient in the most classical environment and also similarly robust nor do we know how to even state, much less prove, such a result.

It is therefore useful to complement these robustness properties by identifying a set of properties or axioms sufficient to uniquely characterize QV. Each of these properties is loosely associated with some of the robustness results, though no direct formal relationship exists. We now briefly discuss the set of axioms we have identified and then state the characterization of QV in terms fo the join of these axioms. While we reserve fully formal definitions of the class of mechanisms in which this characterization result applies to Appendix D, we assume (essentially without loss of generality) that individuals report to the mechanism some scalar real number r_i and make a net transfer $t_i(\boldsymbol{r})$ to the mechanism.

Axiom 1 (Detail-Freeness)**.** The rules of the mechanism are fixed independent of the value distriubtion.

One of the most fundamental desirable features of QV is that it does not rely on the planner having knowledge of the underlying distribution of valuations, which we call Detail-Freeness. It is this lack of detail-dependence that gives QV a reasonable prospect of being robust to alternative informational environments involving aggregate uncertainty or common values as discussed respectively in Subsection 6.2 and 6.4. However in the context of the standard models we have thus far been considering, this property has no meaning unless other claims are made about the mechanism

in their opinion the decision should be given according to the amount of property. In both principles there is some inequality and injustice. For if justice is the will of the few, any one person who has more wealth than all the rest of the rich put together, ought, upon the oligarchical principle, to have the sole power- but this would be tyranny; or if justice is the will of the majority, as I was before saying, they will unjustly confiscate the property of the wealthy minority. To find a principle of equality which they both agree we must inquire into their respective ideas of justice...For example, suppose that there are ten rich and twenty poor, and some measure is approved by six of the rich and is disapproved by fifteen of the poor, and the remaining four of the rich join with the party of the poor, and the remaining five of the poor with that of the rich; in such a case the will of those whose qualifications, when both sides are added up, are the greatest, should prevail...But, although it may be difficult in theory to know what is just and equal, the practical difficulty of inducing those to forbear who can, if they like, encroach, is far greater...

which hold over a range of distributions, because otherwise there would be no domain over which the rules would be held fixed. These claims are embodied in our second axiom.

Axiom 2 (Robust Limit Efficiency)**.** For *any* fixed distribution of values f satisfying the properties discussed in Subsection 3.2, the mechanism yields expected inefficency bounded above for all equilibria by a quantity that approaches 0 as $N \to \infty$. Furthermore this remains true even if value are scaled upwards by an arbitrary function $s(N) > 0$, that is the value distribution is not fixed as f but varies with N as $f_N(v) = f\left(\frac{v}{s(N)}\right)$ for a fixed f.

Section 4 show that QV is robustly efficient in large populations in the sense that its expected inefficiency is bounded above, in all equilibria, by a quantity that approaches 0 as N grows large *regardless of the distribution of values* so long as this distribution satisfies some basic regularity properties. If all values are scaled up by a constant that depends on N this makes no difference to this efficiency result. However, in large populations the density of pivotality declines so that if values are not scaled up then the votes bought by any individual dwindle and thus only the properties of a mechanism over a small number of votes matter. Thus if we were not to scale up values as N grows, limit efficiency would only require a quadratic shape *local to the origin* of 0 votes. However, if efficiency is required even if values are scaled up by an arbitrary factor of N, then shapes that fail to be quadratic globally will lead to inefficiency for the same reasons outlined in the previous subsection: they fail to have a derivative proportional to the number of votes purchased over the range individuals purchase votes. Thus we use this stronger notion of robust lmit efficiency to tie down QV precisely.

Another reason why efficiency under scaled-up distributions is desirable is that there are many reasons (such as imperfect voter rationality or expressive motivations for voting) to believe the number of votes individuals purchase may not dwindle to zero even for a fixed value distribution in practice as we discuss in Subsection 6.3. The global linear derivative property of QV is importance to its efficient performance in such settings.

Axiom 3 (Separability)**.** $t_i(\boldsymbol{r}) = t_i^i(r_i) + t_i^{-i}(\boldsymbol{r}_{-i})$. That is transfers can be additively separated into two components, one of which depends on an individual's actions and one of which depends only on the actions of other individuals.

A basic source of the extreme suceptibility of the VCG mechanism to trivial collusive manipulations is that the payments each individual makes, based on her actions, depend on the actions of others. This allows two individuals to each make extreme reports that insulate the other from making any payments based on these reports. A simple design feature detering such behavior is requiring that the payments each individual makes depending on her own report be additively separable from any payments that she makes or receives depending on others' reports. Many standard auction formats that are more resilient against collusion, such as the all-pay auction and (conditional on the allocation) the first-price auction, are separable. Thus separability plays a role in QV's relative resilience to collusion discussed in Subsection 6.1 below.

However, this connection is tenuous as the Expected Externality mechanism discussed in Subsection 7.3 below is also highly suceptible to collusion and is separable. We believe, but have not yet shown, that Axiom 3 can be replaced with a combination of ex-post symmetry (all individuals are ex-post treated symmetrically given their actions) and budget balance. The reason is that these properties rule out VCG or similar mechanisms that can only be robustly limit efficient and budget balanced through using ex-post asymmetric refund procedures (Green, Kohlberg and Laffont, 1976). However, we have not been able yet to rule out the existence of non-VCG-like schemes that are robustly limit efficient and budget balanced in an ex-post symmetric fashion.

Axiom 4 (Scale-invarance)**.** t_i^i is independent of N.

Transfer rules that asymptotically approach a quadratic rule will be just as robustly limit efficient as QV is regardless of their behavior in small populations. In fact, the rule can be arbitrary up to any finite N and be equivalent to QV above this N and still be robustly limit efficient. Thus to tie down QV precisely requires scale invariance. Obviously other arbitrary rules may perform very poorly in small (or even quite large but finite) populations, while QV performs quite well as discussed in Subsection 6.5 below. Scale invariance is thus related to QV's relative resilience to population sizes.

Theorem 2. *Any mechanism satisfying Axioms 1-4 is equivalent (up to renaming reports) to QV with some multiplier on the quadratic cost. More precisely, any mechanism satisfying these axioms has corresponding to it a collection of functions* $\{v_i(r)\}_{i=1}^N$ *such that* $t_i^i(r_i) = a_i + bv_i^2(r)$ *for some set of constants* $\{(a_i)\}_{i=1}^N$ *and a common-across-individuals constant* b *and the decision rule is* $A = 1$ *if* $\sum_i v_i(r) > 0$ *and* $A = 0$ *if* $\sum_i v_i(r) < 0$*. Furthermore QV satisfies the axioms.*

Axioms 1-4 thus define QV uniquely up to the arbitrary common-across-individuals cost of votes (per unit of vote squared) and the design of the budget-balancing transfers. In addition to uniquely defining QV up to these factors, it clarifies which aspects of QV are arbitrary (budget-balancing transfers and common multiplicative constants) and which are necessary (the quadratic shape and the common-across-individuals cost of voting).

6 Robustness

While QV is essentially uniquely limit-efficient in the class of mechanisms we discussed in the previous section, there are many other efficient mechanisms, as we discuss in the next section, that do not fall into this class. The fundamental virtue that therefore recommends QV to us is not its efficiency in the narrow settings we focus on above but rather the robustness of its efficiency to a wide range of changes in the economic environment. We have analyzed the mechanism in a range of environments with varying degrees of formalism to understand this robustness. In this section we report the results of our analysis in what we consider to be the five most important robustness checks. While this is still a small number compared to the wide range of concerns one could have about the performance of QV in practice, the robustness of QV across this range of settings, especially when compared to other mechanisms (see the next section), has lead us by Occam's Razor (Blumer, anad David Haussler and Warmuth, 1987) to the view that QV is a plausible, practical and robust mechanism for collective decision-making in a range of contexts.

6.1 Collusion and other manipulations

One of the most severe limitations of existing efficient mechanisms for binary collective decisions is their sensitivity to collusion. As we discuss in the next section, any two individuals can, at approximately or sometimes literally zero cost, obtain their desired outcome under the VCG and Expected Externality mechanisms if their behavior is properly coordinated. An important question therefore is whether QV is more robust to such coordination.

The first thing to note is that any efficacious collusive group, one that maximizes the joint payoffs of its members, will always have all members purchasing the same number of votes. The reason is that the utility of the group depends only on the aggregate number of votes it buys and the aggregate payments it makes. Conditional on the first, the second is minimized when all individuals split evenly the aggregate votes because the quadratic function is convex.

The second thing to note is that a collusive group of size m will buy m times as many votes per unit of aggregate utility as its individual members would buy if they had such a utility. This can be interpreted in two ways. In the first interpretation, the cost that they face for a marginal vote is lower than for an individual because the aggregate votes are spread out more thinly and thus run less quickly into the increasing marginal cost of votes. Alternatively, suppose that every individual in the group had the same utility. Each would create positive externalities on the others for each vote she purchased of the same magnitude of the value she obtains from a vote. Collusion internalizes these externalities, magnifying the optimal amount the group would vote.

These two different effects of collusion, leveling votes within the group and magnifying the votes it purchases, have very different impacts on efficiency. The first impacts revenue raised and redistribution, but will never bias the decision, and is greater to the extent that the vote purchases of individuals prior to collusion would have been very heterogeneous. The second may well bias the decision and occurs most strongly in the reverse circumstance: when the vote purchases of individuals are relatively homogeneous. Because this second form of collusion has greater potential harm, we focus on it below.

This sort of collusion is most harmful when it involves a large number with large values pointing in the same direction. The results we obtain therefore depend on which individuals are involved in collusion. We consider two different cases. In the first, "worst-case" scenario individuals are systematically the m most extreme individuals in some direction. In the second, "average case" scenario individuals colluding are drawn randomly from the same distribution as the whole population.

In either case, there are (at least) three basic challenges that limit the efficacy of such a collusive group. First, the group must be relatively large to have a significant impact as its optimal vote purchases are only magnified by a factor of the group size. Second, individual members of the group will face strong unilateral incentives for deviation. Third, collusion will meet with offsetting reactions by other non-colluding agents that may undermine the efficacy of the collusion and even make it counter-productive. We now informally discuss and state propositions formally establishing the limits each of these forces places on the possibility of collusion.

First, consider the necessary size of a coalition that can significantly impact efficiency. Suppose, just to chose one possible case, that $\mu = 0$, that the utility distribution is bounded and that the coalition is composed of the most extreme individuals in one direction. The magnitude of the total value of the group will be approximately proportional to its size m and they will thus optimally buy votes proportional to m^2. On the other hand, the total value of all other individuals in the aggregate is of size $\sqrt{N}$ as this is the order of magnitude of the standard deviation of the sum of mean-zero random variables. Thus as long as $m = O(\sqrt[4]{N})$ it is unlikely that this group will outweigh the direction of the preferences of other individuals. Similar calculations to these may be made for other scenarios and are summarized in the following proposition.

Proposition 1. *If there is a single perfect collusive group of size m and other individuals play as in equilibrium, $\mathbb{E}[I] \to 0$ as N grows large as long as*

1. *$\mu = 0$, colluders are drawn as in the average case and $m = O\left(\sqrt[3]{N}\right)$,*

2. *$\mu = 0$, colluders are drawn as in the worst case and $m = O\left(N^{\frac{\alpha-2}{2(2\alpha-1)}}\right)$,*

3. *$\mu \neq 0$, colluders are drawn as in the average case and $m = O\left(N^{\frac{4}{3(1+\alpha)}}\right)$ or*

4. *$\mu \neq 0$, colluders are drawn as in the worst case and $m = O\left(N^{\frac{\alpha-1}{(\alpha+1)(2\alpha+1)}}\right)$.*

Here α denotes the smaller of the two α values in 2) and that from the side opposite to μ in sign in 3) and 4).

These results show that in many cases successful collusion requires large coalitions that will be hard to form in the face of authorities attempting to detect collusion. This contrasts with the VCG and Expected Externality mechanism where a coalition of two can nearly costlessly achieve arbitrary efficiency.

However, these results are quite weak in some cases. When α is small, $\mu = 0$ and the coalition is drawn from the most extreme individuals or when $\mu \neq 0$ and α is large, even small coalitions are quite dangerous for efficiency. In these cases we must consider the other challenges to collusion discussed above.

As Theorem 1 shows, QV's efficiency occurs at *all* equilibria; thus its efficiency is "coalition-proof" in the sense of Bernheim, Peleg and Whinston (1987), again unlikely VCG. In fact, unilateral incentives for deviation (without any other equilibrating reactions) rule out most collusion under QV. The only collusion that is unilaterally incentive compatible (assuming all other, non-colluding individuals act as in unilateral equilibrium) is an "extremist conspiracy" of individuals seeking to overrule the will of the mean when $\mu \neq 0$. Such a conspiracy may, if sufficiently large, may be incentive compatible for the conspiring individuals because the fact that they expect the conspiracy to succeed raises for them the chance that the election will be tied and thus their unilaterally optimal vote purchases.

Proposition 2. *When $\mu = 0$ there are always unilateral deviation incentives for any collusive behavior for large N and the size of marginal deviation incentives is at least proportional to the deviation from unilateral behavior. When $\mu \neq 0$ a collusive agreement causing inefficiency with no unilateral deviation incentives is possible only when $m = \Omega\left(N^{\frac{1}{1+\alpha}}\right)$ in the worst case and $m = \Omega\left(N^{\frac{4}{1+\alpha}}\right)$ in the average case.*

Thus unilateral deviation incentives are likely to create important challenges to collusion, especially when $\mu = 0$, because such collusion actually makes a tie more likely and thus decreases the optimal unilateral vote purchases for each colluding individual. Such unilateral deviation incentives also provide some deterrence to collusion when $\mu \neq 0$ and α is not too large. However, when α is large and $\mu \neq 0$ in either the average or worst case collusion may be individually incentive compatible.

However, the mechanism through which it incentive compatible is that it raises the chance of extremist behavior and therefor of a tie. If this prospect were known not only by the colluding individuals but also by other individuals it would trigger them also to buy more votes, making collusion partly or wholly self-defeating. In particular, consider a game where the collusive group can act in perfect coordination but is known by all other voters to do so. In order to overrule the social interest, the group will have to have a total utility which, when multiplied by m, is on the order of $N\mu$ because everyone will anticipate their collusive actions and buy votes in the same proportion the collusive group does to their utility, except that the collusive group is able to by votes by a factor of m more cheaply than others. The conditions under which this can occur are summarized in the following proposition.

Proposition 3. *If there is known to be a single perfect collusive group of size m, in any equilibrium $\mathbb{E}[I] \to 0$ as N grows large as long as*

1. *$\mu = 0$, colluders are drawn as in the average case and* $m = O\left(\sqrt[3]{N}\right)$,

2. *$\mu = 0$, colluders are drawn as in the worst case and* $m = O\left(N^{\frac{\alpha-2}{2(2\alpha-1)}}\right)$,

3. *$\mu \neq 0$, colluders are drawn as in the average case and* $m = O\left(N^{\frac{2}{3}}\right)$,

4. *$\mu \neq 0$, colluders are drawn as in the worst case and* $m = O\left(N^{\frac{\alpha-1}{2\alpha-1}}\right)$.

Here α denotes the larger of the two α values in 2) and that from the side opposite to μ in sign in 3) and 4).

Note that in the remaining problematic case, when $\mu \neq 0$ and α is large, m must be of size $N^{\frac{2}{3}}$ or $\sqrt{N}$ depending on whether we are in the average or worst case for collusion to prevent limiting efficiency if the collusion is "detected" by other voters. Thus either private or public monitoring of collusion is likely to make it difficult for collusive agreements to succeed unless they are implausibly and dangerously large. We thus conclude that while collusion, if not adequately policed, may damage the efficiency of QV, it is unlikely to do so in a devastating manner in large populations.

A similar analysis may be applied to a single individual who fraudulently "de-mergers", representing herself as more than a single individual. Such de-merger attacks are known to be highly effective against VCG as discussed in the next section. Logic very similar to that above yields the following limits on the efficacy on such fraudulent activity.

Proposition 4. *If a single individual can fraudulently misrepresent herself as l individuals, in any equilibrium $\mathbb{E}[I] \to 0$ as N grows large as long as*

1. *$\mu = 0$, the fraudulent individual is drawn randomly, other individuals behave either as if they were aware or not aware of the fraudulent behavior and* $l = O\left(\sqrt{N}\right)$,

2. *$\mu = 0$, the fraudulent individual is drawn randomly, other individuals behave either as if they were aware or not aware of the fraudulent behavior and and* $l = O\left(N^{\frac{\alpha-1}{2\alpha}}\right)$,

3. *$\mu \neq 0$, the fraudulent individual is drawn randomly, other individuals behave as in the equilibrium without fraud and* $l = O\left(N^{\frac{2}{1+\alpha}}\right)$,

4. *$\mu \neq 0$, the fraudulent individual is the most extreme individual in the population, other individuals behave as in the equilibrium without fraud and* $l = O\left(N^{\frac{\alpha}{\alpha(\alpha+1)}}\right)$,

5. *$\mu \neq 0$, the fraudulent individual is drawn randomly, all other individuals are aware of the fraudulent behavior, an equilibrium is played given this common knowledge of fraud and* $l = O(N)$ *or*

6. *$\mu \neq 0$, the fraudulent individual is the most extreme individual in the population, all other individuals are aware of the fraudulent behavior, an equilibrium is played given this common knowledge of fraud and* $l = O\left(N^{\frac{\alpha-1}{\alpha}}\right)$.

Here α denotes the smaller of the two α values in 2) and that from the side opposite to μ in sign in 3) and 4).

Thus the number of identities that a perpetrator of a fraud would have have to take on to significantly impact efficiency is even larger than the size of an effective collusive group. Such large-scale fraud is likely to be detected and thus is unlikely to be a serious threat.

6.2 Aggregate uncertainty

Our baseline model assumes that the distribution of utilities is commonly known and thus that, except in the knife-edge case when $\mu = 0$, there is no uncertainty about the optimal action. This both seems unrealistic and makes the informational problem somewhat trivial in large populations.[7] It thus seems natural to consider how QV performs when the distribution of valuations is uncertain amongst the voters.

Consider the simplest possible case of aggregate uncertainty, when there is an unknown scalar parameter $\gamma \in \left(\underline{\gamma}, \overline{\gamma}\right) \subseteq \mathbb{R}$ that determines the density of valuations, $f(u|\gamma)$, has a prior density distribution g and and is affiliated with u, that is it orders f by first-order stochastic dominance (Milgrom, 1981; Milgrom and Weber, 1982). We maintain all of our assumptions on the distribution from above also assume that g is non-atomic and that our assumptions apply to the unconditional distribution of u. Assume that $\exists \gamma_-, \gamma_+ \in \left(\underline{\gamma}, \overline{\gamma}\right) : \mathbb{E}[u|\gamma_+] > 0 > \mathbb{E}[u|\gamma_-]$. The following lemma characterizes the nature of equilibrium in a large population in an intuitive way: there exists a threshold $\gamma^\star$ such that if $\gamma > \gamma^\star$ then the alternative is chosen with probability near 1 and if $\gamma < \gamma^\star$ then the alternative is chosen with probability near 0.

Lemma 4. *Under the assumptions of this section, there exists a unique $\gamma^\star \in \left(\underline{\gamma}, \overline{\gamma}\right)$ such that in any equilibrium as $N \to \infty$, $\mathbb{P}(V > 0|\gamma) \to 1$ if $\gamma > \gamma^\star$ and $\mathbb{P}(V > 0|\gamma) \to 0$ if $\gamma < \gamma^\star$.*

This lemma greatly simplifies the analysis of equilibrium for several reasons. First, note that there is also a unique $\gamma_0 : \mathbb{E}[u|\gamma_+] > 0 > \mathbb{E}[u|\gamma_-]$ whenever $\gamma_+ > 0 > \gamma_-$. As a result, perfect limiting efficiency is achieved if and only if $\gamma_0 = \gamma^\star$. Second, by the analysis of Good and Mayer (1975) and Chamberlain and Rothschild (1981), for large N all ties occur when γ is very close to $\gamma^\star$. This leads to a very simple description of equilibrium behavior.

[7] However, as McLean and Postelwaite (2013) argue, it may be the existence of an efficient mechanism given aggregate certainty that provides correct incentives for individuals to reveal their information to the group and thus create this aggregate certainty. Thus aggregate certainty may be the appropriate framework for analysis of a robust mechanism like QV even if it would admit other, non-robust mechanisms as described in Subsection 7.4 below.

Lemma 5. *At any equilibrium* $v_i(u) = \left[\frac{g(\gamma^\star|u)+\epsilon_i(u,N)}{N}\right] u$ *where* $|\epsilon_i(u,N)| \leq$*FILL IN.*

This characterization states that, in large populations individuals buy votes in proportion to the chance they perceive of $\gamma^\star$ realizing. This in turn leads to a simple integral equation for $\gamma^\star$:

$$\mathbb{E}\left[g\left(\gamma^\star|u\right)u|\gamma^\star\right] = 0. \tag{6.1}$$

We have not been able to derive from this fully general results about efficiency. However, we have studied several examples that admit an analytic solution of Equation 6.1; others can easily be studied by solving Equation 6.1 computationally. A common thread running throughout these analyses is the "Bayesian Underdog Effect" identified by Myatt (2012) in the context costly, but otherwise standard, voting. Suppose, without loss of generality, that $\mathbb{E}[u] > 0$ so that the alternative is the ex-ante "favorite" in welfare terms and that the status quo is the ex-ante "underdog". If efficiency were to result, that is if $\gamma^\star = \gamma_0$, then individuals with $u < 0$ would tend to put a higher probability on $\gamma^\star$ intuitively because their own utility is a poll of one person indicating a lower value of γ. Because the alternative is the favorite lowering γ increases the chance of a tie: Republicans believed that in 2012 a close election was more likely than did Democrats. This Bayesian Underdog Effect thus raises the votes of the ex-ante underdog and thus $\gamma^\star$, leading to inefficiency because there are some values of $\gamma \in (\gamma_0, \gamma^\star)$ when the favorite should win but the alternative does. We have not been able to identify general conditions under which this logic is valid as it is based on a frequentist intuition, while it is the Bayesian probability of $\gamma^\star$ that is relevant. However, it plays an important role in all of the examples we have explored.

In the following examples, efficiency and inefficiency are computed as Section 4, but with an additional average taken over all possible realizations of γ according to the measure over γ. All calculations underlying the following examples, and more details about them, appear in Appendix F.

Example 1. Suppose that u is equal to γ plus normally distributed noise with standard deviation σ_1^2 and that γ is normally distributed with mean μ and variance σ_2^2. Then majority-rules voting is limit-efficient and QV is not. $\frac{\sigma_1^2}{2\left(\sigma_1^2+\sigma_2^2\right)}\mu = \gamma^\star > \gamma_0 = 0$. For large N and a fixed σ_2^2 maximal limit-EI occurs as $\sigma_1^2 \to \infty$; globally maximal limit-EI occurs when $\frac{\mu}{\sigma_2^2} \approx \pm 1.6$ and equals approximately 2.2%. Typically it is much less; for example if $\sigma_1^2 \to \infty$ but $\frac{\mu}{\sigma_2}$ is less than 75 or greater than 3 inefficiency is below 1% and if $\sigma_1^2 = \sigma_2^2$ then inefficiency is always below .5%.

Because the normal distribution is symmetric, standard voting, which always selects the preference of the median voter, achieves perfect efficiency in this example, while QV is not perfectly efficient. However, even in the worst case, QV still achieves more than 97% efficiency; usually it does much better. We now consider an example based on a set-up proposed by Krishna and Morgan (2012) to study costly voting.

Example 2. Suppose that γ is the fraction of individuals who have positive value, but that the distribution of the magnitude of value conditional on its sign is fixed and commonly known. Let μ_+, μ_- be respectively the mean magnitude of values for those with positive and negative values respectively. Assume that γ follows a Beta distribution. Then EI is never greater than 5% for QV and it can be arbitrarily large for majority-rules voting. In the special case when γ has a uniform distribution, QV dominates voting, which may have EI as high as 25% while for QV it is never greater than 3%. For "most" parameter ranges QV appears to outperform majority-rules voting, often quite significantly.

Finally we consider an example similar to the previous one but calibrated to the evaluation of Proposition 8 in California discussed in Section 2 above.

Example 3. Suppose that 4% of the electorate is commonly known to oppose the alternative and is willing to pay on average \$34k to see it defeated. Suppose that the other 96% of the electorate is willing to pay on average \$5k to either support or oppose the alternative with the intensity of their values being independent of γ. γ is the fraction of the 96% that support the alternative and is assumed to have a Beta distribution with parameters set so that on the mean fraction of the population in favor of the alternative is 52%. Then QV is always superior to majority rule and the gap is larger the smaller is the standard deviation of the vote share. When the standard deviation of the population share supporting the alternative is 5 percentage points (well above the margin of error in most individual polls), QV has 4% EI and majority rule 47% EI. Even when the standard deviation is 20 percentage points QV achieves 1.2% EI while majority rule has 7%.

Thus QV appears to perform quite well in a reasonable large range of scenarios with aggregate uncertainty and "usually" outperforms majority rule quite significantly. However, it is certainly not perfectly efficient. Much of this small inefficiency, however, arises from the fact that we considered fairly large amounts of aggregate uncertainty relative to what seems plausible given the extensive polling that typically leads up to large collective decisions like elections. McLean and Postelwaite (2013) argue that it is generically possible, by giving small monetary or social reward for correct predictions, to incentivize sufficient information revelation about the aggregate state using such polling to effectively restore aggregate certainty in a large population.

To the extent that this is true, we may expect QV to be closer to perfect efficiency in practice than the analysis above indicates: people mostly form their views about the chance of a tie in real elections based on public data, not personal preference. This will tend to homogenize views about the chance of a poll and reduce the distortive Bayesian Underdog Effect. In a future draft we plan to prove a formal result about the size of a poll relative to the population size required to restore general limit efficiency.

6.3 Voter behavior

The famous "paradox of voting" (Downs, 1957) holds that in a large society voting is irrational if it requires even a small cost because the chances of being pivotal are miniscule. Yet in practice we observe a large fraction of the population turning out to vote. This suggests weaknesses in the simplistic instrumental models of voter behavior we have thus far employed. In this subsection we consider how QV would perform if individuals behaved according to models capable of explaining observed turnout.

A basic challenge in doing this is that many models that aim at explaining turnout aim only at that; they are not comprehensive models that make predictions about behavior conditional on turning out. It is thus hard to determine what implications they have for the operation of QV. Some formulations, such as voting to tell others that one has voted (DellaVigna et al., 2013), would imply that behavior conditional on voting will follow standard rational choice. In fact, significant evidence exists that, conditional on voting at all, voters *do* behave quite strategically and even fairly consistently with rational choice theory (Kawai and Watanabe, 2013; Spenkuch, 2013). Thus it may be that our preceding analysis is an accurate prediction of voting behavior in QV, given that we assumed universal turnout.

However, we consider the two models discussed in a survey on the paradox of voting by Blais (2000) that would clearly lead to different behavior even conditional on voting.[8] In the first model individuals overestimate the chance of their being pivotal. In the second, individuals gain some direct, "expressive" utility for each of their votes in addition to a chance of changing the outcome. Both of these mechanisms will tend to raise the number of aggregate votes purchased. However, they will not typically endanger the efficiency of QV; in fact they will typically make it more robust by deterring extremists and making the multiple of utility that individuals buy in votes more homogeneous than it is in our baseline analysis.

First, suppose that individuals misestimate $p(-v)$ as $e\left(p(-v), \epsilon\right)$ where e is smooth and $e, e_1, e_2(0) > 0 > e_{11}, e_{12}$ and, but that individuals never take a strictly dominated actions. ϵ is a random variable drawn from $(\underline{\epsilon}, \overline{\epsilon}) \subseteq \mathbb{R}$ according to a smooth distribution h that is independent of u and embodies the extent to which individuals over-estimate the chance of being pivotal. Our assumptions on e ensure that over-estimation is greater the small is the chance of being pivotal; individuals may even underestimate the chance when it is very large.[9] This is consistent with experimental evidence

[8]There are many other potential models of voter behavior that may have extrapolations to behavior in QV, such as others surveyed by Dhillon and Peralta (2002). However, we could not figure out a natural way to operationalize these other theories and suspected they would yield similar results. Analysis of QV under these other theories would be an interesting direction for future research.

[9]An alternative model that we have also considered by do not report here for brevity is one where individuals overestimate the chance of their being pivotal unless they pay a cost to obtain a better estimate. In this case QV behaves more like majority rule, thus losing some of its efficiency benefits over majority rule. However, it may perform better for finite populations as this is the case that most effectively deters extremists and it always continues to outperform majority rule, at least if the costs of acquiring information about p are excluded. If these are included,

reported by Blais. Our assumption of independence of ϵ from u ensures that no type of individual systematically over-estimates more than others *conditional on the true chance of their being pivotal with a marginal vote*. However, individuals with endogenously different chances of being pivotal (because of the number of votes their value induces them to buy) may over-estimate to differing extents.

FILL IN.

Second, suppose that, in addition to the instrumental utility each individual earns, each also receives a benefit $\left(x(\epsilon) + \frac{\xi(\epsilon)}{\sum_i |v_i|}\right) uv$ where x, ξ are smooth and $x, \xi, x', \xi' > 0$. ϵ is again an independent-of-u random variable. x represents a per-unit-of-value expressive utility for each vote she purchases and ξ is a per-unit-of-value expressive utility she earns from the *fraction of total votes cast* that she represents. These two possibilities correspond to two different interpretations of expressive utility in the literature. The first corresponds to traditional expressivist accounts, such as that of Fiorina (1976), where expression creates a personal psychological benefit for the voter. The second corresponds to a more semi-instrumental motive, suggested by Myerson (2000), where voters vote to influence perceptions of political support, assuming only aggregate vote shares are reported by the media. As we will see, which force operates is irrelevant from the perspective of our analysis.

FILL IN.

Thus neither overestimation of the chance of being pivotal nor expressive utility is likely to reduce, and may even enhance, the efficiency of QV. We thus do not think that non-instrumental voter behavior is a significant threat to QV's efficiency. Yet it is also not clear that the sort of voter behavior that exists under current institutions would carry over to a society that implemented QV. The sense of civic duty, expressive value, signaling, etc. that supports voting under present institutions is arguably an outgrowth of the historical development of democratic social institutions (Lipset, 1960) that were geared towards ensuring the implementation of majority rule despite the fact that under voluntary voting formal democratic institutions can implement very different outcomes.[10] A society that adopted QV might evolve political institutions and values that were more similar to the instrumental and individualistic values that developed along side market economies (Greif, 1994; Bruni and Sugden, 2013), as such values would be conducive to the success and spirit of QV .

6.4 Common values

In our analysis above we assumed that elections served to aggregate preferences. Yet in his pioneering work on the aggregation of preferences, de Condorcet (1785) argued that voting also works

majority rule may perform better.

[10] In fact, under some conditions voluntary, costly voting implements more efficient outcomes than if everyone votes, as the cost screens intensity of preference (Ledyard, 1984; Borgers, 2004). See Subsection 7.5 for more detail.

to aggregate information, even when preferences are known to be aligned across individuals. In this subsection we consider how effective QV is at aggregating information, instead of or along side preferences.

We begin by considering the simplest information aggregation setting, that of pure common interest considered by de Condorcet. McLennan (1998) shows that the optimal information aggregation procedure using the actions available to agents is always an equilibrium. By essentially the same argument this is also true of QV. However, because QV allows expression of cardinal values, it allows the expression of strictly more information and thus, for generic information structures, achieves more efficient information aggregation in some equilibrium than voting does in any equilibrium. The one complicating factor is that under QV individuals must pay for their votes and thus interests are not perfectly aligned. However, anything that can be implemented using cardinal reports can be implemented through arbitrarily small cardinal reports that are therefore incentive compatible. Furthermore, the wider gradations of deviations possible under cardinal utilities makes it harder to sustain inefficient equilibria, implying that under fairly broad conditions the optimal first-best information aggregation is the unique equilibrium of QV even in finite populations.

More explicitly, suppose that there is a finite-dimensional parameter $\boldsymbol{\theta}$ drawn from a commonly known distribution smooth and non-atomic density g from a set $\left(\underline{\theta_1}, \overline{\theta_1}\right) \times \cdots \times \left(\underline{\theta_T}, \overline{\theta_T}\right) \subseteq \mathbb{R}^T$. Individual i has a utility function, $u(\boldsymbol{\theta})1_{V>0} - v_i^2 + \frac{1}{N-1}\sum_{j\neq i} v_j^2$; thus individuals have common preferences over the choice, though conflicting preferences over transfers for vote purchases. Individual i receives a signal $\boldsymbol{s}_i \in \left(\underline{s_1}, \overline{s_1}\right) \times \cdots \times \left(\underline{s_S}, \overline{s_S}\right) \subseteq \mathbb{R}^S$ drawn independently and identically according to a smooth and non-atomic density h conditional on $\boldsymbol{\theta}$.

Theorem 3. *Suppose that $T = S = 1$ and that $h(s|\theta)$ forms an exponential family with a single-dimensional sufficient statistic. Then there exists an equilibrium under QV in which $1_{V>0} = 1_{\mathbb{E}[u(\boldsymbol{\theta})|s_1,\ldots,s_N]>0}$ but no such equilibrium exists under majority rule. More broadly, there always exists an equilibrium under QV that outperforms the best equilibrium under majority rule in EI. Under FILL IN CONDITIONS, this best QV equilibrium is unique.*

As Feddersen and Pesendorfer (1997) argue, most collective decisions involve a mixture of conflicting preferences and dispersed information. They show that in large, majority-rule elections this mixture does not prevent information aggregation because a large number of individuals who constitute a small fraction of the population and are close to being indifferent conditional on information leading to an expected tie vote on the basis of their information. The fact that all information aggregation occurs through the votes of a narrow segment of the population, however, does put important limits on information aggregation. If, for example, all individuals have some minimum intensity of preferences, information does not aggregate and if those who are nearly indifferent also have very poor information, information aggregates very slowly. Under QV, by contrast, information aggregates by small adjustments to *all individuals' vote quantities* rather than large adjustments to a small fraction of the population's votes. This leads to information aggregation in settings where it does not under majority rule and faster aggregation even when it does, but very slowly, under majority rule, as the next two example illustrate.

Example 4. Suppose that each individual's value $v_i = \mu + \epsilon_i$ where μ is a common value component and ϵ_i is the individual's idiosyncratic preference. μ and ϵ_i are drawn identically and independently (in the latter case across individuals) from a distribution that equals -1 with .5 probability and 1 with .5 probability. Individuals receive signals s_i that are drawn independently and identically conditional on μ, taking on the same value as μ with probability $p \in (.5, 1)$ and $-\mu$ with probability $1 - p$. Then as $N \to \infty$ in any QV equilibrium the alternative is implemented if and only if $\mu = 1$, which is efficient. Under majority rule the alternative is implemented with probability .5 regardless of μ.

FILL IN EXAMPLES.

More generally, following theorem describes the conditions under and rates at which information aggregates with preference heterogeneity under QV.

Theorem 4. *FILL IN*

6.5 Small populations

All of the preceding analysis is relevant to large or very large populations. The efficiency bounds we provide for finite populations are very weak for small N. It is therefore unclear from the above analysis whether QV performs well, and better than standard mechanisms like majority rule, when N is small.

To address this concern, we computationally solved for an approximate equilibrium of QV using standard computational techniques for various value distributions as described in Appendix I.[11] The value distributions we considered were the Normal distribution and the Pearson Type I distribution, a generalization of the Beta distribution with two additional parameters representing the support of the distribution. We considered very small populations, with N ranging from 2 to 10. We summarize our qualitative findings here and discuss the quantitative results in Appendix I.

1. The worst relative performs of QV comes in cases with $N = 10$ when the value distributions have small variance and a mean and median of these same sign. These are cases when majority rule would and QV would yield the same efficient answer in large populations as "the majority is right". Majority rule appears to approximate this large population behavior more quickly than QV does and thus to outperform QV. Over all the specifications we tried the largest

[11] We never discovered a case when a given value distribution led to convergence to two different equilibria depending on initialization or other execution details. However it is possible the results below reflect only the behavior at one equilibrium and not at others.

gap was a case of the Pearson distribution where QV achieve EI of 15% and majority rule had an EI of about half that, 7.5%. Deviating from these conditions in any way led QV to outperform majority rule, as the next points discuss.

2. Even for the value distributions in the previous bullet point, when $N$was very small, QV outperformed majority rule. With $N = 2$, QV never had an EI above 12.5% for any specification we considered. Majority rule typically had EI between 15% and 25%. For the Pearson distribution giving the worst case result from 1), EI was 22% under majority rule. Under a Uniform distribution on $[-1, 1]$ QV had EI of 8% while majority rule had EI of 22%. Under a standard Normal distribution QV had EI indistinguishable from 0 while majority rule had EI of 24%. This appears to be the result of QV accounting for intensity as majority rule does not, given that majority rule simply flips a coin when a tie between the two individuals occurs, while QV typically decides in favor of the more intense individual.

3. Regardless of N when the distribution had mean 0, QV outperformed majority rule. In all cases we considered with a mean of 0 QV never had EI exceeding 7% and only exceeded 3% in two cases (Uniform distributions with $N < 5$). In these cases majority rule never had EI below 15% and it ranged between 15% and 25% fairly evenly, suffering from greater EI when$N$ was even and thus ties were possible than when N was odd and thus ties impossible. This appears to be because in these cases there is not particular tendency of the majority to be right or wrong and QV's accounting for intensity allows it to typically outperform somewhat arbitrary majority rule. Goeree and Zhang (2013) find similar and related results in a laboratory experiment using a variant on QV in this case discussed further in Subsection 7.3 below.

4. Even when, in large populations, the majority should be right as in 1), QV outperformed majority rule when the variance was sufficiently large relative to the mean. When values followed a Normal distribution with mean .2 and standard deviation 1.5, even in the most relatively unfavorable case ($N = 10$) for QV, QV had an EI of 7% while majority rule had an EI of 18%. When values followed a Pearson distribution with mean of approximately .059 and standard deviation of approximately .44 even in the most unfavorable case for QV ($N = 10$), QV had EI of 8% while majority rule had 19%. This appears to be because when the standard deviation is large relative to the mean, at least in sufficiently small populations, behavior is close to the case when the mean equals 0.

5. QV outperformed majority rule most starkly when the mean and the median of the distribution had opposite signs. In these cases QV never had EI above 5% while majority rule never hand EI below 14%. This smallest gap occurred when $N = 3$ and thus it was quite likely that the majority was right. For most population sizes the gap was significantly larger. This suggests

that the large population prediction of QV's solving the "tyranny of the majority" problem arising under majority rule hold in small populations as well.

The above results are strikingly consistent with the results we obtained under aggregate uncertainty in Subsection 6.2 above. QV certainly is not perfectly efficient, and may fall even farther short of perfect efficiency than it did there. In some cases, particularly when the majority is right in large populations, QV underperforms majority rule. But QV appears much more robustly efficient than majority rule and outperforms it in most cases. However, in cases where there is a strong prior reason to believe that the majority is right, majority rule appears to perform better in both contexts.

7 Relationship to Other Mechanisms

Thus QV appears not only, in some sense uniquely, efficient under standard conditions but also robust to a range of deviations from these standard conditions. We now compare these conclusions to the properties of other mechanisms economists have proposed; these comparisons are summarized in Table 1.

7.1 Voting

Because most of the paper has been devoted to comparing QV to majority voting, we only briefly summarize our findings here and discuss comparisons to super- and sub-majority voting. Like QV, voting, in all its forms, is simple, budget balanced, detail-free and in fact requires no transfers. Also, unlike QV, every individual has a simple dominant strategy to vote in favor of their preferred alternative. However, under the standard conditions discussed in Subsection 3.2 it may be extremely inefficient, achieving an expected efficiency approaching 0 in many cases. In small populations it nearly always continues to be dominated by QV. It is somewhat more robust to collusion than is QV along some dimensions (a colluding group may accomplish less unless it is even larger than needed under QV), but along others it is less robust (incentives for unilateral deviation are smaller and reactions by those outside the colluding group are non-existent). Under aggregate uncertainty majority rule may sometimes outperform QV but typically does far worse. When values are common and interests aligned, QV allows much more information to be communicated than does voting.

Ledyard and Palfrey (1994) and Ledyard and Palfrey (2002) show that in large populations if the distribution of valuations is known then by choosing the threshold for voting equal to the quantile corresponding to its mean, the limit-efficiency of voting may be restored. This mechanism suffers from the same limitations of the mechanisms, discussed in Subsections 7.3 and 7.4 below, that require the planner to know the distribution of valuations. For any fixed super- or sub-majority rule, voting performs worse than it would under a simple majority rule.

	Detail-dependence	Standard efficiency	Aggregate uncertainty	Strategic complexity	Collusion and de-mergers	Budget	Other problems
QV	Detail-free	Limit-efficient	Limit-efficient with large polls or if $\mathbb{E}[u] = 0$; inefficiency small otherwise	Non-direct	Moderately robust	Balanced	
Majority rule	Detail-free	Limit-efficient when mean and median have same definite sign, typically inefficient	Typically limit-inefficient	Strategy-proof	Moderately robust	Transfer-free	Inefficient with common values
VCG	Detail-free	Efficient	Efficient	Strategy-proof	Extremely sensitive	Surplus	Bankruptcy, uncertain payments
Expected Externality	Detail-based, requires aggregate certainty	Efficient	Ill-defined	Direct but Bayesian	Highly sensitive	Balanced	Ill-defined with common values

Table 1: Comparison of binary collective choice mechanisms

7.2 The Vickrey-Clarke-Groves mechanism

The most canonical mechanism that has been suggested by economists as an alternative to democratic decision-making is the Vickrey (1961)-Clarke (1971)-Groves (1973) (VCG) mechanism. First applied to discrete collective decisions by Tideman and Tullock (1976), under this mechanism individuals report their cardinal value for the alternative and the decision is chosen based on the sign of the sum of the reports. Any individual who is pivotal in the sense that, had she reported 0, the decision would have gone the other way pays the amount by which all those other than her preferred the decision she opposed. In addition to sharing with QV its detail-freeness, this system is appealing because if every individual plays the weakly-dominant strategy of reporting her valuation truthfully then, in an extremely broad range of circumstances, VCG implements the efficient outcome. VCG is fully efficient in a sense that is very robust to the information structure and the number of participating individuals, unlike QV.

Despite this, somewhat narrow but much remarked-upon (Bergemann and Morris, 2005; Chung and Ely, 2007) sense of robustness, the VCG mechanism has almost never been used for collective choices. The reason that VCG is "lovely, but lonely" (Ausubel and Milgrom, 2005) is that a number of other severe failures of robustness make it "not practical" (Rothkopf, 2007). These flaws were recognized by the originators of the mechanism from the start (Vickrey, 1961; Groves and Ledyard, 1977*b*), though their severity and implications for implementing VCG were not well-understood until more recently (Tideman and Tullock, 1977) when the first laboratory experiments based on the simplest Tideman and Tullock (1976) environments gave disastrous results (Attiyeh, Franciosi and Isaac, 2000).

Perhaps the most severe defect of the VCG is that, in addition to its efficient equilibria, VCG has a very large number of other equilibria, including, for any two individuals, equilibria where they attain their desired outcome and make no payments. In particular, any two individuals may announce sufficiently large values in the same direction so that neither is individually pivotal. Anticipating this, other individuals can do no better than to report 0. Similarly, any individual who can pretend to be two individuals can "break" the mechanism. Thus VCG is extremely sensitive to any deviation from the supposition that individuals will play their unilaterally weakly dominant strategies and experimental results, reviewed in the above-cited papers, confirm that in practice VCG rarely behaves this way.

VCG has many other problems as well, which are discussed extensively in the above-cited papers a few of which we briefly list here. Any revenue raised must be destroyed to avoid creating perverse incentives, which may be hard for the government to commit to; absent such a commitment, the scheme falls apart.[12] Even when such commitments are possible, the revenue that must be destroyed

[12] Some have suggested schemes to get around this problem in large populations. See, for example, the work of (Bailey, 1997). However these schemes also eliminate the perfect small population efficiency of VCG. Perhaps more importantly, these variants take an already complex and fragile system and make it more complex and more fragile

can often be greater than the improvement in efficiency over even simple mechanisms like majority rule (Groves and Ledyard, 1977*b*; Attiyeh, Franciosi and Isaac, 2000). VCG requires much larger liquidity among participants than does QV; individuals must have in cash the full magnitude of their value and place this into escrow when submitting their report. Unless this "bankruptcy" problem is addressed, individuals have an incentive to exaggerate their report and fall back on judgement-proofness if called upon to pay (which occurs with very small probability in a large population). Under QV payments are limited with probability 1 to a very small portion of underlying values and are certain conditional on the report made.

7.3 Expected Externality mechanism

The next-most canonical mechanism economists have suggested is the Expected Externality (EE) mechanism of Arrow (1979) and d'Aspremont and Gérard-Varet (1979), which was first applied to binary collective decisions by Goeree and Li (2008). This is similar to VCG, except that individuals pay the planner's ex-ante expectation of their VCG payments rather than their actual payments. Because individuals can thus not affect others' payments, the revenue raised may be refunded much as under QV, though, like QV, the mechanism is Bayesian rather than having a dominant-strategy equilibrium. QV was partly inspired by this mechanism as, in the case when $\mu = 0$, these EE payments are approximately quadratic in large populations. The intuition is much like that underlying the Dupuit (1844)-Harberger (1964) triangle as an approximation to the welfare loss from distortions. Because the distribution of the sum of other individuals' valuations are approximately uniform about 0 when $\mu = 0$ by the arguments in Subsection 4.2, both the probability of being pivotal and the average amount by which individuals are pivotal grow linearly with the individual's report. Thus EE payments grow like a Harberger triangle.

In fact, it was from this logic that Weyl originally derived QV. However, this is at most a starting point for QV, as when $\mu \neq 0$ EE payments are nothing like quadratic and in the richer information environments we consider they are not even well-defined. However, Goeree and Zhang (2013) take this logic more literally and propose a mechanism, limited to the case when $\mu = 0$, in which individuals pay the exact quadratic approximation to their EE payments. This quadratic schedule has a coefficient on it that depends on the number of individuals and the standard deviation of their value and thus, like EE, depends on the planner knowing the distribution of valuations. Thus both the EE and Goeree and Zhang mechanisms are not detail-free and only apply under aggregate certainty (and the later also requires $\mu = 0$). Additionally, when when $\mu \neq 0$, the EE mechanism suffers from essentially the same collusion problem as VCG in large populations, though this is less well-known.[13] Thus, outside the extremely non-generic case when it is common knowledge that

along other dimensions.

[13]Suppose two individuals report $-\frac{2\mu}{3}$. Each will make vanishingly small EE payments as the probability of either of these reports being pivotal is exponentially small in N by standard large deviation theory results. However,

$\mu = 0$, neither EE nor the Goeree and Zhang mechanism may plausibly be used.

7.4 Other mechanisms proposed by economists

The mechanisms discussed in the previous three subsections are those taken most seriously because other mechanisms proposed by economists are known to suffer many of the same weaknesses more severely or without as many corresponding benefits. However, for the sake of completeness, we briefly list a few other well-known proposals noting their defects:

1. *Implement the alternative if and only if* $\mu > 0$: In addition to requiring the planner to know the distribution of valuations, this suggestion places great and easily-abused power in the hands of the authority charged with determining the sign of μ (Maskin, 1999).

2. *Maskin (1999)'s mechanism*: In this mechanism, all individuals are asked to report μ and the alternative is implemented if and only if all report $\mu > 0$. If all do not agree, all agents are "killed" (pay a very large and inefficient penalty). In addition to requiring aggregate certainty, this mechanism is likely to be difficult to commit to and has a very large multiplicity of inefficient equilibria.

3. *Crémer and McLean (1988)-McAfee and Reny (1992)-style mechanisms:* These mechanisms are conceptually similar to Maskin's, but adapted to the context with aggregate uncertainty. Roughly, individuals are asked (via their report of their type) to guess other individuals' report of their type and are given large rewards for guessing correctly. Like the Maskin mechanism, this mechanism has a large multiplicity of equilibria (McLean and Postelwaite, 2013) and, perhaps more importantly, depends both on the mechanism designer having a very precise knowledge of the distribution of types and on individuals having preferences that are "appropriately" correlated with their beliefs about other individuals' types (Heifetz and Neeman, 2006).

A number of other mechanisms requiring individuals to have complete information have been proposed (Hurwicz, 1977; Walker, 1981). As a result, these mechanisms are even more fragile than are those discussed above (Bailey, 1994). Other mechanisms rely very sensitively on the absence of heterogeneity in risk attitudes and beliefs and on detailed knowledge by the planner (Thompson, 1966). All of these mechanisms are also quite complicated to explain and strain credibility along a variety of other practical dimensions. While there is not the space here to discuss all of these in detail, a large literature has established their impracticability; see Tideman (2006) for a detailed and excellent survey from which much of the discussion here is derived. An important reason, we believe, for their limitations, and for those of the VCG and EE mechanisms, is that they are tightly

together they will ensure the outcome goes in the inefficient direction.

"fitted" to the particular modeling environments they were designed to optimize within. Like a high degree polynomial fitted to a small number of data points, their ability to achieve good outcomes in tightly defined modeling environments is perhaps not surprising .

7.5 Mechanisms used in practice

What gives us confidence in the applicability of QV is not, or at least not only, its strong asymptotic performance under standard conditions. It is instead or in addition that an extremely simple mechanism, derivable in a variety of ways but not obviously related to any particular modeling, performs so well in a variety of environments. Continuing the analogy, we view QV as being like a linear (or perhaps quadratic) regression which turns out to have a very high degree of fit to a large number of data points, an event unlikely to happen by accident or because of over-fitting. Occam's Razor (Blumer, anad David Haussler and Warmuth, 1987) therefore suggests to us that QV is likely to be a useful mechanism in practice.

One measure of the simplicity that gives us this confidence is the ease of applying QV is that, like the double auction, it is straightforward to compare it to existing social institutions and look for institutions that may approximate it. One suggestion of an approximating institution arises from the work of Ledyard (1984); see also Myerson (2000) and Krishna and Morgan (2012). Ledyard considers a model where individuals have non-pecuniary costs of voting. If these costs are all strictly positive and follow a non-atomic distribution independent of voters' values, then the "representative voter" with a given value effectively faces a quadratic cost as a function of the fraction of cost types she represents who vote by the same argument underlying the analogy to Harberger's triangle in Subsection 7.3 in the limit as the population size grows large and thus both the density of pivotality and thus turn out grow very small. Thus non-pecuniary costs of voting may approximate QV and thus efficiency.

While there may be some truth in this argument, Ledyard argues that it is unlikely to provide a good approximation to reality as it requires turn out to approach zero in large populations, which is rarely observed in practice. If some voters have zero or negative non-pecuniary costs of voting or overestimate their chance of being pivotal to such an extent that turnout remains large in large elections, as much empirical research on voting suggests (Blais, 2000), the result clearly fails while we showed in Subsection 6.3 that QV remains as efficient or even may be more efficient than if voters are "standard". Thus costly voting does not seem a promising approximation to QV in practice.

More plausible are genuinely costly activities undertaken to influence collective decisions, such as log-rolling on committees and legislative bodies and lobbying, campaigning or out-right illegal vote-buying in elections. At first blush, most of these appear fairly linear and thus unlikely to well-approximate QV, though they may somewhat mitigate the tyranny of the majority problem of democracy. However, at least in some contexts, costs appear to be convex. In votes on committees it is usually easy for any individual to influence one vote of a colleague, harder to obtain the second, even harder to sway the third and so forth. This is shown clearly in the recent film *Lincoln* in the context of legislative log-rolling and vote-buying. Similar effects may occur campaigning, where it is typically much easier for individuals to influence their close friends or at least those they are acquainted with than to influence those further away from them. Exactly how close various contexts come to approximating QV is an interesting empirical question, discussed further in Section 9.

8 Applications

What is clear, however, is that none of the institutions discussed in the previous subsection approximates QV with much precision. The situation might be compared to the institutions of informal trading barter and reciprocity that are thought to have preceded more formalized exchange economies and to persist in less developed societies (Sahlins, 1972). While much of the famous substantivist-formalist debate concerns the degree to which such institutions approximated the market, there is near-consensus that the emergence of formalized market economies allowed exchange to occur more efficiently, between larger numbers of individuals, over a broader temporal and spatial extent and at lower cost (Greif, 1994). Similarly we believe that, to whatever extent some informal practical institutions resemble QV, its formalization is likely to improve welfare.

Therefore, in this section we discuss several contexts where we believe it could improve the allocation of public goods compared with existing mechanisms. We begin with the most modest applications that are likely to be feasible over the shortest time horizons and gradually build to those that, if implemented, could have the greatest impact on social welfare. Weyl and Eric Posner (see below) have founded a start-up venture, Collective Decision Engines, that is designing software to ease implementation in these applications. As demand for the various applications becomes clearer, tailored versions of the software fitted to particular applications will be developed. Posner and Weyl (In Preparation) will address wider range of concerns and questions related to these applications than is feasible here or in other academic work.

8.1 Private sector

The easiest, near-term application is likely that to the decision-making of committees that repeatedly interact such as recruitment committees in academic departments or honors-granting committees. Such groups would likely use artificial currency in lieu of actual money, which would allow Pareto improvements over majority rule internal to the decision process but would not permit global Kaldor-Hicks efficiency (Budish, 2011). We have had detailed discussions with a number of such committees and expect implementations in this setting to begin in several locations by year's end.

Somewhat more ambitious are applications to decision-making within small- and medium-sized start-up firms on issues such as appointment scheduling, product develop, amenities, etc. Such firms often make many decisions collectively and are open to experimentation with innovative means of organization. Either artificial or actual currency seems feasible in such contexts, depending on the firm culture. Weyl presented at one such firm, Applico (a mobile strategy firm with approximately 100 employees that Weyl sits on the board of directors of), and the firm plans to implement QV in several areas in the coming months.

More ambitious is the use of QV for the governance of public corporations, as advocated by Posner and Weyl (Forthcoming). Posner and Weyl propose an alternative implementation, *Square Root Voting* (*SRV*), in which investors obtain votes in proxy battles equal to the square root of the number of shares they own. Such an implementation is different more in design than in type from existing minority shareholder protections, such as poison pills, and thus has attracted some attention on Wall Street. Similar discussions are underway about the use of QV for corporate restructuring (Posner and Weyl, 2013). More serious discussions have taken place in South Korea, where concerns about tunneling by the families that control large the large "chaebol" conglomerates have become a focal point of national politics under the banner of "economic democracy" under which the new president was elected (Posner, Weyl and Yi, 2013). In fact, Weyl and Yi presented in front of the Action for Economic Democracy Caucus, a group of 11 members of the ruling Saenuri party, who expressed interest in legislation to amend the South Korean Commercial Code to allow the use of SRV for corporate governance.

8.2 Public sector

As important as such private applications are, the most important public goods are provided by various levels of the public sector. Probably the most plausible applications in the short term are ones where money already plays an important role, such as in the assembly of complementary goods that would otherwise be subject to holdout by coercive means (Kominers and Weyl, 2012*b*). Kominers and Weyl (2012*a*) advocate allowing any potential purchaser of a large number of complementary goods (such as land, spectrum, patents, etc.) to make an offer for the package of goods and having the current owners decide by QV whether to accept the offer, thus avoiding holdout. They show this procedure has many advantages in terms of ex-post efficiency, property right protection and fairness over traditional "eminent domain" procedures.

The next most modest such applications would be to decisions about local public goods and amenities, such as New York's bike share (Posner, 2013). State, local or national referenda, as we used to motivate the paper, are another promising, though much more ambitious, application. At least as important but perhaps somewhat more plausible in the medium term is the use of QV to formalize log-rolling in representative bodies. Rather than inefficient expenditures on pork-barrel projects being used to buy votes close to linearly, representatives could directly use funds from their district to (at quadratic cost) purchase influence on national legislation. Weyl, in collaboration with Hylland and Zeckhauser, plans to show that if representatives act to maximize the interests of those in their district, as they should if elected by QV, then the resulting decisions in a QV-governed representative body will be efficient as well.

Likely the most important application of such a representative property of QV would be to governance of international organizations, such as the European Union, the United Nations, the International Monetary Fund and the World Trade Organization. Such organizations have long been plagued by an inability to make decisions that respect national sovereignty when particular nation's vital interests are at stake while simultaneously allowing efficient decisions on issues that primarily impact great powers (Posner and Sykes, 2014). This has led to great disappointment about the capacity for large-scale international cooperation (Mearsheimer, 1994–1995). QV could offer a method for overcoming these stumbling blocks.

Of course prejudices in democratic culture and theory against the use of money will make such implementations an uphill battle. However, we believe that in the long-term such concerns can be overcome for several reasons:

1. Much of the resistance comes from linear vote buying, which we have shown is highly inefficient and dictatorial in effect. Once the public becomes acquainted with quadratic rules this resistance may fade.

2. It is not clear that QV would actually increase the eventual influence of the wealthy relative to the status quo, as presently wealth can buy significant influence close to linearly by advertising, lobbying and, in some countries with weak institutions, vote-buying campaigns designed to influence nearly indifferent swing voters. Such campaigns would have far less impact if votes were more costly and the difference between a small vote in favor and against was smaller, as it would be under QV. This process might more generally elevate the quality of public discourse.

3. It seems little more than prejudice that wealth has such a powerful influence over the allocation of private goods in most societies while not (at least formally) impacting the allocation of public goods. Allowing wealth to formally influence public goods allocation would increase the incentive to work, which could then be offset by increasing social insurance both directly through the redistribution of QV proceeds and by raising progressive taxes (Kaplow, 2002). Many philosophers have argued (Walzer, 1983; Satz, 2010; Sandel, 2012) that some private goods should be walled off from the influence of wealth; perhaps that is right. But we see no reason why the technological constraints of non-rivalry and non-excludability should imply that all public goods fall into this category while nearly all private goods should be governed by the market.

4. If the concern about the influence of wealth persists, versions of QV, using artificial currency or coefficients in front of the costs of votes that are proportional to income, can easily be used to remove concerns about undue influence of wealth while still improving efficiency. However, some efficiency gains in the exchange between public and private goods would obviously be lost.

In their collaboration, Hylland et al. plan to explore these arguments further, as well as the broader relationship between QV and ideas in democratic theory.

9 Conclusion

Economics as a discipline has typically been skeptical of the ability of political and public decisions to be made efficiently in the way many economists believe private decision can be. While economists disagree about the degree of market failures and thus the extent to which inherently inefficient collective action must be substituted for potentially inefficient private action, there is broad consensus that limits exist on the efficiency of allocation of public goods in practice that do not exist for private goods. This is reflected in formal results such as those of Arrow (1951), Gibbard (1973) and Satterthwaite (1975) that contrast with the fundamental welfare theorems and in informal attitudes in work such as Friedman (1962).

In this paper we have argued that this, if accurate regarding existing institutions, is only an artifact of those institutions. Public goods pose no fundamentally different challenges to those posed by private goods. In fact, the mechanism we propose, has a number of symmetries with market mechanisms for the allocation of private goods, such as the double auction. It is a Bayesian, separable, budget balanced that is simple to explain rather than a VCG mechanism which is dominant strategy, generally not budget balanced and opaque to most participants. It is perfectly efficient only with large numbers of voters, but out-performs natural alternatives when there are a small number of voters. Collusion is possible and potentially profitable, but requires a significant fraction of the market to participate in order to be viable, generates unilateral incentives for deviation and will be exploited by to the harm of the colluders by other market participants if discovered. Equilibrium is summarized by a small-dimensional price-like object (the density of pivotality and the threshold for extreme behavior) and it is the constancy of these across voters that induces efficiency. In fact the symmetries between competitive markets and QV extend further when one considers issues, such as incentives for information acquisition, that are beyond the scope of our analysis here.

This symmetry suggests both applied and theoretical directions for future research. On the applied side, experimental work on QV would be highly complementary with the practical implementations of QV we discussed in Section 8, helping to both shape practical application and providing a more realistic setting for testing than might be available in lab experiments such as those considered by Goeree and Zhang (2013). This parallels the way in which mechanisms for the allocation of private goods have co-evolved with laboratory and field testing of those mechanisms (Roth, 2002; Milgrom, 2004). We are aware of some such experiments in progress at the Universidad de los Andes. Another standard analytic approach in the allocation of private goods is to study how well markets conform to the idealized conditions under which the welfare theorems apply; this is much of the focus of the field of industrial organization. A natural analog is to consider empirically how well practical institutions, such as lobbying, log-rolling and other informal institutions discussed in Subsection 7.5 correspond to an efficient quadratic rule mapping costly expenditures to influence on the decision and thus how efficient the allocation of public goods is in practice in various settings. Also interesting would be work on projecting structurally how QV might change outcomes, as we did in an extremely rough way for the case of Proposition 8; we are aware of work in this direction related to the governance of chaebol in South Korea that is currently in progress.

On the theoretical side, a more formal statement of the sense in which the Arrow, Gibbard and Satterthwaite results may be misleading about the distinctions between private and public goods would be useful. In particular, building off of Hylland and Zeckhauser (1980), Weyl plans a collaboration with Hylland and Zeckhauser to show that the welfare theorems apply to an economy where public goods are allocated by QV and that if strategy-proofness is relaxed in precisely the same way that makes the double auction to be "approximately strategy-proof" in large markets the welfare theorems can be approximated using QV in finite populations. In the process, this collaboration plans to develop variants on QV that are valid without external quasilinear numeraires, which allow for direct reporting of types and approximate strategy-proofness in a sense similar to, but stronger than, that in Azevedo and Budish (2013) and allow for multiple alternatives (rather than a binary choice) as well as endogenous agenda setting. Exploration of robustness of QV beyond the limited settings we were able to formally study here would also be valuable. Also valuable would be analysis of the likely and socially desirable impact of implementation of formal rules based on QV on informal ethical and social institutions. A large theoretical and cultural edifice has been built to complement, and cushion the failings of, democracy. The institutions that could grow around a more formally efficient paradigm would be of great interest.

References

Aristotle. c. 350 B.C.E.*a*. *Ethics.* Athens.

Aristotle. c. 350 B.C.E.*b*. *Politics.* Athens.

Arrow, Kenneth J. 1951. *Social Choice and Individual Values.* New York: Wiley.

Arrow, Kenneth J. 1979. "The Property Rights Doctrine and Demand Revelation under Incomplete Information." In *Economics and Human Welfare.* , ed. Michael Boskin, 23–39. New York: Academic Press.

Attiyeh, Greg, Robert Franciosi, and R. Mark Isaac. 2000. "Experiments with the Pivot Process for Providing Public Goods." *Public Choice*, 102(1–2): 95–112.

Ausubel, Lawrence M., and Paul Milgrom. 2005. "The Lovely but Lonely Vickrey Auction." In *Combinatorial Auctions.* , ed. Peter Cramton, Richard Steinberg and Yoav Shoham, 17–40. Cambridge, MA: MIT Press.

Azevedo, Eduardo M., and Eric Budish. 2013. "Strategyproofness in the Large." http://faculty.chicagobooth.edu/eric.budish/research/Azevedo-Budish-SPL.pdf.

Azevedo, Eduardo M., E. Glen Weyl, and Alexander White. 2013. "Walrasian Equilibrium in Large, Quasilinear Markets." *Theoretical Economics*, 8(2): 281–290.

Bailey, Martin J. 1994. "Lindahl Mechanisms and Free Riders." *Public Choice*, 80(1/2): 35–39.

Bailey, Martin J. 1997. "The Demand Revealing Process: To Distribute the Surplus." *Public Choice*, 91(2): 107–126.

Bergemann, Dirk, and Stephen Morris. 2005. "Robust Mechanism Design." *Econometrica*, 73(6): 1771–1813.

Bernheim, B. Douglas, Bezalel Peleg, and Michael D. Whinston. 1987. "Coalition-Proof Nash Equilibria I. Concepts." *Journal of Economic Theory*, 42(1): 1–12.

Blais, André. 2000. *To Vote or Not to Vote: The Merits and Limits of Rational Choice Theory.* Pittsburgh, PA: University of Pittsburgh Press.

Blumer, Anselm, Andrzej Ehrenfeucht anad David Haussler, and Manfred K. Warmuth. 1987. "Occam's Razor." *Information Processing Letters*, 24(6): 377–380.

Borgers, Tilman. 2004. "Costly Voting." *American Economic Review*, 94(1): 57–66.

Bruni, Luigino, and Robert Sugden. 2013. "Reclaiming Virtue Ethics for Economics." *Journal of Economic Perspectives*, 27(4): 141–164.

Budish, Eric. 2011. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes." *Journal of Political Economy*, 119(6): 1061–1103.

Casella, Alessandra, Aniol Llorente-Saguer, and Thomas R. Palfrey. 2012. "Competitive Equilibrium in Markets for Votes." *Journal of Political Economy*, 120(4): 593–658.

Chamberlain, Gary, and Michael Rothschild. 1981. "A Note on the Probability of Casting a Decisive Vote." *Journal of Economic Theory*, 25(1): 152–162.

Chung, Kim-Sau, and J. C. Ely. 2007. "Foundations of Dominant Strategy Mechanisms." *The Review of Economic Studies*, 74(2): 447–476.

Clarke, Edward H. 1971. "Multipart Pricing of Public Goods." *Public Choice*, 11(1): 17–33.

Coate, Stephen, and Stephen Morris. 1995. "On the Form of Transfers to Special Intersts." *Journal of Political Economy*, 103(6): 1210–1235.

Crémer, Jacques, and Richard P. McLean. 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica*, 56(6): 1247–1257.

Cripps, Martin W., and Jeroen M. Swinkels. 2006. "Efficiency of Large Double Auctions." *Econometrica*, 74(1): 47–92.

d'Aspremont, Claude, and Louis-André Gérard-Varet. 1979. "Incentives and Incomplete Information." *Journal of Public Economics*, 11(1): 25–45.

de Condorcet, Marquis. 1785. *Essai sur L'Application de L'Analyse à la Probabilité des Décisions Rendue à la pluralité des Voix.* Paris: L'Imprimerie Royale.

Dekel, Eddie, Matthew O. Jackson, and Asher Wolinsky. 2008. "Vote Buying: General Elections." *Journal of Political Economy*, 116(2): 351–380.

DellaVigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2013. "Voting to Tell Others." http://elsa.berkeley.edu/ sdellavi/wp/turnout13-03-04.pdf.

Dhillon, Amrita, and Susana Peralta. 2002. "Economic Theories of Voter Turnout." *Economic Journal*, 112(480): F332–F352.

Diamond, Peter, and Emmanuel Saez. 2011. "The Case for a Progressive Tax: From Basic Research to Policy Recommendations." *Journal of Economic Perspectives*, 25(4): 165–190.

Downs, Anthony. 1957. *An Economic Theory of Democracy.* New York: Harper.

Dupuit, Arsène Jules Étienne Juvénal. 1844. *De la Mesure de L'utilité des Travaux Publics.* Paris.

Feddersen, Timothy, and Wolfgang Pesendorfer. 1997. "Voting Behavior and Information Aggregation in Elections With Private Information." *Econometrica*, 65(5): 1029–1058.

Fiorina, Morris P. 1976. "The Voting Decision: Instrumental and Expressive Aspects." *Journal of Politics*, 38(2): 390–413.

Friedman, Milton. 1962. *Capitalism and Freedom.* Chicago: University of Chicago Press.

Gibbard, Alan. 1973. "Manipulation of Voting Schemes: A General Result." *Econometrica*, 41(4): 587–602.

Goeree, Jacob K., and Jingjing Zhang. 2013. "Electoral Engineering: One Man, One ~~Vote~~ Bid." http://goo.gl/s2SCsP.

Goeree, Jacob K., and Nixon Li. 2008. "Inefficient Voting." http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.192.4763&rep=rep1&type=pdf.

Good, I. J., and Lawrence S. Mayer. 1975. "Estimating the Efficacy of a Vote." *Behavioral Science*, 20(1): 25–33.

Green, Jerry, Elon Kohlberg, and Jean-Jacques Laffont. 1976. "Partial Equilibrium Approach to the Free-Rider Problem." *Journal of Public Economics*, 6(4): 375–394.

Greif, Avner. 1994. "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies." *Journal of Political Economy*, 102(5): 912–950.

Groves, Theodore. 1973. "Incentives in Teams." *Econometrica*, 41(4): 617–631.

Groves, Theodore, and John Ledyard. 1977*a*. "Optimal Allocation of Public Goods: A Solution to the "Free Rider" Problem." *Econometrica*, 45(4): 783–809.

Groves, Theodore, and John O. Ledyard. 1977*b*. "Some Limitations of Demand Revealing Processes." *Public Choice*, 29(2S): 107–124.

Harberger, Arnold C. 1964. "The Measurement of Waste." *American Economic Review*, 54(3): 58–76.

Heifetz, Aviad, and Zvika Neeman. 2006. "On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design." *Econometrica*, 74(1): 213–233.

Hurwicz, Leonid. 1977. "On Informationally Decentralized Systems." In *Studies in Resource Allocation Processes*. , ed. Kenneth J. Arrow and Leonid Hurwicz, 425–459. Cambridge, UK: Cabmridge University Press.

Hylland, Aanund, and Richard Zeckhauser. 1980. "A mechanism for selecting public goods when preferences must be elicited." *Kennedy School of Government Discussion Paper D*, 70.

Kaplow, Louis. 2002. "On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy." *Journal of Economic Perspectives*, 18(4): 159–175.

Kawai, Kei, and Yasutora Watanabe. 2013. "Inferring Strategic Voting." *American Economic Review*, 103(2): 624–662.

Kominers, Scott Duke, and E. Glen Weyl. 2012*a*. "Concordance among Holdouts." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1591466.

Kominers, Scott Duke, and E. Glen Weyl. 2012*b*. "Holdout in the Assembly of Complements: A Problem for Market Design." *American Economic Review Papers & Proceedings*, 102(3): 360–365.

Krishna, Vijay, and John Morgan. 2012. "Majority Rule and Utilitarian Welfare." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2083248.

Ledyard, John. 1984. "The Pure Theory of Large Two-Candidate Elections." *Public Choice*, 44(1): 7–41.

Ledyard, John O., and Thomas R. Palfrey. 1994. "Voting and Lottery Drafts as Efficient Public Goods Mechanisms." *Review of Economic Studies*, 61(2): 327–355.

Ledyard, John O., and Thomas R. Palfrey. 2002. "The Approximation of Efficient Public Good Mechanisms by Simple Voting Schemes." *Journal of Public Economics*, 83(2): 153–171.

Lipset, Seymour Martin. 1960. *Political Man: The Social Bases of Politics.* Garden City, NY: Doubleday.

Mailath, George J., and Andrew Postelwaite. 1990. "Asymmetric Information Bargaining Problems with Many Agents." *Review of Economic Studies*, 57(3): 351–367.

Maskin, Eric. 1999. "Nash Equilibrium and Welfare Optimality." *Review of Economic Studies*, 66(1): 23–38.

McAfee, R. Preston, and Philip J. Reny. 1992. "Correlated Information and Mechanism Design." *Econometrica*, 60(2): 395–421.

McLean, Richard P., and Andrew Postelwaite. 2013. "Implementation with Interdependent Values." http://www.ssc.upenn.edu/%7Eapostlew/paper/pdf/iiv.pdf.

McLennan, Andrew. 1998. "Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents." *American Political Science Review*, 92(2): 413–418.

Mearsheimer, John J. 1994–1995. "The False Promise of International Institutions." *International Security*, 19(3): 5–49.

Milgrom, Paul. 2004. *Putting Auction Theory to Work.* Cambridge, UK: Cambridge University Press.

Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics*, 12(2): 380–391.

Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50(5): 1089–1122.

Myatt, David P. 2012. "A Rational Choice Theory of Voter Turnout." http://www2.warwick.ac.uk/fac/soc/economics/news_events/conferences/politicalecon/turnout-2012.pdf.

Myerson, Roger B. 2000. "Large Poisson Games." *Journal of Economic Theory*, 94(1): 7–45.

Olson, Mancur. 1965. *The Logic of Collective Action.* Cambridge, MA: Harvard University Press.

Posner, Eric. 2013. "The Good Way to Buy Votes." *Slate.*

Posner, Eric A., and Alan O. Sykes. 2014. "Voting Rules in International Organizations." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2383469.

Posner, Eric A., and E. Glen Weyl. 2013. "A Solution to the Collective Action Problem in Corporate Reorganizations." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2321904.

Posner, Eric A., and E. Glen Weyl. Forthcoming. "Quadratic Voting as Efficient Corporate Governance." *University of Chicago Law Review.*

Posner, Eric A., and E. Glen Weyl. In Preparation. *Democracy Squared.* Pending.

Posner, Eric, Glen Weyl, and Sang-Seung Yi. 2013. "A New Approach to Regulation of Chaebol, Keiretsu, and Other Conglomerate Organizations in Asia." *CLS Blue Sky Blog.*

Reny, Philip J. 1999. "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games." *Econometrica*, 67(5): 1029–1056.

Reny, Philip J. 2011. "On the Existence of Monotone Pure-Strategy Equilibria in Bayesian Games." *Econometrica*, 79(2): 499–553.

Roth, Alvin E. 2002. "The Economist as Engineer: Game Theory, Experimentation and Computation as Tools for Design Economics." *Econometrica*, 70(4): 1341–1378.

Rothkopf, Michael H. 2007. "Thirteen Reasons Why the Vickrey-Clarke-Groves Process is Not Practical." *Operations Research*, 55(2): 191–197.

Rustichini, Aldo, Mark A. Satterthwaite, and Steven R. Williams. 1994. "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica*, 62(5): 1041–1063.

Sahlins, Marshall. 1972. *Stone Age Economics.* Hawthorne, NY: Aldine de Gruyter.

Sandel, Michael J. 2012. *What Money Can't Buy: The Moral Limits of Markets.* New York: Farrar, Straus and Giroux.

Satterthwaite, Mark A., and Steven R. Williams. 1989. "The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large." *Review of Economic Studies*, 56(4): 477–498.

Satterthwaite, Mark Allen. 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.

Satz, Debra. 2010. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets.* New York: Oxford University Press.

Spenkuch, Jörg L. 2013. "On the Extent of Strategic Voting." http://kellogg.northwestern.edu/faculty/spenkuch/research/voting.pdf.

Thompson, Earl A. 1966. "A Pareto-Efficient Group Decision Process." *Papers on Non-Market Decision-Making.*

Tideman, Nicolaus. 2006. *Collective Decisions and Voting: The Potential for Public Choice.* Hampshire, UK: Ashgate.

Tideman, N. Nicolaus, and Gordon Tullock. 1976. "A New and Superior Process for Making Social Choices." *Journal of Political Economy*, 84(6): 1145–1159.

Tideman, T. Nicolaus, and Gordon Tullock. 1977. "Some Limitations of Demand Revealing Processes: Comment." *Public Choice*, 29(2S): 125–128.

Vickrey, William. 1961. "Counterspeculation, Auctions and Competitive Sealed Tenders." *Journal of Finance*, 16(1): 8–37.

Walker, Mark. 1981. "A Simple Incentive Compatible Scheme for Attaining Lindahl Allocations." *Econometrica*, 49(1): 65–71.

Walzer, Michael. 1983. *Spheres of Justice: A Defence of Pluralism and Equality.* New York: Basic Books.

Appendix

A Existence

We establish Lemma 1 by means of two sublemma that establish intermediate results. A type-symmetric Bayesian Nash equilibrium is a Bayesian Nash equilibrium in which each individual uses the same strategy $v : \mathbb{R} \to \mathbb{R}$. We refer to such an equilibrium simply by the equilibrium strategy v.

Sublemma 1. Suppose that $\underline{u}, \overline{u} \in \mathbb{R}$ and that individuals can buy votes in increments of $\iota > 0$ (permissible vote purchase are integers multiplied by ι), rather than in any continuous amount. Then a monotone-increasing, type-symmetric Bayesian Nash equilibrium v exists for any N.

Proof. We simply verify conditions for Reny (2011)'s Theorem 4.5 which guarantees the existence of a monotone type-symmetric pure strategy Bayesian Nash equilibrium. First, note that the game is clearly symmetric in the sense of Reny's definition: type spaces, action spaces, probability distributions, etc. are all interchangeable across agents. Second, we show that values satisfy the Single Crossing Condition (SCC) for games of incomplete information. By the reasoning in Section 1 of the paper, expected utility of an individual i with value u_i is

$$u_i \text{Prob}\left(v_i + V_{-i} > 0\right) - v_i^2 + \gamma\left(\boldsymbol{v}_{-i}\right),$$

where γ is some, irrelevant, smooth function of $\boldsymbol{v}_{-i} \equiv (v_1, v_2, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$ and $V_{-i} \equiv \sum_{j \neq i} v_j$. Note that $\text{Prob}\left(v_i + V_{-i} > 0\right)$ is a weakly increasing function and that u_i enters this expression nowhere except multiplying it. This immediately implies SCC. Third, note that no individual will ever buy votes less than $-\sqrt{-\underline{u}}$ nor bigger than $\sqrt{\overline{u}}$ as such actions are strictly dominated by buying 0 votes given that $\text{Prob}\left(v_i + V_{-i} > 0\right)$ is bounded between 0 and 1. Thus, without loss of generality, we may restrict attention to vote values defined by integers between $\left\lfloor \frac{-\sqrt{-\underline{u}}}{\iota} \right\rfloor$ and $\left\lceil \frac{\sqrt{\overline{u}}}{\iota} \right\rceil$ multiplied by ι. This is a finite set of actions, as $\iota > 0$, and thus we have a finite action space and a bounded, single-dimensional, real type space. Thus, fourth, by Reny's Proposition 4.4 (and in particular Remark 2), because the action set is finite and SCC holds, each individual's set of monotone pure strategies is non-empty and join-closed. Fifth, Reny's conditions G.1-5 follow from his Proposition 3.1 because each player's type space is a closed interval in $\mathbb{R}$, type-distributions are atomless and their action sets are a finite (and thus compact) subset of the real line. Finally, condition G.6 is trivially satisfied as the action set is finite. □

Sublemma 2. Suppose that $\underline{u} < 0 < \overline{u}$. Then a monotone-increasing, type-symmetric Bayesian Nash equilibrium v exists for any N.

Proof. We prove this result by showing existence in an alternative game with $2N$ players; in particular, for each player i in the original game we consider a game with two players i_+ and i_-. Every player i_+ has utility u_{i_+} drawn from $(0, \overline{u})$ iid according to the pdf $f(u|u > 0)$ and every player i_- has utility u_{i_-} drawn from $(\underline{u}, 0)$ iid $f(u|u < 0)$. Let $q_- = \text{Prob}(u < 0) > 0$ by the full support assumption and $q_+ = \text{Prob}(u > 0) > 0$ by the same assumption; note $q_- + q_+ = 1$ because f has no atom at 0. The game is exactly like QV, except that, iid across i, player i_+ "exists" with probability q_+ and player q_- "exists" with probability $q_- = 1 - q_+$.

More formally, each player i_+ has type-conditional expected utility

$$q_+ \left[u_{i_+} \text{Prob}\left(v_{i_+} + \widetilde{V}_{-i} > 0 \right) - v_{i_+}^2 + \gamma\left(\boldsymbol{v}_{-i}\right) \right]$$

where $\widetilde{V}_{-i} = \sum_{j \neq i} 1_{c_j} v_{j_+} + \left(1 - 1_{c_j}\right) v_{j_-}$ and $\{c_i\}_{i=1}^n$ are iid Bernoulli random variables with success probability q_+. Note that any (type-symmetric-among-players-of-the-same-sign) monotone, pure strategy equilibrium (SPSSMPSE) in this game (v_-, v_+) corresponds to a (type-symmetric) MPSE (SMPSE) in the original game v where v is the function that has value v_- when $u < 0$ and v_+ when $u > 0$ and value 0 when $u = 0$ because the payoff of player i in the original game is just the sum of the payoffs of players i_- and i_+. The converse is also true: a SMPSE in the original game corresponds to one in this game, simply splitting players.

Now we show that this auxiliary game has a SPSSMPSE for any N. First, note that the game is payoff secure. Each player i_+ can ϵ-secure her payoff if at $\left(v_{i_+}, \boldsymbol{v}_{-i_+}\right)$, where $\boldsymbol{v}_{-i_+}$ is the strategy of all other players, by playing $\sqrt{\frac{\epsilon}{2\overline{u}} + 1} v_{i_+}$ as long as $v_{i_+} \leq \sqrt{\overline{u}}$ and we can restrict our attention to strategies such that this is the case by the reasoning in the proof of Sublemma A and *mutatis mutandis* for player i_-. Second note that the game is reciprocally upper semi-continuous as whenever the payoff of any player i_+ jumps upward the payoffs of all players $-i_-$ jump downward at the same point and *mutatis mutants* for players i_- (their payoffs from the unique potential discontinuity driven by the vote threshold always have opposite signs). Therefore by Proposition 3.2 of Reny (1999), this game is better-reply secure. Thus, following the strategy of Reny (2011)'s proof of Corollary 5.2 in Appendix A.7, Remark 3.1 in Reny (1999) and Lemma A.13 of Reny (2011) imply that it suffices for existence of a SPSSMPSE to show that, for any $\epsilon > 0$, this game has an ϵ equilibrium in type-symmetric-among-players-of-the-same-sign monotone pure strategies.

Fix ϵ. By Sublemma A and the above observation of equivalence between equilibria in the original and modified game, if the bid space is discretized by any ι an equilibrium exists. Let

$$\iota^\star \equiv \min \left\{ \frac{\epsilon}{4\sqrt{\max\{\overline{u}, -\underline{u}\}}}, \frac{\sqrt{\epsilon}}{8} \right\}.$$

We claim that the exact equilibrium for the $\iota^\star$ discretization of the game is an ϵ-equilibrium of the continuous game; clearly this equilibrium is symmetric-among-players-of-the-same-sign, monotone and pure. To see that it constitutes an ϵ-equilibrium, consider any best response of any type, $u > 0$, of a player i_+ if she had access to all continuous votes to the $\iota^\star$-discretized equilibrium play of all other players and call this $v^\star$. If, rather than playing $v^\star$, i_+ instead plays $\iota^\star \left\lceil \frac{v^\star}{\iota^\star} \right\rceil$ she achieves utility no worse than

$$\left(\iota^\star \left\lceil \frac{v^\star}{\iota^\star} \right\rceil\right)^2 - (v^\star)^2 \leq \left(v^{\star + \iota^\star}\right)^2 - (v^\star)^2 = (\iota^\star)^2 + 2\iota^\star v^\star$$

below the payoff if she plays $v^\star$. But $\iota^\star \left\lceil \frac{v^\star}{\iota^\star} \right\rceil$ is a strategy available in the $\iota^\star$-discretization of the game and thus, at type u, i_+ must earn a utility at least as high as if she played $\iota^\star \left\lceil \frac{v^\star}{\iota^\star} \right\rceil$ in the $\iota^\star$ discretization equilibrium. But note the any strategy greater than $\sqrt{\overline{u}}$ is strictly dominated in the continuous game and thus $(\iota^\star)^2 + 2\iota^\star v^\star < (\iota^\star)^2 + 2\iota^\star \sqrt{\overline{u}}$. Thus the $\iota^\star$ discretization equilibrium strategy for type u of player i_+ is a

$$(\iota^\star)^2 + 2\iota^\star \sqrt{\overline{u}} \leq \min\left\{\frac{\epsilon^2}{16\overline{u}} + \frac{\epsilon}{4}, \frac{\sqrt{\epsilon \overline{u}}}{4} + \frac{\epsilon}{4}\right\} \leq \frac{\epsilon}{4} + \frac{\epsilon}{2^{\frac{4}{3}}} < \epsilon$$

best-reply for type u of player i_+ of the continuous game. But this is true for any type u of player i_+ and thus the $\iota^\star$ discretization strategy is an ϵ-best reply for player i_+. Applying the same logic to players i_- establishes that the claim and completes the proof. □

In some ways constructing this auxiliary game is a bit artificial. But the existence of types with both strictly positive and strictly negative utility is not only necessary to make this proof work, but also necessary for equilibrium to exist in the continuous game, as the following example shows.

Example 5. Consider QV with values drawn from an atomless distribution on $(0, \overline{u})$ where $\overline{u} > 0$ and with continuous votes. Then note that no pure strategy equilibrium exists if ties are broken in any way other than in favor of the alternative (which is the opposite of the way in the text we assume ties are broken). To see this suppose otherwise. First imagine there is an equilibrium in which all individuals buy 0 votes. Clearly this is not an equilibrium as every type would do better to buy a sufficiently small number of votes to break the tie. Instead suppose there is an equilibrium in which any type of any player chose to buy a positive number of votes with any probability. This clearly is not an equilibrium either as this player could always buy fewer votes and still break the tie; in fact, no player every has a best response in this game, other than possibly 0 and that only when, with probability 1 all other players are buying positive votes. Thus no equilibrium exists.

Players with values of both signs are ensure reciprocal upper semi-continuity, or more intuitively, that there is no problem with exact ties as there are players pulling (with at least some probability) in both directions. The proof strategy we use in the proof of Sublemma 2 may be of some broader interest; it shows that rather than having to show that another *player* must have a payoff jump down whenever one player's payoff jumps up all one need show is that whenever *one type of player i's payoff jumps up there is some positive probability mass of player $j \neq i$ types that have payoffs* *that jump down.* This holds more broadly and we invoke this strategy in the proof of the main theorem now without repeating the construction.

Proof of Lemma 1. This game is (effectively) better-reply secure by the same arguments as in the proof of Sublemma 2 and thus it suffices to show that an equilibrium in symmetric, monotone, pure strategies exists for every .

Fix $\epsilon > 0$. Because the first moment of the distribution exists, there exists a $u^\star \in \mathbb{R}_+$ such that

$$n \int_{u^\star}^{\infty} uf(u)du, -n \int_{-\infty}^{-u^\star} uf(u)du < \frac{\epsilon}{10}. \tag{A.1}$$

Consider the auxiliary game where with bounded values and distribution over these values $f\left(u|u \in (-u^\star, u^\star)\right)$. This game has a monotone, pure strategy, symmetric equilibrium by Subemma 2; call this equilibrium $v^\star$. Consider the strategy, $\hat{v}$ in the original game which plays $v^\star$ in this auxiliary game on $(-u^\star, u^\star)$, $v^\star(-u^\star)$ on $(\underline{u}, 0) \setminus (\underline{u}^\star, 0)$ and $v^\star(u^\star)$ on $(0, \overline{u}) \setminus (0, u^\star)$. Note that this strategy is monotone. We claim that $\hat{v}$ is an ϵ-equilibrium of the original game.

To see this, first note that the probability of any type in $U_E \equiv (\underline{u}, \overline{u}) \setminus (-u^\star, u^\star)$ being drawn by any of $n-1$ individuals is strictly less than $\frac{\epsilon}{5E[|u||u \in U_E]}$ by Bonferroni's inequality and the construction of inequality A.1. Fix a player i and call the event that any other player draws a type in this set E. Then player i's payoff expected utility is

$$\text{Prob}(E)\pi_{i,E}(v_i) + (1 - \text{Prob}(E))\,\pi_{i,\neg E}(v_i) + E\left[\gamma\left(\boldsymbol{v}_{-i}\right)\right], \tag{A.2}$$

where $\pi_{i,E}$ is the player's expected payoff (less γ) in the state that E occurs and and $\pi_{i,\neg E}$ is her payoff if E does not occur. Call a strategy v *$u^\star$-undominated* if $-\sqrt{-\underline{u}^\star} \leq v \leq \sqrt{\overline{u}}$. Note that $\pi_{i,E}$ is bounded below by $\text{Prob}\left(u < 0\right) E\left[u|u < 0\right] - u^\star$ and above by $\text{Prob}\left(u > 0\right)\left(E\left[u|u > 0\right]\right)$ assuming the player uses a strategy that is $u^\star$-undominated. The reason is that the worst she can do with a $u^\star$-undominated strategy is to pay $u^\star$ and received her less desired outcome every time, in which case her expected utility is $E\left[u|u < 0\right]$. The best she can ever do is to pay nothing and always achieve her desired outcome, in which case her expected utility is $\text{Prob}\left(u > 0\right)\left(E\left[u|u > 0\right]\right)$. Thus the maximum gain in $\pi_{i,E}$ that player i could achieve by moving from a strategy that is $(\underline{u}^\star, \overline{u}^\star)$-undominated is

$$\text{Prob}\left(u > 0\right)\left(E\left[u|u > 0\right]\right) - \left(\text{Prob}\left(u < 0\right) E\left[u|u < 0\right] - u^\star\right) = E\left[|u|\right] + u^\star \leq 2E\left[|u||u \in U_E\right],$$

as $E\left[|u||u \in U_E\right] \geq E\left[|u|\right]$ as U_E only contains values that are sufficiently large in absolute value and in particular bigger than $u^\star$ so that $E\left[|u||u \in U_E\right] \geq u^\star$. Clearly $\hat{v}$ is $u^\star$-undominated as it is an equilibrium in a game where no player has a type greater in absolute magnitude than $u^\star$. Thus the first term of equation A.2 can never be more than $\frac{2}{5}\epsilon$ greater than the value it achieves at $\hat{v}$.

Next, consider the second term of expression A.2. Note that this is precisely the payoff to player i if all other players acted as if they were in the auxiliary game; conditional on no other individual having a type in U_E the distribution of types of other players is precisely $f(u|u \in (-u^\star, u^\star))$. The only improvements, therefore, over $\hat{v}$ possible in this term are on U_E, given that $\hat{v}$ agree with $v^\star$ on $(-u^\star, u^\star)$ and, conditional on $\neg E$, $v^\star$ is a best response for these types by construction.

Consider the maximum possible gain that could be achieved on U_E compared to the strategy of always playing 0 on these states. By the same logic as above, this gain is no greater than

$$\text{Prob}\,(U_E)\, E\left[|u||u \in U_E\right] < \frac{2}{5}\epsilon$$

by construction, adding the two terms in inequality A.1. However, by the SCC, which we know applies to this game by the same logic as in the proof of Sublemma A and the fact that $v^\star(-u^\star), v^\star(u^\star)$ are best responses[14] for types $-u^\star, u^\star$ respectively, $\hat{v}$'s prescribed strategy on U_E always is weakly better than 0 for types in U_E. Thus the maximum gain to the second term of expression A.2 that can be achieved by a strategy other than $\hat{v}$ is $\frac{2}{5}\epsilon$. Thus the maximum gain that can be achieved by moving to any other strategy to the full expression A.2 is $\frac{4}{5}\epsilon < \epsilon$ and thus if all players play $\hat{v}$ this is an ϵ equilibrium, completing the proof.

□

This argument provides a generic strategy for extending Reny (2011)'s argument to the case of an unbounded type space if the first moment of the types exist and utility takes a standard multiplicative form. This approach may be of some use in structural auction and other mechanism design problems where an unbounded type-space is more convenient than a bounded type-space.

B Characterization

C Efficiency

D Uniqueness

E Collusion

F Aggregate Uncertainty

Proof of Lemma 4. FILL IN

[14] Except, of course, in the case when $(\underline{u}, 0) \setminus (\underline{u}^\star, 0)$ or $(0, \overline{u}) \setminus (0, u^\star)$ are empty, in which case the result holds trivially for that side and thus can be ignored.

□

Proof of Lemma 5. FILL IN.

□

Before beginning the specific examples, we derive some general facts about Equation 6.1. By Bayes's rule

$$\mathbb{E}\left[g\left(\gamma^{\star}|u\right)u|\gamma^{\star}\right]=\int_{u}u\frac{f\left(u|\gamma^{\star}\right)g\left(\gamma^{\star}\right)}{f(u)}f\left(u|\gamma^{\star}\right)du=g\left(\gamma^{\star}\right)\int_{u}u\frac{f^{2}\left(u|\gamma^{\star}\right)}{f(u)}du.$$

Thus any $\gamma^{\star}$ solving $\int_{u}u\frac{f^{2}(u|\gamma^{\star})}{f(u)}du=0$ also solves $\mathbb{E}\left[g\left(\gamma^{\star}|u\right)u|\gamma^{\star}\right]=0$ and thus is an equilibrium value of $\gamma^{\star}$.

Proof of Example 1. We assume, throughout and without loss of generality given the symmetry of the normal distribution, that $\mu>0$. The marginal distribution of u is $\mathcal{N}\left(\mu,\sigma_{1}^{2}+\sigma_{2}^{2}\right)$ while the γ-conditional distribution is $\mathcal{N}\left(\gamma,\sigma_{1}^{2}\right)$ by standard properties of the normal distribution. We use this to solve out for $\gamma^{\star}$:

$$\int_{u}u\frac{f^{2}\left(u|\gamma^{\star}\right)}{f(u)}du=0\iff\int_{u}ue^{\frac{(u-\mu)^{2}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)}-\frac{(u-\gamma^{\star})^{2}}{\sigma_{1}^{2}}}=0\iff\frac{(u-\mu)^{2}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)}-\frac{(u-\gamma^{\star})^{2}}{\sigma_{1}^{2}}=au^{2}+b$$

for some constants a and b independent of u as this is the only quadratic form symmetric about 0, and symmetry about 0 is clearly necessary to yield a 0 expectation given the normal form of the density.

$$\frac{(u-\mu)^{2}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)}-\frac{(u-\gamma^{\star})^{2}}{\sigma_{1}^{2}}=\frac{\sigma_{1}^{2}\left(u-\mu\right)^{2}-2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\left(u-\gamma^{\star}\right)^{2}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\sigma_{1}^{2}}=au^{2}+b-2\frac{\sigma_{1}^{2}\mu-2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\gamma^{\star}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\sigma_{1}^{2}}u.$$

Thus $\gamma^{\star}$ solves

$$2\frac{\sigma_{1}^{2}\mu-2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\gamma^{\star}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\sigma_{1}^{2}}=0\iff\gamma^{\star}=\frac{\sigma_{1}^{2}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)}\mu.$$

In a large population, the first-best welfare is proportional to $\mathbb{E}\left[|\gamma|\right]=\sigma_{2}\sqrt{\frac{2}{\pi}}e^{-\frac{\mu^{2}}{2\sigma_{2}^{2}}}+\mu\left[1-2\Phi\left(-\frac{\mu}{\sigma_{2}}\right)\right]$. Welfare loss relative to this occurs in a large population when $\gamma\in\left(0,\gamma^{\star}\right)$ and, in these cases, is proportional to $|\gamma|$. This loss equals

$$\int_{0}^{\frac{\sigma_{1}^{2}}{2\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)}\mu}\frac{\gamma e^{-\frac{(\gamma-\mu)^{2}}{2n\sigma_{2}^{2}}}}{\sigma_{2}\sqrt{2\pi}}d\gamma$$

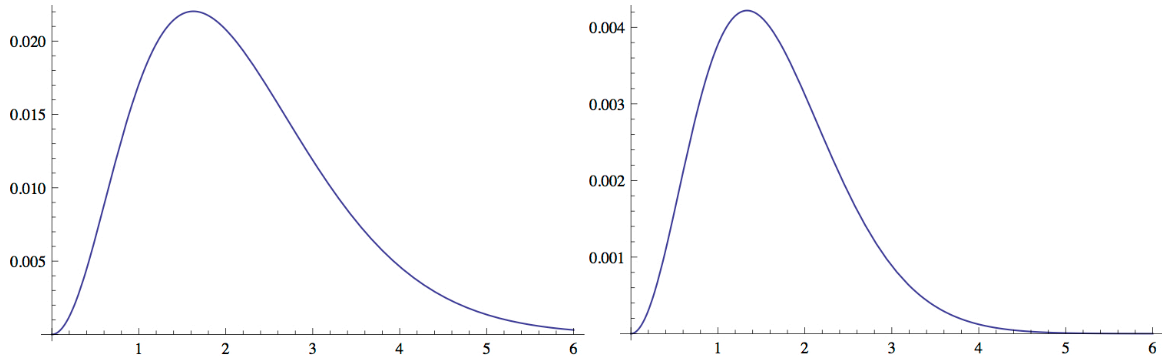

Figure 3: EI for the joint normal example when $\sigma_1 \to \infty$ (left) and $\sigma_1 = \sigma_2$ (right) as a function of $\frac{\mu}{\sigma_2}$.

which is clearly monotonically increasing in $\frac{\sigma_1^2}{2\left(\sigma_1^2+\sigma_2^2\right)}\mu$, which in turn monotonically increases in σ_1^2. We can further compute analytically using Mathematica that

$$\int_0^{\frac{\sigma_1^2}{2\left(\sigma_1^2+\sigma_2^2\right)}\mu} \frac{\gamma e^{-\frac{(\gamma-\mu)^2}{2\sigma_2}}}{\sigma_2\sqrt{2\pi}} d\gamma = \mu\left[\Phi\left(\frac{\mu}{\sigma_2}\right) - \Phi\left(\frac{\mu\left(\sigma_1^2+2\sigma_2^2\right)}{2\sigma_2\left(\sigma_1^2+\sigma_2^2\right)}\right)\right] - \frac{\sigma_2\left(e^{-\frac{\mu^2\left(\sigma_1^2+2\sigma_2^2\right)^2}{8\sigma_2^2\left(\sigma_1^2+\sigma_2^2\right)^2}} - e^{-\frac{\mu^2}{2\sigma_2^2}}\right)}{\sqrt{2\pi}}.$$

Thus EI is

$$\frac{x\left[\Phi(x) - \Phi\left(\frac{x\left(\sigma_1^2+2\sigma_2^2\right)}{2\left(\sigma_1^2+\sigma_2^2\right)}\right)\right] - \frac{e^{-\frac{x^2\left(\sigma_1^2+2\sigma_2^2\right)^2}{8\left(\sigma_1^2+\sigma_2^2\right)^2}} - e^{-\frac{x^2}{2}}}{\sqrt{2\pi}}}{2\left(\sqrt{\frac{2}{\pi}}e^{-\frac{x^2}{2}} + x\left[1-2\Phi(-x)\right]\right)}$$

where $x \equiv \frac{\mu}{\sigma_2}$. In the limit as $\sigma_1 \to \infty$ this becomes

$$\frac{x\left[\Phi(x) - \Phi\left(\frac{x}{2}\right)\right] - \frac{e^{-\frac{x^2}{8}} - e^{-\frac{x^2}{2}}}{\sqrt{2\pi}}}{2\left(\sqrt{\frac{2}{\pi}}e^{-\frac{x^2}{2}} + x\left[1-2\Phi(-x)\right]\right)}$$

and when $\sigma_1 = \sigma_2$

$$\frac{x\left[\Phi(x) - \Phi\left(\frac{3x}{4}\right)\right] - \frac{e^{-\frac{9x^2}{32}} - e^{-\frac{x^2}{2}}}{\sqrt{2\pi}}}{2\left(\sqrt{\frac{2}{\pi}}e^{-\frac{x^2}{2}} + x\left[1-2\Phi(-x)\right]\right)}.$$

Figure 3 shows the EI expression in both of these cases. The results reported in the text can be read directly off these plots.

□

Proof of Example 2. The average value conditional on γ is $\gamma\mu_+ - (1-\gamma)\mu_-$ so that $\gamma_0 = \frac{\mu_-}{\mu_- + \mu_+}$. $f(u|\gamma) = \gamma$ for $u > 0$ and $f(u|\gamma) = 1-\gamma$ for $u < 0$. As a result, $f(u) = \mathbb{E}[\gamma]$ for $u > 0$ and $1 - \mathbb{E}[u]$ for $u < 0$. Thus $\gamma^\star$ solves

$$\mu_+ \frac{\gamma^2}{\mathbb{E}[\gamma]} - \mu_- \frac{(1-\gamma)^2}{1-\mathbb{E}(\gamma)} = 0 \implies \gamma^2 k = (1-\gamma)^2,$$

where $k \equiv \frac{\mu_+(1-\mathbb{E}[\gamma])}{\mu_- \mathbb{E}[\gamma]}$. Solving this quadratic equation yields

$$\gamma = \frac{-1 \pm \sqrt{k}}{k-1}.$$

The solution must be in the interval $[0,1]$, which the negative solution never is and the positive solution always is. Thus

$$\gamma^\star = \frac{\sqrt{k}-1}{k-1} = \frac{1}{\sqrt{k}+1}.$$

Efficiency results if and only if $\gamma_0 = \gamma^\star$, that is if

$$\frac{1}{1+\sqrt{k}} = \frac{\mu_-}{\mu_- + \mu_+} \iff \sqrt{k} = \frac{\mu_+}{\mu_-} \iff \frac{\mu_+^2}{\mu_-^2} = \frac{\mu_+ (1-E[\gamma])}{\mu_- E[\gamma]} \iff \mu_+ E[\gamma] = \mu_- (1 - E[\gamma]),$$

hat is the election is an expected welfare tie ex-ante.

$$\frac{\mu_+}{\mu_- \sqrt{k}} = \sqrt{\frac{\mu_+ E[\gamma]}{\mu_- (1-E[\gamma])}}$$

so that $\frac{\mu_+}{\mu_-} > (<) \sqrt{k} \iff \mu_+ E[\gamma] > (<) \mu_- (1 - E[\gamma])$. Thus

$$\gamma_0 < (>) \gamma^\star \iff \mu_+ E[\gamma] > (<) \mu_- (1-E[\gamma]),$$

the Bayesian Underdog Effect.

Under majority rules the threshold in γ for implementing the alternative with high probability is $1/2$. For each regime we compute EI as

$$1 - \frac{\int_0^{\gamma_t} [\mu_- (1-\gamma) - \mu_+ \gamma] h(\gamma) d\gamma + \int_{\gamma_t}^1 [\mu_+ \gamma - \mu_- (1-\gamma)] h(\gamma) d\gamma}{\int_0^{\gamma_0} [\mu_- (1-\gamma) - \mu_+ \gamma] h(\gamma) d\gamma + \int_{\gamma_0}^1 [\mu_+ \gamma - \mu_- (1-\gamma)] h(\gamma) d\gamma},$$

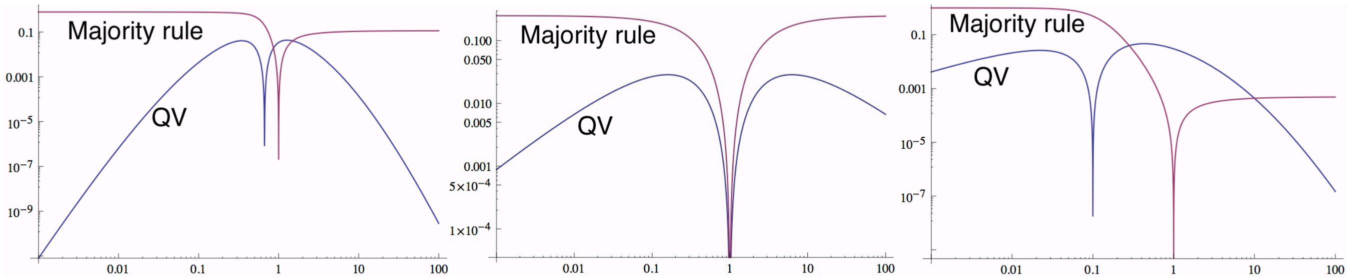

Figure 4: EI of QV compared to majority rule in the Krishna and Morgan (2012) example with γ distributed uniform (left), $\beta(15, 10)$ (center) and $\beta(10, 1)$ (right). Both axes are on a log-scale, though labeled linearly; the x-axis, measures $r = \frac{\mu_+}{\mu_-}$.

where γ_t is the appropriate threshold value f γ. Using this method and explicit integration on Mathematica, we computed the relative (to the first best) efficiency of QV and majority-rules assuming g follows a Beta distribution. Note that, if one divides the numerator and denominator by μ_-,

$$\frac{\int_0^{\gamma_t} \left[\mu_- (1-\gamma) - \mu_+ \gamma\right] h(\gamma) d\gamma + \int_{\gamma_t}^1 \left[\mu_+ \gamma - \mu_- (1-\gamma)\right] h(\gamma) d\gamma}{\int_0^{\gamma_0} \left[\mu_- (1-\gamma) - \mu_+ \gamma\right] h(\gamma) d\gamma + \int_{\gamma_0}^1 \left[\mu_+ \gamma - \mu_- (1-\gamma)\right] h(\gamma) d\gamma} =$$

$$\frac{\int_0^{\gamma_t} \left[(1-\gamma) - \frac{\mu_+}{\mu_-}\gamma\right] h(\gamma) d\gamma + \int_{\gamma_t}^1 \left[\frac{\mu_+}{\mu_-}\gamma - (1-\gamma)\right] h(\gamma) d\gamma}{\int_0^{\gamma_0} \left[(1-\gamma) - \frac{\mu_+}{\mu_-}\gamma\right] h(\gamma) d\gamma + \int_{\gamma_0}^1 \left[\frac{\mu_+}{\mu_-}\gamma - (1-\gamma)\right] h(\gamma) d\gamma},$$

So EI depends only on the ratio $r \equiv \frac{\mu_+}{\mu_-}$ and on the parameters of the Beta distribution, not on both μ_+ and μ_- independently

Figure 4 shows three examples that are representative of the more than 100 cases we experimented with. Whenever $\alpha = \beta$ (the distribution of γ is symmetric), QV always dominates majority rule as it does in the left panel shown, which is $\alpha = \beta = 1$, the uniform distribution. Majority obviously performs best when r, shown on the horizontal axis, is near to unity. When α is larger than β, majority rule may out-perform QV near $r = 1$. This is shown in the center and right panels where $(\alpha, \beta) = (15, 10)$ and $(10, 1)$ respectively. The larger α is relative to β, the larger the region over which voting outperforms QV. However, it is precisely in these cases where, if r is very small, voting is most dramatically inefficient. Intuitively majority may outperform QV by blindly favoring the majority which is almost always in favor of taking the action for$\alpha \gg \beta$, while QV may be a bit too conservative in favoring the action because of the Bayesian Underdog Effect. However this blind favoritism towards the majority view can be highly destructive under majority rule, but not under QV, when the minority has an intense preference. In fact, while voting becomes highly inefficient when the minority preference becomes very intense, QVB actually becomes closer to the first best. In all cases (shown here and that we have sampled) QV's efficiency is above 90\% and usually it is well above this. □

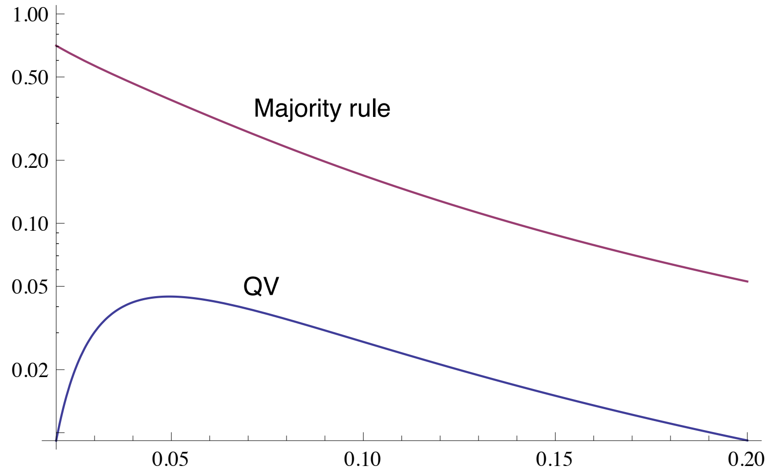

Figure 5: EI as a function of the standard deviation of the total vote share in favor of Proposition 8. The vertical axis is on a logarithmic scale with linear labels.

Proof of Example 3. 4% of the population is certainly opposed to the alternative and willing to pay on average \$34k to oppose it. A fraction γ of the remaining 96% of the population opposes supports the alternative. On average, independent of γ, the proponents and opponents in the 96% are willing to pay \$5k to get their way. The average value from implementing the alternative is therefore $\$4800(2\gamma - 1) - 1360$ and $\gamma_0 = .64$.

Individuals in the 4% receive no signal about γ and thus $f(u|\gamma)$ for this group is simply $f(u)$. For proponents of the alternative among the 96%, $f(u|\gamma)$ is, by the logic of the previous proof, $.96f(u)\gamma$ and for opponents among the 96% $.96f(u)(1-\gamma)$; $f(u)$ is formed by taking expectations over γ as in the previous proof. Using the same techniques derivations as there we can solve for $\gamma^\star$.

To calibrate, we assume that a Beta distribution of γ and that

$$.96\frac{\alpha}{\alpha+\beta} = .96\mathbb{E}[\gamma] = .52.$$

Solving this out implies that $\alpha = 1.18\beta$. The variance of γ is given by the standard formula for the variance of a Beta distributed variable:

$$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)} = \frac{1.18\beta^2}{2.18^2\beta^2(2.18\beta+1)} = \frac{.25}{2.18\beta+1}.$$

Thus the standard deviation of the total fraction of the population supporting the alternative is $\frac{.96\cdot.5}{\sqrt{2.18\beta+1}} = \frac{.48}{\sqrt{2.18\beta+1}}$ and thus if the standard deviation of the vote share for the alternative under standard voting is σ then $\beta = \frac{.45\left(.23-\sigma^2\right)}{\sigma^2}$. Figure 5 shows the dominance of QV.

□

G Voter Behavior

H Common Values

I Small Populations

To calculate approximate equilibria for small populations we used standard computational game theory techniques for solving for equilibria.We began by initializing voting functions for each individual separately to $v_{i0}(u) = \frac{u}{2}$, so as to allow for the potential identification of asymmetric equilibria.[15] We then entered a loop to calculate equilibrium values of v_{i0} for each individual until it "converged" in the sense that the "update error" ϵ_t defined in the loop below being less than .005 or until $8N$ loop iterations had passed, in which case it was determined that the loop was not converging. These and all other numbers below were obtained by trial-and-error to involve the minimum computation time necessary for consistent and reliable results. The loop in period t ran the following steps:

1. If t mod N=0 check to see if all individuals have converged or if $t = 8N$. If either has occurred, terminate. If not, continue the loop.

2. For individual $i = t \mod N + 1$, draw $500,000$ random values for each of the other $N - 1$ individuals. Use each set of draws to calculate the sum of all other votes using v_{it-1}. Tabulate a histogram of these values on 5000-point, evenly-spaced grid from the lowest to the highest observed value of the sum of other votes. To this pure empirical PDF, fit a 13th degree polynomial approximation for smoothing that minimizes mean-squared error to the pure empirical PDF.

3. Divide the support of u, or in the case of the normal distribution the mean of u plus and minus 5.8 standard deviations, into 5000 evenly-spaced grid points. For each grid point, numerically solve for the number of votes maximizing expected utility using Newton's Method on the first-order condition $v_i(u) = \frac{p(-v)}{2}u$, with an approximation accurate two significant digits. Approximate this function using a piecewise cubic spline interpolation (except in the case when $N = 2$ and the distribution is uniform, in which case a 10th degree polynomial offers a better approximation consistently).

4. Calculate the Euclidean norm between this function and v_{it-1}. If this is less than .005 label individual i as "converged"; if it is greater than .005 label individual i as "not converged".

5. Store this as v_{it} and for all $j \neq t \mod N$replace set $v_{jt} \equiv v_{jt-1}$.

[15]Other initialization values were tried, including asymmetric ones, but very similar results were typically obtained though often after more interactions.

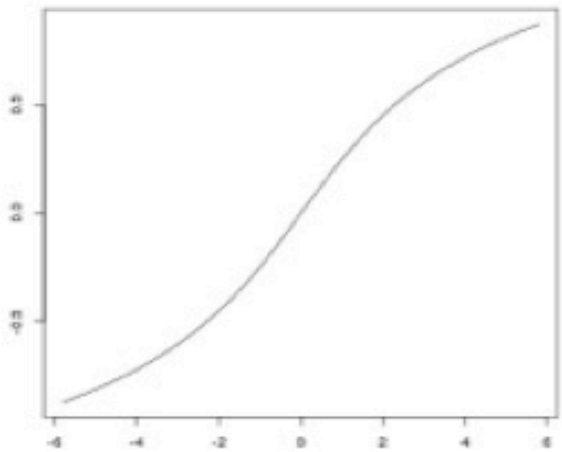

Figure 6: Approximate equilibrium voting function for $N = 10$ given values are drawn from a standard Normal distribution.

6. Loop.

Before moving forward, the distribution of the sum of votes was inspected visually to ensure that there was only one solution to the first-order conditions and that it corresponded to a maximum in the final results. Then, using the output values of v_{iT}, where T was the final period of the loop, we calculated Expected Inefficiency (EI) as follows

1. Draw 500,000 random values of each individual. Let u_i^j but the$j$th utility values for individual i. Let $U_j \equiv \sum_{i=1}^{N} u_i^j$, $V_j \equiv \sum_{i=1}^{N} v_{iT}\left(u_i^j\right)$ and let $M_j \equiv 2\sum_{i=1}^{N} 1_{u_i^j \geq 0} - N$.

2. EI of QV is

$$\frac{1}{2} + \frac{1}{500,000} \sum_{j=1}^{500,000} \frac{U_j 1_{V_j \geq 0} - U_j 1_{V_j < 0}}{2\left(U_j 1_{U_j \geq 0} - U_j 1_{U_j < 0}\right)}$$

and of majority rule[16] is

$$\frac{1}{2} + \frac{1}{500,000} \sum_{j=1}^{500,000} \frac{U_j \left(1_{M_j > 0} + \frac{1_{M_j = 0}}{2}\right) - U_j \left(1_{M_j < 0} + \frac{1_{M_j = 0}}{2}\right)}{2\left(U_j 1_{U_j \geq 0} - U_j 1_{U_j < 0}\right)}.$$

The distributions considered were the Normal distribution, with parameters μ and σ^2, and the Pearson Type I distribution with parameters $\underline{u}$ and $\overline{u}$ for the bounds of the support, α and β for the shape and position within these bounds. We tried many values for these parameters and N ranging from 2 to 10. In this appendix we only show results that directly support claims in the text.

[16]Note that ties are non-generic under QV and thus we simply assume they are broken in favor of the alternative, but under majority rule they matter and thus are broken by a coin flip.

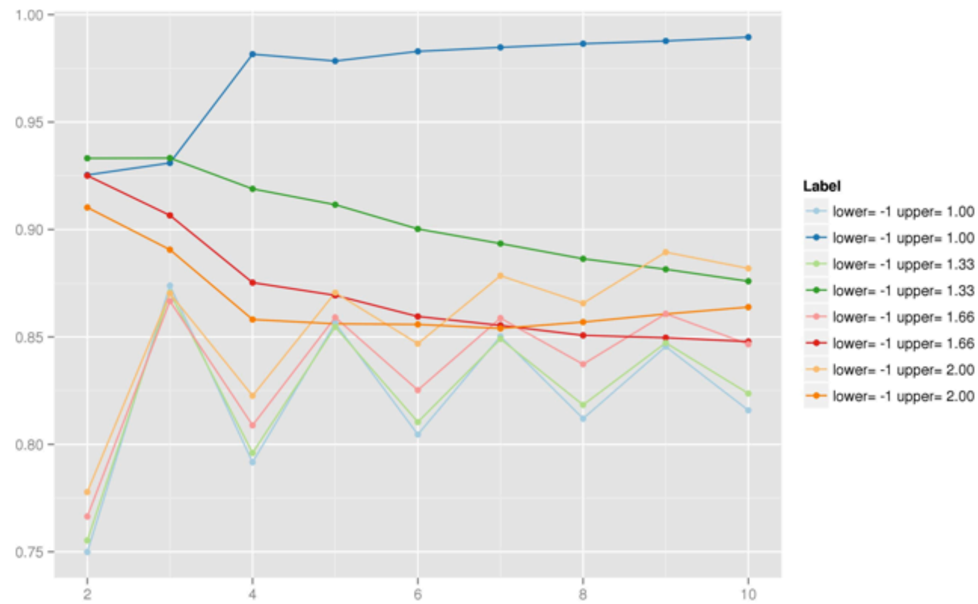

Figure 7: Expected efficiency of QV and majority rule under a uniform value distribution for the various values of N (horizontal axis) and different upper and lower bounds for the distribution (the colors, correspond to the legend at right). The dark plots are QV, the light plots majority rule.

Figure 6 shows the converged voting rule when $N = 10$ and values are drawn from a standard normal distribution. The shape is close to linear but has a gentle version of the S-shape that would be predicted by the characterization in Subsection 4.2: individuals with large values in either direction buy fewer votes per unit of value because they are less likely to be pivotal with a marginal vote.

Figure 7 shows our first set of results, for the uniform distribution varying over different ranges of the bounds and different values of N. Expected efficiency, rather than expected inefficiency, is graphed. For N below 4 and/or an upper bound on the distribution less than 1.33 QV always outperforms majority rule. However as the mean and the median shift up and N becomes large, majority rule outperforms QV, though never by a large amount.

Figure 8 shows results for Normal distributions with varying μ and σ^2 (as well as N) in the left panel and varying μ holding σ^2 fixed at 1 in the right panel. In the left panel, all cases with $\mu = 0$ have QV dominating. Even when the $\mu = .2$ QV dominates except in the case when $N = 10$ and $\sigma = .5$; for larger σ or smaller N, QV again dominates. The right panel varies the mean over a wider range and exhibits a wider range of behaviors as a consequence. For $N = 10$, where the gap between QV and majority rule is largest, QV's performance is non-monotone in the mean, falling and then rising, while majority rule monotonically improves as the mean grows larger, leading majority rule to outperform QV by 4-5 percentage points in some cases.

Finally, Figure 9 highlights our most striking results, which occur under various version of the Pearson Type I distribution. In the left panel is pictured a case where the support of the distribution

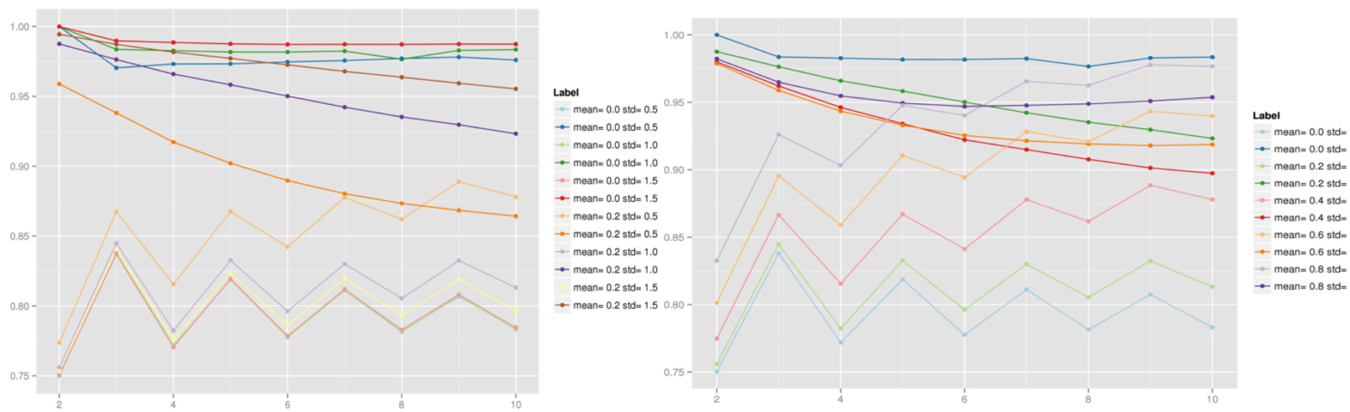

Figure 8: Expected efficiency of QV and majority rule under a Normal value distribution for the various values of N (horizontal axis) and means and variances (the colors, correspond to the legend at right). The dark plots are QV, the light plots majority rule.

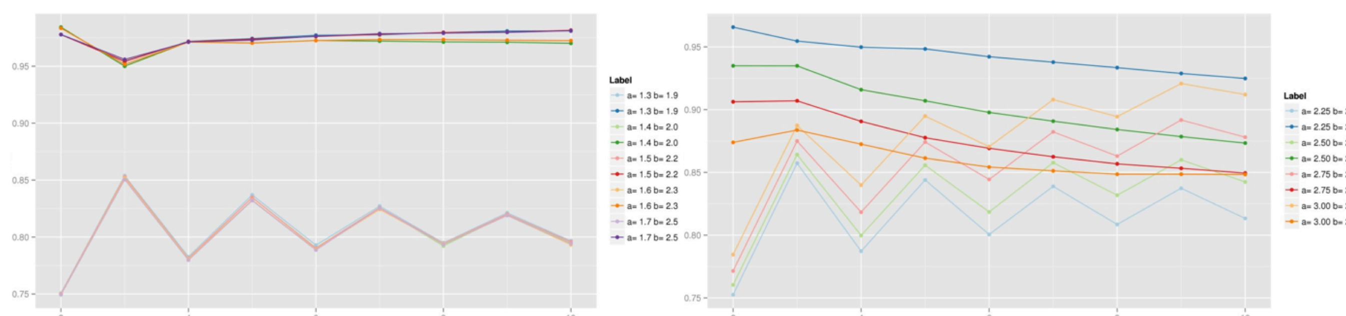

Figure 9: Expected efficiency of QV and majority rule under a Pearson Type I. The left panel shows cases with, a negative median but positive mean, where QV consistently outperforms majority rule. These have $\underline{u} = -1$ and $\overline{u} = 1.5$, with the values of α and β varying by colors as indicated at the legend at right. The right panel shows cases, with a positive mean and median, where QV sometimes underperforms majority rule. There $\underline{u} = -\overline{u} = -1$.

is right-skewed but $\beta > \alpha$ so that the mean of the distribution is positive while its median is negative. The different colors represent different values of α and β, which increase in tandem across curves thereby holding the mean approximately constant while reducing the variance. In all cases in this class, QV dramatically outperforms majority rule, achieving near-perfect efficiency, though this is most extreme when the variance is smallest. In the right panel is pictured a case when the support of the distribution is symmetric: $[-1, 1]$. β is fixed at 2 and $\alpha > 2$ varies. As it increases (moving towards brighter/lighter colors, the mean increases and the variance declines, making the setting less favorable to QV. When α is only 2.25 QV still significantly dominates majority rule for all values of N. However as α rises, especially for large N, QV's performs declines (though in a concave fashion; if α increase further beyond this point it begins to improve again) while majority rule improves. Once $\alpha = 3$ QV only outperforms majority rule for small even numbers of N where tie breaking is important.